DOC.
000 233

CPL/AC

PREGÃO 050/2003

LOCAÇÃO DE EQUIPAMENTOS DE INFORMÁTICA INCLUINDO ASSISTÊNCIA TÉCNICA E TREINAMENTO

COBRA TECNOLOGIA – MANUAL VOLUME 3

2003
PASTA 31

```
02/04/00 13:38:56 0a110210:389ac85b:3c6e:0001 "[CBF502]
DLT_ARC_8" [INITIALIZATION]


02/29/00 16:04:25 0a110210:38bbdff4:6d85:0026 "NULL_33"
[AUTOINITIALIZATION]


03/02/00 10:03:25 0a110210:385a24bf:410b:0002 "[CW1231]
BMW_DLT_15" [IMPORT]
```

upgrade.log

```
03/15/01 09:15:38

UCP session started.


03/15/01 09:20:55

UCP session finished.

total running time: 317 seconds


03/15/01 10:00:09

UDP session started.


03/15/01 10:02:54

Abort request from CLI/GUI on handle 0. Terminating
session


03/15/01 10:03:06

UDP session started.


03/15/01 10:26:47

Abort request from CLI/GUI on handle 0. Terminating
session
```

```
03/15/01 12:40:43

Database check error! Can not proceed with upgrade.


03/15/01 13:24:15

System error


03/15/01 13:24:15

Session was aborted by child ASM, marked error=1026


03/15/01 15:27:22

OmniBack II 3.x database not found.


03/15/01 16:33:19

[12:10904] Open of detail catalog binary file failed.


03/16/01 08:39:31

Internal error: Invalid Ct function argument specified.


03/20/01 10:56:57

[12:1165] Database network communication error.


03/22/01 14:38:21

[12:10953] Database is in incorrect state. Database must
be empty before critical upgrade can start.
```

18.862

Windows Manual Disaster Recovery Preparation Template

The template on the next page can be used to prepare for Windows Assisted Manual Disaster Recovery, as described in the Chapter 10, “Disaster Recovery,” on page 435.

19.858

Windows Manual Disaster Recovery Preparation Template

Table A-1

client properties	computer name	
	hostname	
drivers		
Windows Service Pack		
TCP/IP properties	IP address	
	default gateway	
	subnet mask	
	DNS order	
medium label / barcode number		
partition information and order	1st disk label	
	1st partition length	
	1st drive letter	
	1st filesystem	
	2nd disk label	
	2nd partition length	
	2nd drive letter	
	2nd filesystem	
	3rd disk label	
	3rd partition length	
	3rd drive letter	
	3rd filesystem	

18.860

Changing Block Size on Windows Media Agent

In order to increase the maximum block size on a Windows Media Agent client, you have to modify its Registry. After modifying the Registry, restart the computer. Drivers read `MaximumSGList` at boot time. The actual formula that a Windows class driver uses to determine the maximum transfer size is:

```
maximum size = ((number of supported scatter/gather
elements - 1)* 4096)
```

For the typical aic78xx case, it renders the following:

```
((17-1) * 4096) = 64k (which corresponds to 56k usable
data for Data Protector)
```

Windows provides a mechanism to support more scatter/gather elements via the Registry. Start the `regedit32` and add a `DWORD` value in the following Registry key:

```
\\HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\aic78
xx\Parameters\Device0\MaximumSGList
```

Use the following formula to calculate the value of the `MaximumSGList`:

$$MaximumSGList = \left(\frac{BlockSize}{4096}\right) + 1$$

Example

The `MaximumSGList` value for a 256k block size is 65:
`MaximumSGList` = (265k/4k) + 1 = 64 + 1 = 65

If you have, for example, 3 aic78xx based SCSI channels on your system, change the appropriate ...`\Device0`, ...`\Device1` or ...`\Device2` value. If you want to set all adapters at the same time, specify `MaximumSGList` for ...`\Device\`.... Omitting the numerical reference sets the value for all aic78xx adapters.

18.859

18.858

Glossary

access rights
See **user rights.**

ACSLS *(StorageTek specific term)*
The Automated Cartridge System Library Server (ACSLS) software that manages the Automated Cartridge System (ACS).

Active Directory *(Windows specific term)*
The directory service in a Windows network. It contains information about resources on the network and makes them accessible to users and applications. The directory services provide a consistent way to name, describe, locate, access and manage resources regardless of the physical system they reside on.

AML *(EMASS/GRAU specific term)*
Automated Mixed-Media library.

application agent
A component needed on a client to back up or restore online database integrations.
See also **Disk Agent**.

application system *(ZDB specific term)*
A system the application or database runs on. The application or database data is located on original units.
See also **backup system** and **original unit**.

archived redo log *(Oracle specific term)*
Also called offline redo log. If the Oracle8/9 database operates in the ARCHIVELOG mode, as each online redo log is filled, it is copied to one (or more) archived log destination(s). This copy is the Archived Redo Log. The presence or absence of an Archived Redo Log is determined by the mode that the database is using:

- ARCHIVELOG - The filled online redo log files are archived before they are reused. The database can be recovered from an instance and disk failure. The "hot" backup can be performed only when the database is running in this mode.
- NOARCHIVELOG - The filled online redo log files are not archived.

See also **online redo log.**

archive logging *(Lotus Domino Server specific term)*
Lotus Domino Server database mode where transaction log files are overwritten only after they have been backed up.

ASR Set
A collection of files stored on several diskettes required for proper reconfiguration of the replacement disk

48.857

Glossary

(disk partitioning and logical volume configuration) and automatic recovery of the original system configuration and user data that was backed up during the full client backup. You need these diskettes to perform ASR.

autochanger
See **library**

autoloader
See **library**

BACKINT *(SAP R/3 specific term)*
SAP R/3 backup programs can call the Data Protector backint interface program via an open interface, which enables them to communicate with Data Protector software. For backup and restore, SAP R/3 programs issue orders for the Data Protector backint interface.

backup API
The Oracle interface between the Oracle backup/restore utility and the backup/restore media management layer. The interface defines a set of routines to allow the reading and writing of data to the backup media, the creation, searching and removing the backup files.

backup chain
This relates to a situation where full and incremental backups are performed. Based on the level of the incremental backups used (Incr, Incr 1, Incr 2, and so on), simple or rather complex dependencies of incrementals to previous incrementals can exist. The backup chain are all backups, starting from the full backup plus all the dependent incrementals up to the desired point in time.

backup device
A device configured for use with Data Protector, which can write data to and read data from storage media. This can be, for example, a standalone DDS/DAT drive or a library.

backup generation
One backup generation includes one full backup and all incremental backups until the next full backup.

backup object
Any data selected for backup, such as a disk, a file, a directory, a database, or a part of it. During the backup session, Data Protector reads the objects, transfers the data (through the network), and writes them to the media residing in the devices.

backup owner
Each backup object in the IDB has an owner. The default owner of a backup is the user who starts the backup session.

backup session
A process that creates a copy of data on storage media. The activities are

specified in a backup specification or an interactive session. All clients configured in one backup specification are backed up together in one backup session using the same backup type (full or incremental). The result of a backup session is a set of media, which was written to, also called the backup or media set.
See also **incremental backup** and **full backup.**

backup set
See **media set.**

backup set *(Oracle specific term)*
Backup for (one or more) Oracle8/9 files, where the files are multiplexed together. The reason for multiplexing is to give performance benefits. Files in backup sets have to be extracted using a restore command. There are two types of backup sets: data file backup set and archive log backup set.

backup specification
A list of objects to be backed up, together with a set of devices or drives to be used, backup options for all objects in the specification, days and time that you want backups to be performed. The objects are entire disks/ volumes or parts of them such as files, directories, or even the Windows Registry for example. File selection lists such as include-lists and exclude-lists can be specified.

backup system *(ZDB specific term)*
A system connected to replica units of one or multiple application systems. The backup system is typically connected to a backup device to perform the backup of the data in a replica storage version.
See also **application system** and **replica unit.**

backup types
See **incremental backup**, **differential backup**, **transaction backup**, **full backup** and **delta backup**.

backup view
Data Protector provides different views for backup specifications: By Type - according to the type of data available for backups/templates. Default view. By Group - according to the group to which backup specifications/templates belong. By Name - according to the name of backup specifications/templates. By Manager - if you are running MoM, you can also set the Backup view according to the Cell Manager to which backup specifications/templates belong.

BC *(EMC Symmetrix specific term)*
Business Continuance are processes that allow customers to access and manage instant copies of EMC Symmetrix standard devices.
See also **BCV**.

18.855

Glossary

BC *(HP StorageWorks Disk Array XP specific term)*
The Business Copy XP allows to maintain internal copies of HP StorageWorks Disk Array XP LDEVs for purposes such as data backup or data duplication. The copies (secondary volumes or S-VOLs) can be separated from the primary volumes (P-VOLs) and connected to a different system for various purposes, such as backup and development. For backup purposes, P-VOLs should be connected to the application system, and one of the S-VOL mirror sets to the backup system.
See also **CA** *(HP StorageWorks Disk Array XP specific term)*, **Main Control Unit** and **HP StorageWorks Disk Array XP LDEV**.

BC *(HP StorageWorks Virtual Array specific term)*
Business Copy VA allows you to maintain internal copies of HP StorageWorks Virtual Array LUNs for data backup or data duplication. The copies (secondary volumes or S-VOLs) can be separated from the primary volumes (P-VOLs) and connected to a different system for various purposes, such as backup and development. For backup purposes, P-VOLs should be connected to the application system and one of the S-VOL sets should be connected to the backup system.
See also **HP StorageWorks Virtual Array** LUN.

BC Process *(EMC Symmetrix specific term)*
A protected storage environment solution that has defined specially configured EMC Symmetrix devices as mirrors or Business Continuance Volumes to protect data on EMC Symmetrix standard devices.
See also **BCV**.

BC VA *(HP StorageWorks Virtual Array specific term)*
Business Copy VA allows you to maintain internal copies of HP StorageWorks Virtual Array LUNs for data backup or data duplication within the same virtual array. The copies (child or Business Copy LUNs) can be used for various purposes, such as backup, data analysis or development. When used for backup purposes, the original (parent) LUNs are connected to the application system and the Business Copy (child) LUNs are connected to the backup system.

BCV *(EMC Symmetrix specific term)*
Business Continuance Volumes, or BCV devices, are dedicated SLDs that are pre-configured in the ICDA on which the business continuation operation runs. BCV devices are assigned separate SCSI addresses, differing from the addresses used by the SLDs they mirror. The BCV devices are used as splittable mirrors of the primary

18.854

EMC Symmetrix SLDs that need to be protected.
See also **BC** and **BC Process**.

boolean operators
The boolean operators for the full text search functionality of the online Help system are AND, OR, NOT, and NEAR. Used when searching, they enable you to define your query precisely by creating a relationship between search terms. If no operator is specified in a multi-word search, AND is used by default. For example, the query manual disaster recovery is equivalent to manual AND disaster AND recovery.

boot volume/disk/partition
A volume/disk/partition with files required for the initial step of the boot process. Microsoft terminology defines the boot volume/disk/partition as a volume/disk/partition containing the operating system files.

BRARCHIVE *(SAP R/3 specific term)*
An SAP R/3 backup tool that allows you to archive redo log files. BRARCHIVE also saves all the logs and profiles of the archiving process.
See also **SAPDBA**, **BRBACKUP** and **BRRESTORE**.

BRBACKUP *(SAP R/3 specific term)*
An SAP R/3 backup tool that allows an online or offline backup of the control file, of individual data files, or of all tablespaces and, if necessary, of the online redo log files.
See also **SAPDBA**, **BRARCHIVE** and **BRRESTORE**.

BRRESTORE *(SAP R/3 specific term)*
An SAP R/3 tool that can be used to restore files of the following type:

- Database data files, control files, and online redo log files saved with BRBACKUP
- Redo log files archived with BRARCHIVE
- Non-database files saved with BRBACKUP

You can specify files, tablespaces, complete backups, log sequence numbers of redo log files, or the session ID of the backup.
See also **SAPDBA**, **BRBACKUP** and **BRARCHIVE**.

BSM
The Data Protector Backup Session Manager controls the backup session. This process always runs on the Cell Manager system.

CA *(HP StorageWorks Disk Array XP specific term)*
Continuous Access XP allows you to create and maintain remote copies of HP StorageWorks Disk Array XP LDEVs

for purposes such as data duplication, backup, and disaster recovery. CA operations involve the main (primary) disk arrays and the remote (secondary) disk arrays. The main disk arrays contain the CA primary volumes (P-VOLs), which contain the original data and are connected to the application system. The remote disk arrays contain the CA secondary volumes (S-VOLs) connected to the backup system.
See also **BC** *(HP StorageWorks Disk Array XP specific term)*, **Main Control Unit** and **HP StorageWorks Disk Array XP LDEV**.

CAP *(StorageTek specific term)*
Cartridge Access Port is a port built into the door panel of a library. The purpose is to enter or eject media.

catalog protection
Defines how long information about backed up data (such as file names and file versions) is kept in the IDB.
See also **data protection**.

CDB
The Catalog Database is a part of the IDB that contains information about backup sessions, restore sessions, and backed up data. Depending on the selected log level, it also contains file names and file versions. This part of the database is always local to the cell.
See also **MMDB**.

CDF file *(UNIX specific term)*
A Context Dependent File is a file consisting of several files grouped under the same pathname. The system ordinarily selects one of the files using the context of the process. This mechanism allows machine dependent executables, system data, and device files to work correctly from all hosts in a cluster while using the same pathname.

cell
A set of systems that are under the control of a Cell Manager. The cell typically represents the systems on a site or an organizational entity, which are connected to the same LAN. Central control is available to administer the backup and restore policies and tasks.

Cell Manager
The main system in the cell where the essential Data Protector software is installed and from which all backup and restore activities are managed. The GUI used for management tasks can be located on a different system. Each cell has one Cell Manager system.

centralized licensing
Data Protector allows you to configure centralized licensing for the whole enterprise environment consisting of several cells. All Data Protector licenses are installed and kept on the Enterprise Cell Manager system. You can then

18.852

allocate licenses to specific cells to suit your needs.
See also **MoM**.

Centralized Media Management Database (CMMDB)
See CMMDB.

channel *(Oracle specific term)*
An Oracle8/9 Recovery Manager resource allocation. Every allocated channel starts a new Oracle8/9 process, which performs backup, restore, and recovery actions. The type of channel allocated determines the type of media used:

- type “disk”
- type ‘SBT_TAPE’

If the specified channel is type ‘SBT_TAPE’ and Oracle8/9 is integrated with Data Protector, the server process will attempt to read backups from or write data files to Data Protector.

circular logging *(MS Exchange and Lotus Domino Server specific term)*
Microsoft Exchange database and Lotus Domino Server database mode in which transaction log files are automatically overwritten as soon as the data they contain is committed to the database.

client backup
A backup of all filesystems mounted on a client. Filesystems mounted on the client after the backup specification was created are not automatically detected.

client backup with disk discovery
A backup of all filesystems mounted on a client. When the backup starts, Data Protector discovers the disks on the clients. Client backup with disk discovery simplifies backup configuration and improves backup coverage of systems that often mount or dismount disks.

client or **client system**
Any system configured with any Data Protector functionality and configured in a cell.

CMD Script for OnLine Server
(Informix specific term)
Windows CMD script that is created in INFORMIXDIR when Informix OnLine Server is configured. The CMD script is a set of system commands that export environment variables for OnLine Server.

CMMDB
The Data Protector Centralized Media Management Database (CMMDB) is the result of merging MMDBs from several cells in the MoM environment. It allows you to share high-end devices and media across multiple cells in a MoM

18.852

environment. One cell can control the robotics, serving the devices that are connected to systems in other cells. The CMMDB must reside on the MoM Manager. A reliable network connection between the MoM cell and the other Data Protector cells is highly recommended .
See also **MoM**.

COM+ Registration Database
The COM+ Registration Database and the Windows Registry store COM+ application attributes, class attributes, and computer-level attributes. This guarantees consistency among these attributes and provides common operation on top of these attributes.

command-line interface
A set of DOS and UNIX like commands that you can use in shell scripts to perform Data Protector configuration, backup, restore, and management tasks.

concurrency
See **Disk Agent concurrency**.

control file *(Oracle and SAP R/3 specific term)*
An Oracle data file that contains entries specifying the physical structure of the database. It provides database consistency information used for recovery.

CRS
The Cell Request Server process (service) runs on the Data Protector Cell Manager. It starts and controls the backup and restore sessions. The service is started as soon as Data Protector is installed on the Cell Manager.
CRS runs under the account root on UNIX systems, and under any Windows account. By default, it runs under the account of the user, specified at installation time.

data file *(Oracle and SAP R/3* specific term)
A physical file created by Oracle that contains data structures such as tables and indexes. A data file can only belong to one Oracle database.

data protection
Defines how long the backed up data on media remains protected, that is, Data Protector will not overwrite it. When the protection expires, Data Protector will be able to reuse the media in one of the next backup sessions.
See also **catalog protection**.

Data Protector Event Log
A central repository of all Data Protector related notifications. By default, all notifications are sent to the Event Log. The Event Log is accessible only to Data Protector users in the Admin group and to Data Protector

18.850

Glossary

users who are granted the Reporting and notifications user rights. You can view or delete all events in the Event Log.

Data Protector user account
You can use Data Protector only if you have a Data Protector user account, which restricts unauthorized access to Data Protector and to backed up data. Data Protector administrators create this account specifying a user logon name, the systems from which the user can log on, and a Data Protector user group membership. This is checked whenever the user starts the Data Protector user interface or performs specific tasks.

data stream
Sequence of data transferred over the communication channel.

database library
A Data Protector set of routines that enables data transfer between Data Protector and a server of an online database integration, for example, the Oracle8/9 Server.

database parallelism
More than one database is backed up at a time if the number of available devices allows you to perform backups in parallel.

database server
A computer with a large database stored on it, such as the SAP R/3 or Microsoft SQL database. A server has a database that can be accessed by clients.

Dbobject *(Informix specific term)*
An Informix physical database object. It can be a blobspace, dbspace, or logical-log file.

DC directory
The Detail Catalog (DC) directory consists of DC binary files, which store information about file versions. It represents the DCBF part of the IDB occupying approximately 80% of the IDB. The default DC directory is called the dcbf directory and is located in the *<Data_Protector_home>*\db40 directory. You can create more DC directories and locate them as appropriate to you. Up to 10 DC directories are supported per cell. The default maximum size of a DC directory is 2 GB.

DCBF
The Detail Catalog Binary Files (DCBF) are a part of the IDB. The files in store information about file versions and attributes occupying approximately 80% of the IDB. By default, DCBF consist of one DC directory with a maximum size of 2 GB. You can create more DC directories.

18.849

Glossary

delta backup
A delta backup is a backup containing all the changes made to the database from the last backup of any type.
See also **backup types**

device
A physical unit which contains either just a drive or a more complex unit such as a library.

device chain
A device chain consists of several standalone devices configured for sequential use. When a medium in one device gets full, the backup automatically continues on a medium in the next device in the device chain.

device group *(EMC Symmetrix specific term)*
A logical unit representing several EMC Symmetrix devices. A device cannot belong to more than a single device group. All devices in a device group must be on the same EMC Symmetrix unit. You can use a device group to identify and work with a subset of the available EMC Symmetrix devices.

device streaming
A device is streaming if it can feed enough data to the medium to keep it moving forward continuously. Otherwise, the tape has to be stopped, the device waits for more data, reverses the tape a little and resumes to write to the tape, and so on. In other words, if the data rate written to the tape is less or equal the data rate which can be delivered to the device by the computer system, then the device is streaming. Streaming significantly improves the performance of the device and use of space.

DHCP server
A system running the Dynamic Host Configuration Protocol (DHCP) providing dynamic configuration of IP addresses and related information. Data Protector can back up DHCP server data as part of the Windows configuration.

differential backup
An incremental backup (incr) based on any previous Data Protector backup (full or any incremental), which must still be protected.
See **incremental backup**.

differential backup *(MS SQL specific term)*
A database backup that records only the data changes made to the database after the last full database backup.
See also **backup types**.

differential database backup
A differential database backup records only those data changes made to the database after the last full database backup.

18.848

direct backup A SAN-based backup solution in which data movement directly from disk to tape (or to other secondary storage) is facilitated by the SCSI Extended Copy (Xcopy) command. Direct backup lessens the backup I/O load on systems in a SAN environment. The data movement is facilitated directly from disk to tape (or to other secondary storage) by the SCSI Extended Copy (XCopy) command. The command is provided by any element of the infrastructure including bridges, switches, tape libraries, and disk subsystems. *See also* **XCopy engine.**

directory junction *(Windows specific term)*
Directory junctions use the reparse point concept of Windows. An NTFS 5 directory junction allows you to redirect a directory/file request to another location.

Directory Store (DS) *(MS Exchange specific term)*
A part of the Microsoft Exchange Server directory. The Microsoft Exchange Server directory contains objects used by Microsoft Exchange applications in order to find and access services, mailboxes, recipients, public folders, and other addressable objects within the messaging system.
See also **Information Store (MDB).**

disaster recovery
A process to restore a client's main system disk to a state close to the time when a (full) backup was performed.

Disk Agent
A component needed on a client to back it up and restore it. The Disk Agent controls reading from and writing to a disk. During a backup session, the Disk Agent reads data from a disk and sends it to the Media Agent, which then moves it to the device. During a restore session the Disk Agent receives data from the Media Agent and writes it to the disk.

Disk Agent concurrency
The number of Disk Agents that are allowed to send data to one Media Agent concurrently.

disk discovery
The detection of disks during client backup with disk discovery. During this backup, Data Protector discovers (detects) the disks that are present on the client — even though they might not have been present on the system when the backup was configured — and backs them up. This is particularly useful in dynamic environments, where configurations change rapidly. After the disks are expanded, each inherits all options from its master client object. Even if pre- and post-exec commands are specified once, they are started many times, once per each object.

18.847

Glossary

disk group *(Veritas Volume Manager specific term)*
The basic unit of data storage in VxVM system. A disk group can consist of one or more physical volumes. There can be more than one disk group on the system.

disk image (rawdisk) backup
A high-speed backup where Data Protector backs up files as bitmap images. A disk image (rawdisk) backup does not track the files and directory structure stored on the disk, but stores a disk image structure on byte level. You can perform a disk image backup of either specific disk sections or a complete disk.

disk quota
A concept to manage disk space consumption for all or a subset of users on a computer system. This concept is used by several operating system platforms.

Distributed File System (DFS)
A service that connects file shares into a single namespace. The file shares can reside on the same or on different computers. DFS provides client access to the resources in a location-transparent manner.

DMZ
The Demilitarized Zone (DMZ) is a network inserted as a "neutral zone" between a company's private network (intranet) and the outside public network (Internet). It prevents outside users from getting direct access to company servers in the intranet.

DNS server
In the DNS client-server model, this is the server containing information about a portion of the DNS database that makes computer names available to client resolvers querying for name resolution across the Internet.

domain controller
A server in a network that is responsible for user security and verifying passwords within a group of other servers.

DR image
Data required for temporary disaster recovery operating system (DR OS) installation and configuration.

DR OS
A disaster recovery operating system is an operating system environment in which disaster recovery runs. It provides Data Protector a basic runtime environment (disk, network, tape, and filesystem access). The OS has to be installed and configured before the Data Protector disaster recovery can be performed. The DR OS can be either temporary or active. A temporary DR OS is used exclusively as a host environment for some other operating

system restore along with the target operating system configuration data. It is deleted after the target system is restored to the original system configuration. Active DR OS not only hosts the Data Protector disaster recovery process but is also a part of the restored system because it replaces its own configuration data with the original configuration data.

drive
A physical unit that receives data from a computer system and can write it onto a magnetic medium (typically a tape drive). It can also read the data from the medium and send it to the computer system.

drive index
A number that identifies the mechanical position of a drive inside a library device. This number is used by the robotic control to access a drive.

dynamic client
See **client backup with disk discovery**.

EMC Symmetrix Agent (SYMA)

(EMC Symmetrix specific term)
See **Symmetrix Agent (SYMA)**

EMC Symmetrix Application Programming Interface (SYMAPI)

(EMC Symmetrix specific term)
See **Symmetrix Application Programming Interface (SYMAPI)**

EMC Symmetrix CLI Database File

(EMC Symmetrix specific term)
See **Symmetrix CLI Database File**

EMC Symmetrix Command-Line Interface (SYMCLI) *(EMC Symmetrix specific term)*
See **Symmetrix Command-Line Interface (SYMCLI)**

emergency boot file *(Informix specific term)*
An Informix configuration file that resides in the <*INFORMIXDIR*>\etc directory (on HP-UX) or <*INFORMIXDIR*>/etc directory (on Windows) and is called ixbar.<*server_id*>, where <*INFORMIXDIR*> is the OnLine Server home directory and <*server_id*> is the value of the SERVERNUM configuration parameter. Each line of the emergency boot file corresponds to one backup object.

Enterprise Backup Environment
Several cells can be grouped together and managed from a central cell. The enterprise backup environment includes all clients located in several Data

18.845

Glossary

Protector cells which are managed and administered from a central cell using the Manager-of-Managers concept.
See also **MoM**.

EVA Agent *(HP StorageWorks Enterprise Virtual Array specific term)*
A Data Protector software module that executes all tasks required for the HP StorageWorks Enterprise Virtual Array snapshot integration on the application system and the backup system. It communicates with the HSV Element Manager to control the HP StorageWorks Enterprise Virtual Array.

Event Logs
Files in which Windows logs all events, such as the starting or stopping of services and the logging on and off of users. Data Protector can back up Windows Event Logs as part of the Windows configuration backup.

exchanger
Also referred to as SCSI II Exchanger.
See also **library**.

exporting media
A process that removes all data about backup sessions, such as systems, objects, and file names, which reside on the media from the IDB. Information about the media and their relation to a pool is also removed from the IDB. The data on the media remains unchanged.
See also **importing media.**

Extensible Storage Engine (ESE) *(MS Exchange specific term)*
A database technology used as a storage system for information exchange by Microsoft Exchange 2000 Server.

failover
Transferring of the most important cluster data, called group (on Windows) or package (on Unix) from one cluster node to another. A failover can occur mostly because of software or hardware failures or maintenance on the primary node.

FC bridge
See **Fibre Channel bridge**

Fibre Channel
An ANSI standard for high-speed computer interconnection. Using either optical or copper cables, it allows the high speed bidirectional transmission of large data files and can be deployed between sites kilometers apart.
Fibre Channel connects nodes using three physical topologies: point-to-point, loop, and switched.

Fibre Channel bridge
A Fibre Channel bridge or multiplexer provides the ability to migrate existing parallel SCSI devices, like RAID arrays, solid state disks (SSD), and tape libraries to a Fibre Channel environment. On one side of the bridge or multiplexer there is a Fibre Channel

18.844

interface while on the other side there are parallel SCSI ports. The bridge or multiplexer enables SCSI packets to be moved between the Fibre Channel and parallel SCSI devices.

File Replication Service (FRS)
A Windows service that replicates the domain controller store logon scripts and group policies. FRS also enables replication of Distributed File System (DFS) shares between systems and allows any server to perform replication activity.

file version
The same file can be backed up multiple times in case of full backups and incremental backups (if the file changed). If the log level ALL is selected for backup, Data Protector retains one entry in the IDB for the filename itself and one for each version (date/time) of the file.

filesystem
The organization of files on a hard disk. A filesystem is backed up so that the file attributes and the file contents are stored on the backup media.

first level mirror *(HP StorageWorks Disk Array XP specific term)*
HP StorageWorks Disk Array XP allows up to three mirror copies of a Primary Volume and each of these copies can have additional two copies. The three mirror copies are called first level mirrors.
See also **Primary Volume**, and **MU numbers**.

fnames.dat
The fnames.dat files of the IDB contain information on the names of the backed up files. Typically, these files occupy about 20% of the IDB, if filenames are stored.

formatting
A process that erases any data contained on a medium and prepares it for use with Data Protector. Information about media (media ID, description, and location) is saved in the IDB as well as on the respective media (media header). Data Protector media with protected data are not formatted until the protection expires or the media are unprotected/recycled.

free pool
An auxiliary source of media for use by media pools when they run out of media. The media pools must be configured to use free pools.

full backup
A backup in which all selected objects are backed up, whether or not they have been recently modified.
See also **backup types**.

18.843

Glossary

full database backup
A backup of all data in a database regardless of whether it has changed after the last database backup was created. This means that the full database backup does not depend on any other backup media.

full mailbox backup
A full mailbox backup is a backup of the entire mailbox content.

global options file
A file that allows you to customize Data Protector. It explains the global options, which cover various aspects of Data Protector, typically time-outs and limits, and affect the entire Data Protector cell. The file is located in the /etc/opt/omni/options directory on HP-UX and Solaris systems and in the *<Data_Protector_home>*\config\options directory on Windows systems.

group *(Microsoft Cluster Server specific term)*
A collection of resources (for example disk volumes, application services, IP names and addresses) that are needed to run a specific cluster-aware applications.

GUI
A cross-platform (X11/Motif and Windows) graphical user interface, provided by Data Protector for easy access to all configuration and administration tasks.

hard recovery *(MS Exchange specific term)*
Recovery of data on the level of the database engine (Extensible Storage Engine 98).

heartbeat
A cluster data set with a time stamp carrying information about the operational status of a particular cluster node. This data set or packet is distributed among all cluster nodes.

Hierarchical Storage Management (HSM)
A method for optimizing the use of expensive hard disk storage by migrating less frequently used data to less expensive optical platters. When needed, the data is migrated back to hard disk storage. This balances the need for fast retrieval from hard disk with the lower cost of optical platters.

Holidays file
A file that contains information about holidays. You can set different holidays by editing the Holidays file: /etc/opt/omni/Holidays on the UNIX Cell Manager and *<Data_Protector_home>*\Config\Holidays on the Windows Cell Manager.

host backup
See **client backup with disk discovery**.

18.842

Glossary

hosting system
A working Data Protector client used for Disk Delivery Disaster Recovery with a Data Protector Disk Agent installed.

HP ITO
See **OVO**.

HP OpC
See **OVO**.

HP OpenView SMART Plug-In (SPI)
A fully integrated, out-of-the-box solution which "plugs into" HP OpenView Operations, extending the managed domain. Through the Data Protector integration, which is implemented as an HP OpenView SMART Plug-In, a user can have an arbitrary number of Data Protector Cell Managers monitored as an extension to HP OpenView Operations (OVO).

HP OVO
See **OVO**.

HP StorageWorks Disk Array XP LDEV *(HP StorageWorks Disk Array XP specific term)*
A logical partition of a physical disk within an HP StorageWorks Disk Array XP. LDEVs are entities that are mirrored using Continuous Access XP (CA) and Business Copy XP (BC) configurations.
See also **BC** *(HP StorageWorks Disk Array XP specific term)* and **CA** *(HP StorageWorks Disk Array XP specific term.*

HP StorageWorks Virtual Array LUN *(HP StorageWorks Virtual Array specific term)*
A logical partition of a physical disk within an HP StorageWorks Virtual Array. LUNs are entities that are replicated using the HP StorageWorks Business Copy VA configuration.
See also **BC** *(HP StorageWorks Virtual Array specific term).*

HP VPO
See **OVO**.

HSV Element Manager *(HP StorageWorks Enterprise Virtual Array specific term)*
The HSV Element Manager is used by the Data Protector HP StorageWorks Enterprise Virtual Array integration to provides the features that enable virtualization technology and the management interface for the HP StorageWorks Enterprise Virtual Array environment.

ICDA *(EMC Symmetrix specific term)*
EMC's Symmetrix Integrated Cached Disk Arrays (ICDA) is a disk array device that combines a set of physical disks, a number of FWD SCSI channels,

18.841

an internal cache memory, and control and diagnostic software commonly referred to as the microcode.

IDB
The Data Protector Internal Database is an embedded database located on the Cell Manager that keeps information regarding which data is backed up, on which media it is backed up, how backup and restore sessions are run, and which devices and libraries are configured.

importing media
A process that re-reads all data about backup sessions which are on the medium back into the IDB. This then allows for fast and convenient access to the data on the media.
See also **exporting media.**

incremental backup
A backup that selects only files that have changed since a previous backup. Several levels of incremental backup are available, allowing selective backup of only files that have changed since the last incremental backup.
See also **backup types**.

incremental backup *(MS Exchange specific term)*
A backup of changes since the last full or incremental backup. Only transaction logs are backed up.
See also **backup types**.

incremental mailbox backup
An incremental mailbox backup backs up all the changes made to the mailbox after the last backup of any type.

incremental1 mailbox backup
An incremental1 mailbox backup backs up all the changes made to the mailbox after the last full backup.

incremental (re)-establish *(EMC Symmetrix specific term)*
A BCV or SRDF control operation. In BCV control operations, an incremental establish causes the BCV device to be synchronized incrementally and to function as an EMC Symmetrix mirrored medium. The EMC Symmetrix devices must have been previously paired.
In SRDF control operations, an incremental establish causes the target (R2) device to be synchronized incrementally and to function as an EMC Symmetrix mirrored medium. The EMC Symmetrix devices must have been previously paired.

incremental restore *(EMC Symmetrix specific term)*
A BCV or SRDF control operation.
In BCV control operations, an incremental restore reassigns a BCV device as the next available mirror of the standard device in the pair. However, the standard devices are updated with only the data that was written to the BCV

18.840
J

device during the time of the original pair split, and the data that was written to the standard device during the split is overwritten with data from the BCV mirror. In SRDF control operations, an incremental restore reassigns a target (R2) device as the next available mirror of the source (R1) device in the pair. However, the source (R1) devices are updated with only the data that was written to the target (R2) device during the time of the original pair split, and the data that was written to the source (R1) device during the split is overwritten with data from the target (R2) mirror.

Inet
A process that runs on each UNIX system or service that runs on each Windows system in the Data Protector cell. It is responsible for communication between systems in the cell and for starting other processes needed for backup and restore. The Inet service is started as soon as Data Protector is installed on a system. The Inet process is started by the inetd daemon.

Information Store (MDB) *(MS Exchange specific term)*
This is the default message store provider for the Microsoft Exchange Server. The information store consists of the following stores:

- Public information store (MS Exchange 5.5 Server) or Public folder store (MS Exchange 2000 Server)
- Private information store (MS Exchange 5.5 Server) or Mailbox store (MS Exchange 2000 Server)
- Personal folder store
- Offline information store.

The public information store contains public folders and messages that can be shared among multiple users and applications. A single public store is shared by all users within a Microsoft Exchange Server organization, even if multiple Servers are used. The private information store consists of mail boxes that can belong to users or to applications. The mail boxes reside on the server running the Microsoft Exchange Server.
See also **Directory Store (DS)**.

Initialization Parameter File
(Oracle specific term)
An Oracle8/9 file that contains information on how to initialize a database and instance.

initializing
See **formatting**.

Installation Server
A computer system that holds a repository of the Data Protector

software packages for a specific architecture. The Installation Server is used for remote installation of Data Protector clients. In mixed environments at least two Installation Servers are needed: one for UNIX systems and one for Windows systems.

instant recovery (*ZDB specific term*)
A process where data replicated during the ZDB disk backup or ZDB disk/tape backup is restored at high speed using split mirror or snapshot technology. The restore takes place within the disk array and there is no restore from the standard backup media involved. Full recovery of a database application may require further steps, such as applying the log files, to be performed afterwards. Instant recovery restores the user-selected replica storage version to the original storage.
See also **zero downtime backup (ZDB), ZDB disk backup, ZDB tape backup, ZDB disk/tape backup** and **replica storage pool.**

integrated security *(MS SQL specific term)*
Integrated security allows the Microsoft SQL Server to use Windows authentication mechanisms to validate Microsoft SQL Server logins for all connections. Using integrated security means that users have one password for both Windows and Microsoft SQL Server. Integrated security should be used in environments where all clients support trusted connections. Connections validated by Windows Server and accepted by Microsoft SQL Server are referred to as trusted connections. Only trusted connections are allowed.

Internet Information Server (IIS)

(Windows specific term)
Microsoft Internet Information Server is a network file and application server that supports multiple protocols. Primarily, IIS transmits information in Hypertext Markup Language (HTML) pages by using the Hypertext Transport Protocol (HTTP).

IP address
Internet Protocol address is a numeric address of a system used to identify the system on the network. The IP address consists of four groups of numbers separated by periods (full stops).

ISQL *(Sybase specific term)*
A Sybase utility used to perform system administration tasks on Sybase SQL Server.

ITO
See **OVO**.

jukebox
See **library**.

15.838

Glossary

LBO *(EMC Symmetrix specific term)*
A Logical Backup Object (LBO) is an object of data storage/retrieval in the EMC Symmetrix environment. It is stored/retrieved by EMC Symmetrix as one entity and can only be restored as a whole.

library
Also called autochanger, jukebox, autoloader, or exchanger. A library contains media in repository slots. Each slot holds one medium (for example, DDS/DAT). Media are moved between slots and drives by a robotic mechanism, allowing random access to media. The library can contain multiple drives.

lights-out operation or **unattended operation**
A backup or restore operation that takes place outside of normal business hours without an operator. This implies that no operator personnel is present to work with the backup application or service mount requests, for example.

LISTENER.ORA *(Oracle specific term)*
An Oracle configuration file that describes one or more Transparent Network Substrate (TNS) listeners on a server.

load balancing
By default, Data Protector automatically balances the usage of devices selected for backup, so that they are used evenly. Load balancing optimizes the usage by balancing the number and the size of the objects backed up to each device. Since load balancing is done automatically during backup time, you do not need to manage how the data is actually backed up. You just specify the devices to be used. If load balancing is not selected, you select which device will be used for each object in your backup specification. Data Protector will access the devices in the specified order.

local and remote recovery
Remote recovery is performed if all Media Agent hosts specified in the SRD file are accessible. If any of them fails, the disaster recovery process fails over to the local mode. This means that the target system is searched for locally attached devices. If only one device is found, it is automatically used. Otherwise, Data Protector prompts you to select the device, which will be used for restore.

lock name
You can configure the same physical device several times with different characteristics, by using different device names.

The lock name is a user specified string that is used for locking all such device configurations to prevent collision if several such devices (device names) are used concurrently.

log_full shell script *(Informix UNIX specific term)*
A script provided by ON-Bar that you can use to start backing up logical-log files when OnLine Server issues a log-full event alarm. The Informix ALARMPROGRAM configuration parameter defaults to the <*INFORMIXDIR*>/etc/log_full.sh, where <*INFORMIXDIR*> is the OnLine Server home directory. If you do not want logical logs to be backed up continuously, set the ALARMPROGRAM configuration parameter to <*INFORMIXDIR*>/etc/no_log.sh.

logging level
The logging level determines the amount of details on files and directories written to the IDB during backup. You can always restore your data, regardless of the logging level used during backup. Data Protector provides four logging levels: Log All, Log Directories, Log Files, and No Log. The different logging level settings influence IDB growth, backup speed, and the convenience of browsing data for restore.

logical-log files
This applies to online database backup. Logical-log files are files in which modified data is first stored before being flushed to disk. In the event of a failure, these logical-log files are used to roll forward all transactions that have been committed as well as roll back any transactions that have not been committed.

login ID *(MS SQL Server specific term)*
The name a user uses to log on to Microsoft SQL Server. A login ID is valid if Microsoft SQL Server has an entry for that user in the system table syslogin.

login information to the Oracle Target Database *(Oracle and SAP R/3 specific term)*
The format of the login information is <user_name>/<password>@<service>, where:

- <user_name> is the name by which a user is known to Oracle Server and to other users. Every user name is associated with a password and both have to be entered to connect to an Oracle Target Database. This user must have been granted Oracle SYSDBA or SYSOPER rights.

- <password> is a string used for data security and known only to its owner. Passwords are entered to connect to an operating system or software application. The password has to be the same as the password specified in the Oracle password file (orapwd), which is used for authentication of users performing database administration.

- <service> is the name used to identify an SQL*Net server process for the target database.

login information to the Recovery Catalog Database *(Oracle specific term)*
The format of the login information to the Recovery (Oracle8/9) Catalog Database is <user_name>/<password>@<service>, where the description of the user name, password, and service name is the same as in the Oracle SQL*Net V2 login information to the Oracle target database. In this case, <service> is the name of the service to the Recovery Catalog Database, not the Oracle target database.

Note that the Oracle user specified here has to be the owner of the Oracle Recovery (Oracle8/9) Catalog.

Lotus C API *(Lotus Domino Server specific term)*
An interface for the exchange of backup and recovery information between Lotus Domino Server and a backup solution, like Data Protector.

LVM
A Logical Volume Manager is a subsystem for structuring and mapping physical disk space to logical volumes on UNIX systems. An LVM system consists of several volume groups, where each volume group has several volumes.

Magic Packet
See **Wake ONLAN**.

mailbox *(MS Exchange specific term)*
The location to which email is delivered, which is set up by the administrator for each user. If a set of personal folders is designated as the email delivery location, email is routed from the mailbox to this location.

Mailbox Store *(MS Exchange 2000 Server specific term)*
A part of the Information Store that maintains information about user mailboxes. A mailbox store consists of a binary rich-text .edb file and a streaming native internet content .stm file.

Main Control Unit (MCU) *(HP StorageWorks Disk Array XP specific term)*
An HP StorageWorks XP disk array that contains the primary volumes for the Continuous Access configuration and acts as a master device.
See also **BC** *(HP StorageWorks Disk Array XP specific term)*, **CA** *(HP*

18835

StorageWorks Disk Array XP specific term) and **HP StorageWorks Disk Array XP LDEV**.

Manager-of-Managers (MoM)
See **Enterprise Cell Manager**.

Media Agent
A process that controls reading from and writing to a device, which reads from or writes to a medium (typically a tape). During a backup session, the Media Agent receives data from the Disk Agent and sends it to the device for writing it to the medium. During a restore session, the Media Agent locates data on the backup medium and sends it to the Disk Agent. The Disk Agent then writes the data to the disk. The Media Agent also manages the robotics control of a library.

MAPI *(MS Exchange specific term)*
The MAPI (Messaging Application Programming Interface) is the programming interface that lets applications and messaging clients interact with messaging and information systems.

media allocation policy
Determines in which sequence media are used for backup. The Strict allocation policy directs Data Protector to prompt for a specific medium. The Loose policy directs Data Protector to prompt for any suitable medium. The Formatted First policy directs Data Protector to give preference to unknown media, even if unprotected media are available in the library.

media condition
The quality of a medium as derived from the media condition factors. Heavy usage and age result in an increased number of read and write errors with tape media. Media need to be replaced when they are marked as POOR.

media condition factors
The user-assigned age threshold and overwrite threshold used to determine the state of a medium.

media ID
A unique identifier assigned to a medium by Data Protector.

media label
A user-defined identifier used to describe a medium.

media location
A user-defined physical location of a medium, such as "building 4" or "off-site storage".

media management session
A session performing some action on a medium, such as initializing, scanning the content, verifying data on a medium, or copying a medium.

18.834

Glossary

media pool
A set of media of the same type (such as DDS) used and tracked as a group. Media are formatted and assigned to a media pool.

media set
The result of a backup session is data backed up on a group of media called media set. Depending on the media usage policy, several sessions can share the same media.

media type
The physical type of media, such as DDS or DLT.

media usage policy
The media usage policy controls how new backups are added to the already used media. It can be Appendable, Non-Appendable, or Appendable for incrementals only.

merging
This defines one mode to resolve file conflicts during restore. If the file to be restored already exists at the destination, the one with the more recent modification date is kept. Files not present on the disk are always restored.
See also **overwrite**.

MFS
The Migrating File System enables a standard JFS filesystem with migration capabilities (on HP-UX 11.00). The MFS is accessed via a standard filesystem interface (DMAPI), it is mounted to a directory the same way as any HP-UX filesystem. In an MFS, only the superblock, the inode and the 'extended attribute' information remain permanently on the hard disk and are never migrated.
See also **VBFS**.

Microsoft Exchange Server
A “client-server” messaging and a workgroup system that offers a transparent connection to many different communication systems. It provides users with an electronic mail system, individual and group scheduling, online forms, and workflow automation tools. It provides a developer with a platform on which to build custom information-sharing and messaging-service applications.

Microsoft Management Console (MMC) *(Windows specific term)*
An administration model for Windows-based environments. It provides a simple, consistent, and integrated administration user interface allowing management of many applications through the same GUI, provided that the applications adhere to the MMC model.

Microsoft SQL Server 7.0/2000
A database management system designed to meet the requirements of distributed "client-server" computing.

18.833

Microsoft Volume Shadow Copy service (VSS)
A software service that provides a unified communication interface to coordinate backup and restore of a VSS-aware application regardless of its specific features. This service collaborates with the backup application, writers, shadow copy providers, and the operating system kernel to implement the management of volume shadow copies and shadow copy sets.
See also **shadow copy, shadow copy provider, writer**.

mirror *(ZDB specific term)*
See **replica unit**.

mirror rotation *(HP StorageWorks Disk Array XP specific term)*
See **replica storage rotation**.

MMD
The Media Management Daemon process (service) runs on the Data Protector Cell Manager and controls media management and device operations. The process is started when Data Protector is installed on the Cell Manager.

MMDB
The Media Management Database (MMDB) is a part of the IDB that contains information about media, media pools, devices, libraries, library drives, and slots configured in the cell, as well as the Data Protector media used for backup. In an enterprise backup environment, this part of the database can be common to all cells.
See also **CMMDB**, **CDB**.

MoM
Several cells can be grouped together and managed from a central cell. The management system of the central cell is the Manager-of-Managers (MoM). The MoM allows you to configure and manage multiple cells from a central point.

mount request
A screen prompt that tells you to insert a specific medium into a device. Once you respond to the mount request by providing the required medium and confirm the mount request, the session continues.

mount point
The access point in a directory structure for a disk or logical volume, for example /opt or d:. On UNIX the mountpoints are displayed using the bdf or df command.

MSM
The Data Protector Media Session Manager, which runs on the Cell Manager and controls media sessions, such as copying media.

MU number (*HP StorageWorks Disk Array XP specific term*)
A Mirror Unit number is an integer number (0, 1 or 2), used to indicate a first level mirror.
See also **first level mirror**.

multi-drive server
A license that allows you to run an unlimited number of Media Agents on a single system. This license, which is bound to the IP address of the Cell Manager, is no longer available.

obdrindex.dat
An IDB file with information about IDB backups, media, and devices used for the backup. This information can significantly simplify IDB recovery. It is recommended to relocate the file, together with IDB transaction logs, on a separate physical disk from other IDB directories, and, additionally, to make a copy of the file and locate it where you want.

OBDR capable device
A device that can emulate a CD-ROM drive loaded with a bootable disk and can thus be used as a backup or boot device for disaster recovery purposes.

object
An object can be one of the following:

- for Windows clients, an object is a logical disk (such as d:);
- for UNIX clients, an object is a mounted filesystem or a mount point;
- for Novell Netware clients, an object is a volume.

The scope of the data can be further reduced by selecting files or directories. Additionally, an object can be a database entity.

Object ID *(Windows specific term)*
The object IDs (OIDs) enable access to NTFS 5 files no matter where in the system the files reside. Data Protector treats the OIDs as alternate streams of the files.

offline recovery
Offline recovery is performed if the Cell Manager is not accessible, for example, due to network problems. Only standalone and SCSI-II library devices can be used for offline recovery. Recovery of the Cell Manager is always offline.

offline redo log
See **archived redo log**

OmniStorage
Software providing transparent migration of less frequently used data to the optical library while keeping more frequently used data on the hard disk. HP OmniStorage runs on HP-UX systems.

18.831

On-Bar *(Informix specific term)*
A backup and restore system for OnLine Server. ON-Bar enables you to create a copy of your OnLine Server data and later restore the data. The ON-Bar backup and restore system involves the following components:

- onbar utility
- Data Protector, as the backup solution
- XBSA interface
- ON-Bar catalog tables, which are used to back up dbobjects and track instances of dbobjects through multiple backups.

onbar utility *(Informix specific term)*
The Informix utility that communicates backup and restore requests to OnLine Server. The utility uses XBSA to exchange control data and back up and restore data with Data Protector.

ONCONFIG *(Informix specific term)*
An environment variable that specifies the name of the active ONCONFIG configuration file. If the ONCONFIG environment variable is not present, OnLine uses the configuration values from the file <*INFORMIXDIR*>\etc\onconfig (on HP-UX) or <*INFORMIXDIR*>/etc/onconfig (on Windows).

online backup
A backup that is performed while the application (or database) is available for use. Application-specific interfaces allow backup products, like Data Protector, to back up logical units of the database while retaining access for the application. In simple configurations (non ZDB), the application remains in a backup mode for the entire duration of the backup. In contrast to that, for ZDB configurations, the backup mode lasts only for the duration of the split/snapshot operation. After that, the application can resume to the standard mode. Depending on the configuration, resource requirements vary significantly.

online redo log *(Oracle specific term)*
Redo logs that have not been archived, but are either available to the instance for recording database activity or are filled and waiting to be archived or reused.
See also **archived redo log**.

OnLine Server *(Informix specific term)*
Refers to INFORMIX-OnLine Dynamic Server.

OpC
See **OVO**.

18.830
3

Oracle instance *(Oracle specific term)* Each installation of an Oracle database on one or more systems. One computer system can have several instances of a database running.

ORACLE_SID *(Oracle specific term)* A unique name for an Oracle Server instance. To switch among Oracle Servers, specify the desired *<ORACLE_SID>*. The *<ORACLE_SID>* is included in the CONNECT DATA parts of the connect descriptor in a TNSNAMES.ORA file and in the definition of the TNS listener in the LISTENER.ORA file.

original system
The system configuration backed up by Data Protector before a computer disaster hits the system.

original unit *(ZDB specific term)*
A logical unit that is used as a source for data replication using snapshot or split mirror technologies. Depending on the vendor and technology used, an original unit denotes P-VOL on HP StorageWorks Disk Array XP, parent LUN on HP StorageWorks Virtual Array, logical drive on HP StorageWorks Modular SAN Array 1000, or virtual disk on HP StorageWorks Enterprise Virtual Array. Data in an original unit is replicated to data in a replica unit. Original units are on systems interpreted as physical drives (Windows) or physical volumes (UNIX).
See also **replica unit**, **original storage**, and **replica storage version**.

original storage *(ZDB specific term)*
A set of original units that contain the backup objects selected in one Data Protector backup specification. Data in an original storage is replicated to data in a replica storage version by replicating the set of original units. An original storage is typically used by the application system.
See also **original unit**, **replica unit**, and **replica storage version.**

overwrite
An option that defines one mode to resolve file conflicts during restore. All files are restored from a backup even if they are older than existing files.
See also **merging**.

OVO
HP OpenView Operations for Unix provides powerful capabilities for operations management of a large number of systems and applications on in a network. Data Protector provides an integration into this management product. This integration is implemented as a SMART Plug-In for OVO management servers on HP-UX and Solaris. Earlier versions of OVO were

called IT/Operation, Operations Center and Vantage Point Operations.
See also **merging**.

ownership
The ownership of a backup determines who can restore from the backup. The user who starts an interactive backup is the session owner. If a user starts an existing backup specification without modifying it, the session is not considered interactive. In that case, if the backup owner has been defined in the backup specification, they remain the session owner. Otherwise, the session owner becomes the user who started the backup in question. For the scheduled backups, by default, the session owner is for the UNIX Cell Manager: root.sys@<Cell Manager>, and for the Windows Cell Manager, the user that was specified during the installation of the Cell Manager. It is possible to modify the ownership, so that the specific user becomes the session owner.

package *(MC/ServiceGuard and Veritas Cluster specific term)*
A collection of resources (for example volume groups, application services, IP names and addresses) that are needed to run a specific cluster-aware application.

pair status *(HP StorageWorks Disk Array XP specific term)*
A mirrored pair of disks can have various status values depending on the action performed on it. The three most important status values are:

- COPY - The mirrored pair is currently resynchronizing. Data is transferred from one disk to the other. The disks do not contain the same data.
- PAIR - The mirrored pair is completely synchronized and both disks (the primary volume and the mirrored volume) contain identical data.
- SUSPENDED - The link between the mirrored disks is suspended. That means that both disks are accessed and updated independently. However, the mirror relationship is still maintained and the pair can be resynchronized without transferring the complete disk.

parallel restore
Restoring backed up data to multiple disks at the same time (that is, in parallel) by running multiple Disk Agents, that receive data from one Media Agent. For the parallel restore to work, select data that is located on different disks or logical volumes and during backup, the data from the different objects must have been sent to the same device using a concurrency of 2 or more. During a parallel restore, the

data for multiple objects selected for restore is read from media at the same time, thereby improving performance.

parallelism
The concept of reading multiple data streams from an online database.

physical device
A physical unit that contains either a drive or a more complex unit such as a library.

post-exec
A backup option that executes a command or script after the backup of an object or after the entire session completes. Post-exec commands are not supplied by Data Protector. You need to create your own. They can be written as executables or batch files on Windows and as shell scripts on UNIX.
See also **pre-exec**.

pre- and post-exec commands
Pre- and post-exec commands are used to perform additional action before and after a backup or restore session. They are not supplied by Data Protector. You need to create your own commands. They can be written as executables or batch files on Windows and as shell scripts on UNIX.

prealloc list
A subset of media in a media pool that specifies the order in which media are used for backup.

pre-exec
A backup option that executes a command or script before the backup of an object or before the entire session is started. Pre-exec commands are not supplied by Data Protector. You need to create your own. They can be written as executables or batch files on Windows and as shell scripts on UNIX.
See also **post-exec**.

Primary Volume (P-VOL) *(HP StorageWorks Disk Array XP specific term)*
Standard HP StorageWorks Disk Array XP LDEVs that act as a primary volume for the CA and BC configurations. The P-VOL is located in the MCU.
See also **Secondary Volume (S-VOL)**.

Private Information Store *(MS Exchange 5.5 Server specific term)*
A part of the Information Store that maintains information in user mailboxes. A mailbox store consists of a binary rich-text .edb file.

protection
See **data protection** and also **catalog protection**.

18.827

public folder store *(MS Exchange 2000 Server specific term)*
The part of the Information Store that maintains information in public folders. A public folder store consists of a binary rich-text .edb file and a streaming native internet content .stm file.

public/private backed up data
When configuring a backup, you can select whether the backed up data will be:

- public, that is visible (and accessible for restore) to all Data Protector users
- private, that is, visible (and accessible for restore) only to the owner of the backup and administrators

RAID
Redundant Array of Inexpensive Disks.

RAID Manager Library *(HP StorageWorks Disk Array XP specific term)*
The RAID Manager Library is used internally by Data Protector on Solaris systems to allow access to HP StorageWorks Disk Array XP configuration, status, and performance data and to key HP StorageWorks Disk Array XP features through the use of function calls translated into a sequence of low level SCSI commands.

RAID Manager XP *(HP StorageWorks Disk Array XP specific term)*
The RAID Manager XP application provides an extensive list of commands to report and control the status of the CA and BC applications. The commands communicate through a RAID Manager instance with the HP StorageWorks Disk Array XP Disk Control Unit. This instance translates the commands into a sequence of low level SCSI commands.

rawdisk backup
See **disk image backup**.

RCU *(HP StorageWorks specific term)*
The Remote Control Unit acts as a slave of an MCU in a CA configuration. In bidirectional configurations, the RCU can act as an MCU.

RDBMS
Relational Database Management System.

RDF1/RDF2 *(EMC Symmetrix specific term)*
A type of SRDF device group. Only RDF devices can be assigned to an RDF group. An RDF1 group type contains source (R1) devices and an RDF2 group type contains target (R2) devices.

RDS
The Raima Database Server process (service) runs on the Data Protector Cell

18.826

Manager and manages the IDB. The process is started when Data Protector is installed on the Cell Manager.

Recovery Catalog *(Oracle specific term)*
A set of Oracle8/9 tables and views that are used by Recovery Manager to store information about Oracle8/9 databases. This information is used by Recovery Manager to manage the backup, restore, and recovery of Oracle8/9 databases. The recovery catalog contains information about:

- The physical schema of the Oracle8/9 target database
- Data file and archivelog backup sets
- Data file copies
- Archived Redo Logs
- Stored scripts.

Recovery Catalog Database *(Oracle specific term)*
An Oracle database that contains a recovery catalog schema. You should not store the recovery catalog in your target database.

RecoveryInfo
When backing up Windows configuration files, Data Protector collects the information about the current system configuration (information on disk layout, volume, and network configuration). This information is needed for disaster recovery.

Recovery Manager (RMAN) *(Oracle specific term)*
An Oracle8/9 command-line interface that directs an Oracle8/9 Server process to back up, restore, or recover the database it is connected to. RMAN uses either the recovery catalog or the control file to store information about backups. This information can be used later in restore sessions.

recycle
A process that removes the data protection from all backed up data on a medium, allowing Data Protector to overwrite it during one of the next backups. Data that belongs to the same session(s) but resides on other media is also unprotected. Recycling does not actually alter the data on the medium.

redo log *(Oracle specific term)*
Every Oracle database has a set of two or more redo log files. The set of redo log files for the database is known as the database's redo log. Oracle uses the redo log to record all changes made to data.

Remote Control Unit *(HP StorageWorks Disk Array XP specific term)*
The Remote Control Unit (RCU) acts as a slave of an MCU in a CA

configuration. In bidirectional configurations, the RCU can act as an MCU.

Removable Storage Management Database *(Windows specific term)*
A Windows service used for managing removable media (such as tapes and disks) and storage devices (libraries). Removable Storage allows applications to access and share the same media resources.

reparse point *(Windows specific term)*
A system-controlled attribute that can be associated with any directory or file. The value of a reparse attribute can have user-controlled data. The format of the data is understood by the application that stored the data and a filesystem filter that was installed to interpret the data and process such files. Whenever the filesystem encounters a file with a reparse point, it attempts to find the filesystem filter associated with the data format.

replica unit *(ZDB specific term)*
A logical unit that is used as a target for data replication using snapshot or split mirror technologies. Depending on the vendor and technology used, a replica unit denotes S-VOL on HP StorageWorks Disk Array XP, child (BC) LUN on HP StorageWorks Virtual Array, logical drive on HP StorageWorks Modular SAN Array 1000, or virtual disk on HP StorageWorks Enterprise Virtual Array. Data in an original unit is replicated to data in a replica unit. Replica units are on systems interpreted as physical drives (Windows) or physical volumes (UNIX). A replica unit is also referred to as snapshot or mirror.
See also **original unit**, **original storage**, and **replica storage version.**

replica storage version *(ZDB specific term)*
A set of replica units, created or reused during one ZDB backup session, which contain replica copies of the backup objects selected in one Data Protector backup specification. Data in an original storage is replicated to data in a replica storage version. A replica storage version is typically used by the backup system.
See also **original unit**, **replica unit**, and **original storage**.

replica storage pool *(ZDB specific term)*
A number or group of replica storage versions produced during ZDB sessions to be used for the purpose of replica storage rotation, instant recovery, and split mirror restore. The replica storage versions in the replica storage pool are all created using the same backup specification. The size of a replica storage pool is defined for each backup specification as the maximum number

18.824
J

of replica storage versions that are to be kept on a disk array before the oldest replica storage version for the backup specification is reused.
See also **replica storage rotation.**

replica storage rotation *(ZDB specific term)*
A ZDB process that denotes either a reuse of the oldest replica storage version in the replica storage pool whenever the size of the replica storage pool is reached or, if the size of the replica storage pool is not reached, a creation of a new replica storage version in the replica storage pool.
See also **replica storage pool.**

restore session
A process that copies data from backup media to a client.

RMAN *(Oracle specific term)*
See **Recovery Manager**.

RSM
The Data Protector Restore Session Manager controls the restore session. This process always runs on the Cell Manager system.

RSM *(Windows specific term)*
Removable Storage Manager (RSM) includes a media management service that facilitates communication among applications, robotic changers, and media libraries. It enables multiple applications to share local robotic media libraries and tape or disk drives and to manage removable media.

SA Agent *(HP StorageWorks Modular SAN Array 1000 specific term)*
A Data Protector software module that executes all tasks required for the HP StorageWorks Modular SAN Array 1000 snapshot integration on the application system and the backup system. It communicates with the HP StorageWorks Modular SAN Array 1000 Business Copy Manager to control the HP StorageWorks Modular SAN Array 1000.

SAPDBA *(SAP R/3 specific term)*
An SAP R/3 user interface that integrates the BRBACKUP, BRARCHIVE, and BRRESTORE tools.

scan
A function that identifies the media in a device. This synchronizes the MMDB with the media that are actually present at the selected locations (for example, slots in a library).

scanning
A function which identifies the media in a device. This synchronizes the MMDB with the media that are actually present at the selected locations (for example, slots in a library). It is useful to perform a scan and check the actual media in the

18.823
J

device if someone has manually manipulated media without using Data Protector to eject or enter, for example.

Scheduler
A function that controls when and how often automatic backups occur. By setting up a schedule, you automate the start of backups.

Secondary Volume (S-VOL) *(HP StorageWorks Disk Array XP specific term)*
Secondary Volumes, or S-VOLs, are XP LDEVs that act as a secondary CA or BC mirror of another LDEV (a P-VOL). In the case of CA, S-VOLs can be used as failover devices in a MetroCluster configuration. The S-VOLs are assigned separate SCSI addresses, different from the addresses used by the P-VOLs.
See also **Primary Volume (P-VOL).**

session
See **backup session, media management session,** and **restore session.**

session ID
This environment variable is set by Data Protector during actual backup sessions (not during preview). It identifies a session and is recorded in the database.

session key
This environment variable for the Pre- and Post-exec script is a Data Protector unique identification of any session, including preview sessions. The session key is not recorded in the database, and it is used for specifying options for the omnimnt, omnistat and omniabort CLI commands.

shadow copy *(MS VSS specific term)*
A volume that represents a duplicate of the original volume at a certain point in time. The data is then backed up from the shadow copy and not from the original volume. The original volume continues to changes as the backup process continues, but the shadow copy of the volume remains constant.
See also **Microsoft Volume Shadow Copy service.**

shadow copy provider *(MS VSS specific term)*
An entity that performs the work on creating and representing the volume shadow copies. Providers own the shadow copy data and expose the shadow copies. Providers can be software (e.g. system providers) or hardware (local disks, disk arrays).
See also **shadow copy.**

shadow copy set *(MS VSS specific term)*
A collection of shadow copies created at the same point in time.
See also **shadow copy**.

18.522

shared disks
A Windows disk on another system that has been made available to other users on the network. Systems with shared disks can be backed up without a Data Protector Disk Agent installed.

SIBF
The Serverless Integrations Binary Files (SIBF) is a part of the IDB that stores raw NDMP meta data. This data is necessary to perform restore of NDMP objects.

slot
A mechanical position in a library. Each slot can hold a medium, such as a DLT tape. Data Protector references each slot with a number. To read a medium, a robotic mechanism moves the medium from a slot into the drive.

SMB
See **split mirror backup**.

SMBF
The Session Messages Binary Files (SMBF) is a part of the IDB that stores session messages generated during backup and restore sessions. One binary file is created per session. The files are grouped by year and month.

snapshot *(ZDB specific term)*
See **replica unit**.

snapshot backup *(ZDB specific term)*
A ZDB term encompassing ZDB disk backup, ZDB tape backup and ZDB disk/tape backup utilizing snapshot technology.
See also **zero downtime backup (ZDB)**.

source (R1) device *(EMC Symmetrix specific term)*
An EMC Symmetrix device that participates in SRDF operations with a target (R2) device. All writes to this device are mirrored to a target (R2) device in a remote EMC Symmetrix unit. An R1 device must be assigned to an RDF1 group type.
See also **target (R2) device**.

source medium
When copying media, the source medium is the medium that contains backed up data and is being copied.

sparse file A file that contains data with portions of empty blocks. Examples are: -A matrix in which some or much of the data contains zeros -files from image applications -high-speed databases If sparse file processing is not enabled during restore, it might be impossible to restore this file.

split mirror backup *(EMC Symmetrix specific term)*
See **ZDB tape backup**.

Glossary

split mirror backup *(HP StorageWorks Disk Array XP specific term)*
See **ZDB tape backup**, **ZDB disk/tape backup** and **ZDB disk backup**.

split mirror restore *(HP StorageWorks Disk Array XP specific term)*
A process where data backed up using the ZDB tape backup or ZDB disk/tape backup process is restored from tape media to the replica storage version selected by the replica rotation process or by the user. The replica storage version is then synchronized to the original storage. Split mirror restore is limited to filesystem restore.
See also **ZDB tape backup**, **ZDB disk/ tape backup**, and **replica storage rotation**.

sqlhosts file *(Informix specific term)*
An Informix connectivity-information file that contains the names of each of the database servers and any aliases to which the clients on a host computer can connect.

SRD file
The Data Protector System Recovery Data (SRD) file contains system information required for installing and configuring the operating system in case of a disaster. The SRD file is an ASCII file, generated when a CONFIGURATION backup is performed on a Windows client and stored on the Cell Manager.

SRDF *(EMC Symmetrix specific term)*
The EMC Symmetrix Remote Data Facility is a business continuation process that enables effective, real-time data replication of SLDs between dislocated processing environments. These environments could be situated within the same root computer environment or separated by long distances.

SSE Agent *(HP StorageWorks Disk Array XP specific term)*
A Data Protector software module that executes all tasks required for a split mirror backup integration. It communicates with the HP StorageWorks Disk Array XP storing system using the RAID Manager XP utility (HP-UX and Windows systems) or RAID Manager Library (Solaris systems).

sst.conf file
The file /usr/kernel/drv/sst.conf is required on each Data Protector Sun Solaris client to which a multi-drive library device is connected. It must contain an entry for the SCSI address of the robotic mechanism of each library device connected to the client.

st.conf file
The file /kernel/drv/st.conf is required on each Data Protector Solaris client with a backup device connected. It must contain device information and a SCSI

address for each backup drive connected to that client. A single SCSI entry is required for a single-drive device and multiple SCSI entries are required for a multi-drive library device.

stackers
Devices with multiple slots for media storage usually with only one drive. A stacker selects media from the stack sequentially. In contrast, a library can randomly select media from its repository.

standard security *(MS SQL specific term)*
Standard security uses the login validation process of the Microsoft SQL Server for all connections. Standard security is useful in network environments with a variety of clients, some of which may not support trusted connections. It also provides backward compatibility for older versions of the Microsoft SQL Server.
See also **integrated security**.

Storage Group
(MS Exchange 2000 specific term)
A collection of databases (stores) that share a common set of transaction log files. Exchange manages each storage group with a separate server process.

StorageTek ACS library
(StorageTek specific term)
Automated Cartridge System is a library system (also known as Silo) consisting of one Library Management Unit (LMU) and one to 24 Library Storage Modules (LSM) connected to the unit.

switchover
See **failover**

Sybase Backup Server API *(Sybase specific term)*
An industry-standard interface developed for the exchange of backup and recovery information between a Sybase SQL Server and a backup solution like Data Protector.

Sybase SQL Server *(Sybase specific term)*
The server in the Sybase “client-server” architecture. Sybase SQL Server manages multiple databases and multiple users, keeps track of the actual location of data on disks, maintains mapping of logical data description to physical data storage, and maintains data and procedure caches in memory.

Symmetrix Agent (SYMA) *(EMC Symmetrix specific term)*
The Data Protector software module that prepares the EMC Symmetrix environment for backup and restore operations.

38.819

Symmetrix Application Programming Interface (SYMAPI) *(EMC Symmetrix specific term)*
A linkable library of functions that can interface with EMC Symmetrix units attached to the Data Protector clients. Provided by EMC.

Symmetrix CLI Database File *(EMC Symmetrix specific term)*
The EMC Symmetrix database file that stores EMC Symmetrix configuration data on each system with a configured EMC Symmetrix ICDA and installed SYMCLI.

Symmetrix Command-Line Interface (SYMCLI) *(EMC Symmetrix specific term)*
An application written using the Symmetrix Application Programming Interface (SYMAPI) that retrieves data from an EMC Symmetrix unit using special low-level SCSI commands. The SYMCLI allows you to run commands on the client to obtain configuration, status, and performance data from the EMC Symmetrix units attached to clients that are running in an open systems environment.

System Backup to Tape *(Oracle specific term)*
An Oracle interface that handles the actions required to load, label, and unload correct backup devices when Oracle issues a backup or restore request.

system databases *(Sybase specific term)*
The four system databases on a newly installed Sybase SQL Server are the:

- master database (master)
- temporary database (tempdb)
- system procedure database (sybsystemprocs)
- model database (model).

system disk
A system disk is a disk containing operating system files. Microsoft terminology defines the system disk as a disk containing the files required for initial step of boot process.

system partition
A system partition is a partition containing operating system files. Microsoft terminology defines a system partition as a partition containing the files required for initial step of boot process.

System State *(Windows specific term)*
The System State data comprises the Registry, COM+ Class Registration database, system startup files, and the Certificate Services database (if the server is a certificate server). If the

18.818

server is a domain controller, Active Directory directory services and the Sysvol directory are also contained in the System State data. If the server is running the Cluster service, the System State data also includes resource registry checkpoints and the quorum resource recovery log, which contains the most recent cluster database information.

system volume/disk/partition
A volume/disk/partition containing operating system files. Microsoft terminology defines the system volume/disk/partition as the volume/disk/partition containing files required for the initial step of the boot process.

SysVol *(Windows specific term)*
A shared directory that stores the server copy of the domain's public files, which are replicated among all domain controllers in the domain.

tablespace
A part of a database structure. Each database is logically divided into one or more tablespaces. Each tablespace has data files or raw volumes exclusively associated with it.

tapeless backup *(HP StorageWorks Disk Array XP specific term)*
See **ZDB disk backup**.

target database *(Oracle specific term)*
In RMAN, the target database is the database that you are backing up or restoring.

target (R2) device *(EMC Symmetrix specific term)*
An EMC Symmetrix device that participates in SRDF operations with a source (R1) device. It resides in the remote EMC Symmetrix unit. It is paired with a source (R1) device in the local EMC Symmetrix unit and receives all write data from its mirrored pair. This device is not accessed by user applications during normal I/O operations. An R2 device must be assigned to an RDF2 group type.
See also **source (R1) device**

target medium
When copying media, the target medium is the medium to which data is copied.

target system *(Disaster Recovery specific term)*
A system after a computer disaster has occurred. The target system is typically in a non-bootable state and the goal of disaster recovery is to restore this system to the original system configuration. The difference between a crashed system and a target system is that a target system has all faulty hardware replaced.

Glossary

Terminal Services *(Windows specific term)*
Windows Terminal Services provide a multi-session environment that allows clients to access a virtual Windows desktop session and Windows-based programs running on the server.

thread *(MS SQL Server 7.0/2000 specific term)*
An executable entity that belongs to only one process. It comprises a program counter, a user-mode stack, a kernel-mode stack, and a set of register values. Several threads can run at a time within one process.

TimeFinder *(EMC Symmetrix specific term)*
A business continuation process that creates an instant copy of single or multiple Symmetrix Logical Devices (SLDs). The instant copy is created on specially preconfigured SLDs called BCVs and is accessible via a separate device address to the system(s).

TLU
Tape Library Unit.

TNSNAMES.ORA *(Oracle and SAP R/3 specific term)*
A network configuration file that contains connect descriptors mapped to service names. The file may be maintained centrally or locally, for use by all or individual clients.

transaction
A mechanism for ensuring that a set of actions is treated as a single unit of work. Databases use transactions to keep track of database changes.

transaction backup
Transaction backups generally use fewer resources than database backups, so they can be created more frequently than database backups. By applying transaction backups, you can recover the database to a specific point in time prior to when a problem occurred.

transaction backup *(Sybase and SQL specific term)*
A backup of the transaction log providing a record of changes made since the last full or transaction backup.

transaction log backup
Transaction log backups generally use fewer resources than database backups so they can be created more frequently than database backups. By applying transaction log backups, you can recover the database to a specific point in time.

transaction log files *(MS Exchange and Lotus Domino Server specific term)*
Files in which changes made to a database are recorded.

transaction logs *(Data Protector specific term)*
Keeps track of IDB changes. The

18.826

archiving of transaction logs should be enabled to prevent you from losing the transaction log files that are created after the last IDB backup and are necessary for IDB recovery.

transaction log table *(Sybase specific term)*
A system table in which all changes to the database are automatically recorded.

TSANDS.CFG file *(Novell NetWare specific term)*
A file that allows you to specify the names of containers where you want backups to begin. It is text file located in the SYS:SYSTEM\TSA directory on the server where TSANDS.NLM is loaded.

unattended operation
See **lights-out operation**.

user account
You can use Data Protector only if you have a Data Protector user account, which restricts unauthorized access to Data Protector and to backed up data. Data Protector administrators create this account specifying a user logon name, the systems from which the user can log on, and a Data Protector user group membership. This is checked whenever the user starts the Data Protector user interface or performs specific tasks.

user disk quotas
NTFS quota-management support enables an enhanced tracking mechanism and control over disk space usage on shared storage volumes. Data Protector backs up user disk quotas on the whole system and for all configured users at a time.

user group
Each Data Protector user is member of a User Group. Each User Group has a set of user rights that are given to every user in that User Group. The number of User Groups with their associated user rights can be defined as desired. Data Protector provides three default user groups: admin, operator, and user.

user profile *(Windows specific term)*
Configuration information retained on a user basis. This information includes desktop settings, screen colors, network connections, and so on. When the user logs on, the user profile is loaded and the Windows environment is set accordingly.

user rights
User rights or access rights are the permissions needed to perform specific Data Protector tasks. Configuring a backup, starting a backup session, or starting a restore session are typical user rights. Users have the access rights of the user group to which they belong.

vaulting media
The process of storing media to a safe and remote place. The media are brought back to the data center when they are needed for restore or are ready for reuse in subsequent backups. The vaulting procedure depends on your company's backup strategy and policies for data protection/reliability.

VBFS *(OmniStorage specific term)*
A Very Big File System is an extension of the standard HP-UX file system on HP-UX 9.x. It is mounted to a directory the same way as any HP-UX file system. In a VBFS, only the superblock, the inode and the 'extended attribute' information remain permanently on the hard disk and are never migrated.
See also **MFS**.

verify
A function that lets you check whether the Data Protector data on a specified medium is readable. Additionally, consistency within each block can be checked if the backup was performed with the cyclic redundancy check (CRC) option ON.

Virtual Device Interface *(MS SQL Server 7.0/2000 specific term)*
This is a SQL Server 7.0/2000 programming interface that allows fast backup and restore of large databases.

virtual disk *(HP StorageWorks Enterprise Virtual Array specific term)*
A unit of storage allocated from an HP StorageWorks Enterprise Virtual Array storage pool. Virtual disks are the entities that are replicated using the HP StorageWorks Enterprise Virtual Array snapshot functionality.
See also **original unit** and **replica unit.**

virtual server
A virtual machine in a cluster environment defined in a domain by a network IP name and address. Its address is cached by the cluster software and mapped to the cluster node that is currently running the virtual server resources. This way all requests for a particular virtual server are cached by a specific cluster node.

volser *(ADIC and STK specific term)*
A VOLume SERial number is a label on the medium to identify the physical tape used in very large libraries. A volser is a naming convention specific to ADIC/ GRAU and StorageTek devices.

volume group
A unit of data storage in an LVM system. A volume group can consist of one or more physical volumes. There can be more than one volume group on the system.

18.824

volume mountpoint (*Windows specific term*)
An empty directory on a volume that can be used to mount another volume. The volume mount point acts as a gateway to the target volume. Provided that the volume is mounted, users and applications can refer to the data on the mounted volume by the full (merged) filesystem path as if both volumes are one integral part.

Volume Shadow Copy service
See **Microsoft Volume Shadow Copy service.**

VPO
See **OVO**.

VSS
See **Microsoft Volume Shadow Copy service.**

VxFS
Veritas Journal Filesystem.

VxVM (Veritas Volume Manager)
A Veritas Volume Manager is a system for managing disk space on Solaris platforms. A VxVM system consists of an arbitrary group of one or more physical volumes organized into logical disk groups.

Wake ONLAN
Remote power-up support for systems running in power-save mode from some other system on the same LAN.

Web reporting
The Data Protector functionality that allows you to view reports on backup status and Data Protector configuration using the Web interface.

Windows CONFIGURATION backup
Data Protector allows you to back up Windows CONFIGURATION, including Windows Registry, user profiles, Event Logs, and WINS and DHCP server data (if configured on a system) in one step.

Windows Registry
A database repository about a computer's configuration.

WINS server A system running Windows Internet Name Service software that resolves Windows networking computer names to IP addresses. Data Protector can back up WINS server data as part of the Windows configuration.

writer
(MS VSS specific term)
A process that initiates change of data on the original volume. Writers are typically applications or system services

that write persistent information on a volume. Writers also participate in the shadow copy synchronization process by assuring data consistency.

XBSA interface *(Informix specific term)*
The onbar utility and Data Protector communicate with each other through the X/Open Backup Specification Services Programmer's Interface (XBSA).

XCopy engine *(direct backup specific term)*
A SCSI-3 copy command that allows you to copy data from a storage device having a SCSI source address to a backup device having a SCSI destination address, thus enabling direct backup. The data flows from a source device (either block or streaming, that is, disk or tape) to the destination device (either block or streaming) through XCopy. This releases the controlling server of reading the data from the storage device into memory and then writing the information to the destination device. *See also* **direct backup**.

ZDB
See **zero downtime backup (ZDB)**.

ZDB disk backup *(ZDB specific term)*
The basic concept of ZDB disk backup is to create a copy of data from the original storage at a specific point-in-time, and keep this copy of data in the disk array in the replica storage version selected from or created in the replica storage pool. Data in the replica storage version is not moved to standard backup media. The data backed up utilizing the ZDB disk backup functionality can be either restored by utilizing the instant recovery process or used for data mining and similar purposes.
See also **zero downtime backup (ZDB)**, **ZDB tape backup**, **ZDB disk/tape backup**, **instant recovery**, and **replica storage pool**.

ZDB disk/tape backup *(ZDB specific term)*
The basic concept of ZDB disk/tape backup is to create a copy of data from the original storage at a specific point-in-time, and keep this copy of data in the replica storage version. The copy of data in the replica storage version is additionally used for a backup to a standard backup medium, typically a tape. The data backed up using the ZDB disk/tape backup can be restored using the instant recovery or the standard Data Protector restore procedure. It can also be used for data mining and similar purposes.
See also **zero downtime backup (ZDB)**, **ZDB disk backup**, **ZDB tape backup**, **instant recovery**, and **replica storage pool**.

18.812

ZDB part of the IDB *(ZDB specific term)*
A part of the IDB, storing ZDB related information such as original and replica storage versions, security information and other. The ZDB part of the IDB is used for ZDB, instant recovery, and split mirror restore.
See also **zero downtime backup (ZDB)**.

ZDB tape backup *(ZDB specific term)*
The basic concept of ZDB tape backup is to create a copy of data from the original storage at a specific point-in-time, and use this copy of data in the replica storage version for a backup to a standard backup medium, typically a tape. After the backup is complete, the data in the replica storage version may be overwritten. Instant recovery is not possible from such a backup, the data must be restored following the standard Data Protector restore procedure.
See also **zero downtime backup (ZDB)**, **ZDB disk backup**, **instant recovery**, **ZDB disk/tape backup**, and **replica storage pool**.

zero downtime backup (ZDB)
A backup process utilizing data replication technologies (the split mirror and snapshot technologies) to minimize the backup window for the application system; typically to few minutes. With this technique, application database downtime (offline backup) or backup mode (online backup) is limited to the very short time it takes to split the mirror disks or to create or reuse snapshots. The application is then returned to normal operation, while the data in the replica storage version is either backed up by streaming the data to tape (ZDB tape backup) or kept in the replica storage pool (ZDB disk backup) for the instant recovery or other purposes or both (ZDB disk/tape backup).
See also **ZDB disk backup**, **ZDB tape backup**, **ZDB disk/tape backup**, and **instant recovery**.

18.811

Index

A

- aborting
  - all sessions, 623
  - backup session during the size determination, 312
  - running sessions, 312
  - sessions, elapsed time, 625
  - sessions, using ID, 624
  - user right, 84
- access points
  - Enterprise Event ID, 650
  - Generic Event ID, 651
  - graphical user interface (GUI), 652
  - log files, 652
  - SNMP traps, 649
  - SNMP traps format, 651
  - Specific Event ID, 651
  - system and management applications, 649
  - variables, 651
  - Windows Application Log, 652
- access rights
  - for Data Protector users, 83
- accessing
  - Event Log functionality, 356
  - monitoring functionality, 309
  - notification functionality, 342
  - reporting functionality, 315
  - Web reporting interface, 354
  - Web reporting interface, restricting, 354
- accessing in GUI
  - Event Log, 356
- activating
  - barcode reader support, figure, 67
  - barcode support, 66
  - Cartridge Memory support, 68
  - Cartridge Memory support, figure, 69
- active directory restore, 283
- adding
  - device files, 27
  - library devices, 30
  - magazine devices, 34
  - media to a media pool, 107
  - MoM Administrator, 364
  - multiple reports to the report group, 337
  - reports to a report groups, 335
  - standalone devices, 23
  - unused media, 107
  - unused media to media pool, 107
  - used media, 107
  - used media to a media pool, 107
  - user groups, 88
  - users, 90
- ADIC/GRAU AML, 655
- advanced options
  - setting, defining lock name, figure, 54
- Alarm notification, 346
- allocation policy, media, 103
  - format first policy, 104
  - loose, 103
  - strict, 103
- Allow Fallback, object specific option, 243
- appending backups to media, 117
- Application Response Measurement, 646
- applications
  - cluster-aware, 616
  - system and management, 649
- architecture
  - IDB, 384
- ARM integration, 646
- ASCII report format, 329
- ASR, 480
- Assisted Manual Disaster Recovery
  - limitations, Windows, 451
  - preparation, Windows, 451
  - procedure, Windows, 456
  - Windows system, 450
- ATS configuration file
  - creating, 59
- autoconfiguring devices, SAN, 50
- autoloader
  - configuration, 29
- automated media copying, 145
- Automated System Recovery, 480
  - ASR diskettes, 485
  - ASR set, 483
  - Limitations, 482
  - Preparation, 483
  - Recovery, 486
  - Requirements, 481
- Automated System Recovery set, 483
- automatic drive cleaning, 62
- automating
  - restart of failed sessions, 621
- auxiliary disk
  - creating, 509

B

- backing up
  - clients using disk discovery, 163

cluster (MC/SG), 638
cluster (MSCS), 619
CONFIGURATION, 176
DHCP Server, 178
direct backup environment, 204
disk image, UNIX, 166
disk image, Windows, 190
disks, using NFS, 164
DNS Server, 178
event logs, 183
event logs, Windows, 182
IDB, 398
MC/ServiceGuard local disks, 639
MC/ServiceGuard shared disks, 639
Microsoft Cluster Server local disks, 620
Microsoft Cluster Server shared disks, 620
NetWare Directory Services, NDS, 199
Novell NetWare Cluster local disks, 642
Novell NetWare Cluster shared disks, 643
Novell NetWare filesystems, 194
OpenVMS filesystems, 201
rawdisk, UNIX, 166
rawdisk, Windows, 190
shared Windows disks, 185
System State, 176
UNIX filesystems, 161
user disk quotas, 183
Veritas Cluster local disks, 640
Veritas Cluster shared disks, 641
VxFS, A-3
Windows 2000/XP services, 179
Windows clients, disk discovery, 183
Windows CONFIGURATION, 173
Windows filesystems, 168
Windows Registry, 177
Windows user profiles, 182
WINS Server, 178
backup
aborting session during the size determination, 312
cluster, 619, 638
configuring, 153
creating consistent, 444
failed, managing, 263
full, 155
full or incremental, 213
group specifications, 222
incremental, 155
list options, 236
managing cluster-aware, 620
modifying schedule, 210
ownership, 84
predefined, 209
protection expiration, 586
recurring, 210
restarting failed, 266, 311
right to start, 83
scheduling tips, 212
session concepts, 4
skipping, 211
templates, 216
temporary disabling, cluster environment, 626
troubleshooting, 580
unattended, 207
VSS filesystem, 170
with stacker devices, 36
backup commands
pre- and post-exec, UNIX, 257
pre- and post-exec, Windows, 251
backup devices
adding library, 30
adding standalone, 23
autoloaders, 29
block size, 76, 79
concurrency, 76
concurrency and streaming, 76
configuring, 17
configuring chains, 24
configuring files, 27
configuring magazines, 34
configuring manually, 52
configuring stacker, 35
configuring standalone, 23
disabling, 70
disabling, figure, 71
file device, 26
libraries with multiple systems, 32
library, 29
locking, 46, 74
locking for drives, table, 53
locking mechanism, 46
preparing configuration, 20
relation to backup specifications and media pools, scheme, 22
removing, 72
renaming, 73
restarting, 70

segment size, 78
shared in the SAN, 44
specifying type and name, figure, 24
streaming, 76
used by multiple applications, locking, 46
using, 17
backup environment
setting up, tasks, 15
backup failure
preventing, 264
backup files of size
object specific option, 243
backup objects, 153
selecting, 168
backup options, 225
catalog, 230
configuring, 211
device, 249
frequently used, 227
list, 236
load balancing, 232
log level, 231
logging files to database, 245
ownership, 235
protection, 228
reconnecting broken connections, 239
scheme, 227
backup POSIX hard links as files
object specific option, 243
backup session
concepts, 4
backup specification
checking, 608
concepts, 153
creating, 154
creating for recovery, 509
example, 155
groups, 222
multiple, 155
options, 236
pre- and post-exec commands, 250
right to save, 84
saving groups, 223
backup specifications
relation to devices and media pools, scheme, 22
Backup Specifications reports, 318
backup templates
using for configuring backup, 218
backup types, 213
backup, troubleshooting
mount request for a library device, 582
mount request for a standalone device, 581
protection expiration, 586
starting interactive sessions, 585
starting scheduled sessions, 584
unexpected mounted system detected, 583
Barcode Scan option, 129
barcode support, activating, 66
BDACC
environment variable, 255
block size
backup device options, 79
changing, 79
changing, example, A-51
boot partition, 439
Disk Delivery Disaster Recovery, 459
Enhanced Disaster Recovery, 463
bootable installation CD
disaster recovery, 452
broadcast message send method
notifications, 349
reports, 332
buffer size
Disk Agent, 79
bulk eject of media, 137
busy drive handling, 65

C

calculation of a media condition, changing, 134
media conditions property page, figure, 134
Cartridge Memory
activating support, 68
activating support, figure, 69
data initialization, 109
list, 130, 132
list for specific slot, figure, 130
reformat, 109
reformat for specific slot, figure, 110
catalog
backup, 230
Catalog Database, 385
catalog from media, importing, 114
catalog protection, 230, 388
object specific option, 244
CDB, 600
cell
backup devices, 3
Cell Manager, 3

  concepts, 3
  Disk Agent, 3
  exporting, 377
  importing, 377
  importing, MoM, 363
  Media Agent, 3
  monitoring simultaneously, 314
  moving clients, 378
  reports on multiple, 315
  setting up MoM Manager, 363
Cell Manager
  checking, 607
  concepts, 3
  configuring package, MC/ServiceGuard, 633
  configuring, MC/ServiceGuard, 628
  disaster recovery methods, UNIX, 437
  installation, troubleshooting, 589
  Manual Disaster Recovery, UNIX, 512
  Manual Disaster Recovery, Windows, 487
  Microsoft Cluster Server, 618
  moving the IDB, 412
  on MC/ServiceGuard, 627
  One Button Disaster Recovery, Windows NT, 472
  when not accessible, 590
Centralized Media Management Database (CMMDB)
  configuring, 368
  configuring on the client cell, 370
  configuring on the MoM Manager, 369
  overview
  scheme, 367
certificate services restore, 284
changing
  backup owner, 235
  block size, 79
  block size, example, A-51
  device concurrency, 76
  device type, 578
  encoding, GUI, 580
  message level, 309, 312
  password for the Web reporting interface, 354
  user account, Windows 2000/XP, 188
  user account, Windows NT, 187
  user group rights, 93
changing the start date
  editing backup schedule, 210
checking
  backup specification, 608
  Cell Manager, 607
  failed backups, 263
  IDB consistency, 411
  IDB consistency, manually, 411
  IDB size, 410
  installation, 608
  log files, 608
  media condition, 131
  status of daemons, 571
  TCP/IP setup, 567
checking and maintenance mechanism, 605
cleaning
  drive, 61
  tape, 61
clearing a schedule
  editing backup schedule, 210
CLI. *See* command-line interface
clients
  Assisted Manual Disaster Recovery, Windows, 450
  disaster recovery methods, HP-UX and Solaris, 437
  Disk Delivery Disaster Recovery, UNIX client, 507
  MC/ServiceGuard, 637
  Microsoft Cluster Server, 618
  moving among cells, 378
  One Button Disaster Recovery, Windows NT, 472
cluster
  aborting all running sessions, 623
  aborting running sessions, elapsed time, 625
  aborting running sessions, using ID, 624
  advanced backup specification options, 622
  automating restart of failed session, 621
  backup, 619, 638
  concepts, 613
  disabling backup sessions, 626
  failover, 615
  failover of Data Protector, 621
  failover of other application, 623
  group (MSCS), 615
  heartbeat, 614
  managing backups, 620
  MC/ServiceGuard, 613, 627
  Microsoft Cluster Server, 613, 617

Novell NetWare Cluster Services, 642
omniclus command, 623
package (MC/SG, Veritas Cluster), 615
primary node, 615
secondary node, 615
switchover, 615
Veritas Cluster, 613
virtual server, 615
cluster-aware applications, 616
cluster-aware backups, 620
CM, MA and DA in the DMZ, 540
CMMDB
See Centralized Media Management Database
command-line interface (CLI), 11
commands
pre- and post-exec, 250
pre- and post-exec examples, A-20
pre- and post-exec, UNIX, 257
pre- and post-exec, Windows, 251
communication, troubleshooting, 565
client fails, 567
HOST file resolution problem, 567
host name resolution problems, 565
concurrency
advanced options dialog box, figure, 77
changing, device, 76
device backup option, 249
device properties dialog box, figure, 78
condition
of a media pool, 132
of a media, and device error, 132
of a media, changing calculation, 134
of a media, checking, 131
of a media, checking using Cartridge Memory List, 132
of a media, fair, 133
of a media, good, 133
of a media, influence on how media are selected for backup, 131
of a media, influencing factors, 132
of a media, poor, 133
of a media, property page, figure, 134
condition factors for media, 105, 132
age of a medium, 105
maximum number of overwrites, 132
medium valid for, 132
usage of a medium, 106
CONFIGURATION
backing up, 173, 176
restoring Windows, 280
Windows 2000/XP, 174
Windows NT, 173
configuration
of user rights, 83
Configuration reports, 320
configuring
automatic drive cleaning, 62
automatically, devices, 50
backup devices, 17
backup devices for direct backup, 38
backups, 153
barcode support, 66
Cell Manager package, MC/ServiceGuard, 633
Cell Manager, MC/ServiceGuard, 628
cleaning tape slot, 63
cluster-aware client, MC/SG, 637
cluster-aware client, MSCS, 618
CMMDB, 368
CMMDB on the client cell, 370
CMMDB on the MoM Manager, 369
device chains, 24
device files, 27
Device Flow report, using CLI, example, 341
device streaming, 76
devices, automatically, 50
devices, manually, 52
drives, 52, 59
drives, library, 32
DSI integration, 644
file devices, 26
firewall environment, 528
floating drive, table, 59
floating drives, 59
IDB, 388
libraries with multiple systems, 32
library devices, 29, 30
library for mixed media, 37
library robotics in a cluster, 52
libtab files, manually, 56
magazine devices, 34
Manager-of-Managers, 362
ManageX integration, 648
MC/SG integration in the SAN, 58
media pool, 102
Media Statistics report, using CLI, example, 340

  new Microsoft Exchange Profile, 348
  notifications, 342, 351
  notifications on the Web, 353
  notifications, using Web reporting interface, 355
  report groups, 335
  report groups, using Web reporting interface, 355
  reports, 335
  reports on the Web, 353
  SCSI-II library devices, 29
  separate media pools for the different drives, 43
  Session Flow report, using CLI, example, 340
  SNMP traps, Windows 2000, 333, 350
  SNMP traps, Windows NT, 333, 350
  stacker devices, 35
  stacker devices, example, 35
  standalone devices, 23
  static drive, table, 59
  static drives, 60
  the library robotics, manually, 51
  user rights, 83
  users in MoM, 379
  vaults, 141
  web user password, 354
configuring firewall environment
  CM, MA and DA in the DMZ, 540
  DA and MA in the DMZ, 535
  DA in the DMZ, 538
  examples, 535
  limiting port range, 528
  OB2BAR and MA in the DMZ, 543
  overview, 528
  port usage in Data Protector, 531
configuring the IDB
  backup specification, 398
  catalog protection, 388
  directories, location, 391
  disk space, allocating, 388, 389
  growth factors, 388, 389
  logging level, 388
  notifications, 400
  preparing recovery, 390
  procedure, 388
  recovery file, creating a copy, 395
  reports, 400
  robustness considerations, 390
connection reset by peer
  troubleshooting, 567
consecutive backups
  running, 212
Context List, 9
conventions, xxi
copying
  Data Protector Java programs to the Web Server, 354
  media, 143
  media, automated, 145
corruption
  IDB, 419
CRC check
  device backup option, 249
creating
  ATS configuration file, 59
  auxiliary disk, 509
  backup specification, 154, 509
  backup specification, example, 155
  consistent and relevant backup, 444
critical volumes, 440
CRS debug
  in MS cluster environment, 558
  on Windows, 557
customizing
  notifications, 342
  reports, 315
  the information about the media, 149

D

DA and MA in the DMZ, 535
DA in the DMZ, 538
daemons
  checking the status of, 571
  starting, 571
  starting problems, 571
  stopping, 571
  troubleshooting, 569
daily full backup
  predefined backup schedules, 209
daily intensive backup
  predefined backup schedules, 209
DailyMaintenanceTime global option, 524
data files missing, 596
data protection, 228
  setting, 122
Data Protector internal database
  See IDB
Data Protector Java programs

copying to the Web server, 354
Data Source Integration, 644
database
  See IDB
  backup problems, 598
  import problems, 598
database configuration
  See IDB configuration
database consistency
  See IDB consistency
database directories
  See IDB directories
Database Purge Needed notification, 346
Database Space Low notification, 346
database troubleshooting
  See IDB troubleshooting
DATALIST, definition, 252
DCDirAllocation global option, 523
deactivating centralized licensing, 376
debug syntax, 555
debug.log, 551
debugging
  CRS debug, MS cluster environment, 558
  CRS debug, Windows, 557
  debug syntax, 555
  INET debug, Unix, 557
  INET debug, Windows, 557
  sample, 558
  trace file name, 556
  troubleshooting, 553
default object options, 225
defining
  lock name, 53
deleting
  Event Log Viewer contents, 356
  user groups, 88
  users, 90
density
  setting the same, 42
description
  backup option, 237
description of media, 127
  media label, 127
  modifying using Cartridge Memory, 127
destination
  restoring files to another client, 299
Detail Catalog Binary Files, 386
detect NTFS hardlinks
  object specific option, 244
detection of write-protected media, 147
device
  ejecting a medium from, 137
  entering media into, 136
  error and media condition, 132
  scanning, 129
device backup options, 249
Device Error report, 336
device files, 20
Device Flow report
  configuring, example, 341
devices
  configuration right, 83
  open problem, troubleshooting, 575
  troubleshooting, 574
devices, autoconfiguring, 50
DHCP Server
  backing up, 178
  NT CONFIGURATION, 173
  restoring, 286
dIDB
  growth, 389
Direct Access mechanism
  enabling, 55
  selecting, figure, 56
direct backup
  backup device autodetection, 39
  configuring backup devices, 38
  limitations, 205
  prerequisites, 204
  restoring, 205
  XCopy engine, 39, 204
direct backup environment
  backing up, 204
direct library access, 47
directory junctions, 170, 171
dirty drive detection, enabling, 63
dirty flag, 444
disabling
  automatic check, IDB, 399
  backup device, 70
  backup device, figure, 71
  sessions, cluster environment, 626
disabling a schedule
  editing backup schedule, 210
disaster, 439
disaster recovery
  ASR, 480
  Automated System Recovery set, 483
  concepts, 437
  creating backup specification, 509
  dirty flag, 444

Index

Disk Delivery method, 459, 507
Enhanced Automated method, 463
logging on after, 514
One Button method, 472
overview, 437
Phase 0, 442
Phase 1, 442
Phase 2, 442
Phase 3, 442
planning, 443
preparing, 443
preparing for, 443
troubleshooting, Windows NT, 514, 515
updating SRD, 445
disaster recovery operating system (DR OS), 440
disaster recovery process overview
plan, 443
prepare, 443
recover, 444
disk
restoring disk image (rawdisk), 273
Disk Agent
buffer size, 79
concepts, 3
device streaming and concurrency, 76
Disk Delivery Disaster Recovery
client, Windows, 459
limitations, UNIX client, 507
preparation, UNIX client, 507
preparation, Windows client, 460
procedure, UNIX client, 510
procedure, Windows client, 461
recovered partitions, 459
troubleshooting, Windows NT, 515
UNIX client, 507
disk discovery
Novell NetWare backup, 198
UNIX client backup, 163
when to use, 163
Windows client backup, 183
disk image
backing up, 166, 190
restoring, 273
setting options, 240
disk space
allocating, 388, 389
considerations, 397
display statistical information, 295
distributing
MoM configuration, 378
DNS
troubleshooting, 586
DNS Server
backing up, 178
restoring, 286
DNSServerDatabase, 174
do not preserve access time
object specific option, 244
do not use archive attribute
object specific option, 244
DR OS, 440
drive
backup devices locking for, table, 53
cleaning, 61
configuring, 52, 59
configuring separate media pools for the different, 43
entering media, 136
floating, 59
handling busy, 65
inserting media, 136
SCSI address, 29
static, 59
testing cleaning configuration, 63
using several types in a library, 42
drive cleaning
automatic drive cleaning configured with Data Protector, 61
conditions for automatic cleaning, 62
configuring, 61
configuring automatic, 62
configuring cleaning tape slot, 63
library specific built-in cleaning mechanism, 61
limitations, 61
manual cleaning, 61
testing, 63
drive index
library devices, 29
to SCSI address mapping, scheme, 30
DSI integration
configuring, 644

E

editing backup schedule, 210
ejecting a medium from a device, 137
bulk eject of media, 137
procedure, 137

18.802

scheduled eject, 138
elapsed session time, 625
e-mail send method
  notifications, 348
  notifications, creating a new Microsoft Exchange Profile, 348
  reports, 331
enabling
  Direct Access mechanism, 55
  dirty drive detection, 63
encode
  object specific option, 245
encoding
  changing, GUI, 580
End of Session notification, 347
END_USER_ARCHIVE, 222
Enhanced Automated Disaster Recovery
  client, Windows NT/2000, 463
  disaster recovery CD, 468
  disaster recovery CD ISO image, 463, 468
  DR image, 466
  DR OS image file, 463
  Phase 1 Startup file (P1S), 468
Enhanced Disaster Recovery
  limitations, Windows NT/2000 client, 465, 474
  preparation, Windows NT/2000 client, 466, 475
  procedure, Windows NT/2000 client, 470
  recovered partitions, 463
  troubleshooting, Windows NT, 516
Enterprise Event ID, 650
Event Log, 307, 356
  accessing functionality, 356
  Event Log in GUI, 356
  Event Log message, 357
Event Log message, 357
Event Log send method, notifications, 351
Event Log Viewer
  deleting contents, 356
event logs
  backing up, 183
  backing up Windows, 182
  restoring Windows, 285
EventLog
  NT CONFIGURATION, 173
examples
  changing the block size, A-51
  configuring stacker devices, 35
  creating a Device Flow report, using CLI, 341
  creating a Media Statistics report, using CLI, 340
  creating a Session Flow report, using CLI, 340
  debugging, 558
  Health Check Failed notification, 653
  last night's backup results, 654
  libtab file, 57
  pre- and post-exec commands, A-20
  report groups, 335
  scheduled eject of media, A-14
  user configuration, 94
  verifying Data Protector processes, 653
examples of configuring firewall environments, 535
Exchanger control device
  troubleshooting, 574
exporting
  cells, 377
  copies of media, 144
  media from Data Protector, 125
  media, procedure, 125
Extended List of Media, report, 324
external send method
  notifications, 351
  reports, 334

F

failed backup
  checking, 263
  managing, 263
failed sessions
  automating restart, 621
failover, 615, 621, 623
file devices
  configuring, 26
  handling mount prompts, 27
  specifying pathname for, figure, 28
file name tracing, 556
file ownership
  restoring, 288
FileReplicationService, 174
files
  restoring, 299
  restoring UNIX regular, 276
  restoring Windows, 277
filesystems
  backing up Novell NetWare, 194

18.801

Index

backing up OpenVMS, 201
backing up UNIX, 161
backing up Windows, 168
restore limitations, 278
restoring Novell NetWare, 287
restoring OpenVMS, 291
firewall configurations
CM, MA and DA in the DMZ, 540
DA and MA in the DMZ, 535
DA in the DMZ, 538
examples, 535
OB2BAR and MA in the DMZ, 543
firewall environment
configuring, 528
limiting port range, 528
overview, 528
port usage in Data Protector, 531
firewall support, 528
examples, 535
limiting port range, 528
port usage in Data Protector, 531
floating drives
configuring, 59
format first policy, 104
formats of media, 111
ANSI label, 111
cpio, 111
filesystem, 111
foreign Data Protector (from another cell), 111
OmniBack I, 111
OmniStorage, 111
protected media, 111
tar, 111
unprotected media, 111
written with compression, 111
written without compression, 111
formats of reports
ASCII report format, 329
HTML report format, 329
short report format, 330
tab report format, 330
formatting
stacker devices, 35
formatting media, 108
Cartridge Memory, data initialization, 109
Cartridge Memory, reformat, 109
in a magazine, 110
in full magazine, 110
media format categories, 111
procedure, 109
recognizing other formats, 108, 111
single medium in a magazine, 110
used by other applications, 109
with padding blocks, 108
fortnightly full backup
predefined backup schedules, 210
free pool, 102
full backups, 155
definition, 213
selecting, 215
troubleshooting, 580

G

generating
reports, using omnirpt command, 339
reports, using Web reporting interface, 355
Generic Event ID, 651
global options
overview, 523
usage, 523
variables, 523
global options file, 523
graphical user interface (GUI), 7
access points, 652
changing encoding, 580
Context List, 9
Microsoft Management Console
Navigation Tabs, 10
online Help, 12
Results Area, 10
Results Tab, 10
running problems, 595
Scoping Pane, 10
starting, UNIX, 7
starting, Windows, 7
troubleshooting, 562, 590
GRAU, 655
group (MSCS), 615
grouping backup specifications, 222
GUI. *See* graphical user interface

H

handling busy drive, 65
Health Check Failed notification, 347
heartbeat of the cluster, 614
holiday, skipping backup, 211
HOST file resolution problem, 567
host name resolution problems, 565
hosting system, 440

Index

HP-UX and Solaris client
  disaster recovery methods, 437
HTML report format, 329

I

IDB
  architecture, 384
  backing up, 398
  backing up Windows Registry, 177
  catalog protection, 388
  checking consistency, 411
  checking size, 410
  complete recovery, 421
  configuring, 388
  configuring backup, 398
  corrupted, 421, 422, 423
  corruption, 419
  directories, location, 391
  disabling automatic check, 399
  disk space, 389, 397
  extending size, 408
  fnames.dat file, 409
  growth, 388
  growth, reducing, 405
  logging level, 388
  maintaining, 402
  managing, 383
  moving, 412
  notifications, 400, 401
  obrindex.dat file, 426–428
  overview, 383
  parts, 384
  problems, 402, 404
  purging obsolete filenames, 408
  recovering, 390, 417
  recovery file, 395
  reducing size, 406
  report types, 400
  reports, 321
  restoring, 414
  troubleshooting, 592
IDB configuration
  backup specification, 398
  catalog protection, 388
  directories, location, 391
  disk space, allocating, 388, 389
  growth factors, 388, 389
  logging level, 388
  notifications, 400
  preparing recovery, 390
  procedure, 388
  recovery file, creating a copy, 395
  reports, 400
  robustness considerations, 390
IDB consistency
  checking, 411
  checking manually, 411
  disabling automatic check, 399
IDB directories
  location, 391
  recommended location, 393
  relocating, 393
IDB recovery
  complete, 421
  corrupted (missing) DC binary files, 422
  corrupted filename tablespace, 423
  DC directories, creating, relocating, 396
  directories, location, 391
  methods, 417
  obrindex.dat file, 426–428
  preparing, 390
  recovery file, creating a copy, 395
  relocating directories, 393
  robustness considerations, 390
  to a different disk layout, 431
IDB size
  checking, 410
  extending, 408
  reducing, 406
IDB troubleshooting
  application restore sessions, 586
  backup problems, 598
  data files missing, 596
  import problems, 598
  libraries (executables) missing, 595
  MMDB and CDB not synchronized, 600
  performance problems, 599
  temporary directory missing, 597
  upgrade problems, 592
  user interface running, problems, 595
image objects, 166, 190
importing
  a single medium into a magazine device,
    116
  catalog from medium, procedure, 114
  catalog, figure, 115
  cells, MoM, 363, 377
  magazine, figure, 116

Index

the catalog from media, 114
importing media, 113
  in a magazine device, 115
  multiple, figure, 114
  procedure, 113
Inc backups, definition, 213
Incl-9 backups, definition, 214
incremental backups, 155
  Novell NetWare, 196
  selecting, 215
  troubleshooting, 580
indirect library access, 47
individual reports, running, using GUI, 338
INET debug
  on Unix, 557
  on Windows, 557
Inet service
  setting user account, 187
inet.log, 551
informix.log, 552
InitOnLoosePolicy global option, 523
inserting media in drive, 136
installation,checking, 608
installing
  ARM integration, 646
  Cell Manager on MC/ServiceGuard, 628
  Cell Manager on Microsoft Cluster Server, 618
  Cell Manager, troubleshooting, 589
  clients on Veritas Cluster, 640, 642
  clients, troubleshooting, 588
  cluster-aware client, MC/SG, 637
  cluster-aware client, MSCS, 618
  default users, 90
integrations
  ARM, 646
  Cluster Server, 613
  Data Source, 644
  ManageX, 648
  MC/ServiceGuard, 627
  Microsoft Cluster Server, 617
  Novell NetWare Cluster Services, 642
  Veritas Cluster, 640
interactive backup, troubleshooting, 585
internal locking
  logical devices, 74
IS_install.log, 551

K

keep most recent, 297

L

library
  configuring drives, 32
  configuring for mixed media, 37
  configuring with multiple systems, 32
  SCSI address, 29
  using several drive types in, 42
  when missing, 595
library access concepts
  direct, 47
  indirect, 47
library devices
  configuring, 29, 30
  configuring with multiple systems, 32
  drive index, 29
  mount request for, 581, 582
  SCSI ID, 29
  slot number, 29
  troubleshooting, 581, 582
library robotics
  configuring, 52
  configuring in a cluster, 52
libtab file
  configuring, manually, 56
  examples, 57
licensing
  availability, 584
  deactivating centralized, 376
  Manager-of-Managers, 372
  MC/ServiceGuard, 627
  Microsoft Cluster Server, 617
  moving licenses in the MoM, 375
life-cycle of media, 100
  preparing for backup, 100
  retiring, 101
  using for backups, 101
  vaulting to a safe place, 101
limitations
  Assisted Manual Disaster Recovery, Windows, 451
  direct backup, 205
  Disk Delivery Disaster Recovery, UNIX client, 507
  Enhanced Disaster Recovery, Windows NT/2000 client, 465, 474
  Manual Disaster Recovery, UNIX Cell Manager, 512
  OpenVMS backup, 202
  OpenVMS restore, 291

limiting port range, firewall environment, 528
list
  restored files, 295
List of Pools report, 324
load balancing, 232
  backup option, 237
local disks
  backing up MC/ServiceGuard, 639
  backing up Microsoft Cluster Server, 620
  Novell NetWare Cluster backing up, 642
  Veritas Cluster backing up, 640
location of media, 126
lock files during backup
  object specific option, 245
lock name, 74
  defining, 53
  summary of device definitions using, figure, 55
locked files, 296, 302
locking
  backup devices, 74
  devices used by multiple applications, 46
log files, 652
  backing up Windows event, 182
  checking, 608
  contents, troubleshooting, 551
  format, troubleshooting, 550
  location, troubleshooting, 550
  troubleshooting installation, 588
log level, backup, 231
log to file send method
  notifications, 349
  reports, 332
logging
  object specific option, 245
logging level, IDB, 388
logging on
  problems after disaster recovery, 514
logical devices
  internal locking, 74
logical disk drives
  backing up Windows, 168
logical ID
  aborting sessions, 624
login
  user identity, 90
loose media allocation policy, 103

M

magazine devices
  configuring, 34
magazine support, 105
maintaining
  IDB, 402
Manager-of-Managers
  adding MoM Administrator, 364
  centralized licensing, 372
  configuring, 362
  configuring CMMDB, 368
  configuring users, 379
  deactivating centralized licensing, 376
  distributing configuration, 378
  importing cells, 363
  moving clients, 378
  moving licenses, 375
  overview, 361
  setting up MoM Manager, 363
  tasks, 377
Manager-of-Managers (MoM)
  monitoring multiple cells, 314
ManageX integration, 648
  configuring, 648
managing
  failed backup, 263
  IDB, 383
Manual Disaster Recovery
  Cell Manager, UNIX, 512
  Cell Manager, Windows, 487
  drsetup diskettes, 453
  limitations, UNIX Cell Manager, 512
  preparation, UNIX Cell Manager, 512
  procedure, UNIX Cell Manager, 512
MaxBSession global option, 523
MaxMAperSM global option, 523
MC/ServiceGuard
  backing up, 638
  Cell Manager, 627
  Cell Manager package, 633
  clients, 637
  cluster concepts, 613
  in the SAN, 58
  integration, 627
  licensing, 627
MC/SG. *See* MC/ServiceGuard
measuring
  with ARM integration, 646
  with DSI integration, 644
media

Index

adding to pool, 107
adding unused to a media pool, 107
adding used to a media pool, 107
allocation policy, 103, 120
appendable, 104
appendable of incrementals only, 105
appending backups to, 117
condition, 120
condition factors, 105, 132
configuration right, 83
configuring library for mixed, 37
copying, 143
copying, automated, 145
customizing information about, 149
description, 127
detection of write-protected, 147
ejecting from a device, 137
entering into a device, 136
exporting from Data Protector, 125
exporting, procedure, 125
format types, limitations, 148
formatting, 108
formatting in a magazine, 110
header sanity check, troubleshooting, 577
implementing vaulting, 140
importing, 113
importing in a magazine device, 115
information on, figure, 131
inserting in drive, 136
label, 127
labeling, 107
life-cycle, 100
location, 126
magazine support option, 105
management concepts, 99
managing, 97
modifying descriptions, 127
modifying locations, 126
moving and exporting copies, 144
moving to a vault, 141
moving to another pool, 124
non-appendable, 104
overwrites, 106
pre-allocation list, 120
preparing for backup, 100
quality statistics, troubleshooting, 575
recycling, 123
restoring files, 301
restoring from copy, 144
restoring from in a vault, 141
retiring, 101
scanning in a device, 129
scanning in a device using Barcode Scan option, 129
scanning in a device using Cartridge Memory list, 130
scheduled eject of, 138
searching for, 135
searching for, procedure, 135
selecting, 120, 135
selecting for backup, 120
selecting for backup, table, 121
selecting, procedure, 135
setting data protection for, 122
status, 133
troubleshooting, 574
types, 103
usage, 120
usage policy, 104
using for backups, 101
vaulting, 140
vaulting to a safe place, 101
verifying data on, 128
viewing files from, 301
Media Agent, concepts, 3
media management, 97
adding media to a media pool, 107
appending backups to media, 117
checking the condition of a medium, 131
concepts, 99
copying media, 143
creating a media pool, 102
detection of write-protected media, 147
ejecting a medium from a device, 137
entering media into a device, 136
exporting media from Data Protector, 125
formatting media, 108
importing media, 113
media format types, limitations, 148
media life-cycle, 100
modifying media descriptions, 127
modifying media locations, 126
modifying views in, 149
moving media to another pool, 124
overview, 99
recycling media, 123
relationship between media and other components, scheme, 100

scanning media in a device, 129
searching for and selecting a medium, 135
selecting media for backup, 120
setting data protection, 122
using a pre-allocation list of media for backup, 119
vaulting media, 140
verifying data on a medium, 128
Media Management Database, 385
media pool
adding media to, 107
adding unused media to, 107
adding used media to, 107
allocation, 102
concepts, 102
condition factors, 105
condition of, 132
configuration procedure, 102
configuring, 102
configuring separate for the different drives, 43
deallocation, 102
default, 102
description, 103
device backup option, 249
free pool, 102
labeling media, 107
magazine support, 105
media allocation policy, 103
media types, 103
moving media to, 124
name, 103
properties, 103
relation to backup specifications and devices, scheme, 22
status, 133
use free pool option, 104
Media Statistics report
configuring, example, 340
media.log, 551
MediaView global option, 523
merge option, 297
message level, changing, 309, 312
messages, troubleshooting, 561
Microsoft Cluster Server
backing up, 619
Cell Manager, 618
clients, 618
cluster concepts, 613
installing, 618
integration, 613, 617
licensing, 617
Microsoft Exchange Profile
creating new, 348
Microsoft Management Console, 13
MMC. *See* Microsoft Management Console
MMDB, 600
MODE, definition, 252
modifying
backup schedule, 210
media description, procedure, 127
media descriptions, 127
media locations, 126
message level, 312
user group rights, 93
users, 92
views in the media management window, 149
monitoring, 307
aborting running sessions, 312
accessing functionality, 309
cells simultaneously, 314
currently running sessions, 309
finished sessions, 310
mount requests, 310
restarting failed backup, 311
sessions, 309
user right, 84
monthly full backup
predefined backup schedules, 210
mount prompt
handling, 27
mount request
for a library device, 581, 582
for a standalone device, 581
issuing, 310
responding to, 310
user right, 84
Mount Request report, 336
mountpoint configuration file
Novell NetWare, 200
moving
busy files, 295
clients among cells, 378
copies of media, 144
IDB, 412
licenses in the MoM, 375
media to another pool, 124
media using a free pool, 124
users, 92

Index

MSCS *See* Microsoft Cluster Server
multi host support, 32
multiple backup specifications, 155
multiple reports, adding to the report group, 337

N

Name Space information
  restoring, 287
native tape driver, 20
Navigation Tabs, 10
NDS
  adding objects, 199
  backing up, 199
NDS objects
  restoring, 290
NDS scheme
  restoring, 290
NetWare
  restoring filesystems, 287
NetWare Directory Services (NDS)
  backing up, 199
networking, troubleshooting, 565
  client fails, 567
  HOST file resolution problem, 567
  host name resolution problems, 565
NFS (Network Filesystem)
  backing up disks, 164
non-ASCII characters, troubleshooting, 587
notifications, 307
  accessing functionality, 342
  concepts, 342
  configuring, 342, 351
  configuring on the Web, 353
  configuring, using Web reporting interface, 355
  customizing, 342
  explanation of some, 346
  IDB, 401
  input parameters, 342
  list, 343
  send methods, 347
  triggering a report group by, 337, 352
  types, 342
  user rights, 83
notifications scheduled and started by the Data Protector checking and maintenance mechanism, 343
  Database Purge Needed, 346
  Database Space Low, 346
  End of Session, 347
  Health Check Failed, 347
  User Check Failed, 347
notifications triggered when an event occurs, 342
  Alarm, 346
Novell NDS
  restoring, 289
Novell NetWare
  adding NDS objects, 199
  backing up, 194
  backing up filesystems, 194
  backing up NDS, 199
  restoring filesystems, 287
  restoring NDS, 289
Novell NetWare Cluster Services
  integration, 642
NTFS 5.0 filesystem, 170
number of buffers, 79
number of retries
  object specific option, 245

O

OB2BAR and MA in the DMZ, 543
OB2BLKPADDING omnirc variable, 525
OB2CHECKCHANGETIME omnirc variable, 526
OB2DEVSLEEP omnirc variable, 526
OB2ENCODE omnirc variable, 526
Ob2EventLog.txt, 551
OB2INCRDIFFTIME omnirc variable, 526
OB2OEXECOFF omnirc variable, 526
OB2PORTRANGE omnirc variable, 527, 528
OB2PORTRANGESPEC omnirc variable, 527, 529
OB2RECONNECT_ACK omnirc variable, 526
OB2RECONNECT_RETRY omnirc variable, 526
OB2REXECOFF omnirc variable, 526
OB2SHMEM_IPCGLOBAL omnirc variable, 526
OB2VXDIRECT omnirc variable, 527
object
  pre- and post-exec commands, 260
Object IDs, 170
object options, 239
object specific options
  setting, 241
objects
  restore options, 294

omit deleted files, 294
omniclus, 623
omnirc options
  overview, 525
  usage, 525
  variables, 525
omnirc options file, 525
omnirpt
  generating reports with, 339
omniSRDupdate
  post-exec script, 446
  standalone, 446
OmniStorage, restoring, 276
omnisv.log, 551
One Button Disaster Recovery (OBDR)
  procedure, Windows NT/2000, 477
  Windows NT/2000 system, 472
online Help, 12
  troubleshooting, 603
OpenVMS
  backing up, 201
  backing up filesystems, 201
  backup limitations, 202
  restore limitations, 291
  restoring filesystems, 291
options
  advanced backup specification-clustering, 622
  backup, 225
  backup specification, 236
  global, 523
  omnirc, 525
  restore, 294
oracle8.log, 552
original system, 439
OS partition
  Disk Delivery Disaster Recovery, 459
  Enhanced Disaster Recovery, 463
overview
  CMMDB
  configuring SAN, 47
  disaster recovery, 437
  firewall environment, 528
  global options, 523
  IDB, 383
  Manager-of-Managers, 361
  omnirc options, 525
  system and management applications, 649
  types of reports, 317
overwrite option, 295, 297
OWNER, definition, 252
ownership
  backup, 84, 235
  backup option, 237
  changing, 235
  user rights, 84

P

package (MC/SG, Veritas Cluster), 615
parallel restore, 300
performance considerations, A-8
periodic backup
  starting, 209
permissions
  group, 93
  user, 83
planning
  disaster recovery, 443
  scheduling policies, 208
Pools and Media reports, 324
port range
  limiting with the omnirc variables, 528
port usage in Data Protector, firewall environment, 531
  examples, 535
post-backup media copying, 145
post-exec
  backup option, 238
  commands, 250
pre- and post-exec commands
  examples, A-20
  object, 260
  UNIX, 257
  Windows, 251
prealloc list
  device backup option, 249
pre-allocating media, 120
pre-allocation list of media, using for backup, 119
predefined backup schedule, 209
pre-exec
  backup option, 238
  commands, 250, 296, 297
preparing
  Assisted Manual Disaster Recovery, Windows, 451
  backup devices configuration, 20
  Disk Delivery Disaster Recovery, UNIX client, 507

Index

- Disk Delivery Disaster Recovery, Windows client, 460
- Enhanced Disaster Recovery, Windows NT/2000 client, 466, 475
- for disaster recovery, 443
- Manual Disaster Recovery, UNIX Cell Manager, 512
- media for backup, 100

preparing for a disaster recovery, 443
prerequisites
- direct backup, 204

preventing
- backup failure, 264

PREVIEW, definition, 252
primary node, 615
private object
- who can restore, 235

Private, object specific option, 246
privileges
- group, 93
- user, 83

problems
- IDB, 402, 404

procedure
- copying media, 143
- creating a media pool, 102
- disabling device, 70
- ejecting a medium from a device, 137
- entering media into a device, 136
- exporting media, 125
- formatting media, 109
- importing catalog from medium, 114
- importing media, 113
- modifying media description, 127
- modifying media location, 126
- moving media to another pool, 124
- moving media using a free pool, 124
- scanning media in a device, 129
- searching for and selecting media, 135
- verifying data on a medium, 128

processes
- verifying, 653
- which, when, where they run, 573

profiles
- NT CONFIGURATION, 173
- restoring Windows user, 285

protection
- attributes, 296
- backup, 228
- expiration, 586
- object specific option, 246

Public, object specific option, 246
purge.log, 552
purging
- IDB filenames, 408

Q

QuotaInformation, 174

R

rawdisk, 166, 190
- backing up UNIX, 166
- backing up Windows, 190
- restoring, 273
- sections, 166

RDS.log, 552
recognizing other data formats, 111
- Cartridge Memory enabled recognition, 112
- media format categories, 111
- recognized formats, 111

reconnecting broken connections, 239
recovering
- Cell Manager, UNIX, 512
- complete IDB, 421
- corrupted IDB, 422, 423
- IDB, 390, 417
- IDB, methods, 417

recovering the IDB
- complete, 421
- corrupted (missing) DC binary files, 422
- corrupted filename tablespace, 423
- DC directories, creating, relocating, 396
- directories, location, 391
- methods, 417
- `obrindex.dat` file, 426–428
- preparing, 390
- recovery file, creating a copy, 395
- relocating directories, 393
- robustness considerations, 390
- to a different disk layout, 431

recovery
- disaster recovery, 442

recovery procedure, 512
- Assisted Manual Disaster Recovery, Windows, 456
- Disk Delivery Disaster Recovery, UNIX client, 510
- Disk Delivery Disaster Recovery, Windows client, 461

Index

- Enhanced Disaster Recovery, Windows NT/2000 client, 470
- One Button Disaster Recovery, Windows NT/2000, 477

recurring backup
- configuring, 210

recycling
- media, 123

reducing
- IDB growth, 405
- IDB size, 406

Registry
- backing up Windows, 177
- NT CONFIGURATION, 173
- restoring Windows, 282

reliability
- media condition, 105

RemovableStorageManagementDatabase, 174

removing
- backup devices, 72
- user groups, 88
- users, 90

renaming
- backup devices, 73

reparse points, 170, 171

report groups
- adding multiple reports to, 337
- adding reports to, 335
- configuring, 335
- configuring, using Web reporting interface, 355
- examples, 335
- requirements, 316
- running, using CLI, 339
- running, using GUI, 338
- triggering by a notification, 337, 352

report level
- object specific option, 246

report open locked files as
- object specific option, 247, 248

reporting, 307
- accessing functionality, 315
- adding reports to a report groups, 335
- concepts, 315
- configuring report groups, 335
- configuring reports, 335
- report groups, 315, 335
- report input parameters, 315
- report send methods, 331
- reports format, 329
- reports on multiple cells, 315
- reports types, 317
- starting reports, 315
- user rights, 83

reports
- adding multiple to the report group, 337
- adding to a report group, 335
- configuring, 335
- configuring on the Web, 353
- customizing, 315
- formats, 329
- generating, using omnirpt command, 339
- generating, using Web reporting interface, 355
- groups, 338
- IDB, 400
- input parameters, 315
- on multiple cells, 315
- requirements, 316
- running individual, using GUI, 338
- running, using CLI, 339
- running, using GUI, 338
- send methods, 331
- starting, 315
- types, 317

responding to mount request, 310

RESTARTED, definition, 253

restarting
- backup device, 70
- fail sessions, 621
- failed backup, 266, 311

restore
- concepts, 4
- database application, troubleshooting, 586
- troubleshooting, 580
- with stacker devices, 36

restore options, 294
- display statistical information, 295
- for objects, 294
- keep most recent, 297
- list restored files, 295
- lock files, 296
- move busy files, 295
- no overwrite, 297
- omit deleted files, 294
- omit unrequired incrementals, 295
- overwrite, 297
- pre- exec commands, 296, 297

protection attributes, 296
sparse files, 296
target hostname, 294
time attributes, 296
restoring
bindery, Novell NetWare, 288
data to different client, 299
DHCP Server, 286
direct backup, 205
disk images, 273
file ownerships and trustees, 288
files from media, 301
files in parallel, 300
files in use, 302
files to different paths, 299
from media copy, 144
from media in a vault, 141
IDB, 414
individual files to different paths, 299
Name Space information, 287
NDS scheme, 290
Novell NDS, 289
Novell NetWare filesystems, 287
OmniStorage backups, 276
OpenVMS filesystems, 291
rawdisk, 273
regular files on Windows, 279
regular UNIX files, 276
shared disks, 279
UNIX files, 276
VxFS, A-3
Windows 2000/XP services, 283
Windows 2000/XP System State, 281
Windows CONFIGURATION, 280
Windows Registry, 282
Windows systems, 277
Windows TCP/IP services, 286
WINS server, 286
restoring DNS Server, 286
restoring NDS objects, 290
Results Area, 10
Results Tab, 10
retiring
media, 101
rights, user group, 93
root user rights, 85
running
consecutive backups, 212
report groups, using CLI, 339
report groups, using GUI, 338
reports, using CLI, 339
reports, using GUI, 338

S

SAN, 44
autoconfiguring devices, 50
concepts, 44
configuration goals, 48
configuration methods, 50
configuration overview, 47
configuring library robotics in a cluster, 52
configuring MC/SG, 58
FC-AL and LIP, 46
manually configuring the library robotics, 51
MC/Service Guard, 58
multiple system to multiple device connectivity, scheme, 45
simplified configuration for Windows environment, 51
sap.log, 552
scanning
device, 129
media in a device, 129
media in a device using Barcode Scan option, 129
media in a device using Cartridge Memory list, 130
media in a device, procedure, 129
stacker devices, 35
scheduled eject of media, 138
add the report to the report group and configure It, A-14
copy the script to the specified directory, A-15
example, A-14
notification on Mail Slots Full, 138
overview, 138
prerequisite, 138
schedule the report group, A-14
scheduled media copying, 145
scheduling
modifying backup, 210
predefined backup, 209
tips, 212
troubleshooting, 584
unattended backup, 207
scheduling policies
planning, 208

18c 790
J

Scoping Pane, 10
SCSI address, 29
SCSI ID
  library device, 29
SCSI-II Library devices
  SCSI address, 29
searching for media, 135
secondary node, 615
see private objects, 85
segment size
  restore speed, 78
selecting
  backup objects, 168
  Direct Access, figure, 56
  media, 120
  media for backup, 120
  media for backup, table, 121
  medium, 135
  medium, procedure, 135
send methods, notifications, 347
  broadcast message, 349
  e-mail, 348
  Event Log, 351
  external, 351
  log to file, 349
  SNMP, 349
  use report group, 351
send methods, reports, 331
  broadcast message, 332
  e-mail, 331
  external, 334
  log to file, 332
  SNMP trap, 332
SERVER_DR, 222
Serverless Integrations Binary Files, 386
ServiceGuard. *See* MC/ServiceGuard
services
  starting problems, 569
  troubleshooting, 569
Session Flow report
  configuring, example, 340
Session Messages Binary Files, 386
SESSIONID, definition, 253
SESSIONKEY, definition, 253
sessions
  aborting, 312
  aborting backup during the size determination, 312
  aborting, cluster environment, 623, 624, 625
  backup concepts, 4
  monitoring, 309
  monitoring finished, 310
  restart failed, 621
  rights to change ownership, 84
  temporary disabling, cluster environment, 626
  troubleshooting, 580
  viewing currently running, 309
Sessions in Timeframe reports, 326
setting
  advanced options, defining lock name, figure, 54
  backup options, 162
  block size, 79
  data protection, 122
  disk image options, 240
  MoM Manager, 363
  object specific options, 241
  same density, 42
  user account for the Inet, 187
setting up a backup environment
  tasks, 15
shared devices
  in the SAN, 44
shared disks
  backing up MC/ServiceGuard, 639
  backing up Microsoft Cluster Server, 620
  backing up Novell NetWare Cluster, 643
  backing up Veritas Cluster, 641
  backing up Windows, 185
  restoring, 279
short report format, 330
Single Instance Storage (SIS), 170, 171
Single Session reports, 328
slot number
  library devices, 29
sm.log, 552
SMEXIT, definition, 253
SNMP send method
  configuring reports, 333
  notifications, 349, 350
  reports, 332, 333
SNMP traps
  access points, 649
  configuring, Windows 2000, 333, 350
  configuring, Windows NT, 333, 350
  format, 651
software compression
  object specific option, 247
sparse files, 170, 172, 296
specific backup object

18.789

Index

  pre- and post-exec commands, 250, 255
Specific Event ID, 651
specific object, 225
specifying
  pathname for a file device, figure, 28
  type and name of the backup device, figure, 24
stacker devices
  backup and restore with, 36
  configuring, 35
  configuring, example, 35
  scanning, verifying and formatting, 35
standalone devices
  chains, 24
  configuring, 23
  mount request for, 581
  troubleshooting, 581
start backup specification
  user right, 84
starting
  daemons, 571
  daemons, problems, 571
  failed backup, 266
  GUI, UNIX, 7
  GUI, Windows, 7
  notifications checks, 608
  periodic backup, 209
  reports, 315
  services on Windows, problems, 569
  unattended backup, 209
  user interface, problems, 590
static drives
  configuring, 60
STK ACS, 655
stopping daemons, 571
Storage Area Network. *See* SAN
storing
  catalog backup, 230
strict media allocation policy, 103
support, before calling, 549
switch session ownership, 84
switchover, 615
sybase.log, 552
system and management applications
  access points, 649
  Generic Event ID, 651
  graphical user interface (GUI), 652
  overview, 649
  SNMP traps, 649
  SNMP traps format, 651
  Specific Event ID, 651
  variables, 651
  Windows Application Log, 652
system partition, 439
System Recovery Data (SRD), 445
System State
  backing up, 176
  restoring Windows 2000/XP, 281
  services, 174
SystemRecoveryData
  NT CONFIGURATION, 173

T

tab report format, 330
tape drives, 20
target hostname, 294
target system, 439
TCP/IP setup, checking, 567
templates, 216, 218
temporary directory missing, 597
testing
  drive cleaning configuration, 63
time attributes, 296
trace file name, troubleshooting, 556
triggering a report group by a notification, 337, 352
troubleshooting
  backup sessions, 580
  checking and maintenance mechanism, 605
  client fails, 567
  common problems, 562
  daemons, 569
  debugging, 553
  devices, 574
  disaster recovery, 514
  error messages, browsing, 561
  IDB, 592
  installing Cell Manager, Windows, 589
  installing clients, Windows, 588
  licensing, 584
  log files, 550, 588
  media, 574
  networking and communication, 565
  non-ASCII characters, 587
  online help, 603
  restore sessions, 580
  sample debugging, 558
  services, 569
  starting daemons, Unix, 571
  starting services, Windows, 569

troubleshooting file, 562
user interface, 590
when the user interface not accessible, 562
troubleshooting backup sessions
mount request for a library device, 582
mount request for a standalone device, 581
protection expiration, 586
starting interactive sessions, 585
starting scheduled sessions, 584
unexpected mounted system detected, 583
troubleshooting communication
client fails, 567
HOST file resolution problem, 567
host name resolution problems, 565
troubleshooting devices
device open problem, 575
Exchanger control device not accessible, 574
hardware-related problems, 579
using unsupported SCSI adapters, Windows, 575
troubleshooting disaster recovery
Disk Delivery Disaster Recovery, Windows NT, 515
Enhanced Disaster Recovery, Windows NT, 516
troubleshooting file, 562
troubleshooting media
medium header sanity check, 577
medium quality statistics, 575
troubleshooting messages, browsing, 561
troubleshooting networking
client fails, 567
HOST file resolution problem, 567
host name resolution problems, 565
troubleshooting the IDB
application restore sessions, 586
backup problems, 598
data files missing, 596
import problems, 598
libraries (executables) missing, 595
MMDB and CDB not synchronized, 600
performance problems, 599
temporary directory missing, 597
upgrade problems, 592
user interface running, problems, 595
trustees, restoring, 288
TSANDS.CFG, 200
types of notifications, 342
scheduled and started by the Data Protector checking and maintenance mechanism, 343
triggered when an event occurs, 342
types of reports
Backup Specifications, 318
Configuration, 320
Device Error, 336
IDB, 321
Mount Request, 336
overview, 317
Pools and Media, 324
Sessions in Timeframe, 326
Single Session, 328
typographical conventions, xxi

U

unattended backup
scheduling, 207
starting, 209
undoing the clear
editing backup schedule, 210
UNIX
backing up filesystems, 161
disk discovery, client backup, 163
NFS backup, 164
pre- and post-exec commands, 257
restoring disk image (rawdisk), 273
restoring regular files, 276
root user, 90
VxFS snapshot, A-3
UNIX Cell Manager
disaster recovery methods, 437
Manual Disaster Recovery, 512
recovery procedure, 512
UNIX client
Disk Delivery Disaster Recovery, 507
unused media
adding, 107
update SRD File, Wizard, 446
updating system recovery data (SRD), 445
Upgrade.log, 552
usage policy, media, 104
appendable, 104
appendable on incrementals only, 104
non-appendable, 104
use free pool option, 104
use report group send method, notifications, 351
Use Shadow Copy, object specific option, 248

- used media
  - adding, 107
- user account
  - setting for the Inet, 187
- User Check Failed notification, 347, 606
- user class
  - description of access rights, 83
- user configurations
  - examples, 94
  - rights, 83
  - users restoring their own data, 94
- user definable backup variables
  - object specific option, 248
- user disk quotas, 286
  - backing up, 183
- user groups
  - adding new, 88
  - changing rights, 93
  - deleting, 88
  - predefined, 86
- user interfaces
  - command-line interface
  - graphical user interface, 6
  - Microsoft Management Console
  - online Help, 12
- user profiles
  - backing up Windows, 182
  - restoring, 285
  - restoring deleted, 285
- user rights, 83
  - abort, 84
  - clients configuration, 83
  - device configuration, 83
  - for predefined groups, 86
  - media configuration, 83
  - monitor, 84
  - mount request, 84
  - private objects, 85
  - reporting and notifications, 83
  - restore as root, 85
  - restore from other users, 85
  - restore to other clients, 85
  - save backup specification, 84
  - session ownership, 84
  - start backup, 83
  - start backup specification, 84
  - start restore, 85
  - user configuration, 83
  - users backing up their systems, 94
- USER_FILES, 222
- users
  - adding, 90
  - adding groups, 88
  - changing rights, 93
  - configuration example, 94
  - configuring in MoM, 379
  - default, 90
  - deleting, 90
  - deleting groups, 88
  - description of access rights, 83
  - modifying, 92
  - moving, 92
  - predefined groups, 86
  - rights, 83
- using
  - a pre-allocation list of media for backup, 119
  - backup devices, 17
  - different media format types, 148
  - global options, 523
  - media for backups, 101
  - omnirc options, 525
  - several drive types in a library, 42
  - Web reporting interface, 353

V

- variables
  - access points, 651
  - global options file, 523
  - omnirc option files, 525
  - system and management applications, 651
- vault
  - configuring, 141
  - moving media to, 141
  - restoring from media in, 141
- vaulting
  - and Data Protector, 140
  - configuring vaults, 141
  - implementing, 140
  - media, 140
  - media to a safe place, 101
  - moving media to a vault, 141
  - restoring from media in a vault, 141
  - vault, 140
- verifying
  - data on a medium, 128
  - stacker devices, 35
- Veritas Cluster
  - clients, 640, 642
  - integration, 640

viewing
  currently running sessions, 309
  details of a running session, 310
  files from media, 301
  finished sessions, 310
views
  modifying in the media management
    window, 149
virtual server, 615
volume mount points, 170
Volume Shadow Copy service (VSS), 170
volumes
  backing up, 194
VSS
  See Volume Shadow Copy service (VSS),
    170
VSS filesystem backup, 170
VxFS
  snapshot, A-3

W

Wake ONLINE, 265
Web reporting and notifications interface
  accessing, 354
  changing password for, 354
  configuring notifications using, 355
  configuring report groups using, 355
  generating reports using, 355
  limitations, 353
  restricting access, 354
  using, 353
weekly full backup
  predefined backup schedules, 209
Windows
  administrator, 90
  ASR, 480
  Assisted Manual Disaster Recovery, 450
  Assisted Manual Disaster Recovery, client,
    450
  backing up, 168
  backing up DHCP Server, 178
  backing up event logs, 182
  backing up filesystems, 168
  backing up Registry, 177
  backing up shared disk, 185
  backing up System State, 176
  backing up user profiles, 182
  directory junctions, 170
  Disk Delivery Disaster Recovery, client, 459
  login, 90
  pre- and post-exec commands, 251
  restoring disk image (rawdisk), 273
  restoring event logs, 285
  restoring Registry, 282
  restoring regular files, 277, 279
  restoring shared disks, 279
  restoring user profiles, 285
  restoring WINS server, 286
  WINS Server backup, 178
Windows 2000
  Enhanced Automated Disaster Recovery,
    client, 463
  Manual Disaster Recovery, Cell Manager,
    450
  One Button Disaster Recovery, 472
Windows 2000/XP
  active directory restore, 283
  backing up services, 179
  backing up System State, 176
  certificate services restore, 284
  CONFIGURATION, 174
  directory junctions, 171
  restoring services, 283
  restoring System State, 281
  setting user account, 188
Windows Application Log, 652
Windows CONFIGURATION
  restoring, 280
Windows NT
  CONFIGURATION, 173
  configuring SNMP traps, 333
  Enhanced Automated Disaster Recovery,
    client, 463
  Manual Disaster Recovery, Cell Manager,
    450
  One Button Disaster Recovery, 472
  One Button Disaster Recovery, Cell
    Manager, 472
  setting user account, 187
  troubleshooting disaster recovery, 514
Windows TCP/IP services
  restoring, 286
Windows XP
  Automated System Recovery set, 483
WINS Server
  backing up, 178
  NT CONFIGURATION, 173
WINS server
  restoring, 286

18.755

Index

X

XCopy engine, 39, 204

JB.784

ANEXO UNIDADE DE BACKUP ROBOTIZADO
PARTE E/1

fl. 783

User Guide

rp5400 Family of Servers

First Edition

Manufacturing Part Number : A5191-96018
November 2002

USA

18.782

Legal Notices

The information in this document is subject to change without notice.

Hewlett-Packard makes no warranty of any kind with regard to this manual, including, but not limited to, the implied warranties of merchantability and fitness for a particular purpose. Hewlett-Packard shall not be held liable for errors contained herein or direct, indirect, special, incidental or consequential damages in connection with the furnishing, performance, or use of this material.

Restricted Rights Legend. Use, duplication or disclosure by the U.S. Government is subject to restrictions as set forth in subparagraph (c) (1) (ii) of the Rights in Technical Data and Computer Software clause at DFARS 252.227-7013 for DOD agencies, and subparagraphs (c) (1) and (c) (2) of the Commercial Computer Software Restricted Rights clause at FAR 52.227-19 for other agencies.

HEWLETT-PACKARD COMPANY 3000 Hanover Street Palo Alto, California 94304 U.S.A.

Windows“, Windows NT“, Windows 95“, Windows 2000“, and Windows XP“ are registered trademarks of Microsoft in the U.S. and other countries.

Reproduction, adaptation, or translation of this document without prior written permission is prohibited, except as allowed under the copyright laws.

781

Contents

1. Server Overview

2. Server Unpacking and Installation

Factory Integrated rp54xx Cabinet Installation 4
Receive and Unpack A Non-Integrated Server 7
Unpacking the server 7
Install Deskside Server 10
Install Stand-Alone Server in a Cabinet. 13
Stationary L-Bracket Rail Assembly 19
Identifying Approved Non-E-Series HP Cabinets. 19
Identifying E-Series HP Cabinets 20
Identifying Static Rail Kit 20
Installing Stationary Rails. 21

3. Installing Additional Components

Additional Components 24
Installing Memory 25
Memory Configuration Rules 25
Installing Peripheral Component Interconnect (PCI) Cards. 30
rp5400/rp5450 PCI Card Slots 30
rp5470 PCI Card Slots 31
PCI I/O Card Installation Restrictions 32
PCI I/O Card Installation Order 32
Installing a PCI Card 34
Online Addition/Replacement (OLA/R) of PCI I/O cards 36
Installing Graphics 37
Graphics Troubleshooting 43
Installing Disk Drives 44

4. Cable Connections

Core I/O Connections 48
Revision A GSP 48
Revision B GSP 49
Guardian Service Processor (GSP) Overview 51
GSP LAN 51
GSP RS-232. 51
GSP Features 51
Revision A GSP 52
Revision B GSP 52
Configure System Consoles 53
GSP Cables 54
Configure RS-232 Console 54
Configure Remote Console 58
Configure the LAN Console 59
Configure the Web Console 61
Secure Web Console Installation. 66

18.730

Contents

GSP Configurable Parameters 69
Adding Users. 69
Removing Users 70
Return the GSP to Default Configurations 71
rp54xx Server Boot Process 72
Initial Power-up 74

5. Utilities

Configuring the Rev A Guardian Service Processor (GSP) 78
Configuring the GSP LAN Port 78
Adding Users. 78
Removing Users 80
Changing the Default GSP Configuration. 80
Configuring the Rev B Guardian Service Processor (GSP) 81
Configuring the GSP LAN Port 81
Adding Users. 81
Removing Users 83
Changing the Default GSP Configuration. 83

6. Troubleshooting

Determine Current System State 86
Troubleshooting and FRU identification 87
Problem Symptoms and Repair Actions 87
Chassis Code to FRU Decode 91
Cross-Referencing Chassis Log Errors to rp54xx FRUs. 91
Interpreting System Alerts 94
Interpreting Service Processor Error Chassis Logs 96
Interpreting Chassis Logs Using the chassis_code.codes File 98
rp5400 and rp5450 System Block Diagram 100
rp5400 and rp5450 I/O Block Diagram. 101
rp5430 and rp5470 System Block Diagram 102
rp5430 and rp5470 I/O Block Diagram. 103
Run/Attention/Fault LED States 104
PCI I/O LED States 112
Expansion I/O LED States 114
GSP LED States. 116
LAN/SCSI LED States. 118
Fan, Power Supply, and Disk LED States 119

7. Removing and Replacing Components

List of Changeable Parts with Remove and Replace Components 122
Cardcage Fan 122
Core I/O 122
HotSwap Chassis Fan. 122
Disk Drive 123
Display Board 123

18.779

Front Bezel . . . 123
Memory DIMM . . . 123
PCI I/O Card . . . 124
Power Supply . . . 124
HotSwap Power Converter Fan . . . 124
Platform Monitor . . . 124
Processor Support Module . . . 125
Individual Component Remove/Replace Instructions . . . 126
Extend the Server out the Front . . . 126
Insert the Server from the Front . . . 127
Stand-alone Server Cover Removal . . . 128
Stand-alone Server Cover Replacement . . . 130
Top Cover Removal . . . 131
Top Cover Replacement . . . 132
Side Cover Removal . . . 133
Side Cover Replacement . . . 134
Front Bezel Removal (Single Piece) . . . 135
Front Bezel Replacement (Single Piece) . . . 136
Front Bezel Removal (Two Piece) . . . 137
Front Bezel Replacement (Two Piece) . . . 138
Core I/O Removal . . . 139
Core I/O Replacement . . . 140
PCI Card Separator/Extractor Removal . . . 142
PCI Card Separator/Extractor Replacement . . . 143
HotPlug Disk Drive Removal . . . 144
HotPlug Disk Drive Replacement . . . 149
HotSwap Chassis Fan Cover Removal . . . 152
HotSwap Chassis Fan Cover Replacement . . . 153
HotSwap Chassis Fan Removal . . . 154
HotSwap Chassis Fan Replacement . . . 155
HotSwap Card Cage Fan Removal . . . 156
HotSwap Card Cage Fan Replacement . . . 157
HotSwap Power Supply Removal . . . 158
HotSwap Power Supply Replacement . . . 159
HotSwap Power Converter Fan Removal . . . 160
HotSwap Power Converter Fan Replacement . . . 161
Processor Support Module Removal . . . 162
Processor Support Module Replacement . . . 163
Memory DIMM Removal . . . 164
Memory DIMM Replacement . . . 165
Display Board Removal . . . 166
Display Board Replacement . . . 167
Platform Monitor Removal . . . 168
Platform Monitor Replacement . . . 170
PCI I/O Card Removal . . . 172
PCI I/O Card Replacement . . . 173

Contents

A. Parts and Accessories

CRU Physical Location 176
Customer Replaceable Unit Part Numbers 179

B. System Specifications

Dimensions 182
Uncrating Space 182
Space Requirements 182
Computer Room Physical Space Requirements 183
Computer Room Unpacking Space 184
Electrical 185
Office High Availability Requirements 185
Power Protection 185
Modular PDUs 186
System Power Requirements 186
LAN and Telephone 187
Acoustic Safety Standards 188
Altitude Operation Standards 189
Effects of Altitude 189
Temperature and Humidity Operating Standards 190
Thermal Protection Features 190

C. General Site Preparation Guidelines

Electrical Factors 192
Computer Room Safety 192
Power Consumption 193
Electrical Load Requirements (Circuit Breaker Sizing) 193
Power Quality 193
Distribution Hardware 194
Grounding Systems 195
System Installation Guidelines 197
Environmental Elements 198
Computer Room Preparation 198
Cooling Requirements 198
Humidity Level 200
Dust and Pollution Control 200
Metallic Particulate Contamination 201
Electrostatic Discharge (ESD) Prevention 202
Acoustics 202
Facility Characteristics 203
Floor Loading 203
Windows 205
Space Requirements 206
Delivery Space Requirements 206
Operational Space Requirements 206
Floor Plan Grid 206

Contents

Typical Installation Schedule 207
Site Inspection 208
Delivery Survey 212

Tables

Table 6-1. Problem Symptoms and Repair Actions 87
Table 6-2. Chassis Log Error to FRU Decoder 91
Table A-1. Exchange CRUs 179
Table A-2. Non-Exchange CRUs 180
Table B-1. Power Requirements 186
Table C-1. Effect of Humidity on ESD Charge Levels 202
Table C-2. Floor Loading Term Definitions 203
Table C-3. Typical Raised Floor Specifications. 204
Table C-4. Customer and Hewlett-Packard Information 208
Table C-5. Site Inspection Checklist 209

Tables

Figures

Figure A-1. Server Rear View 176
Figure A-2. Side Service Bay 177
Figure A-3. System Board (Access via Top Service Bay) 177
Figure A-4. Server Front 178
Figure C-1. Raised Floor Metal Strip Ground System 196
Figure C-2. Delivery Survey (Part 1) 213
Figure C-3. Delivery Survey (Part 2) 214

18.771

Preface

Printing History

The Printing History below identifies the edition dates of this manual. Updates are made to this publication on an unscheduled, *as needed*, basis. The updates will consist of a complete replacement manual and pertinent on-line or CD-ROM documentation.

First Edition	November 2002

What's New?

The, *Upgrade Guide, rp5400 Family of Servers*, is new and was developed to provide customers with system maintenance information for those components called customer replaceable units (CRUs). Maintenance of CRUs does not require HP customer engineering services, except when specifically cautioned. The cautions are shown primarily to protect customer product warrantees.

fls. 770

18.769

1 Server Overview

The rp5400 family of servers are 1-way to 4-way servers based on the PA-RISC processor architecture. The rp5400 family of servers accommodate up to 16GB of memory and internal peripherals including disks and DVD ROM/Tape. High availability features include HotSwap fans and power supplies, and HotPlug internal disk drives. The supported operating system is HP-UX.

2 Server Unpacking and Installation

18.766

Factory Integrated rp54xx Cabinet Installation

A factory integrated server is one in which the rp54xx server and associated components are pre-assembled and shipped from the factory already installed in a Hewlett-Packard E-Series cabinet. Factory integrated systems reduce the amount of time required to set-up and begin server operation.

1. Carefully remove the carton and anti-static bag from the pallet.
2. Remove the front two (2) L-brackets. Retain the 1/2-inch bolts for later use.

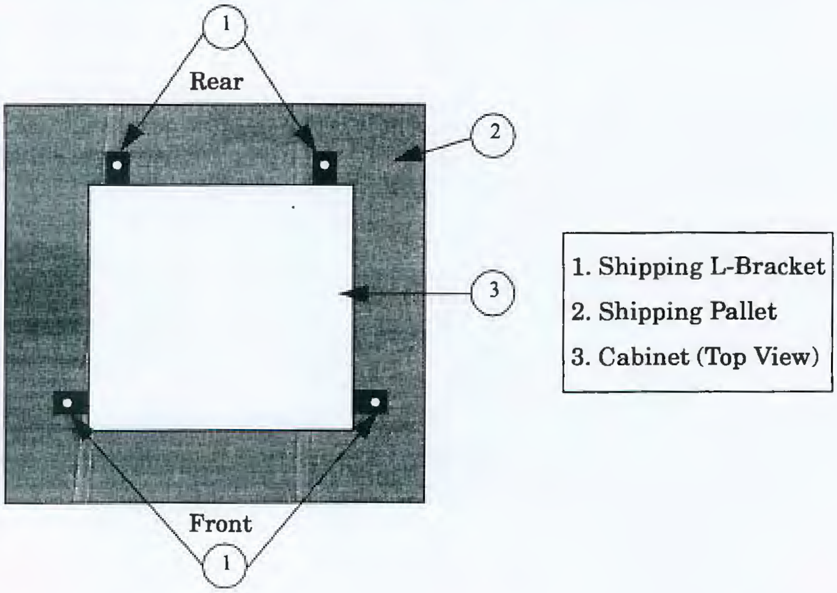

NOTE As viewed from the front, one bracket is located on each side at the base of the cabinet near the front.

fl. 765

3. At the rear of the cabinet:

   a. Open the door.

   b. Remove the anti-tip foot by removing and retaining the two (2) 1/2-inch bolts.

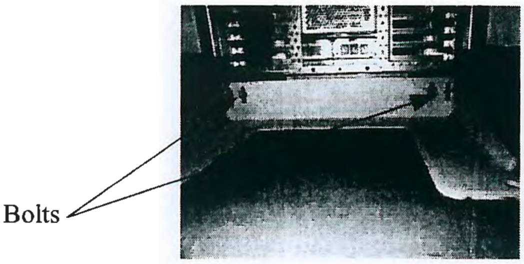

For Shipping:
L-brackets are
mounted behind
anti-tip foot.
Same bolts
secure both.

   c. Remove the two (2) L-brackets (revealed by removing the anti-tip foot).

4. Remove the two ramps from the pallet and carefully place them into the slots at the front of the pallet.

WARNING **Use extreme care when rolling the racked system down the ramps. A rack containing one rp54xx can weigh up to 418 lbs. Do not stand in front of the ramps when rolling the cabinet off the pallet or injury may occur. All but the smallest configurations require two persons to safely remove the rack from the pallet.**

If anti-tip feet or ballast are not installed or are improperly installed the cabinet can tip. Failure to follow this precaution can cause injury to personnel or damage to equipment.

---

5. Straighten the rollers on the cabinet base, if needed, and carefully roll it down the ramps.

WARNING ***After removing the server from the pallet, Do not move the cabinet unless the anti-tip feet are installed!*** **The cabinet can tip if care is not used. Due to their low ground clearance the feet may catch on irregularities on the floor, thresholds, or ramps.**

Do not move the cabinet without first installing the anti-tip feet. **The cabinet may tip if moved without the anti-tip feet or ballast installed.**

Do not move the cabinet after installing the anti-tip feet unless they are in the fully-raised position. **Once installed, the anti-tip feet must be fully raised to allow ground clearance.**

Because of their low ground clearance, the fully-raised anti-tip feet may need to be removed *temporarily* to clear some obstacles such as door jambs, ramps, and other large irregularities or obstructions on the floor.

If you must temporarily remove the anti-tip feet to clear an obstacle, use extreme caution when moving the cabinet. Always reinstall the anti-tip feet as soon as the obstacle has been cleared.

Lower and secure both the anti-tip feet and the cabinet leveling/stabilizer feet once the cabinet is in place.

Failure to follow these precautions can result in equipment damage or personal injury.

6. Install the front and rear anti-tip feet using the 1/2 inch bolts provided. Ensure that the anti-tip feet are installed in the fully up position in the mounting slots. This will provide maximum ground clearance while moving the cabinet to its final position.
7. Carefully move the cabinet to its installation location.
8. Lower the anti-tip feet to the fully down position and adjust the cabinet leveling feet for best cabinet stability.

18.763

Receive and Unpack A Non-Integrated Server

WARNING **The typical rp54xx system can weigh up to 68kg (150lbs). *HP recommends using an an approved lifting device.* Lift and move the server in accordance with all local safety regulations. Failure to follow this precaution can cause injury to personnel or damage to equipment.**

Unpacking the server

The following procedure describes the steps involved in unpacking the server, whether to function as a stand-alone Deskside unit, or to be integrated into a cabinet.

Step 1. Remove the shipping carton and anti-static bag from the server as depicted below.

NOTE The packaging for rp74xx and rp54xx servers is the same, rp74xx is shown.

Step 2. If you are moving the server manually, use three people to lift the server from the packing material and pallet. Carefully move the server to the selected location.

Step 3. If you are moving the server by an approved lifting device (such as Genie Lift ™), remove the tear flap from the front lip of the carton bottom to allow access to the server, as illustrated below. Removal of the tear flap will reveal a slot between the bottom of the server and the inside bottom of the cardboard box.

18.761

Step 4. Carefully raise the lift's platform so that it will slide into the slot located ***under*** the center of the server, but ***over*** the top of the pallet.

NOTE The server's center of gravity will vary with the hardware configuration, but it is generally located slightly behind the middle of the server.

Step 5. Raise the lifting device platform enough for the server to clear the pallet and packing materials, as show below.

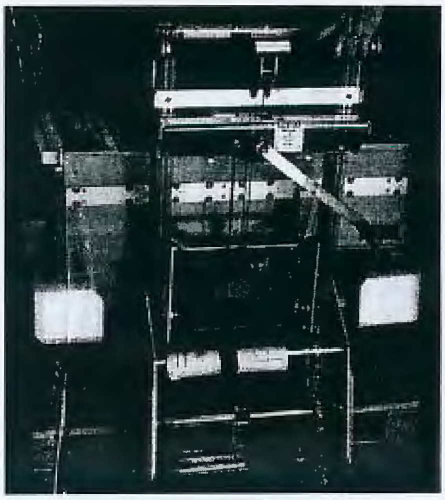

Install Deskside Server

The following section describes the installation of a server into a Deskside enclosure for installation in an office environment.

WARNING **The typical rp54xx system can weigh up to 68kg (150lbs).** ***HP recommends using an approved lifting device.***

- **Lift and move the server in accordance with all local safety regulations.**
- **Do not attempt to lift the server by the plastic handles on the top and side covers.**

Failure to follow these precautions can cause injury to personnel or damage to equipment.

Step 1. Unpack the server.

Step 2. Unpack the deskside enclosure.

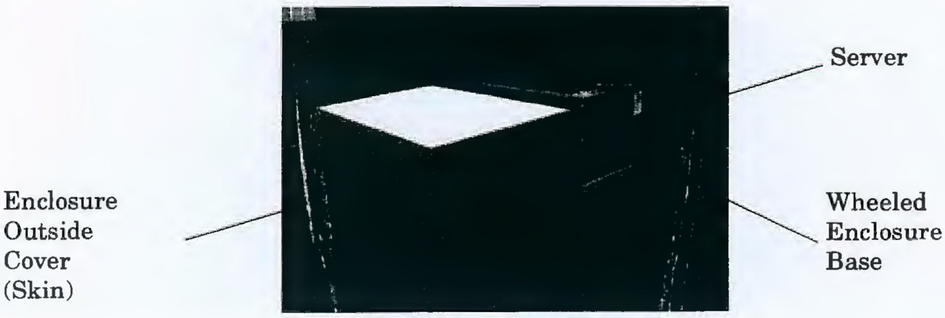

NOTE Ensure that the positioning spring pins in the enclosure base align with the alignment holes in the bottom of the server.

18.759
J

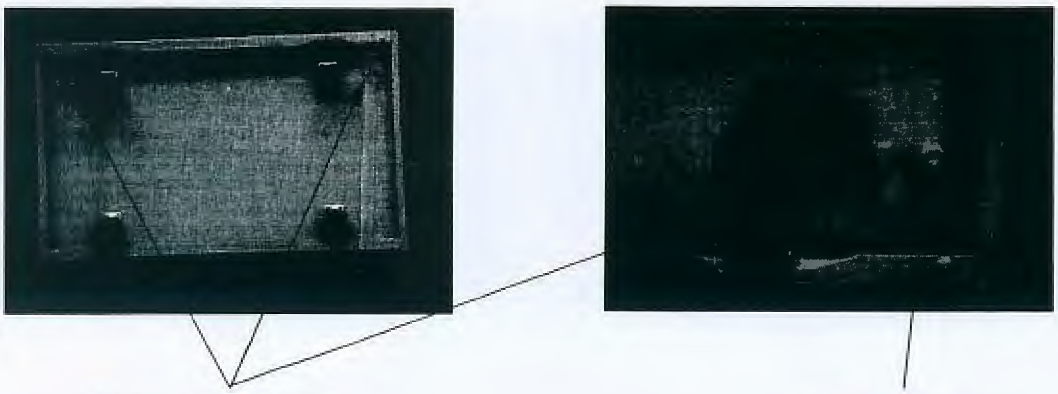

Step 3. Position the server on the wheeled enclosure base.

Step 4. Tighten the two captive screws in the enclosure base to secure the server to the base.

Step 5. Position the enclosure cover (outside skin) over the server and install and tighten the screws to secure it to the base.

NOTE The perforations and the lip of the outside skin should be toward the rear of the server.

WARNING **Stacking rp54xx servers in deskside enclosures is *not* supported.**

Stacking rp54xx servers in deskside enclosures can damage equipment, may cause injury to personnel, and may void your warranty or service contract.

Step 6. Install the Front Bezel.

Step 7. Locate the two pull-tabs. One pull-tab is longer than the other. The shorter pull-tab is blank on both sides. The back of the shorter pull-tab provides a writable surface for Customer use.

18.758

Step 8. Locate the plastic bag containing the label sheet (taped to the server).

Step 9. Remove the label containing serial number, base product, processor product, and model information from the label sheet and apply to the back of the longer pull-tab.

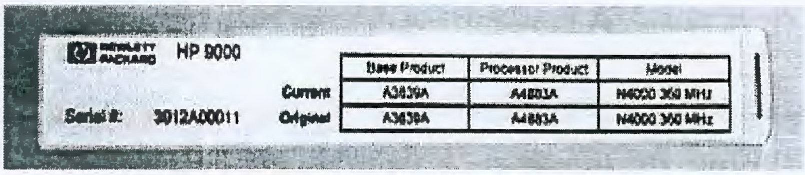

HEWLETT PACKARD HP 9000

Serial #: 3012A00011

	Base Product	Processor Product	Model
Current	A3639A	A4883A	N4000 360 MHz
Original	A3639A	A4883A	N4000 360 MHz

NOTE Pull-tab and label shown above is for an rp74xx server. rp54xx uses the same style label and similar pull-tab.

Step 10. Insert the pull-tabs into the front bezel. Install the longer pull-tab in the left side plastic window in such a way that the rp54xx logo is visible. Install the shorter pull-tab in the right side plastic window with either surface visible. Refer to the diagram above for pull-tab locations.

Install Stand-Alone Server in a Cabinet

The following describes how to install the A5556A slide-tray assembly into an approved HP cabinet in preparation for installing an rp54xx server.

This slide-tray assembly can be installed in an HP E-Series cabinet or other HP cabinets approved for rp54xx system installation. To install the A5556A slide-tray assembly in an approved HP equipment cabinet, proceed as follows:

Step 1. Determine what type of cabinet you are installing the slide-tray assembly into.

a. E-Series cabinets have:

- Parchment white, plastic, sectional, side panels
- Black painted vertical frame posts with a partial return flange.

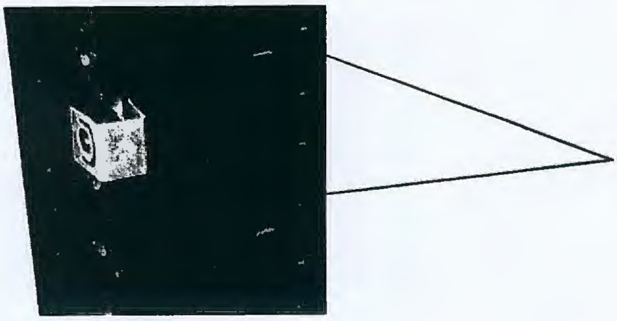

Vertical, Rectangular, Mounting Slots

b. Approved, non-E-Series, cabinets have:

- Single piece metal side panels
- Gray painted verticle frame posts with full return flanges.

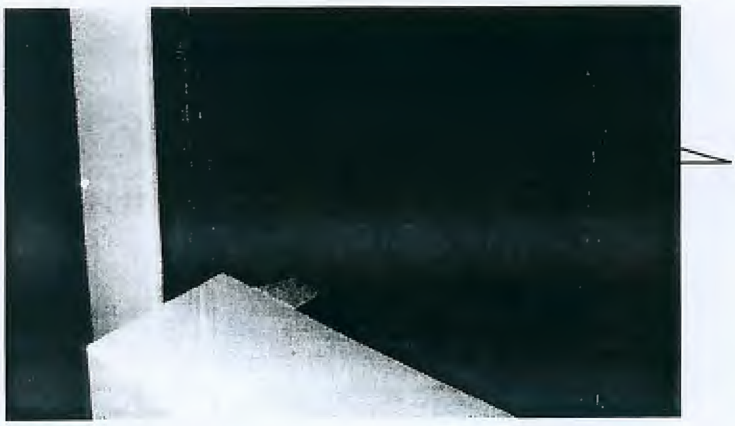

Vertical, Rectangular, Mounting Slots

Step 2. Note the vertical, rectangular, slots in the return flanges on the vertical mounting posts. Determine into which of these vertical slots the slide/tray kit will be installed. This is done by counting down eight rectangular slots from the top of the cabinet or the bottom of the equipment above.

Step 3. On the front vertical mounting posts *only*, slide M5 sheet metal nuts onto the posts over the holes immediately adjacent to the vertical slots determined in the previous step. Also place M5 sheet metal nuts on the holes directly above these. Orient the sheet metal nuts so that the threaded portion faces towards the outside of the cabinet. There should now be a total of four (4) sheet metal nuts installed.

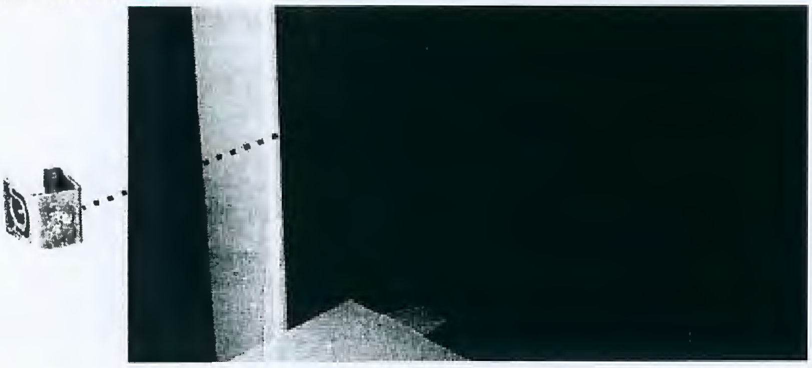

Step 4. If the cabinet is a non-E-Series cabinet, discard the left hand and right hand aluminum spacers and two of the M5 x 16 screws with cress-cup washers and proceed to step 12.

Step 5. If the cabinet is an E-Series cabinet, place the hook of the aluminum spacer marked "L" (5183-1864) into the appropriate vertical, rectangular slot on the front, left hand mounting post. The hook points downward. Similarly, place the spacer marked "R" (5183-1863) into the appropriate slot on the right hand mounting post.

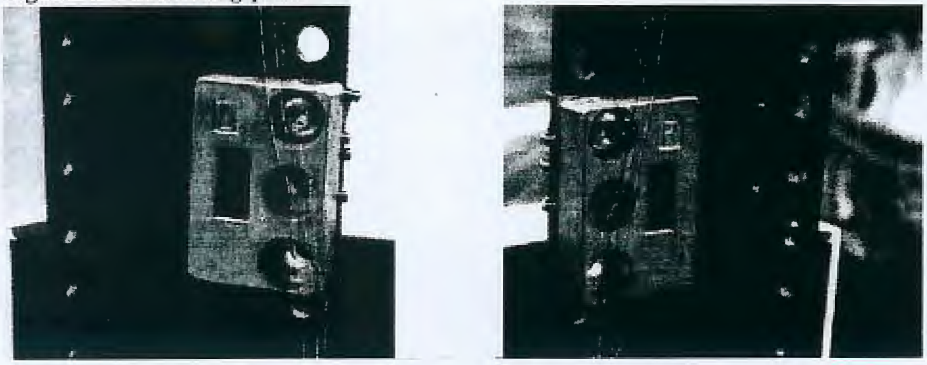

Step 6. Use one M5 x 16 screw with cress-cup washer to attach each spacer to its vertical post. Do this by inserting the screw through the top hole in the spacer, through the mounting rail and tightening it into the sheet metal nut located at that position.

Step 7. Take the left hand slide/bracket assembly (marked 337079-1L) and install it into the left hand vertical mounting posts. This is done by inserting the pin at the rear of the slide's mounting bracket into the 23rd hole in the rear vertical mounting post and inserting the hook at the front of the

18.755

bracket into the vertical, rectangular slot in the aluminum spacer. The slide should be positioned in the cabinet so that it is horizontal and level.

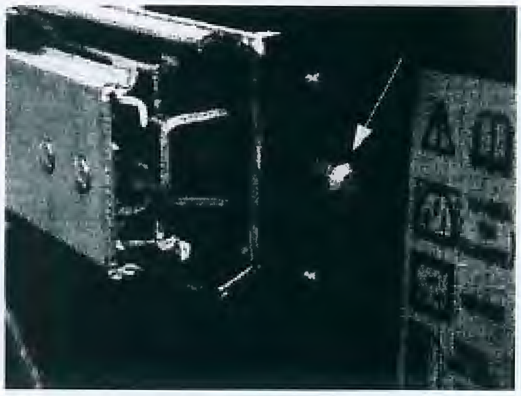

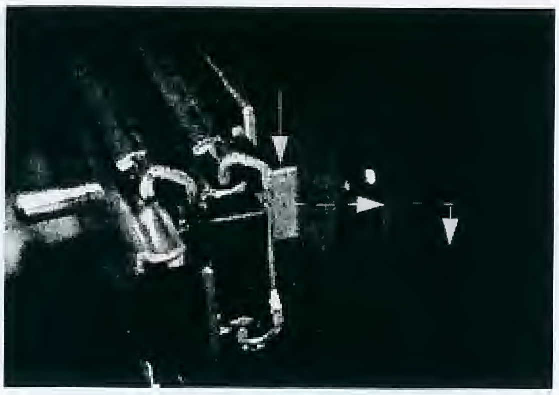

Step 8. Securely fasten the rear of the slide's mounting bracket to the rear vertical mounting post by installing and tightening two of the M5 x 16 screws with cress-cup washers thorough the mounting post, through the slides mounting bracket and into the threaded nuts attached to the mounting bracket.

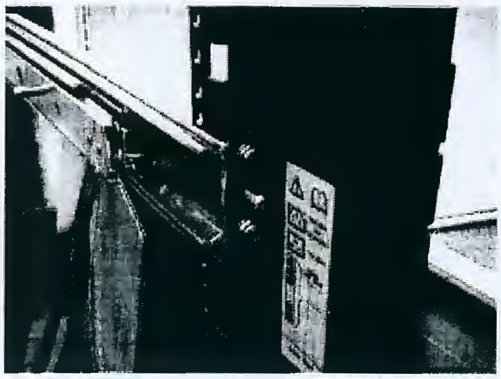

Step 9. Fully extend the slide so that it is locked in the fully open position.

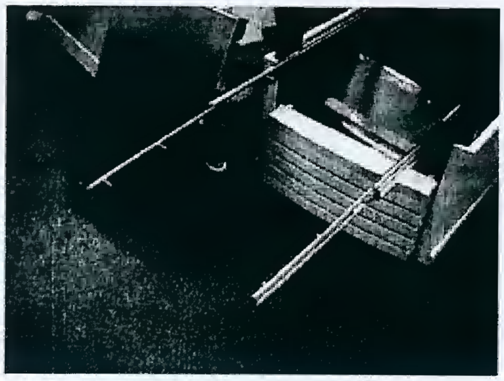

Step 10. Use an M5 x 30 screw with a cress cup washer to attach the front of the slide to the vertical mounting post. Insert the screw through the slide, through the center hole of the aluminum spacer, through the vertical mounting post, and tighten into the sheet metal nut located at that position.

Step 11. Use a procedure similar to steps 7 through 10 to install the right hand slide/bracket assembly (marked 337079-1R) and then proceed to step 12.

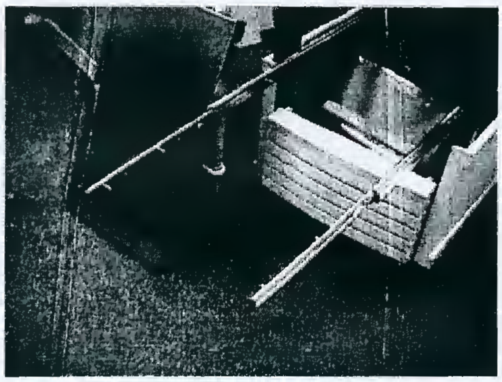

18.753

Step 12. Take the tray and place it onto the pins that extend from the slides' inner members. The slots with wide lead-in guides on the side of the tray fit down onto the slides' pins. The flat part of the tray will be on top, and the mounting holes in the top of the tray will be located to the right of the center of the tray. Slide the tray all the way down on both sides so that the pins reach the top of the slots in the side of the tray.

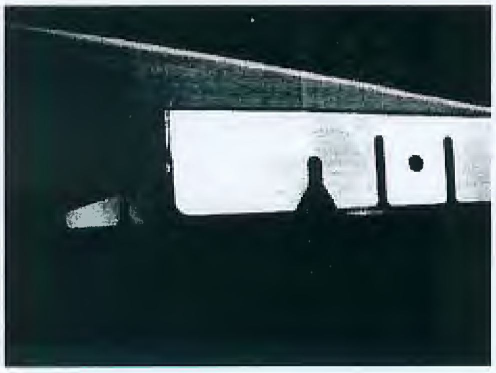

Step 13. Use six, M5 x 12 screws (without washers) to attach the tray to the slides. Three screws are used to attach each slide. Insert the screws through the slides, through the tray and tighten into the threaded nuts located on the inside of the sides of the tray.

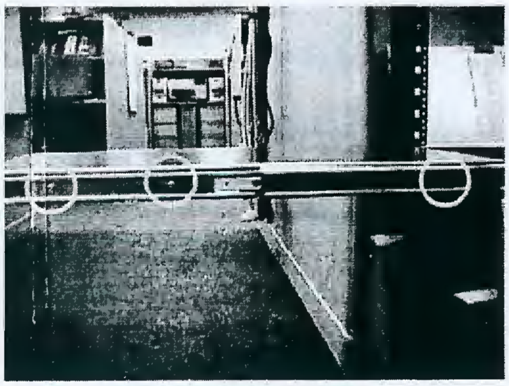

3 Pan Head
M5x12 T25 screws
on each side

Step 14. From the bottom of the tray pull the plunger pin down and give it a 1/4 turn to hold it in place.

18.753

Step 12. Take the tray and place it onto the pins that extend from the slides' inner members. The slots with wide lead-in guides on the side of the tray fit down onto the slides' pins. The flat part of the tray will be on top, and the mounting holes in the top of the tray will be located to the right of the center of the tray. Slide the tray all the way down on both sides so that the pins reach the top of the slots in the side of the tray.

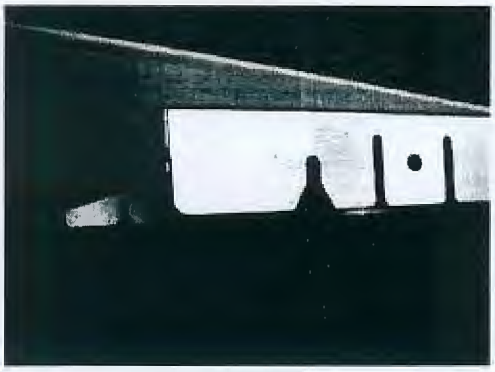

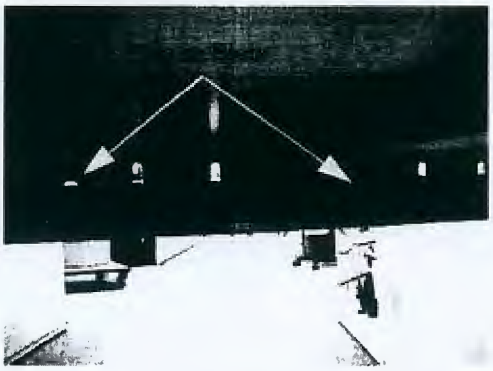

Step 13. Use six, M5 x 12 screws (without washers) to attach the tray to the slides. Three screws are used to attach each slide. Insert the screws through the slides, through the tray and tighten into the threaded nuts located on the inside of the sides of the tray.

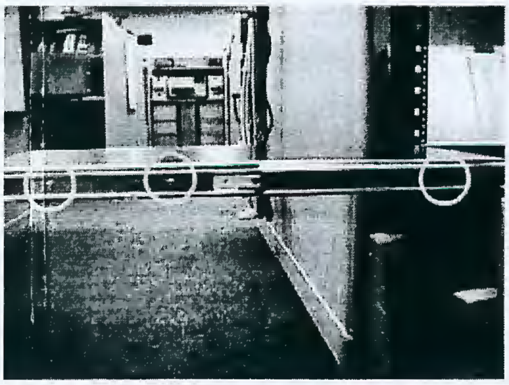

3 Pan Head
M5x12 T25 screws
on each side

Step 14. From the bottom of the tray pull the plunger pin down and give it a 1/4 turn to hold it in place.

18.752
J

Step 15. Position the server on the tray aligning the plunger pins with the alignment holes in the chassis.

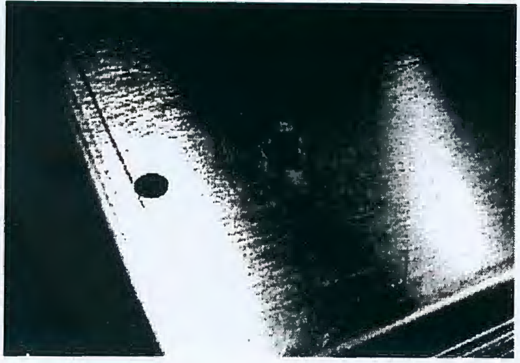

Step 16. Release the plunger pins to secure the server.

18.751

Stationary L-Bracket Rail Assembly

rp54xx servers may be installed into E-Series and approved Non- E-Series cabinets using stationary L-bracket rail assembly kits listed below.

NOTE rp54xx servers are supported in Hewlett-Packard E-series and approved Non- E-series Hewlett-Packard cabinets, and approved rail kits.

For information on additional qualified 3rd party cabinets and rail kits, contact the nearest Hewlett-Packard Response Center.

Cabinet Type	Rail Kit Product Number
E-Series HP Cabinet	A5575A
Other Approved HP Cabinet	A5562A

Identifying Approved Non-E-Series HP Cabinets

Approved Non- E-Series cabinets have black frames, one piece outside sheet metal skins, a partial return flange, and *requires* the installation of the aluminum spacer blocks, supplied with the rail kits.

Approved Non- E-Series cabinets include the following product numbers: A1883A, A1884A, A1896A, A1897A, C1897A, C2785A, C2786A, and C2787A.

Identifying E-Series HP Cabinets

E-Series cabinets have light gray frames, sectioned, plastic outside "skins", a full return flange, and does *not* require the installation of the aluminum spacer block supplied, with the rail kits.

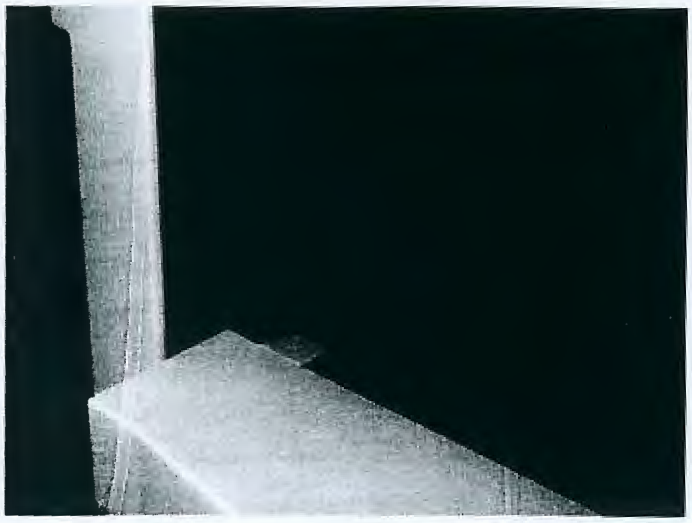

E-Series cabinets include the following product numbers: A5134A, A5136A, A5136A, A4900A, A4901A, A4902A,, J1500A, J1502A, and J1502A.

Identifying Static Rail Kit

Hewlett-Packard has currently approved two static rail kits for use in cabinet mounting the rp54xx server. They are illustrated below.

A5562A Kit Rail

A5575A Kit Rail

Installing Stationary Rails

The installation of stationary rails is similar for most cabinet and rail combinations. The key considerations to are:

- Ensure that all safety precautions are read, understood, and observed
- Follow all installation instructions provided with the cabinet and rail kits, and
- Ensure that the rails extend out from the cabinet posts sufficiently to properly and safely support the equipment being installed.

To install an rp54xx server on stationary rails in an approved cabinet proceed as follows:

Step 1. Locate the rail mounting height in the cabinet. Allow for the following space requirements:

- For each rp54xx server, allow 31.8cm (12.5 inches) vertically (7 EIAs or Rack Units (RUs).
- If installing the A5575A rail kit, allow an *additional* vertical 4.45cm (1.75 inches (1 EIA) each set of rails.

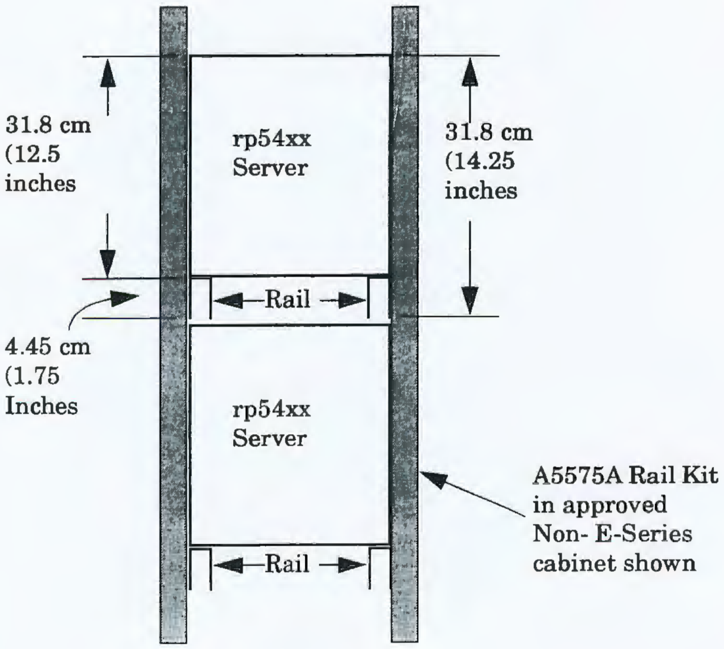

Step 2. Install sheet metal nut(s) in the vertical cabinet posts at the required height for the kit being installed:

- Install the first nut either:
  — 4.45 cm (1.75 Inches) above the top, or
  — 31.8 cm (12.5 inches) below the bottom of the last server.
- If installing a A5562A rail kit, install the second nut in the next frame hole below the first.

Step 3. Hold the rail in place and insert and tighten the screws.

For installation of other qualified cabinet and rail combinations refer to the safety precautions and instructions accompanying them.

3 Installing Additional Components

18.746

Additional Components

Some internal components are too delicate to be installed in the server prior to shipping. These internal components are shipped with the server, but are packed separately. They can be installed after the cabinet has been unpacked and positioned.

Some of the internal components that are packed separately are not user-installable. To maintain warranty validation, these items *must* be installed by a Hewlett-Packard Customer Engineer.

If you received either (or both) of the components listed below, contact your Hewlett-Packard provider to arrange for installation.

- Central Processing Units (CPUs)
- Power Distribution Units (PDUs)

18.745

Installing Memory

Memory Configuration Rules

rp54xx servers have 16 slots (8 DIMM pairs) for memory DIMMs. These slots are numbered 0a/b, 1a/b,... 7a/b. 8 of these slots (4a/b - 7a/b) are disabled on rp5400 servers. rp5450 servers can access all slots. rp5400 and rp5450 servers have DIMM slots located on the System Board.

rp5470 servers install DIMMs using Memory Carriers. The Memory Carriers fit into slots on the System Board.

The following rules govern the installation of memory DIMMs for rp5400, rp5450, and rp5470 servers:

- Memory must be installed in DIMM pairs.
- The capacity of DIMMs within a pair must be the same.
- Install DIMMs with the greatest capacity in the lowest slot numbers.
- Install DIMMs the following slot order: 0a/b, 1a/b, 2a/b, 3a/b, and so on.

Installing rp5400 and/or rp5450 DIMMs

Step 1. Power down and *unplug* the rp54xx server.

CAUTION DC voltages are present when the server is connected to AC power. Do not install or service rp54xx internal components while DC voltage is present. Failure to observe this precautior can result in damage to the server.

Step 2. Loosen the captive T-15 screws that hold the top cover in place, then grasp the strap handle, raise the cover slightly, and pull the cover toward the front of the server to free the cover tabs from the slots in the chassis. The air baffle will be exposed.

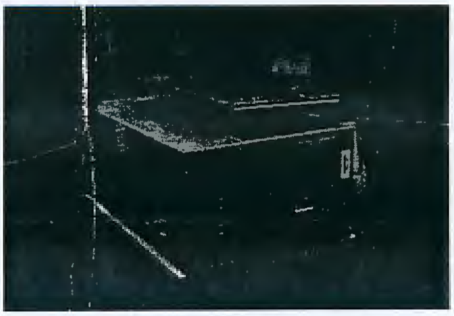

Step 3. Make the top of the server accessible for service.

Step 4. Loosen the captive T-15 screws on the air baffle. Grasp the two handles on the baffle, and lift the baffle remove it.

CAUTION Observe all ElectroStatic Discharge (ESD) precautions Do not touch internal components. Failure to observe ESD precautions can cause damage to components.

Step 5. Observe Electrostatic Discharge (ESD) precautions.

Step 6. Refer to the following graphic for memory slot locations.

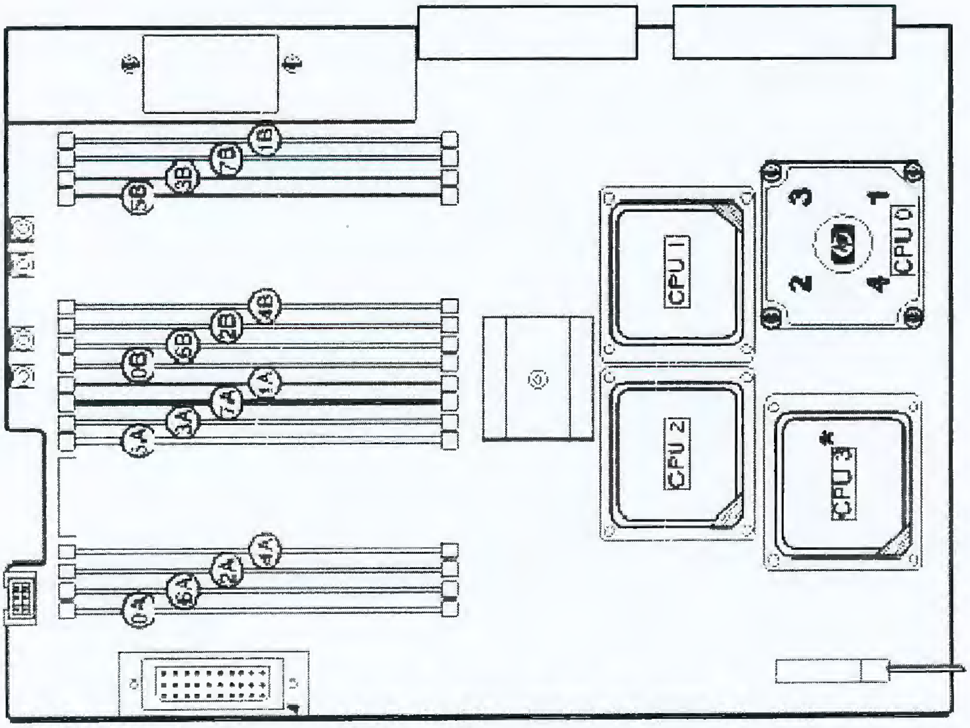

Locate the correct DIMM pair slots. Insert the DIMM connectors into the guides until the card snaps firmly in place. It may be necessary to apply downward force using the palm of your hand on the DIMM. Observe the top of the DIMM to make sure one side is not higher than the other.

NOTE It may be necessary to remove PSM 1 when installing a DIMM in slot 0a and PSM 0 when installing a DIMM in slot 1b. If either PSM is removed to install memory, ensure it is re-installed.

Step 7. Replace the air baffle. Tighten the four captive screws to secure the air baffle in place.

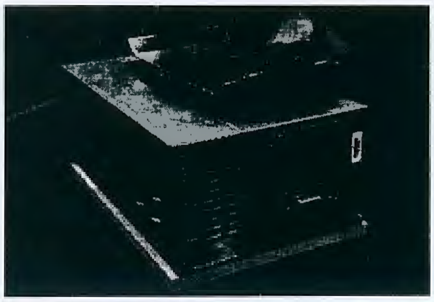

Step 8. Replace the top cover. Tighten the four captive screws to secure the top cover in place.

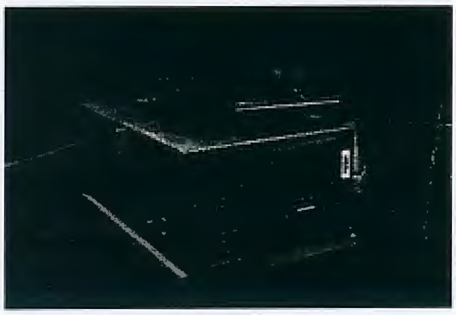

Step 9. For rack configurations, insert the rp54xx server back into the rack.

Step 10. For deskside enclosure configurations, replace the deskside enclosure cover.

Step 11. Power the rp54xx server on.

Step 12. Use the BCH command `in me` to verify the system recognizes the memory that you have just added.

18.742

Installing rp5470 DIMMs

DIMMs for the rp5470 system are installed in memory carriers instead of the system board, as are the other rp54xx systems. However, rp5470 memory carriers are also located on the system board, so the method for opening and closing the system is the same. Procedures for removing and replacing the server top and baffle are listed below, without the pictures shown in the section titled, *"Installing rp5400 and/or rp5450 DIMMs."* If you wish to reveiw the pictures, please refer to the aforementioned section.

Step 1. Power down and unplug the rp54xx server.

NOTE DC voltages are present when the server is connected to AC power. Do not attempt to install or service: CPUs, Memory, PSMs, the Platform Monitor or PCI I/O cards installed in non-Turbo slots (1-6) while DC voltage is present. Failure to observe this warning may result in damage to the server.

Step 2. Make the top of the server accessible for service.

Step 3. Loosen the captive T-15 screws that hold the top cover in place, then grasp the strap handle, raise the cover slightly, and pull the cover toward the front of the server to free the cover tabs from the slots in the chassis. The air baffle will be exposed.

Step 4. Loosen the four (4) captive T-15 screws on the air baffle. Grasp the two handles on the baffle, and lift and remove the baffle.

Step 5. Observe Electrostatic Discharge (ESD) precautions.

Step 6. Refer to the following graphic for Memory Carrier locations.

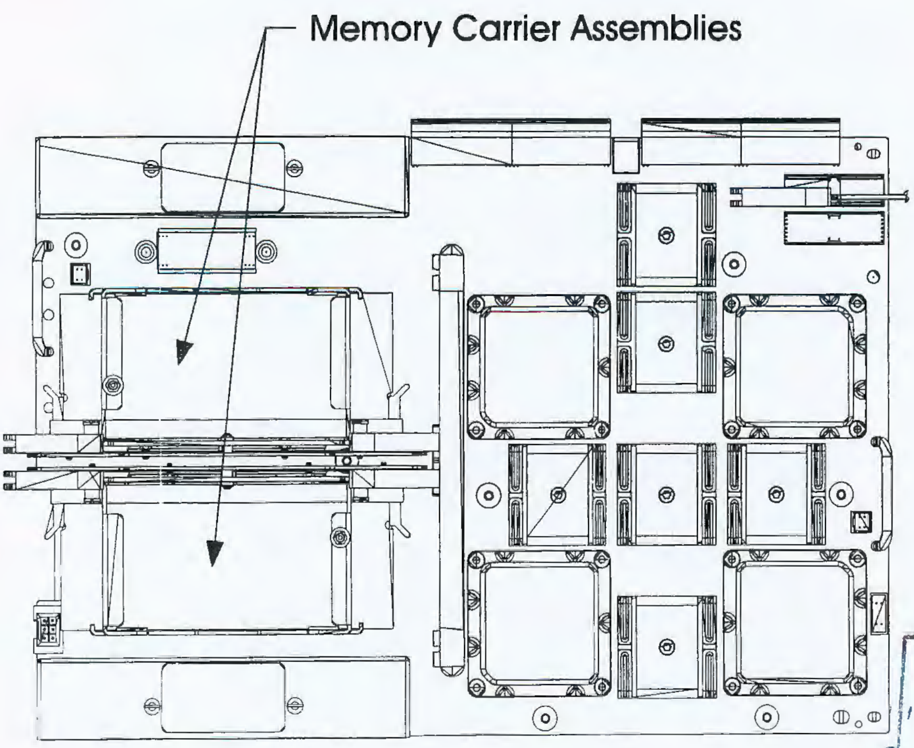

18.745

a. Locate the Memory Carrier and pull up on the extractor levers on each end of the Memory Carrier to unseat the Memory Carrier from its socket.

b. When the Memory Carrier unseats from the socket, pull it away from the System Board.

c. Loosen the captive screws that secure the DIMM Clip and remove the DIMM Clip from the Memory Carrier.

d. Seat the memory DIMM into its socket on the Memory Carrier.

e. Press the extractor levers on each end of the memory DIMM slot inward until the levers snap into place.

f. Attach the Memory Clip to the Memory Carrier with the DIMM slot markings on the top of the Memory Clip aligned with the DIMM slot markings on the Memory Carrier.

g. Secure the Memory Clip using the captive screws.

h. Seat the Memory Carrier into the appropriate slot on the System Board.

i. Push down on the extractor levers and snap them into place.

Step 7. Replace the air baffle. Tighten the four captive screws to secure the air baffle in place.

Step 8. Replace the top cover. Slide the cover tabs into the slots in the chassis and close the cover. Tighten the two captive screws to secure the top cover in place.

Step 9. For rack configurations, insert the rp54xx server back into the rack.

Step 10. For deskside enclosure configurations, replace the deskside enclosure cover.

Step 11. Power the rp54xx server on.

Step 12. Use the BCH command `in me` to verify the system recognizes the memory that you have just added.

18.710

Installing Peripheral Component Interconnect (PCI) Cards

rp54xx servers have a total of 12 PCI I/O slots. Slots 1 and 2 are reserved for the LAN/SCSI and GSP Core I/O cards, leaving 10 PCI I/O slots available for Customer use.

rp5400/rp5450 PCI Card Slots

For rp5400 and rp5450 models, 10 PCI I/O slots consist of Turbo and non-Turbo slots. Server PCI slots are shown below.

rp5400/rp5450 PCI Card Slots

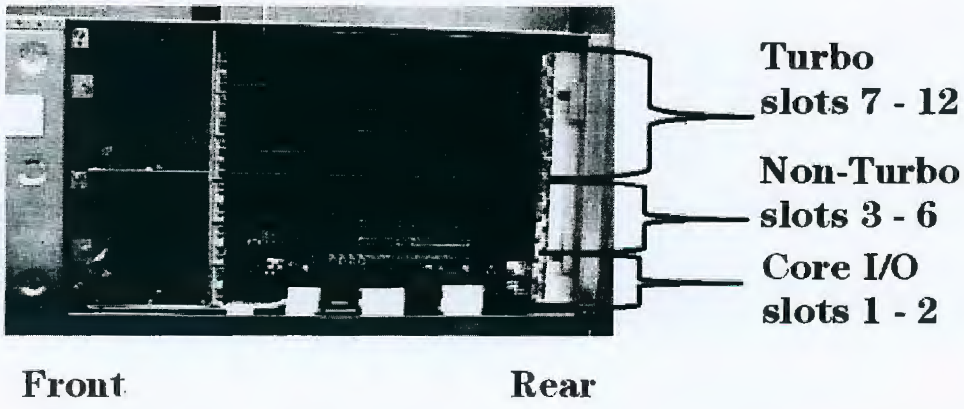

- Slots 1 and 2 are reserved for the rp54xx LAN/SCSI and GSP (Guardian Service Processor) Core I/O cards, respectively. Slots 1 and 2 are non-Turbo slots. Non-Turbo slots share a single 250MB/s PCI bus and are incapable of HotPlug functionality. The server must be turned off prior to removing or installing the LAN/SCSI or GSP cards in these slots.
- Slots 3 - 6 are non-Turbo slots. These four Non-Turbo slots share a single 250MB/s PCI bus, run at 33MHz and support 32 and 64-bit PCI cards. Non-Turbo slots are incapable of HotPlug functionality. The server must be turned off prior to removing or installing PCI cards in these slots.
- Slots 7 - 12 are Turbo slots. Each Turbo slot has a dedicated 250MB/s PCI bus, run at 66MHz and support 32 and 64-bit PCI cards. Turbo slots are HotPlug capable. Below each Turbo slot is a plastic PCI card separator. The PCI card separator has two LEDs and a pull tab on the outer edge. The LED's provide power and status for the slot. The pull tab allows the PCI card to be easily removed.

rp5400 servers have access to slots 1, 2 and 8-12 while rp5450 servers have access to all (1-12) slots.

NOTE Slot 3 will become enabled on rp5400 servers with server firmware versions later than 40.48.

A slot 3 enabled label (A5576-84009) is available for rp5400 systems.

fl. 739

rp5470 PCI Card Slots

For rp5470 models, the 10 PCI I/O slots consist of Twin Turbo, Turbo, and non-Turbo slots. The following illustration shows the PCI card slot layout.

rp5470 PCI Slots

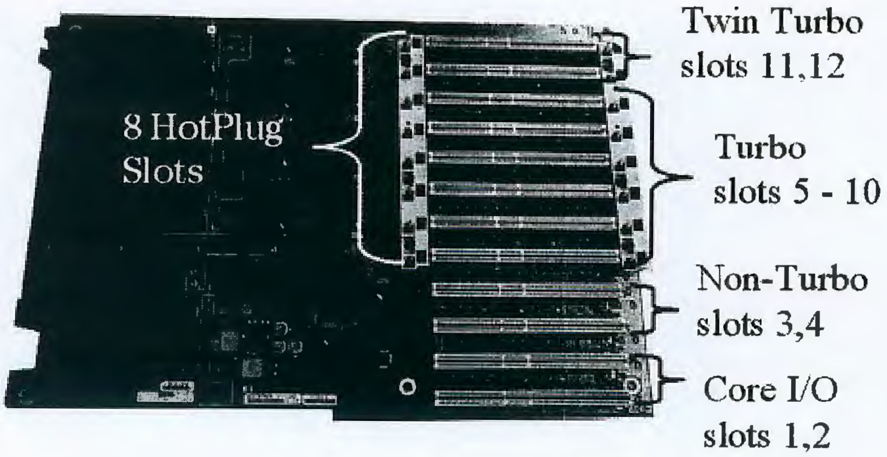

- Slots 1 and 2 are reserved for the rp54xx LAN/SCSI and GSP (Guardian Service Processor) Core I/O cards, respectively. Slots 1 and 2 are non-Turbo slots. Non-Turbo slots share a single 250MB/s PCI bus and are incapable of HotPlug functionality. The server must be turned off prior to removing or installing the LAN/SCSI or GSP cards in these slots.
- Slots 3 and 4 are non-Turbo slots. These two Non-Turbo slots share a single 250MB/s PCI bus, run at 33MHz and support 32 and 64-bit PCI cards. Non-Turbo slots are incapable of HotPlug functionality. The server must be turned off prior to removing or installing PCI cards in these slots.
- Slots 5 - 10 are Turbo slots. Each Turbo slot has a dedicated 250MB/s PCI bus, run at 66MHz and support 32 and 64-bit PCI cards. Turbo slots are HotPlug capable. Below each Turbo slot is a plastic PCI card separator. The PCI card separator has two LEDs and a pull tab on the outer edge. The LED's provide power and status for the slot. The pull tab allows the PCI card to be easily removed.
- Slots 11 and 12 are Twin Turbo slots. Each Twin Turbo slot has a dedicated 500MB/S PCI bus, runs at 66 MHz, and supports 32- and 64-bit PCI cards. Twin Turbo slots are HotPlug capable. Below each Twin Turbo slot is a plastic PCI card separator. The PCI card separator has two LEDs and a pull tab on the outer edge. The LED's provide power and status for the slot and the pull tab allows the PCI card to be easily removed.

rp5470 servers have access to all (1-12) slots.

PCI I/O Card Installation Restrictions

Restrictions apply regarding the installation of PCI I/O cards which contain a PCI-to-PCI bridge:

- HP-UX boot is currently not supported for cards that contain a PCI-to-PCI bridge.
- HP-UX patches are required when more than one card containing a PCI-to-PCI bridge is installed in non-Turbo slots.

PCI I/O Card Installation Order

The following table shows a standard factory PCI card installation that begins with slot 12. Use this table as a guideline for installing PCI I/O cards in the field.

NOTE A system shipped from the factory may have a different configuration than the same system built in the field. For example: The factory will install the graphics card in slot 12 and add other cards below. In the field, slot 12 may already be occupied by another PCI card. It is acceptable for the graphics card to be installed in any available Turbo slot.

Product Number	Description (all are PCI cards)	Max	Boot	Load Order *	Part Number	Notes
A6150A	Graphics, Graphics Card	1	No	1	A4982-66501	3,8
A5838A	Combo		No	3	A5838-60001	9
A5483A	ATM 622Mbps MMF Adapter	10	No	4	A5483-60001	10
A4926A	1000Base SX PCI LAN Adapter	10	No	5	A4926-60001	
A4926A	1000Base TX PCI LAN Adapter	10	No	6	A4926-60001	
A6092A	HYPERFabric		No	7	A6092-60001	11
A5158A	FC Taclite		Y	8	A5846-60001	
A5486A	Praesidium Speed Card	10	No	9	A5486-60001	
A5506A	4 Port 100Base TX LAN Adapter	7/10	No	10	A5506-60101	1,2,6
A5506B	4 Port 100Base TX LAN Adapter	7/10	No	10	A5506-60102	
A5150A	Dual Port Ultra 2 SCSI adapter	10	Yes	11	A5150-60001	4
A5149A	Single Port Ultra 2 SCSI HBA	10	Yes	12	A5149-60001	
J3526A	High Perf 4 Ports Synchronous Adapter	10	No	13	5063-1322	7,5
A4800A	FWD SCSI-2 adapter	10	Yes	14	A4800-67002	
A5230A	100Base-T LAN Adapter	10	No	15	B5509-66001	
A3738A	10/100Base-T LAN Adapter	10	No	16	A3738-60001	
A3739A	Dual FDDI LAN Adapter	10	No	17	A3739-60001	
A5783A	Token Ring 4/16/100 Hardware Adapter	10	No	18	A5783-60101	

Product Number	Description (all are PCI cards)	Max	Boot	Load Order *	Part Number	Notes
J3525A	Dual Port Synchronous Adapter	10	No	19	J3525-60001	
J3593A	64 port Serial MUX system card	10	No	20	J3593-60001	
J3592A	8 Port PCI Serial MUX card	4	No	21	J3592-60101	
A6150A	Graphics, USB Card	1	No	22	A6150-60001	
A6150BX	Pinnacle 2 Graphics	1	No	1	A6150-60003	12,13
A6386A	Hyper Fabric 2 Interconnect	10	No	6	A3686-60001	
A5506A	Quad Port 10/100B-TX LAN	10	No	10	A5506-60102	14
A6749A	3.3v 64 Port Terminal MUX	10	No	24	A6749-60001	
A6748A	3.3v 8 Port Terminal MUX	10	No	25	A6748-60001	

*In top down order.

Notes:

1. Card contains a PCI-to-PCI bridge.
2. Requires PHKL_20123, PHKL_20629 and PHNE_19826 or their superseded equivalents.
3. Not supported in non-Turbo slots. Install in Turbo slots only.
4. Requires server firmware revision 39.46 or later.
5. Requires HP-UX 11.1
6. Maximum is 7 for HP-UX versions prior to 11.0. Maximum is 10 for HP-UX version 11.1 and later.
7. Requires PHKL_19543 and PHKL_19544 or their superseded equivalents.
8. Requires HP-UX 11.0 Support Plus (IPR) 0006, June 2000 or later. This product to be released 6/00.
9. Not supported in a shared slot (slots 3-4 for rp5470, slots 3-6 for rp5450, not applicable for rp5400).
10. If you are installing ATM 622 cards in an rp5470 configuration, do not install them in slots 3 and 4 (shared slots).
11. Requires 768 MB for first card and 512 MB for each additional card.
12. Not supported in shared slots.
13. Max of 1. Needs USB card for keyboard and mouse.
14. Contains PCI bridge.

Installing a PCI Card

Follow these procedures to install a PCI card.

Step 1. Power down and unplug the rp54xx server.

NOTE DC voltages are present when the server is connected to AC power. Do not attempt to install or service: CPUs, Memory, PSMs, the Platform Monitor or PCI I/O cards installed in non-Turbo slots (1-6) while DC voltage is present. Failure to observe this warning may result in damage to the server.

Step 2. Make the right side of the server accessible for service.

Step 3. Using a Torx 15 screwdriver, loosen the captive screws on the right side panel. This panel has a label which shows which PCI I/O slots are available and the corresponding paths. The PCI I/O slot paths for rp5400, rp5450, and rp5470 are shown below.

Slot	**rp5400**		**rp5450**		**rp5430/rp5470**	
	Slot Type	**Path**	**Slot/Type**	**Path**	**Slot Type**	**Path**
12	**Turbo**	**0/4/0**	**Turbo**	**0/4/0**	**Twin Turbo**	**0/10/0**
11	Turbo	0/7/0	Turbo	0/7/0	Twin Turbo	0/12/0
10	Turbo	0/3/0	Turbo	0/3/0	Turbo	0/8/0
9	Turbo	0/6/0	Turbo	0/6/0	Turbo	0/9/0
8	Turbo	0/2/0	Turbo	0/2/0	Turbo	0/3/0
7	Not Available		Turbo	0/5/0	Turbo[a]	0/1/0
6	Not Available		Shared	0/1/0	Turbo[a]	0/5/0
5	Not Available		Shared	0/1/1	Turbo[a]	0/2/0
4	Not Available		Shared	0/1/2	Shared[a]	0/4/0
3	Not Available[b]		Shared	0/1/3	Shared	0/4/2
2	GSP		GSP		GSP	
1	LAN/SCSI		LAN/SCSI		LAN/SCSI	

a. Slot is NOT AVAILABLE for rp5430.
b. Slot 3 becomes available with server firmware versions later than 40.48.

Step 4. Remove the PCI slot cover from the slot that will receive the PCI card. To remove the PCI slot cover, slide the PCI slot cover away from the server.

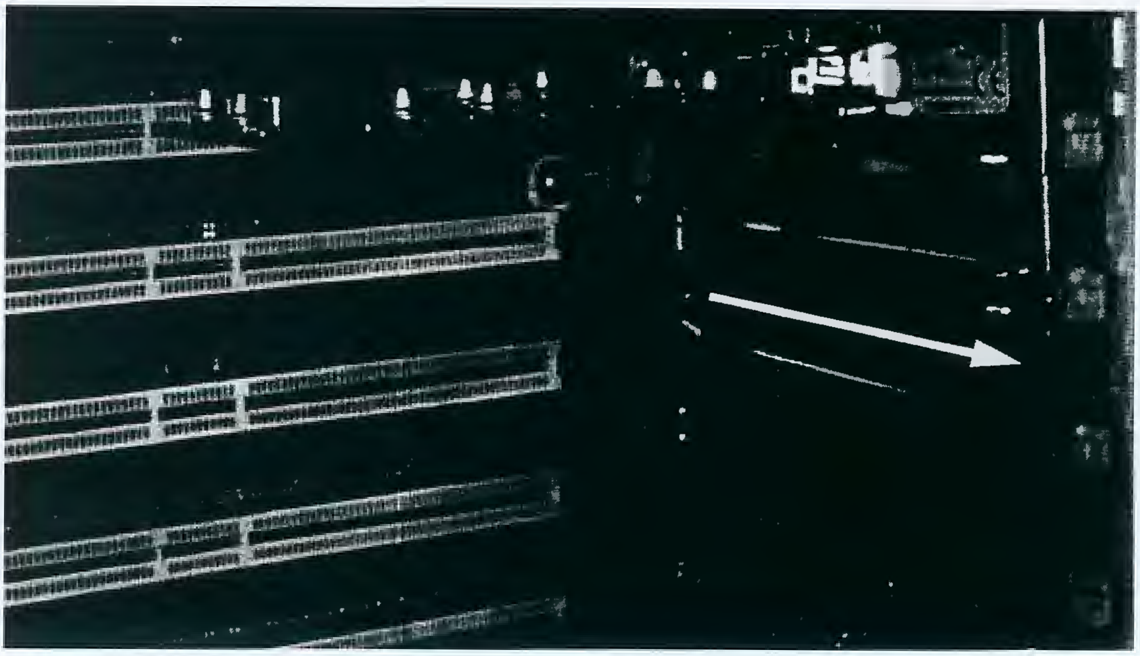

Step 5. Slide the PCI card connectors into the slot, snapping firmly in place. For full length (cards that extend to the left side card guides) PCI cards, use the UPPER card guide.

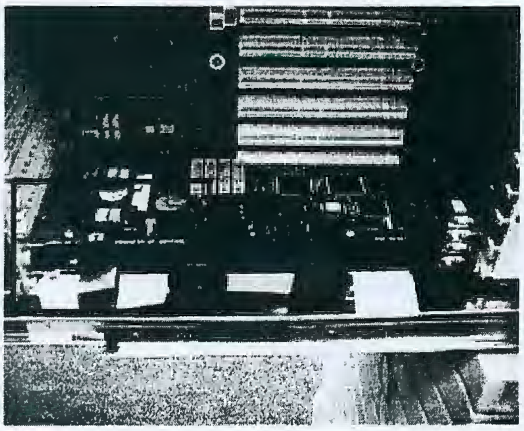

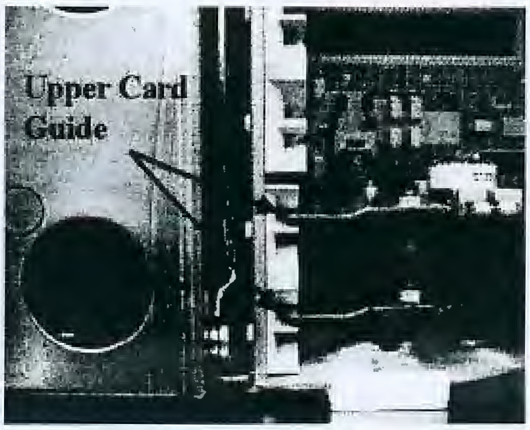

Step 6. At the rear of the chassis, connect the I/O cable to the card just installed.

Step 7. Replace the right side panel and tighten the captive screws.

Step 8. For rack configurations, insert the rp54xx server back into the rack.

18.734

Step 9. For deskside enclosure configurations, replace the deskside enclosure cover.

Step 10. Power the server on.

Step 11. Use the server firmware `in io` command to verify the PCI cards are recognized by the server. If AUTOBOOT is ON, it will be necessary to interrupt the boot process to get to the server firmware `Main Menu: Enter command or menu >` prompt.

Step 12. Boot HP-UX and run the `ioscan` utility to verify the system recognizes the new PCI card.

Online Addition/Replacement (OLA/R) of PCI I/O cards

Beginning with HP-UX 11i (11.11) rp54xx servers support the on-line addition and replacement of PCI I/O cards. In order for this high availability feature to be fully implemented, the following server requirements must be met:

- rp5400A/rp5450A firmware must be later than 40.26 (rp5400B/rp5450B/rp5470A firmware will support OLA/R upon its release).
- HP-UX operating system must be 11i (11.11) or later.

There is a bit that the HP-UX operating system examines to determine if the server hardware and firmware is capable of OLA/R. This bit is controlled by server firmware. If the bit is ON, OLA/R is possible (when requirements have been met). The bit was mistakenly set to ON for all rp5400 and rp5450 revision A (rp5400A and rp5450A) servers. As a result, HP-UX may incorrectly identify these models as being OLA/R capable. In order to avoid this confusion, verify that the correct level of server firmware is installed.

18.733

Installing Graphics

This section explains how to install rp54xx 2D graphics hardware. For a complete graphics solution, three products are required. The products listed below are the only products supported on rp54xx servers.

- A6150A rp54xx Graphics Package
  - — Includes PCI graphics card
  - — Includes PCI USB (Universal Serial Bus) card
- A4983B Keyboard and Mouse Kit
  - — Includes mouse with 114" cable
  - — Includes keyboard with 109" cable
- D8910W (19") or D2847W (21)" Monitor
  - — Includes localized power cord and 75" 15-pin video cable

NOTE rp54xx graphics requires HP-UX 11.0 Support Plus (IPR) 0006, June 2000 or later.

The photo below includes the A6150A, A4983B and D8910W products. The video cable for the monitor is not shown. Black *ESD mat not included.*

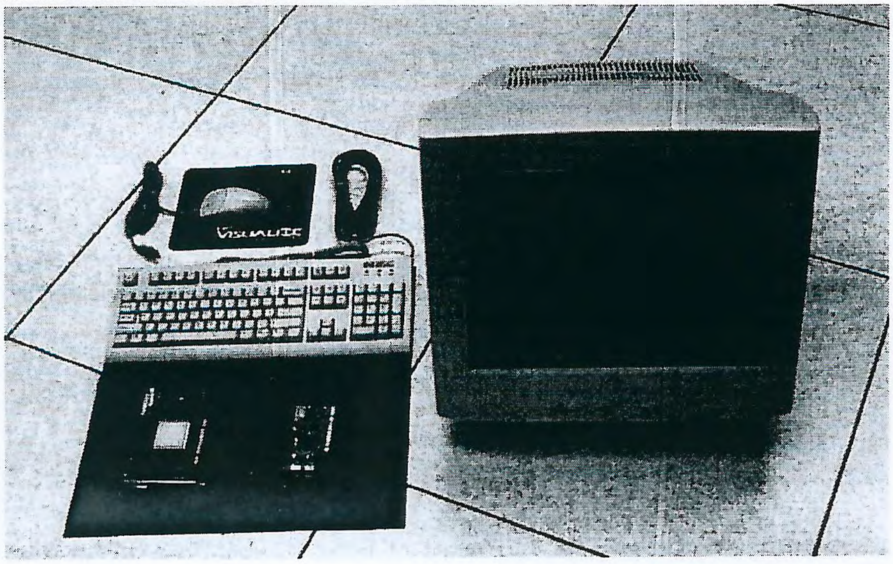

rp54xx servers have a total of 12 PCI I/O slots. Slots 1 and 2 are reserved for the LAN/SCSI and GSP Core I/O cards, leaving 10 PCI I/O slots available for Customer use. These 10 PCI I/O slots consist of Turbo and non-Turbo slots.

rp54xx PCI Slots

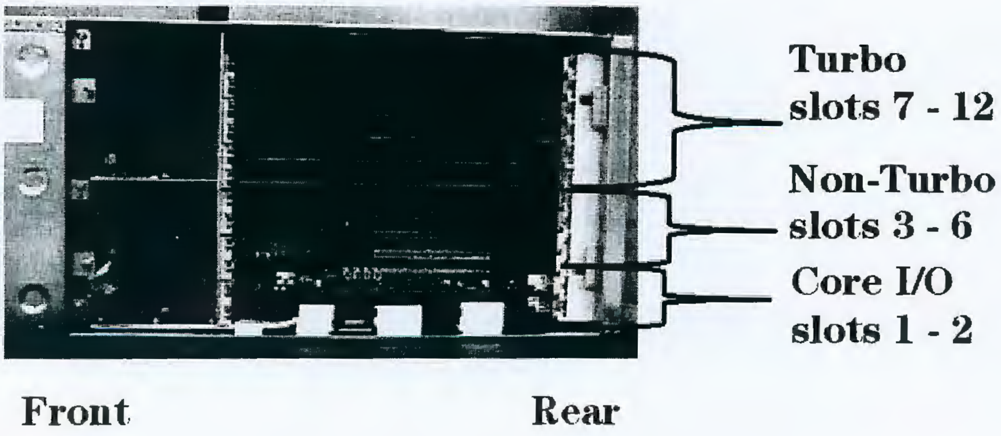

- Slots 1 and 2 are reserved for the rp54xx LAN/SCSI and GSP (Guardian Service Processor) Core I/O cards, respectively. Slots 1 and 2 are non-Turbo slots. Non-Turbo slots share a single 250MB/s PCI bus. Non-Turbo slots are incapable of HotPlug functionality. The server must be turned off prior to removing or installing the LAN/SCSI or GSP cards in these slots.
- Slots 3 - 6 are non-Turbo slots. These four Non-Turbo slots share a single 250MB/s PCI bus, run at 33MHz and support 32 and 64-bit PCI cards. Non-Turbo slots are incapable of HotPlug functionality. The server must be turned off prior to removing or installing PCI cards in these slots.
- Slots 7 - 12 are Turbo slots. Each Turbo slot has a dedicated 250MB/s PCI bus, run at 66MHz and support 32 and 64-bit PCI cards. Turbo slots are HotPlug capable. Below each Turbo slot is a plastic PCI card separator. The PCI card separator has two LEDs and a pull tab on the outer edge. The LED's provide power and status for the slot. The pull tab allows the PCI card to be easily removed.

rp5400 servers can access PCI slots 1,2 and 8-12. rp5450/3000 servers can access all PCI slots.

Follow these procedures to install graphics cards.

Step 1. Install HP-UX 11.0 Support Plus (IPR) 0006, June 2000 or later. This step ensures the appropriate HP-UX drivers are installed.

Step 2. Power down and unplug the rp54xx server.

18.732

NOTE DC voltages are present when the server is connected to AC power. Do not attempt to install or service: CPUs, Memory, PSMs, the Platform Monitor or PCI I/O cards installed in non-Turbo slots (1-6) while DC voltage is present. Failure to observe this warning may result in damage to the server.

Step 3. Make the right side of the server accessible for service.

Step 4. Using a Torx 15 screwdriver, loosen the captive screws on the right side panel. This panel has a label which shows which PCI I/O slots are available and the corresponding paths. The label shown below is for an rp5400.

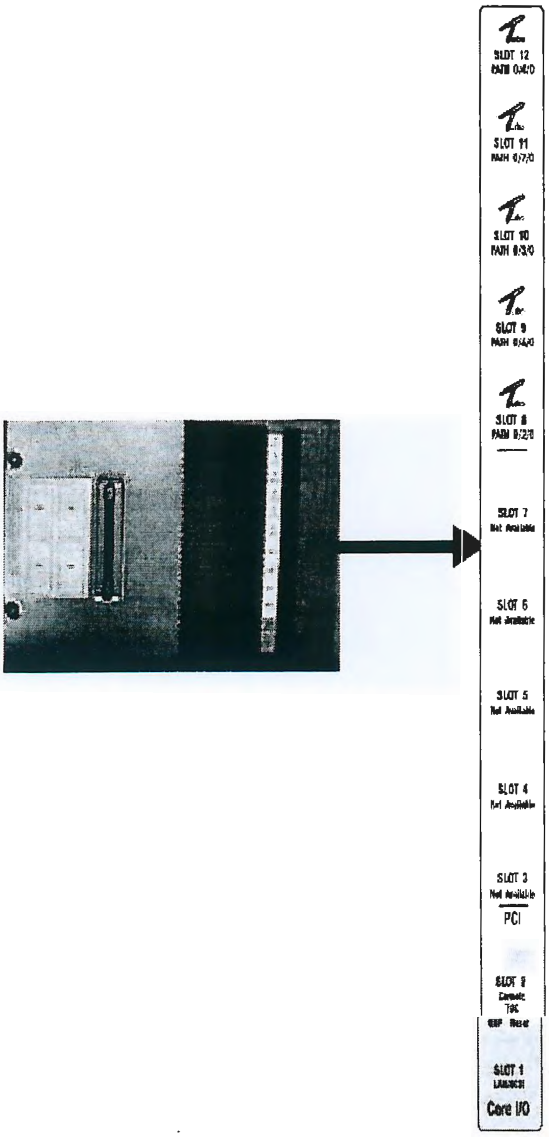

Step 5. Grasp the handle on the right rear panel and remove the panel from the side of the chassis. The 12 PCI slots, numbered 1-12 from bottom to top, will be in view.

Step 6. Remove the PCI slot cover from the slot that will receive the PCI card. To remove the PCI slot cover, slide the PCI slot cover away from the server.

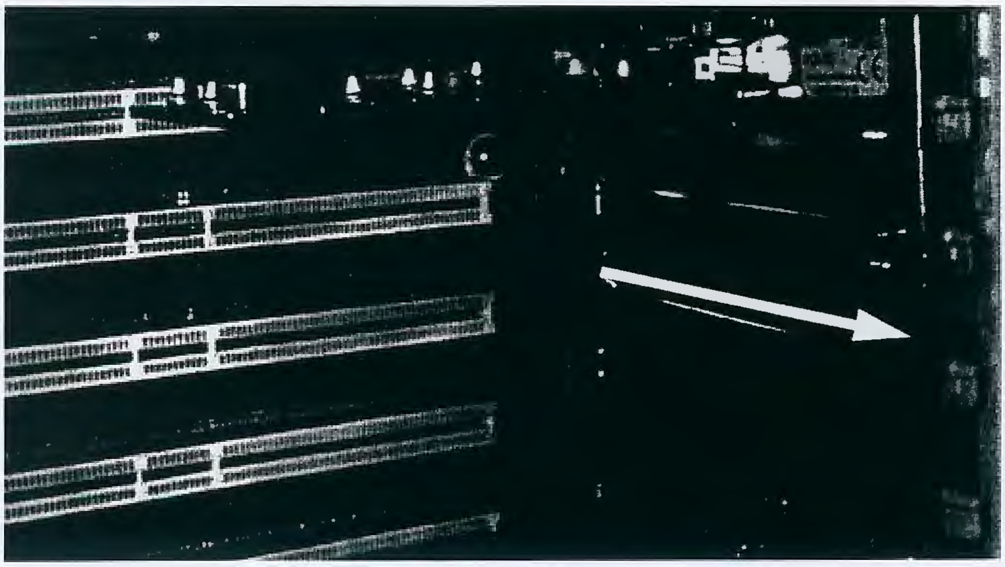

Step 7. Center the graphics card within the space created by removing the PCI I/O slot cover. Slide the card toward the edge connectors. Ensure the edge connectors on the card are in alignment with the connectors of the slot. Apply pressure to the card until it snaps firmly in place. Repeat process for USB card.

NOTE The graphics card must be installed in any Turbo slot while the USB will work in any slot. To reserve Turbo slots for high performance I/O cards, install the USB card in a non-Turbo slot

18.729

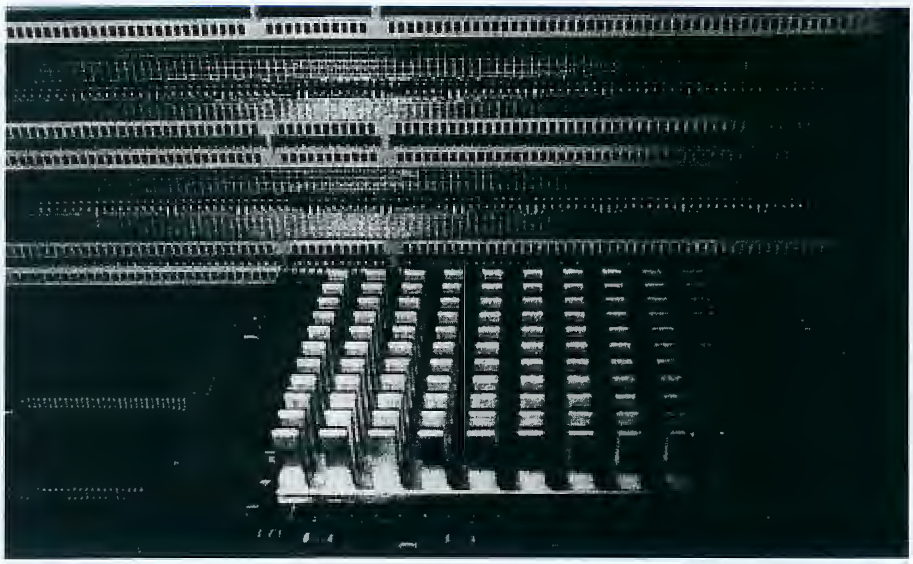

Step 8. At the rear of the chassis, connect the keyboard and mouse cables to the USB card. It does not matter which connector is used for the keyboard or mouse.

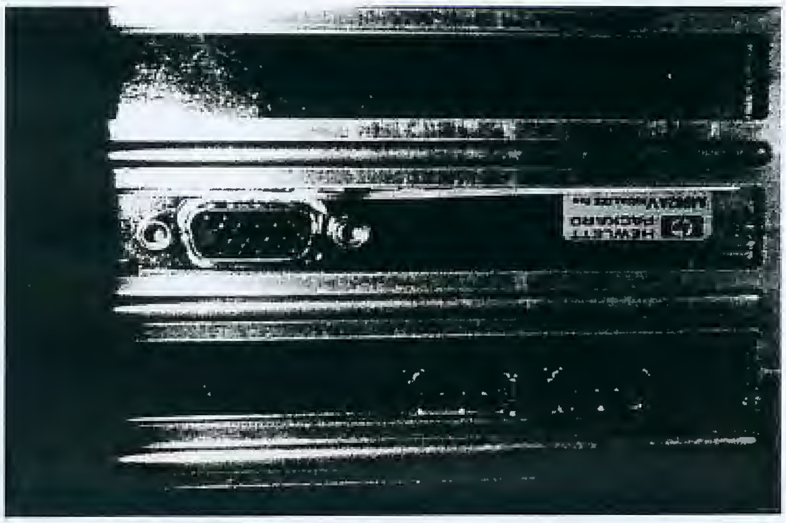

Step 9. Connect one end of the 15-pin video cable connector on the graphics card. This connector is labeled "Graphics Display" and "Video Out". Connect the other end of this cable to the graphics monitor.

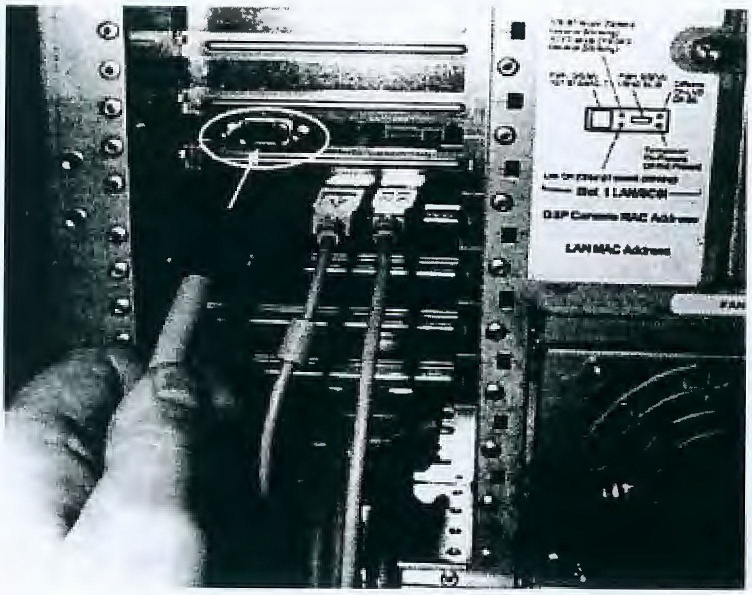

Step 10. Replace the right side panel and tighten the captive screws.

Step 11. For rack configurations, insert the rp54xx server back into the rack.

Step 12. For deskside enclosure configurations, repalce the deskside enclosure cover.

Step 13. Power the server on.

Step 14. Use the server firmware `in io` command to verify the graphics cards are recognized by the server. If AUTOBOOT is ON, it will be necessary to interrupt the boot process to get to the server firmware `Main Menu: Enter command or menu >` prompt.

Step 15. Boot HP-UX and run the `ioscan` utility to verify the system recognizes the new PCI card.

Step 16. Logon as root and install X/CDE/Motif if not already installed.

Graphics Troubleshooting

This section describes how to troubleshoot common problems encountered during installation or attempted use of graphics. The following system utilities can be used to display or set the graphics configuration:

- `/opt/graphics/common/bin/graphinfo` allows you to display the current graphics configuration and the graphics drivers that are being used.
- `/opt/graphics/common/bin/setmon` allows you to reconfigure the monitor type.
- The `display` menu of the HP-UX System Administration Manager (SAM) utility allows you to configure the X-Server and set the monitor type.
- On-line diagnostics provide information, verify and diagnose coverage for the graphics and USB cards.Off-line diagnostics do not exist for either the graphics or USB card.
- The HP-UX `ioscan` utility can be used to verify the HP-UX operating system recognized the hardware.

Symptom: CDE will not come up.

Step 1. Ensure /dev/crt was created. If not created, use `insf -e` to create.

Step 2. Ensure the system is at run level 3. Use `who -r` to determine run level. Use `init 3` to change to run level 3.

Step 3. Ensure dt is enabled. Use `/usr/dt/bin/dtconfig -e` to enable dt.

Step 4. Ensure /etc/dt/config/Xservers exists. If not, use `/usr/dt/config/dtrc.d/20_graph_conf` to create.

Step 5. Ensure the line: * Local local@console /usr/bin/X11/X :0 is not commented out of the /etc/dt/config/Xservers file.

Step 6. Reboot HP-UX.

Symptom: HP-UX does not recognize the graphics cards. `unknown` appears in the ioscan output for these cards

.

Step 1. Examine the output of the `swlist` command to ensure the correct version of HP-UX is installed.

Step 2. Update HP-UX as necessary

18.726

Installing Disk Drives

rp54xx servers support up to four optional internal hard drives. These drives must be installed in the following sequence:

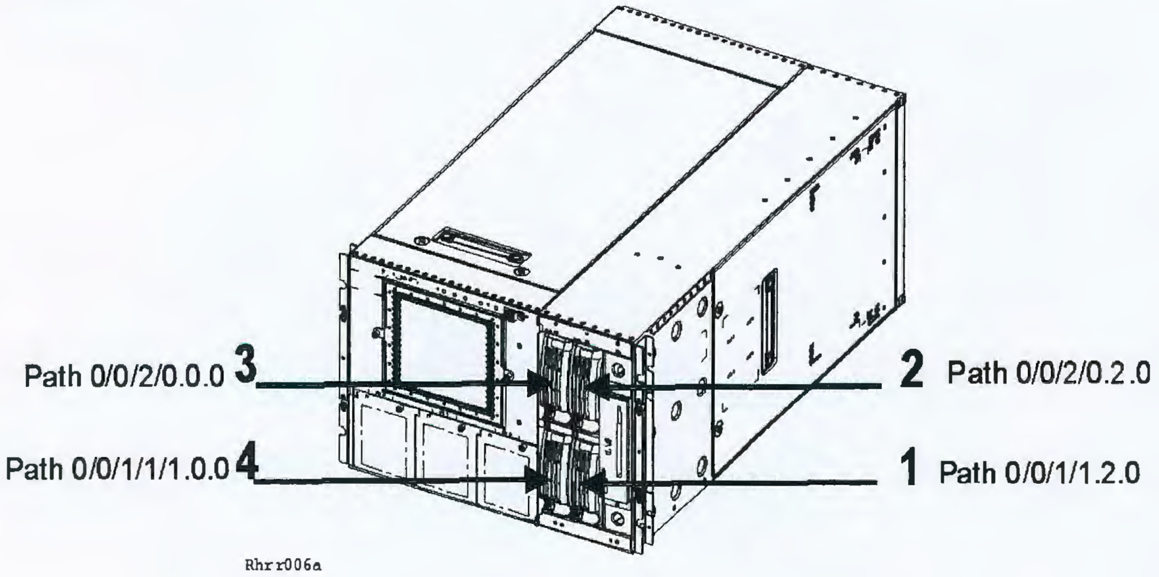

It is not necessary to shutdown the HP-UX operating system or power off the server to install a new disk. Follow this procedure to add internal hard disk drives to your rp54xx server.

Step 1. If a front bezel is installed on the face of the server, open the right-hand panel to gain access to the disk slots.

Step 2. Remove the disk drive slot cover.

Step 3. Insert the new disk drive into the slot until the rear connectors snap into place in the card guide. As shown in the following graphic, the notches at the top of the disk drives must snap over the small brackets in the disk bay to ensure a firm connection.

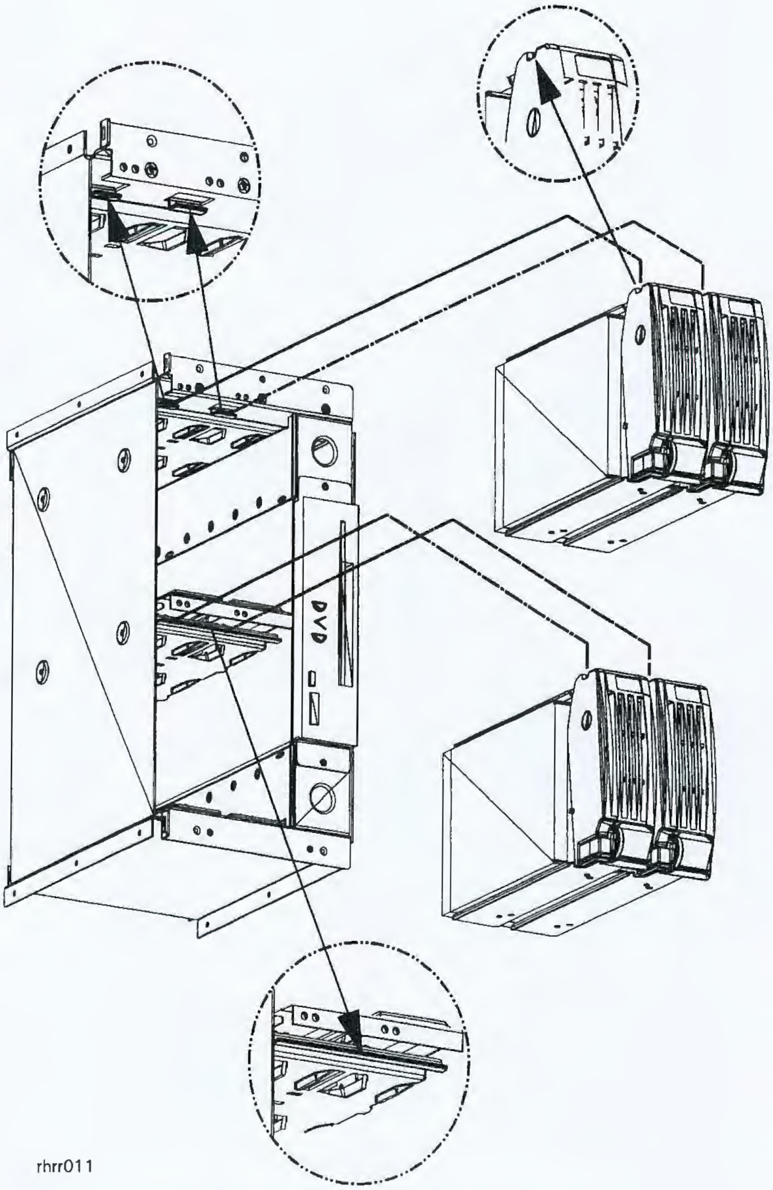

Step 4. Secure the connection by pushing the blue release lever closed.

Step 5. Refer to HP-UX documentation to configure the new disk.

18.724

18.723

4 Cable Connections

Core I/O Connections

The following paragraphs describe the indicators and connections of the rp54xx Core I/O. Core I/O consists of a LAN/SCSI card in slot 1 (lower slot in graphic) and a Guardian Service Processor (GSP) in slot 2 (upper slot in graphic). There are two versions of GSP, revision A and revision B.

Revision A GSP

The following graphic shows the indicators and connectors for the revision A GSP and LAN/SCSI Core I/O boards.

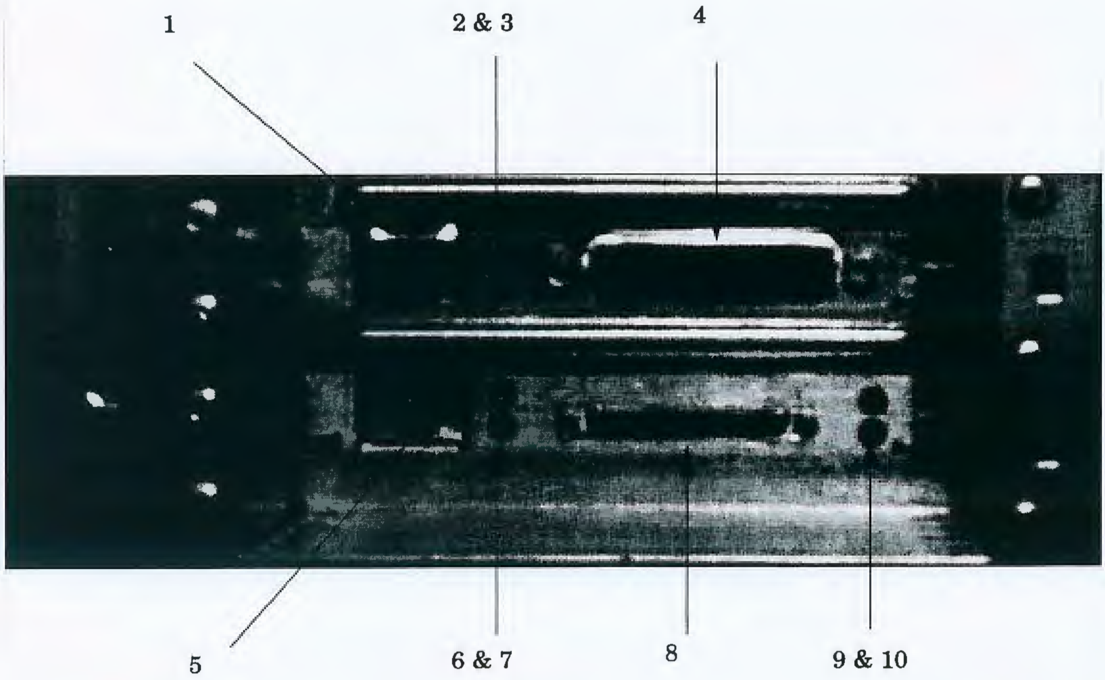

1. 10-Base-T LAN (RJ-45) Connector
   GSP LAN.
2. Green/Red (Upper LED)
   Green = GSP Power On.
   Flashing Green = LAN Receive.
   Red = Guardian Support Processor Test Failed.
3. Green/Red, (Lower LED)
   Green = Link OK.
   Flashing Green = LAN Transmit.
   Red = Guardian Support Processor Test Failed.
4. Console/UPS/Remote Connector (D-Type 25-Pin female).
   Requires an A5191-63001 "W" adapter cable
5. 10/100 Base-T = Primary LAN (RJ-45) Connection
   Path 0/0/0/0

18.721
J

6. Green/Yellow (Upper LED)
   Green = 100 Base-T Mode
   Green Blinking = 100 Base-T Receiving
   Amber = 10 Base-T Mode
   Amber Blinking = 10 Base-T Receiving

7. Green (Lower LED)
   Green = Link OK (10/100 Base-T Mode indicated by LED #6)
   Green Blinking = Transmitting

8. Ultra-2 SCSI Connector (68-Pin VHDCI SCSI)
   Path 0/0/1/0

9. SCSI Mode (Green, Upper LED)
   On = Low Voltage Differential (LVD) Mode.
   Off = Single Ended Mode.

10. SCSI Terminator Power (Amber, Lower LED)
    On = Terminator power present
    Off = Terminator power Not present.

Revision B GSP

The following graphic shows the indicators and connectors for the revision B GSP and LAN/SCSI Core I/O boards.

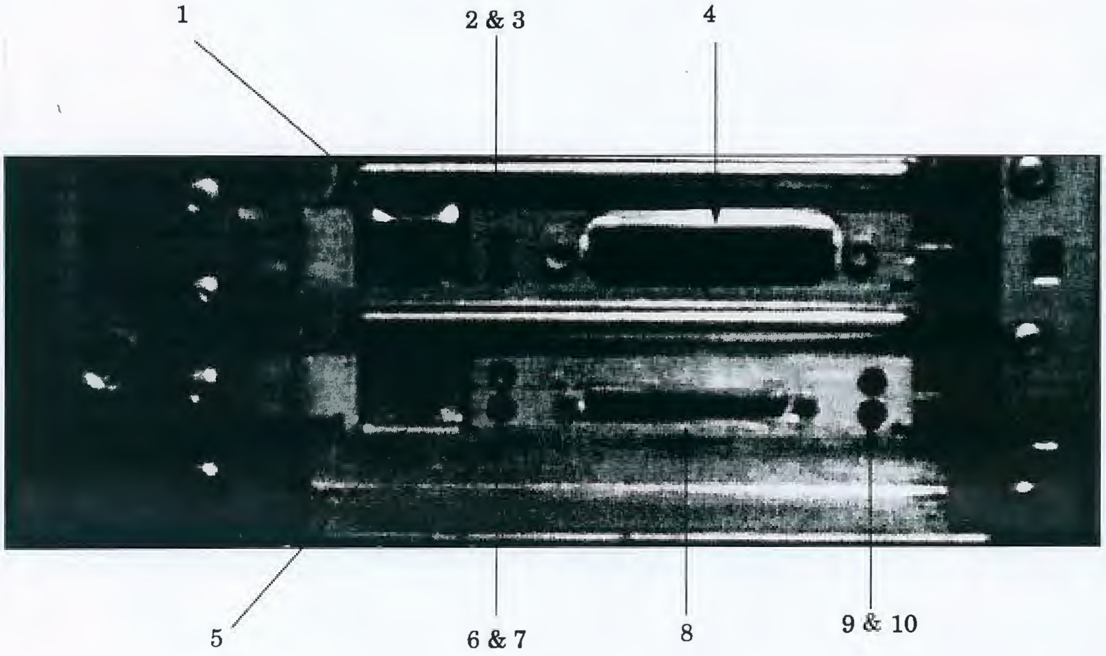

18.720

1. 10/100-Base-T LAN (RJ-45) Connector.
   GSP LAN.
2. Green/Red (Upper LED).
   Green = GSP Power On.
   Red = Guardian Support Processor Test Failed.
3. Green/Yellow, (Lower LED).
   Green = 100 Base-T Link OK.
   Flashing Green = 100 Base-T LAN Activity.
   Yellow = 10 Base-T Link OK.
   Flashing Yellow = 10 Base -T LAN Activity.
4. Console/UPS/Remote Connector (D-Type 25-Pin female).
   Requires an A6144-63001 "M" adapter cable.
5. 10/100 Base-T = Primary LAN (RJ-45) Connection.
   Path 0/0/0/0.
6. Green/Yellow (Upper LED).
   Green = 100 Base-T Mode.
   Green Blinking = 100 Base-T Receiving.
   Amber = 10 Base-T Mode.
   Amber Blinking = 10 Base-T Receiving.
7. Green (Lower LED).
   Green = Link OK (10/100 Base-T Mode indicated by LED #6).
   Green Blinking = Transmitting.
8. Ultra-2 SCSI Connector (68-Pin VHDCI SCSI).
   Path 0/0/1/0.
9. SCSI Mode (Green, Upper LED)
   On = Low Voltage Differential (LVD) Mode.
   Off = Single Ended Mode.
10. SCSI Terminator Power (Amber, Lower LED)
    On = Terminator power present
    Off = Terminator power Not present.

Guardian Service Processor (GSP) Overview

This section provides an overview of the Guardian Service Processor (GSP). The GSP is an always on, dedicated service processor that monitors system power, cooling and configuration, and provides console communications. Power and cooling information is obtained via an interface to the platform monitor card. Configuration information is obtained via connection to the Serial Presence Detect (SPD) bus. The GSP can only be installed in slot 2 and must be present for the server to power up.

The GSP has downloadable firmware which can be updated independent of the HP-UX operating system. GSP firmware updates can occur anytime the GSP is active. If the DC power switch is OFF, the GSP is still operational and GSP firmware updates can still occur. GSP firmware updates may be performed by customers.

If the GSP becomes hung, it is possible to reset the GSP without impacting the server. The GSP may be reset via the GSP RESET button on the right side of the card. The PCI cover panel (right side panel) must first be removed to allow access to the right side if the GSP card.

The GSP has two connectors on the bulkhead. An RJ-45 for LAN connections and a female DB25 connector for RS-232 connections. Attach either a "W" or an "M" cable to the DB25 connector to provide individual output for CONSOLE, REMOTE and UPS.

To access the GSP from the local ASCII console, type control b and the `GSP>` prompt will appear. It may be necessary to type control Ecf first. To exit the GSP, type `GSP>co`.

The GSP was originally a core component of the revision A rp5400 (A5576A) and rp5450 (A5191A) servers. Beginning with introduction of the revision B rp5400 (A5576B), rp5450 (A5191B), and rp5470 servers the GSP became a separate, must order product (A6696A).

There are two revisions of rp54xx GSP: rev A (A6696A) and rev B (A6696B). Both GSPs must be installed in order for the server to power up.

GSP LAN

This LAN is exclusively for LAN console access and is not configurable via HP-UX. The LAN is configured via GSP commands. Hostname, IP, gateway and subnet mask parameters may be set via the `GSP>lc` command. The GSP may also initiate ping via the `GSP>xd` command.

GSP RS-232

The DB25 connector on the GSP is used for RS-232 communications to a local console (via CONSOLE connector), a remote console via modem (REMOTE connector), and a UPS (UPS connector). The baud rate, term type, etc., of the CONSOLE and REMOTE ports are configured via `GSP>ca` command.

The GSP supports VT100 and HPTERM terminal emulation. For correct communications, the GSP and RS-232 device must use the same terminal emulation and baud rates.

GSP Features

The revision A GSP provides a 10 base-T LAN connector for LAN console access and a DB-25 connector to which the A5191-63001 W-cable connects. The W-cable provides REMOTE, UPS, and CQNSOLE DB-9 connectors.

Features of the revision A GSP are:

- 10 Base-T LAN connector for revision A GSP

- 10/100 Base-T LAN connector for revision B GSP
- On-board processor dedicated to GSP functions
- Error logging and notification
- Display of system alerts and selftest chassis codes
- Powered by 15 VDC housekeeping power that is present when the front panel switch is off
- Power and configuration monitoring
- RS-232, LAN, REMOTE and WEB console access
- Administrator and user security
- Alphanumeric paging.

There are two revisions of rp54xx GSP: revision A (A6696A) and revision B (A6696B). Due to significant hardware differences between the revision A and B GSP, each GSP requires it's own firmware. Revision A GSP firmware can only be installed in a revision A GSP and revision B GSP firmware can only be installed in a revision B GSP. The hardware differences are necessary to incorporate the embedded web access, 10/100 Base-t LAN, and faster GSP processor.

The GSP provides four types of console access: RS-232, Remote, LAN and Web. Console information is mirrored to all four console types. Refer to *Configure System Console* for more information.

The GSP was originally a core component of the revision A rp5400 (A5576A) and rp5450 (A5191A) servers. Beginning with introduction of the revision B rp5400 (A5576B), rp5450 (A5191B), and rp5470 servers the GSP became a separate, must order product (A6696A).

Revision A GSP

The revision A GSP is identified by product number A6696A and part numbers: A5191-60012, A5191-69012,. and A5191-69112.

The revision A GSP requires a "W" cable to be attached to the DB25 connector. The part number of the "W" is A5191-63001. The "W" cable provides female DB9 connectors for CONSOLE, REMOTE and UPS. The maximum supported baud rate for the CONSOLE and REMOTE connectors is 19200 baud and 1200 baud for the UPS.

The paths for the CONSOLE, UPS, and REMOTE are 0/0/4/0.0, 0/0/4/0.1, and 0/0/4/0.2 respectively.

For the rev A GSP, the web console is accomplished by shipping one J3591A Secure Web Console with each rp54xx server. The Secure Web Console can be used in place of an ASCII console to provide console access via a web connection. If you are installing an rp54xx server that does not have an ASCII console, you may use the Secure Web Console as the console. However, you must first configure the Secure Web Console. Refer to *Secure Web Console Installation and Configuration* for more information on SWC Installation/Configuration.

Revision B GSP

The revision B GSP is identified by product number A6696B and part numbers: A6144-60012, A6144-69012, and A6144-69112.

The revision B GSP requires an "M" cable to be attached to the DB25 connector. The part number of the "M" cable is A6144-63001. The "M" cable provides female DB9 connectors for CONSOLE, REMOTE, and UPS. The maximum supported baud rate for the CONSOLE and REMOTE connectors is 38400 baud and 1200 baud for the UPS.

The paths for the CONSOLE, UPS, and REMOTE are 0/0/4/1.0, 0/0/4/1.1 and 0/0/4/1.2 respectively.

18.717

Configure System Consoles

rp54xx servers provide RS-232, REMOTE, LAN and WEB console access. All console access involves the Guardian Service Processor (GSP). rp54xx servers use either a revision A or revision B GSP. Below is an illustration of the console access provided by the revision A GSP.

Revision A GSP Console Access

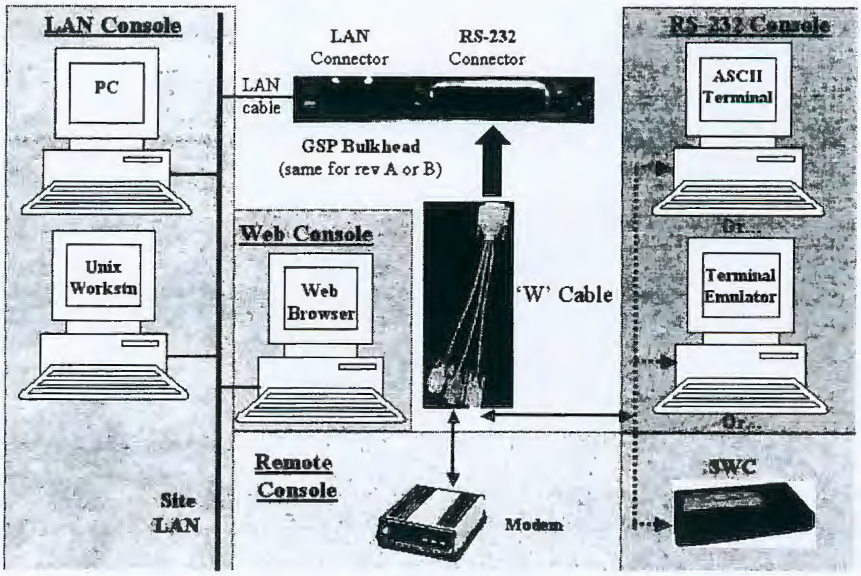

The revision B GSP has embedded web access, eliminating the need for an external Secure Web Console (SWC). Below is an illustration of the console access provided by the revision A GSP.

18.736

Revision B GSP Console Access

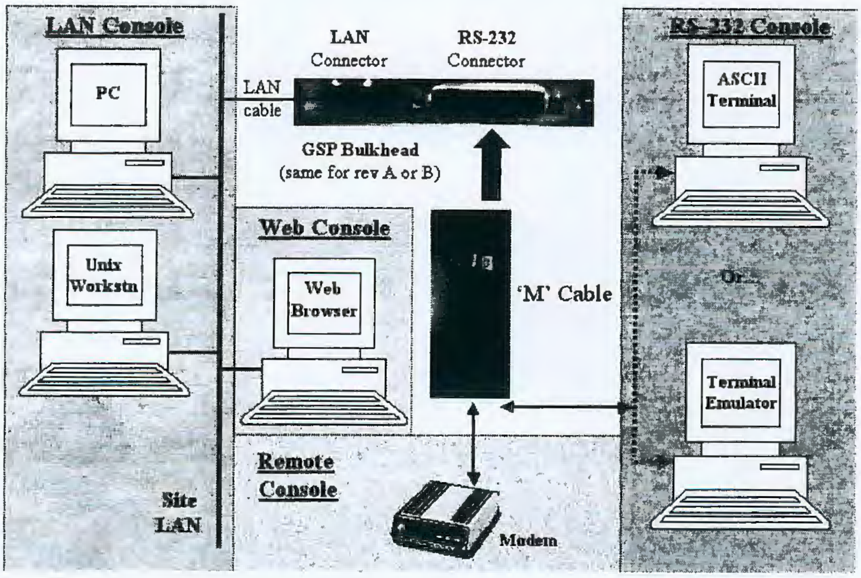

GSP Cables

Both the revision A and B GSPs provide a DB-25 connector for RS-232 communications. Connect the A5191-63001 W- cable to the revision A GSP DB25 connector or connect the A6144-63001 M-cable to the revision B GSP DB25 connector. These cables provide individual DB9 connectors for REMOTE, UPS and CONSOLE.

NOTE Use the A5191-63001 W-cable with revision A GSP and A6144-63001 M-cable with revision B GSP *only*. Failure to use the right cable can result in reduced functionality.

The W and M-cables are slightly different. The W-cable has full RS-232 capability on the REMOTE and CONSOLE connectors and partial RS-232 capability on the UPS connector. The M-cable has full RS-232 capability on the REMOTE and UPS connectors and partial RS-232 capability on the CONSOLE connector. The cable change is to be consistent with rp54xx functionality. The cables are different colors to easily tell them apart. The W-cable is gray and has part number A5191-63001. The M-cable is black and has part number A6144-63001.

Configure RS-232 Console

The physical connections for an RS-232 console include attaching the correct cable to the GSP. Next, connect the 24542G cable (supplied) to the CONSOLE connector and the serial port of the ASCII console. A personal computer (PC) running terminal emulation software may be used in place of an ASCII console. Refer to the illustration below for RS-232 console.

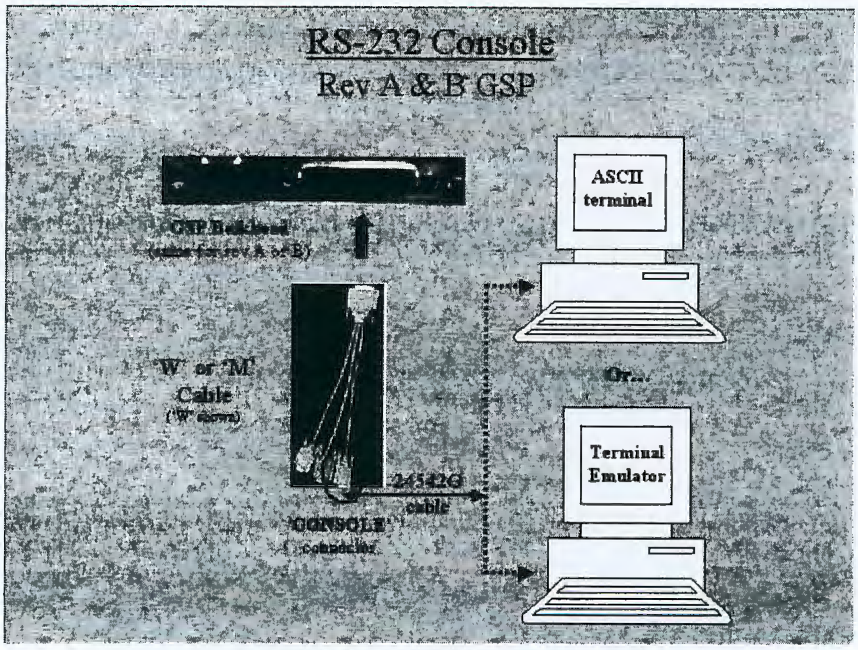

Step 1. The GSP is located in slot 2 of the rp54xx' rear card cage. Connect the 25-pin end of:

- the A5191-63001 W-cable to the 25-pin connector on the revision A GSP card (A5191-60012) OR
- the A6144-63001 M-cable to the 25-pin connector on the revision B GSP card (A6144-60012)

Step 2. Connect the 9-pin "Console" connector of either the W or M-cable to the 9-pin D-type connector of a 24542G RS-232 cable.

Step 3. Connect the 25-pin end of the 24542G serial cable to the serial/RS232 port on the ASCII console. (RS232 Serial Port labeling may vary depending on manufacturer.)

Step 4. Connect the System Console to input AC power.

Step 5. Turn the System Console AC power switch to ON.

After the physical connections have been made, configure the ASCII console. When using the C1099A Terminal Console, the default settings are recommended. Refer to the C1099A Terminal Console operating manual for instructions on how to obtain default settings.

The HP 700 series console may also be used as an ASCII console. Both the C1099A Terminal Console and S700 consoles support HPterm and VT100 emulations. The emulation of the GSP and ASCII console need not match for communications between them to occur. However, to ensure proper communications, HP recommends the ASCII console and GSP use the same emulation. HP also recommends that other configurable parameters on the GSP match those of the ASCII console. Baud rate, start/stop bits, etc... The default emulation of the GSP is VT100.

Below is a procedure to configure a HP 700 serial console for VT100 emulation.

HP 700 Series System Console Configuration

The following describes the steps required to configure the HP 700 series terminal for VT-100 mode for operation with an rp54xx server.

Although any terminal capable of operating in VT-100 mode can be used, the HP700 series terminal is used here as an example because it is fairly common and it's configuration is typical of many terminals currently in use.

HP700 VT-100 Mode Configuration The following procedure outlines the steps to configure the HP700 series terminal for VT100 operation.

NOTE You may use either the arrow keys or the tab key to move between the setting options on the screen.

1. Press **[config keys]** function key. **[f8]**
2. Press **[terminal config]** function key. **[f5]**
3. Move to *Terminal ID* and enter "vt100".
4. Move to *Set TermMode* and, using the **[Prev]** and **[Next]** keys, select "EM100".
5. Press the **[config keys]** function key. **[f8]**
6. Press the **[ansi config]** function key. **[f6]**
7. Move to "multipage" and, using the **[Prev]** and **[Next]** keys, select "yes". (Enables screen scrolling).
8. Move to *Backspace Del* and, using the **[Prev]** and **[Next]** keys, select "Backspace/Del".
9. Move to *EM100 ID* and, using the **[Prev]** and **[Next]** keys, select "EM100".

Configure the Asynchronous Values of the GSP

After the ASCII console has been configured and physical connections made, make any necessary changes to the asynchronous values of the GSP.

1. Access the GSP with the ctrl+b entry. The GSP will respond with a `GSP>` prompt.
2. At the GSP prompt, enter the Configure Asynchronous (ca) command:

The ca command will start a series of prompts. Respond to each prompt with the appropriate information.

Example 4-1 CA command

```
Leaving Console Mode - you may lose write access. When Console Mode returns, type
^Ecf to get console write access.

GSP Host Name:  fesrhapgsp

GSP> ca

CA This command allows you to modify the local and remote modem serial
portconfigurations.

Current configuration settings:

Local Console Serial Port bit rate: 9600 bits/s

Local Console Serial Port Flow Control: Software
Local Console Serial Port Terminal Type: vt100
Remote Console Serial Port Modem Protocol: CCITT  Remote Console Serial Port Modem
bit rate: 19200 bits/s
Remote Console Serial Port Modem Flow Control: Software
Remote Console Serial Port Modem Transmit Configuration Strings: Enabled
Remote Console Serial Port Modem Presence: always connected
Do you want to modify the Local Console Serial Port settings? (Y/[N])
Do you want to modify the Remote Console Serial Port Modem settings? (Y/[N])
GSP Host Name:  fesrhapgsp
```

If necessary, use the GSP help facility by typing `GSP>he`. Once in the help facility, type the command need help with. Use LI for a list of commands.

The following baud rates are recommended for the revision A GSP:

- Console: 19200
- Remote: 19200
- UPS: 1200

The following baud rates are recommended for the revision B GSP:

- Console: 38400
- Remote: 38400
- UPS: 1200

18.732

Configure Remote Console

The remote console allows console access via modem connections. Below is an illustration of the REMOTE console.

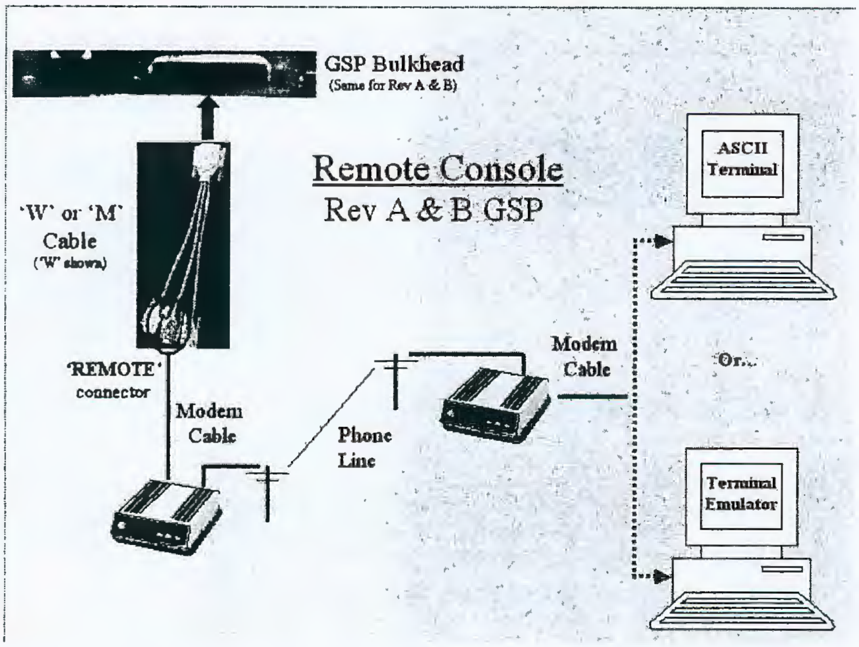

The GSP>ca command is used to configure asynchronous settings for the REMOTE console. Baud rates and emulations should match between the modems, remote ASCII terminal and the GSP. Refer to, "Configure RS232 Console" for information about setting these values.

18.781

Configure the LAN Console

The LAN console allows you to access the console from the LAN using TelNet or http (revision B GSP only) protocols. Below is an illustration of the LAN console.

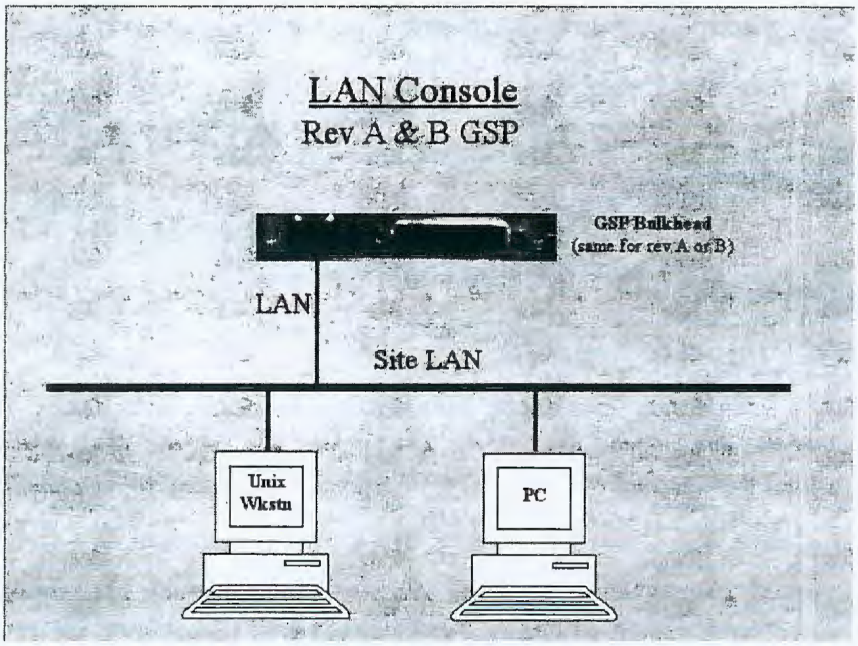

The configuration of the LAN console of both the revision A and B GSPs may be done from either an ASCII console or the external Secure Web Console. For the revision B GSP, an IP may be assigned via LAN by pinging the LAN from a PC or workstation.

Configuring the GSP LAN Port via an ASCII console

The LAN port of the GSP allows connection via TelNet or http connections. Once the LAN parameters are configured, the console may be accessed via a TelNet connection or via a web browser (revision B GSP only). The default IP of the GSP LAN is 127.0.0.1.

NOTE The GSP has a separate LAN port from the system LAN port. It will need a separate LAN drop, IP address, and networking information from the port used by HP-UX.
Before starting this procedure, you will need to know the following information:

- I.P. address (for GSP)
- Subnet mask
- Gateway address
- Hostname (this is used when messages are logged or printed)

To configure the GSP LAN port, perform the following steps:

1. Access the GSP with the ctrl+b entry.
2. At the GSP prompt, enter the LAN Configuration (lc) command:

```
GSP> lc
```

The lc command will start a series of prompts. Respond to each prompt with the appropriate information.

Example 4-2 LC command

```
Leaving Console Mode - you may lose write access. When Console Mode returns, type
^Ecf to get console write access.

GSP Host Name:  fesrhapgsp

GSP> lc

LC This command allows you to modify the LAN configuration.

Current configuration:

MAC Address: 0x00306e050a63

IP Address: 15.8.133.185

GSP Host Name: fesrhapgsp

Subnet Mask: 255.255.248.0

Gateway: 15.8.128.1

Web Console Port Number: 2023

Do you want to modify the LAN configuration? (Y/[N])

GSP Host Name:  fesrhapgsp
```

The revision B GSP introduces a configurable Web Console Port Number parameter. The default value is 2023. Once the GSP LAN is configured, it is accessible via either TelNet or web connections.

Configuring the GSP LAN Port via LAN

The revision B GSP LAN port can be assigned an IP address without using the LAN Configuration (lc) command via an ASCII console. This section describes how to assign the IP address allowing web access. Once web access is accomplished, use the lc command to configure remaining network parameters.

NOTE The GSP LAN port is separate from the system LAN port. It will need a separate LAN drop, IP address, and networking information from the port used by HP-UX.
Before starting this procedure, you will need to know the following information:

- I.P. address (for GSP)
- Subnet mask
- Gateway address
- Hostname (this is used when messages are logged or printed)

To configure the GSP LAN port via LAN, perform the following steps:

NOTE The GSP *must* be on the same subnet as the system being used to remotely configure the LAN port. If it is not, the remote configuration will be unsuccessful.

Step 6. Determine the MAC address of the revision B GSP by examining the GSP MAC address label on the rear of the server.

Step 7. Use the route add command to add the I.P address of the GSP and remote system to the router.

18.709

Step 8. Use the arp command to add an ARP entry for the IP address using the GSP MAC address.

- For HP-UX systems, the format of the MAC address is 00:30:6e:05:0a:ea
- For MS DOS systems, the format of the MAC address is 00-30-6e-05-0a-ea

Step 9. Use the ping command to assign the I.P address for the GSP.

Step 10. The revision B GSP is now accessible via LAN. Access the GSP and configure remaining network parameters using the LAN Configuration (lc) command:

```
GSP> lc
```

The lc command will start a series of prompts. Respond to each prompt with the appropriate information.

Example 4-3 **LAN Configuration from a PC**

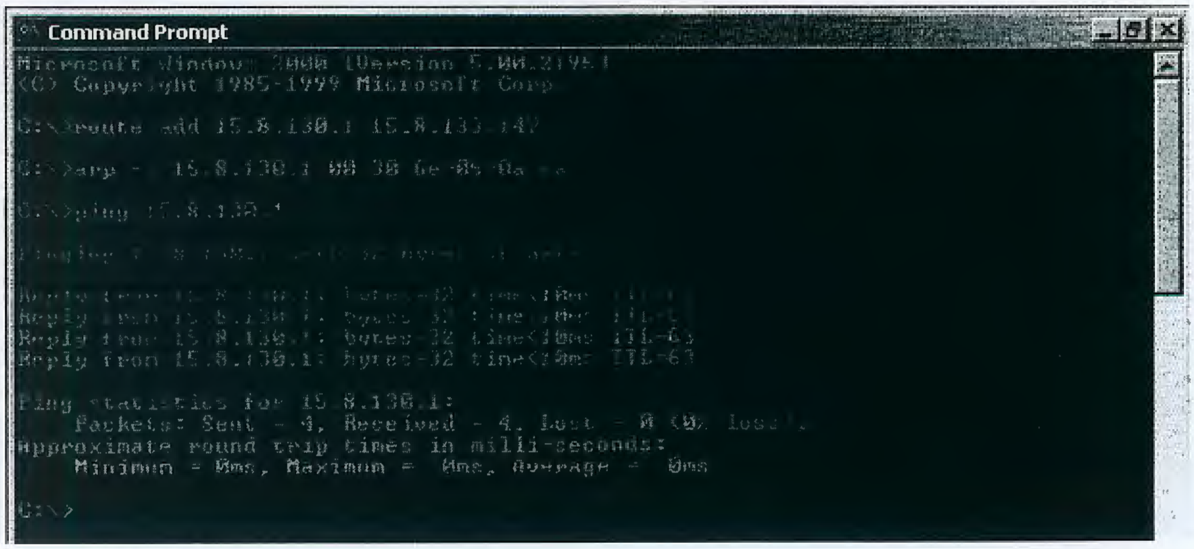

```
Microsoft Windows 2000 [Version 5.00.2195]
(C) Copyright 1985-1999 Microsoft Corp.

C:\>route add 15.8.130.1 15.8.133.147

C:\>arp -s 15.8.130.1 00-30-6e-05-0a-ea

C:\>ping 15.8.130.1

Pinging 15.8.130.1 with 32 bytes of data:

Reply from 15.8.130.1: bytes=32 time<10ms TTL=63
Reply from 15.8.130.1: bytes=32 time<10ms TTL=63
Reply from 15.8.130.1: bytes=32 time<10ms TTL=63
Reply from 15.8.130.1: bytes=32 time<10ms TTL=63

Ping statistics for 15.8.130.1:
    Packets: Sent = 4, Received = 4, Lost = 0 (0% loss),
Approximate round trip times in milli-seconds:
    Minimum = 0ms, Maximum =  0ms, Average =  0ms

C:\>
```

Configure the Web Console

For the revision A GSP, the web console is accomplished via the J3591A Secure Web Console. Below is an illustration of the web console for the revision A GSP.

Refer to, "Install a Secure Web Console" for more information on Secure Web Console installation and configuration.

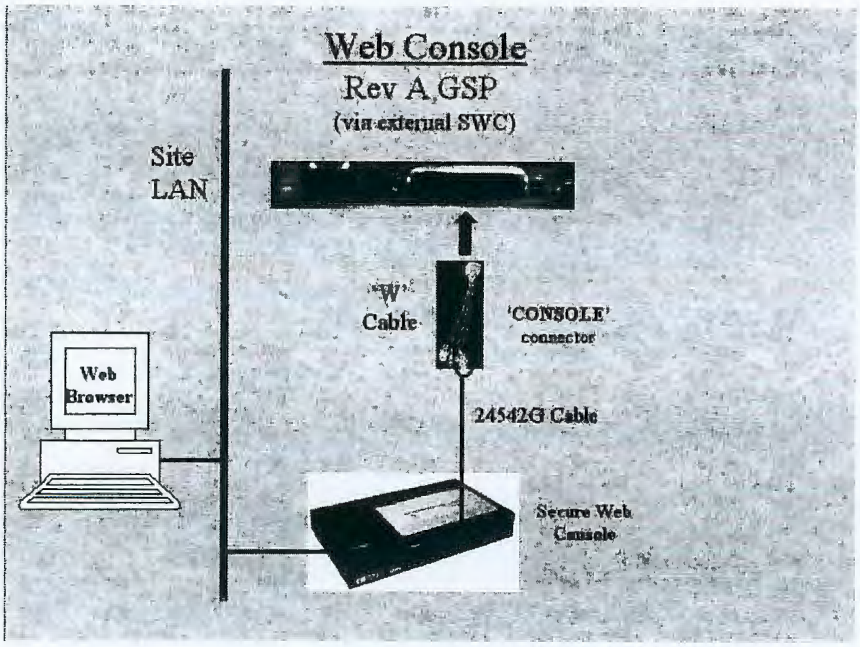

For the revision B GSP, the web console is an embedded feature. The steps to configure a web console are the same as configuring a LAN console. Refer to, "Configure the LAN Console."

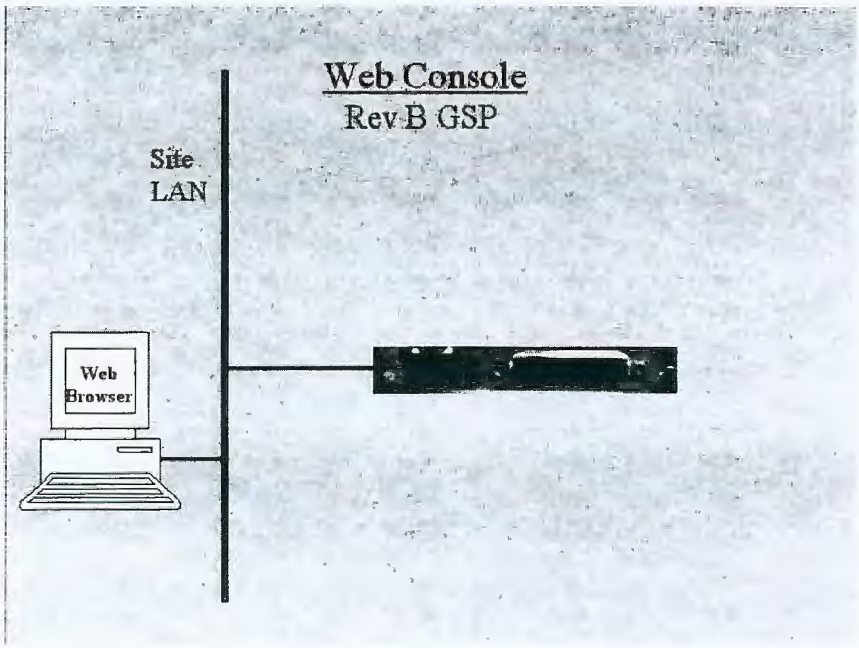

Once the LAN has been configured, access the web console by pointing a web browser, on the same subnet, to the IP of the GSP LAN.

Two browser windows will appear: a window with a white background and the HP invent logo and a separate GSP window with a black background.

Example 4-4 GSP Browser Window

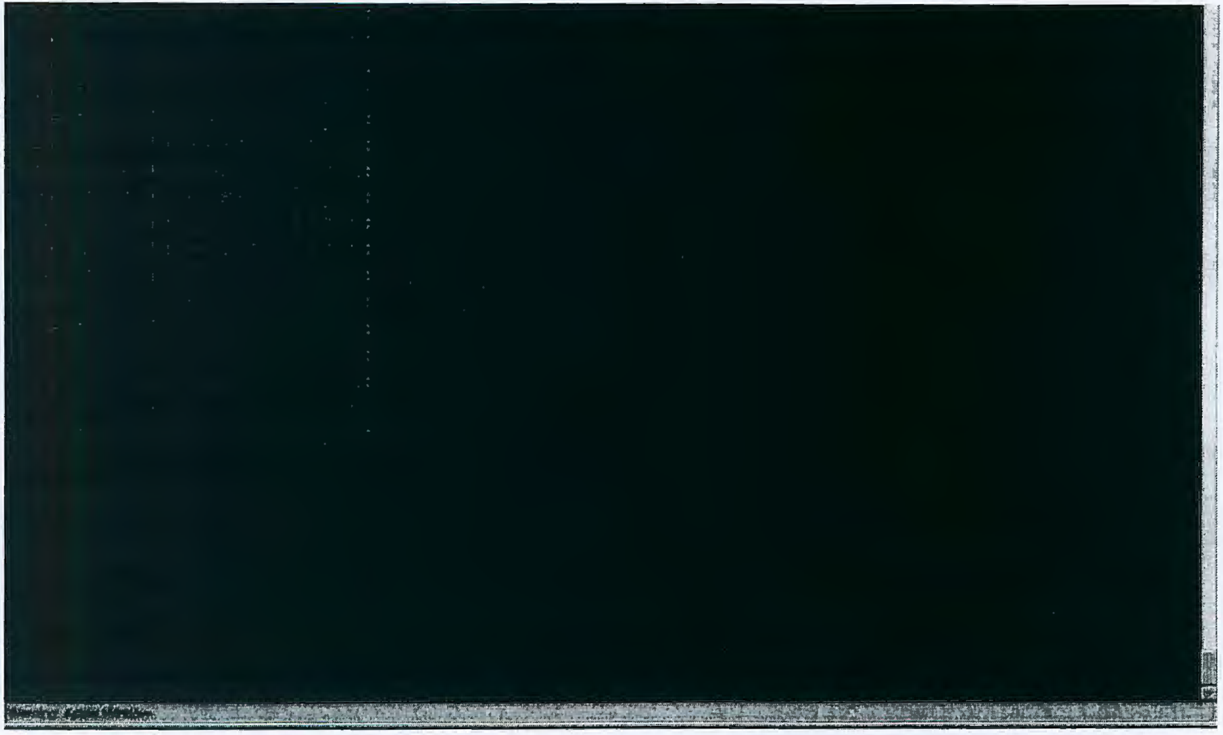

Use the SETTINGS menu bar to configure web browser emulation. The GSP window also has its own HELP facility.

18.706

Example 4-5 GSP Web Browser Help Screen

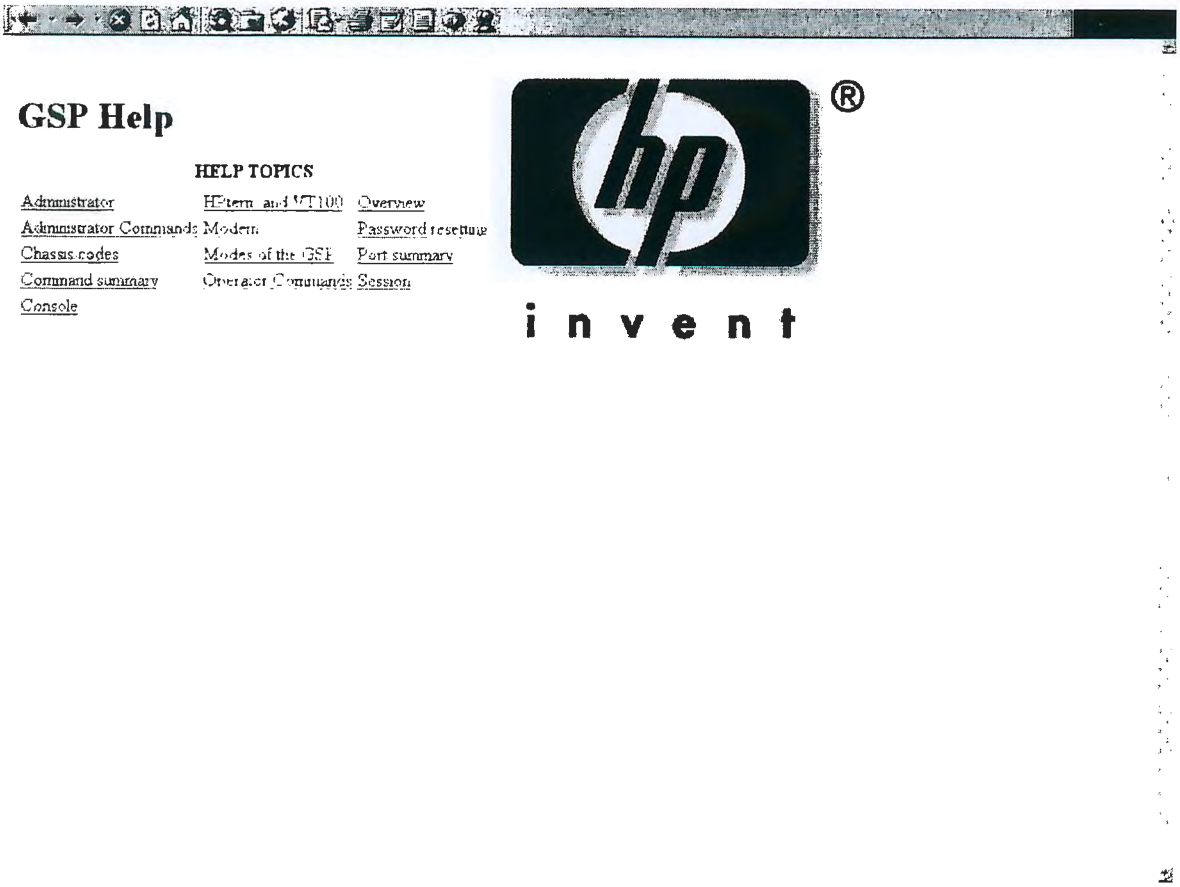

When the separate GSP window is closed, it appears in the HP invent window with Zoom In/Out above it. Click on the Zoom In/Out bar to generate a separate GSP window.

Example 4-6 Combined GSP Browser Window

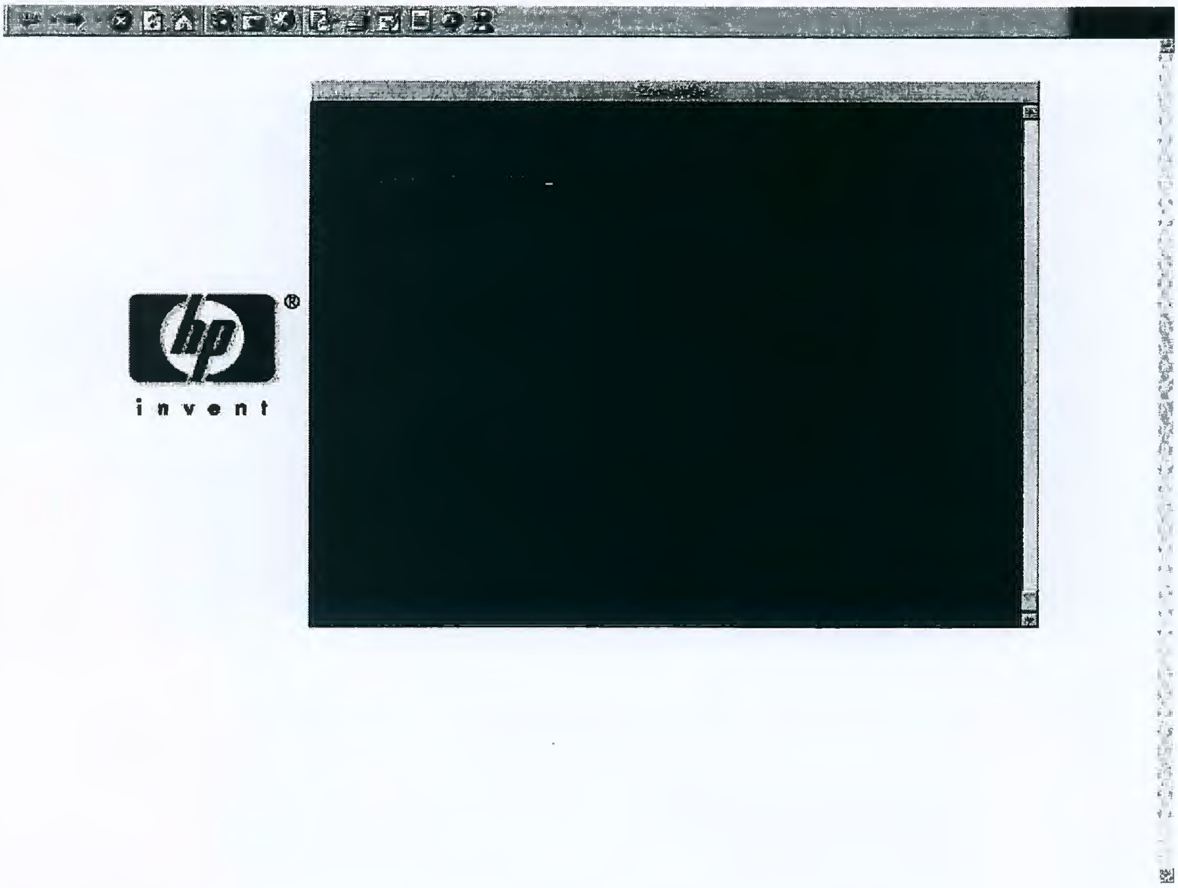

There is not a separate administration "layer" when using the embedded web access of the revision B GSP. Web console access via the external Secure Web Console required that you first logon to the SWC, then click on ACCESS CONSOLE. User configuration was also performed at the Secure Web Console. However, the revision B GSP web console does not require this additional step. When you point the web browser at the IP of the GSP LAN, you are directly connected to the GSP. The web console part of the GSP employs the same users as the GSP.

Secure Web Console Installation

The following section describes installation of the HP Secure Web Console on inside of the rear door of a rack-mounted rp54xx server.

For technical, installation, and configuration instructions for the Secure Web Console, refer to the following URLs on the Internet:

General information:
http://www.hp.com/

Documentation:
http://www.docs.hp.com

NOTE Either the system console (HP series 700 terminal) or the HP Secure Web Console may be installed on an rp54xx server, but not at the same time. Both console types use the same DB9 type LAN *Console* connector.

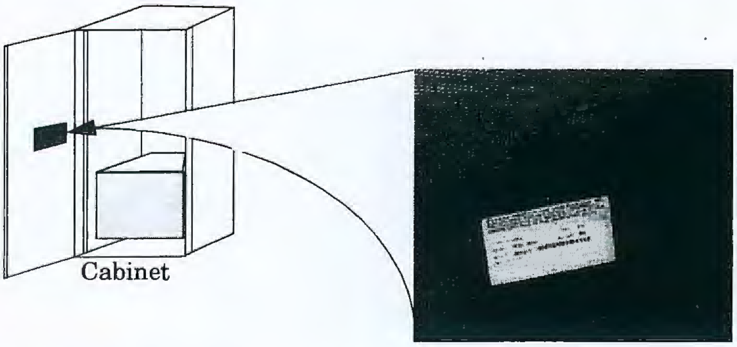

To install the HP Secure Web Console on an rp54xx server, refer to the previous figure and the HP Secure Web Console documentation, then proceed as follows:

Step 1. Install the wire mounting bracket by carefully inserting the two top prongs through the vent grill in the rear door of the cabinet as shown above. Position the bracket toward the hinge side of the door.

Step 2. Place the Secure Web Console power supply into the bottom portion of the wire mounting bracket as shown.

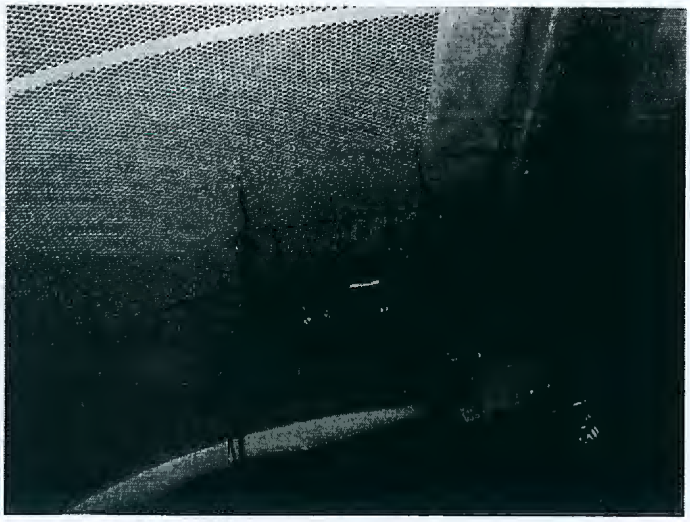

Step 3. Connect one end of the power cable to the power supply where indicated.

Step 4. Position the Secure Web Console unit in the mounting bracket.

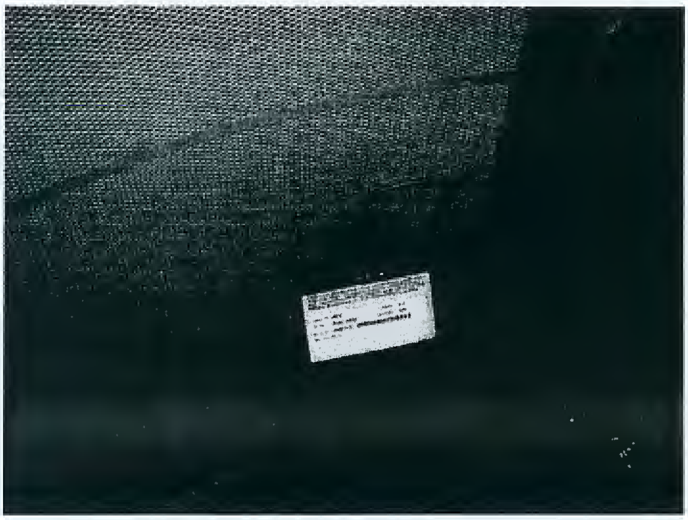

Step 5. Connect the DC out cable from the power supply to the Secure Web Console.

Step 6. Connect one end of the AC power cord (supplied) to the Secure Web Console power supply.

Step 7. Connect the other end of the AC power cord to an available receptacle.
On a PDU if in a cabinet.
In an available wall outlet if in a Deskside enclosure.

18.702

Step 8. Connect the 9-pin end of the RS-232 cable (Supplied) to the connector labeled CONSOLE on the A5591-63002 “W-type” adapter cable.

CAUTION To prevent unauthorized access to your rp54xx system, do not connect the other end of the serial cable to the Secure Web Console until both the server and the Web Console have both been fully configured.

Step 9. Connect one end of a LAN cable to RJ-45 connector labeled 10-Base-T on the Secure Web Console.

Step 10. Connect the other end of the same LAN cable to your site LAN.

Step 11. Configure the Secure Web Console in accordance with the documentation that was provided with it or refer to http://www.docs.hp.com

Step 12. Once the Web Console has been properly configured, the remaining end of the serial cable between the server and the Web Console may be connected.

18.701
3

GSP Configurable Parameters

Once a system console is configured, additional GSP parameters may be set. For a complete list use the `GSP>he` command to access the on-line help facility.

Examples of three configurable parameters follow.

Adding Users

The GSP provides a maximum of 20 users (one administrator and 19 operators). By design, the first user added to the GSP becomes the GSP administrator. Only the GSP administrator can add or remove users or change the GSP configuration.

NOTE Before starting this procedure, you will need to know the following information:

- User's name
- Organization's name
- Login name
- User's password

To add a user, perform the following steps:

1. Access the GSP with the ctrl+b entry.
2. At the GSP prompt, enter the Security options and access control (SO) command:

   GSP> so

3. The first prompt you will see with the so command is for GSP wide parameters:

```
GSP wide parameters are:
Login Timeout: 1 minutes.
. Number of password Faults allowed: 3
. Flow Control Timeout: 5 minutes.
Do you want to modify the GSP wide parameters? (Y / [N]) __
```

   At this point you can modify the GSP wide parameters, or continue with adding users. To add users, respond N for no.

NOTE If this is the first time users are being added, the first user added will be the GSP administrator.

If this is not the first time you are adding users (you are adding additional users), you will need to step through all current users to reach the next available user prompt.

4. The next prompt that appears will ask the following question:

   Do you want to modify the user number 1 parameters? (Y/[N]/Q to quit) __

   Follow the series of prompts to enter all the required fields for adding a user.

CAUTION Be sure to read each prompt carefully and enter the correct response. A missed or incorrect entry could deny entry to that user.

The following is an example of an added user's information:

```
. User's Name: Joe Smith
. Organization's Name: IT Support
. Dial-back configuration: Disabled
. Access Level: Operator
. Mode: multiple
. User's state: enabled
```

For the number 1 user, the Access level is administrator. The Mode entry of single only allows entry for that user one time, then access will be denied. A Mode entry of multiple allows unlimited entries into the GSP.

Removing Users

You can remove (disable) a GSP user with the same Security options and access control (SO) command used to add a user.

To remove a user, perform the following steps:

1. Access the GSP with the ctrl+b entry.
2. At the GSP prompt, enter the Security options and access control (SO) command:

   ```
   GSP> so
   ```

3. The first prompt you will see with the SO command is for GSP-wide parameters:

   ```
   GSP wide parameters are:
   . Login Timeout: 1 minutes.
   . Number of password Faults allowed: 3
   . Flow Control Timeout: 5 minutes.

   Do you want to modify the GSP wide parameters? (Y / [N]) __
   ```

   At this point you can modify the GSP wide parameters, or continue with removing a user. To remove users, respond N for no.

NOTE You will have to step through each user number until you reach the user to be removed.

4. When you access the number of the user to be removed, you must change the data in the prompts for that number.

   It is important that, at a minimum, you need to modify the User's state to Disabled.

Return the GSP to Default Configurations

The Default Configuration (dc) command is used to reset all or some of the GSP values to the default values. To return GSP values to default configurations, perform the following steps:

1. Access the GSP with the ctrl+b entry.
2. At the GSP prompt, enter the Default Configuration (dc) command:

   ```
   GSP> dc
   ```

3. Follow the prompts for the dc command, and be sure to have the change information available.

CAUTION When the Security configuration is reset, all users are removed, including the GSP administrator. It also disables the remote. Remote must be re-enabled through the main console using the Enable Remote (er) command.

rp54xx Server Boot Process

The length of time an rp54xx server will require to complete the boot process depends on the number of processors and the amount of RAM installed. Average configurations can take more than 20 minutes.

The boot process consists of the following main steps:

Typical Boot Process

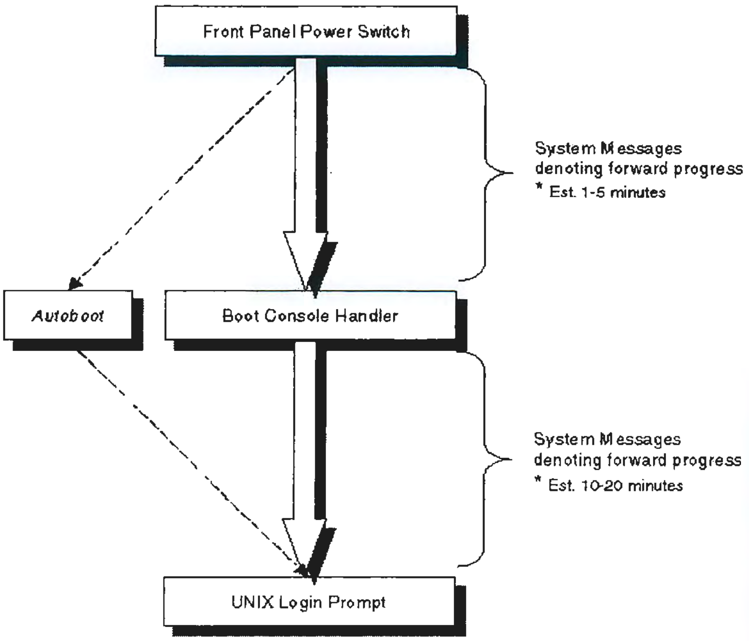

The duration of the full boot process can vary depending on:

- Number of CPUs installed
- Amount of memory installed
- UNIX version installed
- State of the network

All times approximate

18.697

During the Boot process a variety of errors or problems can occur as shown below:

Potential Boot Problems

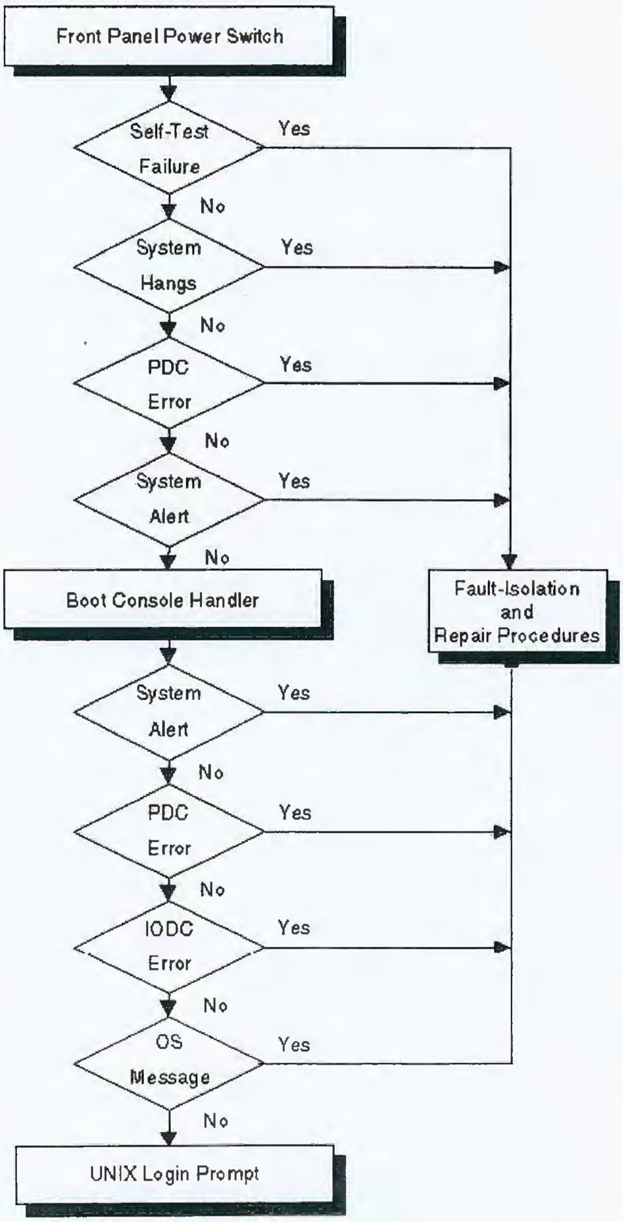

Initial Power-up

The following section describes the process of applying power to the rp54xx server and booting the system to the UNIX Login prompt. The amount of time it takes to go through self-test then boot the system will vary widely depending on hardware configuration. The following provides a “typical” procedure. Yours may vary depending on software and hardware installed:

Step 1. Apply AC Power to the system console.

Step 2. Apply power to the rp54xx server by turning the front panel switch to ON.

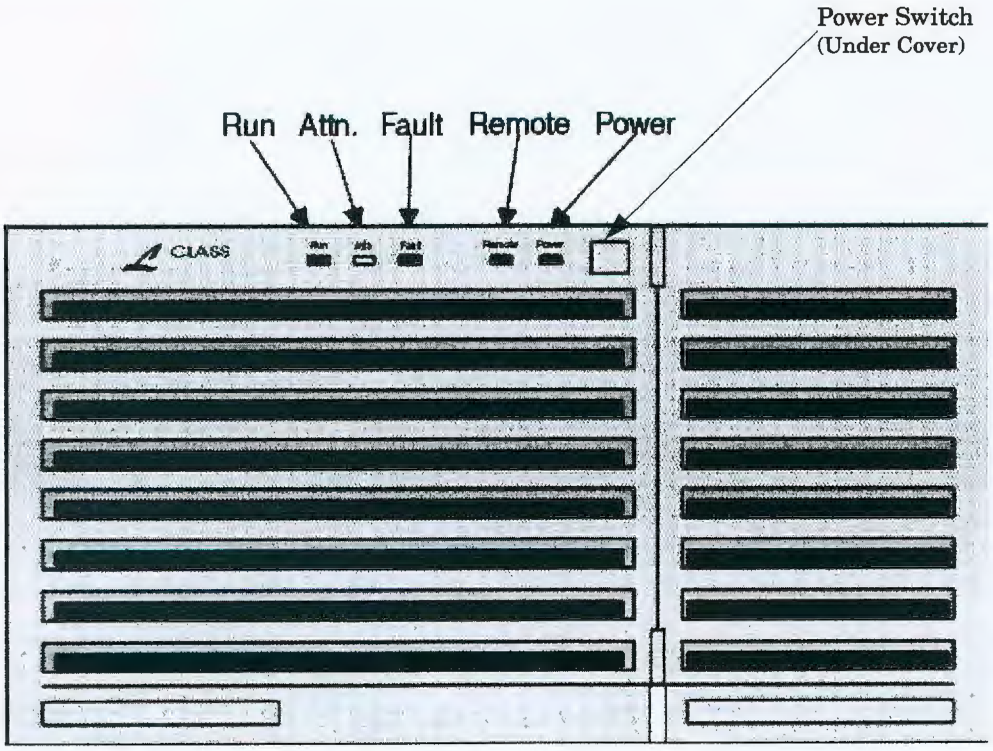

Step 3. Several self-test boot progress screens will be displayed and will scroll rapidly up the screen. Some tests may pause for up to one minute while the test completes.
The following examples of the forward progress screens are typical of the screens displayed.

Brief

Verbose

Step 4. When the initial power-up boot process completes in approximately one to five minutes, the BCH main menu will be displayed:

[redacted]

Step 5. To start the boot process using the primary boot path, enter *BO PRI*, at Boot Console Handler BCH main menu prompt and press **<ENTER>**.

NOTE Booting a system to a UNIX login prompt from BCH main menu can take 20 minutes or longer depending on your software and hardware configuration.

Step 6. Once the system reaches the UNIX login prompt the following will be displayed on the console screen:.

[redacted]

18.693

5 Utilities

18.692

Configuring the Rev A Guardian Service Processor (GSP)

The Rev A Guardian Service Processor (GSP) is a resident processor within the system that allows the local or remote system administrator to monitor and perform administrator functions. This section provides configuration procedures that will instruct you to:

- Configure the LAN port
- Add or delete users (maximum of 20)
- Change the default GSP configuration

Go to the appropriate section for the task that you wish to accomplish.

Configuring the GSP LAN Port

Perform the LAN configuration from the systems local port (either console or the HP secure web console).

NOTE The GSP has a separate LAN port from the system LAN port. It will need a separate LAN drop, IP address, and networking information from the port used by HP-UX.
Before starting this procedure, you will need to know the following information:

- I.P. address (for GSP)
- Subnet mask
- Gateway address
- Hostname (this is used when messages are logged or printed).

To configure the GSP LAN port, perform the following steps:

1. Go into the GSP with the ctrl+b entry.
2. At the GSP prompt, enter the LAN Configuration (lc) command:

```
GSP> lc
```

The lc command will start a series of prompts. Respond to each prompt with the appropriate information.

Adding Users

The GSP can only have a maximum of 20 users (one administrator and 19 operators). By design, the first user added to the GSP becomes the GSP administrator. Only the GSP administrator can add or remove users or change the GSP configuration.

NOTE Before starting this procedure, you will need to know the following information:

- User's name
- Organization's name
- Login name
- User's password

To add a user, perform the following steps:

1. Enter the GSP with the ctrl+b entry.
2. At the GSP prompt, enter the Security options and access control (SO) command:

   GSP> so

3. The first prompt you see with the so command is for GSP-wide parameters:

   ```
   GSP wide parameters are:
   Login Timeout: 1 minutes.
   . Number of password Faults allowed: 3
   . Flow Control Timeout: 5 minutes.
   Do you want to modify the GSP wide parameters? (Y / [N]) __
   ```

   At this point you can modify the GSP-wide parameters or continue adding users. To add users, enter N for no.

NOTE If this is the first time users are being added, the first user added will be the GSP administrator.

If this is not the first time you are adding users (you are adding additional users), you will *need to* step through all current users to reach the next available user prompt.

4. The next prompt that appears will ask the following question:

   Do you want to modify the user number 1 parameters? (Y/[N]/Q to quit) __

   Follow the series of prompts to enter all the required fields for adding a user.

CAUTION Be sure to read each prompt carefully and enter the correct response. A missed or incorrect entry will deny entry to that user.

The following is an example of an added users information:

```
. User's Name: Joe Smith
. Organization's Name: IT Support
. Dial-back configuration: Disabled
. Access Level: Operator
. Mode: multiple
. User's state: enabled
```

For the number 1 user, the Access level is administrator. The Mode entry of single only allows entry for that user one time, then access will be denied. A Mode entry of multiple allows unlimited entries into the GSP.

18.690

Removing Users

You can remove (disable) a GSP user with the same Security options and access control (SO) command used to add a user.

To remove a user, perform the following steps:

1. Enter the GSP with the ctrl+b entry.
2. At the GSP prompt, enter the Security options and access control (SO) command:

   ```
   GSP> so
   ```

3. The first prompt you see with the so command is for GSP-wide parameters:

   ```
   GSP wide parameters are:
   . Login Timeout: 1 minutes.
   . Number of password Faults allowed: 3
   . Flow Control Timeout: 5 minutes.

   Do you want to modify the GSP wide parameters? (Y / [N]) __
   ```

   At this point you can modify the GSP-wide parameters, or continue with removing a user. To remove users, enter N for no.

NOTE You must step through each user number until you reach the user to be removed.

4. When you access the number of the user to be removed, you must change the data in the prompts for that number.

   It is necessary that, at a minimum, you modify the User's state to Disabled.

Changing the Default GSP Configuration

This section describes the process of changing GSP default configurations. To change the GSP default configuration, perform the following steps:

1. Enter the GSP with the ctrl+b entry.
2. At the GSP prompt, enter the Default Configuration (dc) command:

   ```
   GSP> dc
   ```

3. Follow the prompts for the dc command, and have the change information available.

CAUTION When the Security configuration is reset, all users are removed, including the GSP administrator. The remote is disabled. The remote must be re-enabled through the main console using the Enable Remote (er) command.

Configuring the Rev B Guardian Service Processor (GSP)

The Rev B Guardian Service Processor (GSP) is a resident processor within the system that allows the local or remote system administrator to monitor and perform administrator functions. This section provides configuration procedures that will instruct you to:

- Configure the LAN port
- Add or delete users (maximum of 20)
- Change the default GSP configuration

Go to the appropriate section for the task that you wish to accomplish.

Configuring the GSP LAN Port

Perform the LAN configuration from the systems local port (either console or the HP secure web console).

NOTE The GSP has a separate LAN port from the system LAN port. It will need a separate LAN drop, IP address, and networking information from the port used by HP-UX.
Before starting this procedure, you will need to know the following information:

- I.P. address (for GSP)
- Subnet mask
- Gateway address
- Hostname (this is used when messages are logged or printed)

To configure the GSP LAN port, perform the following steps:

1. Go into the GSP with the ctrl+b entry.
2. At the GSP prompt, enter the LAN Configuration (lc) command:

```
GSP> lc
```

The lc command will start a series of prompts. Respond to each prompt with the appropriate information.

Adding Users

The GSP can only have a maximum of 20 users (one administrator and 19 operators). By design, the first user added to the GSP becomes the GSP administrator. Only the GSP administrator can add or remove users or change the GSP configuration.

NOTE Before starting this procedure, you will need to know the following information:

- User's name
- Organization's name
- Login name
- User's password

To add a user, perform the following steps:

1. Go into the GSP with the ctrl+b entry.
2. At the GSP prompt, enter the Security options and access control (SO) command:

```
GSP> so
```

3. The first prompt you will see with the so command is for GSP wide parameters:

```
GSP wide parameters are:
Login Timeout: 1 minutes.
. Number of password Faults allowed: 3
. Flow Control Timeout: 5 minutes.
Do you want to modify the GSP wide parameters? (Y / [N]) __
```

At this point you can modify the GSP wide parameters, or continue with adding users. To add users, respond N for no.

NOTE If this is the first time users are being added, the first user added will be the GSP administrator.

If this is not the first time you are adding users (you are adding additional users), you will need to step through all current users to reach the next available user prompt.

4. The next prompt that appears will ask the following question:

```
Do you want to modify the user number 1 parameters? (Y/[N]/Q to quit) __
```

Follow the series of prompts to enter all the required fields for adding a user.

CAUTION Be sure to read each prompt carefully and enter the correct response. A missed or incorrect entry could deny entry to that user.

An example of an added users information would be:

```
. User's Name: Joe Smith
. Organization's Name: IT Support
. Dial-back configuration: Disabled
. Access Level: Operator
. Mode: multiple
. User's state: enabled
```

For the number 1 user, the Access level is administrator. The Mode entry of single only allows entry for that user one time, then access will be denied. A Mode entry of multiple allows unlimited entries into the GSP.

Removing Users

You can removedisable) a GSP user with the same Security options and access control (SO) command used to add a user.

To remove a user, perform the following steps:

1. Go into the GSP with the ctrl+b entry.
2. At the GSP prompt, enter the Security options and access control (SO) command:

   ```
   GSP> so
   ```

3. The first prompt you will see with the so command is for GSP wide parameters:

   ```
   GSP wide parameters are:
   . Login Timeout: 1 minutes.
   . Number of password Faults allowed: 3
   . Flow Control Timeout: 5 minutes.

   Do you want to modify the GSP wide parameters? (Y / [N]) __
   ```

   At this point you can modify the GSP wide parameters, or continue with removing a user. To remove users, respond N for no.

NOTE You will have to step through each user number until you reach the user to be removed.

4. When you access the number of the user to be removed, you must change the data in the prompts for that number.

   It is important that, at a minimum, you need to modify the User's state to Disabled.

Changing the Default GSP Configuration

This section describes the process of changing GSP default configurations. To change the GSP default configuration, perform the following steps:

1. Go into the GSP with the ctrl+b entry.
2. At the GSP prompt, enter the Default Configuration (dc) command:

   ```
   GSP> dc
   ```

3. Follow the prompts for the dc command, and be sure to have the change information available.

CAUTION When the Security configuration is reset, all users are removed, including the GSP administrator. It also disables the remote. Remote must be re-enabled through the main console using the Enable Remote (er) command.

18.685

6 Troubleshooting

38.684

Determine Current System State

To determine the current system state of an rp54xx server, first note the state of all LED indicators on the front panel. Processing this information using the decoders provided can greatly reduce the amount of time required to repair a suspected system fault.

The following procedure lists the tools available to aid you in determining the current system state.

Step 1. Determine if you can get a system prompt and if so, what kind of prompt.

Software System	Screen Prompt
Boot Console Handler (BCH)	`Main Menu: Enter command or menu>`
Guardian Service Processor (GSP)	`GSP>`
Initial System Loader (ISL)	`ISL>`
HP UNIX (HP-UX)	*Prompt varies depending on UNIX state*

Step 2. Decode the Run/Attention/Fault LED States.

Step 3. Decode the PCI I/O LED States.

Step 4. Decode the Fan, Power Supply, and Disk LED States.

Step 5. Decode the GSP LED States.

Step 6. Decode the LAN/SCSI LED States.

Troubleshooting and FRU identification

Once you have determined the current system state, you must troubleshoot the system to determine what the problem symptoms are and what repair actions to take.

Problem Symptoms and Repair Actions

Use this guide to assist you in repairing the system by matching the problem symptom with the appropriate troubleshooting step.

Table 6-1 Problem Symptoms and Repair Actions

Problem or Symptom	Problem Indicators	Normal Functioning Indicators	Troubleshooting Steps	Potential FRUs
No indication of Housekeeping voltage present when AC connected and power switch in Standby position.	Front Panel Power LED OFF when AC is plugged into system.	• Power switch Off. • Front Panel POWER LED should be FLASHING to indicate presence of Housekeeping voltage. • Power LED on GSP board should be lit solid green.	1. AC must be present. Check that PDU is plugged in. 2. Ensure there are 2 working power supplies (1 supply for rp5400). The LED on each supply should be lit. 3. Check for Service processor prompt (CTRL B at console).	• No AC present • Power Supplies • Power Converter • System Board • Display Board

38.682

Table 6-1 Problem Symptoms and Repair Actions (Continued)

Problem or Symptom	Problem Indicators	Normal Functioning Indicators	Troubleshooting Steps	Potential FRUs
System won't power on when Front Panel Power switch is turned on.	Front Panel Power LED stays BLINKING when Power Switch is turned on. ATTENTION LED may be FLASHING.	• Power switch On. • Power LED on SOLID.	1. Check for remote power shutdown via GSP>PC command. 2. *Check Error Chassis Logs*. Look for Error Chassis Log with a Source Detail = Low Voltage DC Power. This indicates a failure of one of the CPU Support Modules. The failing CPU support module is indicated in the Source ID field. 3. Ensure there are 2 working power supplies (1 supply for rp5400). The LED on each supply should be lit. 4. Check to see if GSP can communicate with platform monitor. Execute the following GSP command: GSP>PC You should get power monitor status information. 5. Housekeeping 1 voltage present indication. Check that platform monitor power LED is lit. 6. Platform Monitor functioning. Check platform monitor heartbeat LED is lit.	• Power Supplies • CPU Support Module • Platform Monitor • System Board

18.681

Table 6-1 Problem Symptoms and Repair Actions (Continued)

Problem or Symptom	Problem Indicators	Normal Functioning Indicators	Troubleshooting Steps	Potential FRUs
No BCH Main Menu prompt.	Front Panel RUN LED is not FLASHING. There is no forward progress chassis codes at the console. There is no BCH Main Menu prompt at the console.	• Flashing RUN LED. • Forward progress chassis codes. • BCH Main Menu prompt.	1. Check for red LED on GSP. If lit red, the problem is with the GSP. 2. Check that the console is properly connected and can communicate with the Service Processor (CTRL B should get you the SP login prompt). 3. Check Service Processor Error logs. Look for entries related to: • Processors • Processor Support Modules (known as low voltage DC supplies in chassis codes. Also known as power pods). • Memory 4. Reduce to minimum configuration and troubleshoot from there.	• Core I/O • Processors • Processor Support Modules • Memory • System Board • Console

Table 6-1 Problem Symptoms and Repair Actions (Continued)

Problem or Symptom	Problem Indicators	Normal Functioning Indicators	Troubleshooting Steps	Potential FRUs
Can't boot to ISL.	Console messages indicating problems booting from the primary or alternate boot path.	• Console messages and prompt indicating you are at ISL.	1. Use BCH commands to verify I/O and presence of valid LIF devices. 2. Use BCH "Warn"command to determine if Boot is disabled. 3. Check for IODC tombstones. 4. Check SP chassis error logs.	• Disk Drive • Disk Media Backplane • LAN/SCSI Board • I/O Backplane
Can't boot HP-UX.	HP-UX boot error messages. RUN LED BLINKING.	• HP-UX boot messages. • HP-UX prompt. • RUN LED on SOLID.	1. Check SP chassis error logs. 2. Run ODE diagnostics.	• Processor • Disk Drive • Disk Media backplane • LAN/SCSI • Corrupt HP-UX

18.679

Chassis Code to FRU Decode

This is a guide to identify failing FRUs from System Alerts and Error Chassis Logs. The guide includes the following information:

- Cross-Referencing Chassis Log Errors to rp54xx FRUs
- Interpreting System Alerts
- Interpreting Service Processor Error Chassis Logs

There is a detailed interpretation of Chassis Logs and System Alerts in the *Interpreting Chassis Logs in Detail* guide.

Cross-Referencing Chassis Log Errors to rp54xx FRUs

Use the following table to identify the failing FRU from the Chassis Log information. You can also use the online Error Chassis Log-to-FRU Decoder utility.

1. Read the Chassis Log entry.
2. Match the **SOURCE**, **SOURCE DETAIL**, **SOURCE ID**, and **PROBLEM DETAIL** values (see table below) in the Chassis Log entry with the appropriate values in the table.
3. Read the table from left to right.

Use these examples to understand how to identify failing FRUs with the table:

- Power Supply Failure Example
- Processor Failure Example

Table 6-2 Chassis Log Error to FRU Decoder

Chassis Log Field Values and Descriptions from Log Entry					
Source	**Source Detail**	**Source ID**	**Problem Detail**	**FRU**	**Action to Take**
1 - Processor	Not Applicable (N/A)	N/A	N/A	Processor	From BCH Main Menu go to the Info Menu and execute the PR command to determine which processor is not functioning.
2 - Processor Cache	Not Applicable (N/A)	N/A	N/A	Processor	From BCH Main Menu go to the Info Menu and execute the PR command to determine which processor is not functioning.

18.678
3

Table 6-2 Chassis Log Error to FRU Decoder (Continued)

Chassis Log Field Values and Descriptions from Log Entry					
Source	**Source Detail**	**Source ID**	**Problem Detail**	**FRU**	**Action to Take**
3 - PDH	Not Applicable (N/A)	N/A	N/A	System Board	Replace the System Board.
4 - Power	1 - AC Mains	N/A	9 - Power Off	AC Power	Check that AC is being supplied to all power supplies.
4 - Power	3 - Low Voltage DC Power	CPU Support Module #	N/A	CPU Support Module	Replace the Processor Support Module (on the System Board) referenced in the Source ID.
4 - Power	4 - High Voltage DC Power	Power Supply #	N/A	Power Supply	Replace the Power Supply (in the front of the system behind the bezel) referenced in the Source ID.
6 - Platform	3 - Cabinet Fan	Fan #	N/A	Fan	Replace the Fan referenced in the Source ID.
6 - Platform	6 - Service Processor	N/A	N/A	Core I/O	The Service Processor is on the GSP I/O board. Replace the GSP
6 - Platform	7 - Power Monitor	N/A	N/A	Power Monitor	Replace Platform Monitor card.
7 - Memory	1 - Controller	N/A	N/A	System Board	Replace the System Board.
7 - Memory	4 - SIMM or DIMM	N/A	N/A	Memory DIMM	Isolate to failing DIMM using BCH (IN, ME) and ODE memory diagnostic.
8 - I/O	6 - Disk	N/A	Various Values	Disk Subsyste m	Use BCH commands and ODE diagnostics to check disk subsystem.

18.677

Power Supply Failure Example

```
GSP> sl

SL

Which buffer are you interested in :
Incoming, Activity, Error, Current boot or Last boot ? (I/A/E/C/L) e
e

Do you want to set up filter options on this buffer ? (Y/[N]) n
n

Type + CR and CR to go up (back in time),
Type - CR and CR to go down (forward in time),
Type Q to escape.

Log Entry #    0 :
ALERT LEVEL: 6=Boot possible, pending failure or environmental problem - action required
REPORTING ENTITY TYPE: 2=power monitor - REPORTING ENTITY ID: 00
CALLER ACTIVITY: 4=monitor - CALLER SUBACTIVITY: 04=low voltage power supply
SOURCE: 4=power - SOURCE DETAIL: 4=high voltage DC power - SOURCE ID: 02
PROBLEM DETAIL: A=unexpected - ACTIVITY STATUS: F
Data  0 : Low=00000000 : High=00000000 - type  0 = Data Field Unused
Data  1 : Low=0F152A28 : High=00006303 - type 11 - Timestamp 04/15/1999 21:42:40
```

Problem Analysis

Step 1. Find the Source value. In this example, it is *SOURCE: 4=power.*

Use the **Power** row of the Error Chassis Log-to-FRU Decoder table.

Step 2. Find the Source Detail value. In this example, it is *SOURCE DETAIL: 4=high voltage DC power.*

Use the **High Voltage DC Power** row of the table.

Step 3. Find the Source ID value. In this example, it is *SOURCE ID: 02.*

The failing power supply is Power Supply #2.

Step 4. The Problem Detail for this row is not applicable.

Step 5. The FRU column of the table identifies the FRU as the Power Supply.

The correct action would be to replace Power Supply #2, located in the front of the system.

Processor Failure Example

```
Log Entry #    1 :
ALERT LEVEL: 6=Boot possible, pending failure or environmental problem - action required
REPORTING ENTITY TYPE: 0=system firmware - REPORTING ENTITY ID: 01
CALLER ACTIVITY: 1=test - CALLER SUBACTIVITY: 62=implementation dependent
SOURCE: 1=processor - SOURCE DETAIL: 1=processor general - SOURCE ID: 00
PROBLEM DETAIL: 3=functional failure - ACTIVITY STATUS: 0
Data  0 : Low=00000003 : High=F7000000 - type  0 = Data Field Unused
Data  1 : Low=0F160920 : High=00006303 - type 11 - Timestamp 04/15/1999 22:09:32
```

Problem Analysis

Step 1. Find the Source value. In this example, it is *SOURCE: 1=processor.*

Use the **Processor** row of the Error Chassis Log-to-FRU Decoder table.

18676

Step 2. The Source Detail, the Source ID, and the Problem Detail values are all not applicable for the Processor row of the table.

Step 3. The FRU column of the table identifies the FRU as a failing processor.

Step 4. The Action column of the table instructs us to use the Info Menu and PR command of the BCH Main Menu to identify the failing processor.

Here is the output of Step 4 in our example:

```
Main Menu: Enter command or menu > in

---- Information Menu ---------------------------------------------

     Command                    Description
     -------                    -----------
     ALL                        Display all system information
     BootINfo                   Display boot-related information
     CAche                      Display cache information
     ChipRevisions              Display revisions of major VLSI
     COprocessor                Display coprocessor information
     FRU                        Display FRU information
     FwrVersion                 Display firmware version
     IO                         Display I/O interface information
     LanAddress                 Display Core LAN station address
     MEmory                     Display memory information
     PRocessor                  Display processor information
     WArnings                 Display selftest warning messages

----
Information Menu: Enter command > pr

PROCESSOR INFORMATION

                      HVERSION  SVERSION            Processor
  Processor   Speed     Model   Model/Op  CVERSION    State
  --------- --------  --------  --------  --------  -------------
      1     440  MHz  0x05c4     0x0491     2.  0   Active
      3     440  MHz  0x05c4     0x0491     2.  0   Stopped:Nonresponding

  Central Bus Speed (in MHz)    :         82
  Software ID (dec)             : 1635329341
  Software ID (hex)             : 0x6179253d
  Software Capability           : 0x01100000f0

Information Menu: Enter command >
```

Processor #3 is `Stopped:Nonresponding`. Replace Processor #3.

Interpreting System Alerts

System Alerts are reported to the system console when a problem is detected by the Service Processor. These alerts are stored in the Service Processor Error Logs. When this new alert is added to the log file, it will cause the front panel ATTENTION LED to blink.

18.675

Interpreting System Alerts

Do one of the following:

1. No response: the alert will time out and the system will continue operating.
2. **A** - Responding with the letter **A** will inform the Service Processor that you have seen the entry. The system will continue to operate.
3. **X** - Responding with the letter **X** will inform the Service Processor to disable all future alert messages. This can be re-enabled with a Service Processor command.

Sample System Alert

```
************************SYSTEM ALERT***************************************
ALERT LEVEL: 6=Boot possible, pending failure or environmental problem - action required
PROBLEM DETAIL: 4=fan failure - SOURCE ID: 04
SYSTEM NAME: fesrhapgsp

MODEL NAME:    MODEL STRING:    S/N:
SPU POWER: ON
ACTIVITY/COMPLETION LEVEL:  0%
SYSTEM BOOT IS PENDING

LEDs:  RUN       ATTENTION     FAULT     REMOTE
       FLASH     OFF           OFF       ON

CALLER ACTIVITY: 4=monitor - CALLER SUBACTIVITY: 05=fan
REPORTING ENTITY TYPE: 2=power monitor - REPORTING ENTITY ID: 00
SOURCE: 6=platform - SOURCE DETAIL: 3=cabinet fan

0x002008646304405F 00000000 00000000 - type  0 = Data Field Unused

0x582008646304405F 00006303 0F151D08 - type 11 = Timestamp 04/15/1999 21:29:08
A: ack read of this entry - X: Disable all future alert messages
Anything else skip redisplay the log entry
->Choice:a
```

Key FRU Identification Fields for System Alerts

The following fields are used for FRU identification.

Alert Level: How the problem has affected the system operation.

Source: What major part of the system the alert is referring to (i.e, platform, memory, processor, etc...).

Source Detail: What sub-part of the system the alert is referring to (i.e, cabinet fan, DIMM, high voltage DC power, etc...).

Source ID: Specific FRU referred to in Source and Source Detail (i.e, cabinet fan #4).

Problem Detail: Specific problem information (i.e, power off, functional failure, etc...).

Timestamp: When the problem occurred.

The above sample system alert shows the following:

1. The problem does not affect system boot.
2. The problem is with platform cabinet fan #4.
3. The problem is a fan failure. Replace fan #4 to correct the problem.
4. The fan failed on April 15, 1999 at 9:29 PM.

Interpreting Service Processor Error Chassis Logs

Accessing the Service Processor Error Chassis Logs will turn the **ATTENTION** LED, blinking on the front panel, OFF.

Chassis Logs (located in the Service Processor) contain low level logging information related to the following 5 categories:

- Incoming log: Contains all chassis logs coming into the Service Processor.
- Activity log: Contains all chassis logs related to system activity.
- Error log: Contains all error chassis logs.
- Current boot log: Contains all chassis logs associated with the current boot.
- Last boot log: Contains all chassis logs associated with the last boot.

The Error Chassis Logs are the ones you need to look at.

Accessing Error Chassis Logs

Execute the following steps to access the Error Chassis Logs.

1. At the system console prompt, type **CTRL B**
2. Enter the Service Processor Login and Password
3. The screen will display: GSP>

   At the GSP> prompt: type SL and press enter
4. The screen will display:

   `Which buffer are you interested in:`
   `Incoming, Error, Current boot, Last boot? (I/A/E/C/L)`, type **E**, and press enter
5. The screen will display:

   Do you want to set up filter options on this buffer? (Y/[N]), type **N**, and press enter
6. The most recent Error Log Entry (Log Entry #0) will be displayed. A carriage return after this will display the next log entry. Type Q to stop displaying the log entries. The screen will display: GSP>
7. At the GSP> prompt: type CO, and press enter to return to the console screen.

Example of Accessing Error Logs

Here is an example of accessing the Error Logs from the Boot Console Handler (BCH) Main Menu prompt. User input is shown in *ITALICS*.

```
Main Menu: Enter command or menu > type CTRL B
Service Processor login: System Operator
Service Processor password: ****** (password hidden)

            Welcome to HP Guardian Service Processor
            System Name: fesrhapgsp


fesrhapgsp:
ALERT LEVEL: 6=Boot possible, pending failure or environmental problem - action required
SOURCE: 6=platform - SOURCE DETAIL: 3=cabinet fan
PROBLEM DETAIL: 4=fan failure
GSP>
```

18.673

```
          HP Guardian Service Processor Command Interface
          Type HE to get the list of available commands

fesrhapgsp:
ALERT LEVEL: 6=Boot possible, pending failure or environmental problem - action required
SOURCE: 6=platform - SOURCE DETAIL: 3=cabinet fan
PROBLEM DETAIL: 4=fan failure
GSP> sl

SL

Which buffer are you interested in :
Incoming, Activity, Error, Current boot or Last boot ? (I/A/E/C/L) e
e

Do you want to set up filter options on this buffer ? (Y/[N]) n
n

Type + CR and CR to go up (back in time),
Type - CR and CR to go down (forward in time),
Type Q CR to escape.

Log Entry #   0 :
ALERT LEVEL: 6=Boot possible, pending failure or environmental problem - action required
REPORTING ENTITY TYPE: 2=power monitor - REPORTING ENTITY ID: 00
CALLER ACTIVITY: 4=monitor - CALLER SUBACTIVITY: 05=fan
SOURCE: 6=platform - SOURCE DETAIL: 3=cabinet fan - SOURCE ID: 04
PROBLEM DETAIL: 4=fan failure - ACTIVITY STATUS: F
Data  0 : Low=00000000 : High=00000000 - type  0 = Data Field Unused
Data  1 : Low=0F151D08 : High=00006303 - type 11 = Timestamp 04/15/1999 21:29:08


q

fesrhapgsp:
ALERT LEVEL: 6=Boot possible, pending failure or environmental problem - action required
SOURCE: 6=platform - SOURCE DETAIL: 3=cabinet fan
PROBLEM DETAIL: 4=fan failure
GSP> co

CO

You are now leaving the Guardian Service Processor Command Interface
and returning to the console mode.  Type Ctrl B to reactivate it.

Main Menu: Enter command or menu >
```

Key FRU Identification Fields for Error Chassis Logs

The following fields are for FRU identification.

Alert Level: How the problem has affected the system operation.

Source: What major part of the system the alert is referring to (i.e., platform, memory, processor, etc.).

Source Detail: What sub-part of the system the alert is referring to (i.e., cabinet fan, DIMM, high voltage DC power, etc.).

Source ID: Specific FRU referred to in Source and Source Detail (i.e., fan #4).

Problem Detail: Specific problem information (i.e, power off, functional failure, etc.).

Timestamp: When the problem occurred.

The above sample system alert message shows the following:

1. The problem does not affect booting of the system.

2. The problem is with fan #4.
3. The problem is a fan failure.
4. The fan failed on April 15, 1999 at 9:29 PM.

In this example, fan #4 should be replaced to correct the problem.

Interpreting Chassis Logs Using the chassis_code.codes File

For chassis logs generated by system firmware (Reporting Entity Type 0), use the chassis_code.codes file for chassis code definitions. Each revision of system firmware (AKA Processor Dependent Code or PDC) has a unique chassis_code.codes file. This file is not part of either the PF_Cxxxxx or PHSS_xxxxx server firmware patches. The chassis_code.codes files appear in the appendices of the *Interpreting Chassis Logs in Detail* guide.

The definition of a PDC reported chassis code is determined by locating either the last four digits of a chassis log or the last three digits of a selftest chassis code in the appropriate chassis_code.codes file. Refer to the *Interpreting Chassis Logs in Detail* guide for definition and examples of selftest chassis codes.

NOTE Be sure to use the appropriate appendix as the PDC for rp5400/rp5450 is different than PDC for rp5470. Using the wrong appendix may result in a mis-interpretation of the chassis code.

To quickly learn the definition of a PDC reported chassis code, follow these four steps:

Step 1. Determine either the last 4 digits of the hex chassis code or the last 3 digits of the selftest chassis code.

Step 2. Go to the appropriate appendix in the *Interpreting Chassis Logs in Detail* guide.

Step 3. Locate the chassis code that matches the last 3 or 4 digits. If viewing via web browser or Adobe Acrobat, use the FIND feature to locate the chassis code.

Step 4. Take action as appropriate.

Example 6-1 Chassis Log: Reporting Entity Type = System Firmware

```
Log Entry #   0 :SYSTEM NAME: fesrhapgspDATE: 12/08/2000 TIME: 23:46:22ALERT LEVEL: 6
= Boot possible, pending failure - action requiredSOURCE: 3 = PDHSOURCE DETAIL: 0 =
unknown, no source stated   SOURCE ID: 3PROBLEM DETAIL: 0 = no problem detailCALLER
ACTIVITY: 1 = test   STATUS: 0CALLER SUBACTIVITY: 71 = implementation
dependentREPORTING ENTITY TYPE: 0 = system firmware   REPORTING ENTITY ID:
030x0000306030031710 00000000 000000FE type  0 = Data Field Unused0x5800386030031710
0000640B 08172E16 type 11 = Timestamp 12/08/2000 23:46:22Type CR for next entry, Q CR
to quit.
```

Using the example above:

Step 1. The last 4 digits are 1710.

Step 2. This is an rp5450 so use Appendix B in the *Interpreting Chassis Logs in Detail* guide.

Step 3. Using the FIND feature to look up 1710 in Appendix B, we learn the definition is CC_BOOT_INVALID_SPHYR_SETTINGS.

Step 4. The appropriate action in this example would be to verify the switch settings on the system board are set correctly for the installed CPU's.

rp5400 and rp5450 System Block Diagram

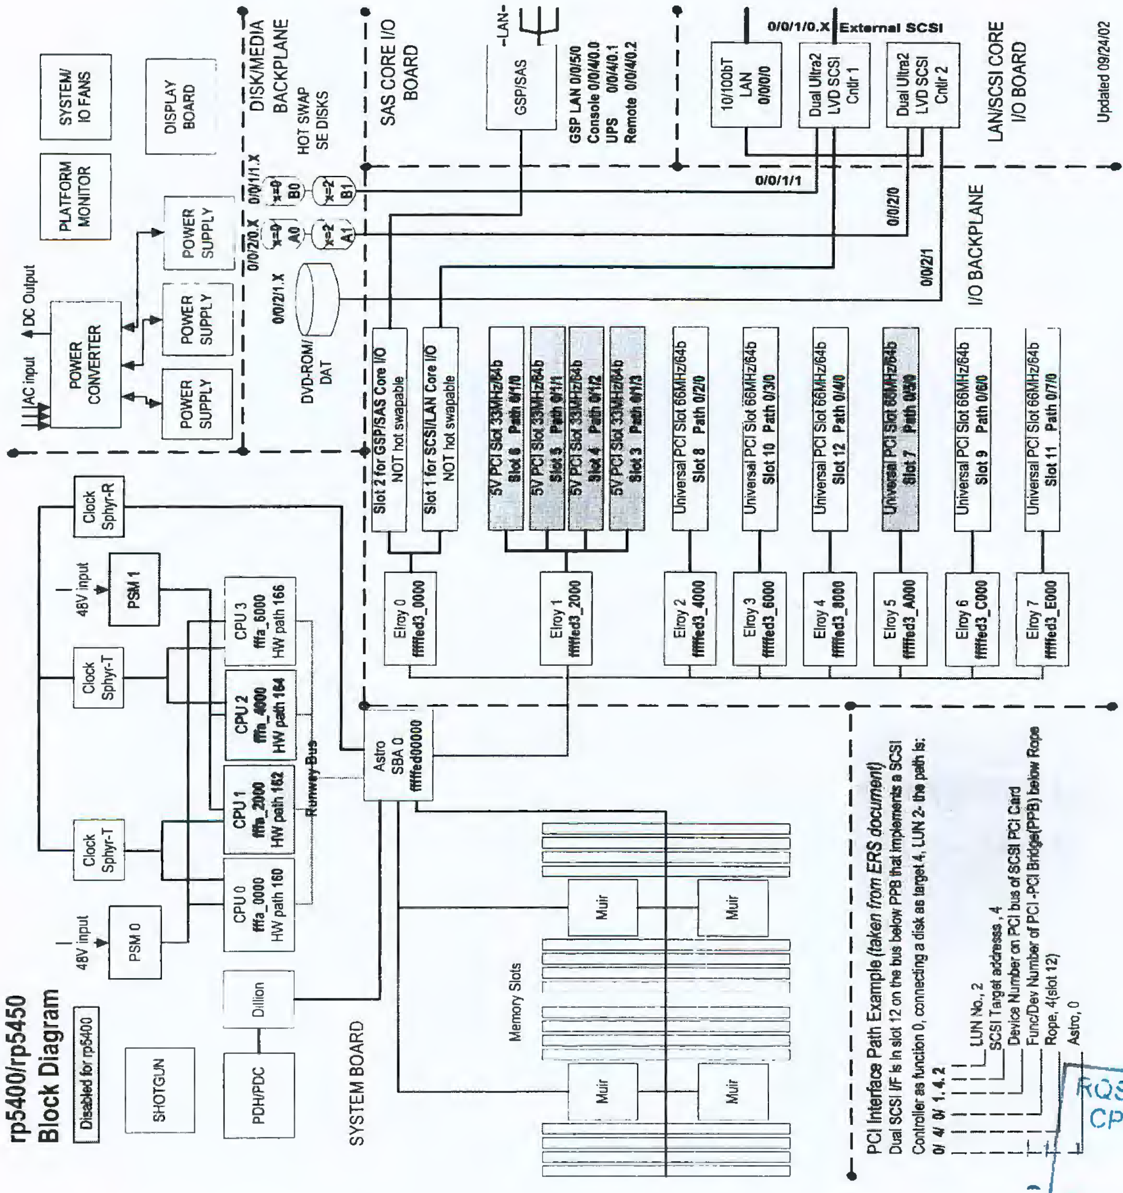

PCI Interface Path Example *(taken from ERS document)*
Dual SCSI I/F is in slot 12 on the bus below PPB that implements a SCSI Controller as function 0, connecting a disk as target 4, LUN 2- the path is:

```
0/ 4/ 0/ 1. 4. 2
|  |  |  |  |  |__LUN No., 2
|  |  |  |  |_____SCSI Target address , 4
|  |  |  |________Device Number on PCI bus of SCSI PCI Card
|  |  |___________Func/Dev Number of PCI -PCI Bridge(PPB) below Rope
|  |______________Rope, 4(slot 12)
|_________________Astro, 0
```

rp5400 and rp5450 I/O Block Diagram

rp5400/rp5450 I/O Block Diagram

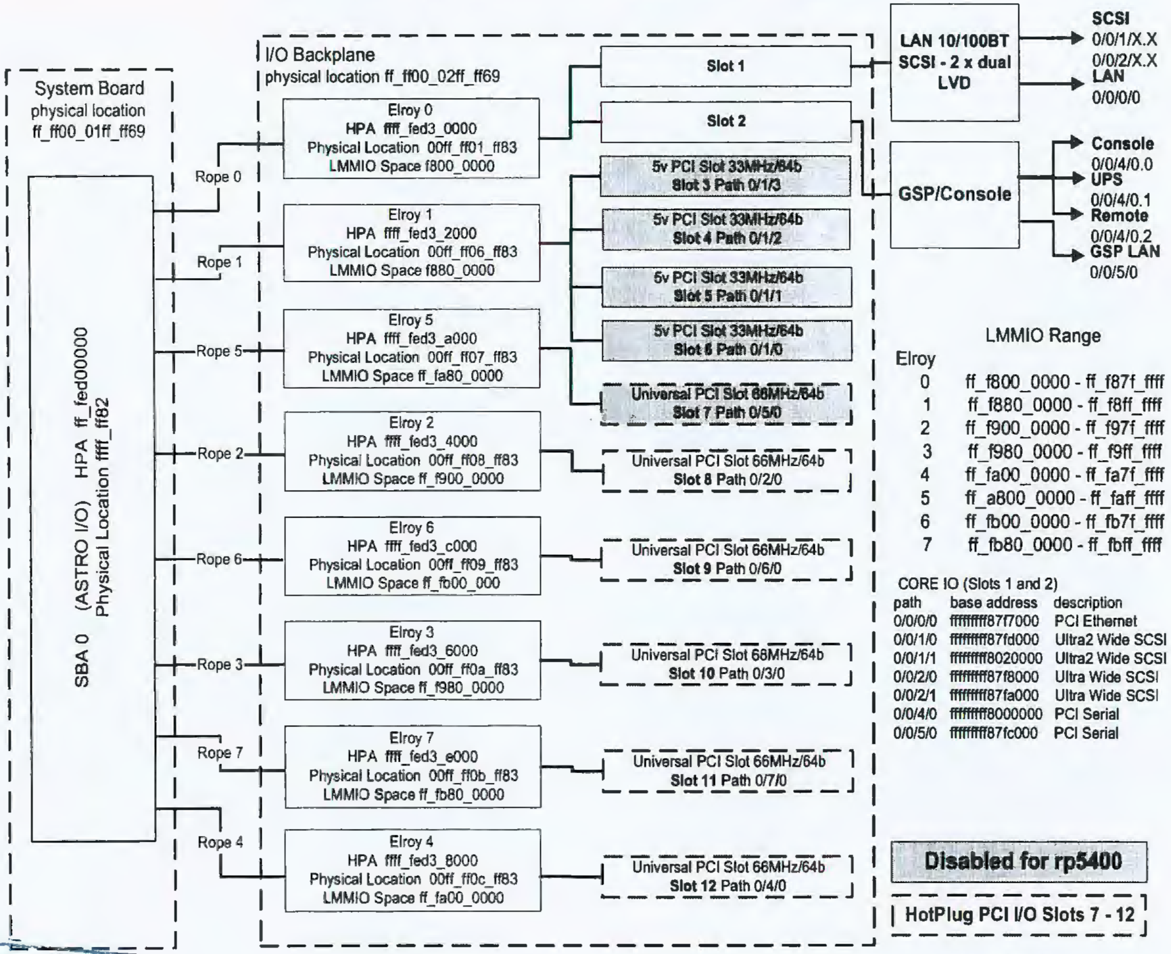

LMMIO Range

Elroy	
0	ff_f800_0000 - ff_f87f_ffff
1	ff_f880_0000 - ff_f8ff_ffff
2	ff_f900_0000 - ff_f97f_ffff
3	ff_f980_0000 - ff_f9ff_ffff
4	ff_fa00_0000 - ff_fa7f_ffff
5	ff_a800_0000 - ff_faff_ffff
6	ff_fb00_0000 - ff_fb7f_ffff
7	ff_fb80_0000 - ff_fbff_ffff

CORE IO (Slots 1 and 2)

path	base address	description
0/0/0/0	fffffffff87f7000	PCI Ethernet
0/0/1/0	fffffffff87fd000	Ultra2 Wide SCSI
0/0/1/1	fffffffff8020000	Ultra2 Wide SCSI
0/0/2/0	fffffffff87f8000	Ultra Wide SCSI
0/0/2/1	fffffffff87fa000	Ultra Wide SCSI
0/0/4/0	fffffffff8000000	PCI Serial
0/0/5/0	fffffffff87fc000	PCI Serial

Disabled for rp5400

HotPlug PCI I/O Slots 7 - 12

Updated 09/24/02

18.669

rp5430 and rp5470 System Block Diagram

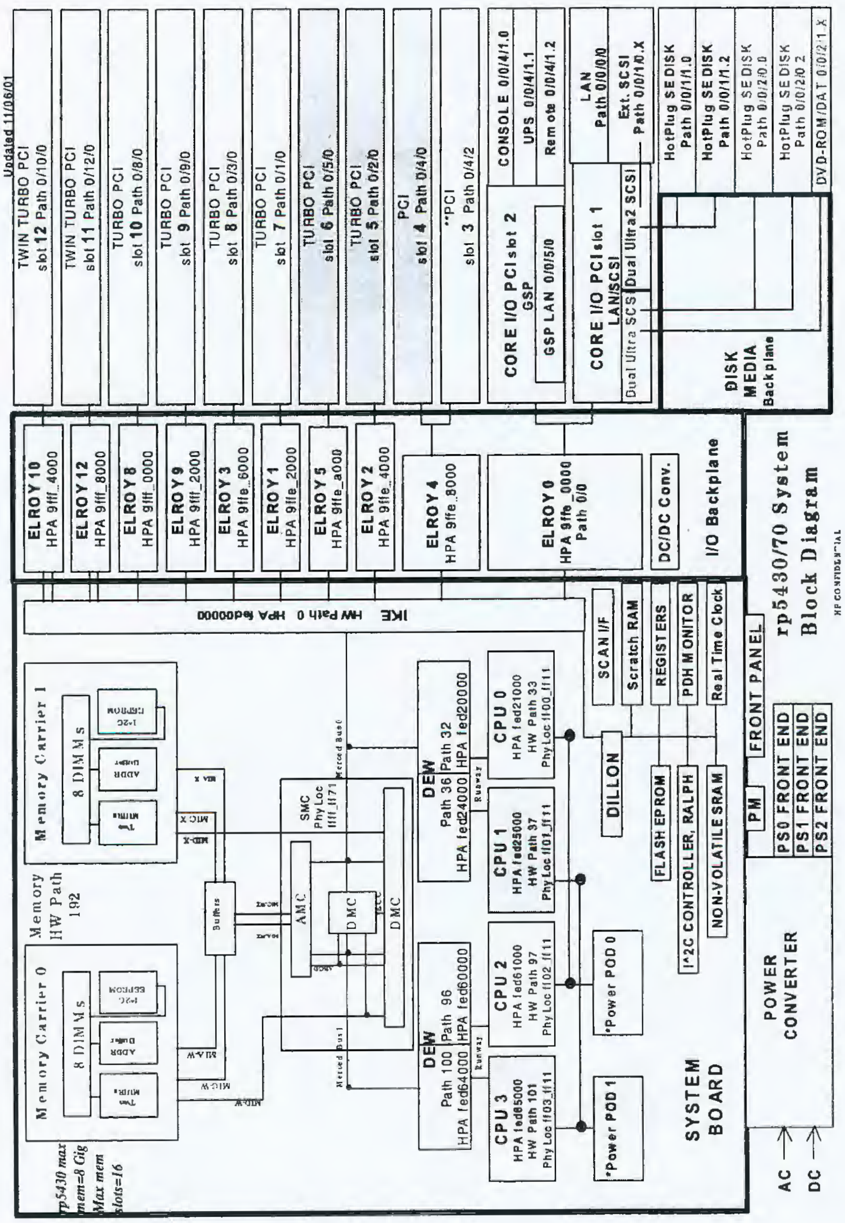

*For 550MHZ use 2.0Vdc Power POD
For 650,750MHZ use 1.6Vdc Power POD

Disabled for rp5430

***Enabled for Internal Raid.*
Other supported cards may be used in slot

rp5430 and rp5470 I/O Block Diagram

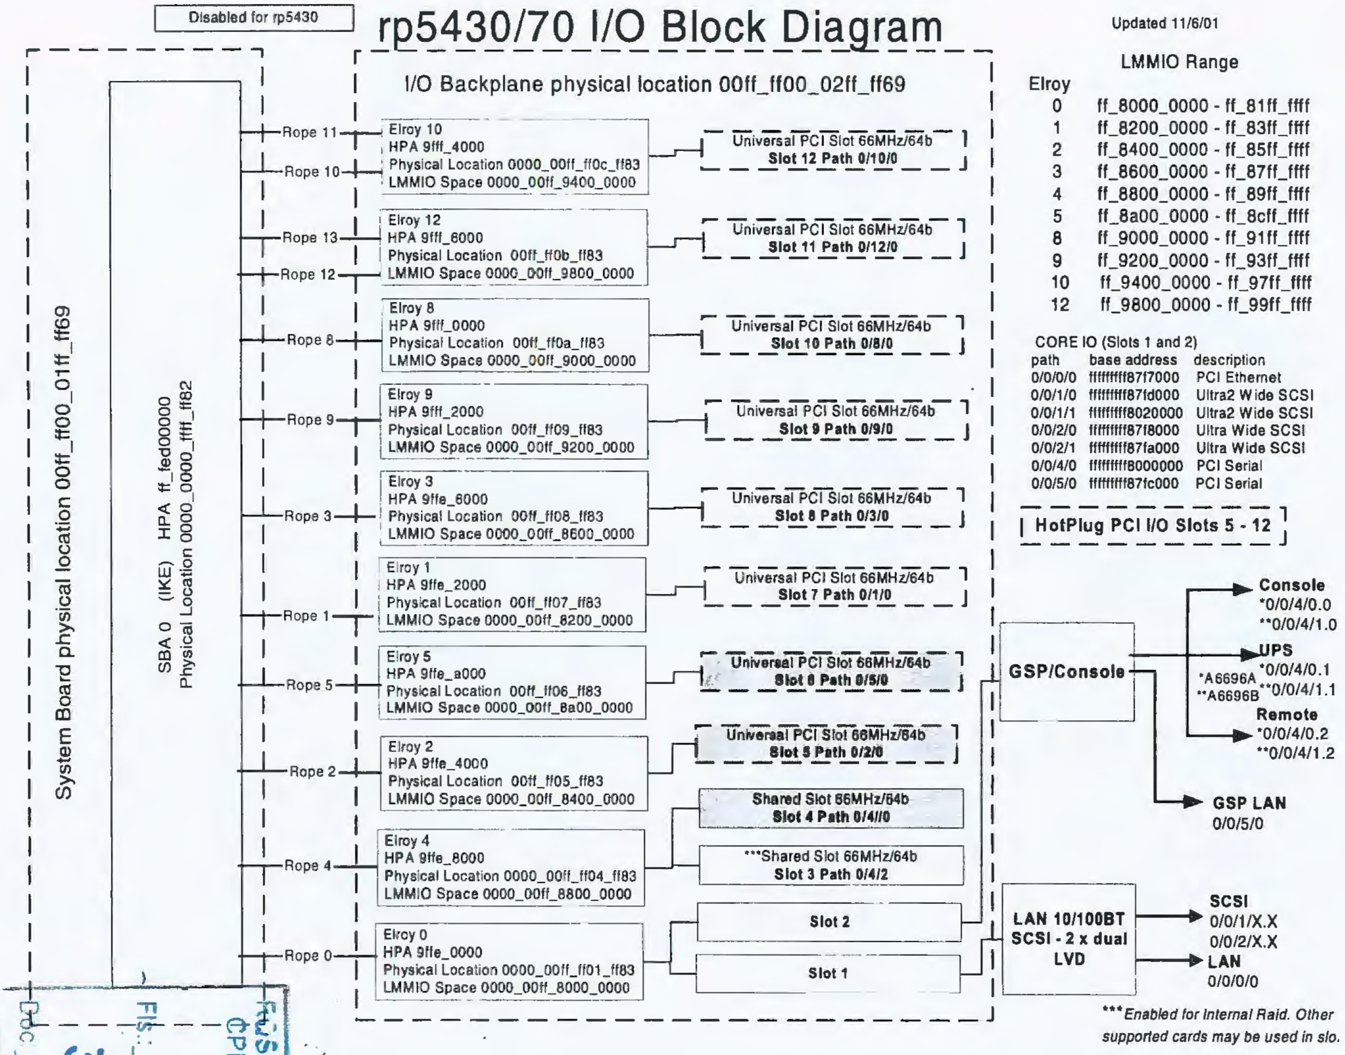
Disabled for rp5430
rp5430/70 I/O Block Diagram
Updated 11/6/01
System Board physical location 00ff_ff00_01ff_ff69
SBA 0 (IKE) HPA ff_fed00000
Physical Location 0000_0000_ffff_ff82
I/O Backplane physical location 00ff_ff00_02ff_ff69
Rope 11
Rope 10
Elroy 10
HPA 9fff_4000
Physical Location 0000_00ff_ff0c_ff83
LMMIO Space 0000_00ff_9400_0000
Universal PCI Slot 66MHz/64b
Slot 12 Path 0/10/0
Rope 13
Rope 12
Elroy 12
HPA 9fff_6000
Physical Location 00ff_ff0b_ff83
LMMIO Space 0000_00ff_9800_0000
Universal PCI Slot 66MHz/64b
Slot 11 Path 0/12/0
Rope 8
Elroy 8
HPA 9fff_0000
Physical Location 00ff_ff0a_ff83
LMMIO Space 0000_00ff_9000_0000
Universal PCI Slot 66MHz/64b
Slot 10 Path 0/8/0
Rope 9
Elroy 9
HPA 9fff_2000
Physical Location 00ff_ff09_ff83
LMMIO Space 0000_00ff_9200_0000
Universal PCI Slot 66MHz/64b
Slot 9 Path 0/9/0
Rope 3
Elroy 3
HPA 9ffe_6000
Physical Location 00ff_ff08_ff83
LMMIO Space 0000_00ff_8600_0000
Universal PCI Slot 66MHz/64b
Slot 8 Path 0/3/0
Rope 1
Elroy 1
HPA 9ffe_2000
Physical Location 00ff_ff07_ff83
LMMIO Space 0000_00ff_8200_0000
Universal PCI Slot 66MHz/64b
Slot 7 Path 0/1/0
Rope 5
Elroy 5
HPA 9ffe_a000
Physical Location 00ff_ff06_ff83
LMMIO Space 0000_00ff_8a00_0000
Universal PCI Slot 66MHz/64b
Slot 6 Path 0/5/0
Rope 2
Elroy 2
HPA 9ffe_4000
Physical Location 00ff_ff05_ff83
LMMIO Space 0000_00ff_8400_0000
Universal PCI Slot 66MHz/64b
Slot 5 Path 0/2/0
Rope 4
Elroy 4
HPA 9ffe_8000
Physical Location 0000_00ff_ff04_ff83
LMMIO Space 0000_00ff_8800_0000
Shared Slot 66MHz/64b
Slot 4 Path 0/4//0
***Shared Slot 66MHz/64b
Slot 3 Path 0/4/2
Rope 0
Elroy 0
HPA 9ffe_0000
Physical Location 0000_00ff_ff01_ff83
LMMIO Space 0000_00ff_8000_0000
Slot 2
Slot 1
HotPlug PCI I/O Slots 5 - 12
GSP/Console
Console
*0/0/4/0.0
**0/0/4/1.0
UPS
*0/0/4/0.1
**0/0/4/1.1
*A6696A
**A6696B
Remote
*0/0/4/0.2
**0/0/4/1.2
GSP LAN
0/0/5/0
LAN 10/100BT
SCSI - 2 x dual
LVD
SCSI
0/0/1/X.X
0/0/2/X.X
LAN
0/0/0/0

LMMIO Range

Elroy	
0	ff_8000_0000 - ff_81ff_ffff
1	ff_8200_0000 - ff_83ff_ffff
2	ff_8400_0000 - ff_85ff_ffff
3	ff_8600_0000 - ff_87ff_ffff
4	ff_8800_0000 - ff_89ff_ffff
5	ff_8a00_0000 - ff_8cff_ffff
8	ff_9000_0000 - ff_91ff_ffff
9	ff_9200_0000 - ff_93ff_ffff
10	ff_9400_0000 - ff_97ff_ffff
12	ff_9800_0000 - ff_99ff_ffff

CORE IO (Slots 1 and 2)

path	base address	description
0/0/0/0	ffffffff87f7000	PCI Ethernet
0/0/1/0	ffffffff87fd000	Ultra2 Wide SCSI
0/0/1/1	ffffffff8020000	Ultra2 Wide SCSI
0/0/2/0	ffffffff87f8000	Ultra Wide SCSI
0/0/2/1	ffffffff87fa000	Ultra Wide SCSI
0/0/4/0	ffffffff8000000	PCI Serial
0/0/5/0	ffffffff87fc000	PCI Serial

****Enabled for Internal Raid. Other supported cards may be used in slo.*

18.667

18.666

Run/Attention/Fault LED States

Run (Green)	Attn (Amber)	Fault (Red)	Description
● On	● Off	● Off	State: □ System running normally. You should expect an OS prompt, if not, system may be hung. Action: □ Attempt to get system prompt to determine if system is hung. □ Talk to customer to determine reason for call.
● On	● Off	● On	This is an invalid indication. Check the server's LEDs and try again.
● On	● Off	● Flashing	State: □ The system crashed and rebooted itself successfully Action: □ Check chassis error logs to determine probable cause of system crash (either HPMC or HP-UX System Panic) □ Check service processor console logs for potential error messages from the OS (ie. Panic messages)
● On	● Flashing	● Off	State: □ There was a system interruption that did not take the system down Action: □ Check chassis error logs to determine probable cause □ Check Sentinel logs for probable cause

18.665

Run (Green)	Attn (Amber)	Fault (Red)	Description
● On	● Flashing	● On	This is an invalid indication. Check the server's LEDs and try again.
● On	● Flashing	● Flashing	State: □ System running, and □ A) unexpected reboot occurred, and □ B) a non-critical error has been detected Action: □ A) Check chassis error logs to determine probable cause of system crash (either HPMC or HP-UX System Panic) □ A) Check GSP console logs for potential error messages from the OS (ie. Panic messages) □ B) Check chassis error logs to determine probable cause of non-critical error □ B) Check Sentinel logs for probable cause of non-critical error

Run (Green)	Attn (Amber)	Fault (Red)	Description
● Flashing	● Off	● Off	State: □ Executing non-OS code - no problems detected □ System may be hung or waiting for BCH response □ Potential causes could be PDC never executed (problem with fetching code from PDH), HPMC while PDC was configuring system. Action: □ Check console for pending responses □ If system appears to be hung, execute TC from GSP prompt and check Last Boot Log for details of previous boot attempt. Pay attention to time stamps to ensure Last Boot Log reflects last boot attempt. □ If necessary, bring system down to minimum configuration, processors, memory, I/O, and troubleshoot.

18.663

Run (Green)	Attn (Amber)	Fault (Red)	Description
● Flashing	● Off	● On	State: □ Boot failed □ Executing non-OS code. Action: □ 1. Check chassis error logs □ 2. Read console messages for indications of problems ie. warnings from PDC
● Flashing	● Off	● Flashing	State: □ Unexpected reboot occurred □ Executing non-OS code. Action: □ 1. Check chassis error logs □ 2. Read console messages for indications of problems ie. warnings from PDC
● Flashing	● Flashing	● Off	State: □ Executing non-OS code. □ Non-critical error detected (ie. fan failure, power supply failure) Action: □ Read chassis error logs starting at entry 0 to determine cause of flashing amber LED and fix problem.

18.662

Run (Green)	Attn (Amber)	Fault (Red)	Description
● Flashing	● Flashing	● On	State: □ Boot failed □ Executing non-OS code. □ Non-critical error detected (ie. fan failure, power supply failure) Action: □ 1. Check chassis error logs □ 2. Read console messages for indications of problems ie. warnings from PDC
● Flashing	● Flashing	● Flashing	State: □ Unexpected reboot/system recovering. □ Executing non-OS code. □ Non-critical error detected (ie. fan failure, power supply failure) Action: □ 1. Check chassis error logs □ 2. Read console messages for indications of problems ie. warnings from PDC

Run (Green)	Attn (Amber)	Fault (Red)	Description
● Off	● Off	● Off	State: □ A) If Power LED Off □ B) If Power LED Flashing □ C) If Power LED On Action: □ A) Check Power Supply switches and LEDs □ Power supply switches should be on and LEDs should be on to indicate presence of AC. If Power Supply LEDs are on and the Front Panel LED is off, replace Power Monitor. □ Check AC power at source □ B) Turn Front Panel Power Switch on - LED should be on solid □ If LED continues to flash, system has been remotely powered off by the service processor or the power monitor has failed. □ Execute service processor PC command: type CTRL B, log into the service processor, and type PC at the prompt. □ If this doesn't cause the LED to go on solid, problem is most likely with the power monitor. □ C) Cannot execute PDC. If more than one processor installed, the problem is most likely a failed system board. If only one processor installed, problem could be either the processor or system board.

18.660

Run (Green)	Attn (Amber)	Fault (Red)	Description
● Off	● Off	● On	State: ▯ Boot failed ▯ OS not up and running, PDC has detected a failure that is preventing boot from occurring. Action: ▯ 1. Check chassis error logs ▯ 2. Read console messages for indications of problems ie. warnings from PDC
● Off	● Off	● Flashing	This is an invalid indication. Check the server's LEDs and try again.
● Off	● Flashing	● Off	State: ▯ No code is executing (PDC, OS or Diagnostics) ▯ Non-critical error detected (ie. fan failure, power supply failure) (Double fault situation) Action: ▯ Cannot execute PDC. If more than one processor installed, the problem is most likely a failed system board. If only one processor installed, problem could be either the processor or system board. ▯ Check chassis error logs to determine the source of non-critical error.

18.659

Run (Green)	Attn (Amber)	Fault (Red)	Description
● Off	◉ Flashing	● On	State: □ Boot failed □ OS not up and running, PDC has detected a failure that is preventing boot from occurring. □ Non-critical error detected (ie. fan failure, power supply failure) Action: □ 1. Check chassis logs □ 2. Read console messages for indications of problems ie. warnings from PDC
● Off	◉ Flashing	● Flashing	This is an invalid indication. Check the server's LEDs and try again.
	◉ On		Any combination with amber on is an invalid indication since amber is never on solid.

18.658

PCI I/O LED States

Power (green)	Attention (amber)	State
● On	● Off	State: ▯ Normal operation
● On	● On	State: ▯ Slot selected <or> slot located ▯ Power on Action: ▯ Not ready for OLRAD
● Off	● On	State: ▯ Slot selected <or> slot located ▯ Power is off Action: ▯ Ready for OLRAD
● On	● Flashing	State: ▯ Fault detected ▯ Power on

Power (green)	Attention (amber)	State
● Off	● Flashing	State: □ Fault detected □ Power off
● Off	● Off	State: □ Slot available

18.656
J

Expansion I/O LED States

Power (green)	Attention (amber)	State
● On	● On	State: □ Slot selected <or> slot located □ Power on Action: □ Not ready for OLRAD
Power (green)	Attention (amber)	State
● On	● Off	State: □ Normal operation
Power (green)	Attention (amber)	State
● On	● Flashing	State: □ Fault detected □ Power on
Power (green)	Attention (amber)	State
● Off	● On	State: □ Slot selected <or> slot located □ Power is off Action: □ Ready for OLRAD

18.655

Power (green)	Attention (amber)	State
● Off	● Off	State: ▯ Slot available
Power (green)	Attention (amber)	State
● Off	● Flashing	State: ▯ Fault detected ▯ Power off

18.654

GSP LED States

GSP - Revision A

Type	Status	State
GSP Upper	● On Green	State: □ GSP Power OK
	● Flashing Green	State: □ GSP LAN receive
	● On Red	State: □ GSP Failure
GSP Lower	● On Green	State: □ Link OK
	● Flashing Green	State: □ GSP LAN transmit
	● On Red	State: □ GSP Failure

GSP - Revision B

Type	Status	State
GSP Upper	● On Green	State: □ GSP Power OK
	● Flashing Green	State: □ GSP LAN receive
	● On Red	State: □ GSP Failure
GSP Lower	● On Amber	State: □ 10 Base-T Link OK
	● Flashing Amber	State: □ 10 Base-T Activity
	● On Green	State: □ 10 Base-T Link OK
	● Flashing Green	State: □ 10 Base-T Activity

18.652

LAN/SCSI LED States

Type	Status			State
LAN Upper	● On Green	- or -	● Flashing Green	State: □ 100bT Mode
	● On Amber	- or -	● Flashing Amber	State: □ 10bT Mode
LAN Lower	● On Green	- or -	● Flashing Green	State: □ LAN Transmit
SCSI Upper	● On Green			State: □ LVD Mode
	● Off			State: □ Single Ended Mode
SCSI Lower	● On Amber			State: □ Termpower Present
	● Off			State: □ Termpower Absent

Fan, Power Supply, and Disk LED States

Type	Status	State
Fan	● On Amber	State: ▯ Fan Failure
	● Off	State: ▯ Normal operation
Power Supply	● On Green	State: ▯ Normal operation
	● Off	State: ▯ Power Supply Failure
Disk Activity	● On Green or ● Flashing Green	State: ▯ Normal Activity
	● Off	State: ▯ No Activity - Normal
Disk Attention	● On Amber	State: ▯ Ready for HotPlug
	● Off	State: ▯ Normal

18.650

18.619

ANEXO UNIDADE DE BACKUP ROBOTIZADO
PARTE E/2

18.648

7 Removing and Replacing Components

The following list of parts can be changed when required to keep the system running properly. The remove/replace components shown under each part indicates the path required for access to each.

List of Changeable Parts with Remove and Replace Components

NOTE When viewed in PDF format, component remove/replace instructions may be accessed directly by clicking on the component title listed under each part.

Cardcage Fan

Extend the Server out the Front (If Racked)

Stand-alone Server Cover Removal (If Not Racked)

Side Cover Removal

HotSwap Card Cage Fan Removal

HotSwap Card Cage Fan Replacement

Side Cover Replacement

Stand-alone Server Cover Replacement (If Not Racked)

Insert the Server from the Front (If Racked)

Core I/O

Extend the Server out the Front (If Racked)

Stand-alone Server Cover Removal (If Not Racked)

Side Cover Removal

Core I/O Removal

Core I/O Replacement

Side Cover Replacement

Stand-alone Server Cover Replacement (If Not Racked)

Insert the Server from the Front (If Racked)

HotSwap Chassis Fan

Front Bezel Removal (Single Piece)

Front Bezel Removal (Two Piece)

HotSwap Chassis Fan Cover Removal

HotSwap Chassis Fan Removal

HotSwap Chassis Fan Replacement

HotSwap Chassis Fan Cover Replacement

Front Bezel Replacement (Single Piece)

Front Bezel Replacement (Two Piece)

18.646

Disk Drive

Front Bezel Removal (Single Piece)

Front Bezel Removal (Two Piece)

HotPlug Disk Drive Removal

HotPlug Disk Drive Replacement

Front Bezel Replacement (Single Piece)

Front Bezel Replacement (Two Piece)

Display Board

Front Bezel Removal (Single Piece)

Front Bezel Removal (Two Piece)

HotSwap Chassis Fan Cover Removal

HotSwap Chassis Fan Removal

Display Board Removal

Display Board Replacement

HotSwap Chassis Fan Replacement

HotSwap Chassis Fan Cover Replacement

Front Bezel Replacement (Single Piece)

Front Bezel Replacement (Two Piece)

Front Bezel

Front Bezel Removal (Single Piece)

Front Bezel Removal (Two Piece)

Front Bezel Replacement (Single Piece)

Front Bezel Replacement (Two Piece)

One Piece Bezel Install

Memory DIMM

Extend the Server out the Front (If Racked)

Stand-alone Server Cover Removal (If Not Racked)

Top Cover Removal

Memory DIMM Removal

Memory DIMM Replacement

Top Cover Replacement

Stand-alone Server Cover Removal (If Not Racked)

Insert the Server from the Front (If Racked)

18.645

PCI I/O Card

Extend the Server out the Front (If Racked)

Stand-alone Server Cover Removal (If Not Racked)

Side Cover Removal

PCI I/O Card Removal

PCI Card Separator/Extractor Removal

PCI Card Separator/Extractor Replacement

PCI I/O Card Replacement

Side Cover Replacement

Stand-alone Server Cover Replacement (If Not Racked)

Insert the Server from the Front (If Racked)

Power Supply

Front Bezel Removal (Single Piece)

Front Bezel Removal (Two Piece)

HotSwap Power Supply Removal

HotSwap Power Supply Replacement

Front Bezel Replacement (Single Piece)

Front Bezel Replacement (Two Piece)

HotSwap Power Converter Fan

HotSwap Power Converter Fan Removal

HotSwap Power Converter Fan Replacement

Platform Monitor

Extend the Server out the Front (If Racked)

Stand-alone Server Cover Removal (If Not Racked)

Top Cover Removal

Platform Monitor Removal

Platform Monitor Replacement

Top Cover Replacement

Stand-alone Server Cover Replacement (If Not Racked)

Insert the Server from the Front (If Racked)

18.644
3

Processor Support Module

Extend the Server out the Front (If Racked)

Stand-alone Server Cover Removal (If Not Racked)

Top Cover Removal

Processor Support Module Removal

Processor Support Module Replacement

Top Cover Replacement

Stand-alone Server Cover Removal (If Not Racked)

Insert the Server from the Front (If Racked)

Individual Component Remove/Replace Instructions

Each component has instructions for removal followed by instructions for replacement.

Extend the Server out the Front

rp54xx servers are available in two housings: rack-mounted or stand-alone. Access to servers mounted in an HP-supported rack is covered in this section.

NOTE Ensure that there is enough area (Approximately 1.5 meters (4.5 ft) to fully extend the server out the front and work on it.

WARNING **Ensure that all anti-tip features (front and rear anti-tip feet installed; adequate ballast properly placed, etc.) are employed prior to extending the server.**

To extend the server, perform the following steps:

1. Remove the four T-25 screws that fasten the server to the rack.
2. Grasp the server chassis and slowly pull forward. The server is fully extended when the rail clips are locked in place. When fully extended, the top and side service bays are fully accessible.

The following graphic shows the server extended and indicates the rail clip location.

Insert the Server from the Front

rp54xx servers are available in two housings: rack-mounted or stand-alone. Access to servers mounted in an HP-supported rack is covered in this section.

To return the server into the rack, press the rail clips on either side of the server in and push the server into the rack until it stops.

The following graphic shows the server extended and indicates the rail clip location.

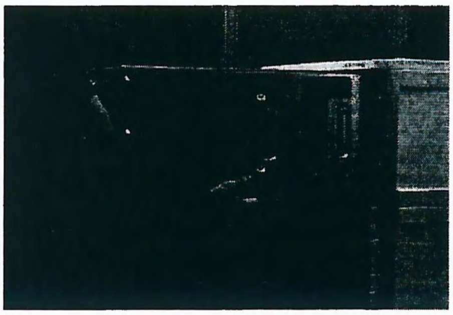

Stand-alone Server Cover Removal

The rp54xx server can be ordered as a stand-alone unit. In this configuration, the server has a one-piece protective cover over it and sits on a platform with locking wheels attached.

To remove the cover from a stand-alone server, perform the following procedures:

1. Unfasten and remove the screws (with captive washers) located near the bottom edge of both sides of the server cover.
2. Lift the protective cover off of the server and set it aside.

WARNING **The stand-alone server weights 69 kg (150 lbs). Removing the stand-alone server from its platform requires three people or a suitable lifting device. Failure to heed this precaution can result in serious personal injury or destruction of the server.**

To remove the server from its wheeled platform, perform the following procedure:

1. Facing the front of the server, reach under the platform on the right side and unfasten the knurl-knobbed, spring-loaded pin (item 1). The pin will retract when it is free. Perform the same step at the rear of the server. The rear knurl-knobbed, spring-loaded pin is aligned behind the front pin, but at the rear of the server.

2. From the back of the server, reach under the platform, behind the knurl-knobbed, spring-loaded pin, and locate the flat head of the second spring-loaded pin (item 2). This pin is spring-loaded to stay up and keep the server from sliding on the platform. Pull the pin down and twist one quarter turn to the left to lock the pin down and out of the way. Perform the same step at the front of the server.
3. The server is now free from its wheeled platform and can be removed from it.

WARNING **The stand-alone server weights 69 kg (150 lbs). Removing the stand-alone server from its platform requires three people or a suitable lifting device. Failure to heed this precaution can result in serious personal injury or destruction of the server.**

The following graphic shows the protective cover.

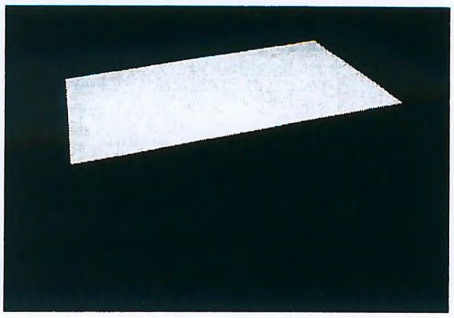

18.639

Stand-alone Server Cover Replacement

The rp54xx server can be ordered as a stand-alone unit. In this configuration, the server has a one-piece protective cover over it and sits on a platform with locking wheels attached.

To place the server on its wheeled platform, perform the following procedure:

WARNING **The stand-alone server weights 69 kg (150 lbs). Placing the stand-alone server on its platform requires three people or a suitable lifting device. Failure to heed this precaution can result in serious personal injury or destruction of the server.**

1. Lift the platform and turn both of the spring-loaded, flat-headed pins (item 1) a quarter turn to the right to lock them out of the way before placing the server on the platform.
2. Using a lifting devise or a minimum of three people, lift the server onto the platform. Align the server and platform so that the pins will lock.
3. Reach under the platform, locate the spring-loaded, flat headed pins and turn them one quarter turn to the left to unlock them. Gently slide the server around on the platform until the spring-loaded flat-head pins snap into their holes.
4. Once again, reach under the platform on the right side and fasten both front and back knurl-knobbed, spring-loaded pins (item 2).

To place the cover on a stand-alone server, perform the following procedures:

1. Set the protective cover on the server and align the holes located near the bottom edge of both sides of the server cover.
2. Fasten the screws (with captive washers) through the cover, into the platform.

The following graphic shows the protective cover.

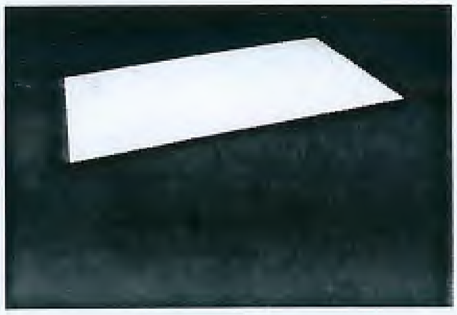

18638

Top Cover Removal

The power to the server does not have to be off to remove the top cover and air baffle. However, operation of the server without the top cover in place can make it susceptible to EMI problems.

Follow the steps listed below to remove the top cover:

1. Loosen the captive T-15 screws that hold the top cover in place.
2. Grasp the strap handles, raise the cover slightly, and pull the cover toward the front of the server to free the cover tabs from the slots in the chassis. The air baffle will be exposed.

Loosen the captive T-15 screws that hold the air baffle in place then lift the air baffle off of the server.

The following graphics show the top service bay cover and the air baffle. The first graphic shows the top service bay cover.

The second shows the air baffle.

18.637

Top Cover Replacement

The power to the server does not have to be off to remove or replace the top service bay cover and air baffle. However, operation of the server without the top cover in place can make it susceptible to EMI problems.

Set the air baffle in place over the opening for the top service bay and tighten the captive T-15 screws.

Follow the steps listed below to replace the top cover:

1. Align the tabs on the end of the top cover with the corresponding slots in the chassis and seat the tabs fully into the slots.
2. Seat the top cover in the top of the service bay and tighten the captive T-15 screws that hold the cover in place.

The following graphics show the air baffle and the top service bay cover. The first graphic shows the air baffle.

The second graphic shows the top service bay cover.

18.636

Side Cover Removal

The Side Cover protects the side service bay. The power to the server does not have to be off to remove the side cover. However, operation of the server without the side cover in place can make it susceptible to EMI problems.

Loosen the captive T-15 screws that hold the side cover in place, then grasp the strap handle and pull the cover away from the server.

The following graphic shows the side cover with captive screw locations.

Side Cover Replacement

The power to the server does not have to be off to replace the side cover. However, operation of the server without the side cover in place can make it susceptible to EMI problems.

Replace the side cover according to the following steps:

1. Grasp the strap handle and insert the tabbed end of the cover into the server chassis slots on the right side of the side service bay.
2. Push the cover into the side service bay opening and fasten the captive T-15 screws that hold the side cover in place.

The following graphic shows the side cover with captive screw locations.

Front Bezel Removal (Single Piece)

The rp54xx server front bezel is hinged on the left (facing the front of the server). The server does not have to be turned off to open or to completely remove the bezel.

Opening the door provides access to the following components:

- HotPlug Disks (A0, A1, and B0, B1).
- Removable Media (CD-ROM drive, etc.).
- HotSwap Chassis Fan cover.
- HotSwap Chassis Fan 0.
- HotSwap Power Supplies (up to three).

To open the front bezel, grasp the right edge of the bezel and pull out. The bezel will swing away from the chassis.

To remove the entire bezel, perform the following steps:

1. Open the front bezel and swing it to the left as far as possible.
2. Pry the hinge cover, located on the left side of the server, off the chassis.
3. While supporting the bezel, remove the screws that secure the bezel hinge to the left side of the server.
4. Grasp the left side of the bezel and pull it loose.

The following graphic shows the front bezel.

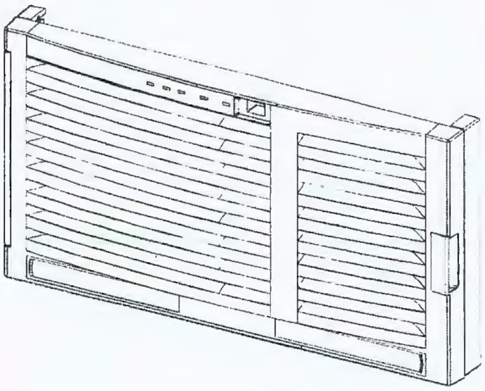

18.633

Front Bezel Replacement (Single Piece)

The rp54xx server front bezel is hinged on the left (facing the front of the server). The server does not have to be turned off to open or to completely remove the bezel.

The bezel encloses the following components:

- HotPlug Disks (A0, A1, and B0, B1).
- Removable Media (CD-ROM drive, etc.).
- HotSwap Chassis Fan cover.
- HotSwap Chassis Fan 0.
- HotSwap Power Supplies (up to three).

To attach the bezel, perform the following steps:

1. Screw the bezel hinge to the threaded inserts on the left side of the server.
2. Install the hinge cover by aligning the notch in the top of the cover with the top of the assembly cover and press firmly into place.

   The following graphic shows the front bezel.

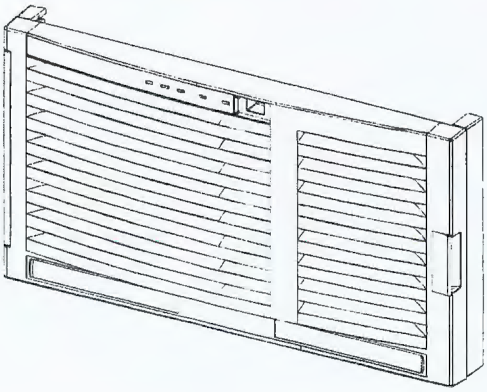

18.632

Front Bezel Removal (Two Piece)

The rp54xx server front bezel is divided and hinged on the right (facing the front of the server) side to provide a door for Disk Media Bay access. The server does not have to be turned off to open the access door or to completely remove the bezel.

Opening the door provides access to the following components:

- HotPlug Disks (A0, A1, and B0, B1).
- Removable Media (CD-ROM drive, etc.).

Removing the entire bezel provides access to the components listed above and the following:

- HotSwap Chassis Fan cover.
- HotSwap Chassis Fan 0.
- HotSwap Power Supplies (up to three).

To open the Disk Media access door, grasp the right edge of the door and pull out. The door will swing away from the chassis, exposing the Disk Media Bay.

To remove the entire bezel, perform the following steps:

1. Open the Disk Media access door, exposing the plastic bezel pins inserted in the clips at the top and bottom of the Disk Media bay on the left side.
2. Grasp the left side of the bezel and pull it loose.

CAUTION DO NOT try to pull the bezel off of the server at this point. The plastic pins holding the bezel on the right side are inserted in metal clips on the Disk Media bay and if the plastic pins are bent, they will break off.

3. Grasp the bezel with both hands and carefully slide the bezel to the left until both plastic pins clear the metal clips on the Disk Media bay.
4. Pull the bezel away from the server and set it aside.

The following graphic shows the front bezel.

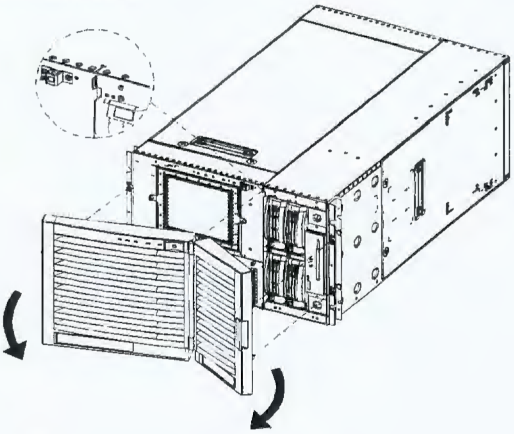

rhrr002

Front Bezel Replacement (Two Piece)

The rp54xx server front bezel is divided and hinged on the right (facing the front of the server) side to provide a door for Disk Media Bay access. The server does not have to be turned off to open the access door or to completely remove the bezel.

The Disk Media Bay door encloses the following components:

- HotPlug Disks (A0, A1, and B0, B1).
- Removable Media (CD-ROM drive, etc.).

The left side of the front bezel encloses the following components:

- HotSwap Chassis Fan cover.
- HotSwap Chassis Fan 0.
- HotSwap Power Supplies (up to three).

To attach the bezel, perform the following steps:

1. Swing the Disk Media access door open so that the plastic pins that go into the metal clips on the Disk Media bay are exposed.

CAUTION The plastic pins holding the bezel on the right side are inserted into metal clips on the Disk Media bay. DO NOT bend the plastic pins or they will break off.

2. Holding the bezel with both hands, align the bezel pins with the metal clips on the Disk Media bay and carefully slide the bezel pins into the clips.
3. Attach the left side of the bezel to the server chassis by aligning the bezel pins with the chassis clipsand press the bezel into the chassis until the pins snap into the clips.

   The following graphic shows the front bezel.

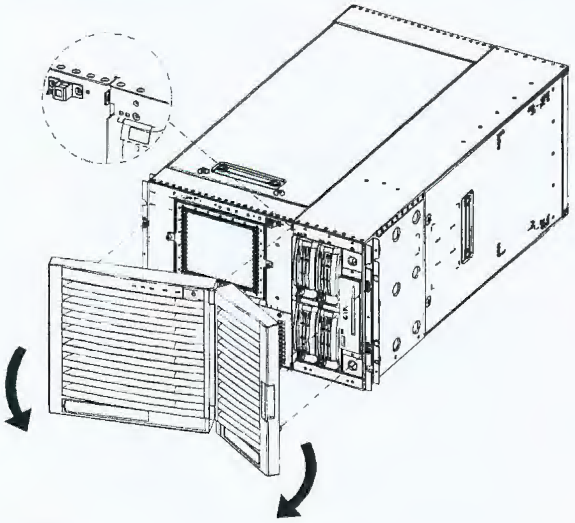

4. Swing the Disk Media access door closed and attach it to the right side of the server chassis.

18.630

Core I/O Removal

rp54xx Core I/O functions are contained on the GSP revision A/B and LAN/SCSI cards. Both cards are located on the PCI Backplane in the side service bay. The LAN/SCSI card is in I/O slot 1 and the GSP revision A/B card is in I/O slot 2.

Before removing either of the Core I/O cards from the server, perform the following tasks:

- Power down the server.
- Detach all power cords from the server.

To remove a Core I/O card from the server, perform the following steps:

1. Remove all cables attached to the Core I/O card at the rear bulkhead.

NOTE Be sure to label the cables before removing them.

2. Disconnect any ribbon cable connectors attached to the Core I/O card in the side service bay.
3. Grasp the edge of the Core I/O card and pull it out of the server.

The following graphic shows both Core I/O cards in the side service bay.

18.629

Core I/O Replacement

rp54xx Core I/O functions are contained on the GSP and LAN/SCSI cards. Both cards are located on the PCI Backplane in the side service bay. The LAN/SCSI card is in I/O slot 1 and the GSP card is in slot 2.

Before replacing either of the Core I/O cards, perform the following tasks:

- Power down the server.
- Detach all power cords from the server.

NOTE For revision B GSP only, remove *one* of the MAC address labels and place on the rear of the server as shown in the following graphics.

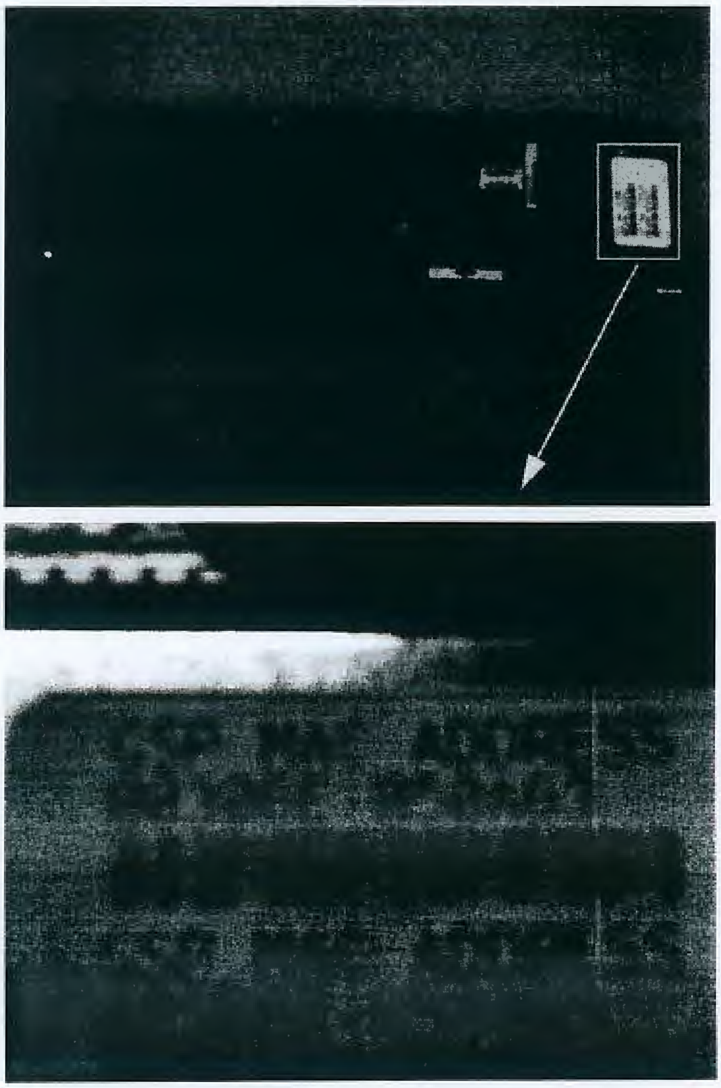

To replace a Core I/O card, perform the following steps:

1. Orient the I/O card in its guide and push it into the server until the card connector seats in the I/O Backplane card slot.
2. Connect the I/O cable attached to the I/O card at the rear PCI bulkhead.
3. Connect any ribbon cable connectors attached to the I/O card in the side service bay.

The following graphic shows an I/O card being replaced.

PCI Card Separator/Extractor Removal

PCI card separator/extractors are located in the Side Service Bay, between the PCI cards. PCI card separator/extractors are plastic cards with two tab handles and two LEDs.

Before removing a PCI card separator/extractor from the server, perform the following tasks:

- Shut down the server.
- Remove the corresponding PCI card.

To remove a PCI card separator/extractor, perform the following steps:

1. Carefully slide a long shaft, medium width, flat blade screw driver (item 1) into the Side Service Bay along the side of the PCI card separator/extractor, and insert it into the slot in the hook tab (item 2).

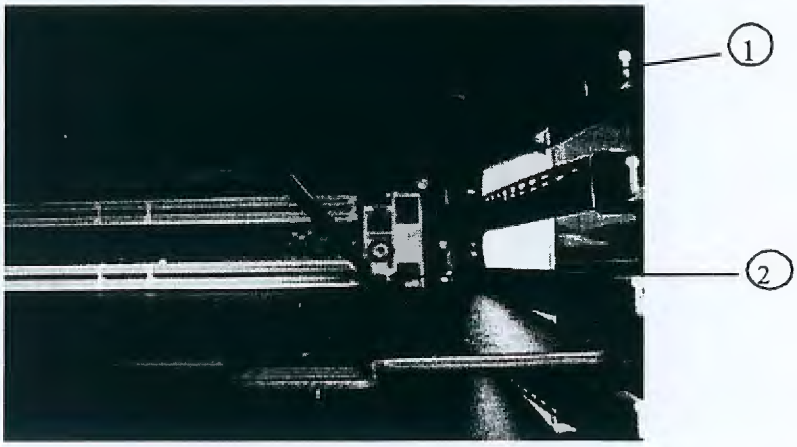

2. At the point where the PCI card separator/extractor inserts into the PCI backplane connector (item 2), there are two slotted, hook tab connectors on the PCI card separator/extractor, one on either side of the PCI Backplane connector. Press one then the other to disengage the card from the backplane.
3. With the two connection points pressed, slide the PCI card separator/extractor out of the PCI card cage.

PCI Card Separator/Extractor Replacement

PCI card separator/extractors are located in the Side Service Bay on the PCI Backplane, between PCI cards. PCI card separator/extractors are plastic cards with two tab handles and two LEDs.

To replace a PCI card separator/extractor, perform the following steps:

1. Insert the PCI card separator/extractor into the available slot and slide it into the PCI backplane connector.
2. Be sure the two hook tabs on the PCI card separator/extractor insert into the connector blocks on either side of the PCI backplane.

Once you have completed replacement of the PCI card separator/extractor into the server, perform the following tasks:

- Replace the PCI card, if necessary.
- Power up the server.

18.625

HotPlug Disk Drive Removal

The internal disk drives (up to four) are located at the front right side of the server (as you are facing it). When proper software and hardware procedures are followed, internal disk drives can be removed and replaced while the server is running.

The procedures in this section are in two parts, a HotSwap Software Procedure and a HotPlug Hardware Procedure, for removing the disk from the server.

Before starting these procedures, you must have an up-to-date configuration backup file. Configuration backup is performed by default each time an LVM command changes the LVM configuration. The default backup files path is:

```
/etc/lvmconf/base_vg_name.conf.
```

The replacement disk drive must be the same product ID as the disk drive that is being replaced.

NOTE HP often uses different manufacturers for disks that have the same product number. The HotSwap and HotPlug procedures will not update the disk drive's internal information to that of the replaced disk drive.

The replacement disk drive will have the same capacity and blocksize as the defective disk because they have the same product number. The only field that can be incorrect is the string specifying the vendor's name. This will not affect the behavior of the LVM. If you desire to update the manufacturer's name, the disks volume group must be deactivated and reactivated.

HotSwap Software Procedure

Perform these software procedural steps to replace a HotPlug disk drive device:

Step 1. Determine whether LVM found the physical volume to be defective when the volume group was activated.

If the volume was defective when the volume group was activated, *vgchange* will generate the following message to the console:

```
WARNING

VGCHANGE: WARNING: COULDN'T ATTACH TO THE VOLUME GROUP PHYSICAL
VOLUME "/DEV/DSK/cXtXdX"

VGCHANGE: WARNING: COULDN'T ATTACH TO THE VOLUME GROUP PHYSICAL
VOLUME "/DEV/DSK/cXtXdX"

THE PATH OF THE PHYSICAL VOLUME REFERS TO A DEVICE THAT DOES NOT
EXIST, OR IS NOT CONFIGURED INTO THE KERNAL.
```

If you are unsure of the status of *vgchange*, check it with the *vgdisplay* command:

#vgdisplay <VG name>

For example:

vgdisplay /dev/vg00

If the disk was defective when the *vgchange* command was entered, the following message will be printed one or more times:

```
WARNING:

VGDISPLAY: WARNING: COULDN'T QUERY PHYSICAL VOLUME "/DEV/DSK/cXtXdX"
THE SPECIFIED PATH DOES NOT CORRESPOND TO PHYSICAL VOLUME ATTACHED TO THE
VOLUME GROUP.

VGDISPLAY: WARNING: COULDN'T QUERY ALL OF THE PHYSICAL VOLUMES.
```

If you see these messages, the disk was defective at the time the volume group was activated. Remove the bad disk as described in step 9 of this procedure, then follow the instructions in for replacing the disk and perform the *HotSwap Procedure for Unattached Physical Volumes* described there.

Otherwise, the disk drive became defective after *vgchange* was run. Proceed to Step 2.

Step 2. Display the names of all the logical volumes on this volume group with the *vgdisplay* command. For example:

```
#vgdisplay /dev/vg00
```

Step 3. Determine which logical volumes have mirrors with the *lvdisplay* command. For example:

```
#lvdisplay /dev/vg00/lvol# | grep -ie "LV Name" -e "Mirror"
```

Step 4. Determine the *pvkey* command status for the mirrored logical volume, again using the *lvdisplay* command with the *-k* option. Compare the output to the *lvdisplay* command with the *-v* option to determine the device file to *pvkey* mapping. For example,

```
# lvdisplay -v -k /dev/vg00/lvol1
--- Logical volumes ---
LV Name                     /dev/vg00/lvol1
VG Name                     /dev/vg00
LV Permission               read/write
LV Status                   available/syncd
Mirror copies               1
Consistency Recovery        MWC
Schedule                    parallel
LV Size (Mbytes)            256
Current LE                  64
Allocated PE                128
Stripes                     0
Stripe Size (Kbytes)        0
Bad block                   off
Allocation                  strict/contiguous
IO Timeout (Seconds)        default

   --- Distribution of logical volume ---
   PV Name            LE on PV  PE on PV
   /dev/dsk/c1t6d0    64        64
   /dev/dsk/c2t6d0    64        64

   --- Logical extents ---
   LE     PV1               PE1   Status 1 PV2                PE2   Status 2
   00000      0             00000 current      1              00000 current
   00001      0             00001 current      1              00001 current
   00002      0             00002 current      1              00002 current
```

```
    00003       0           00003  current       1           00003  current
    00004       0           00004  current       1           00004  current
    00005       0           00005  current       1           00005  current
    00006       0           00006  current       1           00006  current
    00007       0           00007  current       1           00007  current
    00008       0           00008  current       1           00008  current

(etc.)
```

```
# lvdisplay -v /dev/vg00/lvol1
--- Logical volumes ---
LV Name                      /dev/vg00/lvol1
VG Name                      /dev/vg00
LV Permission                read/write
LV Status                    available/syncd
Mirror copies                1
Consistency Recovery         MWC
Schedule                     parallel
LV Size (Mbytes)             256
Current LE                   64
Allocated PE                 128
Stripes                      0
Stripe Size (Kbytes)         0
Bad block                    off
Allocation                   strict/contiguous
IO Timeout (Seconds)         default

   --- Distribution of logical volume ---
   PV Name            LE on PV   PE on PV
   /dev/dsk/c1t6d0    64         64
   /dev/dsk/c2t6d0    64         64

   --- Logical extents ---
   LE    PV1               PE1   Status 1 PV2               PE2   Status 2
   00000 /dev/dsk/c1t6d0   00000 current  /dev/dsk/c2t6d0   00000 current
   00001 /dev/dsk/c1t6d0   00001 current  /dev/dsk/c2t6d0   00001 current
   00002 /dev/dsk/c1t6d0   00002 current  /dev/dsk/c2t6d0   00002 current
   00003 /dev/dsk/c1t6d0   00003 current  /dev/dsk/c2t6d0   00003 current
   00004 /dev/dsk/c1t6d0   00004 current  /dev/dsk/c2t6d0   00004 current
   00005 /dev/dsk/c1t6d0   00005 current  /dev/dsk/c2t6d0   00005 current
   00006 /dev/dsk/c1t6d0   00006 current  /dev/dsk/c2t6d0   00006 current
   00007 /dev/dsk/c1t6d0   00007 current  /dev/dsk/c2t6d0   00007 current
   00008 /dev/dsk/c1t6d0   00008 current  /dev/dsk/c2t6d0   00008 current
```

The *pvkey* stuatus (0 or 1 in this example) shown in the first command, maps to the device file names (*/dev/dsk/c1t6d0* or */dev/dsk/c2t6d0*) in the second command under columns *PV1* and *PV2*, respectively.

Step 5. Reduce any logical volumes that have mirror copies on the faulty disk drive so that they no longer mirror onto that disk drive (note the *-A n* option):

```
# lvreduce -m 0 -A n -k <LV name> /dev/dsk/cXtXcX <pvkey#>&
```

(for 1 way mirroring)

OR

```
# lvreduce -m 1 -A n -k <LV name> /dev/dsk/cXtXcX <pvkey#>&
```

(for 2way mirroring)

For example, enter:

```
# lvreduce -m 0 -A n -k /dev/vg00/lvol4 /dev/dsk/c2t4d0 1&
```

The following message will appear:

```
Logical volume /dev/vg00/lvol4 has been successfully reduced.
lvlnboot: Logical Volume has no extents.
```

NOTE It is important to include the ampersand (&) at the end of the command line. Otherwise, the *lvreduce* process will hang, and you will need terminal control to negate the command. Once the 'successfully reduced' message has been generated, manually end the process using the *kill -9* command.

Step 1. Use the *ps* command to find the PID for the *lvreduce* process.

```
# ps -ef | grep lvreduce
```

Step 2. Manually end the process with the *Kill -9* command (It may take several minutes for the process to end.).

```
# kill -9 <PID>
```

Step 3. Repeat steps 4 and 5 for all logical volumes.

Step 4. With all logical volumes reduced, now reduce the volume group using the *vgreduce* command. For example:

```
# vgreduce /dev/vg00 /dev/dsk/c2t6d0
```

CAUTION The *vgreduce* procedure may take a long time to complete. Do NOT terminate this process.

Step 5. Update the disk BDRA using the `lvlnboot` command.

```
# lvlnboot -R
```

Step 6. Proceed to the *HotPlug Hardware Procedure* to remove the bad disk drive from the server.

HotPlug Hardware Procedure

CAUTION Disk Drives can be removed or installed with the server still powered on. This is referred to as a "manual HotPlug".

To remove a disk drive from the server, perform the following step:

18621

Grasp the tab at the bottom of the cam latch on the selected disk drive, then push the button inside the cam latch and pull the cam latch out and up. The disk drive will unlock. Pull gently until it slides completely free.

The following graphic shows disk features.

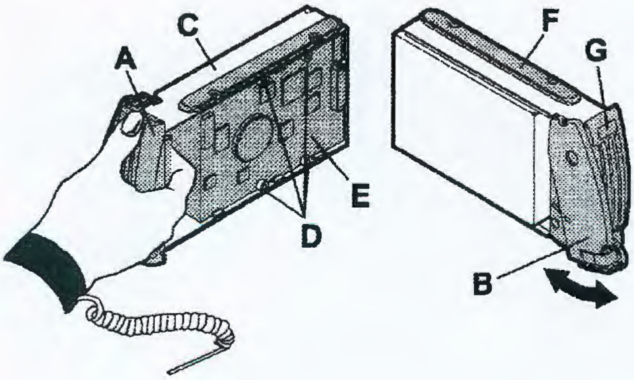

A	bezel handle
B	cam latch
C	carrier frame
D	standoffs
E	circuit board
F	insertion guide
G	capacity label

The next graphic depicts disk removal/replacement.

HotPlug Disk Drive Replacement

The internal disk drives (up to four) are located at the front right side of the server (as you are facing it). When proper software and hardware procedures are followed, internal disk drives can be removed and replaced while the server is running.

CAUTION Disk Drives can be removed or installed with the server still powered on. This is referred to as a “manual HotPlug”.

However, DO NOT replace a HotPlug disk drive until a controlled shutdown of the operating system has been performed.

Hardware HotPlug Procedure

To replace a disk drive in the server, grasp the tab at the bottom of the cam latch on the selected disk drive, push the button inside the cam latch, and pull the cam latch out and up. The disk drive will unlock. Pull gently until it slides completely free.

The following graphic shows disk features.

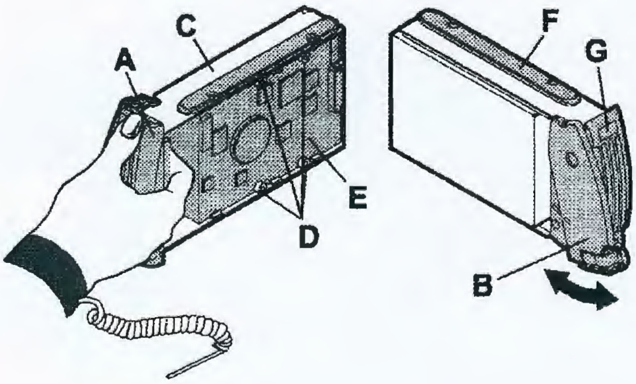

A	bezel handle
B	cam latch
C	carrier frame
D	standoffs
E	circuit board
F	insertion guide
G	capacity label

The next graphic depicts disk removal/replacement.

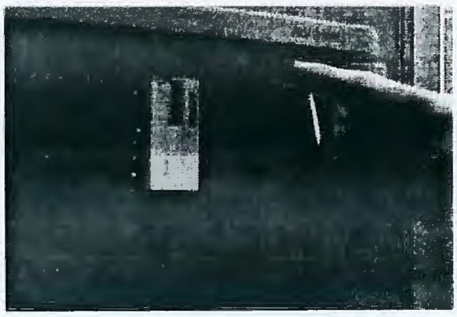

18.659

Hot Swap Software Procedure for Attached Physical Volumes

The following procedure is an example of how to recover from replacing a disk that was recognized as an attached physical volume. Be sure that all the software procedures in the *Disk Drive Removal* section have been completed prior to starting these steps. This example assumes that the disk was mirrored.

NOTE HP often uses different manufacturers for disks, but assigns the same product number. The hot swap manual procedure will not update disk driver internal information to that of the replaced disk drive.

Step 1. Perform an *ioscan* on the replaced disk drive to ensure that it is accessible (claimed), to double check that it is a proper replacement, and that the device files are present. Refer to the above note.

For example: `# ioscan -fnC disk`

Step 2. Use the following procedure to mirror the root disk:

a. Create the new physical volume using the *pvcreate* command.
For example: `# pvcreate -B /dev/rdsk/cXtXdX`

b. Extend the volume group to include the new physical volume using the *vgextend* command:
For example: `# vgextend /dev/vg00 /dev/dsk/cXtXdX`

c. The *mkboot* command must be run to make the device bootable.
For example:`# mkboot /dev/rdsk/cXtXdX`

d. Use the *mkboot* command again to add the HP-UX auto-file-string.
For example: `# mkboot -a "hpux" /dev/rdsk/cXtXdX`

e. Run *lvlnboot* with the following command: `# lvlnboot -R`

f. Run *lvextend* to put a mirror into the replaced disk drive. It may take several minutes to copy the original copy of the data to the mirrored extents. The logical volume(s) will still be accessible to user applications during this operation.

For example:
```
# lvextend -m 1 <LV name> /dev/dsk/cXtXdX
OR
# lvextend -m 2 <LV name> /dev/dsk/cXtXdX (for 3 way mirroring)
```

For example:
```
# lvextend -m 1 /dev/vg00/lvol4 /dev/dsk/cXtXdX
OR
# lvextend -m 1 /dev/vg00/lvol5 /dev/dsk/cXtXdX
```

Repeat this for each logical volume to be mirrored.

Verify that the mirror is bootable and AUTO file is correct.

```
For example:
# lifls -l /dev/rdsk/cXtXdX
OR
# lifcp /dev/rdsk/cXtXdX:AUTO -
```

g. Verify that the mirroring is set up properly.

For example:

```
# vgdisplay -v /dev/vg00
# lvdisplay /dev/vg00/lol1 - lvol8
# lvlnboot -v /dev/vg00
```

Both disks should list as "Boot Disk" and both should appear in the *lvol* lists.

At this point the system will be fully functional.

Hot Swap Procedure for Unattached Physical Volumes

The following steps are an example of how to replace a HotPlug disk drive for unattached physical volumes. This example assumes the disks are mirrored.

NOTE HP often uses different manufacturers for disks, but assigns the same product number. The hot swap manual procedure will not update disk driver internal information to that of the replaced disk drive.

Step 1. Perform an *ioscan* on the replaced disk drive to ensure that it is accessible (claimed), to double check that it is a proper replacement, and that the device files are present. Refer to the above note.

For example: `# ioscan -fnC disk`

Step 2. Restore the LVM configuration/headers onto the replaced disk drive from your backup of the LVM configuration with the following entry:

```
# vgcfgrestore -n <volume group name> /dev/rdsk/cXtXdX
```

For example: `# vgcfgrestore -n /dev/vg00 /dev/rdsk/cXtXdX`

Step 3. Attach the new disk drive to the active volume group with the following *vgchange* command:

```
# vgchange -A y <volume group name>
```

For example: `# vgchange -A y /dev/vg00`

Step 4. Use the *mkboot* command to make the device bootable.

For example: `# mkboot /dev/rdsk/cXtXdX`

Step 5. Use the *mkboot* command again to add the HP-UX auto-file-string.For example: `# mkboot -a "hpux" /dev/rdsk/cXtXdX`

Step 6. Run *lvlnboot* with the following command: `# lvlnboot -R`

Step 7. Resynchronize the mirrors of the replaced disk drive with the following command. It may take several minutes to copy all the data from the original copy of the data to the mirrored extents. The logical volume(s) are still accessible to users' applications during this command.

```
# vgsync <VG name>
```

For example: `# vgsync /dev/vg00`

At this point the system will be fully functional.

HotSwap Chassis Fan Cover Removal

Power to the server does not have to be off to remove or replace a HotSwap Chassis fan cover. Fan number 0 is in the front of the server and fan number 1 is at the rear of the server.

To remove a fan cover from the server, perform the following steps:

1. Loosen the captive T-15 screws from the sides of the cover.
2. Gently pry the cover away from the server and set it aside.

The following graphic shows a Chassis Fan Cover in place.

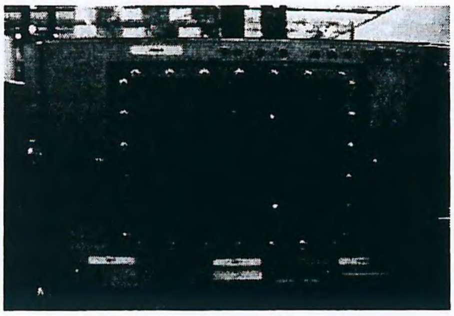

HotSwap Chassis Fan Cover Replacement

The power to the server does not have to be off to remove or replace a HotSwap Chassis Fan cover. Fan number 0 is located in the front of the server and fan number 1 is located at the rear of the server.

To replace a chassis fan cover, perform the following tasks:

1. Insert the cover into position in front of the fan.
2. Tighten the captive T-15 screws on each side of the cover.

The following graphic shows a HotSwap Chassis Fan Cover.

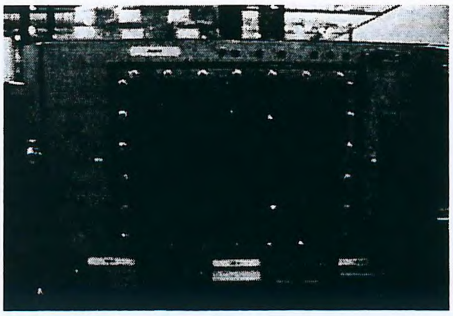

fl. 615

HotSwap Chassis Fan Removal

The power to the server does not have to be off to remove or replace a HotSwap Chassis fan. Fan number 0 is in the front of the server and fan number 1 is at the rear of the server.

To remove a fan from the server, perform the following steps:

1. Check the fan LED located on the fan. If the LED is illuminated, the fan has failed.

NOTE When one fan has failed (or is removed from the server), the system automatically puts the remaining fan into high speed mode. The noise level of the server will increase.

2. Grasp the fan grill and gently pull toward you. The fan assembly will unplug from the electrical outlet and slide out of the server.

CAUTION Running the server for extended periods of time with a cooling fan removed may create hot spots inside the server and possibly shorten component life.

If the other fan fails when one fan is removed, the system will halt.

The following graphic shows a HotSwap Chassis Fan.

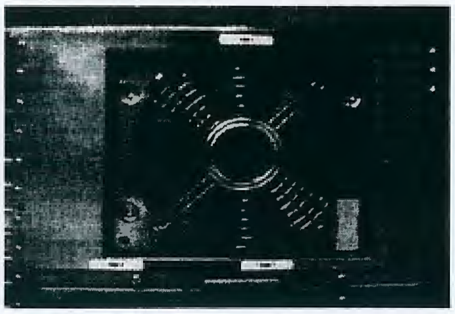

18.684

HotSwap Chassis Fan Replacement

The power to the server does not have to be off to remove or replace a HotSwap Chassis fan. Fan number 0 is located in the front of the server and fan number 1 is located at the rear of the server.

To replace a fan, perform the following steps:

1. Orient the fan assembly to ensure that the chassis outlet connects with the fan assembly. If the fan power connector is covered by the protective cover (see item 1 in the graphic), slide it down to remove it and slide it over the power connection on the other side of the fan. Push the fan firmly into the housing. The fan assembly will plug into the electrical outlet automatically.
2. Check the LED located on the fan.
   - When the fan is functioning normally, the LED is OFF.
   - When the fan fails, the LED is ON.

NOTE when one fan has failed (or is removed from the server), the system automatically puts the remaining fan into high speed mode. The noise level of the server will increase.

3. Replace the HotSwap Chassis fan cover.

The following graphic shows a HotSwap Chassis Fan.

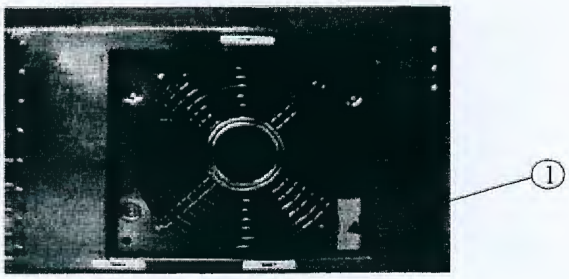

18.613

HotSwap Card Cage Fan Removal

The power to the server does not have to be off to remove or replace a HotSwap Card Cage fan. Fans numbered 2, 3, 4, and 5 are located in a four-fan assembly housing located on the left side of the server's side service bay.

CAUTION Running the server for extended periods of time with a cooling fan removed may create hot spots inside the server and shorten component life.

If other fans fail when one fan is removed, the system will halt.

To remove a fan from the server, perform the following step:

Identify the fan to be removed and pull it out of the Side Fan Assembly Housing. It will automatically disconnect from its electrical outlet on the I/O Backplane.

NOTE When one fan has failed (or is removed from the server) during operation, the system automatically puts the remaining fans into high speed mode. The noise level of the server will increase.

The following graphic shows a card cage fan being removed/replaced.

18.632

HotSwap Card Cage Fan Replacement

The power to the server does not have to be off to remove or replace a HotSwap Card Cage fan. Fans numbered 2, 3, 4, and 5 are located in a four-fan assembly housing located on the left side of the server's side service bay.

To replace a fan in the server, orient the replacement fan into its slot in the Side Fan Assembly Housing and carefully push it in until it connects with its electrical outlet on the I/O Backplane.

The following graphic shows a card cage fan being removed/replaced.

HotSwap Power Supply Removal

Up to three power supplies (0 through 2), located across the bottom front of the server, can be installed in the server without removing power.

To remove a power supply from the server, perform the following step:

1. Remove the T-15 mounting screw located to the right of the handle near the top of the power supply.

CAUTION Be careful when pulling the power supply out of the server. It is heavier than it appears.

2. Grasp the handle and pull the power supply out of the server.

The following graphic shows a front and rear view of a HotSwap Power Supply.

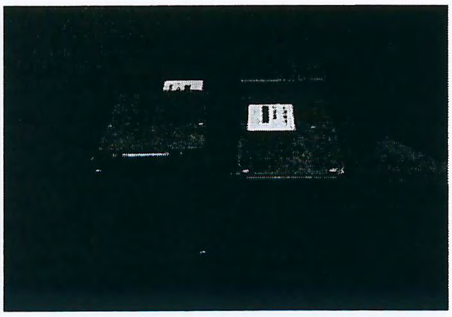

HotSwap Power Supply Replacement

Up to three power supplies (0 through 2), located across the bottom front of the server, can be installed in the server without removing power.

To replace a power supply, perform the following steps:

CAUTION *Be careful when putting the power supply into the server. It is heavier than it appears.*

1. Grasp the handle in one hand and support the power supply with the other. Slide the power supply into the server. The Power Supply LED should illuminate immediately.
2. Replace the T-15 mounting screw located to the right of the handle near the top of the power supply.

The following graphic shows a front and rear view of a HotSwap Power Supply.

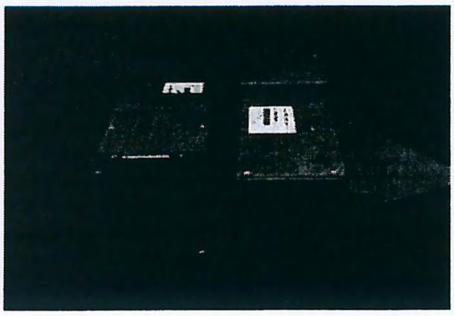

18.609

HotSwap Power Converter Fan Removal

The power to the server does not have to be off to remove or replace a HotSwap Power Converter fan. Fans numbered 6 and 7 are located in the rear of the server.

To remove a fan from the server, perform the following steps:

NOTE When one fan has failed (or is removed from the server), the system automatically puts the remaining fan into high speed mode. The noise level of the server will increase.

1. There are four screws attached to each corner of each fan. Loosen only the captive Torx-head screws located diagonally across the face of the fan (upper left, lower right sides).
2. Grasp the extended screw-heads (or the fan grill) and gently pull toward you. The fan assembly will unplug from the electrical outlet and slide out of the server.

CAUTION Running the server for extended periods of time with a cooling fan removed may create hot spots inside the server and possibly shorten component life.

If the other fan fails when one fan is removed, the system will halt.

The following graphic shows where HotSwap Power Converter Fans are located.

HotSwap Power Converter Fan Replacement

The power to the server does not have to be off to remove or replace a HotSwap Power Converter fan. Fans numbered 6 and 7 are located in the rear of the server.

To replace a fan into the server, perform the following steps:

1. Orient the fan assembly so that the electrical plug will connect, then grasp the extended screw-heads (or the fan grill) and gently push the fan assembly into its housing. The fan assembly plug will connect with the electrical outlet in back of the housing.
2. Tighten the captive T-15 screws located diagonally across the face of the fan (upper left, lower right sides).

The following graphic shows where HotSwap Power Converter Fans are located.

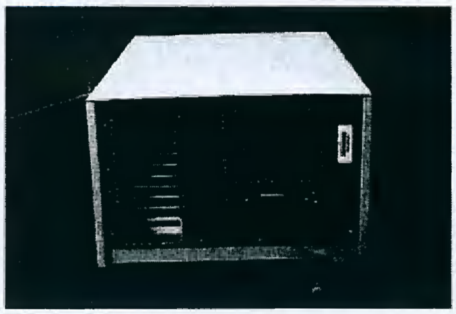

18.607

Processor Support Module Removal

Processor Support Modules (PSMs) reside on the System Board and are accessed via the Top Service Bay. Looking into the Top Service Bay from the front, PSMs are located on either side of the server, at the front. There can be two PSMs, numbered 0 and 1

To remove a PSM from the server, perform the following steps:

1. Loosen the two captive mounting screws that hold the PSM in place.

NOTE For the rp5470, the mounting screws have been replaced by posts and the air baffle is used to secure the PSM's.

2. Grasp the two captive mounting screws and lift the PSM out of the server.

The following graphic shows a PSM in the server.

18.606

Processor Support Module Replacement

Processor Support Modules (PSMs) (there can be two PSMs, numbered 0 and 1) reside on the System Board and are accessed via the Top Service Bay. Looking into the Top Service Bay from the front, PSMs are located on either side of the server, at the front.

To replace a PSM, perform the following steps:

1. Seat the PSM into its socket.
2. Tighten the two captive mounting screws that hold the PSM in place.

NOTE For the rp5470, the mounting screws have been replaced by posts and the air baffle is used to secure the PSM's.

The following graphic shows a PSM in the server.

18.605

Memory DIMM Removal

Memory DIMMs reside in slots (up to eight) located on the System Board. They are loaded in DIMM pairs of equal size.

Removing rp5400 Memory DIMMs

To remove an rp5400 memory DIMM, perform the following steps:

1. Press down on the extractor levers on each end of the selected memory DIMM to unseat the DIMM from its socket.
2. When the memory DIMM unseats from the socket, pull it up and away from the System Board.

The following graphics show a memory DIMM, followed by a display of a DIMM being removed/replaced.

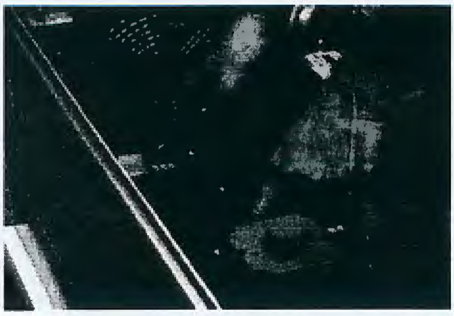

Removing rp5470 Memory DIMMs

To remove an rp5470 memory DIMM, perform the following steps:

1. Pull up on the extractor levers on each end of the Memory Carrier to unseat the Memory Carrier from its socket.
2. When the Memory Carrier unseats from the socket, pull it up and away from the System Board.
3. Loosen the captive screws that secure the DIMM Clip and remove DIMM Clip from the Memory Carrier.
4. Press down on the extractor levers on each end of the selected memory DIMM to unseat the DIMM from its socket.
5. When the memory DIMM unseats from the socket, pull it up and away from the Memory Carrier.

18.604

Memory DIMM Replacement

Memory DIMMs reside in slots located on the System Board (up to eight). They are loaded in DIMM pairs of equal size.

Replacing rp5400 and/or rp5450 Memory DIMMs

To replace a memory DIMM, perform the following steps:

1. Seat the memory DIMM into its socket.
2. Press the extractor levers on each end of the memory DIMM slot inward until the levers snap into place.

The following graphics show a memory DIMM and a display of DIMM removal/replacement.

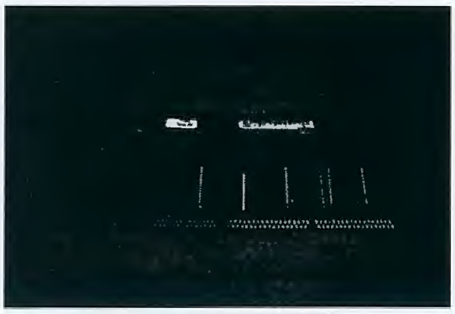

Replacing rp5470 Memory DIMMs

To replace an rp5470 memory DIMM, perform the following steps:

1. Seat the memory DIMM into its socket on the Memory Carrier.
2. Press the extractor levers on each end of the memory DIMM slot inward until the levers snap into place.
3. Attach the Memory Clip to the Memory Carrier with the DIMM slot markings on the top of the Memory Clip aligned with the DIMM slot markings on the Memory Carrier. Secure the Memory Clip using the captive screws.
4. Seat the Memory Carrier into the slot on the System Board.
5. Push down on the extractor levers and snap them into place.

603

Display Board Removal

The Display Board contains the server's ON/OFF switch and five LEDs that indicate server status when power is applied.

CAUTION The Display Board is not a HotSwap or HotPlug unit. Ensure that the server is powered-down prior to removal.

To remove the Display Board, perform the following tasks:

1. Remove the three T-10 screws that hold the Display Board in place near the top of the chassis front.
2. Remove the two T-15 screws that hold the front Chassis Fan in place and extract the fan from the server.
3. Reach up through the top of the Chassis Fan cavity and carefully pull the Display Board back to free the LEDs and the On/Off switch from their chassis openings. Pull the Display Board down through the Chassis Fan cavity.
4. Disconnect the ribbon cable from the Display Board, and place the display board on a suitable work surface.

The following graphic shows the Display Board access location (item 1) (looking up from the bottom of the front of the server).

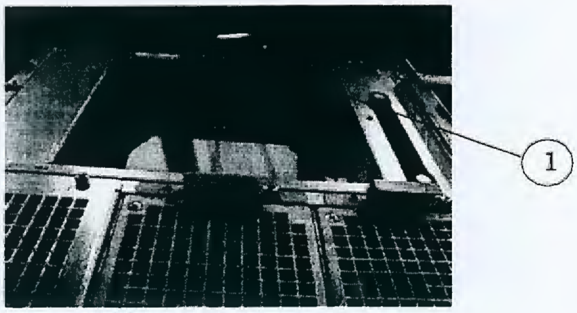

The next graphic shows the Display Board module.

Display Board Replacement

The Display Board contains the server's ON/OFF switch and five LEDs that indicate server status when power is applied.

CAUTION The Display Board is not a HotSwap or HotPlug unit. Ensure that it is powered-down prior to removal.

To replace the Display Board, perform the following tasks:

1. Reconnect the ribbon cable to the Display Board.
2. Carefully push the Display Board up through the top of the Chassis Fan cavity and insert the LEDs and On/Off switch into their respective chassis openings.
3. Replace three T-10 screws.

The following graphic shows the Display Board location (item 1) (looking up from the bottom of the front of the server).

The next graphic shows the Display Board module.

18.601

Platform Monitor Removal

The Platform Monitor resides on the System Board and is accessed via the Top Service Bay.

Removing rp5400 and/or rp5450 Model Platform Monitors

Looking into the Top Service Bay from the front, the rp5400/rp5450 model Platform Monitor is located on the left side at the front of the server.

To remove the rp5400/rp5450 model Platform Monitor, perform the following steps:

1. Pull up on the extractor levers on each end of the Platform Monitor to unseat it from its socket.
2. When the Platform Monitor unseats from the socket, pull it up and away from the System Board.

The following graphics show a Platform Monitor board followed by a display of Platform Monitor removal/replacement.

Removing The rp5470 Model Platform Monitor

Looking into the Top Service Bay from the front, the rp5470 model Platform Monitor is located on the right side at the front of the server.

To remove an rp5470 model Platform Monitor, perform the following steps:

1. Pull up on the extractor levers on each end of the Platform Monitor to unseat it from its socket.
2. When the Platform Monitor unseats from the socket, pull it up and away from the System Board.

18.600

The following graphics show a Platform Monitor board followed by a display of Platform Monitor removal/replacement.

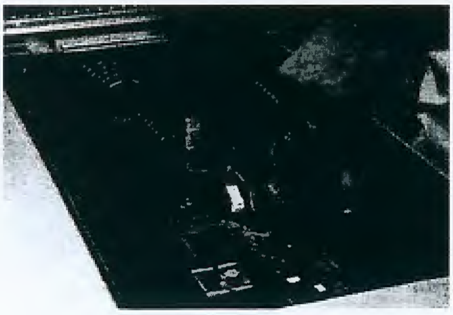

Platform Monitor Replacement

The Platform Monitor resides on the System Board and is accessed via the Top Service Bay.

Replacing rp5400 and/or rp5450 Platform Monitors

Looking into the Top Service Bay from the front, the rp5400/rp5450 model Platform Monitor is located on the left side at the front of the server.

To replace a Platform Monitor, perform the following steps:

1. Seat the Platform Monitor into its socket.
2. Lift the extractor levers and press them onto each end of the Platform Monitor until the levers snap into place.

The following graphics show a Platform Monitor board followed by a display of Platform Monitor removal/replacement.

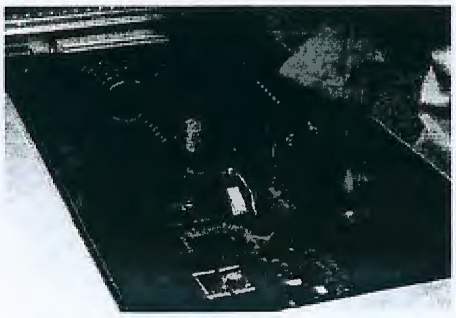

Replacing rp5470 Platform Monitor

Looking into the Top Service Bay from the front, the rp5470 model Platform Monitor is located on the right side at the front of the server.

To replace a Platform Monitor, perform the following steps:

1. Seat the Platform Monitor into its socket.
2. Lift the extractor levers and press them onto each end of the Platform Monitor until the levers snap into place.

18.598

The following graphics show a Platform Monitor board followed by a display of Platform Monitor removal/replacement.

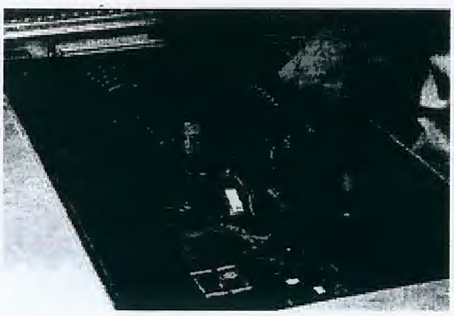

18.597

PCI I/O Card Removal

The side service bay contains card slots for ten PCI I/O cards (slots 3 through 12) and two Core I/O cards (slots 1 and 2).

Perform the following tasks prior to removing PCI I/O cards:

- Power down the server.
- Detach all power cords from the server.

To remove a PCI I/O card from the server, perform the following steps:

NOTE Record the location of all PCI cards as they are removed. Replacing them in a diffferent location will require system reconfiguration and could cause boot failure.

1. Disconnect the I/O cable attached to the I/O card at the rear PCI bulkhead.
2. Disconnect any ribbon cable connectors attached to the I/O card in the side service bay.
3. Grasp the edge of the I/O card and pull it out of the server.

The following graphic shows an I/O card being removed.

PCI I/O Card Replacement

The side service bay contains card slots for 10 PCI I/O cards (slots 3 through 12) and 2 Core I/O cards (slots 1 and 2).

Prior to replacing PCI cards, perform the steps listed below:

- Power down the server.
- Detach all power cords from the server.

To replace a Core or PCI I/O card, perform the following steps:

1. Locate the I/O card guide (item 1) on the outside of the Fan Assembly Housing. Orient the I/O card into its guide slot and push it into the server until the card connector seats in the I/O Backplane card connector.

NOTE Each I/O card guide contains two slots. The top slot is aligned with the I/O Backplane card connector.

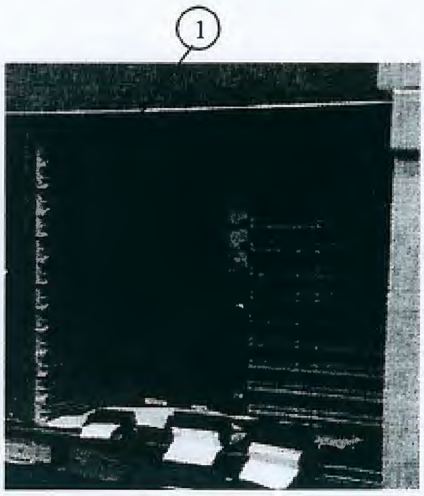

2. Connect the I/O cable attached to the I/O card at the rear PCI bulkhead.
3. Connect any ribbon cable connectors attached to the I/O card in the side service bay.

The following graphic shows an I/O card being replaced.

18.594

A Parts and Accessories

18.593

CRU Physical Location

This section contains views of the rp54xx computer. The locator numbers in the diagrams correspond to the numbers in the CRU Part Number section.

Figure A-1 Server Rear View

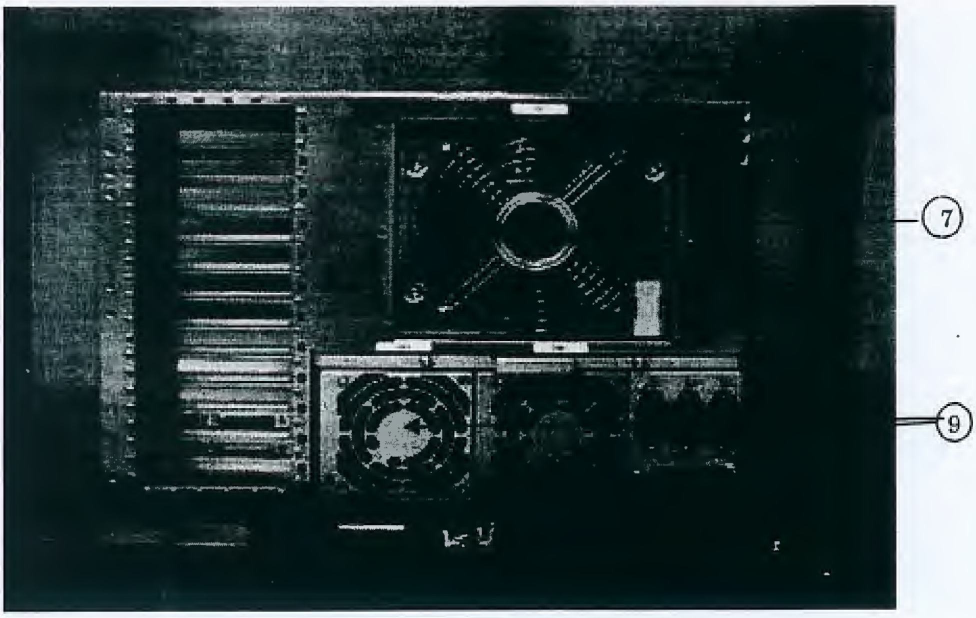

18.592

Figure A-2 Side Service Bay

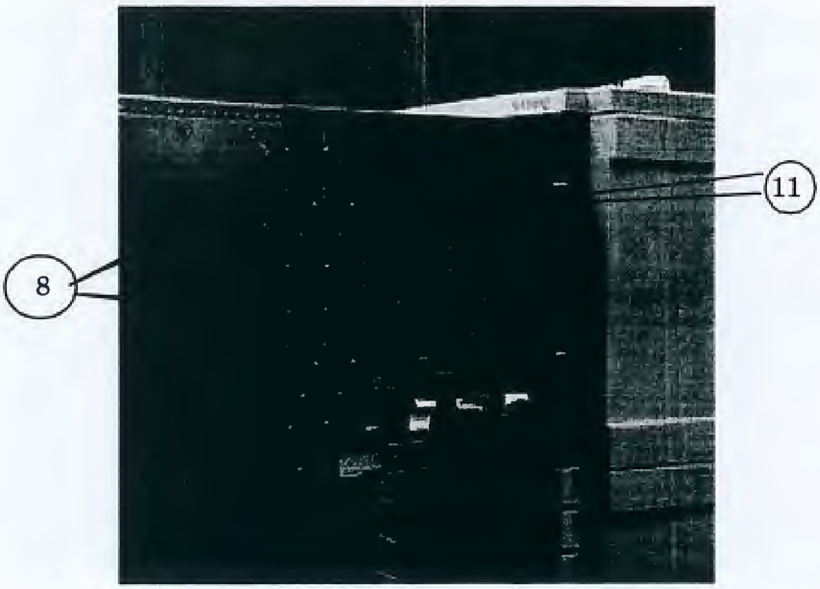

Figure A-3 System Board (Access via Top Service Bay)

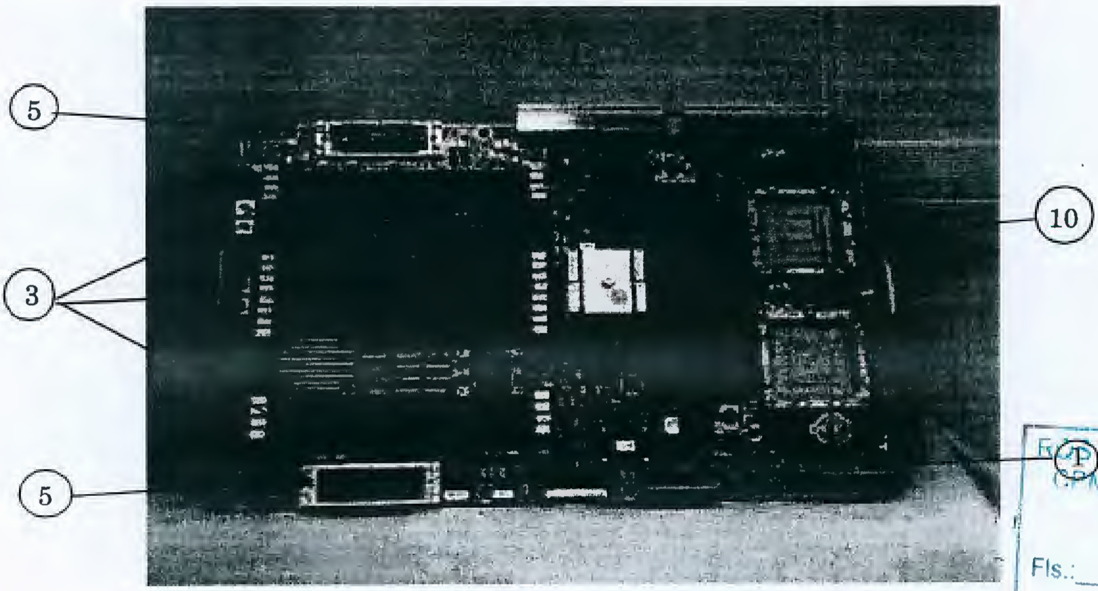

Figure A-4 **Server Front**

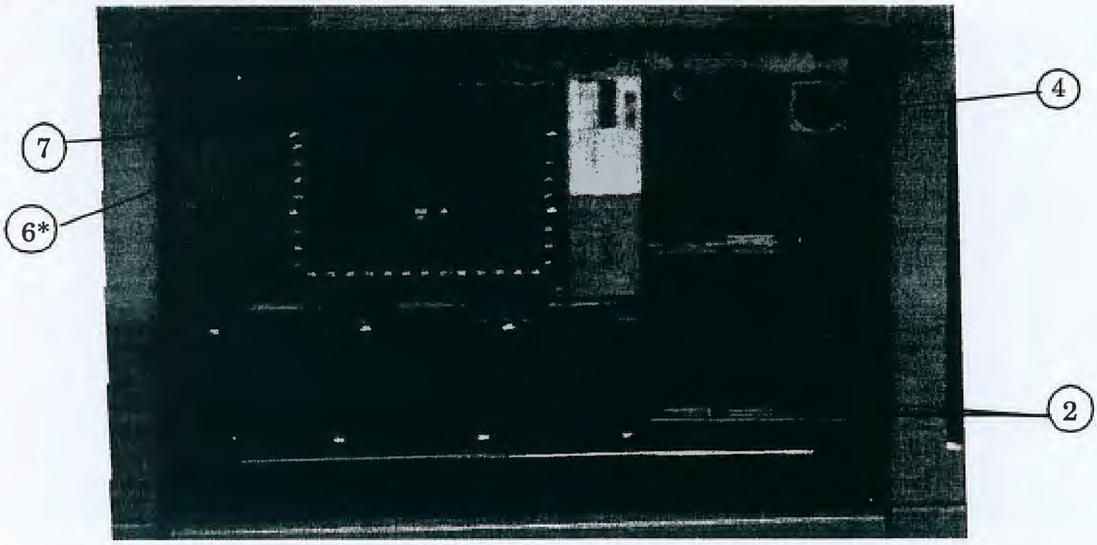

*The Display Board is accessed by removing the Chassis Fan screen and fan (item 7). The Display Board is located through a slot in the inside top of the Chassis Fan cavity.

Customer Replaceable Unit Part Numbers

The following tables list all Customer Replaceable Units (CRUs) for the rp54xx computer system. The following tables list both new and exchange part numbers.

Table A-1 Exchange CRUs

Product #	New Part #	Exch. Part #	Description	Loc #
A5191A/B A5576A/B A6144A/B A6797B	A5191-60010	A5191-69010	Platform Monitor Board	1
A5527A	0950-3471	A5527-69001	Power Supply	2
A6155A	A6155-60001	A6155-69001	Memory Carrier	NA
A6115A	A6115-60001	A6115-69001	1024 MB Memory DIMM	3
A5798A	A5798-60001	A5798-69001	512 MB Memory DIMM	
A5797A	A5797-60001	A5797-69001	256 MB Memory DIMM	
A5554A	A5554-60002	A5554-69002	128 MB Memory DIMM	
A5802A	A5802-67001	A5802-69001	9 GB HotPlug Disk Assembly	4
A5803A	A5803-67001	A5803-69001	18 GB HotPlug Disk Assembly	
A6110A	A6110-67001	A6110-69001	36 GB HotPlug Disk Assembly	
A5796A	A3639-60012	A3639-69012	PA-8500/8600 Processor Support Module	5
A6799A	0950-3908	A3639-69033	PA-8700 Processor Support Module	5
A6696A	A5191-60012	A5191-69012	Revision A GSP	NA
A6696B	A6144-60012	A6144-69012	Revision B GSP	NA

18.589

Table A-2 Non-Exchange CRUs

Product #	Part #	Description	Loc #
A5191A/B A5576A/B A6144A/B A6797B	A5191-60013	Display Board	6
	A5191-04002	Fan, Chassis (172m)	7
	A5191-04003	Fan, Card Cage (119m)	8
	A5191-67006	Fan, Power Converter	9
	0515-0664	Processor Cover Plate	10
	A5236-40024	Disk Filler Panel	NA
	A5191-00107	Power Supply Filler Panel	NA
	A3639-04024	PCI Separator/Extractor	11
A5576A/B A5191A/B A6144A	A5191-04008	Plastic Front Bezel, old style (split door)	NA
A5576A/B A5191A/B A6144A/B	A5191-04012[a]	Plastic Front Bezel, new style rp54xx (solid piece)	NA
A6797B	A5191-04013[a]	Plastic Front Bezel, rp54xx (solid piece)	NA
A5576A/B A5191A/B A6144A/B A6797B	A5191-70010	Bezel Hardware Kit	NA
A6696A	A5191-63001	W-Cable (beige color), use with revision A GSP (A5191-60012/69012) only	NA
A6696B	A6144-63001	M-Cable (black color), use with revision B GSP (A6144-60012/69012) only	NA

a. When replacing the split door bezel (A5191-04008) with a solid piece bezel, the Bezel Hardware Kit (A5191-70010) is required.

18.558 ✓

B System Specifications

18.587

Dimensions

Uncrating Space

rp54xx systems are shipped in boxes on a 34.75 in. x 26 in. (88.25 cm. x 66 cm.) pallet. The combined height of the packaged container and the pallets is 23.25 in. (59 cm.).

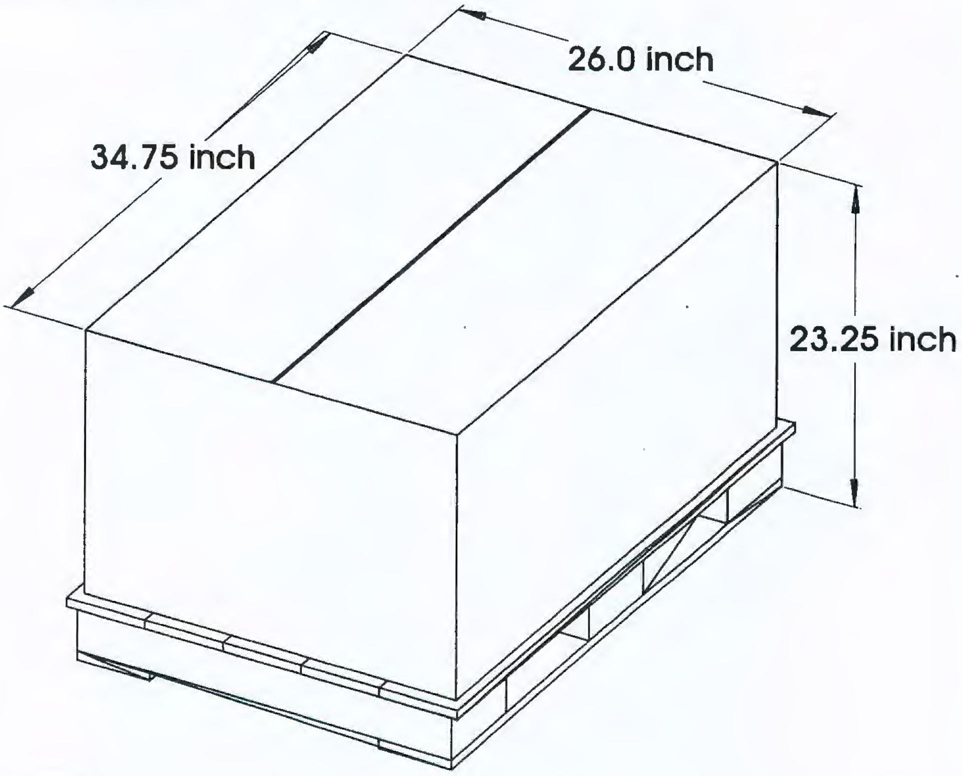

rhrr004

Allow a circular area approximately 5 ft. (2 m.) in diameter room for uncrating the system. Allow additional space for temporary storage of the shipping containers and packing materials.

Space Requirements

A minimum access area of 2 ft. (0.7 m.) in all directions is required for serviceability.

18586
D.

Computer Room Physical Space Requirements

Server

The A5537A, A5538A, or A5539A cabinets in which the server resides is 38.5" deep by 23.5" wide, with heights as follows:

- A5537A 1.2 meter rack - 49.5"
- A5538A 1.6 meter rack - 63.5"
- A5539A 2.0 meter rack - 77.5"

The cabinet anti-tip feet (required for safety) extend the depth an additional 10" in the front and 14" in the back. The minimum standalone physical space for the rp54xx server in a cabinet is shown in the following illustration:

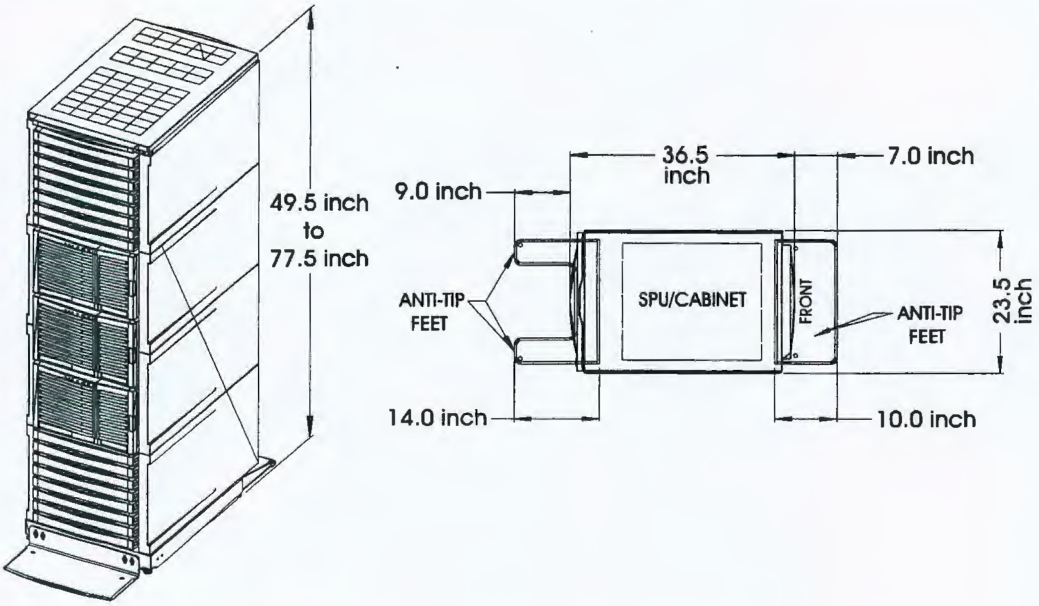

rhrr005

For service access, the server slides on rails 28" beyond the edge of the chassis in the front. An additional minimum of 3' of workspace on all sides is required for servicing the server.

Aisle Space

The minimum aisle space between rows of installed HP rp54xx server/cabinets is 3 feet, front and back, for airflow and serviceability.

Computer Room Unpacking Space

Specify Uncrating Space

rp54xx Systems are shipped in integrated cabinets on a 34.75" (88.25 cm.) x 26" (66 cm.) pallet. The combined height of the packaged container and the pallets are as follows:

- 1.1 meter rack - 60 inches
- 1.6 meter rack - 73 inches
- 2.0 meter rack - 87 inches

Allow room (a circular area approximately 12 feet (3.5 meters) in diameter) for uncrating the system and rolling the cabinet off the pallet on rails. Allow additional space for temporary storage of the shipping containers and packing materials..

18584
D-

Electrical

The HP rp54xx power system is comprised of one, two or three autoranging, 12A/10A/5A, 930W hot-swappable system power supplies, depending on how the system is configured. The maximum power needed by fully-configured server is 1200W. Each power supply requires a dedicated 20A branch circuit.

CAUTION HP does not recommend and does not support the use of "ferro-active" or "ferro-resonant" power correction in conjunction with the rp54xx server. These type of line conditioners represent an older technology that is not compatible with the most recent designs in active PFC power supplies such as those in the HP rp54xx servers. "Ferro-active" or "ferro-resonant" line conditioners may cause an increase in total harmonic distortion and may produce significant and unpredictable voltage regulation anomalies.

Office High Availability Requirements

Server-level Enhanced Power Availability

At the server level, enhanced power availability is achieved through the n+1 hot-swappable power supplies.

One power supply is required for a minimally configured rp54xx system operation and in order to allow the system to boot. If a second power supply is present, one of the two power supplies can fail without the system shutting down. Similarly, two power supplies are required for a more heavily-configured rp54xx server. If a third power supply is present, one of the two power supplies can fail without the system shutting down.

If a third (redundant) power supply is present, all three power supplies become hot-swappable. Any one of the power supplies can fail without affecting system performance, and can be replaced while the system is on-line. Single point of failure is reduced to the local wall circuit or PowerTrust UPS to which the power supply power cords are connected.

Power Protection

Power protection is provided through the use of HP PowerTrust UPSs (Uninterruptible Power Supplies). The only supported models for rp54xx systems are the 3.0kVA and 5.5kVA models. Recommendations for other manufacturers and models are not yet determined.

HP rp54xx server power supplies may be plugged directly into the customer's site UPS. However, customers are advised against plugging the power supply into an HP PowerTrust UPS and then connecting that UPS to the site UPS.

18583 D.

Modular PDUs

NOTE The acronym PDU (Power Distribution Unit) in this document refers to the power strips attached to the HP rp54xx cabinet.

Three modular PDUs are available for use with the rp54xx product:

- 20A modular PDU, HP product number E7674A. This PDU has seven C13 outlets and one C19 outlet. Note that if this PDU is used to power the rp54xx server, any other peripherals must have their power supplied by additional PDUs in the cabinet. The PDU power cord (HP part number 8120-6903) has an L6-20P plug which must be plugged into an L6-20R wall or floor receptacle.

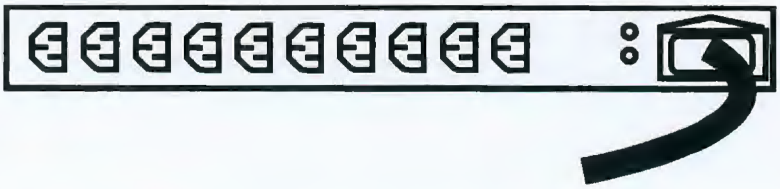

- 30A modular PDU, HP product number E7681A (North America) or E7682A (International). This PDU has eight C13 outlets and two C19 outlets, split across two 20A branch circuits (max. 30A available). The PDU power cord has an L6-30 plug.

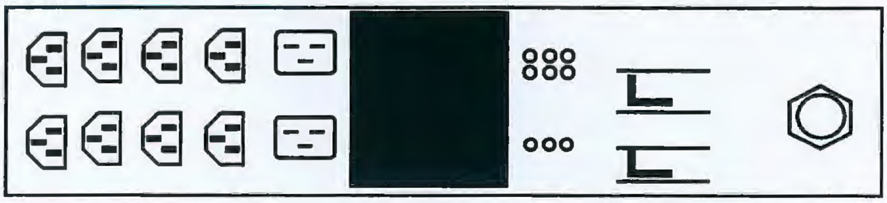

- 60A modular PDU, HP product number E7683A (North America) or E7667A (International). FINAL DESIGN ON THIS PRODUCT IS NOT YET AVAILABLE.

System Power Requirements

Table B-1 Power Requirements

Requirements	Value	Comments
Nominal input voltage	100-240 VAC	
Frequency range (minimum - maximum)	50 - 60 HZ	
Number of phases	1	
Theoretical maximum current	12.0 A	Per line cord
Maximum inrush current	69.0 A	Per line cord
Ground leakage current (mA)	< 3.6 mA	Per line cord

18582

LAN and Telephone

rp54xx servers provide remote console access via the secure web console, which may be connected to the RS-232 serial port or 10Base-T RJ45 LAN port on the GSP (Guardian Service Processor) card in I/O slot #2. The same connectors may also be used to connect directly to a hard console.

NOTE The RJ45 LAN port should be used ONLY for remote console access and not as a production LAN port.

rp54xx servers may require as many as three unique IP addresses:

- The rp54xx server requires its own IP address.
- The Guardian Service Processor, as a separate network device, has to have its own IP address, gateway, and subnet mask configured at the site in order to be separately addressed. The service processor does not support DHCP, so the IP address must be assigned out of a separate pool from any that are assigned dynamically.
- The remote web console, if used, requires its own IP address.

NOTE Check with your local telephone company to be thatsure the telephone service at the site can accommodate modem/data quality transmission.

18581
D.

Acoustic Safety Standards

The acoustic specifications for the rp54xx server are as follows:

Sound power	7.0 Bels LwA maximum at >31° C
Sound pressure	60 dB maximum at > 31° C
	No prominent tones

Altitude Operation Standards

Maximum Altitude	
Operational	3000 meters above sea level
Non-operational	4572 meters above sea level

Effects of Altitude

Some old models of tape drives, including those supplied by Hewlett-Packard, have vacuum column transport mechanisms that are affected by atmospheric pressure. Adjustments to these mechanisms may be required to compensate for the lower atmospheric pressure at higher altitudes.

Temperature and Humidity Operating Standards

The following table lists the temperature and humidity specifications for rp54xx servers.

Temperature and Humidity Specifications

Parameter	Operational Limits	Recommended Operating Range	Maximum Rate of Change (per hour)	Non-Operating Ranges
Temperature[a]	5°C to 35°C (41°F to 95°F)[b]	20°C to 25°C (68°F to 77°F)	10°C (50°F) per hour (With tape media) 20°C (68°F) per hour (Without tape media)	-40°C to +70°C (-40°F to +158°F)
Operating Humidity	15 to 80% Relative humidity (Non-condensing) at 35°C (95°F)	40 to 60% Relative humidity (Non-condensing)	30% Per hour Relative humidity (Non-condensing)	90% Relative humidity (Non-condensing) at 65°C (149°F)

a. The rp54xx has been designed to operate within the above specific temperature and relative humidity operational limits. In general, operating any electronic equipment within the recommended ranges of temperature and humidity will produce optimal reliability.

b. Temperature ranges stated above are at sea level. Maximum operating temperature is derated (reduced) by 2°C for each 1000 meters above sea level up to a maximum of 3000 meters.

NOTE Operating ranges refer to the ambient air temperature and humidity measured at the cabinet cooling air intake vents.

Thermal Protection Features

If the cabinet temperature approaches 35°C, thermal protection will be invoked. At 35°C +/- 2°C a warning message will be displayed on the system console. At 40°C +/- 2°C an ungraceful shutdown will occur.

18578
D

C General Site Preparation Guidelines

The following information provides general principles and practices to consider before the installation or operation of an hp server.

18577
D

Electrical Factors

Proper design and installation of a power distribution system for an hp server requires specialized skills. Those responsible for this task must have a thorough knowledge and understanding of appropriate electrical codes and the limitations of the power systems for computer and data processing equipment.

In general, a well-designed power distribution system exceeds the requirements of most electrical codes. A good design, when coupled with proper installation practices, produces the most trouble-free operation.

A detailed discussion of power distribution system design and installation is beyond the scope of this information. However, electrical factors relating to power distribution system design and installation must be considered during the site preparation process.

The electrical factors discussed in this section are:

- Computer room safety
- Power Consumption
- Electrical load requirements (circuit breaker sizing)
- Power quality
- Distribution hardware
- System installation guidelines

Computer Room Safety

Inside the computer room, fire protection and adequate lighting (for equipment servicing) are important safety considerations. Federal and local safety codes govern computer installations.

Fire Protection

The National Fire Protection Association's Standard for the Protection of Electronic Computer Data Processing Equipment, NFPA 75, contains information on safety monitoring equipment for computer rooms.

Most computer room installations are equipped with the following fire protection devices:

- Smoke detectors
- Fire and temperature alarms
- Fire extinguishing system

Additional safety devices are:

- Circuit breakers
- An emergency power cutoff switch
- Devices specific to the geographic location i.e., earthquake protection

Lighting Requirements for Equipment Servicing

Adequate lighting and utility outlets in a computer room reduce the possibility of accidents during equipment servicing. Safer servicing is also more efficient and, therefore, less costly.

For example, adequate lighting reduces the chances of connector damage when cables are installed or removed.

18576

The minimum recommended illumination level is 70 foot-candles (756 lumens per square meter) when the light level is measured at 30 inches (76.2 cm) above the floor.

Power Consumption

When determining power requirements, you must consider any peripheral equipment that will be installed during initial installation or as a later update. Refer to the applicable documentation for such devices to determine the power required to support these devices.

Electrical Load Requirements (Circuit Breaker Sizing)

NOTE Local authority has jurisdiction [LAHJ] and should make the final decision regarding adherence to country- specific electrical codes and guidelines.

It is good practice to derate power distribution systems for one or more of the following reasons:

- To avoid nuisance tripping from load shifts or power transients, circuit protection devices should never be run above 80% of their root-mean-square (RMS) current ratings.
- Safety agencies derate most power connectors to 80% of their RMS current ratings.

Power Quality

The hp server is designed to operate over a wide range of voltages and frequencies. The server is tested and shown to comply with EMC Specification EN50082. However, damage can occur if these ranges are exceeded. Severe electrical disturbances can exceed the design specifications of the equipment.

Sources of Voltage Fluctuations

Voltage fluctuations, sometimes called glitches, affect the quality of electrical power. Common sources of these disturbances are:

- Fluctuations occurring within the facility's distribution system
- Utility service low-voltage conditions (such as sags or brownouts)
- Wide and rapid variations in input voltage levels
- Wide and rapid variations in input power frequency
- Electrical storms
- Large inductive sources (such as motors and welders)
- Faults in the distribution system wiring (such as loose connections)
- Microwave, radar, radio, or cell phone transmissions

Power System Protection

The hp server can be protected from the sources of many of these electrical disturbances by using:

- A dedicated power distribution system
- Power conditioning equipment
- Over- and under-voltage detection and protection circuits
- Screening to cancel out the effects of undesirable transmissions
- Lightning arresters on power cables to protect equipment against electrical storms

Precautions have been taken during power distribution system design to provide immunity to power outages of less than one cycle. However, testing cannot conclusively rule out loss of service. Therefore, adherence to the following guidelines provides the best possible performance of power distribution systems for hp computer equipment:

- Dedicated power source—Isolates an hp server power distribution system from other circuits in the facility.
- Missing-phase and low-voltage detectors—Shuts equipment down automatically when a severe power disruption occurs. For peripheral equipment, these devices are recommended but optional.
- Online uninterruptible power supply (UPS)—Keeps input voltage to devices constant and should be considered if outages of one-half cycle or more are common. Refer to qualified contractors or consultants for each situation.

Distribution Hardware

This section describes wire selection and the types of raceways (electrical conduits) used in the distribution system.

Wire Selection

Use copper conductors instead of aluminum, as aluminum's coefficient of expansion differs significantly from that of other metals used in power hardware. Because of this difference, aluminum conductors can cause connector hardware to work loose, overheat, and fail.

Raceway Systems (electrical conduits) [LAHJ]

Raceways (electrical conduits) form part of the protective ground path for personnel and equipment. Raceways protect the wiring from accidental damage and also provide a heatsink for the wires.

Any of the following types may be used:

- Electrical metallic tubing (EMT) thin-wall tubing
- Rigid (metal) conduit
- Liquidtight with RFI shield grounded (most commonly used under raised floors)

Building Distribution

All building feeders and branch circuitry should be in rigid metallic conduit with proper connectors (to provide ground continuity) Conduit that is exposed and subject to damage should be constructed of rigid galvanized steel.

18574
Diogo

Grounding Systems

An hp server requires two methods of grounding:

- Power distribution safety grounding
- High frequency intercabinet grounding

Power Distribution Safety Grounding [LAHJ]

The power distribution safety grounding system consists of connecting various points in the power distribution system to earth ground using green (green/yellow) wire ground conductors. Having these ground connections tied to metal chassis parts that may be touched by computer room personnel protects them against shock hazard from current leakage and fault conditions.

Power distribution systems consist of several parts. Hewlett-Packard recommends that these parts be solidly interconnected to provide an equipotential ground to all points.

Main Building Electrical Ground The main electrical service entrance equipment should have an earth ground connection, as required by applicable codes. Connections such as a grounding rod, building steel, or a conductive type cold water service pipe provide an earth ground.

Electrical Conduit Ground All electrical conduits should be made of rigid metallic conduit that is securely connected together or bonded to panels and electrical boxes, so as to provide a continuous grounding system.

Power Panel Ground Each power panel should be grounded to the electrical service entrance with green (green/yellow) wire ground conductors. The green (green/yellow) wire ground conductors should be sized per applicable codes (based on circuit over current device ratings).

NOTE The green wire ground conductor mentioned above may be a black wire marked with green tape. [LAHJ]

Computer Safety Ground Ground all computer equipment with the green (green/yellow) wire included in the branch circuitry. The green (green/yellow) wire ground conductors should be connected to the appropriate power panel and should be sized per applicable codes (based on circuit over current device ratings).

Cabinet Performance Grounding (High Frequency Ground)

Signal interconnects between system cabinets require high frequency ground return paths. Connect all cabinets to site ground.

NOTE In some cases power distribution system green (green/yellow) wire ground conductors are too long and inductive to provide adequate high frequency ground return paths. Therefore, a ground strap (customer-supplied) should be used for connecting the system cabinet to the site grounding grid (customer-supplied). When connecting this ground, ensure that the raised floor is properly grounded for high frequency.

Power panels located in close proximity to the computer equipment should also be connected to the site grounding grid. Methods of providing a sufficiently high frequency ground grid are described in the next sections.

13573
Diego.

Raised Floor "High Frequency Noise" Grounding

If a raised floor system is used, install a complete signal grounding grid for maintaining equal potential over a broad band of frequencies. The grounding grid should be connected to the equipment cabinet and electrical service entrance ground at multiple connection points using a minimum #6 AWG (16mm2) wire ground conductor. The following figure illustrates a metallic strip grounding system.

Hewlett-Packard recommends the following approaches:

- Excellent—Add a grounding grid to the subfloor. The grounding grid should be made of copper strips mounted to the subfloor. The strips should be 0.032 in. (0.08 cm) thick and a minimum of 3.0 in. (8.0 cm) wide.

  Connect each pedestal to four strips using 1/4 in. (6.0 mm) bolts tightened to the manufacturer's torque recommendation.
- Better - A grounded #6 AWG minimum copper wire grid mechanically clamped to floor pedestals and properly bonded to the building/site ground.
- Good—Use the raised floor structure as a ground grid. In this case, the floor must be designed as a ground grid with bolted down stringers and corrosion resistive plating (to provide low resistance and attachment points for connection to service entrance ground and hp computer equipment). The use of conductive floor tiles with this style of grid further enhances ground performance.

Figure C-1 Raised Floor Metal Strip Ground System

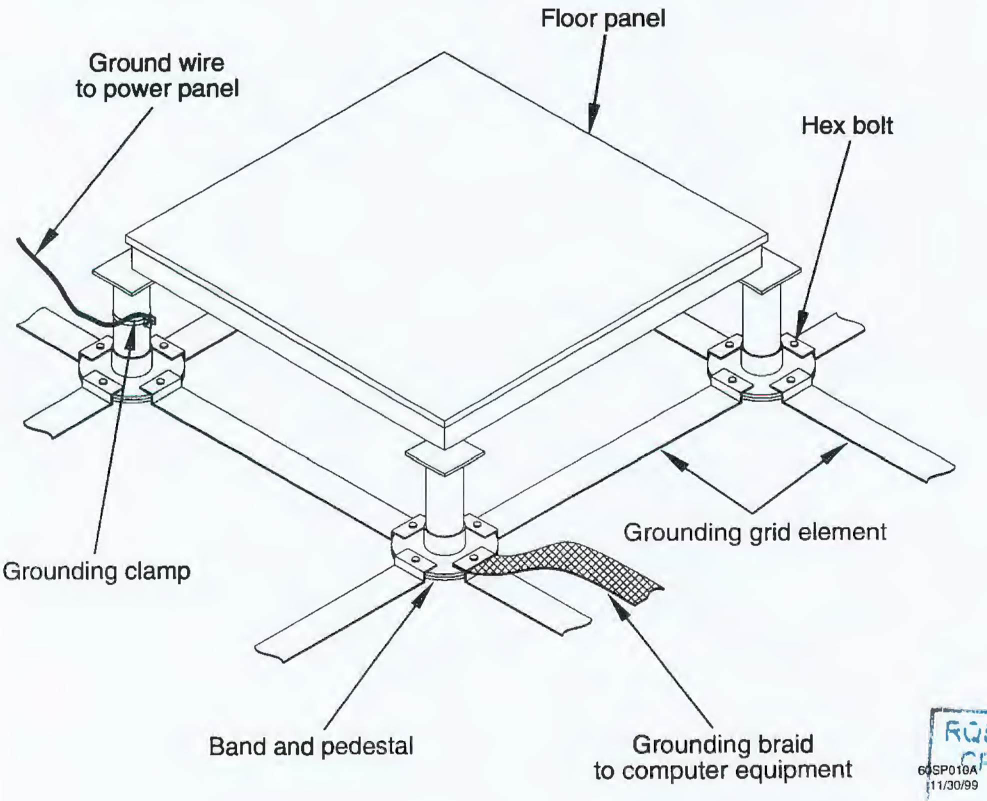

1857-2 Diego.

Equipment Grounding Implementation Details

Connect all Hewlett-Packard equipment cabinets to the site ground grid as follows:

Step 1. Attach one end of each ground strap to the applicable cabinet ground lug.

Step 2. Attach the other end to the nearest pedestal base (raised floor) or cable trough ground point (nonraised floor).

Step 3. Check that the braid contact on each end of the ground strap consists of a terminal and connection hardware (a 1/4-in. (6.0-mm) bolt, nuts, and washers).

Step 4. Check that the braid contact connection points are free of paint or other insulating material and treated with a contact enhancement compound (similar to Burndy Penetrox).

System Installation Guidelines

This section contains information about installation practices. Some common pitfalls are highlighted. Both power cable and data communications cable installations are discussed.

NOTE In domestic installations, the proper receptacles should be installed prior to the arrival of Hewlett-Packard equipment. Refer to the appropriate installation guide for installation procedures.

Wiring Connections

Expansion and contraction rates vary among different metals. Therefore, the integrity of an electrical connection depends on the restraining force applied. Connections that are too tight compress or deform the hardware and causes it to weaken. This usually leads to high impedance preventing circuit breakers from tripping when needed or can contribute to a buildup of high frequency noise.

CAUTION Connections that are too loose or too tight can have a high impedance that cause serious problems, such as erratic equipment operation. A high impedance connection overheats and sometimes causes fire or high temperatures that can destroy hard-to-replace components such as distribution panels or system bus bars.

Wiring connections must be properly torqued. Many equipment manufacturers specify the proper connection torque values for their hardware.

Ground connections must only be made on a conductive, nonpainted surface. When equipment vibration is present, lock washers must be used on all connections to prevent connection hardware from working loose.

Data Communications Cables

Power transformers create high-energy fields in the form of electromagnetic interference (EMI). Heavy foot traffic can create electrostatic discharge (ESD) that can damage electronic components. Route data communications cables away from these areas. Use shielded data communications cables that meet approved industry standards to reduce the effects of external fields.

Environmental Elements

The following environmental elements can affect an hp server installation:

- Computer room preparation
- Cooling requirements
- Humidity level
- Air conditioning ducts
- Dust and pollution control
- Electrostatic discharge (ESD) prevention
- Acoustics (noise reduction)

Computer Room Preparation

The following guidelines are recommended when preparing a computer room for an hp server:

- Locate the computer room away from the exterior walls of the building to avoid the heat gain from windows and exterior wall surfaces.
- When exterior windows are unavoidable, use windows that are double or triple glazed and shaded to prevent direct sunlight from entering the computer room.
- Maintain the computer room at a positive pressure relative to surrounding spaces.
- Use a vapor barrier installed around the entire computer room envelope to restrain moisture migration.
- Caulk and vapor seal all pipes and cables that penetrate the envelope.
- Use at least a 12-inch raised floor system for minimum favorable room air distribution system (underfloor distribution).
- Ensure a minimum clearance of 12 inches between the top of the hp server cabinet and the ceiling to allow for return air flow and ensure that all ceiling tiles are in place.
- Allow 18 inches (or local code minimum clearance) from the top of the hp server cabinet to the fire sprinkler heads.

Cooling Requirements

Air conditioning equipment requirements and recommendations are described in the following sections.

Basic Air Conditioning Equipment Requirements

The cooling capacity of the installed air conditioning equipment for the computer room should be sufficient to offset the computer equipment dissipation loads, as well as any space envelope heat gain. This equipment should include:

- Air filtration
- Cooling or dehumidification
- Humidification
- Reheating

18570

- Air distribution
- System controls adequate to maintain the computer room within specified operating ranges.

Lighting and personnel must also be included. For example, a person dissipates about 450 BTUs per hour while performing a typical computer room task.

At altitudes above 10,000 feet (3048 m), the lower air density reduces the cooling capability of air conditioning systems. If your facility is located above this altitude, the recommended temperature ranges may need to be modified. For each 1000 feet (305 m) increase in altitude above 10,000 feet (up to a maximum of 15,000 feet), subtract 1.5° F (0.83° C) from the upper limit of the temperature range.

Air Conditioning System Guidelines

The following guidelines are recommended when designing an air conditioning system and selecting the necessary equipment:

- The air conditioning system that serves the computer room should be capable of operating 24 hours a day, 365 days a year. It should also be independent of other systems in the building.
- Consider the long-term value of hp server availability, redundant air conditioning equipment or capacity.
- The system should be capable of handling any future hp server expansion.
- Air conditioning equipment air filters should have a minimum rating of 45% (based on “ASHRAE Standard 52-76, Dust Spot Efficiency Test”).
- Introduce only enough outside air into the system to meet building code requirements (for human occupancy) and to maintain a positive air pressure in the computer room.

Air Conditioning System Types

The following three air conditioning system types are listed in order of preference:

- Complete self-contained package unit(s) with remote condenser(s). These systems are available with up or down discharge and are usually located in the computer room.
- Chilled water package unit with remote chilled water plant. These systems are available with up or down discharge and are usually located in the computer room.
- Central station air handling units with remote refrigeration equipment. These systems are usually located outside the computer room.

Basic Air Distribution Systems

A basic air distribution system includes supply air and return air.

An air distribution system should be zoned to deliver an adequate amount of supply air to the cooling air intake vents of the hp server equipment cabinets. Supply air temperature should be maintained within the following parameters:

- Ceiling supply system—From 55° F (12.8° C) to 60° F (15.6° C)
- Floor supply system—At least 60° F (15.6° C)

If a ceiling plenum return air system or a ducted ceiling return air system is used, the return air grill(s) in the ceiling should be above the exhaust area or the exhaust row.

The following three types of air distribution system are listed in order of recommendation:

- Underfloor air distribution system—Downflow air conditioning equipment located on the raised floor of the computer room uses the cavity beneath the raised floor as plenum for the supply air.

28569

Perforated floor panels (available from the raised floor manufacturer) should be located around the front of the system cabinets. Supply air emitted though the perforated floor panels is then available near the cooling air intake vents of the hp server cabinets.

- Ceiling plenum air distribution system—Supply air is ducted into the ceiling plenum from upflow air conditioning equipment located in the computer room or from an air handling unit (remote).

  The ceiling construction should resist air leakage. Place perforated ceiling panels (with down discharge air flow characteristics) around the front of the system cabinets. The supply air emitted downward from the perforated ceiling panels is then available near the cooling air intake vents of the hp server cabinets.

  Return air should be ducted back to the air conditioning equipment though the return air duct above the ceiling.

- Above ceiling ducted air distribution system—Supply air is ducted into a ceiling diffuser system from upflow air conditioning equipment located in the computer room or from an air handling unit (remote).

  Adjust the supply air diffuser system grilles to direct the cooling air downward around the front of the hp server cabinets. The supply air is then available near the cooling air intake vents of the hp server cabinets.

Air Conditioning System Installation

All air conditioning equipment, materials, and installation must comply with any applicable construction codes. Installation of the various components of the air conditioning system must also conform to the air conditioning equipment manufacturer's recommendations.

Air Conditioning Ducts

Use separate computer room air conditioning duct work. If it is not separate from the rest of the building, it might be difficult to control cooling and air pressure levels. Duct work seals are important for maintaining a balanced air conditioning system and high static air pressure. Adequate cooling capacity means little if the direction and rate of air flow cannot be controlled because of poor duct sealing. Also, the ducts should not be exposed to warm air, or humidity levels may increase.

Humidity Level

Maintain proper humidity levels at 40 to 60% RH. High humidity causes galvanic actions to occur between some dissimilar metals. This eventually causes a high resistance between connections, leading to equipment failures. High humidity can also have an adverse affect on some magnetic tapes and paper media.

CAUTION Low humidity contributes to undesirably high levels of electrostatic charges. This increases the electrostatic discharge (ESD) voltage potential. ESD can cause component damage during servicing operations. Paper feed problems on high-speed printers are usually encountered in low-humidity environments.

Low humidity levels are often the result of the facility heating system and occur during the cold season. Most heating systems cause air to have a low humidity level, unless the system has a built-in humidifier.

Dust and Pollution Control

Computer equipment can be adversely affected by dust and microscopic particles in the site environment.

Specifically, disk drives, tape drives, and some other mechanical devices can have bearing failures resulting from airborne abrasive particles. Dust may also blanket electronic components like printed circuit boards causing premature failure due to excess heat and/or humidity build up on the boards. Other failures to power supplies and other electronic components can be caused by metallically conductive particles, including zinc whiskers. These metallic particles are conductive and can short circuit electronic components. Use every effort to ensure that the environment is as dust and particulate free as possible. See the following heading titled "Metallic Particulate Contamination" for additional details.

Smaller particles can pass though some filters and over a period of time, cause problems in mechanical parts. Small dust particles can be prevented from entering the computer room by maintaining the air conditioning system at a high static air pressure level.

Other sources of dust, metallic, conductive, abrasive, and/or microscopic particles can be present. Some sources of these particulates are:

- Subfloor shedding
- Raised floor shedding
- Ceiling tile shedding

These particulates are not always visible to the naked eye. A good check to determine their possible presence is to check the underside of the tiles. The tile should be shiny, galvanized, and free from rust.

The computer room should be kept clean. The following guidelines are recommended:

- Smoking—Establish a no-smoking policy. Cigarette smoke particles are eight times larger than the clearance between disk drive read/write heads and the disk surface.
- Printer—Locate printers and paper products in a separate room to eliminate paper particulate problems.
- Eating or drinking—Establish a no-eating or drinking policy. Spilled liquids can cause short circuits in equipment such as keyboards.
- Tile floors—Use a dust-absorbent cloth mop rather than a dry mop to clean tile floors.

Special precautions are necessary if the computer room is near a source of air pollution. Some air pollutants, especially hydrogen sulfide (H2S), are not only unpleasant but corrosive as well. Hydrogen sulfide damages wiring and delicate sound equipment. The use of activated charcoal filters reduces this form of air pollution.

Metallic Particulate Contamination

Metallic particulates can be especially harmful around electronic equipment. This type of contamination may enter the data center environment from a variety of sources, including but not limited to raised floor tiles, worn air conditioning parts, heating ducts, rotor brushes in vacuum cleaners or printer component wear. Because metallic particulates conduct electricity, they have an increased potential for creating short circuits in electronic equipment. This problem is exaggerated by the increasingly dense circuitry of electronic equipment.

Over time, very fine whiskers of pure metal can form on electroplated zinc, cadmium, or tin surfaces. If these whiskers are disturbed, they may break off and become airborne, possibly causing failures or operational interruptions. For over 50 years, the electronics industry has been aware of the relatively rare but possible threat posed by metallic particulate contamination. During recent years, a growing concern has developed in computer rooms where these conductive contaminants are formed on the bottom of some raised floor tiles.

Although this problem is relatively rare, it may be an issue within your computer room. Since metallic contamination can cause permanent or intermittent failures on your electronic equipment, Hewlett-Packard strongly recommends that your site be evaluated for metallic particulate contamination before installation of electronic equipment.

18567
D.

Electrostatic Discharge (ESD) Prevention

Static charges (voltage levels) occur when objects are separated or rubbed together. The voltage level of a static charge is determined by the following factors:

- Types of materials
- Relative humidity
- Rate of change or separation

Table C-1 Effect of Humidity on ESD Charge Levels

Personnel Activity[a]	Humidity[b] and Charge Levels (voltages)[c]			
	26%	**32%**	**40%**	**50%**
Person walking across a linoleum floor	**6,150 V**	**5,750 V**	**4,625 V**	**3,700 V**
Person walking across a carpeted floor	18,450 V	17,250 V	13,875 V	11,100 V
Person getting up from a plastic chair	24,600 V	23,000 V	18,500 V	14,800 V

[a] Source: B.A. Unger, Electrostatic Discharge Failures of Semiconductor Devices (Bell Laboratories, 1981)

[b] For the same relative humidity level, a high rate of airflow produces higher static charges than a low airflow rate.

[c] Some data in this table has been extrapolated.

Static Protection Measures

Follow these precautions to minimize possible ESD-induced failures in the computer room:

- Maintain recommended humidity level and airflow rates in the computer room.
- Install conductive flooring (conductive adhesive must be used when laying tiles).
- Use conductive wax if waxed floors are necessary.
- Ensure that all equipment and flooring are properly grounded and are at the same ground potential.
- Use conductive tables and chairs.
- Use a grounded wrist strap (or other grounding method) when handling circuit boards.
- Store spare electronic modules in antistatic containers.

Acoustics

Computer equipment and air conditioning blowers cause computer rooms to be noisy. Ambient noise level in a computer room can be reduced as follows:

- Dropped ceiling—Cover with a commercial grade of fire-resistant, acoustic rated, fiberglass ceiling tile.
- Sound deadening—Cover the walls with curtains or other sound deadening material.
- Removable partitions—Use foam rubber models for most effectiveness.

18666
D.

Facility Characteristics

This section contains information about facility characteristics that must be considered for the installation or operation of an hp server. Facility characteristics are:

- Floor loading
- Windows
- Altitude effects

Floor Loading

The computer room floor must be able to support the total weight of the installed hp server as well as the weight of the individual cabinets as they are moved into position.

Floor loading is usually not an issue in nonraised floor installations. The information presented in this section is directed toward raised floor installations.

NOTE An appropriate floor system consultant should verify any floor system under consideration for an hp server installation.

Raised Floor Loading

Raised floor loading is a function of the manufacturer's load specification and the positioning of the equipment relative to the raised floor grid. While Hewlett-Packard cannot assume responsibility for determining the suitability of a particular raised floor system, it does provide information and illustrations for the customer or local agencies to determine installation requirements.

The following guidelines are recommended:

- Because many raised floor systems do not have grid stringers between floor stands, the lateral support for the floor stands depends on adjacent panels being in place. To avoid compromising this type of floor system while gaining under floor access, remove only one floor panel at a time.
- Larger floor grids (bigger panels) are generally rated for lighter loads.

CAUTION Do not position or install any equipment cabinets on the raised floor system until you have carefully examined it to verify that it is adequate to support the appropriate installation.

Floor Loading Terms

Table C-2 Floor Loading Term Definitions

Term	Definition
Dead load	The weight of the raised panel floor system, including the understructure. Expressed in lb/ft2 (kg/m2).
Live load	The load that the floor system can safely support. Expressed in lb/ft2 (kg/m2).

Table C-2 Floor Loading Term Definitions (Continued)

Term	Definition
Concentrated load	The load that a floor panel can support on a 1-in2 (6.45 cm2) area at the panel's weakest point (typically the center of the panel), without the surface of the panel deflecting more than a predetermined amount.
Ultimate load	The maximum load (per floor panel) that the floor system can support without failure. Failure expressed by floor panel(s) breaking or bending. Ultimate load is usually stated as load per floor panel.
Rolling load	The load a floor panel can support (without failure) when a wheel of specified diameter and width is rolled across the panel.
Average floor load	Computed by dividing total equipment weight by the area of its footprint. This value is expressed in lb/ft2 (kg/m2).

Average Floor Loading

The average floor load value is not appropriate for addressing raised floor ratings at the floor grid spacing level. However, it is useful for determining floor loading at the building level, such as the area of solid floor or span of raised floor tiles covered by the hp server footprint.

Typical Raised Floor Site

This section contains an example of a computer room raised floor system that is satisfactory for the installation of an hp server.

Based on specific information provided by Hewlett-Packard, Tate Access Floors has approved its Series 800 all-steel access floor with bolt-together stringers and 24 in. (61.0 cm) by 24 in. (61.0 cm) floor panels.

In the event that the flooring is being replaced or a new floor is being installed, Tate Access Floors recommends its Series 1250 all-steel access floor with bolt-together stringers and 24 in. (61.0 cm) by 24 in. (61.0 cm) floor panels be used to support the hp installation.

NOTE If the specific floor being evaluated or considered is other than a Tate Series 800 floor, the specific floor mar ufacturer must be contacted to evaluate the floor being used.

The following table lists specifications for the Tate Access Floors Series 800 raised floor system.

Table C-3 Typical Raised Floor Specifications

Item[a]	Rating
Dead load	7 lb/ft 2 (34.2 kg/m2)
Live load	313 lb/ft 2 (1528.3 kg/m2)
Concentrated load[b]	1250 lb (567 kg)

Table C-3 Typical Raised Floor Specifications (Continued)

Item[a]	Rating
Ultimate load	4000 lb (1814 kg) per panel
Rolling load	400 lb (181 kg)
Average floor load	500 lb (227 kg)

a. From Table C-2 on page 203

b. With 0.08 in (0.2 cm) of span maximum deflection

Windows

Avoid housing computers in a room with windows. Sunlight entering a computer room may cause problems. Magnetic tape storage media is damaged if exposed to direct sunlight. Also, the heat generated by sunlight places an additional load on the cooling system.

18563
D

Space Requirements

This section contains information about space requirements for an hp server. This data should be used as the basic guideline for space plan developments. Other factors, such as airflow, lighting, and equipment space requirements must also be considered.

Delivery Space Requirements

There should be enough clearance to move equipment safely from the receiving area to the computer room. Permanent obstructions, such as pillars or narrow doorways, can cause equipment damage.

Delivery plans should include the possible removal of walls or doors.

Operational Space Requirements

Other factors must be considered along with the basic equipment dimensions. Reduced airflow around equipment causes overheating, which can lead to equipment failure. Therefore, the location and orientation of air conditioning ducts, as well as airflow direction, are important. Obstructions to equipment intake or exhaust airflow must be eliminated.

The locations of lighting fixtures and utility outlets affect servicing operations. Plan equipment layout to take advantage of lighting and utility outlets. Do not forget to include clearance for opening and closing equipment doors.

Clearance around the cabinets must be provided for proper cooling airflow through the equipment.

If other equipment is located so that it exhausts heated air near the cooling air intakes of the hp server cabinets, larger space requirements are needed to keep ambient air intake to the hp server cabinets within the specified temperature and humidity ranges.

Space planning should also include the possible addition of equipment or other changes in space requirements. Equipment layout plans should also include provisions for the following:

- Channels or fixtures used for routing data cables and power cables
- Access to air conditioning ducts, filters, lighting, and electrical power hardware
- Power conditioning equipment
- Cabinets for cleaning materials
- Maintenance area and spare parts

Floor Plan Grid

A floor plan grid is used to plan the location of equipment in the computer room. In addition to its use for planning, the floor plan grid should also be used when planning the locations of the following items:

- Air conditioning vents
- Lighting fixtures
- Utility outlets
- Doors
- Access areas for power wiring, air conditioning filters and equipment cable routing.

18562
D.

Typical Installation Schedule

The following schedule lists the sequence of events for a typical system installation:

- 60 days before installation
  - — Floor plan design completed and mailed to Hewlett-Packard (if required to be an HP task)
- 30 days before installation
  - — Primary power and air conditioning installation completed
  - — Telephone and data cables installed
  - — Fire protection equipment installed
  - — Major facility changes completed
  - — Special delivery requirements defined
  - — Site inspection survey completed
  - — Delivery survey completed
  - — A signed copy of the site inspection and delivery survey mailed to Hewlett-Packard
  - — Site inspection and predelivery coordination meeting arranged with a Hewlett-Packard representative to review the inspection checklist and arrange an installation schedule.
- 7 days before installation
  - — Final check made with an Hewlett-Packard site preparation specialist to resolve any last minute problems

NOTE Not all installations follow a schedule like the one noted above. Sometimes, an hp server is purchased through another vendor which can preclude a rigid schedule. Other conditions could also prevent following this schedule. For those situations, consider a milestone schedule.

- Site Preparation - schedule with the customer as soon as possible after the order is placed.
- Site Verification - schedule with the customer a minimum of one to two days before the hp server is scheduled to be installed.

18561
D

Site Inspection

Table C-4 Customer and Hewlett-Packard Information

Customer Information	
Name:	Phone No:
Street Address:	City or Town:
State or Province:	Country
Zip or postal code:	
Primary customer contact:	Phone No.:
Secondary customer contact:	Phone No.:
Traffic coordinator:	Phone No.:
Hewlett-Packard information	
Sales representative	Order No:
Representative making survey	Date:
Scheduled delivery date	

18560
D

Table C-5 Site Inspection Checklist

Please check either Yes or No. If No, include comment# or date				**Comment or Date**
Computer room				
No.	Area or condition	Yes	No	
1.	Is there a completed floor plan?			
2.	Is there adequate space for maintenance needs? Front 36 in (91.4 cm) minimum, Rear 36 in (91.4 cm) minimum are recommended clearances.			
3.	Is access to the site or computer room restricted?			
4.	Is the computer room structurally complete? Expected date of completion?			
5.	Is a raised floor installed and in good condition?			
6.	Is the raised floor adequate for equipment loading?			
7.	Are there channels or cutouts for cable routing?			
8.	Is there a remote console telephone line available with an RJ11 jack?			
9.	Is a telephone line available?			
10.	Are customer supplied peripheral cables and LAN cables available and of the proper type?			
11.	Are floor tiles in good condition and properly braced?			
12.	Is floor tile underside shiny or painted? If painted, judge the need for particulate test.			
Power and lighting				
No.	Area or condition	Yes	No	
13.	Are lighting levels adequate for maintenance?			
14.	Are there AC outlets available for servicing needs? (i.e. vacuuming)			
15.	Does the input voltage correspond to equipment specifications?			
15A	Is dual source power used? If so, identify type(s) and evaluate grounding.			

18559
D.

Table C-5 Site Inspection Checklist (Continued)

Please check either Yes or No. If No, include comment# or date				**Comment or Date**
16	Does the input frequency correspond to equipment specifications?			
17.	Are lightning arrestors installed inside the building?			
18.	Is power conditioning equipment installed?			
19.	Is there a dedicated branch circuit for equipment?			
20.	Is the dedicated branch circuit less than 250 feet (72.5 meters)?			
21.	Are the input circuit breakers adequate for equipment loads?			
Safety				
No.	Area or condition	Yes	No	
22.	Is there an emergency power shut-off switch?			
23.	Is there a telephone available for emergency purposes?			
24.	Is there a fire protection system in the computer room?			
25.	Is antistatic flooring installed?			
26.	Are there any equipment servicing hazards (loose ground wires, poor lighting, etc.)?			
Cooling				
No.	Area or condition	Yes	No	
27.	Can cooling be maintained between 20 °C and 55 °C (up to 5000 ft.)? Derate 1 °C/1000 ft. above 5000 ft. and up to 10,000 ft.			
28.	Can temperature changes be held to 10 °C per hour with tape media? Can temperature changes be held to 20 °C per hour without tape media?			
29.	Can humidity level be maintained at 40% to 60% at 35 °C noncondensing?			
30.	Are air conditioning filters installed and clean?			
Storage				

18558
D.

Table C-5 Site Inspection Checklist (Continued)

Please check either Yes or No. If No, include comment# or date				Comment or Date
No.	Area or condition	Yes	No	
31.	Are cabinets available for tape and disc media?			
32.	Is shelving available for documentation?			
Training				
No.	Area or Condition			
33	Are personnel enrolled in the System Administrator's Course?			
34	Is on-site training required?			

1855?
D.

Delivery Survey

The delivery survey forms list delivery or installation requirements. If any of the items on the list apply, enter the appropriate information in the areas provided on the form.

Special instructions or recommendations should be entered on the special instructions or recommendations form. The following list gives examples of special instructions or issues:

- Packaging restrictions at the facility, such as size and weight limitations
- Special delivery procedures
- Special equipment required for installation, such as tracking or hoists
- What time the facility is available for installation (after the equipment is unloaded)
- Special security requirements applicable to the facility, such as security clearance

18556
D.

Figure C-2 Delivery Survey (Part 1)

DELIVERY CHECKLIST

DOCK DELIVERY

Is dock large enough for a semitrailer? Yes ______ No ______

Circle the location of the dock and give street name if different than address.

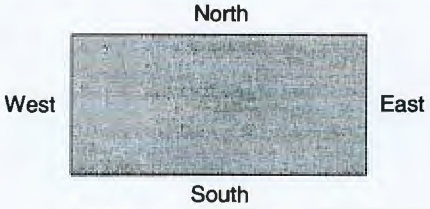

STREET DELIVERY

Circle the location of access door and list street name if different than address.

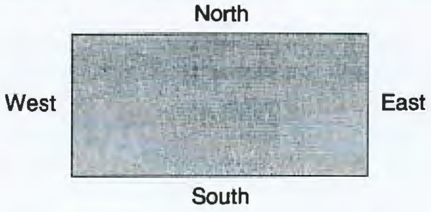

List height ________ and width _______ of access door.

List special permits (if required) for street delivery.

Permit type: ______________________

____

Agency obtained from: ______________________

____

60SP018A
12/7/99

18556
D

Figure C-3 Delivery Survey (Part 2)

ELEVATOR

Fill in the following information if an elevator is required to move equipment.

Capacity (lb or kg) ____________

Depth ____________

Height ____________

Width ____________

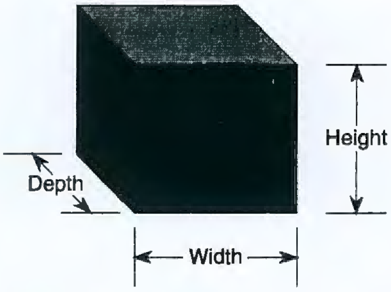

STAIRS

Please list number of flights and stairway dimensions.

Number of flights ____________

Width ____________

Width ____________

Number of flights ____________

Width ____________

Width ____________

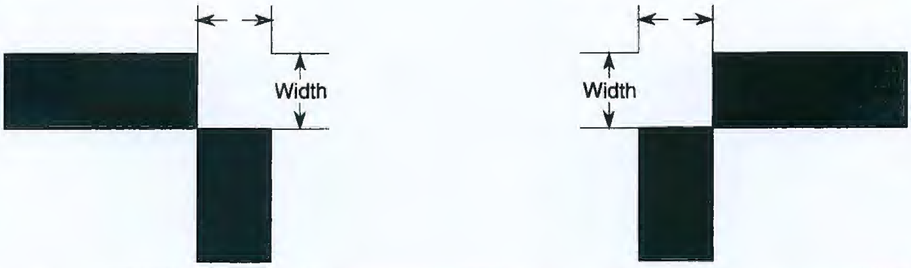

60SP019A
11/24/99

L8554
D

Index

A

Accessing Error Chassis Logs, 98
Acoustic Safety Standards, 248
Adding Users, 71, 80, 83
Additional Components, 22
Altitude Operation Standards, 249
  Effects of Altitude, 249

C

Cabinet Performance Grounding (High Frequency Ground), 255
Changing the Default GSP Configuration, 82, 85
Chassis Code to FRU Decode, 93
Computer Room Physical Space Requirements, 243
  Aisle Space, 243
  Server, 243
Computer Room Unpacking Space, 244
Configure Remote Console, 60
Configure RS-232 Console, 56
Configure System Consoles, 55
Configure the LAN Console, 61
Configure the Web Console, 63
Configuring the GSP LAN Port, 80, 83
Configuring the GSP LAN Port via an ASCII console, 61
Configuring the Rev A Guardian Service Processor (GSP), 80
Configuring the Rev B Guardian Service Processor (GSP), 83
Core I/O Connections, 50
Core I/O Removal, 146
Core I/O Replacement, 147
CPU Removal, 182
CPU Socket Replacement, Cleaning, and CPU Replacement, 183
  Cleaning Kit Components, 183
  CPU Cleaning for System Board Replacement, 185, 186
  Replacing the CPU, 187
  Tools Required, 183
Cross-Referencing Chassis Log Errors to rp54xx FRUs, 93
CRU Physical Location, 216
Customer Replaceable Unit Part Numbers, 219

D

Data Communications Cables, 257
Determine Current System State, 88
Dimensions, 242
Disk Media Backplane Removal, 212
Disk Media Backplane Replacement, 213
Disk Media Housing Removal, 188
Disk Media Housing Replacement, 189
Display Board Removal, 173, 174
  Removing rp5400 and/or rp5450 Model Platform Monitors, 175
  Removing The rp5470 Model Platform Monitor, 175

E

Electrical, 245
  Modular PDUs, 246
  Office High Availability Requirements, 245
    Server-level Enhanced Power Availability, 245
  Power Protection, 245
  System Power Requirements, 246
electrical and environmental guidelines
  electrical conduit ground, 255
  electrostatic discharge prevention, 248
  lighting requirements, 252
  main building electrical ground, 255
  power panel grounds, 255
Electrical Conduit Ground, 255
electrical load requirements, 253
Environment Elements
  Acoustics, 262
Environmental Elements, 258
  Computer Room Preparation, 258
  Cooling Requirements, 258
    Air Conditioning Ducts, 260
    Air Conditioning System Guidelines, 259
    Air Conditioning System Installation, 260
    Air Conditioning System Types, 259
    Basic Air Conditioning Equipment Requirements, 258
    Basic Air Distribution Systems, 259
  Delivery Survey, 272
  Dust and Pollution Control, 260
  Electrostatic Discharge (ESD) Prevention, 262
  Facility Characteristics, 263
    Floor Loading, 263
    Average Floor Loading, 264
    Floor Loading Terms, 263
    Raised Floor Loading, 263
    Typical Raised Floor Site, 264
    Windows, 265
  Humidity Level, 260
  Metallic Particulate Contamination, 261
  Site Inspection, 268
  Space Requirements, 266
    Delivery Space Requirements, 266
    Floor Plan Grid, 266
    Operational Space Requirements, 266
  Typical Installation Schedule, 267
Equipment Grounding Implementation Details, 257
Example of Accessing Error Logs, 98
Expansion I/O LED States, 116
Extend the Server out the Front, 133

F

Factory Integrated Cabinet Installation, 25
Fan, Power Supply, and Disk LED States, 121
Field Replaceable Unit Part Numbers, 235
Front Bezel Removal (Single Piece), 142
Front Bezel Removal (Two Piece), 144
Front Bezel Replacement (Single Piece), 143
Front Bezel Replacement (Two Piece), 145
FRU Physical Location, 221

G

Graphics Troubleshooting, 46
grounding systems

electrical conduit ground, 255
GSP Cables, 56
GSP Configurable Parameters, 71
GSP Features, 53
GSP LAN, 53
GSP LED States, 118
GSP RS-232, 53
Guardian Service Processor (GSP) Overview, 53

H

hot swap procedure
  software, 151
HotPlug
  Hardware HotPlug Procedure, 156
  Hot Swap Procedure for Unattached Physical Volumes, 158
  Hot Swap Software Procedure for Attached Physical Volumes, 157
HotPlug Disk Drive Removal, 151
HotPlug Disk Drive Replacement, 156
HotSwap Card Cage Fan Removal, 163
HotSwap Card Cage Fan Replacement, 164
HotSwap Chassis Fan Cover Removal, 159
HotSwap Chassis Fan Cover Replacement, 160
HotSwap Chassis Fan Removal, 161
HotSwap Chassis Fan Replacement, 162
HotSwap Power Converter Fan Removal, 167
HotSwap Power Converter Fan Replacement, 168
HotSwap Power Supply Removal, 165
HotSwap Power Supply Replacement, 166

I

Identifying Approved Non-E-Series HP Cabinets, 17
Identifying E-Series HP Cabinets, 18
Identifying Static Rail Kit, 18
Individual Component Remove/Replace Instructions, 133
Initial Power-up, 76
Insert the Server from the Front, 134
Install Deskside server, 8
Install Stand-Alone Server in a Cabinet, 11
Installing a PCI Card, 36
Installing Central Processing Units (CPUs), 23
Installing Disk Drives, 47
Installing Graphics, 39
Installing Memory, 28
Installing Peripheral Component Interconnect (PCI) Cards, 32
Installing rp5400 and/or rp5450 DIMMs, 28
Installing rp5470 DIMMs, 30
Installing Stationary Rails, 19
Interpreting Chassis Logs Using the chassis_code.codes File, 100
Interpreting Service Processor Error Chassis Logs, 98
Interpreting System Alerts, 96, 97

K

Key FRU Identification Fields for Error Chassis Logs, 99
Key FRU Identification Fields for System Alerts, 97

L

Power Distribution Safety Grounding, 255
Site Preparation
  Distribution Hardware
    Raceway Systems (electrical conduits), 254
LAN and Telephone, 247
LAN/SCSI LED States, 120
lighting requirements, 253

M

Memory Configuration Rules, 28
Memory DIMM Removal, 171

    Removing rp5400 Memory DIMMs, 171
  Removing rp5470 Memory DIMMs, 171
Memory DIMM Replacement, 172
  Replacing rp5400 and/or rp5450 Memory DIMMs, 172
  Replacing rp5470 Memory DIMMs, 172

O

Online Addition/Replacement (OLA/R) of PCI I/O cards, 38

P

PCI Backplane Removal, 206
  Removing rp5400 and/or rp5450 PCI Backplanes, 206
  Removing rp5470 PCI Backplane, 207
PCI Backplane Replacement, 209
  Replacing rp5400 and/or rp5450 PCI Backplanes, 209
  Replacing rp5470 PCI Backplane, 211
PCI Card Separator/Extractor Removal, 149
PCI Card Separator/Extractor Replacement, 150
PCI I/O Card Installation Order, 34
PCI I/O Card Installation Restrictions, 34
PCI I/O Card Removal, 179
PCI I/O Card Replacement, 180
PCI I/O LED States, 114
Platform Monitor Replacement, 177
  Replacing rp5400 and/or rp5450 Platform Monitors, 177
  Replacing rp5470 Platform Monitor, 177
Power Converter Removal, 193
Power Converter Replacement, 194
Power Panel Ground, 255
Power Supply Failure Example, 95
Problem Symptoms and Repair Actions, 89
Processor Failure Example, 95
Processor Support Module Removal, 169
Processor Support Module Replacement, 170

R

raised floor
  ground system, illustrated, 256
Raised Floor “High Frequency Noise” Grounding, 256
Receive/Unpack A Non-Integrated Server, 5
Removable Media Removal, 190
Removable Media Replacement, 191

SCSI ID Settings, 191
Remove and Replace Procedures by Part, 124
Removing Users, 72, 82, 85
Return the GSP to Default Configurations, 73
Revision A GSP, 50, 54
Revision B GSP, 51, 54
rp5400 and rp5450 I/O Block Diagram, 103
rp5400 and rp5450 System Block Diagram, 102
rp5400/rp5450 PCI Card Slots, 32
rp5430 and rp5470 System Block Diagram, 104, 105
rp5470 PCI Card Slots, 33
rp54xx Server Boot Process, 74
Run/Attention/Fault LED States, 106

S

Sample System Alerts, 97
Secure Web Console Installation, 68
Side Cover Removal, 140
Side Cover Replacement, 141
Site Preparation
Computer Room Safety, 252
Fire Protection, 252
Lighting Requirements for Equipment Servicing, 252
Distribution Hardware, 254
Building Distribution, 254
Wire Selection, 254
Electrical Factors, 252
Electrical Load Requirements (Circuit Breaker Sizing), 253
Grounding Systems, 255
Power Consumption, 253
Power Quality, 253
Power System Protection, 254
Sources of Voltage Fluctuations, 253
Space Requirements, 242
Stand-alone Server Cover Removal, 135
Stand-alone Server Cover Replacement, 137
Stationary L-Bracket Rail Assembly, 17
System Board Removal, 195
Removing rp5400 and/or rp5450 Model System Boards, 195
Removing rp5470 Model System Board, 196
System Board Replacement, 197
rp5400 or rp5450 System Board Replacement, 197
Configuring System Board, 198
rp5470 System Board Replacement, 201
Power up procedure., 204
System Installation Guidelines, 257

T

Temperature and Humidity Operating Standards, 250
Thermal Protection Features, 250
Top Cover Removal, 138
Top Cover Replacement, 139
Troubleshooting and FRU identification, 89

U

Uncrating Space, 242
Unpacking the SPU, 5

V

V-Class system
temperature and humidity specifications, 250

W

Wiring Connections, 257

18551
D.

ANEXO UNIDADE DE BACKUP ROBOTIZADO
PARTE F

18550

hp invent

A4902A- HP Rack System/E, 41U, 19" quartz ▨
A4902D- HP Rack System/E, 41U, 19" graphite ■

Technical Data

Save valuable floor space with the vertically extendible HP Rack System/E

Designed and built to the highest HP quality standards, the Rack System/E delivers leading edge protection in the simplest form. Ease of use, integration and installation characterizes this 41U rack that is comprised of:

- 63% Perforated, locking rear door (ordered as A5213AZ/A5213DZ)
- Bolt-on front/back anti-tip feet
- Numbered columns
- Fully perforated top cap
- 3-inch urethane casters
- Leveling feet
- Side panels

Multiple racks may be tied together to create continuous data center rack space. Individual racks may be expanded an additional 8Us of vertical space.

Standards

Conforms to the Electronic Industries Association (EIA) standard 310-D. It is a Type A cabinet with 41U of vertical mounting space. One 'U' is equal to 44.45 mm (1.75 in).

Customer must order rear door, A5213AZ(quartz)

A5213DZ (graphite)

Features

- Ability to move and ship fully integrated racks
- Optimized ventilation with fully perforated top, and rear door
- Extendibility can add 8Us of vertical mounting space
- Easy, bolt-on (front and back) anti-tip feet
- Numbered 12-gauge steel columns for easy installation and secure racking of up to 907 kg (2000 lbs) of equipment
- Columns include threaded inserts (AVKs) at strategic locations for quick installation of common accessories such as the tie kit, front door and PDUs

Shipping/Setup

- Can fit through most doorways around the world
- Packaging designed for integrated rack shipment
- Self-tuning pallet adjusts for variable integrated rack weights
- Shipping pallet includes ramp for easy set-up

Product Number
A4902D

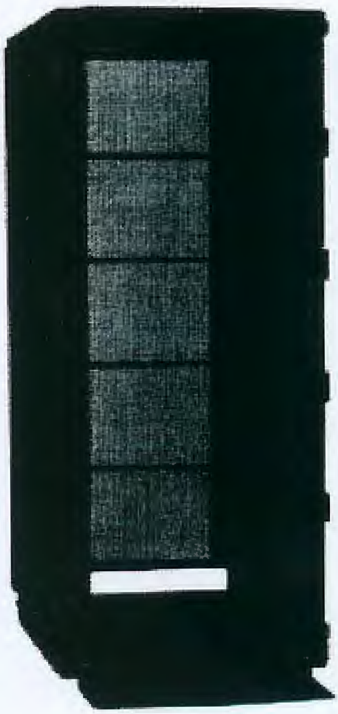

Tools required for setup:

- Torx T25 screwdriver
- Phillips #2 screwdriver
- 13mm Socket wrench

Warranty

One- year replacement

18549
D.

Specifications

Color (A4902A- quartz)
Columns and base: Slate gray
Top: Quartz gray
Side panels: Quartz gray
Rear door: Quartz gray

Color (A4902D- graphite)
Columns and base: Graphite metallic
Top: Graphite metallic
Side panels: Graphite
Rear door: Graphite metallic

Material

Columns: 12-gauge, cold-rolled steel
Base: 10-gauge, cold-rolled steel
Top cap: 18-gauge, cold-rolled steel

A4902D

A	1961.1 mm
B	1824.7 mm
C	111.2 mm

A4901D

A	1605.5 mm
B	1469.1 mm
C	111.2 mm

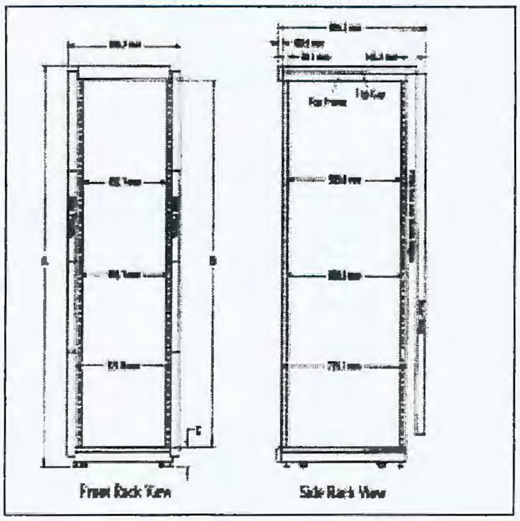

* Dimensions are for reference only

Weight

Rack (empty): 100.45 kg (221 lbs)
Rack (empty) on shipping pallet: 169.3 kg (372.5 lbs)
Rear door (unpacked): 10.68 kg (23.5 lbs)
Anti-tip foot: 16.14 kg (35.5 lbs)

Supported weight

Load capacity:
On shipping pallet: 816 kg (1800 lbs);
Off shipping pallet: 907 kg (2000 lbs)
Casters rating: 453.6 kg (1000 lbs) per caster

Related Products

Other sizes available are:

- **A4901A (quartz)**
- **A4901D (graphite)**
- 33U of vertical mounting space, includes side panels, bolt-on (front and back) anti-tip feet

- **A4900A (quartz)**
- 25U of vertical mounting space,includes side

Optional Accessories

J1506A/ J1506D	Side Panel Kit (1 kit per rack) included in standard rack
J1509A/ J1509D	Front Door: Perforated, lockable
J1512A/ J1512D	Tie Kit
J1514A/ J4387A	Filler Panels (set of 6)
J1518A/ J1518D	Keyboard Kit, retractable
J1519A/ J1519D	Monitor Kit
J1520A/ J1520D	Plain Shelf, static
J1521A	Lift Hooks (set of 4)
J1522A	Mounting Hardware

COBRA TECNOLOGIA
Integradora de Soluções em TI e Serviços

18548
D-

ANEXO UNIDADE DE BACKUP ROBOTIZADO
PARTE G/1

R547
D

Using ServiceGuard Extension for Real Application Cluster (RAC)

First Edition

i n v e n t

Manufacturing Part Number : T1859-90006

June 2003

Legal Notices

The information in this document is subject to change without notice.

Hewlett-Packard makes no warranty of any kind with regard to this manual, including, but not limited to, the implied warranties of merchantability and fitness for a particular purpose. Hewlett-Packard shall not be held liable for errors contained herein or direct, indirect, special, incidental or consequential damages in connection with the furnishing, performance, or use of this material.

Warranty

A copy of the specific warranty terms applicable to your Hewlett-Packard product and replacement parts can be obtained from your local Sales and Service Office.

Restricted Rights Legend

Use, duplication or disclosure by the U.S. Government is subject to restrictions as set forth in subparagraph (c) (1) (ii) of the Rights in Technical Data and Computer Software clause at DFARS 252.227-7013 for DOD agencies, and subparagraphs (c) (1) and (c) (2) of the Commercial Computer Software Restricted Rights clause at FAR 52.227-19 for other agencies.

HEWLETT-PACKARD COMPANY
3000 Hanover Street
Palo Alto, California 94304 U.S.A.

Use of this document and any supporting software media (CD-ROMs, flexible disk, and tape cartridges) supplied for this pack is restricted to this product only. Additional copies of the programs may be made for security and back-up purposes only. Resale of the programs, in their present form or with alterations, is expressly prohibited.

Copyright Notice

Trademark Notices

ActivePerl ® is a registered trademark of ActiveState Tool Corporation.

13545
D.

Java™ and all Java-based trademarks and logos are trademarks or registered trademarks of Sun Microsystems, Inc. in the U.S. and other countries. Hewlett-Packard is independent of Sun Microsystems.

MS-DOS® and Microsoft® are U.S. registered trademarks of Microsoft Corporation.

Netscape ® is a registered trademark of Netscape Communications Corporation.

Oracle ® is a registered trademark of Oracle Corporation.

Oracle8 ™ is a trademark of Oracle Corporation.

UNIX® is a registered trademark in the United States and other countries, licensed exclusively through The Open Group.

VERITAS® is a registered trademark of VERITAS Software Corporation.

VERITAS File System™ is a trademark of VERITAS Software Corporation.

X Window System™ is a trademark of the Massachusetts Institute of Technology.

18544
D

18343
D

Contents

1. Introduction to ServiceGuard Extension for RAC

What is a ServiceGuard Extension for RAC Cluster? 16
Group Membership 17
Using Packages in a Cluster 18
ServiceGuard Extension for RAC Architecture 19
Group Membership Daemon 19
How ServiceGuard Works with Oracle Real Application Clusters 20
Configuring Packages for Oracle RAC Instances 21
Node Failure 22
Larger Clusters 24
Up to Four Nodes with SCSI Storage 24
Point to Point Connections to Storage Devices 25

2. ServiceGuard Configuration for Real Application Clusters

Planning Database Storage 28
Volume Planning with SLVM 28
Volume Planning with CVM 30
Installing ServiceGuard Extension for RAC 32
Configuration File Parameters 33
Creating a Storage Infrastructure with LVM 34
Building Volume Groups for RAC on Mirrored Disks 34
Building Mirrored Logical Volumes for RAC with LVM Commands 36
Creating RAC Volume Groups on Disk Arrays 38
Creating Logical Volumes for RAC on Disk Arrays 40
Oracle Demo Database Files 40
Displaying the Logical Volume Infrastructure 42
Exporting the Logical Volume Infrastructure 42
Installing Oracle Real Application Clusters 45
Cluster Configuration ASCII File 46
Creating a Storage Infrastructure with CVM 49
Initializing the VERITAS Volume Manager 49
Preparing the Cluster for Use with CVM 50
Starting the Cluster and Identifying the Master Node 51
Converting Disks from LVM to CVM 51
Initializing Disks for CVM 51
Creating Disk Groups for RAC 51
Creating Volumes 52
Oracle Demo Database Files 53

1854
D.

Contents

Adding Disk Groups to the Cluster Configuration 55
Configuring Packages that Access the Oracle Database 56

3. Maintenance and Troubleshooting

Reviewing Cluster and Package States with the cmviewcl Command 58
Types of Cluster and Package States. 58
Examples of Cluster and Package States 62
Managing the Shared Storage. 70
Making LVM Volume Groups Shareable 70
Activating an LVM Volume Group in Shared Mode 71
Making Changes to Shared Volume Groups 72
Adding Additional Shared LVM Volume Groups 73
Changing the VxVM or CVM Storage Configuration 74
Removing ServiceGuard Extension for RAC from a System. 75
Monitoring Hardware 76
Using Event Monitoring Service 76
Using EMS Hardware Monitors 76
Using HP Predictive Monitoring 77
Adding Disk Hardware 78
Replacing Disks 79
Replacing a Mechanism in a Disk Array Configured with LVM 79
Replacing a Mechanism in an HA Enclosure Configured with Exclusive LVM 79
Offline Replacement of a Mechanism in an HA Enclosure Configured with Shared LVM (SLVM) 80
Replacing a Lock Disk 81
On-line Hardware Maintenance with In-line SCSI Terminator 82
Replacement of I/O Cards 85
Replacement of LAN Cards. 86
Off-Line Replacement. 86
On-Line Replacement. 86
After Replacing the Card 87

A. Blank Planning Worksheets

LVM Volume Group and Physical Volume Worksheet. 90
VxVM Disk Group and Disk Worksheet 91
Oracle Logical Volume Worksheet. 92

18341
D

Contents

Index **93**

18540
Gomes

Contents

11539
Scu3A

Printing History

Table 1

Printing Date	Part Number	Edition
June 2003	T1859-90006	First Edition Print, CD-ROM (Instant Information), and Web (`http://www.docs.hp.com/`)

The last printing date and part number indicate the current edition, which applies to the 11.14.03 and 11.15 version of ServiceGuard Extension for RAC (Oracle Real Application Cluster).

The printing date changes when a new edition is printed. (Minor corrections and updates which are incorporated at reprint do not cause the date to change.) The part number is revised when extensive technical changes are incorporated.

New editions of this manual will incorporate all material updated since the previous edition. To ensure that you receive the new editions, you should subscribe to the appropriate product support service. See your HP sales representative for details.

HP Printing Division:

Business Critical Computing
Hewlett-Packard Co.
19111 Pruneridge Ave.
Cupertino, CA 95014

13538

18537
5026

Preface

This guide describes how to use the ServiceGuard Extension for RAC (Oracle Real Application Cluster) to configure ServiceGuard clusters for use with Oracle Real Application Cluster software on HP 9000 High Availability clusters running the HP-UX operating system. The contents are as follows:

- Chapter 1, "Introduction," describes a ServiceGuard cluster and provides a roadmap for using this guide. This chapter should be used as a supplement to Chapters 1–3 of *Managing MC/ServiceGuard*.
- Chapter 2, "Configuration," describes the additional steps you need to take to use ServiceGuard with Real Application Clusters. This chapter should be used as a supplement to Chapters 4–6 of *Managing MC/ServiceGuard*.
- Chapter 3, "Maintenance and Troubleshooting," describes tools and techniques necessary for ongoing cluster operation. This chapter should be used as a supplement to Chapters 7–8 of *Managing MC/ServiceGuard*.

Related Publications

The following documents contain additional useful information:

- *Clusters for High Availability: a Primer of HP Solutions.* Hewlett-Packard Professional Books: Prentice Hall PTR, 2001 (ISBN 0-13-089355-2)
- *Managing MC/ServiceGuard* (B3936-90065)
- *Using High Availability Monitors* (B5736-90022)
- *Using the Event Monitoring Service* (B7612-90009)
- *Using Advanced Tape Services* (B3936-90032)
- *Designing Disaster Tolerant High Availability Clusters* (B7660-90009)
- *Managing ServiceGuard Extension for SAP/R3* (B7885-90004)
- *Managing Systems and Workgroups* (B2355-90157)
- *Managing Highly Available NFS* (B5125-90001)
- *HP Auto Port Aggregation Release Notes*

18536
Gomes.

Before attempting to use VxVM storage with ServiceGuard, please refer to the following:

- *VERITAS Volume Manager Administrator's Guide*. This contains a glossary of VERITAS terminology.
- *VERITAS Volume Manager Storage Administrator Administrator's Guide*
- *VERITAS Volume Manager Reference Guide*
- *VERITAS Volume Manager Migration Guide*
- *VERITAS Volume Manager for HP-UX Release Notes*

Use the following URL to access HP's high availability web page:

- http://www.hp.com/go/ha

Use the following URL for access to a wide variety of HP-UX documentation:

- http://docs.hp.com/hpux

Problem Reporting

If you have any problems with the software or documentation, please contact your local Hewlett-Packard Sales Office or Customer Service Center.

Conventions

We use the following typographical conventions.

audit (5) — An HP-UX manpage. *audit* is the name and *5* is the section in the *HP-UX Reference*. On the web and on the Instant Information CD, it may be a hot link to the manpage itself. From the HP-UX command line, you can enter "`man audit`" or "`man 5 audit`" to view the manpage. See *man* (1).

Book Title — The title of a book. On the web and on the Instant Information CD, it may be a hot link to the book itself.

KeyCap — The name of a keyboard key. Note that **Return** and **Enter** both refer to the same key.

Emphasis — Text that is emphasized.

Emphasis — Text that is strongly emphasized.

Term — The defined use of an important word or phrase.

`ComputerOut` — Text displayed by the computer.

13335
D-

`UserInput`	Commands and other text that you type.
`Command`	A command name or qualified command phrase.
`Variable`	The name of a variable that you may replace in a command or function or information in a display that represents several possible values.
[ ]	The contents are optional in formats and command descriptions. If the contents are a list separated by |, you must choose one of the items.
{ }	The contents are required in formats and command descriptions. If the contents are a list separated by |, you must choose one of the items.
...	The preceding element may be repeated an arbitrary number of times.
|	Separates items in a list of choices.

18334
D

18533
D.

1 Introduction to ServiceGuard Extension for RAC

ServiceGuard Extension for RAC (SGeRAC) enables the Oracle Real Application Cluster (RAC), formerly known as Oracle Parallel Server RDBMS, to run on HP 9000 high availability clusters under the HP-UX operating system. This chapter introduces ServiceGuard Extension for RAC and shows where to find different kinds of information in this book. The following topics are presented:

- What is a ServiceGuard Extension for RAC Cluster?
- ServiceGuard Extension for RAC Architecture
- How ServiceGuard Works with Oracle Real Application Clusters
- Configuring Packages for Oracle RAC Instances
- Node Failure
- Larger Clusters

If you are ready to start setting up RAC clusters, skip ahead to the chapter “ServiceGuard Configuration for Real Application Clusters.”

19532
D

What is a ServiceGuard Extension for RAC Cluster?

A **high availability cluster** is a grouping of HP 9000 series 800 servers having sufficient redundancy of software and hardware components that a single point of failure will not disrupt the availability of computer services. High availability clusters configured with Oracle Real Application Cluster software are known as **RAC clusters**. Figure 1-1 shows a very simple picture of the basic configuration of a RAC cluster on HP-UX.

Figure 1-1 **Overview of Oracle RAC Configuration on HP-UX**

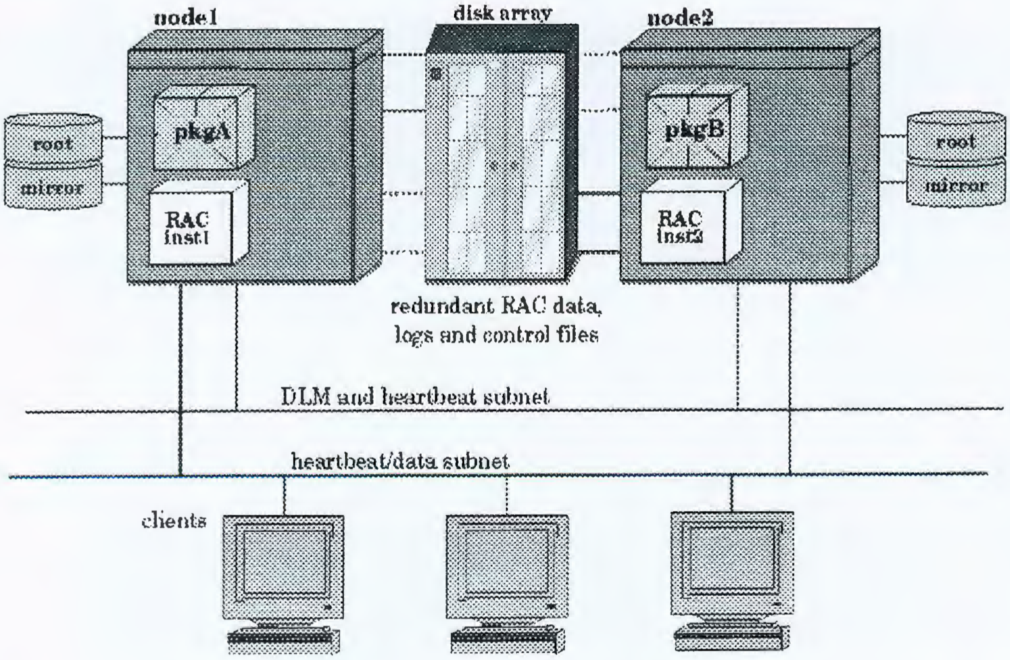

In the figure, two loosely coupled HP 9000 series 800 systems (each one known as a **node**) are running separate instances of Oracle software that read data from and write data to a shared set of disks. Clients connect to one node or the other via LAN.

18533
D.

RAC on HP-UX lets you maintain a single database image that is accessed by the HP 9000 servers in parallel, thereby gaining added processing power without the need to administer separate databases. Further, when properly configured, ServiceGuard Extension for RAC provides a highly available database that continues to operate even if one hardware component should fail.

Group Membership

Oracle RAC 8.1.x and later systems implement the concept of **group membership**, which allows multiple instances of RAC to run on each node. Related processes are configured into **groups**. Groups allow processes in different instances to choose which other processes to interact with. This allows the support of multiple databases within one RAC cluster.

A Group Membership Service (GMS) component provides a process monitoring facility to monitor group membership status. GMS is provided by the `cmgmsd` daemon, which is an HP component installed with ServiceGuard Extension for RAC.

Figure 1-2 shows how group membership works. Nodes 1 through 4 of the cluster share the Sales database, but only Nodes 3 and 4 share the HR database. Consequently, there is one instance of RAC each on Node 1 and Node 2, and there are two instances of RAC each on Node 3 and Node 4. The RAC processes accessing the Sales database constitute one group, and the RAC processes accessing the HR database constitute another group.

J3530
D

Figure 1-2 **Group Membership Services**

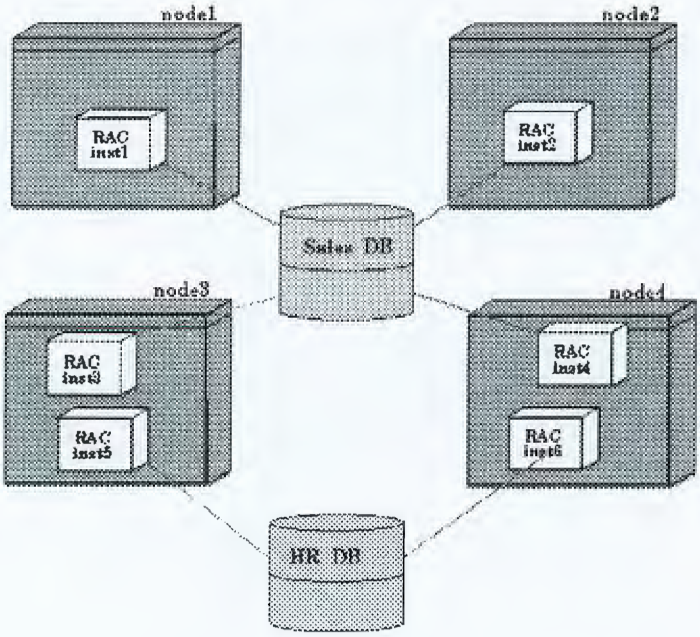

Using Packages in a Cluster

In order to make other important applications highly available (in addition to the Oracle Real Application Cluster), you can configure your RAC cluster to use **packages**. Packages group applications and services together; in the event of a service, node, or network failure, ServiceGuard Extension for RAC can automatically transfer control of all system resources in a designated package to another node within the cluster, allowing your applications to remain available with minimal interruption.

NOTE In RAC clusters, you create packages to start and stop RAC itself as well as to run applications that access the database instances. For details on the use of packages with RAC, refer to the chapter "Configuring Packages and Their Services."

18529
D

ServiceGuard Extension for RAC Architecture

This chapter discusses the main software components used by ServiceGuard Extension for RAC in some detail. The components are:

- Oracle Components
  - — Custom Oracle Applications
- ServiceGuard Extension for RAC Components
  - — Group Membership Services (RAC)
  - — Package Manager
  - — Cluster Manager
  - — Network Manager
- Operating System
  - — Volume Manager Software
  - — HP-UX Kernel

Group Membership Daemon

In addition to the ServiceGuard daemon processes mentioned in Chapter 3 of *Managing MC/ServiceGuard*, there is another daemon that is used for each version of Oracle to enable communication with ServiceGuard Extension for RAC:

- cmgmsd—Group Membership Daemon for RAC (OPS) 8.1 or later

releases

This HP daemon provides group membership services for Oracle Real Application Cluster (OPS) 8.1.x. Group membership allows multiple Oracle instances to run on the same cluster node. GMS is illustrated in Figure 1-2 on page 18.

How ServiceGuard Works with Oracle Real Application Clusters

ServiceGuard provides the cluster framework for Oracle, a relational database product in which multiple database instances run on different cluster nodes. A central component of Real Application Clusters is the distributed lock manager (DLM), which provides parallel cache management for database instances. Each node in a RAC cluster starts an instance of the DLM process when the node joins the cluster, and the instances then communicate with each other over the network.

The Distributed Lock Manager is an internal component of the Real Application Clusters software. The **group membership service** (GMS) is the means by which Oracle instances communicate with the ServiceGuard cluster software. GMS runs as a separate daemon process that communicates with the cluster manager. This daemon is an HP component known as `cmgmsd`.

The cluster manager starts up, monitors, and shuts down the GMS daemon. When an Oracle instance starts, the instance registers itself with GMS; thereafter, if an Oracle instance fails, GMS notifies other cluster nodes to perform recovery. If GMS dies unexpectedly, ServiceGuard will fail the node with a TOC (Transfer of Control).

Configuring Packages for Oracle RAC Instances

Oracle instances can be configured as packages with a single node in their node list. Package configuration is described in Chapter 2.

NOTE

Packages that start and halt Oracle instances (called **instance packages**) do *not* fail over from one node to another; they are single-node packages. You should include only one node name in the package ASCII configuration file. The `AUTO_RUN` setting will determine whether the RAC instance will start up as the node joins the cluster. Your cluster may include RAC and non-RAC packages in the same configuration.

18526
D.

Node Failure

RAC cluster configuration is designed so that in the event of a node failure, another node with a separate instance of Oracle can continue processing transactions. Figure 1-3 shows a typical cluster with instances running on both nodes.

Figure 1-3 **Before Node Failure**

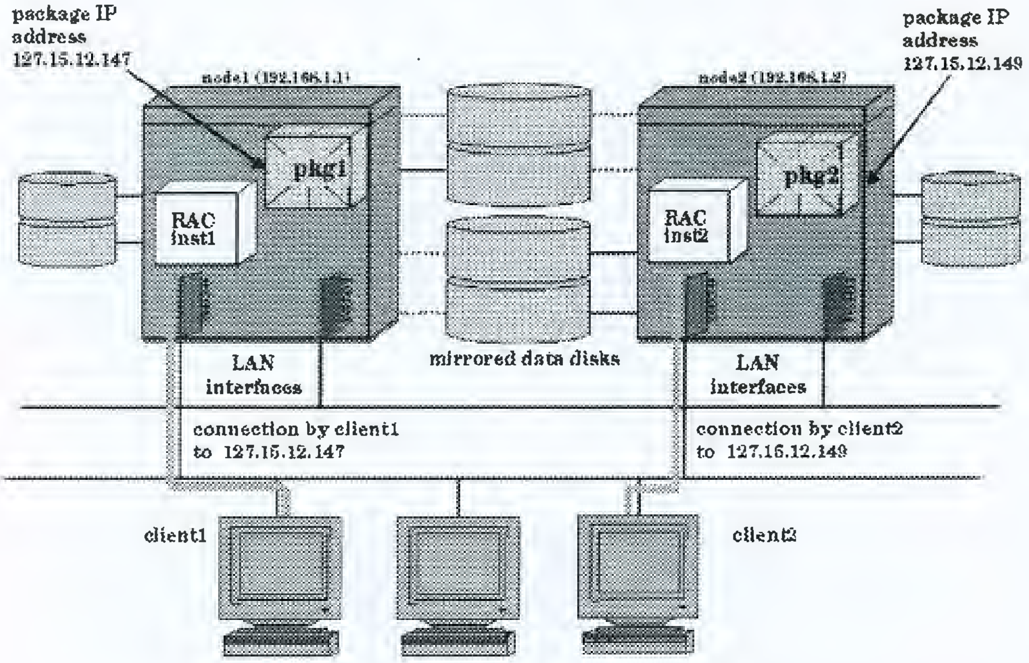
package IP address 127.15.12.147
node1 (192.168.1.1)
pkg1
RAC inst1
LAN interfaces
mirrored data disks
node2 (192.168.1.2)
package IP address 127.15.12.149
pkg2
RAC inst2
LAN interfaces
connection by client1 to 127.15.12.147
connection by client2 to 127.15.12.149
client1
client2

Figure 1-4 shows the condition where Node 1 has failed and Package 1 has been transferred to Node 2. Oracle instance 1 is no longer operating, but it does *not* fail over to Node 2. Package 1's IP address was transferred to Node 2 along with the package. Package 1 continues to be available and is now running on Node 2. Also note that Node 2 can now access both Package 1's disk and Package 2's disk. Oracle instance 2 now handles all database access, since instance 1 has gone down.

18525
D

Figure 1-4 After Node Failure

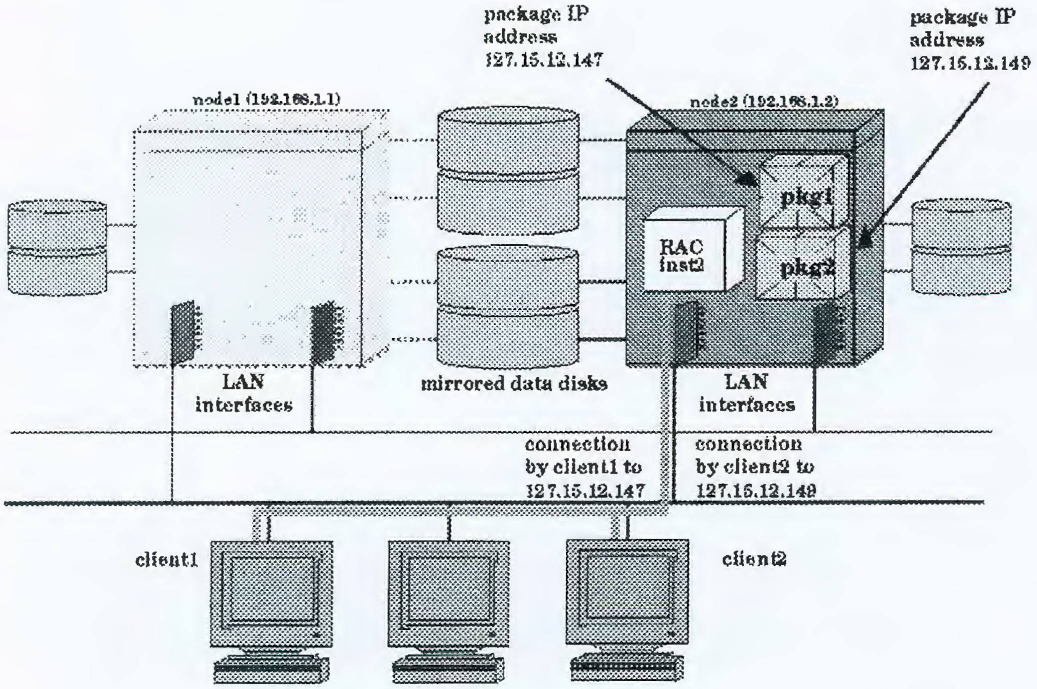

In the above figure, pkg1 and pkg2 are not instance packages. They are shown to illustrate the movement of packages in general.

28524
▽

Larger Clusters

ServiceGuard Extension for RAC supports clusters of up to 16 nodes. The actual cluster size is limited by the type of storage and the type of volume manager used.

Up to Four Nodes with SCSI Storage

You can configure up to four nodes using a shared F/W SCSI bus; for more than 4 nodes, FibreChannel must be used. An example of a four-node RAC cluster appears in the following figure.

Figure 1-5 **Four-Node RAC Cluster**

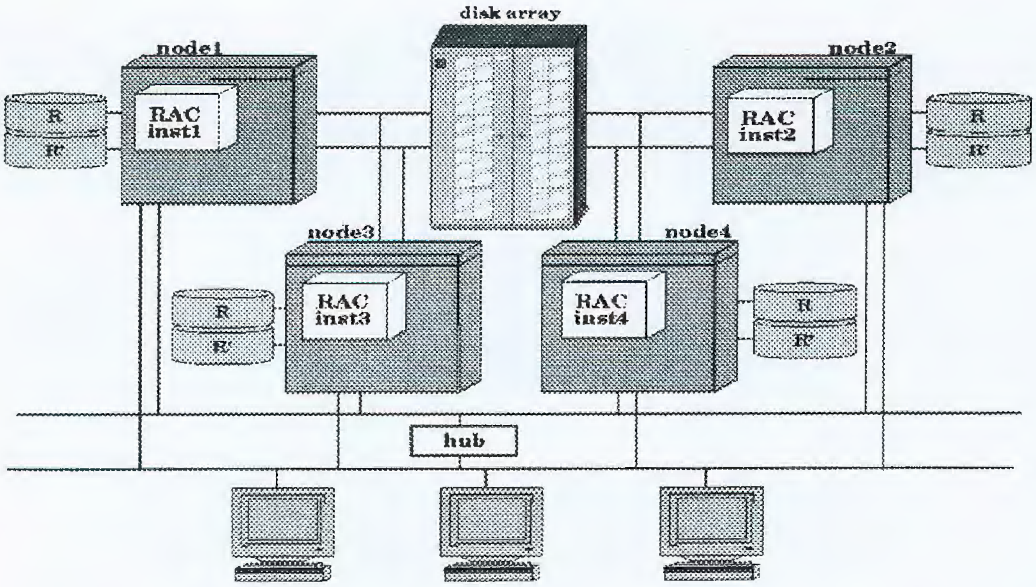

In this type of configuration, each node runs a separate instance of RAC and may run one or more high availability packages as well.

The figure shows a dual Ethernet configuration with all four nodes connected to a disk array (the details of the connections depend on the type of disk array). In addition, each node has a mirrored root disk (R and R'). Nodes may have multiple connections to the same array using

18523
D.

alternate links (PV links) to take advantage of the array's use of RAID levels for data protection. Alternate links are further described in the section "Creating RAC Volume Groups on Disk Arrays" on page 38.

Point to Point Connections to Storage Devices

Some storage devices allow point-to-point connection to a large number of host nodes without using a shared SCSI bus. An example is shown in Figure 1-6, a cluster consisting of eight nodes with a FibreChannel interconnect. (Client connection is provided through Ethernet.) The nodes access shared data on an HP SureStore XP series or EMC disk array configured with 16 I/O ports. Each node is connected to the array using two separate F/W SCSI channels configured with PV Links. Each channel is a dedicated bus; there is no daisy-chaining.

Figure 1-6 **Eight-Node Cluster with XP or EMC Disk Array**

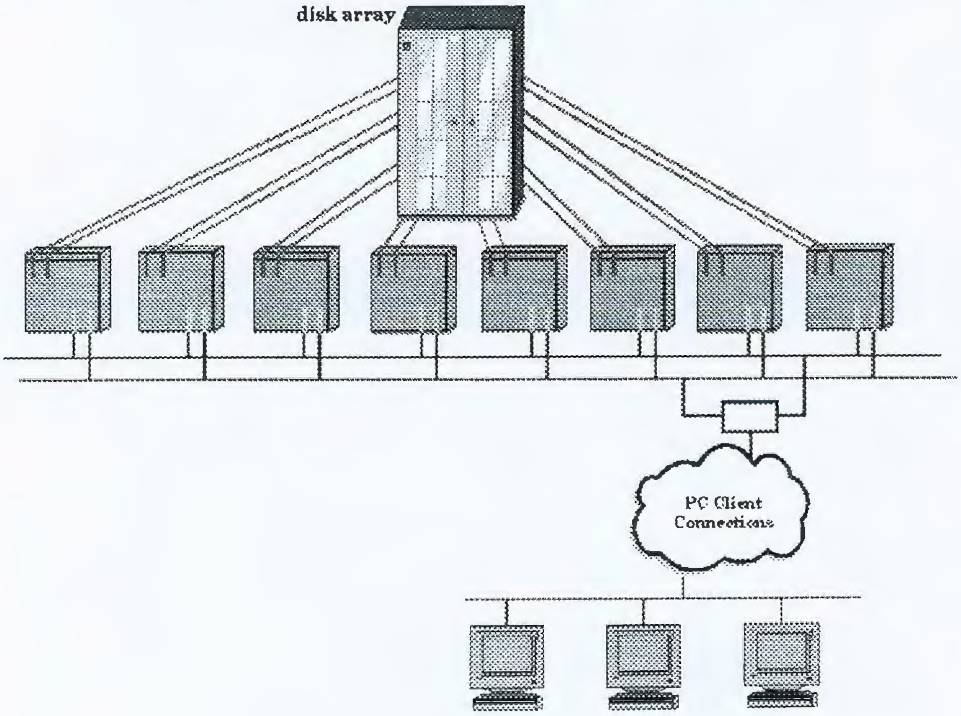

18522
D.

FibreChannel switched configurations also are supported using either an arbitrated loop or fabric login topology. For additional information about supported cluster configurations, refer to the *HP 9000 Servers Configuration Guide*, available through your HP representative.

13521
D

2 ServiceGuard Configuration for Real Application Clusters

This chapter shows the additional planning and configuration that is needed to use Oracle Real Application Clusters with ServiceGuard. The following topics are presented:

- Planning Database Storage
- Installing ServiceGuard Extension for RAC
- Creating a Storage Infrastructure with LVM
- Installing Oracle Real Application Clusters
- Cluster Configuration ASCII File
- Creating a Storage Infrastructure with CVM
- Configuring Packages that Access the Oracle Database

Planning Database Storage

The files needed by the Oracle database must be placed on physical volumes that are accessible to all RAC cluster nodes. This section shows how to plan the volumes using either SLVM or VERITAS CVM storage groups.

Volume Planning with SLVM

Storage capacity for the Oracle database must be provided in the form of logical volumes located in shared volume groups. The Oracle software requires at least two log files (an one undo tablespace for Oracle9) for each Oracle instance, several Oracle control files and data files for the database itself. For all these files, ServiceGuard Extension for RAC uses HP-UX *raw logical volumes*, which are located in volume groups that are shared between the nodes in the cluster. High availability is achieved by using high availability disk arrays in RAID modes. The logical units of storage on the arrays are accessed from each node through multiple physical volume links (PV links, also known as alternate links), which provide redundant paths to each unit of storage.

Fill out a Logical Volume worksheet to provide logical volume names for logical volumes that you will create with the `lvcreate` command. The Oracle DBA and the HP-UX system administrator should prepare this worksheet together. Create entries for shared volumes only. For each logical volume, enter the full pathname of the raw logical volume device file. Be sure to include the desired size in MB. Following is a sample worksheet filled out. However, this sample is only representative. For different versions of the Oracle database, the size of files are different. Refer to the appendix, "Blank Planning Worksheets" for samples of blank worksheets. Make as many copies as you need. Fill out the worksheet and keep it for future reference.

18519
D.

```
ORACLE LOGICAL VOLUME WORKSHEET FOR LVM                       Page ___ of ____
===============================================================================
                              RAW LOGICAL VOLUME NAME            SIZE (MB)

Oracle Control File _____/dev/vg_ops/ropsctl1.ctl_______100______

Oracle Control File 2: ___/dev/vg_ops/ropsctl2.ctl______100______
Oracle Control File 3: ___/dev/vg_ops/ropsctl3.ctl______100______
Instance 1 Redo Log 1: ___/dev/vg_ops/rops1log1.log_____20______
Instance 1 Redo Log 2: ___/dev/vg_ops/rops1log2.log_____20_______
Instance 1 Redo Log 3: ___/dev/vg_ops/rops1log3.log_____20_______
Instance 1 Redo Log: __________________________________________________
Instance 1 Redo Log: __________________________________________________
Instance 2 Redo Log 1: ___/dev/vg_ops/rops2log1.log____20________
Instance 2 Redo Log 2: ___/dev/vg_ops/rops2log2.log____20________
Instance 2 Redo Log 3: ___/dev/vg_ops/rops2log3.log____20_________
Instance 2 Redo Log: _________________________________________________
Instance 2 Redo Log: __________________________________________________
Data: System ___/dev/vg_ops/ropssystem.dbf___400__________
Data: Temp ___/dev/vg_ops/ropstemp.dbf______100_______
Data: Users ___/dev/vg_ops/ropsusers.dbf_____120_________
Data: Tools ___/dev/vg_ops/ropstools.dbf____15___________
Data: User data ___/dev/vg_ops/ropsdata1.dbf_200__________
Data: User data ___/dev/vg_ops/ropsdata2.dbf__200__________
Data: User data ___/dev/vg_ops/ropsdata3.dbf__200__________
Data: Rollback ___/dev/vg_ops/ropsrollback.dbf__300_________(Oracle8)

parameter: spfile1 /dev/vg_ops/ropsspfile1.ora __5_____(Oracle9)

Instance 1  undotbs1: /dev/vg_ops/ropsundotbs1.dbf___312___(Oracle9)

Instance 2  undotbs2: /dev/vg_ops/ropsundotbs2.dbf___312___(Oracle9)

Data: example1__/dev/vg_ops/ropsexample1.dbf__________160____(Oracle9)

data: cwmlite1__/dev/vg_ops/ropscwmlite1.dbf__100____(Oracle9)

Data: indx1__/dev/vg_ops/ropsindx1.dbf____70___(Oracle9)

Data: drsys1__/dev/vg_ops/ropsdrsys1.dbf___90___(Oracle9)
```

Volume Planning with CVM

Storage capacity for the Oracle database must be provided in the form of volumes located in shared disk groups. The Oracle software requires at least two log files (an one undo tablespace for Oracle9) for each Oracle instance, several Oracle control files and data files for the database itself. For all these files, ServiceGuard Extension for RAC uses HP-UX *raw volumes*, which are located in disk groups that are shared between the nodes in the cluster. High availability is achieved by using high availability disk arrays in RAID modes. The logical units of storage on the arrays are accessed from each node through multiple physical volume links via DMP (Dynamic Multi-pathing), which provides redundant paths to each unit of storage.

Fill out the VERITAS Volume worksheet to provide volume names for volumes that you will create using the VERITAS utilities. The Oracle DBA and the HP-UX system administrator should prepare this worksheet together. Create entries for shared volumes only. For each volume, enter the full pathname of the raw volume device file. Be sure to include the desired size in MB. Following is a sample worksheet filled out. Refer to the appendix, "Blank Planning Worksheets" for samples of blank worksheets. Make as many copies as you need. Fill out the worksheet and keep it for future reference.

```
ORACLE LOGICAL VOLUME WORKSHEET FOR CVM                        Page ___ of ____
===============================================================================
                             RAW LOGICAL VOLUME NAME           SIZE (MB)

Oracle Control File 1: ___/dev/vx/rdsk/ops_dg/opsctl1.ctl______100______

Oracle Control File 2: ___/dev/vx/rdsk/ops_dg/opsctl2.ctl______100______

Oracle Control File 3: ___/dev/vx/rdsk/ops_dg/opsctl3.ctl______100______

Instance 1 Redo Log 1: ___/dev/vx/rdsk/ops_dg/ops1log1.log_____20______

Instance 1 Redo Log 2: ___/dev/vx/rdsk/ops_dg/ops1log2.log_____20_______

Instance 1 Redo Log 3: ___/dev/vx/rdsk/ops_dg/ops1log3.log_____20_______

Instance 1 Redo Log: ___________________________________________________

Instance 1 Redo Log: ___________________________________________________

Instance 2 Redo Log 1: ___/dev/vx/rdsk/ops_dg/ops2log1.log____20________

Instance 2 Redo Log 2: ___/dev/vx/rdsk/ops_dg/ops2log2.log____20________

Instance 2 Redo Log 3: ___/dev/vx/rdsk/ops_dg/ops2log3.log____20_________

Instance 2 Redo Log: _________________________________________________

Instance 2 Redo Log: __________________________________________________

Data: System ___/dev/vx/rdsk/ops_dg/system.dbf___400__________

Data: Temp ___/dev/vx/rdsk/ops_dg/temp.dbf______100_______

Data: Users ___/dev/vx/rdsk/ops_dg/users.dbf_____120_________

Data: Tools ___/dev/vx/rdsk/ops_dg/tools.dbf____15___________

Data: User data ___/dev/vx/rdsk/ops_dg/data1.dbf_200__________

Data: User data ___/dev/vx/rdsk/ops_dg/data2.dbf__200__________

Data: User data ___/dev/vx/rdsk/ops_dg/data3.dbf__200__________

Data: Rollback ___/dev/vx/rdsk/ops_dg/rollback.dbf__300_________(Oracle8)


parameter: spfile1 /dev/vx/rdsk/ops_dg/spfile1.ora __5_____(Oracle9)

Instance 1  undotbs1: /dev/vx/rdsk/ops_dg/undotbs1.dbf___312___(Oracle9)

Instance 2  undotbs2: /dev/vx/rdsk/ops_dg/undotbs2.dbf___312___(Oracle9)

Data: example1__/dev/vx/rdsk/ops_dg/example1.dbf__________160____(Oracle9)

data: cwmlite1__/dev/vx/rdsk/ops_dg/cwmlite1.dbf__100____(Oracle9)

Data: indx1  /dev/vx/rdsk/ops_dg/indx1.dbf____70___(Oracle9)

Data: drsys1  /dev/vx/rdsk/ops_dg/drsys1.dbf___90___(Oracle9)
```

18516
D.

Installing ServiceGuard Extension for RAC

Installing ServiceGuard Extension for RAC includes updating the software and rebuilding the kernel to support high availability cluster operation for Oracle Real Application Clusters.

Prior to installing ServiceGuard Extension for RAC, the following must be installed:

- HP-UX 11i
- Correct version of ServiceGuard

To install ServiceGuard Extension for RAC, use the following steps *for each node*:

1. Mount the distribution media in the tape drive or CD ROM reader.
2. Run Software Distributor, using the `swinstall` command.
3. Specify the correct input device.
4. Choose the following bundle from the displayed list:

   `ServiceGuard Extension for RAC`
5. After choosing the bundle, select OK. The software is loaded.
6. Run `ioscan` on each node to validate that disks and drivers have been configured correctly.

18515
D

Configuration File Parameters

You need to code specific entries for all the storage groups that you want to use in an Oracle RAC configuration. If you are using LVM, the `OPS_VOLUME_GROUP` parameter is included in the cluster ASCII file. If you are using VERITAS CVM, the `STORAGE_GROUP` parameter is included in the package ASCII file. Details are as follows:

OPS_VOLUME_GROUP

The name of an LVM volume group whose disks are attached to at least two nodes in the cluster; the disks will be accessed by more than one node at a time using SLVM with concurrency control provided by Oracle RAC. Such disks are considered cluster aware.

Volume groups listed under this parameter are marked for activation in shared mode. The entry can contain up to 40 characters.

STORAGE_GROUP

This parameter is used for CVM disk groups. Enter the names of all the CVM disk groups the package will use.

In the ASCII package configuration file, this parameter is called `STORAGE_GROUP`.

NOTE *Do not* enter the names of LVM volume groups or VxVM disk groups in the package ASCII configuration file.

Creating a Storage Infrastructure with LVM

In addition to configuring the cluster, you create the appropriate logical volume infrastructure to provide access to data from different nodes. This is done with Logical Volume Manager (LVM), VERITAS Cluster Volume Manager (CVM), or VERITAS Volume Manager (VxVM). LVM and VxVM configuration are done before cluster configuration, and CVM configuration is done *after* cluster configuration.

This section describes how to create LVM volume groups for use with Oracle data. Before configuring the cluster, you create the appropriate logical volume infrastructure to provide access to data from different nodes. This is done with Logical Volume Manager. Separate procedures are given for the following:

- Building Volume Groups for RAC on Mirrored Disks
- Building Mirrored Logical Volumes for RAC with LVM Commands
- Creating RAC Volume Groups on Disk Arrays
- Creating Logical Volumes for RAC on Disk Arrays

The Event Monitoring Service HA Disk Monitor provides the capability to monitor the health of LVM disks. If you intend to use this monitor for your mirrored disks, you should configure them in physical volume groups. For more information, refer to the manual *Using HA Monitors*.

Building Volume Groups for RAC on Mirrored Disks

The procedure described in this section uses **physical volume groups** for mirroring of individual disks to ensure that each logical volume is mirrored to a disk on a different I/O bus. This kind of arrangement is known as **PVG-strict mirroring**. It is assumed that your disk hardware is already configured in such a way that a disk to be used as a mirror copy is connected to each node on a different bus than the bus that is used for the other (primary) copy.

For more information on using LVM, refer to the HP-UX *Managing Systems and Workgroups* manual.

18513

Creating Volume Groups and Logical Volumes

If your volume groups have not been set up, use the procedure in the next sections. If you have already done LVM configuration, skip ahead to the section “Configuring the Cluster.”

Selecting Disks for the Volume Group Obtain a list of the disks on both nodes and identify which device files are used for the same disk on both. Use the following command on each node to list available disks as they are known to each system:

```
# lssf /dev/dsk/*
```

In the following examples, we use `/dev/rdsk/c1t2d0` and `/dev/rdsk/c0t2d0`, which happen to be the device names for the same disks on both *ftsys9* and *ftsys10*. In the event that the device file names are different on the different nodes, make a careful note of the correspondences.

Creating Physical Volumes On the configuration node (ftsys9), use the `pvcreate` command to define disks as physical volumes. This only needs to be done on the configuration node. Use the following commands to create two physical volumes for the sample configuration:

```
# pvcreate -f /dev/rdsk/c1t2d0
# pvcreate -f /dev/rdsk/c0t2d0
```

Creating a Volume Group with PVG-Strict Mirroring Use the following steps to build a volume group on the configuration node (ftsys9). Later, the same volume group will be created on other nodes.

1. First, set up the group directory for `vgdatabase`:

   ```
   # mkdir /dev/vgdatabase
   ```

2. Next, create a control file named *group* in the directory `/dev/vgdatabase`, as follows:

   ```
   # mknod /dev/vgdatabase/group c 64 0xhh0000
   ```

   The major number is always 64, and the hexadecimal minor number has the form

   ```
   0xhh0000
   ```

18512
D

where *hh* must be unique to the volume group you are creating. Use the next hexadecimal number that is available on your system, after the volume groups that are already configured. Use the following command to display a list of existing volume groups:

```
# ls -l /dev/*/group
```

3. Create the volume group and add physical volumes to it with the following commands:

```
# vgcreate -g bus0 /dev/vgdatabase /dev/dsk/c1t2d0
# vgextend -g bus1 /dev/vgdatabase /dev/dsk/c0t2d0
```

   The first command creates the volume group and adds a physical volume to it in a physical volume group called *bus0*. The second command adds the second drive to the volume group, locating it in a different physical volume group named *bus1*. The use of physical volume groups allows the use of PVG-strict mirroring of disks and PV links.

4. Repeat this procedure for additional volume groups.

Building Mirrored Logical Volumes for RAC with LVM Commands

After you create volume groups and define physical volumes for use in them, you define mirrored logical volumes for data, logs, and control files. It is recommended that you use a shell script to issue the commands described in the next sections. The commands you use for creating logical volumes vary slightly depending on whether you are creating logical volumes for RAC redo log files or for use with Oracle data.

Creating Mirrored Logical Volumes for RAC Redo Logs and Control Files

Create logical volumes for use as redo log and control files by selecting mirror consistency recovery. Use the same options as in the following example:

```
# lvcreate -m 1 -M n -c y -s g -n redo1.log -L 4 /dev/vg_ops
```

The `-m 1` option specifies single mirroring; the `-M n` option ensures that mirror write cache recovery is set off; the `-c y` means that mirror consistency recovery is enabled; the `-s g` means that mirroring is

L8521
D

ANEXO UNIDADE DE BACKUP ROBOTIZADO
PARTE G/2

18510
D.

PVG-strict, that is, it occurs between different physical volume groups; the `-n redo1.log` option lets you specify the name of the logical volume; and the `-L 4` option allocates 4 megabytes.

NOTE It is important to use the `-M n` and `-c y` options for both redo logs and control files. These options allow the redo log files to be resynchronized by SLVM following a system crash before Oracle recovery proceeds. If these options are not set correctly, you may not be able to continue with database recovery.

If the command is successful, the system will display messages like the following:

```
Logical volume "/dev/vg_ops/redo1.log" has been successfully
created
with character device "/dev/vg_ops/rredo1.log"
Logical volume "/dev/vg_ops/redo1.log" has been successfully
extended
```

Note that the *character* device file name (also called the raw logical volume name) is used by the Oracle DBA in building the RAC database.

Creating Mirrored Logical Volumes for RAC Data Files

```
# lvcreate -m 1 -M n -c n -s g -n system.dbf -L 28 /dev/vg_ops
```

For data files other than the redo logs or control files, choose a mirror consistency policy of "none" by disabling both mirror write caching and mirror consistency recovery. Create logical volumes for use as Oracle data files by using the same options as in the following example:

The `-m 1` option specifies single mirroring; the `-M n` option ensures that mirror write cache recovery is set off; the `-c n` means that mirror consistency recovery is disabled; the `-s g` means that mirroring is PVG-strict, that is, it occurs between different physical volume groups; the `-n system.dbf` option lets you specify the name of the logical volume; and the `-L 28` option allocates 28 megabytes.

If the command is successful, the system will display messages like the following:

18509
D

```
Logical volume "/dev/vg_ops/system.dbf" has been successfully
created
with character device "/dev/vg_ops/rsystem.dbf"
Logical volume "/dev/vg_ops/system.dbf" has been successfully
extended
```

Note that the *character* device file name (also called the raw logical volume name) is used by the Oracle DBA in building the OPS database.

Creating RAC Volume Groups on Disk Arrays

The procedure described in this section assumes that you are using RAID-protected disk arrays and LVM's physical volume links (PV links) to define redundant data paths from each node in the cluster to every logical unit on the array.

On your disk arrays, you should use redundant I/O channels from each node, connecting them to separate controllers on the array. Then you can define alternate links to the LUNs or logical disks you have defined on the array. If you are using SAM, choose the type of disk array you wish to configure, and follow the menus to define alternate links. If you are using LVM commands, specify the links on the command line.

The following example shows how to configure alternate links using LVM commands. The following disk configuration is assumed:

```
8/0.15.0 /dev/dsk/c0t15d0  /* I/O Channel 0 (8/0) SCSI address 15 LUN 0 */
8/0.15.1 /dev/dsk/c0t15d1  /* I/O Channel 0 (8/0) SCSI address 15 LUN 1 */
8/0.15.2 /dev/dsk/c0t15d2  /* I/O Channel 0 (8/0) SCSI address 15 LUN 2 */
8/0.15.3 /dev/dsk/c0t15d3  /* I/O Channel 0 (8/0) SCSI address 15 LUN 3 */
8/0.15.4 /dev/dsk/c0t15d4  /* I/O Channel 0 (8/0) SCSI address 15 LUN 4 */
8/0.15.5 /dev/dsk/c0t15d5  /* I/O Channel 0 (8/0) SCSI address 15 LUN 5 */

10/0.3.0 /dev/dsk/c1t3d0   /* I/O Channel 1 (10/0) SCSI address 3 LUN 0 */
10/0.3.1 /dev/dsk/c1t3d1   /* I/O Channel 1 (10/0) SCSI address 3 LUN 1 */
10/0.3.2 /dev/dsk/c1t3d2   /* I/O Channel 1 (10/0) SCSI address 3 LUN 2 */
10/0.3.3 /dev/dsk/c1t3d3   /* I/O Channel 1 (10/0) SCSI address 3 LUN 3 */
10/0.3.4 /dev/dsk/c1t3d4   /* I/O Channel 1 (10/0) SCSI address 3 LUN 4 */
10/0.3.5 /dev/dsk/c1t3d5   /* I/O Channel 1 (10/0) SCSI address 3 LUN 5 */
```

Assume that the disk array has been configured, and that both the following device files appear for the same LUN (logical disk) when you run the ioscan command:

```
/dev/dsk/c0t15d0
/dev/dsk/c1t3d0
```

13508
D.

Use the following procedure to configure a volume group for this logical disk:

1. First, set up the group directory for vg_ops:

   ```
   # mkdir /dev/vg_ops
   ```

2. Next, create a control file named group in the directory /dev/vg_ops, as follows:

   ```
   # mknod /dev/vg_ops/group c 64 0xhh0000
   ```

   The major number is always 64, and the hexadecimal minor number has the form

   ```
   0xhh0000
   ```

   where *hh* must be unique to the volume group you are creating. Use the next hexadecimal number that is available on your system, after the volume groups that are already configured. Use the following command to display a list of existing volume groups:

   ```
   # ls -l /dev/*/group
   ```

3. Use the pvcreate command on one of the device files associated with the LUN to define the LUN to LVM as a physical volume.

   ```
   # pvcreate -f /dev/rdsk/c0t15d0
   ```

   It is only necessary to do this with *one* of the device file names for the LUN. The -f option is only necessary if the physical volume was previously used in some other volume group.

4. Use the following to create the volume group with the two links:

   ```
   # vgcreate /dev/vg_ops /dev/dsk/c0t15d0 /dev/dsk c1t3d0
   ```

LVM will now recognize the I/O channel represented by /dev/dsk/c0t15d0 as the primary link to the disk; if the primary link fails, LVM will automatically switch to the alternate I/O channel represented by /dev/dsk/c1t3d0. Use the vgextend command to add additional disks to the volume group, specifying the appropriate physical volume name for each PV link.

Repeat the entire procedure for each distinct volume group you wish to create. For ease of system administration, you may wish to use different volume groups to separate logs from data and control files.

18507
D.

NOTE The default maximum number of volume groups in HP-UX is 10. If you intend to create enough new volume groups that the total exceeds ten, you must increase the maxvgs system parameter and then re-build the HP-UX kernel. In SAM, select the `Kernel Configuration` area, then choose `Configurable Parameters`. Maxvgs appears on the list.

Creating Logical Volumes for RAC on Disk Arrays

After you create volume groups and add PV links to them, you define logical volumes for data, logs, and control files. The following are some examples:

```
# lvcreate -n ops1log1.log -L 4 /dev/vg_ops
# lvcreate -n opsctl1.ctl -L 4 /dev/vg_ops
# lvcreate -n system.dbf -L 28 /dev/vg_ops
# lvcreate -n opsdata1.dbf -L 1000 /dev/vg_ops
```

Oracle Demo Database Files

The following set of files is required for the Oracle demo database which you can create during the installation process.

Table 2-1 **Required Oracle File Names for Demo Database**

Logical Volume Name	LV Size (MB)	Raw Logical Volume Path Name	Oracle File Size (MB)*
opsctl1.ctl	32	/dev/vg_ops/ropsctl1.ctl	100
opsctl2.ctl	32	/dev/vg_ops/ropsctl2.ctl	100
opsctl3.ctl	32	/dev/vg_ops/ropsctl3.ctl	100
system.dbf	80	/dev/vg_ops/rsystem.dbf	400
ops1log1.log	10	/dev/vg_ops/rops1log1.log	20
ops1log2.log	10	/dev/vg_ops/rops1log2.log	20
ops1log3.log	10	/dev/vg_ops/rops1log3.log	20

18506
D.

Table 2-1 Required Oracle File Names for Demo Database (Continued)

Logical Volume Name	LV Size (MB)	Raw Logical Volume Path Name	Oracle File Size (MB)*
system.dbf	15	/dev/vg_ops/rsystem1.dbf	400
temp.dbf	4	/dev/vg_ops/rtemp.dbf	100
users.dbf	4	/dev/vg_ops/rusers.dbf	120
tools.dbf	25	/dev/vg_ops/rtools.dbf	15
opsdata1.dbf	15	/dev/vg_ops/ropsdata1.dbf	200
opsdata2.dbf	15	/dev/vg_ops/ropsdata2.dbf	200
opsdata3.dbf	15	/dev/vg_ops/ropsdata3.dbf	200
ops2log1		/dev/vg_ops/rops2log1.log	20
ops2log2		/dev/vg_ops/rops2log2.log	20
ops2log3		/dev/vg_ops/rops2log3.log	20
ops3log1		/dev/vg_ops/rops2log1.log	20
ops3log2		/dev/vg_ops/rops2log2.log	20
ops3log3		/dev/vg_ops/rops2log3.log	20
opsdata1		/dev/vg_ops/ropsdata1.dbf	200
opsdata2		/dev/vg_ops/ropsdata2.dbf	200
opsdata3		/dev/vg_ops/ropsdata3.dbf	200

* The size of the logical volume is larger than the Oracle file size because Oracle needs extra space to allocate a header in addition to the file's actual data capacity.

Create these files if you wish to build the demo database. The three logical volumes at the bottom of the table are included as additional data files, which you can create as needed, supplying the appropriate sizes. If your naming conventions require, you can include the Oracle SID and/or the database name to distinguish files for different instances and different databases. If you are using the ORACLE_BASE directory structure, create symbolic links to the ORACLE_BASE files from the appropriate directory. Example:

```
# ln -s /dev/vg_ops/ropsctl1.ctl
/u01/ORACLE/db001/ctrl01_1.ctl
```

For more information about Oracle directories, refer to the *Oracle Server for HP 9000 Installation and Configuration Guide*. More information about the maximum sizes of data files is found in the Oracle 8i *Reference* (Oracle Part Number A76961-01), Chapter 4, "Database Limits."

After creating these files, set the owner to oracle and the group to dba with a file mode of 660. The logical volumes are now available on the primary node, and the raw logical volume names can now be used by the Oracle DBA.

Displaying the Logical Volume Infrastructure

To display the volume group, use the vgdisplay command:

```
# vgdisplay -v /dev/vg_ops
```

Exporting the Logical Volume Infrastructure

Before the Oracle volume groups can be shared, their configuration data must be exported to other nodes in the cluster. This is done either in SAM or by using HP-UX commands, as shown in the following sections.

NOTE A volume group using an HP High Availability Disk Array cannot be designated as a cluster lock volume group. You need to use another volume group containing an independent disk for the cluster lock.

Exporting with SAM

In SAM, choose Disks and File Systems, then choose Volume Groups. Open the Actions menu and choose Distribute. Follow the prompts in the dialog box to select the volume group that is to be distributed to one or more additional nodes. Enter the name of each node that is to receive

18504
D.

the volume group and select Add. When the list is complete, press OK. SAM automatically configures the volume group for use on the other nodes.

NOTE In distributing the volume group, SAM does not import PV links on other nodes. Therefore, to complete the distribution, you use the vgextend command on all nodes to add the alternate PV links to the volume group you have distributed. Use the vgdisplay command to verify the existence of the correct groups of PV links on all nodes.

Exporting with LVM Commands

Use the following commands to set up the same volume group on another cluster node. In this example, the commands set up a new volume group on a system known as ftsys10. This volume group holds the same physical volume that was created on a configuration node known as ftsys9.

To set up the volume group on ftsys10 (and other nodes), use the following steps:

1. On ftsys9, copy the mapping of the volume group to a specified file.

   ```
   # vgexport -s -p -m /tmp/vg_ops.map  /dev/vg_ops
   ```

2. Still on ftsys9, copy the map file to ftsys10 (and to additional nodes as necessary.)

   ```
   # rcp /tmp/vg_ops.map ftsys10:/tmp/vg_ops.map
   ```

3. On ftsys10 (and other nodes, as necessary), create the volume group directory and the control file named group:

   ```
   # mkdir /dev/vg_ops
   # mknod /dev/vg_ops/group c 64 0xhh0000
   ```

   For the group file, the major number is always 64, and the hexadecimal minor number has the form

   0x*hh*0000

   where *hh* must be unique to the volume group you are creating. If possible, use the same number as on ftsys9. Use the following command to display a list of existing volume groups:

   ```
   # ls -l /dev/*/group
   ```

LPS03
D

4. Import the volume group data using the map file from node `ftsys9`. On node `ftsys10` (and other nodes, as necessary), enter:

```
# vgimport -s -m /tmp/vg_ops.map /dev/vg_ops \
  /dev/dsk/c3t12d0 /dev/dsk/c4t3d0
```

CAUTION

Special care should be used when importing volume groups that use PV links with the A3231A (Model 10) and A3232A (Model 20) disk arrays. On these arrays, each disk device file name represents a link to a disk through a separate controller, and all primary links must use the same controller. Even though the device file names may be different on different nodes, it is very important to specify the PV links in the `vgimport` command in *the same controller order on all nodes* that they were specified in creating the volume group on the configuration node. Failure to do this with PV links results in poor performance.

Installing Oracle Real Application Clusters

Before installing the Oracle Real Application Cluster software, make sure the cluster is running. Log in as the `oracle` user on one node and then use the Oracle `installer` to install Oracle software and to build the correct Oracle runtime executables. The Oracle `installer` also copies the executables to the other nodes in the cluster. Select the following installation option to install RAC software and to create the demo database:

```
COMPLETE SOFTWARE/DATABASE FRESH INSTALL
```

Refer to your Oracle installation documentation for Oracle installation details. As part of this installation, the Oracle installer builds the Oracle demo database on the primary node, using the character (raw) device file names for the logical volumes created earlier. For the demo database, create fourteen logical volumes as shown in the table "Required Oracle File Names for Demo Database" earlier in this chapter. As the installer prompts for database file names, enter the pathnames of the raw logical volumes instead of using the defaults. If you do not wish to install the demo database, select:

```
SOFTWARE INSTALL ONLY
```

In this case, create an appropriate number of raw logical volumes to build your development or production system. Be sure to create enough log files for all instances.

1850d
D

Cluster Configuration ASCII File

The following is an example of an ASCII configuration file generated with the cmquerycl command using the -w full option on a system with ServiceGuard Extension for RAC. The OPS_VOLUME_GROUP parameters appear at the end of the file.

```
# **********************************************************************
# ********* HIGH AVAILABILITY CLUSTER CONFIGURATION FILE ***************
# ***** For complete details about cluster parameters and how to    ****
# ***** set them, consult the cmquerycl(1m) manpage or your manual. ****
# **********************************************************************

# Enter a name for this cluster.  This name will be used to identify the
# cluster when viewing or manipulating it.

CLUSTER_NAME            lpcluster

# Cluster Lock Device Parameters.  This is the volume group that
# holds the cluster lock which is used to break a cluster formation
# tie.  This volume group should not be used by any other cluster
# as cluster lock device.

FIRST_CLUSTER_LOCK_VG   /dev/vg01

# Definition of nodes in the cluster.
# Repeat node definitions as necessary for additional nodes.

NODE_NAME               ftsys9
  NETWORK_INTERFACE     lan0
    HEARTBEAT_IP        15.13.171.32
  NETWORK_INTERFACE     lan3
    HEARTBEAT_IP        192.6.7.3
  NETWORK_INTERFACE     lan4
  NETWORK_INTERFACE     lan1
    HEARTBEAT_IP        192.6.143.10
  FIRST_CLUSTER_LOCK_PV /dev/dsk/c1t2d0

# List of serial device file names
# For example:
# SERIAL_DEVICE_FILE    /dev/tty0p0

# Primary Network Interfaces on Bridged Net 1: lan0.
#   Warning: There are no standby network interfaces on bridged net 1.
```

18500

```
# Primary Network Interfaces on Bridged Net 2: lan3.
#   Possible standby Network Interfaces on Bridged Net 2: lan4.
# Primary Network Interfaces on Bridged Net 3: lan1.
#   Warning: There are no standby network interfaces on bridged net 3.

# Cluster Timing Parameters (microseconds).

# The NODE_TIMEOUT parameter defaults to 2000000 (2 seconds).
# This default setting yields the fastest cluster reformations.
# However, the use of the default value increases the potential
# for spurious reformations due to momentary system hangs or
# network load spikes.
# For a significant portion of installations, a setting of
# 5000000 to 8000000 (5 to 8 seconds) is more appropriate.
# The maximum recommended value for NODE_TIMEOUT is 3000000
# (30 seconds).


HEARTBEAT_INTERVAL        1000000
NODE_TIMEOUT              2000000

# Configuration/Reconfiguration Timing Parameters (microseconds).

AUTO_START_TIMEOUT        600000000
NETWORK_POLLING_INTERVAL 2000000

# Package Configuration Parameters.
# Enter the maximum number of packages which will be configured in the cluster.
# You can not add packages beyond this limit.
# This parameter is required.

MAX_CONFIGURED_PACKAGES            10

# List of cluster aware Volume Groups. These volume groups
# will be used by clustered applications via the vgchange -a e command.
# For example:
# VOLUME_GROUP            /dev/vgdatabase
# VOLUME_GROUP            /dev/vg02

VOLUME_GROUP              /dev/vg01
VOLUME_GROUP              /dev/vg02

# List of OPS Volume Groups.

# Formerly known as DLM Volume Groups, these volume groups
# will be used by OPS cluster applications via
# the vgchange -a s command. (Note: the name DLM_VOLUME_GROUP
```

18.499

```
# is also still supported for compatibility with earlier versions.)
# For example:
# OPS_VOLUME GROUP              /dev/vgdatabase.
# OPS_VOLUME_GROUP              /dev/vg02.

OPS_VOLUME_GROUP /dev/vg_ops
```

18.498
A.

Creating a Storage Infrastructure with CVM

In addition to configuring the cluster, you create the appropriate logical volume infrastructure to provide access to data from different nodes. This is done with Logical Volume Manager (LVM), VERITAS Volume Manager (VxVM), or VERITAS Cluster Volume Manager (CVM). LVM and VxVM configuration are done before cluster configuration, and CVM configuration is done after cluster configuration.

This section shows how to configure storage using the VERITAS Cluster Volume Manager (CVM). The examples show how to configure RAC disk groups, but you can also create CVM disk groups for non-RAC use. For more information, including details about configuration of plexes (mirrors), multi-pathing, and RAID, refer to the HP-UX documentation for the VERITAS Volume Manager.

Initializing the VERITAS Volume Manager

If you are about to create disk groups for the first time, you need to initialize the Volume Manager. This is done by creating a disk group known as *rootdg* that contains at least one disk. Use the following command after installing CVM on each node:

```
# vxinstall
```

This displays a menu-driven program that steps you through the CVM initialization sequence. From the main menu, choose the "Custom" option, and specify the disk you wish to include in *rootdg*.

IMPORTANT

The *rootdg* in the VERITAS Volume Manager is not the same as the HP-UX root disk, and *cannot* be used for the HP-UX root file system (/). Note also that *rootdg* cannot be used for shared storage. However, *rootdg* can be used for other local filesystems (e.g., /export/home), so it need not be wasted. Creating a *rootdg* disk group is only necessary the first time you use the Volume Manager.

13497
A.

Preparing the Cluster for Use with CVM

In order to use the VERITAS Cluster Volume Manager (CVM), you need a cluster that is running with a special CVM package. This means that the cluster must already be configured and running before you create disk groups.

NOTE Cluster configuration is described in the previous section.

To prepare the cluster for CVM disk group configuration, you need to set `MAX_CONFIGURED_PACKAGES` to 1 or greater in the cluster ASCII configuration file, and ensure that only one heartbeat subnet is configured. Then use the following command, which creates the special package that communicates cluster information to CVM:

```
# cmapplyconf -P /etc/cmcluster/cvm/VxVM-CVM-pkg.conf
```

WARNING **This file should never be edited.**

After this command completes successfully, you can start the cluster and create disk groups for shared use as described in the following sections. When you start the cluster, it will now run with a special **system multi-node package** named `VxVM-CVM-pkg`, which is on all nodes. This package is shown in the following output of the `cmviewcl -v` command:

```
CLUSTER        STATUS
bowls          up

NODE           STATUS         STATE
spare          up             running
split          up             running
strike         up             running

SYSTEM_MULTI_NODE_PACKAGES:

PACKAGE        STATUS         STATE
VxVM-CVM-pkg   up             running
```

Starting the Cluster and Identifying the Master Node

Run the cluster, which will activate the special CVM package:

```
# cmruncl
```

When CVM starts up, it selects a master node, and this is the node from which you must issue the disk group configuration commands. To determine the master node, issue the following command from each node in the cluster:

```
# vxdctl -c mode
```

One node will identify itself as the master. Create disk groups from this node.

Converting Disks from LVM to CVM

You can use the `vxvmconvert` utility to convert LVM volume groups into CVM disk groups. Before you can do this, the volume group must be deactivated, which means that any package that uses the volume group must be halted. The procedure is described in Appendix F.

Initializing Disks for CVM

You need to initialize the physical disks that will be employed in CVM disk groups. If a physical disk has been previously used with LVM, you should use the `pvremove` command to delete the LVM header data from all the disks in the volume group (this is not necessary if you have not previously used the disk with LVM).

To initialize a disk for CVM, log on to the master node, then use the `vxdiskadm` program to initialize multiple disks, or use the `vxdisksetup` command to initialize one disk at a time, as in the following example:

```
# /usr/lib/vxvm/bin/vxdisksetup -i /dev/dsk/c0t3d2
```

Creating Disk Groups for RAC

Use the `vxdg` command to create disk groups. Use the -s option to specify shared mode, as in the following example:

```
# vxdg -s init ops_dg c0t3d2
```

Verify the configuration with the following command:

```
# vxdg list
NAME          STATE                ID

rootdg         enabled            971995699.1025.node1
ops_dg         enabled,shared     972078742.1084.node2
```

Creating Volumes

Use the vxassist command to create logical volumes. The following is an example:

```
# vxassist -g log_files make ops_dg 1024m
```

This command creates a 1024 MB volume named *log_files* in a disk group named ops_dg. The volume can be referenced with the block device file /dev/vx/dsk/ops_dg/log_files or the raw (character) device file /dev/vx/rdsk/ops_dg/log_files.

Verify the configuration with the following command:

```
# vxdg list
```

IMPORTANT After creating these files, use the vxedit command to change the ownership of the raw volume files to oracle and the group membership to dba, and to change the permissions to 660. Example:

```
# cd /dev/vx/rdsk/ops_dg
# vxedit -g ops_dg set user=oracle *
# vxedit -g ops_dg set group=dba *
# vxedit -g ops_dg set mode=660 *
```

The logical volumes are now available on the primary node, and the raw logical volume names can now be used by the Oracle DBA.

18.494
A.

Mirror Detachment Policies with CVM

The default CVM disk mirror detachment policy is 'global', which means that as soon as one node cannot see a specific mirror copy (plex), all nodes cannot see it as well. The alternate policy is 'local', which means that if one node cannot see a specific mirror copy, then CVM will deactivate access to the volume for that node only.

This policy can be re-set on a disk group basis by using the `vxedit` command, as follows:

```
# vxedit set diskdetpolicy=[global|local] <DiskGroupName>
```

NOTE The specific commands for creating mirrored and multi-path storage using CVM are described in the HP-UX documentation for the VERITAS Volume Manager.

Oracle Demo Database Files

The following set of volumes is required for the Oracle demo database which you can create during the installation process.

Table 2-2 Required Oracle File Names for Demo Database

Volume Name	Size (MB)	Raw Device File Name	Oracle File Size (MB)
cntrl1		/dev/vx/rdsk/ops_dg/opsctl1.ctl	1
cntrl2		/dev/vx/rdsk/ops_dg/opsctl2.ctl	1
opsctl1.ctl		/dev/vx/rdsk/ops_dg/opsctl1.ctl	100
opsctl2.ctl		/dev/vx/rdsk/ops_dg/opsctl2.ctl	100
opsctl3.ctl		/dev/vx/rdsk/ops_dg/opsctl3.ctl	100
system.dbf		/dev/vx/rdsk/ops_dg/system.dbf	400
ops1log1.log		/dev/vx/rdsk/ops_dg/ops1log1.log	20
ops1log2.log		/dev/vx/rdsk/ops_dg/ops1log2.log	20

18.493
A.

Table 2-2 Required Oracle File Names for Demo Database (Continued)

Volume Name	Size (MB)	Raw Device File Name	Oracle File Size (MB)
ops1log3.log		/dev/vx/rdsk/ops_dg/ops1log3.log	20
system.dbf		/dev/vx/rdsk/ops_dg/system1.dbf	400
temp.dbf		/dev/vx/rdsk/ops_dg/temp.dbf	100
users.dbf		/dev/vx/rdsk/ops_dg/users.dbf	120
tools.dbf		/dev/vx/rdsk/ops_dg/tools.dbf	15
opsdata1.dbf		/dev/vx/rdsk/ops_dg/opsdata1.dbf	200
opsdata2.dbf		/dev/vx/rdsk/ops_dg/opsdata2.dbf	200
opsdata3.dbf		/dev/vx/rdsk/ops_dg/opsdata3.dbf	200
ops2log1		/dev/vx/rdsk/ops_dg/ops2log1.log	20
ops2log2		/dev/vx/rdsk/ops_dg/ops2log2.log	20
ops2log3		/dev/vx/rdsk/ops_dg/ops2log3.log	20
ops3log1		/dev/vx/rdsk/ops_dg/ops2log1.log	20
ops3log2		/dev/vx/rdsk/ops_dg/ops2log2.log	20
ops3log3		/dev/vx/rdsk/ops_dg/ops2log3.log	20
opsdata1		/dev/vx/rdsk/ops_dg/opsdata1.dbf	200
opsdata2		/dev/vx/rdsk/ops_dg/opsdata2.dbf	200
opsdata3		/dev/vx/rdsk/ops_dg/opsdata3.dbf	200

Create these files if you wish to build the demo database. The three logical volumes at the bottom of the table are included as additional data files, which you can create as needed, supplying the appropriate sizes. If your naming conventions require, you can include the Oracle SID and/or the database name to distinguish files for different instances and different databases. If you are using the `ORACLE_BASE` directory structure, create symbolic links to the `ORACLE_BASE` files from the appropriate directory.

18.492
A.

Example, Oracle8:

```
# ln -s /dev/vg_ops/ropsctl1.ctl \
    /u01/ORACLE/db001/ctrl01_1.ctl
```

Example, Oracle9:

1. Create an AASCII file, and define the path for each database object.

   ```
   # control1= /dev/vg_ops/ropsctl1.ctl \
       /u01/ORACLE/db001/ctrl01_1.ctl
   ```

2. Set the following environment variable where *filename* is the name of the ASCII file created.

   ```
   # setenv DBCA_RAW_CONFIG filename
   ```

For more information about Oracle directories, refer to the *Oracle Server for HP 9000 Installation and Configuration Guide*. More information about the maximum sizes of data files is found in the Oracle *Reference*.

Adding Disk Groups to the Cluster Configuration

After creating units of CVM storage with VxVM commands, you need to specify the disk groups in each package configuration ASCII file. Use one DISK_GROUP parameter for each CVM disk group the package will use. You also need to identify the CVM disk groups, file systems, logical volumes, and mount options in the package control script. The package configuration process is described in detail in Chapter 6.

13.491
CA.

Configuring Packages that Access the Oracle Database

You can also use packages to start up applications that access the Oracle RAC instances. If an application is intended to fail over among cluster nodes, then you must set it up as a distinct package, separate from the package that starts and stops the RAC instance. Use the following procedures for packages that contain applications which access the Oracle database:

1. In the ASCII package configuration file, set the `AUTO_RUN` parameter to `NO`, or if you are using SAM to configure packages, set `Automatic Switching` to `Disabled`. This keeps the package from starting up immediately when the node joins the cluster, and before Oracle RAC is running.

2. You can then manually start the package using the `cmmodpkg -e` *`packagename`* command after the RAC instance is started. Alternatively, you can choose to automate the process of package activation by writing your own script, and copying it to all nodes that can run the package. This script should contain the `cmmodpkg -e` command and activate the package after Oracle RAC and the cluster manager have started.

18.490
CA.

3 Maintenance and Troubleshooting

This chapter includes information about carrying out routine maintenance on an Oracle Parallel Server configuration. As presented here, these tasks differ in some details from the similar tasks described in Chapter 7 of *Managing MC/ServiceGuard*.

Tasks include:

- Reviewing Cluster and Package States with the cmviewcl Command
- Managing the Shared Storage
- Removing ServiceGuard Extension for RAC from a System
- Monitoring Hardware
- Adding Disk Hardware
- Replacing Disks
- Replacement of I/O Cards
- Replacement of LAN Cards

Reviewing Cluster and Package States with the cmviewcl Command

A cluster or its component nodes may be in several different states at different points in time. Status information for clusters, packages and other cluster elements is shown in the output of the cmviewcl command and in some displays in SAM. This section explains the meaning of many of the common conditions the cluster or package may be in.

Information about cluster status is stored in the status database, which is maintained on each individual node in the cluster. You can display information contained in this database by issuing the cmviewcl command:

```
# cmviewcl -v
```

The command when issued with the -v option displays information about the whole cluster. See the man page for a detailed description of other cmviewcl options.

TIP Some commands take longer to complete in large configurations. In particular, you can expect ServiceGuard's CPU utilization to increase during cmviewcl -v as the number of packages and services increases.

You can also specify that the output should be formatted as it was in a specific earlier release by using the -r option indicating the release format you wish. Example:

```
# cmviewcl -r A.11.12
```

See the man page for a detailed description of other cmviewcl options.

Types of Cluster and Package States

A cluster or its component nodes may be in several different states at different points in time. The following sections describe many of the common conditions the cluster or package may be in.

Cluster Status

The *status* of a cluster may be one of the following:

13488
A.

- `Up`. At least one node has a running cluster daemon, and reconfiguration is not taking place.
- `Down`. No cluster daemons are running on any cluster node.
- `Starting`. The cluster is in the process of determining its active membership. At least one cluster daemon is running.
- `Unknown`. The node on which the `cmviewcl` command is issued cannot communicate with other nodes in the cluster.

Node Status and State

The *status* of a node is either up (*active* as a member of the cluster) or down (*inactive* in the cluster), depending on whether its cluster daemon is running or not. Note that a node might be down from the cluster perspective, but still up and running HP-UX.

A node may also be in one of the following states:

- `Failed`. A node never sees itself in this state. Other active members of the cluster will see a node in this state if that node was in an active cluster, but is no longer, and is not halted.
- `Reforming`. A node is in this state when the cluster is re-forming. The node is currently running the protocols which ensure that all nodes agree to the new membership of an active cluster. If agreement is reached, the status database is updated to reflect the new cluster membership.
- `Running`. A node in this state has completed all required activity for the last re-formation and is operating normally.
- `Halted`. A node never sees itself in this state. Other nodes will see it in this state after the node has gracefully left the active cluster, for instance with a `cmhaltnode` command.
- `Unknown`. A node never sees itself in this state. Other nodes assign a node this state if it has never been an active cluster member.

Package Status and State

The *status* of a package can be one of the following:

- `Up`. The package control script is active.
- `Down`. The package control script is not active.
- `Unknown`.

13.437
A.

The *state* of the package can be one of the following:

- `Starting`. The start instructions in the control script are being run.
- `Running`. Services are active and being monitored.
- `Halting`. The halt instructions in the control script are being run.

Package Switching Attributes

Packages also have the following switching attributes:

- `Package Switching`. Enabled means that the package can switch to another node in the event of failure.
- `Switching Enabled for a Node`. Enabled means that the package can switch to the referenced node. Disabled means that the package cannot switch to the specified node until the node is enabled for the package using the `cmmodpkg` command.

  Every package is marked Enabled or Disabled for each node that is either a primary or adoptive node for the package.

Status of Group Membership

The state of the cluster for Oracle RAC is one of the following:

- `Up`. Services are active and being monitored. The membership appears in the output of `cmviewcl -l` group.
- `Down`. The cluster is halted and GMS services have been stopped. The membership does not appear in the output of the `cmviewcl -l` group.

The following is an example of the group membership output shown in the cmviewcl command:

```
# cmviewcl -l group

GROUP        MEMBER      PID      MEMBER_NODE
DGop            1        10394      comanche
                0        10499      chinook
DBOP            1        10501      comanche
                0        10396      chinook
DAALL_DB        0        10396      comanche
                1        10501      chinook
IGOPALL         2        10423      comanche
                1        10528      chinook
```

18.486

where the `cmviewcl` output values are:

`GROUP`	the name of a configured group
`MEMBER`	the ID number of a member of a group
`PID`	the Process ID of the group member
`MEMBER_NODE`	the Node on which the group member is running

Service Status

Services have only status, as follows:

- `Up`. The service is being monitored.
- `Down`. The service is not running. It may have halted or failed.
- `Uninitialized`. The service is included in the package configuration, but it was not started with a run command in the control script.
- `Unknown`.

Network Status

The network interfaces have only status, as follows:

- `Up`.
- `Down`.
- `Unknown`. We cannot determine whether the interface is up or down. This can happen when the cluster is down. A standby interface has this status.

Serial Line Status

The serial line has only status, as follows:

- `Up`. Heartbeats are received over the serial line.
- `Down`. Heartbeat has not been received over the serial line within 2 times the `NODE_TIMEOUT` value.
- `Recovering`. A corrupt message was received on the serial line, and the line is in the process of resynchronizing.
- `Unknown`. We cannot determine whether the serial line is up or down. This can happen when the remote node is down.

18.485
A.

Failover and Failback Policies

Packages can be configured with one of two values for the FAILOVER_POLICY parameter:

- CONFIGURED_NODE. The package fails over to the next node in the node list in the package configuration file.
- MIN_PACKAGE_NODE. The package fails over to the node in the cluster with the fewest running packages on it.

Packages can also be configured with one of two values for the FAILBACK_POLICY parameter:

- AUTOMATIC. With this setting, a package, following a failover, returns to its primary node when the primary node becomes available again.
- MANUAL. With this setting, a package, following a failover, must be moved back to its original node by a system administrator.

Failover and failback policies are displayed in the output of the cmviewcl -v command.

Examples of Cluster and Package States

The following sample output from the cmviewcl -v command shows status for the cluster in the sample configuration.

Normal Running Status

Everything is running normally; both nodes in a two-node cluster are running, and each Oracle RAC instance package is running as well. The only packages running are Oracle RAC instance packages.

```
CLUSTER        STATUS
example        up
  NODE           STATUS        STATE
  ftsys9         up            running

    Network_Parameters:
    INTERFACE    STATUS        PATH          NAME
    PRIMARY      up            56/36.1       lan0
    STANDBY      up            60/6          lan1

    PACKAGE      STATUS        STATE         AUTO_RUN      NODE
    ops_pkg1     up            running       disabled      ftsys9
```

18.484

```
          Policy_Parameters:
          POLICY_NAME       CONFIGURED_VALUE
          Start         configured_node
          Failback      manual

          Node_Switching_Parameters:
          NODE_TYPE     STATUS        SWITCHING     NAME          (cur
rent)
          Primary       up            enabled       ftsys9

   NODE          STATUS        STATE
   ftsys10       up            running

      Network_Parameters:
      INTERFACE     STATUS        PATH          NAME
      PRIMARY       up            28.1          lan0
      STANDBY       up            32.1          lan1

      PACKAGE       STATUS        STATE         AUTO_RUN      NODE
      ops_pkg2      up            running       disabled      ftsys1
0

         Policy_Parameters:
         POLICY_NAME       CONFIGURED_VALUE
         Start         configured_node
         Failback      manual

         Node_Switching_Parameters:
         NODE_TYPE     STATUS        SWITCHING     NAME          (cur
rent)
         Primary       up            enabled       ftsys10
         Alternate     up            enabled       ftsys9
```

Quorum Server Status

If the cluster is using a quorum server for tie-breaking services, the display shows the server name, state and status following the entry for each node, as in the following excerpt from the output of `cmviewcl -v`:

```
CLUSTER       STATUS
example       up

   NODE          STATUS        STATE
   ftsys9        up            running
```

```
    Quorum Server Status:
    NAME                 STATUS        STATE
    lp-qs                up            running
...

    NODE         STATUS        STATE
    ftsys10      up            running

    Quorum Server Status:
    NAME                 STATUS        STATE
    lp-qs                up            running
```

CVM Package Status

If the cluster is using the VERITAS Cluster Volume Manager for disk storage, the system multi-node package CVM-VxVM-pkg must be running on all active nodes for applications to be able to access CVM disk groups. This package is shown in the following output of the cmviewcl command:

```
CLUSTER       STATUS
example       up

NODE          STATUS        STATE
ftsys8        down          halted
ftsys9        up            running

SYSTEM_MULTI_NODE_PACKAGES:

PACKAGE       STATUS        STATE
VxVM-CVM-pkg  up            running
```

When you use the -v option, the display shows the system multi-node package associated with each active node in the cluster, as in the following:

```
SYSTEM_MULTI_NODE_PACKAGES:

    PACKAGE       STATUS        STATE
    VxVM-CVM-pkg  up            running

     NODE          STATUS        STATE
     ftsys8        down          halted

     NODE          STATUS        STATE
     ftsys9        up            running
```

18.482
A.

```
      Script_Parameters:
      ITEM         STATUS    MAX_RESTARTS   RESTARTS    NAME
      Service      up                   0          0
VxVM-CVM-pkg.srv
```

Status After Moving the Package to Another Node

After issuing the following command:

```
# cmrunpkg -n ftsys9 pkg2
```

the output of the `cmviewcl -v` command is as follows:

```
CLUSTER        STATUS
example        up

  NODE           STATUS         STATE
  ftsys9         up             running

    Network_Parameters:
    INTERFACE     STATUS         PATH           NAME
    PRIMARY       up             56/36.1        lan0
    STANDBY       up             60/6           lan1

    PACKAGE       STATUS         STATE          AUTO_RUN       NODE
    pkg1          up             running        enabled        ftsys9

      Policy_Parameters:
      POLICY_NAME        CONFIGURED_VALUE
      Failover      min_package_node
      Failback      manual

      Script_Parameters:
      ITEM          STATUS   MAX_RESTARTS    RESTARTS        NAME
      Service       up            0           0         service1
      Subnet        up            0           0         15.13.168.0
      Resource      up                                /example/float

      Node_Switching_Parameters:
      NODE_TYPE     STATUS   SWITCHING      NAME
      Primary       up       enabled        ftsys9         (current)
      Alternate     up       enabled        ftsys10

    PACKAGE       STATUS         STATE          AUTO_RUN       NODE
    pkg2          up             running        disabled       ftsys9
```

18 481
CA.

```
    Policy_Parameters:
    POLICY_NAME       CONFIGURED_VALUE
    Failover      min_package_node
    Failback      manual

    Script_Parameters:
    ITEM        STATUS    NAME        MAX_RESTARTS     RESTARTS
    Service       up     service2.1      0               0
    Subnet        up     15.13.168.0     0               0

    Node_Switching_Parameters:
    NODE_TYPE     STATUS     SWITCHING     NAME
    Primary       up         enabled       ftsys10
    Alternate     up         enabled       ftsys9        (current)

NODE          STATUS         STATE
ftsys10       up             running

  Network_Parameters:
  INTERFACE     STATUS        PATH          NAME
  PRIMARY         up            28.1          lan0
  STANDBY         up            32.1          lan1
```

Now pkg2 is running on node ftsys9. Note that it is still disabled from switching.

Status After Package Switching is Enabled

The following command changes package status back to Package Switching Enabled:

```
# cmmodpkg -e pkg2
```

The output of the cmviewcl command is now as follows:

```
CLUSTER       STATUS
example       up

  NODE          STATUS        STATE
  ftsys9        up            running

    PACKAGE       STATUS        STATE         AUTO_RUN      NODE
    pkg1          up            running       enabled       ftsys9
    pkg2          up            running       enabled       ftsys9
```

18.480
A.

```
  NODE          STATUS        STATE
  ftsys10       up            running
```

Both packages are now running on ftsys9 and pkg2 is enabled for switching. Ftsys10 is running the daemon and no packages are running on ftsys10.

Status After Halting a Node

After halting `ftsys10`, with the following command:

```
# cmhaltnode  ftsys10
```

the output of `cmviewcl` is as follows on `ftsys9`:

```
CLUSTER       STATUS
example       up

  NODE          STATUS        STATE
  ftsys9        up            running

    PACKAGE       STATUS        STATE         AUTO_RUN      NODE
    pkg1          up            running       enabled       ftsys9
    pkg2          up            running       enabled       ftsys9

  NODE          STATUS        STATE
  ftsys10       down          halted
```

This output is seen on both `ftsys9` and `ftsys10`.

Viewing RS232 Status

If you are using a serial (RS232) line as a heartbeat connection, you will see a list of configured RS232 device files in the output of the `cmviewcl -v` command. The following shows normal running status:

```
CLUSTER       STATUS
example       up
  NODE          STATUS        STATE
  ftsys9        up            running

    Network_Parameters:
    INTERFACE     STATUS        PATH          NAME
    PRIMARY       up            56/36.1       lan0
```

18.479

```
        Serial_Heartbeat:
        DEVICE_FILE_NAME      STATUS    CONNECTED_TO:
        /dev/tty0p0            up         ftsys10      /dev/tty0p0


      NODE          STATUS        STATE
      ftsys10        up            running

        Network_Parameters:
        INTERFACE     STATUS        PATH          NAME
        PRIMARY       up            28.1          lan0

        Serial_Heartbeat:
        DEVICE_FILE_NAME      STATUS     CONNECTED_TO:
        /dev/tty0p0             up         ftsys9       /dev/tty0p0
```

The following shows status when the serial line is not working:

```
CLUSTER       STATUS
example       up
  NODE          STATUS        STATE
  ftsys9        up            running

      Network_Parameters:
      INTERFACE     STATUS        PATH          NAME
      PRIMARY       up            56/36.1       lan0

      Serial_Heartbeat:
      DEVICE_FILE_NAME      STATUS        CONNECTED_TO:
      /dev/tty0p0            down            ftsys10   /dev/tty0p0


  NODE          STATUS        STATE
  ftsys10        up            running

      Network_Parameters:
      INTERFACE     STATUS        PATH          NAME
      PRIMARY       up            28.1          lan0

      Serial_Heartbeat:
      DEVICE_FILE_NAME      STATUS     CONNECTED_TO:
      /dev/tty0p0            down        ftsys9       /dev/tty0p0
```

18.478

Viewing Data on Unowned Packages

The following example shows packages that are currently unowned, that is, not running on any configured node. Information on monitored resources is provided for each node on which the package can run; this allows you to identify the cause of a failure and decide where to start the package up again.

```
UNOWNED_PACKAGES

PACKAGE       STATUS        STATE         AUTO_RUN      NODE
PKG3          down          halted        enabled       unowned

Policy_Parameters:
POLICY_NAME       CONFIGURED_VALUE
Failover          min_package_node
Failback          automatic

Script_Parameters:
ITEM        STATUS    NODE_NAME     NAME
Resource    up        manx          /resource/random
Subnet      up        manx          192.8.15.0
Resource    up        burmese       /resource/random
Subnet      up        burmese       192.8.15.0
Resource    up        tabby         /resource/random
Subnet      up        tabby         192.8.15.0
Resource    up        persian       /resource/random
Subnet      up        persian       192.8.15.0

Node_Switching_Parameters:
NODE_TYPE     STATUS        SWITCHING     NAME
Primary       up            enabled       manx
Alternate     up            enabled       burmese
Alternate     up            enabled       tabby
Alternate     up            enabled       persian
```

18.477
A.

ANEXO UNIDADE DE BACKUP ROBOTIZADO
PARTE G/3

Managing the Shared Storage

Making LVM Volume Groups Shareable

Normally, volume groups are marked to be activated in shared mode when they are listed with the `OPS_VOLUME_GROUP` parameter in the cluster configuration file or in SAM. However, in some cases you may want to manually make a volume group sharable. For example, if you wish to add a new shared volume group without shutting down the cluster, you can use the manual method to do it online. However, when convenient, it's a good practice to bring down the cluster and reconfigure it to include the new volume group.

1. Use the vgchange command on each node to ensure that the volume group to be shared is currently inactive on all nodes. Example:

```
# vgchange -a n /dev/vg_ops
```

2. On the configuration node, use the vgchange command to make the volume group shareable by members of the cluster:

```
# vgchange -S y -c y /dev/vg_ops
```

This command is issued from the configuration node only, and the cluster must be running on all nodes for the command to succeed. Note that both the `-S` and the `-c` options are specified. The `-S y` option makes the volume group shareable, and the `-c y` option causes the cluster id to be written out to all the disks in the volume group. In effect, this command specifies the cluster to which a node must belong in order to obtain shared access to the volume group.

Making a Volume Group Unshareable

If you wish to unmark a previously marked shared volume group:

1. Remove the volume group name from the ASCII cluster configuration file.

2. Enter the following command:

```
# vgchange -S n -c n /dev/volumegroup
```

The above example marks the volume group as non-shared and not associated with a cluster.

18.475
A-

Activating an LVM Volume Group in Shared Mode

Activation and deactivation of shared volume groups is normally done through a control script. If you need to perform activation from the command line, you can issue the following command from each node to activate the volume group in shared mode. (The node on which you first enter the command becomes the server node.)

```
# vgchange -a s -p /dev/vg_ops
```

The following message is displayed:

```
Activated volume group in shared mode.
This node is the Server.
```

When the same command is entered on the second node, the following message is displayed:

```
Activated volume group in shared mode.
This node is a Client.
```

NOTE Do *not* share volume groups that are not part of the RAC configuration.

Deactivating a Shared Volume Group

Issue the following command from each node to deactivate the shared volume group:

```
# vgchange -a n /dev/vg_ops
```

Remember that volume groups remain shareable even when nodes enter and leave the cluster.

NOTE If you wish to change the capacity of a volume group at a later time, you must deactivate and unshare the volume group first. If you add disks, you must specify the appropriate physical volume group name and make sure the /etc/lvmpvg file is correctly updated on both nodes.

Making Changes to Shared Volume Groups

You may need to change the volume group configuration of RAC shared logical volumes to add capacity to the data files or to add log files. No configuration changes are allowed on shared LVM volume groups while they are activated. The volume group must be deactivated first on all nodes, and marked as non-shareable. Use the following procedure (examples assume the volume group is being shared by node 1 and node 2, and they use the volume group vg_ops):

1. Ensure that the Oracle RAC database is not active on either node.
2. From node 2, use the vgchange command to deactivate the volume group:
   ```
   # vgchange -a n /dev/vg_ops
   ```
3. From node 2, use the vgexport command to export the volume group:
   ```
   # vgexport -m /tmp/vg_ops.map.old /dev/vg_ops
   ```
   This dissociates the volume group from node 2.
4. From node 1, use the vgchange command to deactivate the volume group:
   ```
   # vgchange -a n /dev/vg_ops
   ```
5. Use the vgchange command to mark the volume group as unshareable:
   ```
   # vgchange -S n -c n /dev/vg_ops
   ```
6. Prior to making configuration changes, activate the volume group in normal (non-shared) mode:
   ```
   # vgchange -a y /dev/vg_ops
   ```
7. Use normal LVM commands to make the needed changes. Be sure to set the raw logical volume device file's owner to oracle and group to dba, with a mode of 660.
8. Next, still from node 1, deactivate the volume group:
   ```
   # vgchange -a n /dev/vg_ops
   ```
9. Use the vgexport command with the options shown in the example to create a new map file:
   ```
   # vgexport -p -m /tmp/vg_ops.map /dev/vg_ops
   ```

Make a copy of /etc/lvmpvg in /tmp/lvmpvg, then copy the file to /tmp/lvmpvg on node 2. Copy the file /tmp/vg_ops.map to node 2.

10. Use the following command to make the volume group shareable by the entire cluster again:

```
# vgchange -S y -c y /dev/vg_ops
```

11. On node 2, issue the following command:

```
# mkdir /dev/vg_ops
```

12. Create a control file named group in the directory /dev/vg_ops, as in the following:

```
# mknod /dev/vg_ops/group c 64 0xhh0000
```

The major number is always 64, and the hexadecimal minor number has the form

```
0xhh0000
```

where *hh* must be unique to the volume group you are creating. Use the next hexadecimal number that is available on your system, after the volume groups that are already configured.

13. Use the vgimport command, specifying the map file you copied from the configuration node. In the following example, the vgimport command is issued on the second node for the same volume group that was modified on the first node:

```
# vgimport -v -m /tmp/vg_ops.map /dev/vg_ops
/dev/dsk/c0t2d0/dev/dsk/c1t2d0
```

14. Activate the volume group in shared mode by issuing the following command on both nodes:

```
# vgchange -a s -p /dev/vg_ops
```

Skip this step if you use a package control script to activate and deactivate the shared volume group as a part of OPS startup and shutdown.

Adding Additional Shared LVM Volume Groups

To add capacity or to organize your disk resources for ease of management, you may wish to create additional shared volume groups for your Oracle RAC databases. If you decide to use additional shared volume groups, they must conform to the following rules:

18.472
A.

- Volume groups should include different PV links to each logical unit on the disk array.
- Volume group names must be the same on all nodes in the cluster.
- Logical volume names must be the same on all nodes in the cluster.

Changing the VxVM or CVM Storage Configuration

You can add VxVM disk groups to the cluster configuration while the cluster is running. To add new CVM disk groups, the cluster *must* be running.

If you are creating new CVM disk groups, be sure to determine the master node on which to do the creation by using the following command:

```
# vxdctl -c mode
```

One node will identify itself as the master. Create disk groups from this node.

Similarly, you can delete VxVM or CVM disk groups provided they are not being used by a cluster node at the time.

NOTE If you are removing a disk group from the cluster configuration, make sure that you also modify or delete any package control script that imports and deports this disk group. If you are removing a CVM disk group, be sure to remove the STORAGE_GROUP entries for the disk group from the package ASCII file.

18.471
A.

Removing ServiceGuard Extension for RAC from a System

If you wish to remove a node from ServiceGuard Extension for RAC operation, use the `swremove` command to delete the software. Note the following:

- The cluster should not be running on the node from which you will be deleting ServiceGuard Extension for RAC.
- The node from which you are deleting ServiceGuard Extension for RAC should not be in the cluster configuration.
- If you are removing ServiceGuard Extension for RAC from more than one node, `swremove` should be issued on one node at a time.

NOTE After removing ServiceGuard Extension for RAC, your cluster will still have MC/ServiceGuard installed. For information about removing MC/ServiceGuard, refer to the user's manual *Managing MC/ServiceGuard* for your version of the product.

18.470
A.

Monitoring Hardware

Good standard practice in handling a high availability system includes careful fault monitoring so as to prevent failures if possible or at least to react to them swiftly when they occur. The following should be monitored for errors or warnings of all kinds:

- Disks
- CPUs
- Memory
- LAN cards
- Power sources
- All cables
- Disk interface cards

Some monitoring can be done through simple physical inspection, but for the most comprehensive monitoring, you should examine the system log file (`/var/adm/syslog/syslog.log`) periodically for reports on all configured HA devices. The presence of errors relating to a device will show the need for maintenance.

Using Event Monitoring Service

Event Monitoring Service (EMS) allows you to configure monitors of specific devices and system resources. You can direct alerts to an administrative workstation where operators can be notified of further action in case of a problem. For example, you could configure a disk monitor to report when a mirror was lost from a mirrored volume group being used in a non-RAC package. Refer to the manual *Using the Event Monitoring Service* (B7612-90009) for additional information.

Using EMS Hardware Monitors

A set of hardware monitors is available for monitoring and reporting on memory, CPU, and many other system values. Refer to the *EMS Hardware Monitors User's Guide* (B6191-90020) for additional information.

Using HP Predictive Monitoring

In addition to messages reporting actual device failure, the logs may accumulate messages of lesser severity which, over time, can indicate that a failure may happen soon. One product that provides a degree of automation in monitoring is called HP Predictive, which gathers information from the status queues of a monitored system to see what errors are accumulating. This tool will report failures and will also predict failures based on statistics for devices that are experiencing specific non-fatal errors over time. In a ServiceGuard cluster, HP Predictive should be run on all nodes.

HP Predictive also reports error conditions directly to an HP Response Center, alerting support personnel to the potential problem. HP Predictive is available through various support contracts. For more information, contact your HP representative.

13.463
A.

Adding Disk Hardware

As your system expands, you may need to add disk hardware. This also means modifying the logical volume structure. Use the following general procedure:

1. Halt packages.
2. Ensure that the Oracle database is not active on either node.
3. Deactivate and mark as unshareable any shared volume groups.
4. Halt the cluster.
5. Deactivate automatic cluster startup.
6. Shutdown and power off system before installing new hardware.
7. Install the new disk hardware with connections on all nodes.
8. Reboot all nodes.
9. On the configuration node, add the new physical volumes to existing volume groups, or create new volume groups as needed.
10. Start up the cluster.
11. Make the volume groups shareable, then import each shareable volume group onto the other nodes in the cluster.
12. Activate the volume groups in shared mode on all nodes.
13. Start up the Oracle RAC instances on all nodes.
14. Activate automatic cluster startup.

NOTE As you add new disks to the system, update the planning worksheets (described in the "Planning" chapter) so as to record the exact configuration you are using.

18.467
A.

Replacing Disks

The procedure for replacing a faulty disk mechanism depends on the type of disk configuration you are using and on the type of Volume Manager software. For a description of replacement procedures using VERITAS VxVM or CVM, refer to the chapter on "Administering Hot-Relocation" in the *VERITAS Volume Manager 3.2 Administrator's Guide*. Additional information is found in the *VERITAS Volume Manager 3.2 Troubleshooting Guide*.

The following paragraphs describe how to replace disks that are configured with LVM. Separate descriptions are provided for replacing a disk in an array and replacing a disk in a high availability enclosure.

Replacing a Mechanism in a Disk Array Configured with LVM

With any HA disk array configured in RAID 1 or RAID 5, refer to the array's documentation for instruction on how to replace a faulty mechanism. After the replacement, the device itself automatically rebuilds the missing data on the new disk. No LVM activity is needed. This process is known as *hot swapping* the disk.

NOTE If your LVM installation requires online replacement of disk mechanisms, the use of disk arrays may be required, because software mirroring of JBODs with MirrorDisk/UX does not permit hot swapping for disks that are activated in shared mode.

Replacing a Mechanism in an HA Enclosure Configured with Exclusive LVM

Non-Oracle data that is used by packages may be configured in volume groups that use exclusive (one-node-at-a-time) activation. If you are using exclusive activation and software mirroring with MirrorDisk/UX and the mirrored disks are mounted in a high availability disk enclosure, you can use the following steps to *hot plug* a disk mechanism:

18.466
A.

1. Identify the physical volume name of the failed disk and the name of the volume group in which it was configured. In the following examples, the volume group name is shown as `/dev/vg_sg01` and the physical volume name is shown as `/dev/c2t3d0`. Substitute the volume group and physical volume names that are correct for your system.
2. Identify the names of any logical volumes that have extents defined on the failed physical volume.
3. On the node on which the volume group is currently activated, use the following command *for each logical volume that has extents on the failed physical volume*:

   ```
   # lvreduce -m 0 /dev/vg_sg01/lvolname /dev/dsk/c2t3d0
   ```
4. At this point, remove the failed disk and insert a new one. The new disk will have the same HP-UX device name as the old one.
5. On the node from which you issued the `lvreduce` command, issue the following command to restore the volume group configuration data to the newly inserted disk:

   ```
   # vgcfgrestore /dev/vg_sg01 /dev/dsk/c2t3d0
   ```
6. Issue the following command to extend the logical volume to the newly inserted disk:

   ```
   # lvextend -m 1 /dev/vg_sg01 /dev/dsk/c2t3d0
   ```
7. Finally, use the `lvsync` command *for each logical volume that has extents on the failed physical volume.* This synchronizes the extents of the new disk with the extents of the other mirror.

   ```
   # lvsync /dev/vg_sg01/lvolname
   ```

Offline Replacement of a Mechanism in an HA Enclosure Configured with Shared LVM (SLVM)

Hot plugging of disks is not supported for Oracle RAC data, which is configured in volume groups with Shared LVM (SLVM). If you need this capability, you should use disk arrays for your Oracle RAC data.

If you are using software mirroring for shared concurrent activation of Oracle RAC data with MirrorDisk/UX and the mirrored disks are mounted in a high availability disk enclosure, use the following steps to carry out *offline* replacement:

18.465
A.

1. Make a note of the physical volume name of the failed mechanism (e.g., `/dev/dsk/c2t3d0`).
2. Deactivate the volume group on all nodes of the cluster:

   ```
   # vgchange -a n vg_ops
   ```

3. Replace the bad disk mechanism with a good one.
4. From one node, initialize the volume group information on the good mechanism using `vgcfgrestore`(1M), specifying the name of the volume group and the name of the physical volume that is being replaced:

   ```
   # vgcfgrestore /dev/vg_ops /dev/dsk/c2t3d0
   ```

5. Activate the volume group on one node in exclusive mode then deactivate the volume group:

   ```
   # vgchange -a e vg_ops
   ```

   This will synchronize the stale logical volume mirrors. This step can be time-consuming, depending on hardware characteristics and the amount of data.

6. Deactivate the volume group:

   ```
   # vgchange -a n vg_ops
   ```

7. Activate the volume group on all the nodes in shared mode using `vgchange - a s`:

   ```
   # vgchange -a s vg_ops
   ```

Replacing a Lock Disk

Replacing a failed lock disk mechanism is the same as replacing a data disk. If you are using a *dedicated* lock disk (one with no user data on it), then you need to issue only one LVM command:

```
# vgcfgrestore /dev/vg_lock /dev/dsk/c2t3d0
```

After doing this, wait at least an hour, then review the syslog file for a message showing that the lock disk is healthy again.

18.464
A.

On-line Hardware Maintenance with In-line SCSI Terminator

ServiceGuard allows on-line SCSI disk controller hardware repairs to all cluster nodes if you use HP's *in-line terminator* (C2980A) on nodes connected to the end of the shared FW/SCSI bus. The in-line terminator cable is a 0.5 meter extension cable with the terminator on the male end, which connects to the controller card for an external bus. The in-line terminator is used *instead of* the termination pack that is attached to the controller card and makes it possible to physically disconnect the node from the end of the F/W SCSI bus without breaking the bus's termination. (Nodes attached to the middle of a bus using a Y cable also can be detached from the bus without harm.) When using in-line terminators and Y cables, ensure that all orange-socketed termination packs are *removed* from the controller cards.

NOTE You cannot use inline terminators with internal FW/SCSI buses on D and K series systems, and you cannot use the inline terminator with single-ended SCSI buses. You must *not* use an inline terminator to connect a node to a Y cable.

Figure 3-1 shows a three-node cluster with two F/W SCSI buses. The solid line and the dotted line represent different buses, both of which have inline terminators attached to nodes 1 and 3. Y cables are also shown attached to node 2.

18463
A.

Figure 3-1 F/W SCSI Buses with In-line Terminators

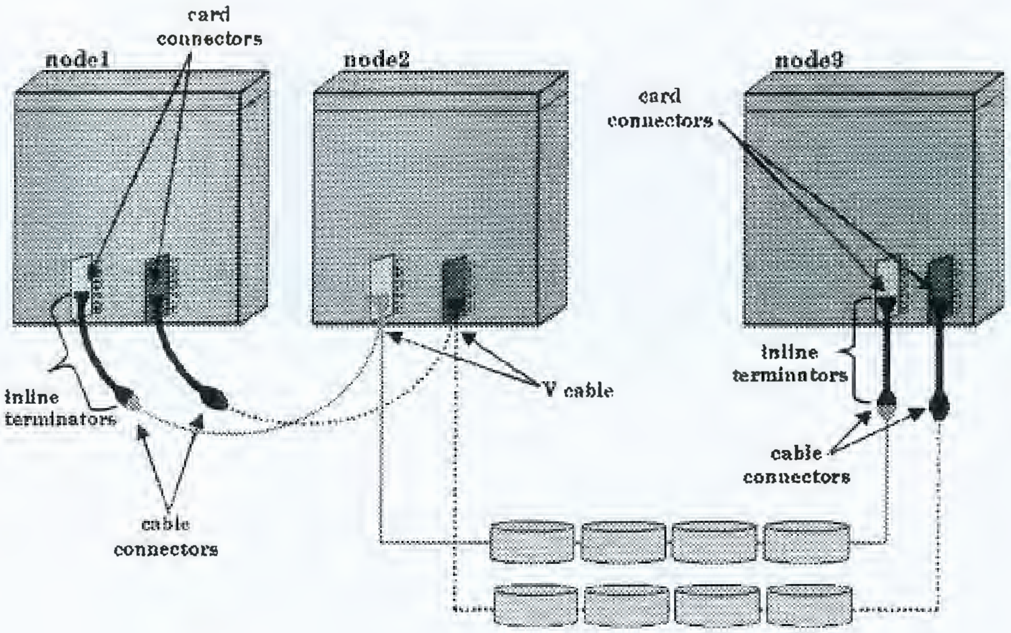

The use of in-line SCSI terminators allows you to do hardware maintenance on a given node by temporarily moving its packages to another node and then halting the original node while its hardware is serviced. Following the replacement, the packages can be moved back to the original node.

Use the following procedure to disconnect a node that is attached to the bus with an in-line SCSI terminator or with a Y cable:

1. Move any packages on the node that requires maintenance to a different node.
2. Halt the node that requires maintenance. The cluster will re-form, and activity will continue on other nodes. Packages on the halted node will switch to other available nodes if they are configured to switch.
3. Disconnect the power to the node.

18.462
A·

4. Disconnect the node from the in-line terminator cable or Y cable if necessary. The other nodes accessing the bus will encounter no problems as long as the in-line terminator or Y cable remains connected to the bus.
5. Replace or upgrade hardware on the node, as needed.
6. Reconnect the node to the in-line terminator cable or Y cable if necessary.
7. Reconnect power and reboot the node. If `AUTOSTART_CMCLD` is set to 1 in the `/etc/rc.config.d/cmcluster` file, the node will rejoin the cluster.
8. If necessary, move packages back to the node from their alternate locations and restart them.

Replacement of I/O Cards

After an I/O card failure, you can replace the card using the following steps. It is not necessary to bring the cluster down to do this if you are using SCSI inline terminators or Y cables at each node.

1. Halt the node by using SAM or the `cmhaltnode` command. Packages should fail over normally to other nodes.
2. Remove the I/O cable from the card. With SCSI inline terminators, this can be done without affecting the disks or other nodes on the bus.
3. Using SAM, select the option to do an on-line replacement of an I/O card.
4. Remove the defective I/O card.
5. Install the new card. The new card must be exactly the same card type, and it must be installed in the same slot as the card you removed.
6. In SAM, select the option to attach the new I/O card.
7. Add the node back into the cluster by using SAM or the `cmrunnode` command.

18.460
A.

Replacement of LAN Cards

If you have a LAN card failure, which requires the LAN card to be replaced, you can replace it on-line or off-line depending on the type of hardware and operating system you are running. It is not necessary to bring the cluster down to do this.

Off-Line Replacement

The following steps show how to replace a LAN card off-line. These steps apply to both HP-UX 11.0 and 11i:

1. Halt the node by using the `cmhaltnode` command.
2. Shut down the system using `/etc/shutdown`, then power down the system.
3. Remove the defective LAN card.
4. Install the new LAN card. The new card must be exactly the same card type, and it must be installed in the same slot as the card you removed.
5. Power up the system.
6. If necessary, add the node back into the cluster by using the `cmrunnode` command. (You can omit this step if the node is configured to join the cluster automatically.)

On-Line Replacement

If your system hardware supports hotswap I/O cards, and if the system is running HP-UX 11i (B.11.11 or later), you have the option of replacing the defective LAN card on-line. This will significantly improve the overall availability of the system. To do this, follow the steps provided in the section "How to On-line Replace (OLR) a PCI Card Using SAM" in the document *Configuring HP-UX for Peripherals*. The OLR procedure also requires that the new card must be exactly the same card type as the card you removed to avoid improper operation of the network driver. ServiceGuard will automatically recover the LAN card once it has been replaced and reconnected to the network.

18459
A.

After Replacing the Card

After the on-line or off-line replacement of LAN cards has been done, ServiceGuard will detect that the MAC address (LLA) of the card has changed from the value stored in the cluster binary configuration file, and it will notify the other nodes in the cluster of the new MAC address. The cluster will operate normally after this.

It is also recommended that you update the new MAC address in the cluster binary configuration file by re-applying the cluster configuration. Use the following steps for on-line reconfiguration:

1. Use the `cmgetconf` command to obtain a fresh ASCII configuration file, as follows:

```
# cmgetconf config.ascii
```

2. Use the `cmapplyconf` command to apply the configuration and copy the new binary file to all cluster nodes:

```
# cmapplyconf -C config.ascii
```

This procedure updates the binary file with the new MAC address and thus avoids data inconsistency between the outputs of the `cmviewconcl` and `lanscan` commands.

18.453
A.

18.457
A.

A Blank Planning Worksheets

This appendix reprints blank planning worksheets used in preparing the RAC cluster. You can duplicate any of these worksheets that you find useful and fill them in as a part of the planning process.

13.456
A.

LVM Volume Group and Physical Volume Worksheet

VG and PHYSICAL VOLUME WORKSHEET Page ___ of ____

==

Volume Group Name: ____

PV Link 1 PV Link2

Physical Volume Name:____

Physical Volume Name:____

Physical Volume Name:____

Physical Volume Name: ____

Physical Volume Name: ____

Physical Volume Name: ____

Physical Volume Name: ____

Volume Group Name: ____

PV Link 1 PV Link2

Physical Volume Name: ____

Physical Volume Name:____

Physical Volume Name: ____

Physical Volume Name: ____

Physical Volume Name: ____

Physical Volume Name: ____

Physical Volume Name: ____

18.455
A.

VxVM Disk Group and Disk Worksheet

```
DISK GROUP WORKSHEET                    Page ___ of ____
=============================================================================

Disk Group Name: ___________________________________________________________

Physical Volume Name:_______________________________________________________

Physical Volume Name:_______________________________________________________

Physical Volume Name:_______________________________________________________

Physical Volume Name: ______________________________________________________

Physical Volume Name: ______________________________________________________

Physical Volume Name: ______________________________________________________

Physical Volume Name: ______________________________________________________

Physical Volume Name: ______________________________________________________


Disk Group Name: ___________________________________________________________

Physical Volume Name: ______________________________________________________

Physical Volume Name:_______________________________________________________

Physical Volume Name: ______________________________________________________

Physical Volume Name: ______________________________________________________

Physical Volume Name: ______________________________________________________

Physical Volume Name: ______________________________________________________

Physical Volume Name: ______________________________________________________

Physical Volume Name: ______________________________________________________
```

18.454
A.

Oracle Logical Volume Worksheet

NAME SIZE

Oracle Control File 1:

Oracle Control File 2:

Oracle Control File 3:

Instance 1 Redo Log 1:

Instance 1 Redo Log 2:

Instance 1 Redo Log 3:

Instance 1 Redo Log:

Instance 1 Redo Log:

Instance 2 Redo Log 1:

Instance 2 Redo Log 2:

Instance 2 Redo Log 3:

Instance 2 Redo Log:

Instance 2 Redo Log:

Data: System

Data: Rollback

Data: Temp

Data: Users

Data: Tools

18.453
A.

Index

A

activation of volume groups
  in shared mode, 71
administration
  cluster and package states, 58
array
  replacing a faulty mechanism, 79, 80
AUTO_START_TIMEOUT
  in sample configuration file, 46

B

building a cluster
  CVM infrastructure, 49
building an RAC cluster
  displaying the logical volume infrastructure, 42
  logical volume infrastructure, 34
building logical volumes
  for RAC, 40

C

cluster
  state, 62
  status options, 58
cluster configuration file, 46
cluster volume group
  creating physical volumes, 35
CLUSTER_NAME (cluster name)
  in sample configuration file, 46
CVM
  creating a storage infrastructure, 49
  use of the VxVM-CVM-pkg, 50

D

deactivation of volume groups, 71
deciding when and where to run packages, 21
deleting nodes while the cluster is running, 74
demo database
  files, 40, 53
disk
  choosing for volume groups, 35
disk arrays
  creating logical volumes, 38
disk storage
  creating the infrastructure with CVM, 49
disks
  replacing, 79

E

eight-node cluster with disk array
  figure, 26
EMS
  for preventive monitoring, 76
enclosure for disks
  replacing a faulty mechanism, 79
Event Monitoring Service
  in troubleshooting, 76
exporting
  shared volume group data, 42
exporting files
  LVM commands, 43
  using SAM, 43

F

figures
  eight-node cluster with EMC disk array, 26
FIRST_CLUSTER_LOCK_PV
  in sample configuration file, 46
FIRST_CLUSTER_LOCK_VG
  in sample configuration file, 46

G

GMS
  group membership services, 20
group membership services
  define, 20

H

hardware
  adding disks, 78
  monitoring, 76
heartbeat subnet address
  parameter in cluster manager configuration, 33
HEARTBEAT_INTERVAL
  in sample configuration file, 46
HEARTBEAT_IP
  in sample configuration file, 46
high availability cluster
  defined, 16
HP Predictive monitoring
  in troubleshooting, 77

I

in-line terminator
  permitting online hardware maintenance, 82

Index

installing software
  MC/LockManager, 32
  Oracle Parallel Server, 45
IP address
  switching, 22

L

lock disk
  replacing a faulty mechanism, 81
logical volumes
  blank planning worksheet, 91, 92
  creating, 40
  creating for a cluster, 36, 51, 52
  creating the infrastructure, 34
  disk arrays, 38
  filled in planning worksheet, 28, 30
lssf
  using to obtain a list of disks, 35
LVM
  creating on disk arrays, 38
LVM commands
  exporting files, 43

M

maintaining an RAC cluster, 57
maintenance
  adding disk hardware, 78
  making changes to shared volume groups, 72
MC/LockManager
  installing, 32
  introducing, 15, 27
monitoring hardware, 76

N

network
  status, 61
NETWORK_INTERFACE
  in sample configuration file, 46
NETWORK_POLLING_INTERVAL (network polling interval)
  in sample configuration file, 46
node
  halting status, 67
  in an RAC cluster, 16
  status and state, 59
NODE_TIMEOUT (heartbeat timeout)
  in sample configuration file, 46

O

online hardware maintenance
  by means of in-line SCSI terminators, 82
opsctl.ctl
  Oracle demo database files, 40, 53
opslog.log
  Oracle demo database files, 40, 53
Oracle
  demo database files, 40, 53
Oracle Parallel Server
  installing, 45

P

package
  basic concepts, 17, 18
  moving status, 65
  state, 62
  status and state, 59
  switching status, 66
package configuration
  service name parameter, 33
packages
  accessing RAC database, 56
  deciding where and when to run, 21
physical volumes
  creating for clusters, 35
  filled in planning worksheet, 90
planning
  worksheets for logical volume planning, 28, 30
  worksheets for physical volume planning, 90
planning worksheets
  blanks, 89
point to point connections to storage devices, 25
Predictive monitoring, 77
PVG-strict mirroring
  creating volume groups with, 35

Q

quorum server
  status and state, 63

R

RAC
  group membership services, 20
  overview of configuration, 16
  packages to access database, 56

18.451
A.

status, 60
RAC cluster
  defined, 16
RAC software
  installing, 45
removing MC/ServiceGuard from a system, 75
replacing disks, 79
rollback.dbf
  Oracle demo database files, 41, 54
RS232 status, viewing, 67

S

SAM
  exporting files, 43
serial line
  status, 61
service
  status, 61
service name
  parameter in package configuration, 33
SERVICE_NAME
  parameter in package configuration, 33
shared mode
  activation of volume groups, 71
  deactivation of volume groups, 71
shared volume groups
  making volume groups shareable, 70
sharing volume groups, 42
SLVM
  making volume groups shareable, 70
state
  cluster, 62
  node, 59
  of cluster and package, 58
  package, 59, 62
status
  cluster, 58
  halting node, 67
  moving package, 65
  network, 61
  node, 59
  normal running RAC, 63
  of cluster and package, 58
  package, 59
  RAC, 60
  serial line, 61
  service, 61
  switching package, 66
switching IP addresses, 22
system multi-node package
  used with CVM, 50
system.dbf
  Oracle demo database files, 40, 53

T

temp.dbf
  Oracle demo database files, 41, 54
tools.dbf
  Oracle demo database files, 41, 54
troubleshooting
  monitoring hardware, 76
  replacing disks, 79

V

viewing RS232 status, 67
volume group
  creating for a cluster, 35
  creating physical volumes for clusters, 35
volume groups
  adding shared volume groups, 73
  displaying for RAC, 42
  exporting to other nodes, 42
  making changes to shared volume groups, 72
  making shareable, 70
  making unshareable, 70
VOLUME_GROUP
  in sample configuration file, 46
VxVM-CVM-pkg, 50

W

worksheet
  logical volume planning, 28, 30
worksheets
  physical volume planning, 90
worksheets for planning
  blanks, 89

18.450
A.

Index

18.449
A.

ANEXO UNIDADE DE BACKUP ROBOTIZADO
PARTE H

COBRA Tecnologia S.A.
Estrada dos Bandeirantes 7966
CEP 22783-110 Rio de Janeiro RJ
Tel. 21 2442-8800
www.cobra.com.br

18.448
A.

invent

hp server rp5400 series
entry-level UNIX servers

August 2002

a technical white paper
from Hewlett-Packard

18.447
A.

table of contents

introduction to the hp server rp5400 series 3
the hp server product line 4
binary compatibility 5
Intel Itanium Processor Family ready 5
Intel Itanium Processor Family transition 6
architecture 7
low-latency memory access 7
speeds and feeds 8
I/O subsystem design 8
internal removable media 10
scalability 10
rp5400 series industrial design and packaging 11
racking density 11
high availability slider rails 11
cabinet spacing requirements 11
standalone/deskside configuration 11
high availability 12
redundant, hot-swap power supplies 12
redundant power input protection 12
redundant, hot-swap cooling 13
main memory—advanced ECC and parity 13
memory chip spare technology 14
hot-plug disk drives 14
hot-plug PCI I/O slots 14
dynamic processor deallocation and resilience 15
manageability and support 16
LED event notification 16
event monitoring service 16
extended fault management system 17
system platform monitor 17
built-in Web console 18
LAN console 18
ASCII consoles 18
remote access 18
self-diagnosis 18
online and offline diagnostics 19
for more information 20

18.416

introduction to the hp server rp5400 series

In today's economy, whether you're managing your own IT infrastructure or hosting someone else's, you have to operate with a faster time-to-solution, within budgetary constraints, and with the highest standards for customer service and operational efficiency.

The HP Server rp5400 series gives you the fastest and most reliable means to succeed in this new business environment. The rp5400 series consists of two products that deliver the proven performance, scalability, and high-availability capabilities of UNIX®—without high maintenance requirements and costs. And they give you plenty of room to grow. You can start with a low-price entry point and scale up to the leading 4-way UNIX performance—in the same form factor.

The rp5400 series is made up of two different servers, each with a unique ability to match your computing needs. The rp5430 is a 2-way system with time-proven, cost-effective PA-RISC technology. It offers the latest PA-8700+ processors and the same high-performance core electronics found in more expandable HP servers. The rp5470 is the high-performance flagship of the lineup. It supports up to 4-way PA-8700+ processing power and industry-leading bandwidth.

Both members of the rp5400 series use the same rack-optimized 7U package. This allows seamless scalability with simple in-box upgrades between servers in the series. Additionally, the rp5400 series was designed for board-swap upgrades to the Itanium 2-based HP Server rx5670. Together, the rp5400 series offers the industry's best 1- to 4-way lineup with unparalleled investment protection.

figure 1.1 a front view of the rp5400 series

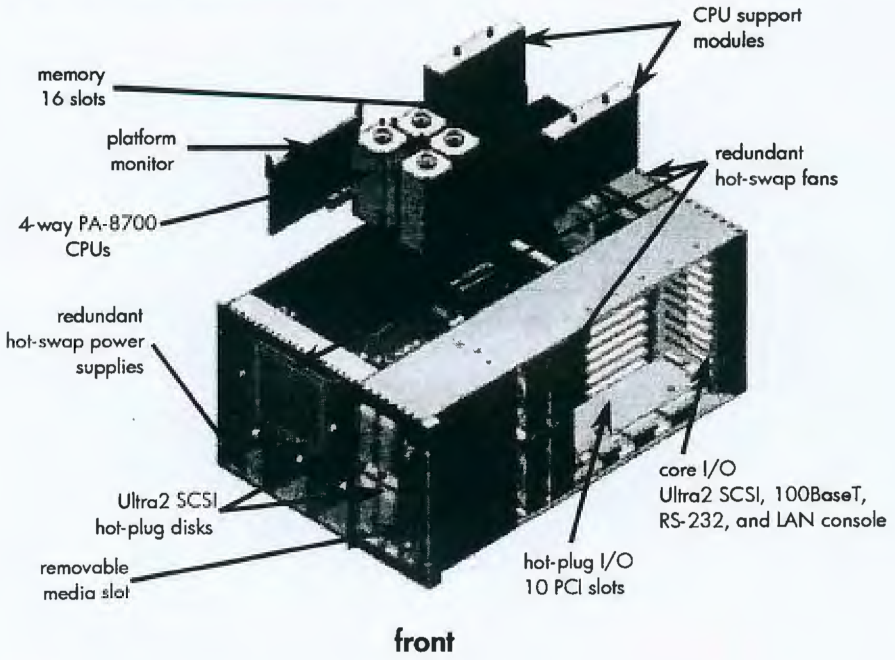
CPU support modules
memory 16 slots
platform monitor
redundant hot-swap fans
4-way PA-8700 CPUs
redundant hot-swap power supplies
core I/O Ultra2 SCSI, 100BaseT, RS-232, and LAN console
Ultra2 SCSI hot-plug disks
removable media slot
hot-plug I/O 10 PCI slots
front

18.445
A.

Figure 1.1 reveals the location of major components, as well as the mechanical and architectural features of the rp5400 series. The server is partitioned into two main electrical assemblies—the system board and the I/O backplane—and into three main volumes—processor and memory, I/O and disk, and power.

Looking at the front face, three hot-swap power supply bays are located in the lower left corner. To the right, a peripheral bay provides space for four hot-plug disks and one removable media device (either DVD-ROM or DDS-3). Directly above the power supply bays is the first of eight hot-swap cooling fans.

The right side of the system houses the I/O card bay. There are ten PCI I/O slots available. Two pairs of fans located here provide cooling for the I/O bay as well as the peripheral bay.

The opening at the top provides access to the system board, which supports the four CPUs, sixteen dual inline memory module (DIMM) slots, two processor support modules, and the platform monitor board. The core I/O is located at the rear of the system.

rp5430 features at-a-glance

- 1 to 2 PA-8700+ or PA-8700 CPUs
- 875MHz and 750MHz CPUs
- high-performance "stretch" core electronics complex (leveraged from rp7400)
- Intel® Itanium® 2 upgradable
- up to 8GB of memory
- 6 PCI I/O slots (5 are hot-plug 66MHz × 64-bit)
- 6 independent PCI buses for I/O slots
- N+1 power and cooling
- 4 hot-plug disk drives
- removable media bay: DVD-ROM or DDS-3
- 4.3GB/s system bus bandwidth
- 2.1GB/s I/O bus bandwidth
- 4.3GB/s memory bus bandwidth
- 64-bit HP-UX 11.0 & 11i
- high-density 7-EIA-unit, 19-inch rackmount or pedestal package

rp5470 features at-a-glance

- 1 to 4 PA-8700+ or PA-8700 CPUs
- 875MHz and 750MHz CPUs
- high-performance "stretch" core electronics complex (leveraged from rp7400)
- Intel Itanium 2 upgradable
- up to 16GB of memory
- 10 PCI I/O slots (8 hot-plug, 2 non-hot-plug; all are 66MHz × 64-bit)
- 9 independent PCI buses for I/O slots
- N+1 power and cooling
- 4 hot-plug disk drives
- removable media bay: DVD-ROM or DDS-3
- 4.3GB/s system bus bandwidth
- 3.2GB/s I/O bus bandwidth
- 4.3GB/s memory bus bandwidth
- 64-bit HP-UX 11.0 & 11i
- high-density 7-EIA-unit, 19-inch rackmount or pedestal package

the hp server product line

The rp5400 series is the entry-level cornerstone of the business-critical proven HP server product line. HP servers are #1 among UNIX servers for reliability, scalability, availability, and price/performance. This robust product line addresses the major computing challenges customers face today in online transaction processing (OLTP), electronic commerce (ECOM), Internet/intranet serving (Web), enterprise resource planning (ERP), supply chain management (SCM), and technical applications.

At the low end, affordable rp2400 and rp5400 series servers effortlessly handle Internet workloads and enterprise-size applications. Both platforms also add leadership price/performance and include bundled Internet software solutions.

In the midrange, the rp8400 and rp7410 deliver the high-performance, compact Internet-era UNIX server platform that today's IS executives are demanding. With up to 16 PA-8700+ processors, the HP server midrange lineup provides the robust performance and scalability needed for the most demanding workloads.

13.4
A

With exceptional OLTP performance, availability, scalability, and manageability, HP Superdome has become the pacesetter for high-end computing. Superdome, coupled with HP's always-on infrastructure strategy, provides UNIX application performance and Internet-critical high availability to help you meet the rigorous demands of e-services and systems consolidation, as well as large-scale, highly complex technical modeling and simulations.

All of HP's UNIX servers provide excellent investment protection with a smooth transition path to future PA-RISC and/or Itanium-based architectures. So whether your business requires cutting-edge e-services, systems consolidation, or a host of other solutions, our power-packed servers are business-critical proven and ready to meet the challenge—today and tomorrow.

figure 1.2 the industry's strongest UNIX lineup—top to bottom

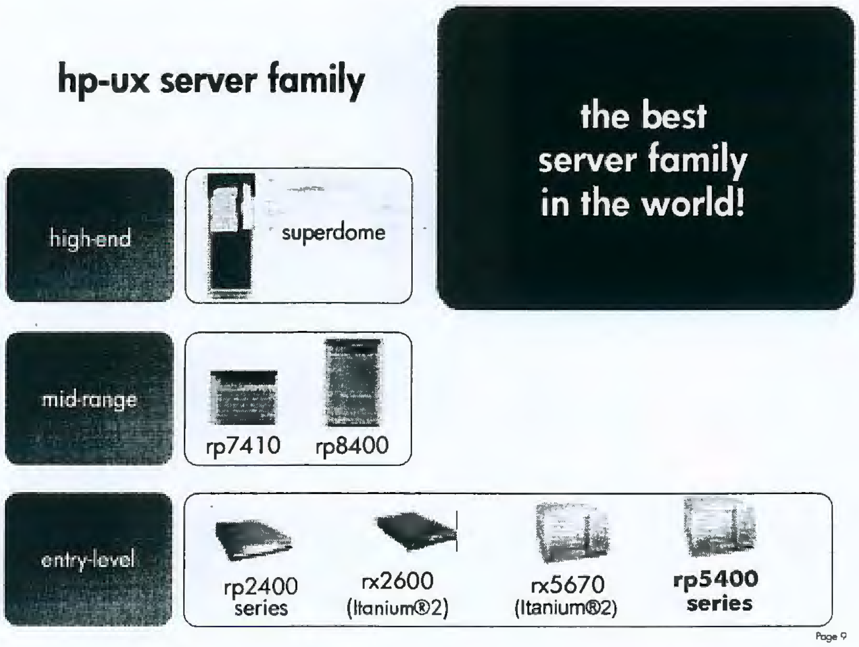

binary compatibility

The rp5400 series supports the 64-bit HP-UX 11 operating system. With HP-UX 11, HP maintains its longstanding tradition of providing the industry's best record of investment protection. HP-UX provides forward binary compatibility, in which a fully bound application developed on an earlier version of HP-UX is ensured to run smoothly on HP-UX 11. Thus, current 32- and 64-bit applications can run without requiring recompilation.

Intel Itanium Processor Family ready

The rp5400 series was designed for several generations of PA-RISC and is upgradable to the Itanium 2-based HP Server rx5670. HP offers a board-swap upgrade to move any rp5400 series product or any legacy HP 9000 L-Class product to the Intel Itanium Processor Family.

The Intel Itanium Processor Family is based on Explicitly Parallel Instruction Computing (EPIC), a new architecture technology invented by HP Labs. The EPIC architecture breaks through the sequential nature of today's RISC and CISC processor architectures by allowing the software to communicate explicitly to the processor when operations can be done in parallel. EPIC serves as the enabler for future high-performance chips by providing explicit parallelism, massive resources, and inherent scalability not available with conventional RISC architectures. Increased performance is realized by reducing the number of branches and branch mispredicts and by reducing the effects of memory-to-processor latency.

18.443
A.

Intel Itanium Processor Family transition

For the vast majority, the transition to the Intel Itanium Processor Family will be simple and seamless. For customers who require additional assistance, HP provides transition services around the world to help make this upgrade as smooth as possible. HP can provide assistance every step of the way, from assessment and design to verification and deployment. Consult the Intel Itanium processor section of HP's Web page for further information.

18.442
A-

architecture

Figure 2.1 shows the relationship of the rp5470 main blocks with the buses that connect them. The rp5470 uses the "stretch" high-performance core electronics complex (CEC), which is also used in the midrange rp7400 server. This CEC, specifically designed for demanding Internet workloads, brings unprecedented levels of bandwidth and performance to the 4-way entry-level market.

Two front-side buses, both running at 133MHz, provide 4.3GB/s of bandwidth to four PA-8700+ or PA-8700 processors. The low-latency memory controller provides 4.3GB/s of memory bandwidth to two 8-slot memory extenders. The I/O controller provides twelve 250MB/s data channels, for an aggregate bandwidth of 3.2GB/s distributed among the 10 PCI slots and multi-function core I/O.

The rp5430 architecture is similar to the rp5470. However, only half of the processor, memory, and I/O slot capacity is utilized.

figure 2.1 rp5470 architecture

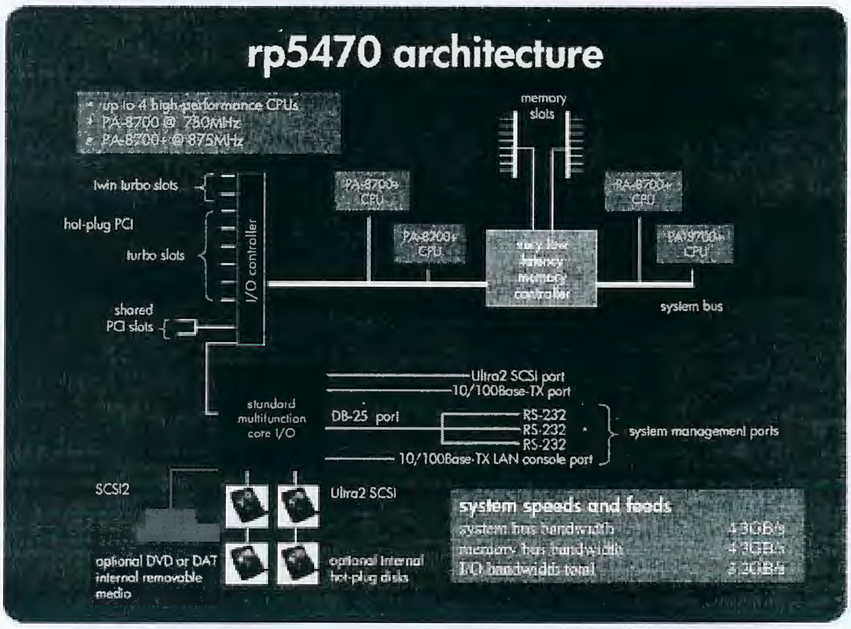

system speeds and feeds	
system bus bandwidth	4.3GB/s
memory bus bandwidth	4.3GB/s
I/O bandwidth total	3.2GB/s

low-latency memory access

Both the rp5430 and rp5470 support one or two 8-memory-slot carrier boards, for a maximum of 16 memory slots. The memory for both systems is connected to the CEC through a low-latency/high-bandwidth bus. With approximately half the latency of HP's previous generation K-Class server, the rp5400 series can supply the CPU with requested data in a fraction of the time of competitive systems.

The rp5400 series uses state-of-the-art synchronous dynamic random access memory (SDRAM) technology, available in 256MB, 512MB, 1GB, and 2GB DIMM pairs, all with advanced error checking and correcting (ECC) protection to detect and correct single-bit errors. The rp5470 supports up to 16GB of total system memory. The rp5430 supports up to 8GB of memory. Although all sixteen memory slots are active in the rp5430, the system will not boot if more than 8GB of memory is loaded. Memory configurations should be planned appropriately.

The "stretch" core electronics complex used in the rp5430 and rp5470 supports memory chip spare. This high-availability technology detects and corrects multiple-bit errors on memory DIMMs. With chip spare, any single DRAM chip can fail and the system will continue to operate normally. Chip spare is not supported on the 256MB DIMM pair, nor is it supported on the older-generation rp5400 and rp5450 servers.

To decrease memory latency and improve performance, the memory address lines are buffered three times: once on the system board to drive each memory carrier, once on the memory carrier to drive banks of DIMMs, and again on each DIMM before driving the memory components.

speeds and feeds

Tables 2.1 and **2.2** show the theoretical maximum bandwidth for various system buses. Theoretical maximum bandwidth is defined as the bus width multiplied by the frequency and number of buses.

table 2.1 maximum bandwidth for rp5470 system buses

	# of buses (or controllers)	**maximum bus bandwidth**	**aggregate bus bandwidth**
twin-turbo PCI slots	2	500MB/s	1GB/s
turbo PCI slots	6	250MB/s	1.5GB/s
shared PCI slots	1	250MB/s	250MB/s
core I/O	1	250MB/s	250MB/s
I/O subsystem	1 (controller)	3.2GB/s	3.2GB/s
memory subsystem	2	2.15GB/s	4.3GB/s
CPU buses	2	2.15GB/s	4.3GB/s

table 2.2 maximum bandwidth for rp5430 system buses

	# of buses (or controllers)	**maximum bus bandwidth**	**aggregate bus bandwidth**
twin-turbo PCI slots	2	500MB/s	1GB/s
turbo PCI slots	3	250MB/s	750MB/s
core I/O	1	250MB/s	250MB/s
I/O subsystem	1 (controller)	2.1GB/s	2.1GB/s
memory subsystem	2	2.15GB/s	4.3GB/s
CPU buses	2	2.15GB/s	4.3GB/s

I/O subsystem design

The rp5470 contains ten PCI I/O slots. The top eight slots have hot-plug capabilities under HP-UX 11i. The eight hot-plug slots all have independent I/O channels. This independent design prevents slow cards from affecting the performance of a fast card. Not only does independence provide great performance, but it also provides error containment. For example, if a card hangs in slot 9, cards in slots 0–8 will still function properly. The first two hot-plug slots are twin-turbo slots, meaning they each have two dedicated 250MB/s channels or a total of 500MB/s per slot. These two slots should be reserved for the highest performing I/O cards, such as Fibre Channel, Gigabit Ethernet, or Hyperfabric controllers. The remaining six hot-plug slots are turbo slots, each with a single 250MB/s channel.

In addition to the eight hot-plug slots, the rp5470 has two shared PCI slots. These slots share a single 250MB/s channel.

All ten of the rp5470 I/O slots use HP-developed adaptive signaling technology to automatically match an I/O card's appropriate speed and data width. Therefore, all slots will accept 64- or 32-bit cards running at either 33MHz or 66MHz.

18.440
A.

The rp5430 I/O subsystem is similar to the rp5470. In the rp5430, however, the shared PCI slots and three of the turbo slots are not active. Both twin turbo slots, four additional turbo slots, and the multifunction core I/O are available in the rp5430.

figure 2.2 rp5470 I/O subsystem

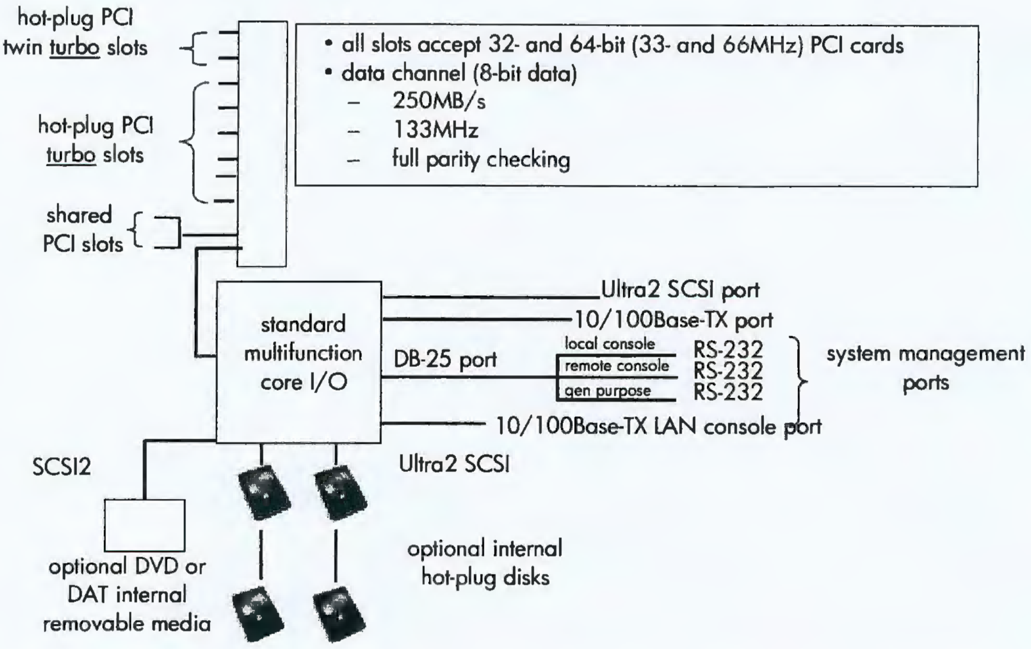
hot-plug PCI twin turbo slots
hot-plug PCI turbo slots
shared PCI slots
• all slots accept 32- and 64-bit (33- and 66MHz) PCI cards
• data channel (8-bit data)
– 250MB/s
– 133MHz
– full parity checking
standard multifunction core I/O
Ultra2 SCSI port
10/100Base-TX port
DB-25 port
local console
remote console
gen purpose
RS-232
RS-232
RS-232
system management ports
10/100Base-TX LAN console port
SCSI2
Ultra2 SCSI
optional DVD or DAT internal removable media
optional internal hot-plug disks

PCI is the optimized, industry-standard I/O bus. **Tables 2.3** and **2.4** summarize the PCI slots for each of the systems in the rp5400 series.

table 2.3 rp5470 PCI I/O

	# of slots	hot plug	bandwidth per channel	bus width	signaling speed	slot keying	adaptive signaling
twin turbo	2	yes	500MB/s	64 bits	66 & 33 MHz	5 volts	yes
turbo	6	yes	250MB/s	64 bits	66 & 33 MHz	5 volts	yes
shared	2	no	250MB/s	64 bits	33MHz	5 volts	yes

table 2.4 rp5430 PCI I/O

	# of slots	hot plug	bandwidth per channel	bus width	signaling speed	slot keying	adaptive signaling
twin turbo	2	yes	500MB/s	64 bits	66 & 33 MHz	5 volts	yes
turbo	4	3 of 4	250MB/s	64 bits	66 & 33 MHz	5 volts	yes

18.439
A.

internal removable media

The rp5400 series contains a single removable media bay that can accommodate either a DVD-ROM or DDS-3. The media bay is supported by one of two SCSI controllers located within the core I/O.

A dedicated single-ended (SE) SCSI channel connects the media bay to the controller. The removable media bay does not support hot-plug capability. The DVD-ROM drive provides access of up to 650MB of data from one disk. The DVD-ROM drive provides enhanced features while preserving backward read compatibility with the CD-ROM. Data transfer rates of up to 6.75MB/s are achieved with the DVD format; 4.8MB/s can be achieved with the CD format.

The DDS-3 drive offered with the rp5400 series provides storage capacity of up to 12GB on a single tape. This drive can store up to 7.2GB of data per hour, and automatic read-after-write verification helps to ensure the integrity of stored data. Read-write backward compatibility with DDS-1 and DDS-2 allows continued use of existing archive tapes.

scalability

The rp5400 series is designed without tradeoffs in CPU, memory, internal storage, or I/O expandability to offer the best scalability in the market.

- **CPU upgrades**—With its entry-level configuration of one CPU and single-CPU increments available up to four processors, the rp5400 series offers great flexibility to cover a wide range of performance points. The rp5430 and rp5470 offer 875MHz PA-8700+ processors, as well as the 750MHz PA-8700 processor.
- **memory upgrades**—The rp5400 series memory subsystem is also designed for scalability. With 16 available slots, the servers range from a minimum of 256MB to a maximum of 16GB of main memory.
- **internal storage**—The rp5400 series supports up to four internal hot-plug disk drives, which can be either half-height or low-profile form factors. Current disk offerings include 18, 36, and 73GB Ultra160 disk drives. The maximum internal storage is 292GB, via four 73GB drives.

18.438

rp5400 series industrial design and packaging

The rp5400 series has been designed to fit into environments ranging from data centers to deskside. The industrial design is coordinated with other HP servers and peripherals for a consistent appearance.

racking density

The rp5400 series is designed to provide unprecedented performance density that easily adapts to different environments. At 7 EIA units (EIA unit= 1.75 inches), up to five rp5430 or rp5470 systems can be installed into a single 2-meter HP cabinet. With the high cost of computer room floor space, this small footprint dramatically lowers total cost of ownership.

The rp5400 series is supported in A490xA and A189xA cabinets. When using the high availability slider rail, bolt-on anti-tip feet are required. When using the slider in A189xA cabinets, ballasts are required (see the HP 9000 Enterprise Servers Configuration Guide for details).

The rp5400 series is also supported in a variety of third-party, non-HP racks and cabinets. Please refer to the HP 9000 Enterprise Servers Configuration Guide for the latest list of qualified third-party racks.

Note—dimensions for rack configuration: H= 12.25 inches (311 mm), D= 30.5 inches (775 mm), W= 19 inches (482 mm).

high availability slider rails

There are two rail options, static or slider, available for racking the rp5400 series into an HP cabinet. The high availability (HA) slider rails were designed to allow easy service access to the system, as well as to enable the hot-plug capability of the I/O slots and the hot-swap of four fans in the side cavity. With the HA slider rail, the rp5400 series can be completely serviced without removing it from the rack, thus allowing side-by-side racks of systems to be completely supported without sacrificing floor space for side access to the system. The slider rails also contribute to a 100% improvement in "mean time to repair" over D- and K-Class servers. The high availability slider rails are highly recommended.

Note—the slider mechanism occupies 1 EIA unit of rack space. When used with the rp5400 series, the combination will occupy 8 EIA units of rack space.

Static rails do not consume EIA space within the cabinet, therefore leaving more EIA space for peripherals. However, using static rails prohibits hot-plug of the I/O cards and hot-swap of the I/O bay fans.

cabinet spacing requirements

The rp5400 series requires a minimum of 24 inches (61 cm) of free space in both the front and rear of the cabinet for proper ventilation. During product installation and servicing, a total of 32 inches (82 cm) of free space is needed at the front of the cabinet.

The depth of HP A490xA cabinets is 39 inches (99 cm). Therefore, a minimum of 87 inches (221 cm) of total space is needed for each cabinet during normal operation. An additional 8 inches (21 cm) is needed during installation and servicing.

standalone/deskside configuration

The rp5400 series is also available in a standalone configuration when a cabinet is not desired. The standalone system is ideal for an office environment, under a desk, or on a shelf. The standalone configuration utilizes the same internal chassis and front plastic bezel as the racked version. However, a sheet metal cover, base, and casters are added for functionality and aesthetics. Casters can be removed when not desired.

Note—dimensions for standalone/deskside configuration: H= 14.5 inches (368 mm), D= 30.5 inches (775 mm), W= 19 inches (482 mm).

18.437
A.

high availability

The rp5400 series has numerous high availability features that are unmatched in the entry-level server market—features such as redundant hot-swap fans and power, hot-plug I/O and disks, memory scrubbing and page deallocation, memory chip spare, independent PCI slots, failure avoidance and notification capability, and MC/Serviceguard support. These features improve the availability level of the total system and are introduced in this section.

redundant, hot-swap power supplies

HP power supplies have a long history of excellent reliability, and the redundant power supply option increases HP's commitment to even higher reliability and availability.

The rp5400 series power subsystem holds a maximum of three hot-swap power supplies. These supplies are located in the very front of the server. Each supply is capable of sustaining 930 watts of output. The rp5430 comes standard with one power supply; a second and even third supply can be added for N+1 or N+2 redundancy. The rp5470 comes standard with two power supplies; a third supply can be ordered for N+1 redundancy. Each power supply has its own power cord, which provides protection against losing the power from a single cord or breaker. To maximize availability, the power cords should be plugged into separate breakers whenever possible.

Because of the hot-swap capability, in the event of a power supply failure, the faulty supply can be removed and replaced without notifying the system. This, of course, is assuming that an N+1 condition exists.

Exchanging a power supply in a running system involves opening the hinged, front plastic bezel. The failed power supply is easily identified and removed. The power supply is exchanged with a good one and the door is then closed to finish the process. The system will log a management code to indicate that redundancy is re-enabled. It is that simple.

There is another advantage for those customers with rigorous preventative maintenance programs. While the server continues to operate, the power supplies can be removed one-at-a-time and dust buildup can be vacuumed using proper electrostatic discharge (ESD) procedures.

redundant power input protection

Figure 3.1 contains a diagram of the rp5400 series power subsystem. This section explains how customers can utilize these capabilities to achieve different levels of power input protection.

figure 3.1 power subsystem

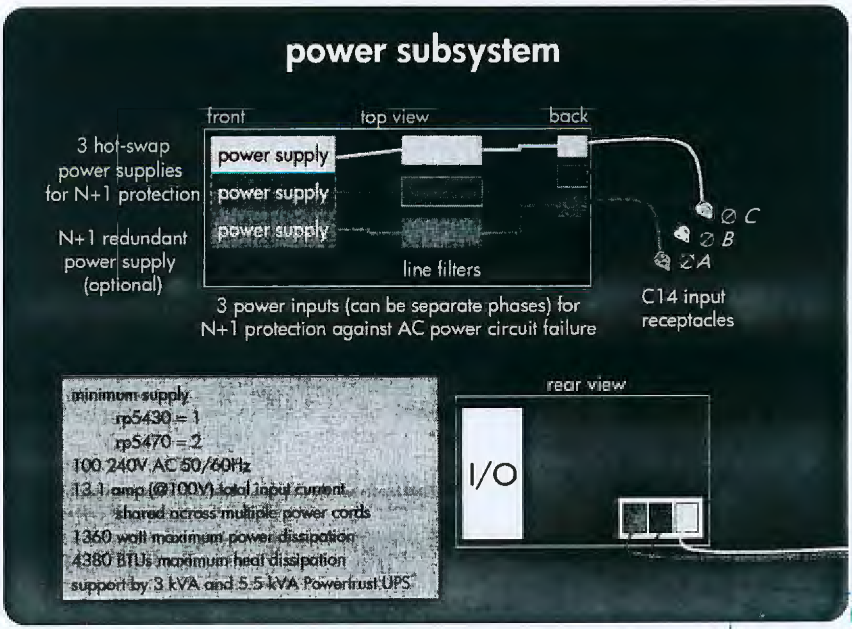

power subsystem

front
top view
back
3 hot-swap power supplies for N+1 protection
N+1 redundant power supply (optional)
power supply
power supply
power supply
line filters
C
B
A
3 power inputs (can be separate phases) for N+1 protection against AC power circuit failure
C14 input receptacles
rear view
I/O

minimum supply
rp5430 = 1
rp5470 = 2
100 240V AC 50/60Hz
13.1 amp (@100V) total input current
shared across multiple power cords
1360 watt maximum power dissipation
4380 BTUs maximum heat dissipation
support by 3 kVA and 5.5 kVA Powertrust UPS

18.436
A.

The server has three AC input line cords to reduce single points of failures. Each line cord supplies power to one of the three internal power supplies. The system is designed to operate on nominal 100–240 VAC and 50- or 60-Hz power without line-select switches. Each power supply can draw up to 930 watts. Because the servers will continue to operate with two of the three supplies functioning, many possibilities exist for the customer to configure the AC input depending on the level of protection desired. If the site has very stable AC power, all three line cords could be plugged into the same power grid. For additional protection, a single uninterruptible power supply (UPS) could be utilized to supply power to all three cords if primary AC power should fail.

- The next higher level of protection is to have three branch AC circuits, one for each AC input. This reduces the dependency on single-point breaker failures and common wiring. Additional protection for this configuration would utilize three smaller UPSs.
- The highest level of protection is three electrical utilities that each supply a branch circuit. This approach is expensive but does greatly reduce single points of failures. Large sites with many systems may find this configuration cost-effective. For the ultimate protection of large sites, install a large UPS on each branch circuit.

redundant, hot-swap cooling

The rp5400 series contains eight hot-swappable fans to cool system components. The eight cooling fans (1 front-access, 4 side-access, 3 rear-access) are arranged in an N+1 configuration so any fan can fail and not affect system uptime. In the event of a fan failure, the faulty fan can simply be removed and replaced while the server continues to run. The design pairs fans together. If one fan fails, the other speeds up to ensure adequate system cooling.

In addition, the server monitors ambient temperature and the power consumed within the box to determine the desired fan speed. By sensing the tachometer outputs from each fan, the actual speed is determined. Digital phase locked loop (DPLL) circuitry is used to individually adjust the speed of each fan to the desired common speed.

These smart algorithms reduce unnecessary fan noise, power consumption, and wear while producing a very clear indication of a working, cooling subsystem. In the unlikely event of a fan failure, it will drop out-of-lock with the DPLL. The server signals a fan failure via chassis codes to the console and will light an LED on the failed fan assembly.

There is another advantage for those customers with rigorous preventative maintenance programs. While the server continues to operate, the fans can be removed one-at-a-time and dust buildup can be vacuumed using proper ESD procedures.

main memory—advanced ECC and parity

Data stored in the main memory is protected by error checking and correcting (ECC) and address/control parity. The ECC design provides memory scrubbing and page deallocation functionality that will tolerate typical hard single-bit SDRAM failures without requiring DIMM replacement.

The data controllers generate ECC bits and store these ECC bits with the data in the DIMMs. The 256MB, 512MB, and 1GB DIMMs use x4 SDRAMs to store each bit of a word, including its ECC bits, in a different SDRAM within the DIMM pair. The 128MB DIMMs use x8 SDRAMs. When reading the data back, the data controllers are able to detect and correct single-bit data errors. Double-bit errors cannot be corrected. Double-bit data errors are highly unlikely because the data and ECC bits are stored one-bit-per-SDRAM, and multiple SDRAMs would have to be involved in the error. Hence, a single SDRAM could fail within each DIMM pair and the system would still function.

The system also detects address and control parity errors to prevent data corruption from reading or writing to the wrong location in main memory. The address controller and each address buffer generate address and control parity. Each address buffer detects address and control parity problems and reports it back to the address controller. There are three levels of address buffers as the address lines fan out. These address buffers are located on the system board and on each memory carrier on each DIMM.

memory chip spare technology

Chip spare is the ability of the system to continue to run in the face of any single or multi-bit chip error on a DRAM. DRAMs are basically N+1 per memory word. This functionality is essential in the design of reliable memory systems. Systems without this functionality are doomed to fail at an alarming rate when compared to HP servers.

Both the rp5430 and rp5470 support chip spare. The 256MB DIMM pair (product A5554A) does not support chip spare. To ensure maximum memory availability, the rp5430 and rp5470 should be configured with 512MB, 1GB, or 2GB memory modules only. The older generation rp5400 and rp5450 do not have chip spare capabilities.

hot-plug disk drives

The rp5400 series has four embedded SCSI disks accessible from the front of the server. These disks can be removed and inserted while the server continues to operate. This operation is called hot-plug, and it is different from hot-swap.

During both hot-plug and hot-swap operations, the power remains on and the system continues to function. However, hot-swap means that the assembly can be removed, added, or replaced without informing the system. Hot-plug requires the assembly to be deconfigured before removal and reconfigured before the system can utilize the newly inserted assembly. Because disks have unique information stored on them, hot-plug methods are used. Fans and power supplies are hot-swap assemblies.

Two dual-channel SCSI controllers manage the four internal hot-plug disks. For added availability, disk pairs are on separate channels as well as separate SCSI controllers. This means that with disk mirroring, a SCSI controller, SCSI channel, or root disk could fail and the server would continue to run properly.

The rp5400 series contains circuitry to properly control the disk's power and reset during the hot-plug operation. Either system administration manager (SAM) or the MESA suite of online diagnostic software can be utilized to effectively deconfigure and reconfigure the disk.

Another advantage for those customers with rigorous security programs is the ability to completely remove and isolate disks in a disaster- and theft-safe environment.

hot-plug PCI I/O slots

The ability to hot-plug PCI cards offers excellent flexibility for adding, reconfiguring, and maintaining I/O functions while the system continues operations. No reboot is required.

The I/O card bay is located at the right rear of the chassis. The I/O bay supports up to 10 PCI cards. Access to the I/O bay in rackmounted systems utilizing the high availability slider is achieved by sliding the server forward. Special features on the chassis, along with custom rack rails, allow the unit to move safely and smoothly during online service with all cables still attached. Once the system is slid into the service position, the I/O bay cover can be removed to gain side access to the PCI cards. In the standalone configuration, the outer shell is removed to gain access to the I/O bay.

The rp5470 has ten PCI I/O slots, and eight of those slots are hot-plug capable. Each hot-plug slot supports 64-bit x 66MHz PCI cards running at full speed and is connected to the I/O controller via independent, high-speed 250MB/s channels. This independent design prevents slow cards from affecting the performance of a fast card. Not only does independence provide great performance, but it also provides error containment. For example, if a card hangs in slot 10, cards in all other slots will still function properly. The highest-performing cards should always be placed in these independent slots.

The rp5430 has six PCI I/O slots. Ten physical slots are available, but only six of the slots are functioning electrically. Five of these six slots support hot-plug actions.

The PCI cards are spaced on a .9-inch pitch to allow for special hot-plug features and increased PCI reliability. Extra airflow holes between bulkheads more than double PCI airflow. Between the PCI slots, I/O card separators prevent electrical shorting and exposure to hazardous energy during hot-plug installation and removal. Locking features are designed into the main chassis to eliminate the need for individual PC board bulkhead screws, thus removing a potential electrical safety hazard.

18.434
A

Hot-plugging I/O cards have both hardware and software components. The hardware requirements are met by the electronics on the I/O backplanes and by mechanical design in the I/O cardcage. Bus idling, slot-to-slot electromechanical isolation, per-slot power and reset control, and visual indicators are all components of the total hot-plug hardware solution. With associated software, any card located in a hot-pluggable PCI slot can be removed, replaced, or added without power cycling, rebooting the system, or impacting the operation of other I/O transactions.

Please note that software support for hot-plug I/O is available in HP-UX 11i, but not in HP-UX 11.0.

dynamic processor deallocation and resilience

Every multi-CPU server with properly loaded HP-UX 11 has the capability for Dynamic Processor Deallocation and Resilience. Incorporated into HP-UX 11 is the capability to take a processor out of service while the system is running, without interruption to applications. This technology is referred to as Dynamic Processor Deallocation. Once a processor is deallocated, the HP-UX operating system will migrate all application processes that are currently scheduled on that processor to other active processors. Note that if the processor has been assigned to handle interrupts for any I/O drivers, it will continue to do so while it is deallocated.

The rp5400 series PA-RISC processors have the ability to detect and correct single-bit cache errors. The embedded event monitoring service (EMS) monitors the rate of correctable errors in each processor's on-board cache. These errors are manifested as low-priority machine checks (LPMCs). While occasional correctable errors are to be expected in the on-board cache, too many of these errors in a short period of time indicate an increased likelihood that a non-correctable cache error could occur. The EMS LPMC monitor will continuously monitor the rate at which LPMCs are occurring and dynamically deallocate a processor, using the Dynamic Processor Deallocation facility. This technology is referred to as Dynamic Processor Resilience.

manageability and support

The rp5400 series has many features to minimize the effort required to manage one system or an entire computer room. The server simplifies system management in several aspects: event notification, automatic error handling, power monitoring, and user interface to system management.

LED event notification

For an operator who is physically present, the simplest and easiest way to check system status is by quickly glancing at the status LEDs on the front of the system. The five LEDs each have a specific meaning:

- power—power is present and on, and power supplies are functioning properly
- remote—remote console is enabled
- run—system is up and running
- attention—occurrence of a non-catastrophic event, e.g., failure of an N+1 component
- fault—occurrence of a catastrophic system event

In addition to the five specific meanings of the LEDs, related system status is encoded based on whether the LED is solid or flashing. Examples include unexpected reboot system recovered, operating system not running, and operator intervention required.

event monitoring service

HP EMS is a system monitoring application designed to facilitate remote/centralized real-time monitoring and error detection for HP products in the enterprise environment. This framework provides centralized management of hardware devices such as the rp5400 series servers and system resources, and it provides immediate notification of hardware failures and system status. HP EMS can receive data on unusual activity, add information on the problem's source, and provide recommendations on problem resolution.

HP EMS consists of a set of system and network monitors within a monitoring environment. This monitoring framework has an easy-to-use interface and provides a mechanism for monitoring resources, registering monitoring requests, and sending notification when resources reach user-defined critical values.

How it works:

- A hardware event monitor detects abnormal behavior in one of the hardware resources (devices) it is monitoring.
- The hardware event monitor creates the appropriate event message, which includes suggested corrective action, and passes it to the EMS.
- EMS sends the event message to the system administrator using the notification method specified in the monitoring request (for example: e-mail, message to the console, entry in a system log).
- The system administrator (or HP service provider) receives the message, corrects the problem, and returns the hardware to its normal operating condition.
- If the peripheral status monitor (PSM) has been properly configured, events are also processed by the PSM. The PSM changes the device status to DOWN if the event is serious enough. The change in device status is passed to EMS, which in turn alerts MC/Serviceguard. The DOWN status will cause MC/Serviceguard to failover any package associated with the failed hardware resource.
- Any of the following consoles can be used with EMS to remotely monitor server farms: HP MC/Serviceguard, CA Unicenter, HP OpenView ITO, HP Secure Web Console, and HP Toptools.

The monitors can also poll hardware, disks, clusters, network interfaces, and system resources and send information to the framework. An "event" can be simply defined as something you want to know about—for example, a disk failure or file space dropping below a predefined level.

The primary EMS benefits include:

- enables efficient and effective system monitoring within a single, comprehensive framework
- delivers the ability to tailor the monitoring system to fit specific needs

18.432
A.

- provides a wide variety of notification methods through multiple protocols (SNMP traps, TCP, UDP, OPC messaging)
- provides immediate alerts if a component fails, enabling proactive replacement
- integrates with HP MC/Serviceguard and Serviceguard OPS Edition to provide a complete high-availability solution

extended fault management system

The rp5400 series employs a dedicated processor to aid system management and diagnosis. The extended fault management system can diagnose a system failure even in the unlikely event that the system is unable to execute code. It allows system power to be remotely turned on or off, and it has battery backup that even allows diagnosis of power failures. The system interfaces with key components via an inter-integrated circuit (I^2C) bus to continually monitor the status of system fans, temperature, and power supplies; it signals the operator if any significant system events occur.

Major features of the extended fault management system include:

- system console redirection
- console mirroring
- configuration of system for automatic restart
- viewing history log of system events
- viewing history log of console activity
- setting inactivity timeout thresholds
- remote system control
- power control—remote power on and off
- viewing system status logs
- configuration of virtual front-panel display
- event notification to system console, e-mail, pager, and/or HP Response Centers
- auto system restart
- virtual front-panel display
- password security (same level as UNIX)

system platform monitor

Closely integrated with the extended fault management system is the system platform monitor. The system platform monitor controls and monitors system power and cooling. Aspects controlled and monitored by the system platform monitor are:

- power supply status and temperature
- system supply voltages—including remote system power on and off
- total system power consumption
- individual Processor Support Module status
- external ambient air temperature
- individual fan speed and status

Various temperatures are monitored to control the system fans, provide thermal warnings, and prevent permanent damage from overheating by graceful shutdown if the temperature is too high. (Note that the system fans are run only as fast as necessary to keep the system cool. The fans are kept in sync with each other, turning at exactly the same rate. This intelligent fan control allows the L-Class to generate as little noise as possible while maintaining an optimum operating environment to maximize reliability.)

The power monitor senses the presence of power supplies and the power consumption of system components to determine if the system is in an N or N+1 power configuration; it can determine:

- number of bulk power supplies
- number of CPUs
- amount of memory present
- number and power consumption of each installed PCI I/O card

System configuration and health is tracked by the system platform monitor and passed via a dedicated I^{2}C bus to the fault management processor. This information can be processed as follows: simply displayed on the system console, logged to an event file, or used to trigger an alert based on a specific threshold (system temperature, fan status, or power supply status, for example).

built-in Web console

The rp5400 series has integrated Web console functionality, which allows management of many systems from a single Internet browser. The Web console is embedded into the fault management processor and can be accessed through the core 10/100Base-TX management LAN. The external Secure Web Console box that shipped with older-generation HP servers is no longer needed. The Secure Web Console allows an Internet browser to be used as a system console, giving total system access to authorized system administrators anywhere, just as if they were at an ASCII console. A high level of password protection is used to control access to the Web console.

Major features of Secure Web Console include:

- system management over the Internet or intranet
- mirrored access—up to four operators can simultaneously share the same screen and keyboard
- security—built-in password encryption, data scrambling, and Java™ download protection
- universal browser-based support for Netscape v.3.0+ and Microsoft® Internet Explorer v.3.0+ Web browsers
- easy updates of Web console software over the network
- easy installation—just connect the L-Class console part to a LAN; there is no client software to install
- support for HTTP, FTP, TFTP, and other key Internet standards

LAN console

The server also provides a LAN console interface using industry-standard telnet connections. Like the Web console, the LAN console can be used remotely for managing many systems from a single control center. The telnet interface allows scripts to be used to vastly simplify multiple system management. Password protection provides a high level of security to control access to the LAN console, ensuring that only authorized personnel perform system management.

ASCII consoles

For users who wish to locally administer their systems, the rp5400 series provides an RS-232 port to use for ASCII terminal console connections. Any VT100-capable terminal or emulator can be used as a local system console.

remote access

As with previous HP server systems, an RS-232 interface for a remote console is useful for obtaining help from HP service experts. Customers need only add a modem to allow remote access via phone; security is ensured by having to explicitly enable remote console access, which is protected with a password, and via dial-back phone verification.

self-diagnosis

Many features have been designed into the server to maximize system uptime. There are several aspects to maximizing uptime: eliminating common single points of failure, allowing the system to continue running after some errors, and allowing quick identification and servicing of hardware faults if they do occur.

Besides using traditional diagnostic software, the server also continuously monitors system health with the platform monitor. Knowing a failure has occurred that reduces N+1 protection is important. It is important to minimize the risk of downtime by replacing a failed component as soon as possible to get back to the safety of an N+1 configuration. To enable this, the server provides several methods of event notification.

The rp5400 series has extensive firmware-based self-tests. These diagnostics are evoked on power-up or reset. The self-tests check for correct system operation prior to booting the operating system. The firmware diagnostics first check the processors, then processor caches and memory, and finally I/O devices. Testing complexity increases as more of the system is proven good and more pieces of the system can be relied upon to increase test coverage on the remaining parts. Self-test failures are reported to the system console and the support processor, along with failure specifics and recommended corrective action.

online and offline diagnostics

The rp5400 also offers traditional online and offline diagnostics to validate system health and provide extensive system fault coverage.

With online diagnostics, the system is tested while the operating system and applications continue to run. This allows basic testing of system components that are not currently being used, or it allows testing in situations where the testing does not prevent continued use of the operating system and applications.

Offline diagnostics provide increased coverage of system components for improved fault isolation and intensive system testing before returning to production.

18.429
A.

for more information

HP product information and technical documentation is available online at:

http://www.hp.com/go/rp5430

http://www.hp.com/go/rp5470

Contact any of our worldwide sales offices or HP Channel Partners (in the U.S., call 1-800-637-7740) or at the following international numbers:

United States of America:
+1 800 637 7740

Canada:
Hewlett-Packard Ltd.
5150 Spectrum Way
Mississauga, Ontario L4W 5G1
+1 905 206 4725

Japan:
Hewlett-Packard Japan, Ltd.
Japan Country H.Q.
3-29-21, Takaido-Higashi, Suginami-ku,
Tokyo, 160-8585 Japan
+81 3 3331 6111

Latin America:
Hewlett-Packard
Latin American Region Headquarters
Waterford Building, 9th Floor
5200 Blue Lagoon Drive
Miami, Florida 33126 USA
+1 305 267 4220
Refer to country phone numbers

Australia/New Zealand:
Hewlett-Packard Australia Ltd.
31-41 Joseph Street
Blackburn, Victoria 3130
Australia (A.C.N. 004 394 763)
+61 3 9272 2895

Asia Pacific:
Hewlett-Packard Asia Pacific Ltd.
17-21/F, Shell Tower
Times Square
1 Matheson Street
Causeway Bay
Hong Kong
+8522 599 7777

Europe/Africa/Middle East:
Hewlett-Packard S.A.
150, Route du Nant-d'Avril
CH-1217 Meyrin 2
Geneva, Switzerland
+41 22 780 81 11
European Multicountry: +41 22 780 81 11
Middle East and Africa: +41 22 780 71 11
European Headquarters: +41 22 780 81 81

For direct country contact call:

Argentina: +541 787 7145

Austria: +43 1 25 000 0

Belgium and Luxembourg: +32 2 778 31 11

Brazil: +5511 7296 8000

Chile: +562 203 3233

Colombia: +571 629 5030

Denmark: +45 45 99 10 00

East Central Europe, CIS, and Yugoslavia: +43 1 25 000 0

Finland: +358 9 887 21

France: +33 1 69 82 60 60

Germany: +49 7031 140

Greece: +30 1 689 644

Hungary: +36 1 252 7300

Iceland: High Performance Systems hf. +354 1 67 10 00

Ireland: +353 1 615 8200

Israel: Computation and Measurement Systems (CMS) Ltd. +972 3 5380 333

Italy: +39 2 92122770

Mexico: +525 326 4600

Netherlands: +31 20 547 6911

Norway: +47 22 7356 00

Poland: +48 22 608 77 00

Portugal: +351 1301 7343

Russia and the CIS, excl. Ukraine: +7 095 923 5001

Slovenia: +38 61 55 84 72

Spain: +34 1 631 1600

Sweden: +46 8 444 2000

Switzerland: +411 735 7111

South Africa: Hewlett-Packard South Africa (Pty) Ltd.+27 11 806 785 1000

Turkey: +90 212 224 5925

United Kingdom: +44 1344 369231

Venezuela: +582 239 4133

UNIX is a registered trademark of The Open Group. Intel and Itanium are U.S. registered trademarks of Intel Corporation in the U.S. and other countries and are used under license. Microsoft is a U.S. registered trademark of Microsoft Corporation. Netscape is a U.S. trademark of Netscape Communications Corporation. Java is a U.S. trademark of Sun Microsystems, Inc.

Technical information contained in this document is subject to change without notice.

08/02

5981-2650EN

18.428
A.

ANEXO UNIDADE DE BACKUP ROBOTIZADO
PARTE I

COBRA Tecnologia S.A.
Estrada dos Bandeirantes 7966
CEP 22783-110 Rio de Janeiro RJ
Tel. 21 2442-8800
www.cobra.com.br

1 / 1

18.427
A.

» hp home | » products & services | » support & drivers | » solutions | » how to b

» contact hp

search: [ ] » more o|

● software depot ○ all of hp US

printable ve

invent

mirrordisk/ux license for servers

» software depot

» electronic (re)download
» how to buy support
» hp-ux OE product information
» hp-ux 11i promotions

product details & specifications

overview

MirrorDisk/UX License for Servers
Introduction

MirrorDisk/UX software prevents data loss due to disk failures by maintaining up to thre data on separate disks. Applications can continue to access data even after a single dis addition, you can perform on-line backups to avoid user and application disruption.

To prevent the failure of a single I/O interface from causing a system failure, HP recomn mirrored disks be connected to separate interface cards.

Features and Benefits

- No single point of failure - separate controllers/power supplies
- Up to 3-way disk mirroring
- On-line backup while maintaining mirroring
- Application transparency
- Dynamic mirror configuration
- Selective mirror of data
- Fast data synchronization
- Menu-driven administration tools

Because high availability solutions require full-time access to data, HP has developed M to provide mirroring capability within the HP-UNIX® Operating System (HP-UX) environ

MirrorDisk/UX, the mirroring component of Logical Volume Manager (LVM) from the Op Foundation (OSF™), prevents data loss by maintaining up to three copies of data on se This enables data to remain intact after a single disk or interface card failure.

MirrorDisk/UX can mirror a disk partition, including the root and swap partitions. It suppo of raw disk access as well as file system access. MirrorDisk/UX can increase input/outp performance depending upon the mix of disk reads and writes.

Because MirrorDisk/UX works with the HP-UX kernel to manage the mirrored disks, it is to the applications, which require no modification.

MirrorDisk/UX allows customers to perform a backup by taking one disk of a mirrored p while the other disk continues to service applications. As the backup is being performed occurring to the online disk are maintained in table memory. Upon completion of the bac procedure, a faxt update is done to synchronize the disks while application continuity is To prevent the possibility of data loss caused by a failure during a backup, three-way mi provides a fully redundant mirrored pair while the third copy is being backed up.

Date: 6/21/99

18.426
A.

additional product information

product #:	B2491BA
version:	-
software specification:	Server, HP-UX11.00 Server, HP-UX11.11
price:	$1017.00 Per Processor License

» support

buy fro

privacy statement using this site means you accept its terms

13.425
A.

ANEXO UNIDADE DE BACKUP ROBOTIZADO
PARTE J

18.424
(A.

hp invent

hp rack-optimized rp5430 and rp5470 servers entry-level UNIX® servers

hp servers product brief

hp rack-optimized servers can help you be always on, always there, always connected

hp servers rp5430 and rp5470: smart, simple, stress-free

hp servers rp5430 and rp5470

In today's economy, whether you're managing your own IT infrastructure or hosting someone else's, you have to operate with a faster time-to-solution, within budgetary constraints, and with the highest standards for customer service and operational efficiency.

To create and run an infrastructure for an always-on business, you need a computing platform that will support the way you—and your customers—do business. The HP Servers rp5430 and rp5470 give your business the fastest—and most reliable—means of succeeding in this new business environment.

The HP Servers rp5430 and rp5470 deliver the proven performance, scalability, and high-availability capabilities of UNIX—without high maintenance requirements and costs. And they give you plenty of room to grow. You can start at a low-price entry point and scale up to the leading 4-way UNIX performance—in the same rack-optimized form factor, without penalty. And with its industry-leading solution partners, HP has developed business solutions surrounding these servers that are tested, easy to deploy, and easy to manage.

With HP Servers rp5430 and rp5470, owning and operating a UNIX server is smart, simple, and stress-free.

smart

HP Servers rp5430 and rp5470 offer leading entry-level server performance, dynamic scalability, and unmatched investment protection—all in a rack-optimized package—making them the smart choice for the most demanding applications.

leading performance—scalable functionality

- industry-leading OLTP performance
- massive bandwidth for I/O-intensive applications
- 7U-height and packed with CPUs, memory, and I/O, plus the ability to scale subsystems without compromise
- rack-optimized to make the best use of valuable data-center floorspace

unparalleled investment protection

- industry's only in-box upgrade from 2-way to 4-way UNIX computing
- built-in growth path to the HP Server rx5670, featuring Intel® Itanium® 2 processors

13.423
A.

proven solutions for business-critical computing

simple

hp makes IT easy

With HP Servers rp5430 and rp5470, HP offers a unique combination of solutions designed to help you get started quickly and manage your IT environment effectively.

hp-ux: robust, proven enterprise operating environment for mission-critical applications

- industry-leading performance, scalability, availability, manageability, and security
- pre-packaged, integrated, and production-ready operating environments
- industry-leading Windows® and Linux interoperability
- centralized software updates that are timely, simple, and efficient
- powerful alliances with industry-leading software vendors and systems integrators to deliver robust solutions from e-commerce to enterprise resource planning (ERP) and beyond

leading-edge management capabilities with hp-ux virtual partitions and hp-ux workload manager

- system resource optimization enabling multiple workloads to run simultaneously on the same server, each with their own instance of HP-UX
- improved security and server availabilty through complete software and operating system isolation
- HP-UX Virtual Partitions integrate with HP-UX Workload Manager for the most efficient resource distribution across partitions, in a single server
- base offering complementary with HP-UX 11i for your HP Server rp5470

integrated management capabilities

- HP Servicecontrol Manager and integrated HP Secure Web Console capability for full remote management, including centralized configuration of multiple servers

the right server for today's applications

ERP (supply-chain management)

High availability, leading-edge manageability, and scalable performance support demanding end-to-end enterprise applications.

broadband

The combination of leading performance, I/O throughput, and capacity and high availability with end-to-end solutions delivers more powerful, reliable broadband services.

Internet infrastructure

Highly scalable, reliable, and manageable Web server, caching server, load balancing, e-commerce server, firewalls, or mail server.

e-commerce

Leading performance and I/O bandwidth, in-box scalability, rack-optimized form factor (five per standard 2-meter rack), high availability, and Internet management features.

technical computing

Leading performance and throughput, N+1 redundant components, extensive memory capabilities, and clustering solutions to meet the demands of computation, NFS file serving and product data management, and Web hosting.

flexible financing

- operating leases with a variety of attractive terms
- the Tech Refresh program for cost-effective upgrades to stay on the leading edge of new technology
- bundled Solution Finance program to consolidate and simplify financing arrangements

instant capacity on demand for your hp server rp5470

iCOD

- instant activation of incremental CPU power when you need it
- pay only for the processing power you use

temporary capacity for iCOD

- temporary activation of incremental CPU power for a limited period
- ideal for short-term, predictable processing demands

stress-free

make your business your focus

The HP Servers rp5430 and rp5470 handle the demands of users efficiently and reliably—so you can concentrate on running your business, not managing your IT resources.

high availability for continuous operations

- a rich set of in-box high-availability features
- affordable high-availability clustering solutions based on industry-leading HP MC/Serviceguard
- self-healing capabilities, a first for entry-level servers—proactively avoid faults to improve uptime
- a critical building block for your always-on e-business needs

total solution support

- options ranging from Web-based services to the industry's only 6-hour call-to-repair commitment
- "one-stop" solution support delivered with partners such as Cisco, Oracle®, SAP, i2, Inktomi, and many more

best UNIX server family—top to bottom

The low-cost entry point to the two-way HP Server rp5430 through the more scalable HP Server rp5470 are part of the powerful HP UNIX server line—servers that set the standards for business-critical computing and total cost of ownership. Simple in-chassis upgrades let you move up the line as your business grows. HP UNIX servers provide the hardware foundation for an Internet infrastructure that is always on. Combining leading technology with proactive and reactive services, HP offers complete, end-to-end solutions that include hardware, software, applications, services, support, consulting, and an extensive portfolio of experienced partners, so you can get to market quickly with a single source of expertise. The HP UNIX server family is robust from the top to bottom—from the high-end HP Superdome; through the HP Servers rp8400 and rp7410, the midrange performance and price/performance leaders; and to the scalable entry-level HP Servers rp5430 and rp5470. Rounding out the family are the hyper-dense HP rp2400 series servers for Internet applications and branch offices.

18.422
A.

hp servers rp5430 and rp5470 features and benefits at-a-glance

features	benefits
smart	
1–4 (rp5470) or 1–2 (rp5430) 875MHz PA-8700+ CPUs with 2.25 MB on-chip cache per CPU or 750MHz PA-8700 CPUs with 2.25 MB on-chip cache per CPU	Superior performance over comparable systems, with plenty of headroom for growth
Up to 16 GB (rp5470) or up to 8 GB (rp5430) SDRAM memory with advanced ECC protection	Fast and reliable processing power for frequently accessed data
Up to 10 (rp5470) or 6 (rp5430) PCI I/O slots with 3.2 GB/s (rp5470) or 2.3 GB/s (rp5430) I/O bandwidth	Easily handles I/O-intensive applications and allows the system to scale I/O, CPUs, and memory without compromise
Core I/O, including 10/100Base-T LAN with auto speed-sensing, a second 10/100Base-T support LAN, Ultra2 LVD SCSI, and RS-232	Provides easy, ready-to-go networking capabilities
Up to 4 internal 36 GB, 73 GB, or 146 GB Ultra320 SCSI hot-plug high-uptime disks	Store critical data with massive internal capacity
1 internal DVD or DAT drive	Protects critical data
7U chassis with up to 5 servers per standard 2-meter rack; also available in standalone (pedestal) configuration	Optimizes use of floor space and delivers high-performance density in a racked configuration
Easy in-box upgrades from the rp5430 to the rp5470	Architectural scalability ensures these servers can grow with the business, maximizing flexibility and investment protection
Support of Intel Itanium Processor Family as well as PA-RISC processors	Provides investment protection through dual growth paths
simple	
Built-in unlimited user license for proven 64-bit HP-UX 11i and 11.0	Proven, enterprise UNIX operating system for mission-critical applications
HP Virtual Partitioning for the rp5470	Maximizes usage of computing resources
Integrated HP Secure Web Console and Servicecontrol Manager for full local, Web, and remote control of servers	Provide complete single-system and multisystem administration capabilities, including a range of security features, from any browser-based PC
Flexible financing programs	Make initial ownership and modular growth easy and affordable
Instant capacity on demand (iCOD); temporary capacity for iCOD	Immediate access to CPU power when you need it, either permanently or temporarily
HP global deployment and partner integration services	Offer guaranteed error-free solution deployment to reduce implementation time and cost
HP On-site Solution consolidated manufacturing, streamlined product assembly and testing, and state-of-the-art integration	Ensure superior quality and faster delivery
stress-free	
Error-correcting cache, parity checking on all buses, memory scrubbing and page de-allocation, dynamic processor resilience, and de-allocation of application processes Dual Ultra 2 SCSI buses and controllers for mirrored storage Hot-swap, redundant power supplies and fans; redundant, hot-plug PCI; Ultra2 SCSI hot-plug disks	Built-in high-availability features deliver superior levels of • error correction, • error containment, • data protection, and • serviceability to help maximize uptime for business-critical workgroups and applications
Integrated Event Monitoring Service (EMS) Built-in fault management system with separate support processor and bus Integrated with HP MC/Serviceguard, HP Toptools for Servers, and enterprise management software such as HP OpenView and CA Unicenter	Provides superior system uptime through constant, proactive fault avoidance, detection, and notification; monitors power, cooling system hardware, processors, memory, HP-UX resources, and external storage
Pre-tested and pre-integrated workgroup clustering solutions based on HP MC/Serviceguard	Deliver complete, ready-to-go solutions for clustered high availability that eliminate all single points of failure, at an affordable price
3-year on-site service warranty and HP services and support options ranging from Web-based support to mission-critical, 6-hour Call-to-Repair commitment; includes full solution support for hardware and software	Reduce risk through worldwide support for business-critical computing; provide "one-stop shopping" for support through partnerships with Cisco, Oracle, SAP, BroadVision, Inktomi, and others

13.4.21
A.

stay ahead of the curve

configuration options	hp server rp5430	hp server rp5470
processor	875MHz PA-8700+ or 750MHz PA-8700	875MHz PA-8700+ or 750MHz PA-8700
SMP configuration	1 to 2 CPUs	1 to 4 CPUs
supported OS versions	HP-UX 11i PA-8700+ HP-UX 11.0, 11i PA-8700	HP-UX 11i PA-8700+ HP-UX 11.0, 11i PA-8700
minimum/maximum memory	512 MB/8 GB	512 MB/16 GB
on-chip cache (data/instr)	1.5 MB/750 KB	1.5 MB/750 KB
total I/O slots	6	10
internal removable media bays	1	1
maximum internal disk capacity	584 GB (4 disk bays)	584 GB (4 disk bays)
standard I/O features	Ultra2 SCSI, 100Base-T LAN, 3 RS-232 ports, 100Base-T support LAN, and Web-based console	
supported I/O connectivity	Ultra2 SCSI RAID, Ultra2 SCSI LVD (single- and dual-port), FWD SCSI (single- and dual-port), Fibre Channel, Gigabit Ethernet, 100Base-TX (single- and quad-port), 100Base-FX, ATM 155 Mb/s (MMF, UTP-5), ATM 622 Mb/s (MMF), FDDI dual-attach LAN, Token Ring 100 Mb/s; X.25/FR/SDLC (dual-port), multiplexer (8- and 16-port)	

environmental specifications	
electromagnetic interference	Complies with FCC Rules and Regulations, part 15, as a Class A digital device; Manufacturer's Declaration to EN 55022 Level A; VCCI Registered, Class I; Korea RLL
AC input power	100–240V 50/60Hz
maximum current requirements	13.8A at 110V
maximum power dissipation	1283 watts
physical dimensions	Depth: 774 mm (30.5 in) Width: 482 mm (19.0 in) Height: 368 mm (14.5 in)/7 EIA units Weight: 68 kg (150 lb)
operating temperature	+5° to 35°C (41° to 95°F)
nonoperating temperature	–40° to 65°C (–40° to 149°F)
maximum rate of temperature change	20°C/hour
operating relative humidity	15% to 80%, noncondensing, max. wet bulb @ 26°C
nonoperating relative humidity	5% to 90%, noncondensing
operating altitude	To 3.0 km (10,000 ft) above sea level
nonoperating altitude	To 4.5 km (15,000 ft) above sea level
compliance model number	RSVL - 0105-A

for more information

Contact any of our worldwide sales offices or HP Channel Partners (in the U.S. call 1-800-637-7740) or visit the HP servers Web site at **www.hp.com/go/servers** or **www.hp.com/go/rp5430** or **www.hp.com/go/rp5470**

HP product information and technical documentation are available online. In addition, configuration tools and pricing information allow registered users to place orders online.

06/12/03

5980-8432EN Rev. 6

hp®
i n v e n t

ANEXO UNIDADE DE BACKUP ROBOTIZADO
PARTE K

18.4.20
A.

18.419
A.

hp serviceguard extension for RAC

invent

proven, flexible high-availability solution

the challenge

To maintain and manage a highly available mission-critical parallel database environment

the solution

HP Serviceguard Extension for RAC

business benefits

- data protection
- application availability
- ease of management
- flexible and easy deployment of applications
- expert consulting and support services

To best meet data center requirements for availability, flexibility, and scalability, HP offers a robust architecture that combines multiple computers into entities called "clusters." The systems—or nodes—of a cluster are connected in a loosely coupled manner, each maintaining its own separate processor(s), memory, and operating system. Special communications protocols and system processors bind these nodes together and allow them to cooperate to provide outstanding levels of availability and flexibility for supporting mission-critical applications.

Clusters maintain strict compliance to the principles of open systems. There are no propriety APIs that force vendor lock-in and require substantial development investment. Most applications will run in a cluster without any modification at all. And only standards-based hardware components such as SCSI or Fibre Channel storage devices and FDDI Ethernet LANs are used to create a cluster.

18.414
A.

Clusters provide a cost-effective, flexible architecture for meeting the demanding requirements of the commercial UNIX® market.

Building on the superior capabilities of HP Serviceguard, HP Serviceguard Extension for RAC (formerly called Serviceguard OPS Edition) allows a group of HP 9000 servers to be configured as a highly available cluster that supports Oracle9*i* Real Application Cluster (RAC) on both HP-UX PA-RISC and Intel® Itanium™ Processor Family platforms. These two products are tightly integrated to provide the best aspects of HP enterprise clusters and Oracle® relational database servers: high availability, data integrity, flexibility, scalability, and reduced database administration costs.

scaling beyond one system

Oracle9*i* RAC-enhancing scalability

With the introduction of Oracle9*i* Real Application Cluster (RAC), the unique architecture of a Serviceguard Extension for RAC cluster enables the full aggregate processing power of up to 16 nodes to access the database, increasing overall throughput for certain kinds of applications. Examples of applications that can benefit from the performance gains and scalability of this cluster include query-intense applications such as decision support, applications that generate random reads and writes to very large databases, and applications that access separate partitions of the database.

enhancing performance

The new HP Hyper-messaging Protocol with Oracle9*i* RAC provides high-bandwidth, low-latency server-to-server communication that promises to deliver superior interconnect performance.

serviceguard extension for RAC

Serviceguard Extension for RAC provides all of the functionality needed to support the Oracle9*i* RAC environment on HP 9000 servers. The major components of Serviceguard Extension for RAC are listed here.

- **cluster manager:** establishes and monitors the cluster members and monitors various components within each node
- **cluster membership:** informs RAC about system failures to facilitate fast database recovery
- **package manager:** monitors and controls packages containing highly available applications
- **network manager:** detects and recovers from card and cable failures
- **shared logical volume manager** (ships with HP-UX): provides the basic functionality to share physical disks and buses between the nodes

Serviceguard Extension for RAC also provides further enhancements to the environment needed to support mission-critical applications. Special functionality is included to significantly enhance the availability of each node within the cluster and to provide extra protection for database integrity. Since high availability is a primary design goal, this cluster has been created with no single point of failure. The data disks are mirrored via a disk array, and multiple LANs are used.

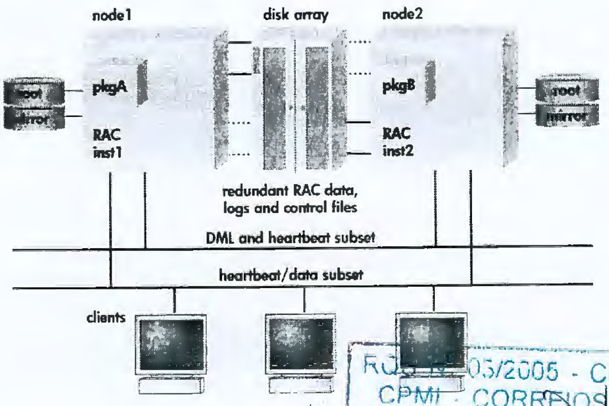

13.413
A.

key features and benefits

data protection

- [illegible] prevent [illegible] from [illegible] the integrity of application data

availability

- local area network (LAN) monitoring ensures quick, transparent recovery to maximize database availability
- quick automatic detection and recovery time maximizes application availability and minimizes operator error
- the ability to survive multiple node failures provides a unique level of protection
- fast failback [illegible] applications in the primary [illegible] node
- rolling upgrades ensure application availability during hardware and software maintenance
- integration with HP Workload Manager ensures service-level objectives (SLOs) are maintained during planned and unplanned downtime
- tight integration with Oracle's specialized RAC HA Extension configuration offers quicker detection and failover
- the Enterprise Cluster Master toolkit provides quick and easy deployment of applications
- extended-distance clustering (up to 10 km) provides disaster-tolerant protection

flexibility

- multiple cluster configurations—active-active, active-standby, and rotating standby—offer flexibility

manageability

- an intuitive graphical user interface, Serviceguard Manager, reduces the total cost of managing multiple clusters from a single console
- support for virtual and hard partitioning addresses the increasing demand for systems consolidation

ROI

- integration with HP Pay per Use offers a cost-effective disaster-tolerant solution that is unique to HP

one call, one voice

hp and Oracle: joint facilities

To provide consulting and support expertise for customers working with enterprise clusters, HP and Oracle maintain joint facilities staffed with technical experts. These facilities are available for developing proofs-of-concept for new projects, executing benchmarks, and performing sizing activities for enterprise clusters.

support services

HP understands that ensuring a highly available mission-critical environment means more than just having the right technology—it is just as critical to have the right IT processes and support services in place. That's why HP has created a comprehensive portfolio of services, ranging from consulting, education and training, and disaster planning and recovery services to mission-critical support for HP hardware and software products—including Serviceguard Extension for RAC.

- **personalized systems support (PSS):** a comprehensive support solution. PSS combines proactive account services with industry-leading technical assistance to help you improve operational effectiveness and successfully manage and implement change within your IT environment.
- **business continuity support (BCS):** emphasizes downtime prevention through continuous improvement of your IT infrastructure, the best mission-critical processes, and constant vigilance. BCS starts with an extensive assessment to identify and analyze areas that put availability and service-level requirements at risk, and then it presents recommendations to minimize these risks.
- **critical systems support (CSS):** provides technical expertise through an integrated combination of proactive services and fast problem resolution to meet the demands of your computing environment. Also available through CSS is a portfolio of technical services designed to minimize system problems and downtime and to help you make more effective use of technology. Areas of interest include high-availability technologies, performance analysis, change planning, security review, and system administration.

18.412
A.

value-added hp services

Because continued proactive management of your cluster is critical to ensuring high availability HP has introduced an additional technical service to help increase cluster availability and stability.

- **hp cluster consistency service (H8395AT)**—delivers a diagnostic tool for spotting potential disruptions to critical applications by identifying cluster configuration problems before they occur.

benefits of hp services

By taking advantage of HP services, you will realize major benefits:

- reduced implementation time with expertise, partnership, and hands-on assistance from your assigned high-availability–certified HP team
- decreased IT crisis risk and exposure through robust proactive services
- increased availability due to improved environment stability and rapid problem resolution through established and proven processes
- maximized end-user productivity thanks to optimized processes and system performance
- comprehensive planning support using proven HP methodology

more information

For more information, please visit our Web sites at
www.hp.com/hps
www.hp.com/go/financialservices

www.hp.com/go/ha

invent

Intel and Itanium are trademarks or registered trademarks of Intel Corporation in the U.S. and in other countries and are used under license. Oracle is a registered U.S. trademark of Oracle Corporation, Redwood City, California. UNIX is a registered trademark of The Open Group.

Technical information contained in this document is subject to change without notice.

© 2002 Hewlett-Packard Company
12/02
5981-4712EN

18.411
A.

ANEXO UNIDADE DE BACKUP ROBOTIZADO
PARTE L

18.410
A.

HP OpenView Storage Data Protector 5.1 software

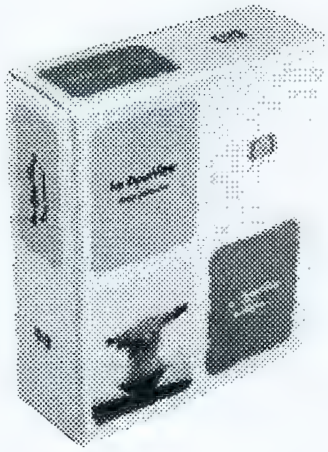

HP OpenView Storage Data Protector 5.1 software provides enterprise data protection and disaster recovery, ensuring recovery from any disruption. Data Protector integrates a variety of techniques to eliminate planned downtime ranging from online backup and backup of open files to zero-downtime, zero-impact backup. The software provides industry-leading instant recovery, and several disaster recovery alternatives to eliminate unplanned downtime, allowing recovery of entire data centers in minutes.

Data Protector enables organizations to reduce costs by automating routine tasks, thus maximizing the effectiveness of data protection staff. It gives you enterprise functionality at an entry-level price. Data Protector software scales from a single server to a distributed enterprise, covering an extensive range of heterogeneous applications, operating environments and storage configurations all in a single solution.

Key features and benefits

Controllable:

- *Efficient:* offers ease of deployment and use with single consistent GUI across platforms and functions
- *Affordable:* offers value-for-price with operational efficiencies for faster ROI; tailored to decrease overhead and costs
- *Automated:* eliminates operator involvement by application-aware automation, e.g., built in software distribution
- *Manageable:* allows central management, scheduling and cataloging of protection features
- *Integrated:* includes the "manager-of-managers" feature that controls multi-site or global operations
- *Comprehensive:* through the HP OpenView software infrastructure provides enterprise-wide management capabilities with a service-oriented environment

Resilient:

- *Available:* increases information availability by eliminating backup and recovery windows
- *Comprehensive:* offers comprehensive approaches for disaster tolerance and disaster recovery
- *Reliable:* provides full cluster-awareness, designed for 24 x 7 protection, including online protection of backup environment

Extensible:

- *Open:* supports a broad range of hosts, applications, storage, and protection approaches, allowing you to grow as your needs grow
- *Flexible:* offers a simple and modular structure, broad compatibility with platforms, operating systems, libraries, disks and topologies
- *Scalable:* offers the most scalable product in the market with a single solution for small, medium and large enterprises

The HP OpenView Storage Data Protector software gives you enterprise data protection and automates routine tasks to maximize the effectiveness of your data protection staff. It ensures recovery from any potential disruption, from instant recovery through system or site disaster recovery.

18.409
A-

HP OpenView Storage Data Protector 5.1 software puts you in control with automated routine tasks — you get maximum protection at the lowest cost.

What's new

Enhanced functional capabilities for more control

- Adds point and click GUI for the restore of Oracle8 and Oracle9i databases: administrators can select all or individual RMAN restore options from the GUI freeing them from dealing with the Oracle RMAN Command Line Interface
- Offers simplified Microsoft® Exchange single Mailbox/Mail restore: Mailbox users can select all or individual mailboxes for single backup and restore; copy individual objects, e.g., single mails, from the backup to the current mailbox
- Offers simplified SAN auto-configuration wizard: automatically detects and configures the backup drives in the SAN on all desired SAN attached systems. The new "group by devices" and "group by hosts" views allow optimized shared device access
- Allows single backup specification: enabling administrators to schedule different protection times for different backup instances using the same data definition. Disk and tape backup can now be scheduled within the same backup specification
- Enables use of external encryption capabilities for enhanced security

Increased resilience with recovery from any disruption; from instant recovery to site or system disaster recovery

- Eliminates backup windows for HP StorageWorks EVA customers with the industry's first fully-integrated out-of-the-box zero-downtime backup solution for the HP StorageWorks EVA array
- Eliminates recovery windows for EVA customers by integrating its instant recovery capability to enable the recovery of even terabytes of application data in minutes, not the hours it would normally take to restore this amount of data from tape
- Incorporates latest Microsoft Windows® Server 2003 backup features into customer's environments; fully integrates Microsoft Volume Shadow Copy Service to streamline backup of open files, databases and applications; incorporates support for Microsoft Automated System Recovery
- Expands the coverage of its unique disk-image delivery recovery approach and now includes AIX and HP Tru64 environments, recovering complete systems with minimal administrator time
- Simplifies preparations for offsite disaster recovery by enabling automated media copy as part of the protection process
- Supports Microsoft 2-8 node Cluster Server including the Network Load Balancing (NLB) to increase availability of application running on Microsoft Windows 2003
- Improves the availability of data through support of multi-pathing software (HP StorageWorks Secure Path, HP StorageWorks Auto Path, EMC PowerPath), to eliminate single point-of-failure of SAN components

Increased extensibility which scales from single server to enterprise

- Incorporates comprehensive protection of emerging IA-64 platforms (local, network and SAN-based): Windows Server 2003 on IA-64; HP-UX 11.23 on IA-64); Linux on IA-64 (Redhat, SuSE, Debian – network-based protection)
- Adds support for applications, operating and storage environments: DB2 online backup for IBM AIX and HP-UX; OpenVMS backup; Windows Server 2003 online backup of Oracle9i, SQL Server 2000; ZDB Backup on Windows Server 2003 for Oracle, Exchange, SQL and SAP

18403
A-

HP OpenView Storage Data Protector 5.1 software

Accessories

HP OpenView Storage Media Operations provides tracking and management of offline storage media, such as magnetic tapes, resulting in more reliable backups, faster data recovery, improved staff efficiency, and reduced costs.

Product number	Product name	Description
B7100AA	Storage Media Operations entry level LTU only	License to use Media Operations for up to 2,000 media pieces (one management server, unlimited clients)
B7101AA	Storage Media Operations enterprise level LTU only	License to use Media Operations for up to 10,000 media pieces (one management server, unlimited clients)
B7102AA	Storage Media Operations unlimited media LTU only	License to use Media Operations for unlimited media pieces (one management server, unlimited clients)
B7128AA	Storage Media Operations manual set	Combined installation, administration and user guides
B7129AA	Storage Media Operations distribution CD	Distribution media (CD)

Ordering information

	Part number				Description
Starter packs and management extensions	*Windows*	*Solaris*	*HP-UX*		
	B6961AA	B6951DA	B6951AA		Starter pack LTU, media and manuals (includes cell manager)
	B6961BA	B6951CA	B6951BA		LTU only
	B6966AA	B6956AA	B6956AA		Manager-of-manager LTU, 1x starter pack
	For any platform				
	B6960MA				Media kit
	B6960LA				Starter pack, manuals
	B6960EA				Functional extensions
Single server edition	*Windows*	*Solaris*	*HP-UX*		
	B7030AA	B7020DA	B7020AA		Single server edition LTU, media and manuals
	B7030BA	B7020CA	B7020BA		LTU only
	B7031AA	B7021DA	B7021AA		Upgrade to starter pack
Drive and library extensions	*SAN, NAS*	*Windows, NetWare, Linux (local attach)*	*UNIX*		
	B6953AA	B6963AA	B6953AA		Additional drive LTU (first drive included in starter pack)
	For any platform				
	B6957BA				Extended library LTU, 1x 61-250 slots
	B6958BA				Extended library LTU, 1x unlimited slots (libraries under 60 slots included in starter pack)
On-line backup and recovery extensions	*Windows*		*UNIX*		
	B6965BA (incl. Linux)		B6955BA		On-line backup LTU, 1x app. server
	B7033AA				Windows open file backup LTU, 1x server
	B7033BA				Windows open file backup LTU, 1x 10 servers
	B7034AA				Windows open file backup LTU, 1x 5 workstations
	B7035AA (incl. NetWare)				Windows open file backup LTU, 1x enterprise server
Disk-based protection	*NDMP*	*HP XP and compatible*	*New HP EVA/VA*	*EMC Symmetrix*	
		B7023CA	B7025CA	B6959CA	Zero downtime backup LTU (ZDB), 1x TB
		B7026CA	B7028AA		Instant recovery LTU (requires ZDB), 1x TB
	B7022BA	B7027AA*			Direct backup LTU (*requires ZDB), 1x TB

3

18.407
A.

HP OpenView Storage Data Protector 5.1 software

What's included

Data Protector software starter pack includes: License for a management system; unlimited number of backup agents (clients) on any platform; one concurrently used backup drive license (directly attached to a system, used in a SAN, NAS, or used for server-less backup); built in media management; library sharing between multiple systems and/or management systems; cluster support; SAN support (no additional licenses needed); disaster recovery alternatives; service-driven management through integrations into HP OpenView software; sophisticated reporting (in Data Protector GUI and via web); libraries with up to 60 slots (including robotic control – for control of tape drives and tape library systems); Customer Care Standard (5x8 unlimited telephone support, timely product upgrades, direct access to eCare. First-year included with each OpenView software product)

Warranty and support

Customer Care Standard (5 x 8 unlimited telephone support, timely product upgrades, direct access to eCare. First-year included with each HP OpenView software product). Coverage is upgradeable to 24x7 telephone support by purchasing Customer Care Extended. Both Standard 8x5 or Extended 24x7 Customer Care include unlimited telephone support during coverage hours, software update service, and access to eCare. Software Update Service proactively sends all major and minor software updates to the customer contact registered with the product. Our web-based eCare adds the following benefits:

Customers wanting a more personalized or mission critical offering can also order the following premier services:

- Customer Care Advantage: Proactive support for software management servers. Provides technical account advocate to proactively maintain customers' software management server(s). Benefits: Proactively prevents problems and reduces risks of unplanned downtime. Minimizes customers' support overhead. Designed for: Enterprise customers whose business is IT-dependent. Require proactive support to ensure stability and availability of software management servers.
- OpenView Premier and Critical Premier: Personalized service for software management environments. Provides named software engineer to partner with customer to manage their environments in the most cost-effective and efficient way. All of CCA benefits, plus: Named contact for reactive and proactive support activities. Trusted relationships, personalized knowledge and attention to customer needs. Maximizes customers' ROI in HP management software. Reduces risks of IT crisis. Minimizes customers' support and operational overhead. Designed for: Enterprise customers with business critical IT (extreme IT dependency). Require personalized support relationship to operate, manage, and evolve complex, business critical software environments.

For a complete list visit: www.hp.com/hps/storage

HP OpenView Storage Data Protector software Implementation Service and Zero Downtime Backup Implementation Service offered through HP Consulting

For more information

For more information on HP OpenView Storage Data Protector software, contact any of our worldwide sales offices or visit: www.hp.com/go/dataprotector

www.hp.com/go/dataprotector

invent

18 406
A

ANEXO UNIDADE DE BACKUP ROBOTIZADO
PARTE M

QuickSpecs

HP StorageWorks ESL9595 Tape Library

Overview

At a Glance

- **The Ultimate in High Capacity Unattended Backup and Restore**
  - The HP StorageWorks ESL9595 tape library provides the ultimate in high capacity unattended backup anc restore for customers with large server pools serviced by network attached storage devices or storage area networks (SAN). The consolidation of the enterprise's data into a single shared storage device requires that these devices provide both high capacity and high performance With the consolidation of backup data into a single device also comes the increased need for redundancy and high availability. The ESL9595 can provide years of fully automated operation for backup, save and restore, and hierarchical storage management of critical data when coupled with a variety of qualified industry standard backup application solutions. The HP StorageWorks ESL9595 tape library comes with up to 400, 500, and 595 slots respectively and up to 16 drives of the following technologies: LTO Ultrium 460, LTO Ultrium 230, SDLT 320 and SDLT 220, the ESL9595 offers up to 119 TB of native data storage and native transfer rates up to 1 7 TB/hr.

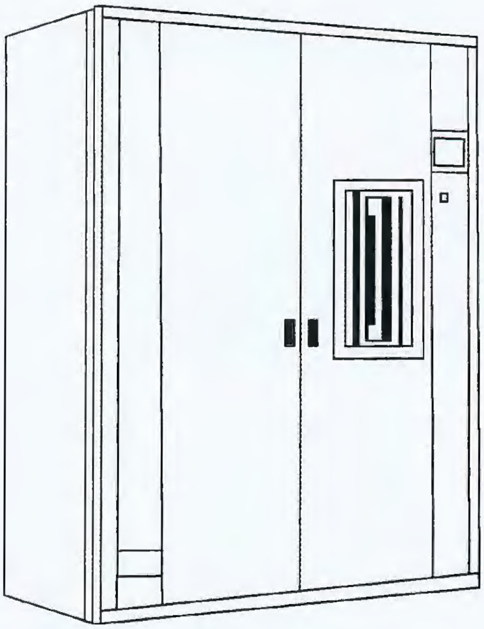

HP StorageWorks ESL9595 Tape Library

hp
invent

18.404
A.

Standard Features

HP StorageWorks ESL9595 Tape Library Models

The HP StorageWorks ESL9595 tape library comes with up to 400, 500 and 595 slots respectively and 16 drives of the following technologies: LTO Ultrium 460 tape drives, LTO Ultrium 230 tape drives, SDLT 320 tape drives, and SDLT 220 tape drives.

Product Numbers associated with Ultrium ESL9595 Enterprise Tape Libraries	
ESL9595, 400 slot, 0 drive, Enterprise Library	330833-B21
ESL9000 400 to 500 slot upgrade*	330842-B21
ESL9000 500 to 595 slot upgrade*	330842-B22
ESL9000 400 to 595 slot upgrade*	330842-B23
ESL9000 Ultrium 460 Drive Upgrade Kit	330834-B21
ESL9000 Ultrium 230 Drive Upgrade Kit	301930-B21
ESL9000 SDLT 320 Drive Upgrade Kit	293414-B21
ESL9000 SDLT 220 Drive Upgrade Kit	234617-B21

NOTE: *Serial number of your existing library is required when ordering a capacity upgrade SKU. Customers are required to also purchase the necessary number of host bus adapters, SCSI cables, media, country specific power cords (pair) and application software licenses in order to complete the library solution. Due to the size and weight of the unit, a pre-sales site survey is required.

NOTE: The ESL9595 is three inches wider than the ESL9326. Installation services can be ordered separately on Hewlett-Packard Services as part of each library order (requires scheduling with local Hewlett-Packard Services provider).

NOTE: Does not include special site work, which may be required as a result of not meeting minimum site survey requirements.

The ESL9000 Enterprise Library Family

The ESL9595 tape library is part of the HP StorageWorks ESL9000 class enterprise library family. With its high availability feature set, the ESL9000 family provides for maximum uptime. Based on PCI backplane architecture, the ESL9000 series libraries also provide a simple upgrade path for future "plug-in" feature upgrades.

High Availability

As data becomes more and more mission critical and downtime becomes more costly to the enterprise, customers are requiring high availability consolidated storage solutions which offer hot-swappable components and redundancies of key operating subsystems (tape drives, controllers, and power supplies). With redundant power distribution, power supplies, and cooling fans, coupled with "hot-pluggable" drive carriers, power supplies and cooling fan modules, the ESL9595 offers the best in high availability for continuous operation.

Hot-Pluggable Drives

The ESL9595 uses "hot pluggable" drive carriers, allowing drives replacement and additional drives to be added for capacity expansion without interrupting backup and restore functions. These hot-pluggable drive trays allow for easy scalability of performance, as well as provide an easy upgrade path towards future drive technologies.

The ESL9000 series offers mixed media capabilities enabling both SDLT and Ultrium technologies to exist within the same library frame.

NOTE: The ESL9595 accepts only library ready drives specifically modified for use in the ESL9000 series libraries

13.403
A.

Standard Features

Scalability	The ESL9595 offers maximum investment protection with its scaling capabilities, from 1 to 16 drives, and 400 to 595 cartridges. Extending the investment, the ESL platform can scale up to 64 drives and 2278 cartridges via a pass-through mechanism. ISV support for multi-unit scaling Includes: HP OpenView Storage Data Protector VERITAS Backup Exec VERITAS Net Backup CA Windows only with Virtual Library Option For a complete compatibility listing please visit http://www.hp.com/go/ebs	
Proven Reliability	The ESL9595 provides the necessary high reliability with proven library robotics rating of **3,000,000** load/unload cycles.	
Drive Technology	**HP LTO Ultrium Technology**	HP's Ultrium tape technology uses a "best of breed" technology approach, taking the best features from other tape technologies and combining them into a single "superdrive" without the need for design compromises. Ultrium 2 builds on the strengths and success of Ultrium 1 to deliver even better resilience than before. The HP StorageWorks Ultrium 460 tape drive has a mean time between failures (MTBF) rating of 250,000 hours at a 100 percent duty cycle, making it a truly enterprise class tape drive. Additionally, Ultrium 460 tape drives are fully backwards compatible with all Ultrium Generation 1 products from all licensed manufacturers, and can both read and write Ultrium 1 tapes.
	HP Ultrium Storage Media	HP's test program for HP branded Ultrium Generation 2 media is probably the most thorough and comprehensive in the industry. HP cartridges must satisfy an exhaustive battery of additional tests that relate directly to real life situations, where real data and real businesses are at stake. For example, each HP Ultrium 2 cartridge is specified for 20,000 load/unloads and must complete a torturous 'shoeshine' test. This measures the ability of the media to withstand repeated passes over the tape head by simulating excessive repositioning or error recovery on a short length of tape. It is particularly relevant to the library environment where media may be frequently loaded or unloaded by the operating system. Both of these procedures are unique to HP.
	SDLT Technology	Digital Linear Tape uses a proven linear recording technology, a highly accurate tape guide system, and a read after write correction system to provide high capacity, high reliability data backup and restoration. SDLT, in addition to enhanced capacity and backup performance, provides several features such as Laser Guided Magnetic Recording (LGMR™) technology, Pivoting Optical Servo (POS), Magneto-Resistive Cluster Heads (MRC), Partial Response Maximum Likelihood (PRML) channel technology and Backward Read Compatible with earlier generation of DLT drives, to be able to read media written in DLT IV format.
Manageability	The ESL9595 library configurations are fully supported through HP Insight Manager, providing information to administrators for both robotic, tape status changes, and fault status of redundant components (high availability fault status requires use of Insight Manager v7.0) HP StorageWorks Library and Tape Tools software provides diagnostics and easy firmware upgrade capabilities. Other data and automation manageability features, including multi-host direct connectivity, are provided via specific third-party applications software packages	

18.402
A.

Standard Features

Server Support	**HP ProLiant Servers**	ProLiant DL760	DL380/G2
		ProLiant ML750	ML370/G2
		ProLiant DL580/G2	ML350
		ProLiant ML 570	ML330G2
		ML530/G2	
	HP-UX Servers	Superdome	rp2450/2470 series (A class)
		rp8400 series	K class series
		rp7400 series (N class)	D Class series
		rp5400 series (L Class)	
	HP NetServer	HP lxr8500	HP lh3000
		HP lh6000	HP lc2000r
		HP lh6000r	HP lp2000r
		HP lt6000r	HP lp1000r
		HP tc4100	HP e200
		HP tc3100	HP e800
SCSI Controllers Supported	**HP ProLiant Adapters**	66MHz Ultra3 LVD, dual channel controller	129803-B21
		Ultra3 LVD, single channel controller	154457-B21
	HP-UX Servers	U2 LVD single port	A5149A
		U2 LVD dual port	A5150A
		U3 LVD single port	A6828A
		U3 LVD dual port	A6829A
		LVD single port	A4999A
	HP NetServer	Adaptec AHA 29160	
		Adaptec AHA 39160	
Fibre Channel Controllers	**HP ProLiant Servers**	FCA-2101 HBA (lp952)	
		KGPSA-CB HBA (NT lp8000)	
	HP-UX Servers	FC	A6684A
		FC	A6685A
		1G FC	A5158A
		2G FC	A6795A
	HP NetServer	Emulex LP8000-F1	
		Emulex LP9000	
		Qlogic QLA-2200F	
Fibre Channel Interface Kits	Network Storage Router E2400		
	E2400-160 FC Interface Controller		
	Network Storage Router M2402		
	For more information please visit the Interface Kits QuickSpecs.		

18.401
CA.

Standard Features

Application Support

The ESL9595 supports a broad range of SAN and direct attach environments including:
Applications supported include:
HP OpenView Storage Data Protector
HP Omniback
CA Brightstor
Legato Networker
Tivoli Storage Manager
VERITAS NetBackup
VERITAS BackupExec
Operating Systems Supported Include:
HP-UX
AIX
Linux
NetWare
OVMS
Tru64
Solaris
Windows
For a full support matrix of HP StorageWorks Enterprise Backup Solution, consult the EBS Compatibility Matrix at: http://www.hp.com/go.ebs.

NOTE: * HP-UX commands fbackup\frecover and backups launched using HP-UX Site Area Manager is not supported.

18.400
(A.

Hewlett-Packard Installation, Care Pack and CarePaq Services

Hewlett-Packard Service & Warranty Support

Hewlett-Packard Global Services provides a one year, limited warranty, fully supported by a worldwide network of resellers and service providers toll-free 7x 24 hardware technical phone support for the duration of the warranty. In addition, available service offerings include a full range of Care Pack/CarePaq packaged hardware and software services:

- Hardware Installation
- Extended coverage hours and enhanced response times
- System management and performance services

Hewlett-Packard Service and Warranty Support

Additional Warranty protection and/or Hewlett-Packard Installation packages can be purchased for the ESL9595 library products:

Software Product Services

- Stand-alone telephone support
- Rights to a new license
- Media and documentation updates

Hardware Product Services

- Installation services
- On-site maintenance (includes warranty support)
- Response time upgrades during the warranty period
- Post-warranty coverage

For additional hardware installation and maintenance information please refer to the URLs listed below:
http://www.compaq.com/services/CarePaq/us/hardware/cp_storage.html
http://www.compaq.com/services/CarePaq/us/install/cp_storage.html

Warranty Upgrade Options

- Response – Upgrade on-site response from next business day to same day 4-hours
- Coverage – Extend hours of coverage from 5 days x 9 hours to 7 days x 24 hours
- Duration – Select duration of coverage for a period of 1 or 3 years

Additional Warranty protection and/or Hewlett-Packard Installation packages can be purchased for the ESL9595 library products:

Care Packs

Description	Option	Part Number
Installation Services		
Tape Library Implementation into SAN (HW install included)	5DQ	HA114A1
Tape Library Installation (SCSI direct attached)	5DP	HA113A1

Description	Band/Option for ESL9322 0 drive library	Band/Option for ESL9595 0 drive library	Band/Option for drive upgrade kits	Band/Option for FC and Pass Thru upgrade kits	Part Number
Warranty Upgrade Services					
HW-1yr Next Day onsite	8G3	8G4	8G5	699	HA101A1
HW-1yr 8x5 onsite	8G3	8G4	8G5	699	HA103A1
HW-1yr 24x7 onsite	8G3	8G4	8G5	699	HA104A1
HW-3yr Next Day onsite	8G3	8G4	8G5	699	HA101A3
HW-3yr 8x5 onsite	8G3	8G4	8G5	699	HA103A3
HW-3yr 24x7 onsite	8G3	8G4	8G5	699	HA104A3

13 399
A.

Related Options

The ESL9595 solution includes the following related options

	Description	UPC Code	Part Number
ESL9595 Library System	ESL9000 400 to 500 slot upgrade	808736892933	330842-B21
	ESL9000 500 to 595 slot upgrade	808736892957	330842-B22
	ESL9000 400 to 595 slot upgrade	808736892971	330842-B23
	ESL9000 Ultrium 460 drive upgrade Kit	808736893039	330834-B21
	ESL9000 Ultrium 230 drive upgrade Kit	613326562840	301930-B21
	ESL 9000 SDLT 320 drive upgrade in hot plug canister	613326517642	293414-B21
	ESL 9000 SDLT 220 drive upgrade in hot plug canister	720591776985	234617-B21
	ESL9000 SDLT magazine (for use with SDLT/DLT cartridges only Will not hold LTO cartridges)	720591696979	229909-B21
	ESL9000 UNIVERSAL LOAD PORT magazine (two 4-slot fixed magazines)* *maximum of one pair per library	613326572603	302254-B21
	Fibre Channel Interface Controllers		
	E2400-160 Single card with 2 FC X 4 Ultra 3 SCSI ports. Installed into the ESL9000 card cage which is ordered separately. One drive per bus configuration, or one controller for every four tape drives for a maximum of four FC interface controllers for 16 tape drives. Installation services will be quoted with this product.	808736892773	330839-B21
	ESL9000 Card Cage Kit cPCI Card cage for ESL9000 libraries for installation of embedded Fibre Channel interface cards. Only one card cage is necessary per library is required. Card cage adapters will ship with this sku to enable both E2400-160 and E2400's to exist in the same cage.	808736892759	330838-B21
	ESL9000 Card Cage Slot Adapter Kit Required ONLY for customers who already have a card cage installed with E2400 Network Storage Routers who wish to add E2400-160's to their configuration. CE install is required and only one adapter kit is necessary per card cage.	808736968515	330838-B22

18.398
A.

	Description	UPC Code	Part Number
ESL9595 Library System *(continued)*	***Network Storage Router (required for SAN configuration)***		
	Network Storage Router M2402 (2FC X 4 LVD SCSI) *for use with Ultrium 460 Ultra3 SCSI tape drives*	613326254516	262653-B21
	E2400: Embedded Data Router with card cage, ESL9000 embedded fibre – two 1 FC X 2 LVD, Ultra2, SCSI ports to be used with ***Ultra2 Drive Kits*** : SDLT 220, SDLT 320 and Ultrium 230 Drives (Includes card cage for installation of embedded router which is required for configuration) (Includes two router cards each with 1 FC X 2 LVD SCSI ports, 4 VHDCI-WIDE SCSI cable, user guide CD, and the ESL9000 cPCI card cage) **NOTE**: It is recommended that the installation CarePaq be ordered with this SKU.	613326254752	262665-B21
	E2400 ESL9000 embedded data router– two (2 FC X 4 LVD Ultra2 SCSI ports) for use with Ultra2 Drive Kits. SDLT 220, SDLT 320 and Ultrium 230 Drives (Includes: two 1FC x 2LVD SCSI router cards, four VHDCI-WIDE SCSI cables, User Guide CD.) This SKU does not include a card cage; it is an upgrade SKU and requires 262665-B21 to be configured first, when configured it increases the number of SCSI ports available to connect drives in an ESL9000 library to the SAN, to a total of eight). **NOTE**: It is recommended that the installation CarePaq be ordered with this SKU	613326254998	262664-B21
	Supplied with each Library Unit Library unit		
	Documentation CD including user guide		
	Safety warnings with translations		
	Warranty card with translations		
	Safety booklet with translations		
	Two keys		
	SCSI terminators (1 per drive, installed)		
	Diagnostic cable (for PC attach)		
	Bar code label samples		
	Supplied with each library 2 Line Cords (IEC 320 C19), US/Canada		
Power Cord (Pair Required)	ESL9000 POWER CORDS 2X, US/CANADA/MEX* *(shipped with all units)	743172935649	146210-001
	ESL9000 POWER CORDS 2X, SE ASIA/INDIA	743172935625	146210-AR1
	ESL9000 POWER CORDS 2X, CENTRAL EUROPE	743172935632	146210-021
	ESL9000 POWER CORDS 2X, AUS/CHINA/NZL	743172935670	146210-011
	ESL9000 POWER CORDS 2X, SWITZERLAND	743172935656	146210-111
	ESL9000 POWER CORDS 2X, UK/HONG KONG	743172935663	146210-031
	ESL9000 POWER CORDS 2X, ITALY	743172935687	146210-061
	ESL9000 POWER CORDS 2X, DENMARK	743172935694	146210-081
	ESL9000 POWER CORDS 2X, JAPAN	743172935700	146210-291
	NOTE: Order Local power cords separately		

18.397
A.

Related Options

	Description	UPC Code	Part Number
Media and Bar Code Labels	HP Ultrium 400 GB Data Cartridge - Generation 2 LTO	808736395847	C7972A
	HP Ultrium 200 GB Data Cartridge - Generation 1 LTO	025184163416	C7971A
	HP LTO Ultrium universal cleaning cartridge	808736038799	C7978A
	HP Ultrium Generation 2 (Ultrium 460) bar code label pack (100 data labels, 10 cleaning labels)	808736617093	Q2002A
	HP Ultrium Generation 1 (Ultrium 230) bar code label pack (100 data labels, 10 cleaning labels)	808736617086	Q2001A
	HP SDLT 220/320 GB data cartridge	725184813436	C7980A
	HP SDLT cleaning cartridge	725184813443	C7982A
	SDLT bar code label pack (100 data labels, 10 cleaning labels)	808736617109	Q2003A
	DLT IV bar code labels (100 data labels, 10 cleaning labels)	808736617116	Q2004A
SCSI Cables	Wide/Wide SCSI cable, 5 m	743172633910	328215-001
	Wide/Wide SCSI cable, 10 m (do not order for multiple device per HBA configurations)	743172633927	328215-002
	Wide/Wide SCSI cable, 15 m (do not order for multiple device per HBA configurations)	743172633934	328215-003
	Wide/Wide SCSI cable, 20 m (do not order for multiple device per HBA configurations)	743172633941	328215-004
	SCSI Cable, VHDCI/Wide, 6 ft	743172660831	341176-B21
	SCSI Cable, VHDCI/Wide, 12 ft	743172660848	341177-B21
	SCSI Cable, VHDCI/Wide, 10 m (do not order for multiple device per HBA configurations)		
	SCSI Cable, VHDCI/Wide, 15 m (do not order for multiple device per HBA configurations)		
	SCSI Cable, VHDCI/Wide, 20 m (do not order for multiple device per HBA configurations)		
	Check the cable-adapter matrix for tape libraries on: ftp://ftp.compaq.com/pub/products/storageworks/tape-and-optical-storage/adapter-cable-matrix.pdf		
SCSI Host Bus Ad	***HP ProLiant Controllers***		
	64-bit/ 66 MHz PCI Dual Channel Wide Ultra3 SCSI adapter	743172876379	129803-B21
	Ultra3 LVD, single channel controller	743172991447	154457-B21
	HP-UX server adapters		
	U2 LVD single port		A5149A
	U2 LVD dual port		A5150A
	U3 LVD single port		A6828A
	U3 LVD dual port		A6829A
	HP NetServer adapters		
	Adaptec AHA 29160		C7474A
	Adaptec AHA 39160		P3413A
Fibre Host Bus Adapters	***HP ProLiant Adapters***		
	fca2101 HBA (lp952)	743172876379	245299-B21
	KGPSA-CB HBA (NT lp8000)	743172991447	176479-B21
	HP-UX Server Adapters		
	FC		A6685A
	1G FC		A5158A
	2G FC		A6795A
	FC		A6684A

18.396
A.

Technical Specifications

SCSI Identification

Product ID	ESL9000 Series
Vendor ID	HP

Control Panel

Library and drive status
Move tape cartridges
Load/unload tape drives
Maintenance (drive cleaning, diagnostics)
Library configuration
Statistics

Robotic SCSI Interface

Ultra3 SCSI LVD, 68HD Interconnect

External cable length between library and HBA limited to 9 m when using more than one device per HBA. Can extend up to 22 m when using exactly one device per HBA (point to point SCSI termination)

Form Factor

Free Standing

Dimensions

	Unit	**Packaged**
Height	75 in/190.5 cm	80 in/203.2 cm
Width	60 in/152.4 cm	71 in/180.3 cm
Depth	29 in/73.7 cm	48 in/121.9 cm

Unit Weight

Without data cartridges, 0 drives 1444 lb/655 kg
Tape Drive in Hot Plug canister 13.5 lb/6.12 kg per drive

Performance (native)

	Ultrium 460	**Ultrium 230**	**SDLT 320**	**SDLT 220**
Capacity	200 GB native	100 GB native	160 GB native	110 GB native
Transfer Rates	30 MB/s	15 MB/s	16 MB/s	11 MB/s

Storage Capacity

	Ultrium 460	**Ultrium 230**	**SDLT 320**	**SDLT 220**
400 slots	80 TB native, 160 TB compressed*	40 TB 80 TB compressed*	64 TB 128 TB compressed*	44 TB 88 TB compressed*
500 slots	100 TB 200 TB compressed*	50 TB 100 TB compressed*	80 TB 160 TB compressed*	55 TB 110 TB compressed*
595 slots	119 TB native 238 TB compressed*	59.5 TB 119TB compressed*	95.2 TB 190.4 TB compressed*	65.4 TB 130.9TB compressed*

* assuming 2:1 compression

Input/Output Port

two 4-slot fixed magazines

Initialization (cartridge inventory) **time**

7 minutes

Cartridge Swap Time

8 seconds average from slot to drive

Cartridge Pass-through time – end-to-end

78 seconds average from first frame to fourth frame

Electrical Specifications

Auto-range to domestic and international AC voltages

AC Voltage Range	90 to 132 VAC or 180 to 264 VAC, single phase
Nominal	100 to 120 VAC or 220 to 240 VAC
Line Frequency Range	47 to 63Hz
AC Receptacle	IEC-320 C19 Type 100-120V/200-240V, 50/60 Hz, 20A/16A (5-20P plug for U.S. and Canada)
Power Consumption	VA Maximum 1,600W 5500 BTU/hr or 1400 Kcal/hr

18.395
A.

Technical Specifications

Acoustic Noise	Noise Power Emission Level	
Regulatory Compliance **Regulatory ID Number ED 1002**	Safety	EN60950/IEC950, UL, CSA/cUL, CE, CB Scheme, Nemko CC, DHHS/FDA (Class 1 Laser), IEC60825-1 (Laser Class 1).
	EMI	EN 60825:1994/A11, FCC Part 15 Class A, CISPR Publication 22 Class A, ICES-003 Class A, VCCI Class 1, SMA AZ/NZS 3548 Class A
	EMC	EN6100-3-2, EN50082-1: 1992, EN50022: 1994
Warranty	One year parts exchange, one year on-site next day service Additional warranty protection can be purchased for all ESL9322 library products. **NOTE**: Certain restrictions and exclusions apply. Consult the Hewlett-Packard Customer Support Center for details.	

Windows NT is a registered trademark or trademark of Microsoft Corporation in the U.S. and/or other countries. UNIX is a registered trademark of The Open Group in the U.S. and other countries. All other product names mentioned herein may be trademarks of their respective companies.

The only warranties for HP products and services are set forth in the express warranty statements accompanying such products and services . Nothing herein should be construed as constituting an additional warranty. HP shall not be liable for technical or editorial errors or omissions contained herein.

18.394
A.

ANEXO UNIDADE DE BACKUP ROBOTIZADO
PARTE N

18.393
A.

invent

hp OpenView storage data protector 5.1
Platform & Integration support matrices

Version: 1.1 Edition date: June 23, 2003

NOTE: For the following Data Protector components, only those "Data Protector component-Operating System" combinations are supported for which the corresponding Operating System version is supported by the respective vendors.

supported operating systems

Data Protector component	supported operating systems
Cell Manager	• Windows NT 4.0[6] • Windows XP PRO • Windows 2000 • Windows 2003 (32-bit) • HP-UX 11.0[3], 11.11[3,5], 11.20[1,3,5] • Solaris 7, 8 & 9
Installation Server	• Windows NT 4.0[6] • Windows XP PRO • Windows 2000 • Windows 2003 (32-bit) • HP-UX 11.0[3], 11.11[3,5], 11.20[1,3,5] • Solaris 7, 8 & 9
Graphical User Interface[7]	• Windows XP HE[2] • Windows XP PRO • Windows XP 64-bit • Windows NT 4.0 • Windows 2000 • Windows 2003 (32-bit) • HP-UX 11.0[3], 11.11[3,5], 11.20[1,3,5] • Solaris 7, 8
Managers-of-Managers (MoM)	• Windows NT 4.0 • Windows 2000 • Windows XP PRO • Windows 2003 (32-bit) • HP-UX 11.0[3], 11.11[3,5], 11.20[1,3,5] • Solaris 7, 8 & 9
Backup Device Server (media agent), including robotic control (For Linux, only 32-bit is supported)	• Windows NT 4.0, Windows 2000 • Windows XP PRO, XP (64-bit) • Windows 2003 (32-bit) • Windows 2003 (64-bit) • Novell NetWare 4.2[2], 5.1[2], 6.0[2] • HP-UX (PA-RISC) – 11.0[3], 11.11[3,5], • HP-UX (IA64) – 11.20[1,3,5], 11.22[3,5]

18.392
A.

	• Tru64 5.xx • OpenVMS 7.3-1 • Sun Solaris[4] 2.6, 7, 8, 9 • MPE/iX 6.0[2], 6.5[2], 7.0[2], 7.5[2] • Linux (32bit) : Redhat 7.x, 8.x Redhat Enterprise Linux – Advanced Server 2.1[8] Linux SuSe 7.x, 8.x Linux Debian 3.0
Backup Device Server (media agent), without robotic control	The above list plus the following: • IBM AIX 4.3.x, 5Lver.5.1, 5.2 • SNI Sinix 5.4.3, 5.4.4 • SCO OpenServer 5.0.5
Backup Agents (disk agents)	• Windows NT 4.0 • Windows 98[2] • Windows ME[2] • Windows 2000 • Windows XP HE[4] • Windows XP PRO • Windows XP 64-bit • Windows 2003 (32 bit) • Windows 2003 (64 bit) • Novell NetWare 4.2[2], 5.1[2], 6.0[2] • HPUX (PA-RISC) - 11.0[3], 11.11[3,5] • HPUX (IA64) -11.20[1,3,5], 11.22[3,5] • HP MPE/iX 6.0[2], 6.5[2], 7.0[2], 7.5[2] • Sun Solaris[4] 2.6, 7, 8, 9 • IBM AIX 4.3.x, 5Lver.5.1, 5.2 • SGI IRIX 6.4, 6.5 • SNI Sinix 5.4.3, 5.4.4 • SCO OpenServer 5.0.5, SCO Unixware 7.x • Tru64 Unix 4.0x, 5.xx • Open VMS 7.3-1 • Linux (32-bit) – Red Hat 6.x, 7.x, 8.x • Red Hat Enterprise Linux – Advanced Server 2.1[4] SuSE 6.x, 7.x, 8.x Caldera OpenLinux 2.4x, 3.1.1 Debian Linux 2.2r3, 3.0 • Linux (64-bit) – RedHat 7.x SuSE 7.x Debian Linux 3.0 • Additional UNIX platforms via NFS • Additional platforms via shared disk

[1] Supported in Compatibility Mode

[2] Push installation not possible, need to be installed locally

[3] NIS+ is supported in a DNS environment

[4] For Solaris – SunOS version compatibility, please refer to the SunOS support matrix

[5] HP-UX 11.11 is HP-UX 11i version 1.0.
HP-UX 11.20 is HP-UX 11i version 1.5.
HP-UX 11.22 is HP-UX 11i version 1.6

[6] Microsoft Windows NT 4.0 Terminal Server is not supported as a Cell Manager

18.391
A.

[7] In case of Reflection X, version 9 & above are supported

[8] Includes support for Red Hat Enterprise Linux Advanced Workstation and Enterprise Server

platform restrictions for supported operating systems	
operating system	**supported processor platform**
HP-UX	PA-RISC and IA-64
Windows[1]	IA-32 and AMD (for 32-bit Windows) & IA-64 (for 64-bit Windows)
Linux	IA-32 and AMD (for 32-bit Linux) & IA-64 (for 64-bit Linux)
Sun Solaris	SPARC

[1] The Windows NT 4.0 Backup Agent (disk agent) is also supported on Alpha Systems.

NOTE: For the following integration matrix, only those combinations of Applications & Operating Systems are supported that are supported by the respective vendors.

supported online database & application integrations	
databases	**supported versions**
Oracle (32-bit)[7] (including Oracle Parallel Server & RAC)[4]	• Oracle 8.0.x/Recovery Manager : HP-UX 11.0 (32 & 64-bit) Windows NT 4.0 Solaris 2.6 (32-bit), 7 (32 & 64-bit) AIX 4.3.x • Oracle 8.1.x/Recovery Manager : HP-UX 11.0 (32 & 64-bit), 11.11(32 & 64-bit) Solaris 7, 8 & 9 (32 & 64-bit) AIX 4.3.x, 5.1, 5.2 Windows NT 4.0, Windows2000 Linux - Red Hat Distribution 6.x , 7.x SuSe 8.x • Oracle 9i/Recovery Manager : Solaris 7, 8 & 9 (32 & 64-bit) Linux (32-bit) – Red Hat 7.x , 8.x Red Hat Advanced Server 2.1[1] SuSE 7.x, 8.x Windows NT 4.0, Windows 2000 Windows 2003 (32-bit)
Oracle (64-bit)[7] (including Oracle Parallel Server & RAC)[4]	• Oracle 8.0.x/Recovery Manager : HP-UX 11.0 (64-bit) • Oracle 8.1.x/Recovery Manager : HP-UX 11.0 (64-bit), 11.11(64-bit) Solaris 7 & 8 & 9 (64-bit) AIX 4.3.3 (64-bit), 5.1 (64-bit), 5.2 (64-bit) Tru64 5.xx

18390
A

	• Oracle 9i/Recovery Manager : HP-UX 11.0 (64-bit), 11.11 (64-bit), 11.22 (64-bit) Solaris 7, 8 & 9 (64-bit) Windows 2003 (64-bit) AIX 4.3.3 (64-bit), AIX 5.1 (64-bit), 5.2 (64-bit) Tru64 5.xx
Informix (32-bit)	• Informix 7.24 : HP-UX 11.0 (32-bit) Sinix 5.4.3 Solaris 7 (32-bit) • Informix 7.3x : 11.0 (32 & 64-bit), 11.11 (32 & 64 bit) Sun Solaris 2.6 (32-bit), 7 & 8 (32 & 64-bit) AIX 4.3.x Windows NT 4.0 Windows 2000 • Informix IDS 9.2x : HP-UX 11.0 (32 & 64-bit), 11.11 (32 & 64-bit) Solaris 7, 8 (32 & 64-bit) Windows 2000 • Informix IDS 9.3x : HP-UX 11.0 (32 & 64-bit), 11.11 (32 & 64-bit) Solaris 2.6 (32-bit), Solaris 7, 8 (32 & 64-bit) AIX 4.3.x (32 & 64-bit) Windows 2000 • Informix IDS 9.40 : HP-UX 11.0 (32 & 64-bit), 11.11 (32 & 64-bit) Solaris 7, 8 & 9 (32 & 64-bit) Windows2000 • Informix XPS 8.2x : HP-UX 11.0 (32 & 64-bit) • Informix XPS 8.3x : HP-UX 11.0 (32 & 64-bit), 11.11 (32 & 64-bit)
Informix (64-bit)	• Informix 7.31 : HP-UX 11.0 (64-bit) Solaris 7 & 8 (64-bit) AIX 4.3.x (64-bit) • Informix IDS 9.2(0,1)[6] : HP-UX 11.0 (64-bit), 11.11 (64-bit) Solaris 7 & 8 (64-bit) AIX 4.3.x (64-bit) Tru64 5.xx (64-bit) • Informix IDS 9.3x : HP-UX 11.0 (64-bit), 11.11 (64-bit) Solaris 7 & 8 (64-bit) AIX 4.3.x(64-bit) Tru64 5.xx (64-bit) • Informix IDS 9.40 : HP-UX 11.0 (64-bit), 11.11 (64-bit) Solaris 7, 8 & 9 (64-bit) • Informix XPS 8.2x : HP-UX 11.0 (64-bit) • Informix XPS 8.3x : HP-UX 11.0 (64-bit), 11.11 (64-bit)
Sybase (32-bit)	• Sybase 12.0 : HP-UX 11.0, 11.11 (32 & 64-bit) Solaris 7 & 8 (32 & 64-bit) Windows 2000 Windows NT 4.0 • Sybase Adaptive Server 12.5 : HP-UX 11.0 (32 & 64-bit), 11.11 (32 & 64-bit)

18.339
A.

	Solaris 7 & 8 (32 & 64-bit) Windows NT 4.0, Windows 2000
Sybase (64-bit)	• Sybase Adaptive Server 11.9.3 : HP-UX 11.0 (64-bit), 11.11 (64-bit) Solaris 7, 8 (64-bit) • Sybase 12.0 : HP-UX 11.0 (64-bit), 11.11 (64-bit) Solaris 7 & 8 (64-bit) • Sybase Adaptive Server 12.5 : HP-UX 11.0 (64-bit), 11.11 (64-bit) Solaris 7 & 8 (64-bit)
Microsoft SQL Server (32-bit)	• Microsoft SQL Server 7.0 : Windows NT 4.0 • Microsoft SQL Server 7.0 : Windows 2000 • Microsoft SQL Server 2000 : Windows 2000 • Microsoft SQL Server 8.0 : Windows 2003(32-bit)
Microsoft SQL Server (64-bit)	• Microsoft SQL Server 8.0 : Windows 2003 (64-bit)
Microsoft Exchange Server	• Microsoft Exchange Server 5.0 : Windows NT 4.0 • Microsoft Exchange Server 5.5 : Windows NT 4.0 • Microsoft Exchange Server 5.5 : Windows 2000 • Microsoft Exchange Server 5.5 – Single mailbox restore : Windows NT 4.0 Windows 2000 • Microsoft Exchange Server 2000 – Single mailbox restore : Windows 2000
IBM DB2 (32-bit)	• IBM DB2 UDB 7.2x : HP-UX 11.0 (32 & 64-bit), 11.11 (32 & 64-bit) AIX 4.3.3 (32 & 64-bit), AIX 5.1, 5.2 (32 & 64-bit) • IBM DB2 UDB 8.1 : HP-UX 11.11 (32 & 64-bit) AIX 4.3.3 (32 & 64-bit), AIX 5.1, 5.2 (32 & 64-bit)
IBM DB2 (64-bit)	• IBM DB2 UDB 7.2x : HP-UX 11.0 (64-bit), 11.11 (64-bit) AIX 4.3.3 (64-bit), AIX 5.1 (64-bit), 5.2 (64-bit) • IBM DB2 UDB 8.1 : HP-UX 11.11 (64-bit) AIX 4.3.3 (64-bit), AIX 5.1 (64-bit), 5.2 (64-bit)
Lotus Notes, Lotus Domino	• Lotus Domino/Notes R5 : HP-UX 11.0 (32 & 64-bit), HPUX 11.11 (32 & 64-bit) AIX 4.3.x Windows NT 4.0, Windows 2000
SAPDB[5]	• SAPDB version 7.3x : HP-UX 11.0, 11.11 Windows NT4.0, Windows2000 Linux – Red Hat 7.x – SuSe 7.x • SAPDB version 7.4x : HP-UX 11.0, 11.11 Windows NT4.0, Windows2000 Linux – Red Hat 7.x, 8.x – SuSe 7.x, 8.x
SAP [2] (32-bit)	• SAP R/3 3.1x : 11.0 (32-bit) AIX 4.3.x (32-bit) Windows NT 4.0 • SAP R/3 4.0 : 11.0 (32 & 64-bit) AIX 4.3.x Windows NT 4.0 • SAP R/3 4.5, using RMAN mode : 11.0 (32 & 64-bit)

18.388
A.

	Windows NT 4.0, Windows 2000 • SAP R/3 4.5, using backint mode : 11.0 (32 & 64-bit) AIX 4.3.x Solaris 2.6 (32-bit), 7, 8 (32 & 64-bit) Windows NT 4.0 • SAP R/3 4.6x, using RMAN mode : 11.0 (32 & 64-bit) HP-UX 11.11 (32 & 64-bit) Windows NT 4.0, Windows 2000 Solaris 7 & 8 (32 & 64-bit) • SAP R/3 4.6x, using backint mode: 11.0 (32 & 64-bit) HP-UX 11.11 (32 & 64-bit) Linux Suse 7.x, 8.x AIX 4.3.x, 5.1, 5.2 Solaris 7, 8 (32 & 64-bit) Windows NT 4.0 Windows 2000 • SAP 4.7, using backint & RMAN mode: Windows 2000 Windows 2003 (32-bit) Linux SuSe 8.x • SAP Brtools 6.10, using backint & RMAN mode: Windows NT4.0 Windows 2000 Windows 2003(32-bit) Linux Suse 7.x • SAP Brtools 6.20, using backint & RMAN mode: Windows 2000 Windows 2003 (32-bit) Linux SuSe 8.x
SAP[2] (64-bit)	• SAP R/3 4.0 : HP-UX 11.0 (64-bit) • SAP R/3 4.5, using backint and RMAN mode : HP-UX 11.0 (64-bit) • SAP R/3 4.6x, using backint mode : HP-UX 11.0 (64-bit), 11.11(64-bit) • SAP R/3 4.6x, using RMAN mode : HP-UX 11.0 (64-bit), HP-UX 11.11(64-bit) Solaris 7 & 8 (64-bit) Tru64 5.xx (64-bit) AIX 4.3.x (64-bit) • SAP 4.7, using backint & RMAN mode : HP-UX 11.0 (64-bit) HP-UX 11.11(64-bit) Solaris 9 (64-bit) AIX 4.3.3 (64-bit) AIX 5.1 (64-bit), 5.2 (64-bit) Tru64 5.xx (64-bit) • SAP Brtools 6.10, using backint & RMAN mode : HP-UX 11.0 (64-bit) HP-UX 11.11(64-bit) Solaris 7 & 8 (64-bit) Tru64 5.xx (64-bit) AIX 4.3.3 (64-bit), 5.x (64-bit) • SAP Brtools 6.20, using backint & RMAN mode : HP-UX 11.0 (64-bit)

13.387
A.

	HP-UX 11.11(64-bit) Solaris 9 (64-bit) AIX 4.3.3 (64-bit) AIX 5.1(64-bit), 5.2 (64-bit) Tru64 5.xx (64-bit)
Baan IV [3]	• Baan IV on Oracle, Informix, and Microsoft SQL Server (see versions above)

[1] Includes support for Red Hat Enterprise Linux Advanced Workstation and Enterprise Server

[2] Data Protector uses the official SAP backup/restore API (BRTOOLS), which is available only in conjunction with the Oracle database. Therefore, Data Protector support is only dependent on the SAP brtools version (as supported by SAP in combination with various SAP Kernels) but is independent of the Oracle version.

[3] Baan users with Oracle, Informix, or Sybase databases can back up their application data by following the instructions for the integration of Data Protector and their database in the HP OpenView Storage Data Protector 5.1 Integration Guide.

[4] This is true for the platforms on which Data Protector has Cluster support

[5] This is a white paper based solution. The whitepaper is available in the "*Product_Information\Whitepapers & Service Deployment*" directory on the Data Protector CD.

[6] Chain restore is not working for Informix ODS 9.20.FC1 (64-bit) and earlier versions.

[7] Oracle 9i covers all released versions of Oracle 9.x.x.x

Following are the details of the Openview Applications that can be integrated with Data Protector. This integration however is an application level integration & not a database level integration unless specifically mentioned by means of a footnote.

supported openview applications

application	supported versions
HP OpenView Operations UNIX and WINDOWS	• OVO/Unix Management Server English and Japanese : OVO/Unix 6.0 including Service Navigator 6.x is supported on HP-UX 11.0, Solaris 8 OVO/Unix 7.10 including Service Navigator 7.x is supported on HP-UX 11.0, 11.11, Solaris 8 • OVO/Unix Managed Node (= Data Protector Cell Server) : OVO/Unix Agent 6.xx : HP-UX 11.0, 11.11, 11.20 Solaris 8, Solaris 9 Windows NT 4.0, Windows 2000 Win XP PRO (32bit) OVO/Unix Agent 7.10 : HP-UX 11.0, HP-UX 11.11, 11.20 Solaris 8, Solaris 9 Windows NT 4.0, Windows 2000, Win XP PRO (32bit) • OVO/Windows Management Server English and Japanese : OVO/Windows 7.1 is supported on

13.386
A.

	Win 2000 Server/ Advanced Server /Data Center • OVO/Windows Managed Node (= Data Protector Cell Server) : OVO/Windows 7.1 Agent: HP-UX 11.0, 11.11, 11.20 Solaris 8 Windows NT 4.0, Windows 2000, Win XP PRO (32bit) • OV Performance Agent on OVO/Unix or OVO/Windows Managed Node (= Data Protector Cell Server) : OVO Performance Agent C.03.70 : HP-UX 11.0, 11.11 OVO Performance Agent C.03.75: Solaris 8, Solaris 9 OVO Performance Agent C.03.65 : Windows NT 4.0, Windows 2000, Win XP PRO (32bit)
HP OpenView ManageX	• HP OpenView ManageX 3.5, 4.0, 4.2 : Windows 4.0 Windows 2000
HP OpenView OmniStorage	• HP Openview Omnistorage 3.12 : HP-UX 11.0
HP Openview Network Node Manager[1]	• HP Openview Network Node Manager 6.2 :11.0 Windows NT 4.0, Windows 2000
HP OpenView Service Desk	• HP OpenView Service Desk 4.5 : Windows NT 4.0, Windows 2000
HP OpenView Reporter	• HP OpenView Reporter 3.0 : Windows NT 4.0, Windows 2000
HP OpenView SIP	• HP OpenView SIP 3.0 : HP-UX 11.0, 11.11 Windows NT 4.0, Windows 2000

[1] Data Protector integrates with the Solid database in Network Node Manager for an online backup of the database.

supported zero downtime solutions	
array	**type supported**
HP StorageWorks XP [1]	• HP StorageWorks Business Copy XP/Continuous Access XP
EMC [2]	• EMC TimeFinder/SRDF
HP StorageWorks Virtual array[3]	• HP Stirageworks Business Copy VA
HP StorageWorks EVA[4]	• HP StorageWorks Vsnap

[1] Please refer to the "hp OpenView Storage Data protector 5.1 Zero Downtime (Split-mirror) backup & Instant Recovery support matrix for HP StorageWorks Disk Array XP" for further details on the integration.

[2] Please refer to the "hp OpenView Storage Data Protector 5.1 EMC split-mirror backup integration" for further details on the integration.

[3] Please refer to the "hp OpenView Storage Data Protector 5.1 Zero Downtime Backup & Instant

18.385
A.

Recovery support matrix for HP StorageWorks Virtual Array" for further details on the integration.

4 Please refer to the "hp OpenView Storage Data Protector 5.1 Zero Downtime Backup support matrix for HP StorageWorks Enterprise Virtual Array" for further details on the integration.

supported NAS applications/appliances

NAS application	Supported versions
NDMP[1]	• NetApp Filer[2] ONTAP version 5.3.x, 6.x, NDMP (v2, v3 & v4) : HP-UX 11.0 Solaris 7 & 8 Windows NT 4.0 Windows 2000 Windows XP • EMC Celerra (Software Version: 2.2.49 or higher, 4.x or higher) : HP-UX 11.0 Solaris 7 & 8 Windows NT 4.0 Windows 2000 Windows XP
HP NAS 8000[1]	• NAS 8000 Linux 2.4.17
HP NAS B3000/e7000/B2000	• Windows 2000

1 For supported devices, please refer to the device support matrices of the respective NAS vendors. Only those devices that are supported by the respective NAS vendors and are also listed in HP OpenView Storage Data Protector 5.1 Device Support Matrix as supported, are supported.

2 If you have NetApp ONTAP version 6.1.2, please refer to the Data Protector 5.1 Release Notes for support details.

other supported applications

application agents	supported databases / integrations
HP OpenView Storage Data Protector Open File Manager	• HP OpenView Storage Data Protector Open File Manager 8.x : Windows NT 4.0 Windows 2000 Novell NetWare 4.x, 5.x, 6.0

18.384
A.

supported clusters

cluster software	supported cluster software version
HP MC/ServiceGuard	• HP MC/ServiceGuard 11.xx : HP-UX 11.0 • HP MC/ServiceGuard 11.xx : HP-UX 11.11 • HP MC/ServiceGuard 11.xx : HP-UX 11.22
Microsoft Cluster Server	• Microsoft Cluster Server : Windows NT 4.0 Windows 2000
Veritas Cluster Server	• Veritas Cluster Server 1.3 (Disk Agent Support only) : Solaris 7 & 8 • Veritas Cluster Server 2.0 (Disk Agent Support only) : Solaris 7 & 8
Netware Cluster	• Netware Cluster Services 1.6 : Netware 6.0

supported file systems

operating system	file systems
Windows NT	• NTFS 1.1 • FAT • CIFS
Windows 2000	• NTFS 3.0(Including Image Backup) • FAT16, FAT32 • CIFS
Windows 98	• FAT • VFAT • FAT32 • CIFS
Windows XP Pro/2003	• NTFS 3.1(Including Image Backup) • FAT16, FAT32 • CIFS
Windows XP 64/2003	• NTFS 3.1 (Including Image Backup of MBR disks only) • FAT16, FAT32 • CIFS
Novell NetWare	• NetWare FS supported namespaces: MAC, NFS, OS/2 (long namespace), DOS • NSS[1]
HP-UX	• HFS (Including Rawdisk) • NFS • LOFS (Loopback FS) • VxFS (Including Rawdisk) • DCE DFS
Solaris[2]	• UFS • NFS • PC FS (MSOS compatible FS)

18 383

	• HSFS • VxFS[3] • Tmp FS • LOFS (Loopback FS)
IBM AIX[2]	• AIX physical FS "oaix" • JFS (Journaling FS)
SGI IRIX[2]	• EFS (Extent File System) • XFS (IRIX Journaling FS)
SNI Sinix[2]	• HSFS (High Sierra FS) • UFS (UNIX File System) • VxFS
SCO OpenServer[2]	• HTFS (High Throughput FS) • DTFS • S51K • S52K
SCO Unixware[2]	• memfs • s5 • sfs (Note: ACLs are not backed up and restored.) • ufs • bfs • vxfs (Note: ACLs are not backed up and restored.)
Tru64 UNIX[2]	• UFS (UNIX FS) • AdvFS (Advanced File System)
Linux (Red Hat, SuSe, Caldera & Debian)[2]	• ext3, ext2, ext, minix, xiafs, ReiserFS

[1] Novell NetWare NSS is supported with the following limitations:

- The new NSS volumes on NetWare 5.x allow creation of deep directory structures (deeper than 100). Data Protector can back up only to a depth of 100.
- NetWare 5.x NSS volumes allow 8 TB file size, but the maximum file size that Data Protector backs up is 4 GB.

NetWare 5.x has an NSS_ADMIN volume on every system with NSS volumes present. This is a special read-only volume that contains information about NSS volume configuration. NSS_ADMIN is not backed up.

[2] Raw Disk Backup is supported. On Tru64, this is not true if LSM is present

[3] VxFS attributes backup is only supported for Solaris 2.6

ACL support in data protector 5.1		
operating system	**file system**	**number of ACL entries (basic/extended)**
Windows NT	NTFS	Unlimited

18.382
A.

Windows NT	FAT	Not Applicable
Windows 2000	NTFS	Unlimited
Windows XP Pro	NTFS	Unlimited
Windows XP 64	NTFS	Unlimited
Novell Netware	NetWare FS (MAC, NFS, OS/2, DOS) [1]	512
Novell Netware	NSS	Unlimited
HP-UX	HFS	3/16
HP-UX	VxFS (10.x, 11.x)	4/17 (JFS 3.3)
IBM AIX	AIX (Physical FS 'oaix')	1024 (4096 bytes)

[1] Limited with Data Protector record size and SMS API.

With Windows 2003, Volume Shadow Copy service (VSS) is a software service introduced by Microsoft. Data Protector integration with Microsoft Volume Shadow Copy service provides support of the VSS-aware writers. This includes automatic detection of the supported VSS-aware writers, backup and restore functionality. Following is a list or writers supported by Data Protector 5.10:

VSS writers support in data Protector 5.1

writer name	writer GUID	restore method
Certificate Authority	6f5b15b5-da24-4d88-b737-63063e3a1f86	REPLACE_AT_REBOOT
Dhcp Jet Writer	be9ac81e-3619-421f-920f-4c6fea9e93ad	REPLACE_AT_REBOOT
Event Log Writer	eee8c692-67ed-4250-8d86-390603070d00	REPLACE_AT_REBOOT
IIS Metabase Writer	59b1f0cf-90ef-465f-9609-6ca8b2938366	REPLACE_AT_REBOOT
NTDS	b2014c9e-8711-4c5c-a5a9-3cf384484757	RESTORE_IF_CAN_BE_REPLACED
MSDEWriter	f8544ac1-0611-4fa5-b04b-f7ee00b03277	RESTORE_IF_CAN_BE_REPLACED
Remote Storage	b959d2c3-18bb-4607-b0ca-688cd0d41a50	REPLACE_AT_REBOOT
Removable Storage Manager	5d3c3e01-0297-445b-aa81-a48d7151e235	REPLACE_AT_REBOOT
System Writer	e8132975-6f93-4464-a53e-1050253ae220	REPLACE_AT_REBOOT
TermServLicensing	5382579c-98df-47a7-ac6c-98a6d7106e09	REPLACE_AT_REBOOT
WINS Jet Writer	f08c1483-8407-4a26-8c26-6c267a629741	REPLACE_AT_REBOOT
WMI Writer	a6ad56c2-b509-4e6c-bb19-49d8f43532f0	REPLACE_AT_REBOOT

18 331
A-

ANEXO UNIDADE DE BACKUP ROBOTIZADO
PARTE O

invent

Enterprise Backup Solution Hardware/Software Compatibility Matrix June 2003

This matrix includes tape automation compatibility details for Direct Attach Storage (DAS), Direct Attach Fibre (DAF or point-to-point fibre), and Storage Area Network (SAN) connection options.

For information about tape drive compatibility and support, go to http://www.hp.com/go/connect

start here

18.330

STORAGE HARDWARE / SOFTWARE *COMPATIBILITY MATRIX*

June 2003

DAS, DAF, and SAN Support	HP Data Protector	Atempo Time Navigator	BakBone NetVault	CA BrightStor ARCserve	CA BrightStor EB	CommVault Galaxy	Legato NetWorker	Syncsort Backup Express	Tivoli Storage Manager	VERITAS Backup Exec	VERITAS NetBackup
Enterprise Libraries											
ESL 9595 (Ultrium 460)	✓			✓[4]			✓		✓	✓[5]	✓
ESL 9595 (Ultrium 230)	✓			✓		✓	✓	✓	✓	✓	✓
ESL 9595 S2	✓			✓	✓	✓	✓	✓	✓	✓	✓
ESL 9595 SL	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
ESL 9326 SL	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
ESL 9326 DX	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
ESL 9322 (Ultrium 460)	✓			✓[4]			✓		✓	✓[5]	✓
ESL 9322 (Ultrium 230)	✓			✓		✓	✓	✓	✓	✓	✓
ESL 9322 S2	✓			✓	✓	✓	✓	✓	✓	✓	✓
ESL 9198 SL	✓	✓	✓	✓[4]	✓	✓	✓	✓	✓	✓	✓
ESL 9198 DLX	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
ESL Multi-Unit Scalability[6]	✓						✓				✓
Mid-Range Libraries											
MSL 6060 (Ultrium 460)	✓			✓			✓			✓[2]	✓
MSL 5060 (Ultrium 230)	✓			✓			✓		✓	✓	✓
MSL 5052 S2	✓			✓	✓	✓	✓	✓		✓	✓
MSL 5052 SL	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
MSL 6030 (Ultrium 460)	✓			✓			✓			✓[2]	✓
MSL 5030 (Ultrium 230)	✓			✓			✓		✓	✓	✓
MSL 5026 S2	✓			✓	✓	✓	✓	✓		✓	✓
MSL 5026 SL	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
MSL 5026 DLX	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
MSL Multi-Unit Scalability[6]	✓			✓			✓			✓	✓
SSL 2020 TL[3]	✓	✓	✓	✓		✓	✓[1]	✓	✓[1]	✓	✓[1]

NOTES:

1) Windows NT/2000 and Solaris support in homogeneous SAN environments only
2) Support is with Windows NT/2000 only
3) NSR fibre-SCSI bridge does not support SSL 2020 TL library
4) Support is with CA BrightStor ARCserve for Windows NT/2000 v9.0 and CA BrightStor ARCserve for NetWare v7.0 in a homogeneous environment only
5) SAN and DAF support is with Windows NT/2000 only and DAS support is with Windows NT/2000 and NetWare
6) ESL multi-unit scalability is with a horizontal pass-through mechanism, and the MSL multi-unit scalability is with a vertical pass-through mechanism

18.349

STORAGE HARDWARE / SOFTWARE COMPATIBILITY MATRIX

June 2003

DAS Support	HP Data Protector	Atempo Time Navigator	BakBone NetVault	CA BrightStor ARCserve	CA BrightStor EB	CommVault Galaxy	Legato NetWorker	Syncsort Backup Express	Tivoli Storage Manager	VERITAS Backup Exec	VERITAS NetBackup
AutoLoaders[1]											
SSL1016 Ultrium 460 Tape Autoloader	✓[6]			✓			✓[3]			✓[10]	✓
1/8 Ultrium 230 Tape Autoloader	✓[4]			✓[8]			✓[11]		✓[3]	✓[5]	✓[9]
1/8 DLT vs80 Tape Autoloader	✓[3]	✓[3]	✓[3]	✓	✓[3]	✓[3]	✓[11]	✓[3]		✓[5]	
SSL1016 SDLT160/320 Tape Autoloader	✓[4]			✓[3]			✓[10]			✓[3]	✓[2]
SSL1016 DLT1 Tape Autoloader	✓[3]	✓[3]	✓[3]	✓[5]	✓[3]	✓[3]	✓[10]	✓[3]		✓[3]	✓[3]

NOTES:

1) Autoloaders are only supported in Direct Attach SCSI configurations

2) Support with Windows NT/2000, HP-UX 11.0, and HP-UX 11.11 only

3) Support is with Windows NT/2000 only

4) Support for Windows NT/2000, NetWare, and HP-UX 10.2, 11.0 & 11.11 only

5) Support for Windows NT/2000 and NetWare only

6) Support for Windows NT/2000/2003, NetWare, and HP-UX 10.2, 11.0 & 11.11 only

7) HP-UX commands fbackup/frecover and backups launched using HP-UX Site Area Manager are not supported

8) Support for Windows NT/2000, Linux, and NetWare only

9) Support for Windows NT/2000, Linux, Solaris, and HP-UX 10.2, 11.0 & 11.11 only

10) Not supported with Windows 2003

11) Support for Windows NT/2000, Linux, and Solaris only

18.378
SA.

STORAGE HARDWARE / SOFTWARE COMPATIBILITY MATRIX

June 2003

DAS, DAF, and SAN Support	HP Data Protector	Atempo Time Navigator	BakBone NetVault	CA BrightStor ARCserve	CA BrightStor EB	CommVault Galaxy	Legato NetWorker	Syncsort Backup Express	Tivoli Storage Manager	VERITAS Backup Exec	VERITAS NetBackup
Operating Systems											
HP-UX 11.0	✓		✓[2]		✓[2]	✓[2]	✓	✓[2]	✓		✓
HP-UX 11.11	✓	✓[2]					✓	✓[2]	✓		✓
HP-UX 11.22	✓										✓
WIN NT 4.0	✓	✓	✓	✓[7]	✓	✓	✓	✓	✓	✓	✓
WIN 2000 Server/AS	✓	✓	✓	✓[7]	✓	✓	✓	✓	✓	✓	✓
WIN Server 2003 (32-bit)				✓[7]			✓			✓	
Tru64 5.1[3]							✓				✓
Tru64 5.1A[3]	✓	✓	✓		✓	✓	✓	✓			✓
Tru64 5.1B[3]							✓				✓
Linux Redhat Advanced Server 2.1							✓				
NetWare 5.1	✓		✓	✓[7]				✓		✓	
NetWare 6.0	✓			✓[7]				✓		✓	
Solaris 7	✓						✓		✓		✓
Solaris 8	✓	✓	✓		✓	✓	✓	✓	✓		✓
Solaris 9	✓						✓				✓
AIX 4.3.3 (32-bit)	✓[5]		✓				✓		✓		✓
AIX 5.1L (32-bit)							✓[1]		✓		✓
High Availability[8]											
HP-UX (Service Guard)[4]	✓[6]										✓[6]
Windows (MSCS)		✓	✓			✓	✓	✓	✓	✓	✓
Tru64 (TruCluster)		✓					✓	✓			✓
NetWare 5.1 (NWCS)										✓	
NetWare 6.0 (NWCS)								✓		✓	
Solaris											✓

NOTES:

1) DLT8000 products not supported
2) A-Class, L-Class, N-Class, and V-Class servers only
3) Ultrium 460 (LTO) support only
4) Supported on XP and VA raid array products only
5) Robot control is performed on Windows platform
6) ATS is not supported
7) Supports ESL9000 libraries in a homogeneous environment only
8) Refer to the EBS Design Guide for details on failover configurations

13.317
A.

STORAGE HARDWARE / SOFTWARE COMPATIBILITY MATRIX

June 2003

SAN Support	HP Data Protector	Atempo Time Navigator	BakBone NetVault	CA BrightStor ARCserve	CA BrightStor EB	CommVault Galaxy	Legato NetWorker	Syncsort Backup Express	Tivoli Storage Manager	VERITAS Backup Exec	VERITAS NetBackup
RAID Arrays											
XP 1024/128[1,5]	✓			✓			✓		HP-UX and NT/WIN2K only	✓	✓
XP 48/512[1,5]	✓			✓			✓		HP-UX and NT/WIN2K only	✓	✓
Zero Downtime with XP	✓										
Instant Recovery with XP	✓										
Enterprise Virtual Array (EVA)[1]	Solaris, Tru64, HP-UX, and NT/WIN2K	NT/WIN2K only	NT/WIN2K only	✓[3]	NT/WIN2K only	NT/WIN2K only	✓	NT/WIN2K only	✓	✓[3]	✓
VA 7410/7110[1,5]	✓			✓			✓			✓	✓
VA 7400/7100[1,5]	✓			✓			✓			✓	✓
EMA 16000[1]	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
EMA 12000/8000[1]	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
MSA1000[2, 4]	NT/WIN2K only	NT/WIN2K only	NT/WIN2K only	✓	NT/WIN2K only	NT/WIN2K only	NT/WIN2K & LINUX	NT/WIN2K only	✓	✓	NT/WIN2K only

NOTES:

1) Consult with your hp solutions provider for disk array SAN compatibility information

2) For compatibility information about the optional MSA Fabric Switch, refer to the Infrastructure page within this document for support details

3) Netware 5.1 & 6.0 supported in single path configuration with FCA 2210 HBA only

4) NetWare 5.1 & 6.0 supported with FCA2210 HBA only

5) For XP and VA array configurations, disk and tape must be on separate HBAs

18.376
A.

STORAGE HARDWARE / SOFTWARE COMPATIBILITY MATRIX

June 2003

Libraries | AutoLoaders | Operating Systems | RAID A | Infrastructure >>Next | Servers | Fibre HBAs
SCSI HBAs | Software | Disk Array Multipath | mware & Drivers | Backup App Details

SAN Support	HP Data Protector	Atempo Time Navigator	BakBone NetVault	CA BrightStor ARCserve	CA BrightStor EB	CommVault Galaxy	Legato NetWorker	Syncsort Backup Express	Tivoli Storage Manager	VERITAS Backup Exec	VERITAS NetBackup
Interconnects[1]											
Director Switch 2/140	✓			✓	✓		✓		✓	✓	✓
Director Switch 2/64	✓			✓	✓		✓		✓	✓	✓
Edge Switch 2/32	✓			✓	✓		✓		✓	✓	✓
Edge Switch 2/24	✓			✓	✓		✓		✓	✓	✓
Edge Switch 2/16	✓			✓	✓		✓		✓	✓	✓
Core Switch 2/64	✓			✓	✓		✓		✓	✓	✓
SAN Switch 2/32	✓			✓	✓		✓		✓	✓	✓
SAN Switch 2/16	✓			✓	✓	✓	✓	✓	✓	✓	✓
SAN Switch 2/8	✓			✓	✓	✓	✓	✓	✓	✓	✓
SAN Switch 2/16-EL	✓			✓	✓	✓	✓	✓	✓	✓	✓
SAN Switch 2/8-EL	✓			✓	✓	✓	✓	✓	✓	✓	✓
SAN Switch 16	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
SAN Switch 8	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
SAN Switch16EL	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
SAN Switch 8EL	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
FC Switch 16B	✓			✓	✓	✓	✓	✓	✓	✓	✓
FC Switch 8B	✓			✓	✓	✓	✓	✓	✓	✓	✓
MSA Fabric Switch 2/8	✓			✓	✓		✓		✓	✓	✓
MSA Fabric Switch 6[3]										✓	
FC-AL Switch[2]	✓	✓	✓	✓		✓		✓		✓	

NOTES:

1) Refer to Heterogeneous Open SAN Design Reference Guide for switch configuration rules and guidelines located at: http://h18000.www1.hp.com/products/storageworks/san/documentation.html or consult with your hp solutions provider for assistance

2) NSR products do not support FC-AL Switch

3) Supported with Windows and NSR – no NetWare

18.375
Paula

>>Next Page

STORAGE HARDWARE / SOFTWARE *COMPATIBILITY MATRIX*

braries | AutoLoaders | Operating Systems | RAID Array <<Previous Infrastructure | Servers | Fibre HBAs
SCSI HBAs | Software | Disk Array Multipath | ware & Drivers | Backup App Details

June 2003

SAN Support	HP Data Protector	Atempo Time Navigator	BakBone NetVault	CA BrightStor ARCserve	CA BrightStor EB	CommVault Galaxy	Legato NetWorker	Syncsort Backup Express	Tivoli Storage Manager	VERITAS Backup Exec	VERITAS NetBackup
Fibre Channel Interface Controller											
E2400-160 FC Interface Controller	✓			✓			✓[4]				✓
NSR E1200-160	✓			✓			✓			✓	✓
NSR E2400	✓			✓			✓		✓	✓	✓
NSR E1200	✓			✓			✓		✓	✓	✓
NSR M2402	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
NSR N1200	✓			✓			✓		✓	✓	✓
MDR[1,3]	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Serverless (3PC) Backup Support											
NSR E2400				Windows 2000 Serverless opt w/ HSG only	Windows 2000 Serverless opt w/ HSG only					Windows 2000 ServerFree opt with HSG, HSV, MSA1000 only	Sol 8 NBU 4.5 nbu_snap w/ HSG only[2]
NSR E1200				Windows 2000 Serverless opt w/ HSG only	Windows 2000 Serverless opt w/ HSG only					Windows 2000 ServerFree opt with HSG, HSV, MSA1000 only	Sol 8 NBU 4.5 nbu_snap w/ HSG only[2]
NSR M2402				Windows 2000 Serverless opt w/ HSG only	Windows 2000 Serverless opt w/ HSG only					Windows 2000 ServerFree opt with HSG, HSV, MSA1000 only	Sol 8 NBU 4.5 nbu_snap w/ HSG only[2]
NSR N1200				Windows 2000 Serverless opt w/ HSG only	Windows 2000 Serverless opt w/ HSG only					Windows 2000 ServerFree opt with HSG, HSV, MSA1000 only	Sol 8 NBU 4.5 nbu_snap w/ HSG only[2]

NOTES:

1) MDR does not support the following libraries:
- LTO-based libraries MSL5060L1, MSL5030L1, ESL9595L1 and ESL9322L1
- SDLT2-based libraries MSL5026S2, MSL5052S2, ESL9322S2 and ESL9595S2

2) Master router configuration not supported

3) MDR does not support Direct Attach Fibre (DAF) configurations

4) Not supported with Windows 2003

19.374
Paula

<<Previous Page

STORAGE HARDWARE / SOFTWARE COMPATIBILITY MATRIX

June 2003

Libraries | AutoLoaders | Operating Systems | Arrays | Infrastructure | **Servers** | Fibre HBAs
SCSI HBAs | Software | Disk Array Multip | Firmware & Drivers | Backup App Details

DAS, DAF, and SAN Support	HP Data Protector	Atempo Time Navigator	BakBone NetVault	CA BrightStor ARCserve	CA BrightStor EB	CommVault Galaxy	Legato NetWorker	Syncsort Backup Express	Tivoli Storage Manager	VERITAS Backup Exec	VERITAS NetBackup
Servers											
HP SuperDome (pa-risc)	✓						✓				✓
HP rp series (pa-risc)	✓						✓		✓		✓
HP A, L, N, and V-Class	✓	✓	✓		✓	✓	✓	✓	✓		✓
HP ProLiant	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
HP NetServer	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
HP AlphaServer	✓	✓	✓		✓	✓	✓	✓			✓
IBM RS/6000	✓		✓				✓		✓		✓
Sun UltraSPARC	✓	✓	✓		✓	✓	✓	✓	✓		✓
Third Party X86	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Workstations											
B-Class B2000	✓										✓
C-Class C3600, C3700	✓										✓
J-Class J6700, J6750	✓										✓
NAS											
NAS 8000	✓										
NAS Executor 7000	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
NAS B3000	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
NAS B2000	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓

18.373
Paula

STORAGE HARDWARE / SOFTWARE COMPATIBILITY MATRIX

June 2003

Libraries | AutoLoaders | Operating Systems | Arrays | Infrastructure | Servers | Fibre HBAs >>Next
SCSI HBAs | Software | Disk Array M th | Firmware & Drivers | Backup App Details

SAN and DAF Support	HP Data Protector	Atempo Time Navigator	BakBone NetVault	CA BrightStor ARCserve	CA BrightStor EB	CommVault Galaxy	Legato NetWorker	Syncsort Backup Express	Tivoli Storage Manager	VERITAS Backup Exec	VERITAS NetBackup
Fibre Host Bus Adapters[1]											
HP-UX											
A6795A (2G)	✓						✓		✓		✓
A6685A (1G)	✓						✓				✓
A5158A (1G)	✓	✓	✓		✓	✓	✓	✓	✓		✓
Tru64											
FCA2354 (2G)							✓				✓
KGPSA -CA (1G)	✓	✓	✓		✓	✓	✓	✓			✓
Solaris											
FCA2257P (2G - PCI)											
FCA2257C (1G - cPCI)											
SWSA4-SC 64-bit (1G - Sbus)	✓	✓	✓		✓	✓	✓	✓	✓		✓
SWSA4-PC 32-bit (1G - PCI)							✓		✓		✓
JNI FCE6460 (2G - PCI)[2]	✓						✓				✓
JNI FCE6410 (1G - PCI)	✓						✓				✓
AIX											
IBM 6228 (2G)	✓						✓		✓		✓
PCI-to-FC HBA for AIX (1G) 197819-B21	✓		✓				✓		✓		✓

NOTES:

1) Refer to Heterogeneous Open SAN Design Reference Guide for configuration rules and guidelines located at:
http://h18000.www1.hp.com/products/storageworks/san/documentation.html or consult with your hp solutions provider for assistance

2) The FCE6460 is not supported with tape. An additional HBA needs to be used for access to shared tape.

18.372
Paula

>>Next Page

STORAGE HARDWARE / SOFTWARE *COMPATIBILITY MATRIX*

June 2003

Libraries | AutoLoaders | Operating Systems | [illegible] Arrays | Infrastructure | Servers | << Previous Fibre HBAs
SCSI HBAs | Software | Disk Array [illegible]ath | Firmware & Drivers | Backup App Details

SAN and DAF Support	HP Data Protector	Atempo Time Navigator	BakBone NetVault	CA BrightStor ARCserve	CA BrightStor EB	CommVault Galaxy	Legato NetWorker	Syncsort Backup Express	Tivoli Storage Manager	VERITAS Backup Exec	VERITAS NetBackup
Fibre Host Bus Adapters[1]											
Windows											
FCA2355 (2G)	✓			✓			✓		✓	✓	✓
ISP2312 (2G - BL20pG2 Blade Server)	✓									✓	✓
FCA2101 (2G)	✓			✓	✓	✓	✓	✓	✓	✓	✓
64-bit/33MHz PCI to FC (1G) KGPSA-CB	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
LINUX											
FCA2214 (2G)											
QLA2200 (1G)	✓						✓	✓			
NetWare											
FCA2210 (2G)	✓[3]			✓[3]						✓[3]	
64-bit / 66 MHz PCI to FC HBA[2, 4] (1G)	✓		✓	✓				✓		✓	
32-bit / 33 MHz HBA[2, 4] (1G)	✓			✓				✓		✓	

NOTES:

1) Refer to Heterogeneous Open SAN Design Reference Guide for configuration rules and guidelines located at: http://h18000.www1.hp.com/products/storageworks/san/documentation.html or consult with your hp solutions provider for assistonce

2) COMPAQ HBAs

3) Supports Enterprise Virtual Array (single path only) and MSA1000 (with external HP StorageWorks SAN switch only).

4) Hetergeneous Windows and NetWare configurations require FCA-2210 HBA for access to shared tape. If the configuration includes NetWare 6, HSG80-based disk storage (MA8000, EMA16000) and shared tape, servers must be configured with Compaq HBA's for disk storage and FCA-2210 HBA's for tape storage (separate switch zones for disk and tape).

18.371
Paula

<<Previous Page

STORAGE HARDWARE / SOFTWARE COMPATIBILITY MATRIX

June 2003

SCSI Host Bus Adapters				
HP-UX	**Tru64**	**Solaris**	**Windows**	**NetWare**
A6829A (LVD Ultra-160 PCI)	3X-KZPCA-AA (LVD Ultra)	LSI 6603 (LVD Dual port)	154457-B21 (Ultra3 LVD, single channel)	348757-B21 (dual channel wide)
A6828A (LVD Ultra-160 PCI)		LSI 6602 (LVD Single port)	129803-B21 (Ultra3 LVD, dual channel)	268351-B21 (U320, dual channel)[1]
A5149A (LVD Ultra 2 PCI)			348757-B21 (dual channel wide)	C7474A
A5150A (LVD Ultra 2 PCI)			Adaptec 2940U2W	
A4974A (PCI Ultra SCSI Single Ended Adaptor)			Adaptec 2940U2B	
A4999A (PCI Ultra2 Low Voltage Diff SCSI adaptor)			Adaptec 29160	
A5838A (PCI 2 Port 100Base-T 2-Port Ultra2 SCSI)			Adaptec 39160	
			Qlogic 1280	
			Symbios/LSI 8951	
			Symbios/LSI 8952U	
			Symbios/LSI 22902	
			Symbios/LSI 22903	
			Symbios/LSI 22915	
			268351-B21 (U320, dual channel)[1]	
			C7474A	

NOTES:
1) Supported with driver 1.08.18 only

18.570
Paula

STORAGE HARDWARE / SOFTWARE *COMPATIBILITY MATRIX*

June 2003

	HP Data Protector	Atempo Time Navigator	BakBone NetVault	CA BrightStor ARCserve	CA BrightStor EB	CommVault Galaxy	Legato NetWorker	Syncsort Backup Express	Tivoli Storage Manager	VERITAS Backup Exec	VERITAS NetBackup
OpenView											
Storage Virtual Replicator				✓			✓		✓	✓	✓
OV-SAM	✓			✓			✓		✓	✓	✓
StorageWorks											
Data Replication Manager				✓			✓		✓	✓	NT/WIN2K only
Enterprise Volume Manager		✓		✓		✓	✓		✓	NT/WIN2K only	NT/WIN2K only
Tools and Utilities											
Smart Start	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Insight Manager 7				✓			✓		✓	✓	✓

18.369
Paula

STORAGE HARDWARE / SOFTWARE COMPATIBILITY MATRIX

June 2003

Libraries | AutoLoaders | Operating Syste RAID Arrays | Infrastructure | Servers | Fibre HBA:
SCSI HBAs | Software | **Disk Array M path** | Firmware & Drivers | Backup App Detail:

SAN Support	HP Data Protector	Atempo Time Navigator	BakBone NetVault	CA BrightStor ARCserve	CA BrightStor EB	CommVault Galaxy	Legato NetWorker	Syncsort Backup Express	Tivoli Storage Manager	VERITAS Backup Exec	VERITAS NetBackup
Disk Array Multipath[1]											
AutoPath for HP-UX[1]	✓						✓				✓
Secure Path for HP-UX[1]	EVA Only						EVA Only		✓		EVA Only
Secure Path for WIN NT[1]	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Secure Path for WIN 2000[1]	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
AutoPath for Windows 2000[1]										✓	
Tru64 Multipath[2]							✓	✓			✓
Secure Path for NetWare[1]										✓[3]	
Secure Path for Solaris[1]							✓		✓		✓
Secure Path for AIX[1]											
IBM DPO									✓		

NOTES:

1) Refer to Heterogeneous Open SAN Design Reference Guide for configuration rules and guidelines located at:
http://h18000.www1.hp.com/products/storageworks/san/documentation.html or consult with your hp solutions provider for assistance

2) Multipath to tape supported

3) MSA1000 requires 3.0c with FCA-2210, HSG80's require 3.0b with COMPAQ HBAs

18.368
Paula

STORAGE HARDWARE / SOFTWARE COMPATIBILITY MATRIX

June 2003

Components & Firmware	Current Release	HP Data Protector	Atempo Time Navigator	BakBone NetVault	CA BrightStor ARCserve	CA BrightStor EB	CommVault Galaxy	Legato NetWorker	Syncsort Backup Express	Tivoli Storage Manager	VERITAS Backup Exec	VERITAS NetBackup
Enterprise Libraries												
ESL 9595 (Ultirum 460)												
Robot FW	3.40	3.40						3.40		3.40	3.40	3.40
Drive FW Ultrium 460	F3CW	F3CW						F3CW		F3CW	F3CW	F3CW
ESL 9595 (Ultirum 230)												
Robot FW	3.40	3.40			3.31		3.31	3.40	3.31	3.40	3.40	3.40
Drive FW Ultrium 230	E33W	E33W			E33W		E30W (005.452)	E33W	E30W (005.452)	E33W	E33W	E33W
ESL 9595 S2												
Robot FW	3.40	3.40			3.31	3.20	3.20	3.40	3.20	3.40	3.40	3.40
Drive FW SDLT 160/320	v52	v52			v52	V35	V35	v52	V35	v52	v52	v52
ESL 9595 SL												
Robot FW	3.40	3.40	3.20p5	3.20p5	3.31	3.20	3.20	3.40	3.20	3.40	3.40	3.40
Drive FW SDLT 110/220	v51	v51	V35	V35	v51	V35	V35	v51	V35	v51	v51	v51
ESL 9326 SL												
Robot FW	3.40	3.40	3.20p5	3.20p5	3.31	3.20	3.20	3.40	3.20	3.40	3.40	3.40
Drive FW SDLT 110/220	v51	v51	V35	V35	v51	V35	V35	v51	V35	v51	v51	v51
ESL 9326 DX												
Robot FW	3.40	3.40	3.20p5	3.20p5	3.31	3.20	3.20	3.40	3.20	3.40	3.40	3.40
Drive FW 40/80	V80	V80	V80	V80	V80	V80	V80	V80	V80	V80	V80	V80
ESL 9322 (Ultirum 460)												
Robot FW	3.40	3.40						3.40		3.40	3.40	3.40
Drive FW Ultrium 460	F3CW	F3CW						F3CW		F3CW	F3CW	F3CW
ESL 9322 (Ultirum 230)												
Robot FW	3.40	3.40			3.31		3.31	3.40	3.31	3.40	3.40	3.40
Drive FW Ultrium 230	E33W	E33W			E33W		E30W (005.452)	E33W	E30W (005.452)	E33W	E33W	E33W
ESL 9322 S2												
Robot FW	3.40	3.40			3.31	3.20	3.20	3.40	3.20	3.40	3.40	3.40
Drive FW SDLT 160/320	v52	v52			v52	V35	V35	v52	V35	v52	v52	v52
ESL 9198 SL												
Robot FW	3.40	3.40	3.20p5	3.20p5	3.31	3.20	3.20	3.40	3.20	3.40	3.40	3.40
Drive FW SDLT 110/220	v51	v51	V35	V35	v51	V35	V35	v51	V35	v51	v51	v51
ESL 9198 DLX												
Robot FW	3.40	3.40	3.20p5	3.20p5	3.31	3.20	3.20	3.40	3.20	3.40	3.40	3.40
Drive FW 40/80	v80	v80	V80	V80	v80	V80	V80	v80	V80	v80	v80	v80

18.367
Paula

>>Next Page

STORAGE HARDWARE / SOFTWARE COMPATIBILITY MATRIX

June 2003

Libraries | AutoLoaders | Operating ems | RAID Arrays | Infrastructure | Servers | Fibre HBAs
SCSI HBAs | Software | Disk Array Multip | <<Previous Firmware & Drivers >>Next | Backup App Details

Components & Firmware	Current Release	HP Data Protector	Atempo Time Navigator	BakBone NetVault	CA BrightStor ARCserve	CA BrightStor EB	CommVault Galaxy	Legato NetWorker	Syncsort Backup Express	Tivoli Storage Manager	VERITAS Backup Exec	VERITAS NetBackup
Mid-Range Libraries												
MSL 6060 (Ultirum 460)												
Robot FW	4.14	4.14			4.14			4.14			4.14	4.14
Drive FW Ultrium 460	F3CW	F3CW			F3CW			F3CW			F3CW	F3CW
MSL 5060 L1												
Robot FW	4.14	4.14			4.14			4.14		4.14	4.14	4.14
Drive FW Ultrium 230	E33W	E33W			E33W			E33W		E33W	E33W	E33W
MSL 5052 S2												
Robot FW	4.14	4.14			4.14	3.11	3.11	4.14	3.11		4.14	4.08
Drive FW SDLT 160/320	v52	v52			v52	v35	V35	v52	V35		v52	v52
MSL 5052 SL												
Robot FW	4.14	4.14	3.11	3.11	4.14	3.11	3.11	4.14	3.11	4.14	4.14	4.14
Drive FW SDLT 110/220	v51	v51	V35	V35	v51	v35	V35	v51	v35	v51	v51	v51
MSL 6030 (Ultirum 460)												
Robot FW	4.14	4.14			4.14			4.14			4.14	4.14
Drive FW Ultrium 460	F3CW	F3CW			F3CW			F3CW			F3CW	F3CW
MSL 5030 L1												
Robot FW	4.14	4.14			4.14			4.14		4.14	4.14	4.14
Drive FW Ultrium 230	E33W	E33W			E33W			E33W		E33W	E33W	E33W
MSL 5026 S2												
Robot FW	4.14	4.14			4.14	3.11	3.11	4.14	3.11		4.14	4.14
Drive FW SDLT 160/320	v52	v52			v52	v35	V35	v52	V35		v52	v52
MSL 5026 SL												
Robot FW	4.14	4.14	3.11	3.11	4.14	3.11	3.11	4.14	3.11	4.14	4.14	4.14
SDLT 110/220 FW	v51	v51	V35	V35	v51	V35	V35	v51	v35	v51	v51	v51
MSL 5026 DLX												
Robot FW	4.14	4.14	3.11	3.11	4.14	3.11	3.11	4.14	3.11	4.14	4.14	4.14
Drive FW 40/80	v80	v80	V80	V80	v80	v80	V80	v80	v80	v80	v80	v80
SSL 2020 TL												
Robot FW	4.16	4.16	4.15	4.15	4.16		4.16	4.16	4.16	4.16	4.16	4.16
Drive FW AIT50	1.32	1.32	1.30	1.30	1.32		1.30	1.32	1.27	1.32	1.32	1.32

18.366
Paula

<<Previous Page >>Next Page

STORAGE HARDWARE / SOFTWARE COMPATIBILITY MATRIX

June 2003

Libraries | AutoLoaders | Operating Systems | RAID Arrays | Infrastructure | Servers | Fibre HBAs
SCSI HBAs | Software | Disk Array Multip <<Previous Firmware & Drivers >>Next | Backup App Details

Components & Firmware	Current Release	HP Data Protector	Atempo Time Navigator	BakBone NetVault	CA BrightStor ARCserve	CA BrightStor EB	CommVault Galaxy	Legato NetWorker	Syncsort Backup Express	Tivoli Storage Manager	VERITAS Backup Exec	VERITAS NetBackup
Autoloaders												
SSL1016 Ultrium 460 Tape Autoloader												
Robot FW	v10	v10			v10			v10			v10	v10
Drive FW Ultrium 460	F3CW	F3CW			F3CW			F3CW			F3CW	F3CW
1/8 Ultrium 230 Tape Autoloader												
Robot FW	H33R	H33R			H33R			H33R		H33R	H33R	H33R
Drive FW Ultrium 230	E00P	E00P			E00P			E00P		E00P	E00P	E00P
1/8 DLT vs80 Tape Autoloader												
Robot FW	0110	0110	0110	0110	0110	0110	0110	0110	0110		0110	
Drive FW vs80	v56	v56	v56	v56	v56	v56	v56	v56	v56		v56	
SSL1016 SDLT160/320 Tape Autoloader												
Robot FW	v31	v31			v31			v31			v31	v31
Drive FW SDLT 160/320	v52	v52			v52			v52			v52	v52
SSL1016 DLT1 Tape Autoloader												
Robot FW	v31	v31	v31	v31	v31	v13	v13	v31	v13		v31	v31
Drive FW DLT 1	v53	v53	v53	v53	v53	v46	v46	v53	v46		v53	v53
RAID Arrays												
XP 1024/128	21.04.04.00/00	21.04.04.00/00			21.04.04.00/00			21.04.04.00/00		21.04.04.00/00	21.04.04.00/00	21.04.04.00/00
XP 48/512	01.18.48.00/00	01.18.48.00/00			01.18.48.00/00			01.18.48.00/00		01.18.48.00/00	01.18.48.00/00	01.18.48.00/00
Zero Downtime with XP												
Instant Recovery with XP												
Enterprise Virtual Array (EVA) FW	2.003	2.003	v2_w020214-2000	v2_w020214-2000	2.003	2.0	2.0	2.003	2.0	2.003	2.003	2.003
VA 7410/7110	ha A100	ha A100			ha A100			ha A100			ha A100	ha A100
VA 7400/7100	hp18	hp18			hp18			hp18			hp18	hp18
EMA 16000	ACS 8.7	ACS 8.7	8.6	8.6	ACS 8.7	8.6	8.6	ACS 8.7	8.6	ACS 8.7	ACS 8.7	ACS 8.7
EMA 12000/8000	ACS 8.7	ACS 8.7	8.6	8.6	ACS 8.7	8.6	8.6	ACS 8.7	8.6	ACS 8.7	ACS 8.7	ACS 8.7
MSA1000												
MSA1000 Controller	2.38	2.38	1.14	1.14	2.38	1.14	1.14	2.38	1.14	2.38	2.38	2.38
MSA1000 GEM	1.86	1.86			1.86			1.86		1.86	1.86	1.86
Fibre Channel Interface Controllers												
NSR M2402	4.03.21	4.03.21	4.02	4.02	4.03.21	4.02.aw	4.02.aw	4.03.21	4.02.aw	4.03.21	4.03.21	4.03.21
NSR N1200	4.03.19	4.03.19			4.03.19			4.03.19		4.03.19	4.03.19	4.03.19
NSR E2400	4.03.19	4.03.19			4.03.19			4.03.19		4.03.19	4.03.19	4.03.19
NSR E2400-160	5.01.03	5.01.03			5.01.03			5.01.03				5.01.03
NSR E1200-160	4.03.24	4.03.24			4.03.24			4.03.24			4.03.24	4.03.24
NSR-E1200	4.03.19	4.03.19			4.03.19			4.03.19		4.03.19	4.03.19	4.03.19
MDR	v1187	v1187	v1186	v1186	v1187	v1187	v1187	v1187	V1187	V1187	v1187	v1187

18.365
Paula

<<Previous Page >>Next Page

STORAGE HARDWARE / SOFTWARE COMPATIBILITY MATRIX

June 2003

Components & Firmware	Current Release	HP Data Protector	Atempo Time Navigator	BakBone NetVault	CA BrightStor ARCserve	CA BrightStor EB	CommVault Galaxy	Legato NetWorker	Syncsort Backup Express	Tivoli Storage Manager	VERITAS Backup Exec	VERITAS NetBackup
Interconnects												
Director Switch 2/140	v4.01.02, HAFM (required) v06.03.01	v4.01.02, HAFM (required) v06.03.01			v4.01.02, HAFM (required) v06.03.01	v4.01.02, HAFM (required) v06.03.01		v4.01.02, HAFM (required) v06.03.01		v4.01.02, HAFM (required) v06.03.01	v4.01.02, HAFM (required) v06.03.01	v4.01.02, HAFM (required) v06.03.01
Director Switch 2/64	v4.01.02, HAFM (required) v06.03.01	v4.01.02, HAFM (required) v06.03.01			v4.01.02, HAFM (required) v06.03.01	v4.01.02, HAFM (required) v06.03.01		v4.01.02, HAFM (required) v06.03.01		v4.01.02, HAFM (required) v06.03.01	v4.01.02, HAFM (required) v06.03.01	v4.01.02, HAFM (required) v06.03.01
Edge Switch 2/32	v4.01.02, HAFM (required) v06.03.01	v4.01.02, HAFM (required) v06.03.01			v4.01.02, HAFM (required) v06.03.01	v4.01.02, HAFM (required) v06.03.01		v4.01.02, HAFM (required) v06.03.01		v4.01.02, HAFM (required) v06.03.01	v4.01.02, HAFM (required) v06.03.01	v4.01.02, HAFM (required) v06.03.01
Edge Switch 2/24	v4.01.02, HAFM (required) v06.03.01	v4.01.02, HAFM (required) v06.03.01			v4.01.02, HAFM (required) v06.03.01	v4.01.02, HAFM (required) v06.03.01		v4.01.02, HAFM (required) v06.03.01		v4.01.02, HAFM (required) v06.03.01	v4.01.02, HAFM (required) v06.03.01	v4.01.02, HAFM (required) v06.03.01
Edge Switch 2/16	v4.01.02, HAFM (required) v06.03.01	v4.01.02, HAFM (required) v06.03.01			v4.01.02, HAFM (required) v06.03.01	v4.01.02, HAFM (required) v06.03.01		v4.01.02, HAFM (required) v06.03.01		v4.01.02, HAFM (required) v06.03.01	v4.01.02, HAFM (required) v06.03.01	v4.01.02, HAFM (required) v06.03.01
Core Switch 2/64	4.0.2b	4.0.2b			4.0.2b	4.0.2b		4.0.2b		4.0.2b	4.0.2b	4.0.2b
SAN Switch 2/32	3.0.2k	3.0.2k			3.0.2k	3.0.2e		3.0.2k		3.0.2k	3.0.2k	3.0.2k
SAN Switch 2/16	3.0.2k	3.0.2k			3.0.2k	3.0.2e	3.0.2e	3.0.2k	3.0.2e	3.0.2k	3.0.2k	3.0.2k
SAN Switch 2/8	3.0.2k	3.0.2k			3.0.2k	3.0.2e	3.0.2e	3.0.2k	3.0.2e	3.0.2k	3.0.2k	3.0.2k
SAN Switch 2/16-EL	3.0.2k	3.0.2k			3.0.2k	3.0.2e	3.0.2e	3.0.2k	3.0.2e	3.0.2k	3.0.2k	3.0.2k
SAN Switch 2/8-EL	3.0.2k	3.0.2k			3.0.2k	3.0.2e	3.0.2e	3.0.2k	3.0.2e	3.0.2k	3.0.2k	3.0.2k
SAN Switch 16	2.6.0h	2.6.0h	2.6	2.6	2.6.0h	a2.6.0.c	a2.6.0.c	2.6.0h	a2.6.0.c	2.6.0h	2.6.0h	2.6.0h
SAN Switch 8	2.6.0h	2.6.0h	2.6	2.6	2.6.0h	a2.6.0.c	a2.6.0.c	2.6.0h	a2.6.0.c	2.6.0h	2.6.0h	2.6.0h
SAN Switch 16EL	2.6.0h	2.6.0h	2.6	2.6	2.6.0h	a2.6.0.c	a2.6.0.c	2.6.0h	a2.6.0.c	2.6.0h	2.6.0h	2.6.0h
SAN Switch 8EL	2.6.0h	2.6.0h	2.6	2.6	2.6.0h	a2.6.0.c	a2.6.0.c	2.6.0h	a2.6.0.c	2.6.0h	2.6.0h	2.6.0h
FC Switch 16B	3.0.2k	3.0.2k			3.0.2k	3.0.2e	3.0.2e	3.0.2k	3.0.2e	3.0.2k	3.0.2k	3.0.2k
FC Switch 8B	3.0.2k	3.0.2k			3.0.2k	3.0.2e	3.0.2e	3.0.2k	3.0.2e	3.0.2k	3.0.2k	3.0.2k
MSA Fabric Switch 2/8	3.0.2k	3.0.2k			3.0.2k	3.0.2k	3.0.2k	3.0.2k		3.0.2k	3.0.2k	3.0.2k
MSA Fabric Switch 6	101G12										101G12	
FC-AL Switch												
FC Loop Switch FW	1.20-g07	1.20-g07	1.20-g07	1.20-g07	1.20-g07		1.20-g07		1.20-g07		1.20-g07	
FC Loop Switch ROM	V1.2-b2	V1.2-b2	V1.2-b2	V1.2-b2	V1.2-b2		V1.2-b2		V1.2-b2		V1.2-b2	
HP OpenView Storage & StorageWorks Software												
OpenView												
Storage Virtual Replicator	3.0b sp1				3.0b sp1			3.0b sp1		3.0b sp1	3.0b sp1	3.0b sp1
OVSAM	3.0.1	3.0.1			3.0.1			3.0.1		3.0.1	3.0.1	3.0.1
StorageWorks												
Data Replication Manager	2.0a				2.0a			2.0a		2.0a	2.0a	2.0a
Enterprise Volume Manager	V2.0		V2.0		V2.0		V2.0	V2.0		V2.0	V2.0	V2.0

<<Previous Page >>Next Page

1844
Paula

STORAGE HARDWARE / SOFTWARE COMPATIBILITY MATRIX

June 2003

Libraries | AutoLoaders | Operating S | RAID Arrays | Infrastructure | Servers | Fibre HBAs
SCSI HBAs | Software | Disk Array Multipa <<Previous Firmware & Drivers >>Next | Backup App Details

Components & Firmware	Current Release	HP Data Protector	Atempo Time Navigator	BakBone NetVault	CA BrightStor ARCserve	CA BrightStor EB	CommVault Galaxy	Legato NetWorker	Syncsort Backup Express	Tivoli Storage Manager	VERITAS Backup Exec	VERITAS NetBackup
Tools and Utilities												
Smart Start	6.2	6.2	5.2, 5.3	5.2, 5.3	6.2	5.2, 5.3, 5.4	5.2, 5.3	6.2	5.2, 5.3, 5.4	6.2	6.2	6.2
Insight Manager 7	7				7			7		7	7	7
Disk Array Multipath Support												
AutoPath for HP-UX	2.01	2.01						2.01				2.01
Secure Path for HP-UX	3.0a, SP1	3.0a, SP1						3.0a, SP1		3.0a, SP1		3.0a, SP1
Secure Path for WIN NT	4.0a	4.0a	v3.1a sp2	v3.1a sp2	4.0a	4	4	4.0a	4	4.0a	4.0a	4.0a
Secure Path for WIN 2000	4.0a	4.0a	v3.1a sp2	v3.1a sp2	4.0a	4	4	4.0a	4	4.0a	4.0a	4.0a
AutoPath for Windows 2000											02.00.00	
Tru64 Multipath	Native							Native	Native			Native
Secure Path for NetWare	3.0c										3.0c	
Secure Path for Solaris	3.0a, SP1							3.0a, SP1		3.0a, SP1		3.0a, SP1
Secure Path for AIX	2.0b											
IBM DPO												
Operating System												
HP-UX	11.11 w/SP56 (HWE 11.11) PHKL_22857 11.00 w/SP56 (HWE 11.0) PHKL_25769	11.11 w/SP56 (HWE 11.11) PHKL_22857 11.00 w/SP56 (HWE 11.0) PHKL_25769	11.11	11.0		11.0	11.0	11.11 w/SP56 (HWE 11.11) PHKL_22857 11.00 w/SP56 (HWE 11.0) PHKL_25769	11.0, 11.11	11.11 w/SP56 (HWE 11.11) PHKL_22857 11.00 w/SP56 (HWE 11.0) PHKL_25769		11.11 w/SP56 (HWE 11.11) PHKL_22857 11.00 w/SP56 (HWE 11.0) PHKL_25769
Win NT 4.0 Server	sp6a	sp6a	sp6a	sp6a	sp6a	sp6a	sp6a	sp6a	sp6a	sp6a	sp6a	sp6a
Win 2000 Server	sp3	sp3	sp2	sp2	sp3	sp2	sp2	sp3	sp2	sp3	sp3	sp3
Tru64 UNIX	5.1 pk6, 5.1A pk4 5.1B pk1 (ipk)	5.1A pk4	5.1A	5.1A		5.1A	5.1A pk1	5.1 pk6, 5.1A pk4 5.1B pk1 (ipk)	5.1A pk2			5.1 pk6, 5.1A pk4
Linux – Red Hat	Advanced Server 2.1							Advanced Server 2.1				
NetWare	5.1 sp5 6.0 sp2	5.1 sp5 6.0 sp2		5.1 sp5	5.1 sp5 6.0 sp2				5.1 sp5 6.0 sp2		5.1 sp5 6.0 sp2	
Solaris UNIX	7, 8, and 9	7, 8, and 9	8	8		8	8	7, 8, and 9	8	7,8		7, 8, and 9
AIX	4.3.3 and 5.1L	4.3.3		4.3.3				4.3.3 and 5.1L		4.3.3 and 5.1L		4.3.3 and 5.1L

18.363
Paula

<<Previous Page >>Next Page

STORAGE HARDWARE / SOFTWARE COMPATIBILITY MATRIX

June 2003

Libraries | AutoLoaders | Operating Sy | RAID Arrays | Infrastructure | Servers | Fibre HBAs
SCSI HBAs | Software | Disk Array Multipath Previous Firmware & Drivers >>Next | Backup App Details

Components & Firmware	Current Release	HP Data Protector	Atempo Time Navigator	BakBone NetVault	CA BrightStor ARCserve	CA BrightStor EB	CommVault Galaxy	Legato NetWorker	Syncsort Backup Express	Tivoli Storage Manager	VERITAS Backup Exec	VERITAS NetBackup
Fibre Host Bus Adapters												
Windows												
FCA2355 (2G)												
FW	3.82a1	3.82a1			3.82a1			3.82a1		3.82a1	3.82a1	3.82a1
WIN NT 4.0 Driver	4-4.82a14	4-4.82a14			4-4.82a14			4-4.82a14		4-4.82a14	4-4.82a14	4-4.82a14
WIN 2000 Driver	5-4.82a14	5-4.82a14			5-4.82a14			5-4.82a14		5-4.82a14	5-4.82a14	5-4.82a14
WIN 2003 Driver	5-4.82a16	5-4.82a16										
ISP2312 (2G - BL20pG2 Blade Server)												
FW	1.33	1.33									1.33	1.33
WIN Driver	8.2.0.72	8.2.0.72									8.2.0.72	8.2.0.72
FCA2101 (2G)												
FW	3.82a1	3.82a1			3.82a1	3.81a3	3.81a3	3.82a1	3.81a3	3.82a1	3.82a1	3.82a1
WIN NT 4.0 Driver	4-4.82a14	4-4.82a14			4-4.82a14			4-4.82a14		4-4.82a14	4-4.82a14	4-4.82a14
WIN 2000 Driver	5-4.82a14	5-4.82a14			5-4.82a14	5-4.54a7	5-4.81a9	5-4.82a14	5-4.53a9	5-4.82a14	5-4.82a14	5-4.82a14
WIN 2003 Driver	5-4.82a16	5-4.82a16										
64-bit/33MHz PCI to FC (1G - Windows) KGPSA-CB												
FW	3.82a1	3.82a1	3.82a1	3.82a1	3.82a1	3.82a1	3.82a1	3.82a1	3.82a1	3.82a1	3.82a1	3.82a1
WIN NT 4.0 Driver	4-4.82a14	4-4.82a14			4-4.82a14	4-4.53a7	4-4.53a7	4-4.82a14	4-4.53a7	4-4.82a14	4-4.82a14	4-4.82a14
WIN 2000 Driver	5-4.82a14	5-4.82a14	5-4.52a9	5-4.52a9	5-4.82a14	5-4.53a7	5-4.52a9	5-4.82a14	5-4.53a7	5-4.82a14	5-4.82a14	5-4.82a14
WIN 2003 Driver	6-4.52a9				6-4.52a9			6-4.52a9			6-4.52a9	
LINUX												
FCA2214 (2G)												
FW	1.29											
Linux Driver (Intel)	6.04.00											
QLA2200 (1G)												
FW	1.76	1.76						1.76	1.76			
Linux Driver (Intel)	6.0.1c	6.0.1c						6.0.1c	6.0.1			
NetWare												
FCA2210 (2G)												
FW	1.29	1.29			1.29						1.29	
NetWare Driver	6.50s3	6.50s3			6.50s3						6.50s3	
64 Bit / 66 MHz PCI to FC HBA (1G)												
NetWare Driver	2.52	2.52		2.52	2.52				2.52		2.52	
32-bit / 33 MHz HBA (1G)												
NetWare Driver	2.52	2.52			2.52				2.52		2.52	
AIX												
IBM 6228 (2Gb)												
FW												
AIX Driver												
PCI to FC HBA for AIX (1G) 197819-B21												

18.202
Paula

<<Previous Page >>Next Page

STORAGE HARDWARE / SOFTWARE
COMPATIBILITY MATRIX

June 2003

Components & Firmware	Current Release	HP Data Protector	Atempo Time Navigator	BakBone NetVault	CA BrightStor ARCserve	CA BrightStor EB	CommVault Galaxy	Legato NetWorker	Syncsort Backup Express	Tivoli Storage Manager	VERITAS Backup Exec	VERITAS NetBackup
Fibre Host Bus Adapters												
HP-UX												
A6795A (2G)												
HP-UX Driver	11.11.09 11.0.10	11.11.09 11.0.10						11.11.09 11.0.10		11.11.09 11.0.10		11.11.09 11.0.10
A6685A (1G)												
HP-UX Driver	11.11.09	11.11.09						11.11.09 11.0.10				11.11.09 11.0.10
A5158A (1G)												
HP-UX Driver	11.11.09 11.0.10	11.11.09 11.0.10	11.00.06	11.00.06		11.00.06	11.00.06	11.11.09 11.0.10	11.00.06	11.11.09 11.0.10		11.11.09 11.0.10
Tru64												
FCA2354												
FW	3.81a4							3.81a4				3.81a4
Tru64 Driver	2.06							2.06				2.06
KGPSA -CA												
FW	3.81a4	3.81a4	3.81a4	3.81a4		3.81a4	3.81a4	3.81a4	3.81a4			3.81a4
Tru64 Driver	2.06	2.06	1.32a	1.32a		1.32a	1.32a	2.06	2.02			2.06
Solaris												
FCA2257P (2G - PCI)												
Fcode	3.1.2											
Solaris Driver	3.26											
FCA2257C (1G - PCI)												
Fcode	3.1.2											
Solaris Driver	3.26											
FCA2257S (1G - Sbus)												
Fcode	3.1.2											
Solaris Driver	3.26											
SWSA4-SC 64-bit Sbus												
Fcode	13.3.7	13.3.7	13.3.7	13.3.7		13.3.7	13.3.7	13.3.7	13.3.7	13.3.7		13.3.7
Solaris Driver	2.5.9	2.5.9	2.5.9	2.5.9		2.5.9	2.5.9	2.5.9	2.5.9	2.5.9		2.5.9
SWSA4-PC 32-bit PCI												
Fcode	3.0.3							3.0.3		3.0.3		3.0.3
Solaris Driver	2.5.9							2.5.9		2.5.9		2.5.9
JNI FCE6460 (2G - PCI)												
Fcode	3.89	3.89						3.89				3.89
Solaris Driver	5.2.1	5.2.1						5.2.1				5.2.1
JNI FCE6410 (1G - PCI)												
Fcode	3.9	3.9						3.9				3.9
Solaris Driver	4.13	4.13						4.13				4.13

18.361
Paula

<<Previous Page

STORAGE HARDWARE / SOFTWARE COMPATIBILITY MATRIX

June 2003

Libraries | AutoLoaders | Operating Syste RAID Arrays | Infrastructure | Servers | Fibre HBAs
SCSI HBAs | Software | Disk Array M ath | Firmware & Drivers | **Backup App Details**

ISV DETAILS

Name	Version	Date Last Tested	Recommended Tape Driver	Application Information/Support
HP OpenView Storage Data Protector	5.0	Current	Use Data Protector Driver for Windows	Data Protector Product Information Data Protector Device Support Matrix
Atempo Time Navigator	3.6	Mar-02		www.atempo.com
BakBone NetVault	6.5.1	Mar-02		www.bakbone.com
Computer Associates BrightStor ARCserve for Windows NT/2000	ARCServe v9.0 build 1868	Current	CABA Driver for Windows	support.ca.com
Computer Associates BrightStor ARCserve for NetWare	9.0	Current	HP Driver for Netware	support.ca.com
Computer Associates BrightStor Enterprise Backup	BEB v10.0 SP4 Bd. 1400	Aug-02		support.ca.com/
CommVault Galaxy	4.1.0	Aug-02		www.commvault.com
Legato NetWorker	7.0 and 6.1.3	Current	HP Driver for Windows	www.legato.com
Syncsort Backup Express	2.1.5B	Aug-02	Use Native Drivers	www.syncsort.com
Tivoli Storage Manager	5.2	Current	Tivoli Driver for Windows	www.tivoli.com
VERITAS Backup Exec for Windows Servers	9.0 build 4367.2	Current	BE Driver for Windows	www.veritas.com
VERITAS Backup Exec for NetWare	9.0 build 4175.1	Current	BE Driver for Netware	www.veritas.com
VERITAS NetBackup Windows	4.5 FP3 and 3.4.1 MP3	Current	HP Driver for Windows	www.veritas.com
VERITAS NetBackup UNIX	4.5 FP4 and 3.4 MP4	Current		www.veritas.com

NOTE: All ISV certifications have been completed with a minimum configuration of 16 servers by 16 drives.
This does not imply any limitations of the ISV backup and restore application or HP StorageWorks SAN products.
Refer to the ISV website or your ISV representative for maximum configuration details.

18.360
Paula

18.354
Paula

ANEXO UNIDADE DE BACKUP ROBOTIZADO
PARTE Q

18.358
Paula

hp rack-optimized
rp5430 and rp5470
servers
entry-level UNIX® servers

hp servers
product brief

hp rack-optimized servers can help you be always on, always there, always connected

hp servers rp5430 and rp5470: smart, simple, stress-free

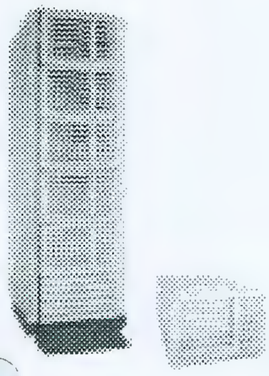

hp servers rp5430 and rp5470

In today's economy, whether you're managing your own IT infrastructure or hosting someone else's, you have to operate with a faster time-to-solution, within budgetary constraints, and with the highest standards for customer service and operational efficiency.

To create and run an infrastructure for an always-on business, you need a computing platform that will support the way you—and your customers—do business. The HP Servers rp5430 and rp5470 give your business the fastest—and most reliable—means of succeeding in this new business environment.

The HP Servers rp5430 and rp5470 deliver the proven performance, scalability, and high-availability capabilities of UNIX—without high maintenance requirements and costs. And they give you plenty of room to grow. You can start at a low-price entry point and scale up to the leading 4-way UNIX performance—in the same rack optimized form factor, without penalty. And with its industry-leading solution partners, HP has developed business solutions surrounding these servers that are tested, easy to deploy, and easy to manage.

With HP Servers rp5430 and rp5470, owning and operating a UNIX server is smart, simple, and stress-free.

smart

HP Servers rp5430 and rp5470 offer leading entry-level server performance, dynamic scalability, and unmatched investment protection—all in a rack-optimized package—making them the smart choice for the most demanding applications.

leading performance—scalable functionality

- industry-leading OLTP performance
- massive bandwidth for I/O-intensive applications
- 7U-height and packed with CPUs, memory, and I/O, plus the ability to scale subsystems without compromise
- rack-optimized to make the best use of valuable data-center floorspace

unparalleled investment protection

- industry's only in-box upgrade from 2-way to 4-way UNIX computing
- built-in growth path to the HP Server rx5670, featuring Intel® Itanium® 2 processors

18357
Paula

proven solutions for business-critical computing

simple

hp makes IT easy

With HP Servers rp5430 and rp5470, HP offers a unique combination of solutions designed to help you get started quickly and manage your IT environment effectively.

hp-ux: robust, proven enterprise operating environment for mission-critical applications

- industry-leading performance, scalability, availability, manageability, and security
- pre-packaged, integrated, and production-ready operating environments
- industry-leading Windows® and Linux interoperability
- centralized software updates that are timely, simple, and efficient
- powerful alliances with industry-leading software vendors and systems integrators to deliver robust solutions from e-commerce to enterprise resource planning (ERP) and beyond

ding-edge management capabilities with hp-ux virtual partitions and hp-ux workload manager

- system resource optimization enabling multiple workloads to run simultaneously on the same server, each with their own instance of HP-UX
- improved security and server availabilty through complete software and operating system isolation
- HP-UX Virtual Partitions integrate with HP-UX Workload Manager for the most efficient resource distribution across partitions, in a single server
- base offering complementary with HP-UX 11i for your HP Server rp5470

integrated management capabilities

- HP Servicecontrol Manager and integrated HP Secure Web Console capability for full remote management, including centralized configuration of multiple servers

the right server for today's applications

ERP (supply-chain management)

High availability, leading-edge manageability, and scalable performance support demanding end-to-end enterprise applications.

broadband

The combination of leading performance, I/O throughput, and capacity and high availability with end-to-end solutions delivers more powerful, reliable broadband services.

Internet infrastructure

Highly scalable, reliable, and manageable Web server, caching server, load balancing, e-commerce server, firewalls, or mail server.

e-commerce

Leading performance and I/O bandwidth, in-box scalability, rack-optimized form factor (five per standard 2-meter rack), high availability, and Internet management features.

technical computing

Leading performance and throughput, N+1 redundant components, extensive memory capabilities, and clustering solutions to meet the demands of computation, NFS file serving and product data management, and Web hosting.

flexible financing

- operating leases with a variety of attractive terms
- the Tech Refresh program for cost-effective upgrades to stay on the leading edge of new technology
- bundled Solution Finance program to consolidate and simplify financing arrangements

instant capacity on demand for your hp server rp5470

iCOD

- instant activation of incremental CPU power when you need it
- pay only for the processing power you use

temporary capacity for iCOD

- temporary activation of incremental CPU power for a limited period
- ideal for short-term, predictable processing demands

stress-free

make your business your focus

The HP Servers rp5430 and rp5470 handle the demands of users efficiently and reliably—so you can concentrate on running your business, not managing your IT resources.

high availability for continuous operations

- a rich set of in-box high-availability features
- affordable high-availability clustering solutions based on industry-leading HP MC/Serviceguard
- self-healing capabilities, a first for entry-level servers—proactively avoid faults to improve uptime
- a critical building block for your always-on e-business needs

total solution support

- options ranging from Web-based services to the industry's only 6-hour call-to-repair commitment
- "one-stop" solution support delivered with partners such as Cisco, Oracle®, SAP, i2, Inktomi, and many more

best UNIX server family—top to bottom

The low-cost entry point to the two-way HP Server rp5430 through the more scalable HP Server rp5470 are part of the powerful HP UNIX server line—servers that set the standards for business-critical computing and total cost of ownership. Simple in-chassis upgrades let you move up the line as your business grows. HP UNIX servers provide the hardware foundation for an Internet infrastructure that is always on. Combining leading technology with proactive and reactive services, HP offers complete, end-to-end solutions that include hardware, software, applications, services, support, consulting, and an extensive portfolio of experienced partners, so you can get to market quickly with a single source of expertise. The HP UNIX server family is robust from the top to bottom—from the high-end HP Superdome; through the HP Servers rp8400 and rp7410, the midrange performance and price/performance leaders; and to the scalable entry-level HP Servers rp5430 and rp5470. Rounding out the family are the hyper-dense HP rp2400 series servers for Internet applications and branch offices.

18.356
Paula

hp servers rp5430 and rp5470 features and benefits at-a-glance

features	benefits
smart	
1–4 (rp5470) or 1–2 (rp5430) 875MHz PA-8700+ CPUs with 2.25 MB on-chip cache per CPU or 750MHz PA-8700 CPUs with 2.25 MB on-chip cache per CPU	Superior performance over comparable systems, with plenty of headroom for growth
Up to 16 GB (rp5470) or up to 8 GB (rp5430) SDRAM memory with advanced ECC protection	Fast and reliable processing power for frequently accessed data
Up to 10 (rp5470) or 6 (rp5430) PCI I/O slots with 3.2 GB/s (rp5470) or 2.3 GB/s (rp5430) I/O bandwidth	Easily handles I/O-intensive applications and allows the system to scale I/O, CPUs, and memory without compromise
Core I/O, including 10/100Base-T LAN with auto speed-sensing, a second 10/100Base-T support LAN, Ultra2 LVD SCSI, and RS-232	Provides easy, ready-to-go networking capabilities
Up to 4 internal 36 GB, 73 GB, or 146 GB Ultra320 SCSI hot-plug high-uptime disks	Store critical data with massive internal capacity
1 internal DVD or DAT drive	Protects critical data
7U chassis with up to 5 servers per standard 2-meter rack; also available in standalone (pedestal) configuration	Optimizes use of floor space and delivers high-performance density in a racked configuration
Easy in-box upgrades from the rp5430 to the rp5470	Architectural scalability ensures these servers can grow with the business, maximizing flexibility and investment protection
Support of Intel Itanium Processor Family as well as PA-RISC processors	Provides investment protection through dual growth paths
simple	
Built-in unlimited user license for proven 64-bit HP-UX 11i and 11.0	Proven, enterprise UNIX operating system for mission-critical applications
HP Virtual Partitioning for the rp5470	Maximizes usage of computing resources
Integrated HP Secure Web Console and Servicecontrol Manager for full local, Web, and remote control of servers	Provide complete single-system and multisystem administration capabilities, including a range of security features, from any browser-based PC
Flexible financing programs	Make initial ownership and modular growth easy and affordable
Instant capacity on demand (iCOD); temporary capacity for iCOD	Immediate access to CPU power when you need it, either permanently or temporarily
HP global deployment and partner integration services	Offer guaranteed error-free solution deployment to reduce implementation time and cost
HP On-site Solution consolidated manufacturing; streamlined product assembly and testing, and state-of-the-art integration	Ensure superior quality and faster delivery
stress-free	
Error-correcting cache, parity checking on all buses, memory scrubbing and page de-allocation, dynamic processor resilience, and de-allocation of application processes Dual Ultra 2 SCSI buses and controllers for mirrored storage Hot-swap, redundant power supplies and fans; redundant, hot-plug PCI; Ultra2 SCSI hot-plug disks	Built-in high-availability features deliver superior levels of • error correction, • error containment, • data protection, and • serviceability to help maximize uptime for business-critical workgroups and applications
Integrated Event Monitoring Service (EMS) Built-in fault management system with separate support processor and bus Integrated with HP MC/Serviceguard, HP Toptools for Servers, and enterprise management software such as HP OpenView and CA Unicenter	Provides superior system uptime through constant, proactive fault avoidance, detection, and notification; monitors power, cooling system hardware, processors, memory, HP-UX resources, and external storage
Pre-tested and pre-integrated workgroup clustering solutions based on HP MC/Serviceguard	Deliver complete, ready-to-go solutions for clustered high availability that eliminate all single points of failure, at an affordable price
3-year on-site service warranty and HP services and support options ranging from Web-based support to mission-critical, 6-hour Call-to-Repair commitment; includes full solution support for hardware and software	Reduce risk through worldwide support for business-critical computing; provide "one-stop shopping" for support through partnerships with Cisco, Oracle, SAP, BroadVision, Inktomi, and others

18.353
Paula

hp surestore tape array 5300

getting started guide

rack enclosure

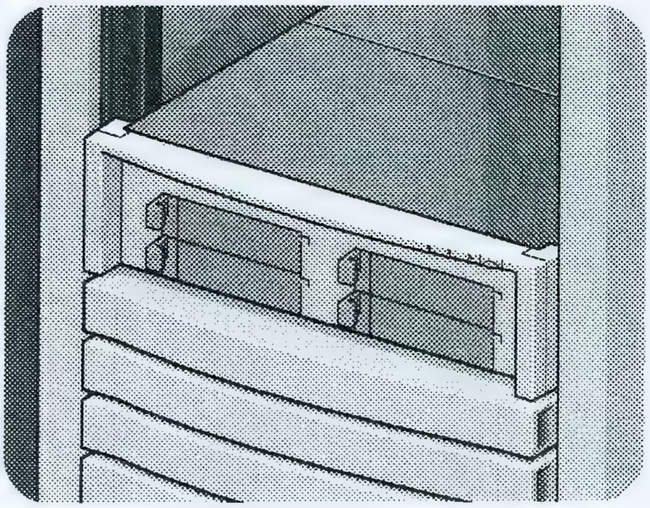

for tape drives

18.357
Paula

open this flap for quick start installation poster

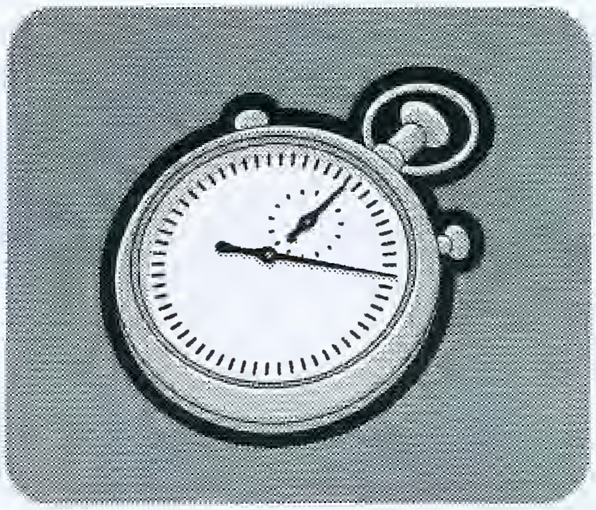

18.351
Paula

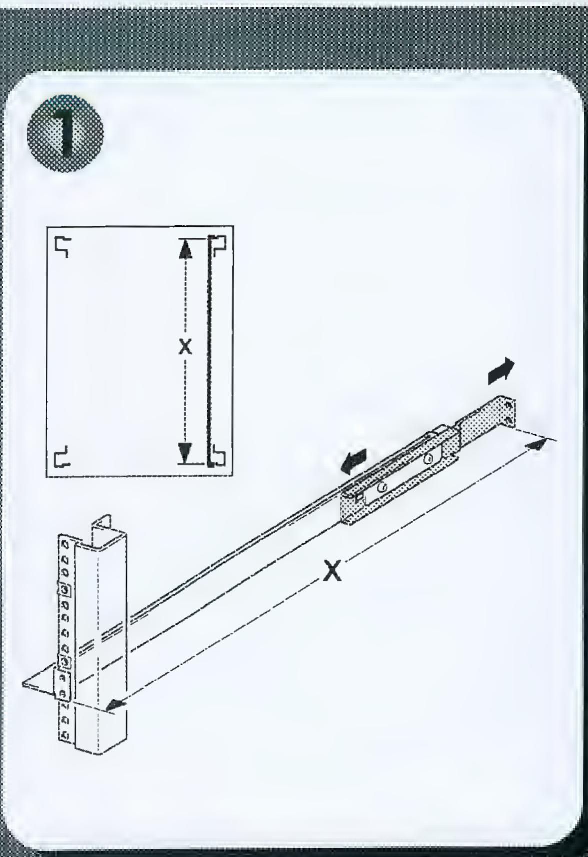

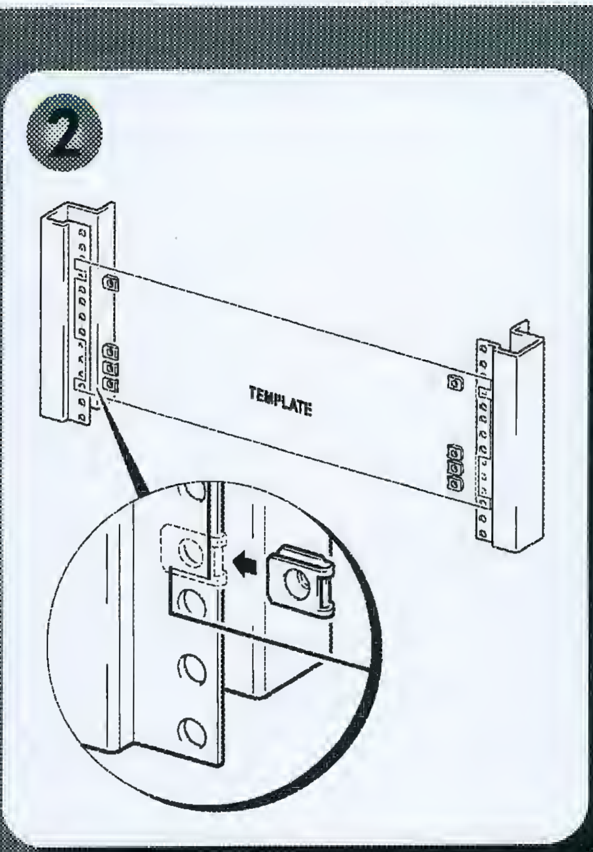

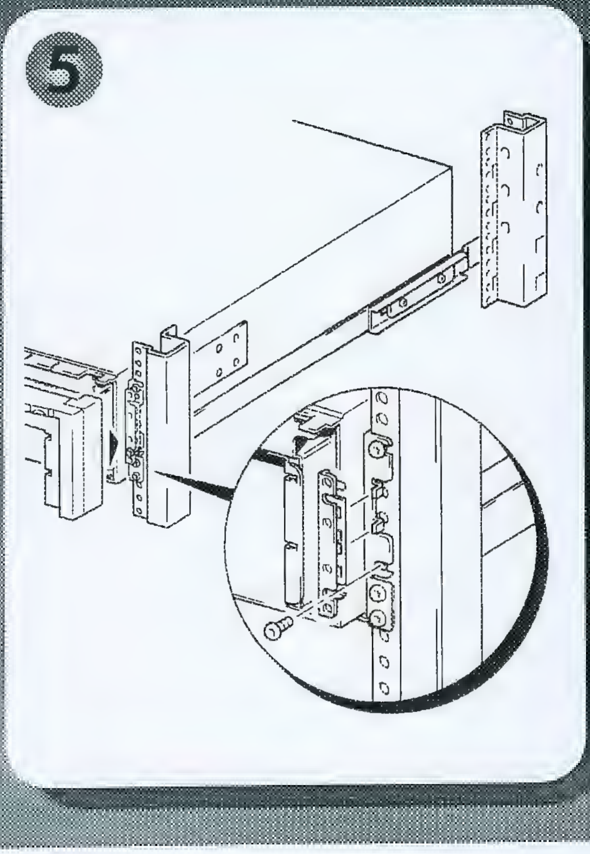

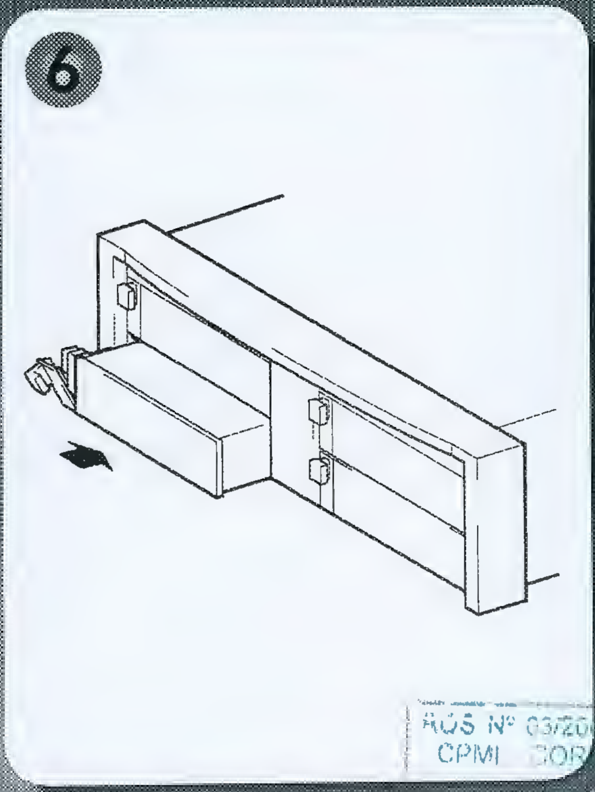

18.350
Paula

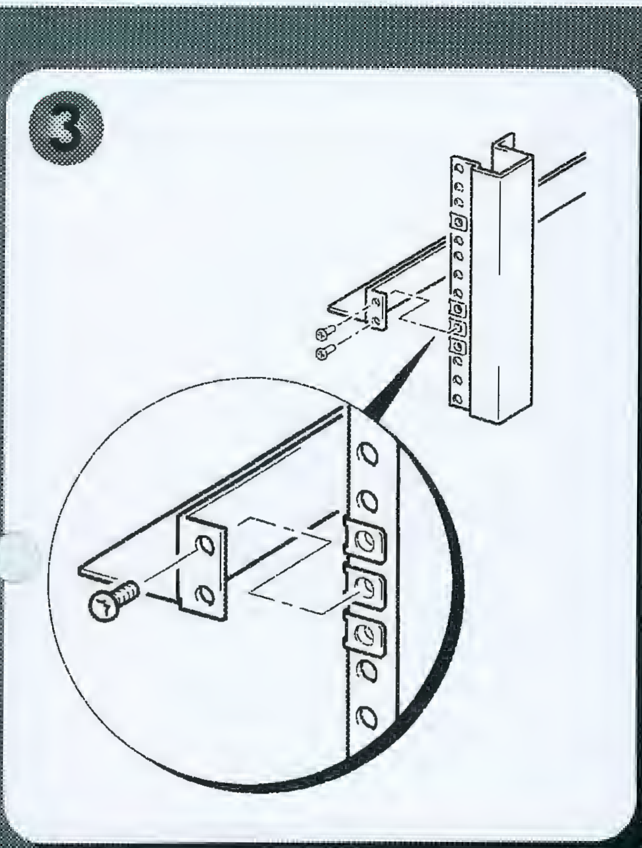

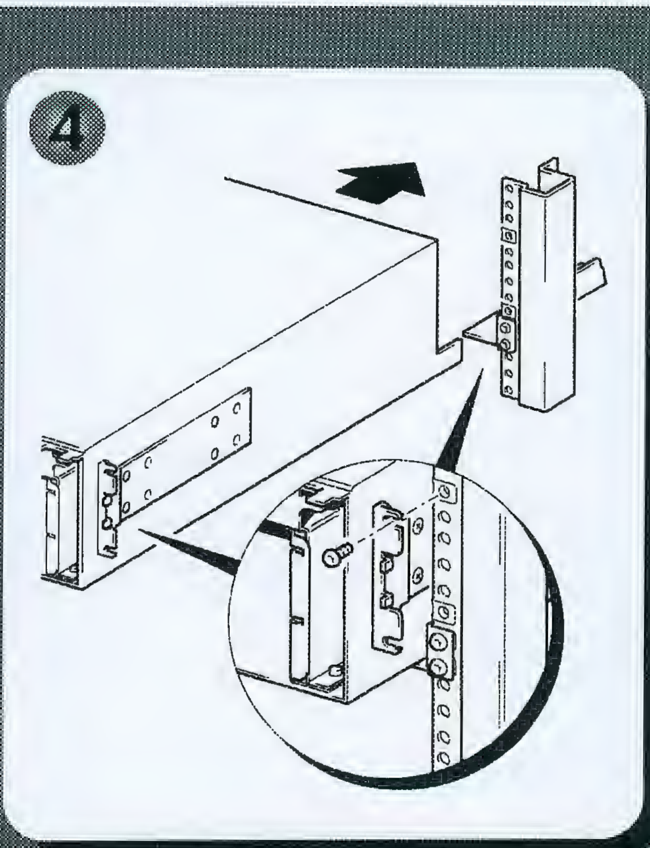

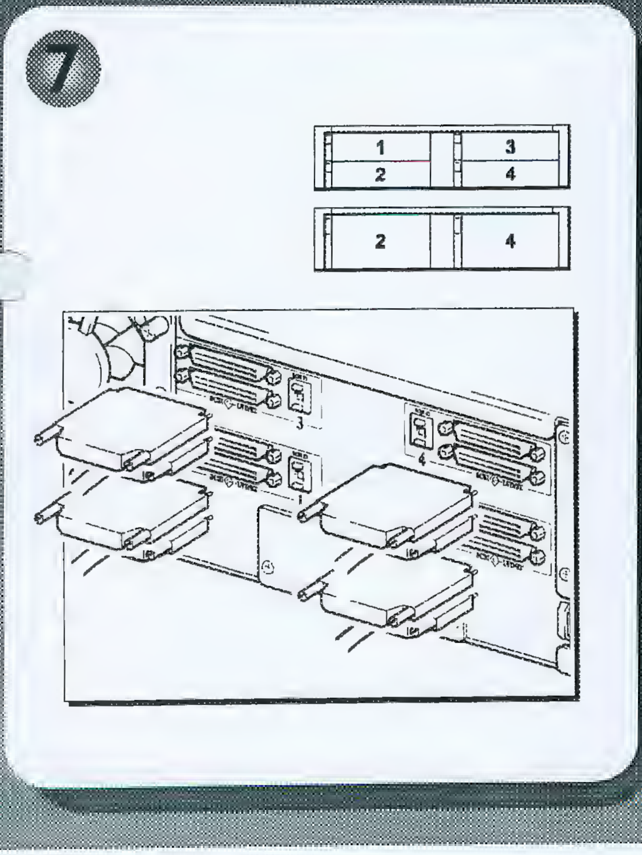

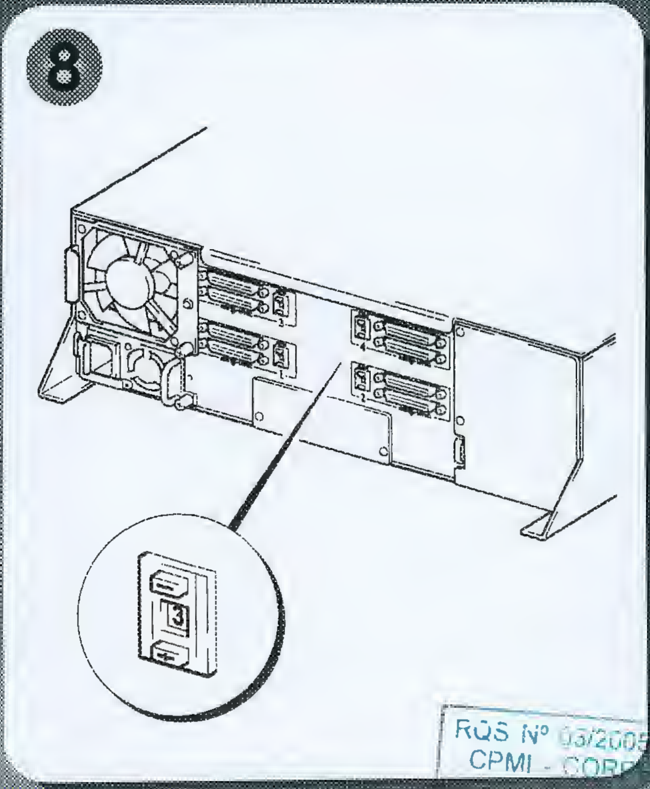

18.349
Paula

HP Surestore Tape Array 5300 - overview

For further details, see:

Introduction page 3

Step 1: Adjust the length of the support rails page 5

Step 2a: Install M5 clip nuts in HP racks page 7

2b: Install M6 cage nuts in non-HP racks page 9

Step 3a: Install support rails in HP racks page 11

3b: Install support in non-HP racks page 13

Step 4a: Install the tape array enclosure in HP racks page 15

4b: Install the tape array enclosure in non-HP racks page 17

Step 5a: Install the front bezel on HP racks page 19

5b: Install the front bezel in non-HP racks page 21

Step 6: Install the tape drive module page 23

Step 7: Connect the SCSI cables, terminator and power cord page 25

Step 8: Set the SCSI ID page 27

Your HP Surestore Tape Array 5300 page 29
Advice on daisy-chaining page 31
Replacing a tape drive module, fan, or power supply page 33
Installing high availability/management card upgrade kits page 35
Troubleshooting the HP Surestore Tape Array 5300 page 36
Product Specifications page 39
Warranty page 40

18.318
Paula

March 2001
Part Number: C7508-90901

The information contained in this document is subject to change without notice.

Hewlett-Packard shall not be liable for errors contained herein or for incidental or consequential damages (including lost profits) in connection with the furnishing, performance, or use of this material whether based on warranty, contract, or other legal theory.

Product Details

Write your HP Surestore Tape Array 5300 details here so you can find them easily if you need them

Serial number:		
Date purchased/installed:		
SCSI ID:	Bay 1	
	Bay 2	
	Bay 3	
	Bay 4	

18.347
Paula

Introduction

The HP Surestore Tape Array 5300 is designed as a low-cost backup solution that merges high performance of backup and recovery in one compact, modular, and easily serviced package.

The HP Surestore Tape Array 5300 supports up to four half-height, or two full-height tape drive modules and up to four independent SCSI buses. The backup capacity of one high-density enclosure depends on the type of tape drive installed; refer to the documentation shipped with your tape drive modules for full details.

Features

The HP Surestore Tape Array 5300 provides:

- High-density storage: the HP Surestore Tape Array 5300 supports up to four half-height tape drive modules in the footprint (3U).
- Rack optimized design: the HP Surestore Tape Array 5300 is designed for installation in HP, Compaq, and other compatible19-inch racks.
- Easy serviceability: the tape drive modules are offline hot-swappable[1] for easy in-rack servicing without downtime.
- Supportability: the HP Surestore Tape Array 5300 supports full-height and half-height tape drive modules, across a full range of technologies and capacities.

High availability/management card upgrade kit options

The HP Surestore Tape Array 5300 has a single fan and power supply unit. An optional upgrade kit is available to fit an additional fan and power supply unit to provide high-availability, hot-swap functionality whereby should one unit suffer component failure, the other unit will maintain correct operation. A further option allows the addition of a SCSI management card to provide a remote management and diagnostic functionality. When used with HP Library and Tape Tools software, this provides reporting outputs via SCSI for over temperature, fan failure and power supply unit failure.

Tools needed

You need the following tools for the installation procedure:

- Large or small Philips screwdrivers depending on the type of rack.
- T15 and T25 Torx® wrenches (provided).
- Template (provided) for marking the clip nut or cage nut location on the rack columns.

SCSI cables and terminators

You will also need to order the correct size SCSI cables and terminators for your tape drive modules to suit your particular application (see page 25). Cable and terminator options are listed at the HP support web site http://www.hp.com/support/tapearray. You should ensure you have these cables and terminators before starting the installation.

[1] Offline hotswap is defined as the ability to plug and unplug drive modules with the power still connected, but not while data transfers are taking place onto other drives on the same SCSI bus.

18.346
Paula

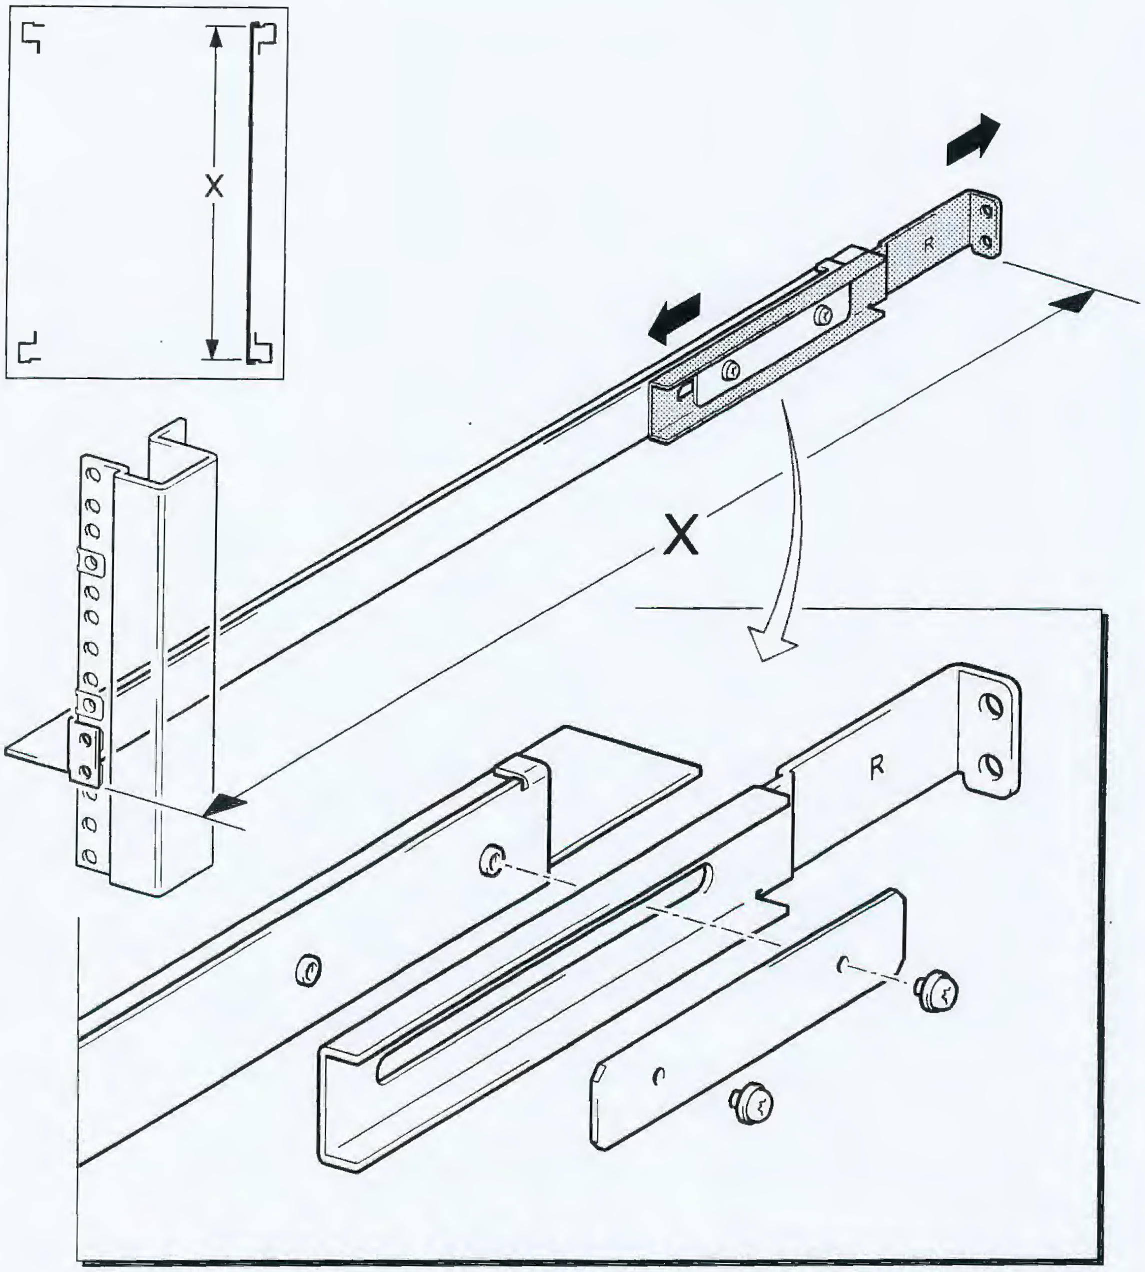

Figure 1: Adjust the length of the support rails

18.305
Paula

Step 1: Adjust the length of the support rails

The side support rails are pre-assembled to fit HP rack systems. The length of these rails may need adjusting to fit different manufacturer's racks.

To check and adjust the length of the rails:

1 Open the front and rear door panels of the rack.

Warning The HP Surestore Tape Array 5300 may need to be installed in a rack system that already contains other items of equipment where the power is turned on. Take care when working alongside other equipment in the rack.

In most cases, adequate access to install the HP Surestore Tape Array 5300 can be gained from the front and rear of the rack without having to remove the side panels. It is not necessary to disconnect the power to other items of equipment already installed in the rack but care must be taken to avoid damaging or disturbing connections.

2 Carefully offer up a support rail with the telescopic section of the rail towards the rear of the rack. It is much easier if this is done by two persons. The telescopic section is stamped with the letter "R" for the support rail to be installed on the right hand side of the rack, when the rack is viewed from the front, and with the letter "L" for the left hand side of the rack.

3 Compare the length of the rail with the distance between the front and rear rack columns.

Note To aid installation the support rail is inscribed with two lines, stamped "A" and "B". When the telescopic section is aligned with scribe line "B" (as pre-assembled) the support rail is the correct nominal length to fit HP rack systems. When adjusted to align with scribe line "A", the support rail is the correct nominal length to fit Compaq rack systems. However, the length of the support rail should always be checked as described below and may need adjusting to take account of the rack's manufacturing tolerances.

4 Check that the length is correct to allow the flange on the front of the rail to fit over the front rack column and the flange at the rear of the rail to fit over the rear rack columns, see Figure 1.

If adjustment is not required, proceed to install the HP Surestore Tape Array 5300 as described on page 7 (installing in HP racks) or page 9 (installing in non-HP racks).

5 Use the T15 Torx® wrench to loosen the two securing screws on the support rail's mounting plate.

6 Extend the telescoping section of the rail to the maximum length allowed by the slotted hole.

7 Re-measure the support rail against the rack columns and adjust to the correct length.

Take care to maintain the measured length of the rail and remove the rail from the rack.

8 Tighten the two screws.

9 Repeat paragraph 2 to 8 for the other support rail.

18.344
Paula

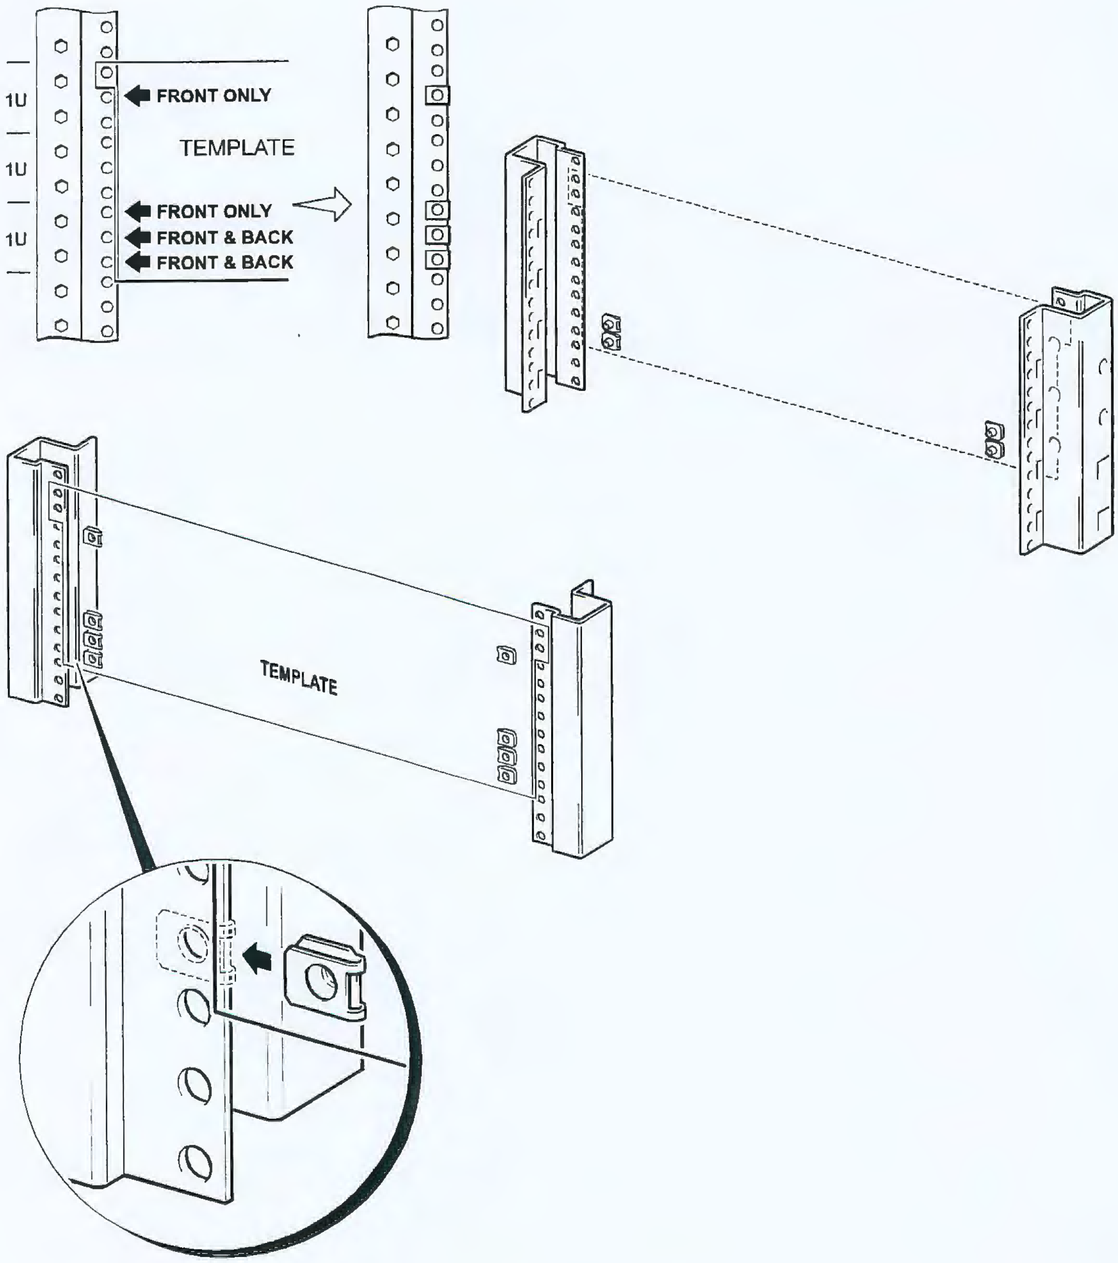

Figure 2a: Install M5 clip nuts in HP racks

18.343
Paula

Step 2a: Install M5 clip nuts in HP racks

1 Locate the installation template supplied with your HP Surestore Tape Array 5300.

2 Align the template with the front columns of the rack so that the '#' markings on the template align with the required EIA unit[1] numbers on the rack column. Make sure the template is in the position you require for installing the HP Surestore Tape Array 5300.

Note The round profile holes in the rack's front columns are not uniformly spaced.

3 Place the first eight clip nuts, as shown in Figure 2a, at the locations indicated by the black arrows marked either "FRONT ONLY" or "FRONT & BACK" on the template (four clips to each side). This is done by sliding the clip nut over the hole of the rack column until it snaps into place.

4 Note which numbered EIA units were used on the front columns.

5 Repeat paragraph 2 and 3 for each of the rear columns using the lower arrow locations indicated on the template marked "FRONT & BACK" (two clips to each side).

[1] Vertical space within a rack is measured by industry-standard EIA units. (1 EIA unit = 1.75 inches or 44.5 millimeters.) Equipment height is also specified in EIA units. System configuration is made easier by counting EIA from the base of the rack. EIA units are numbered on the vertical columns of many (but not all) rack systems.

18.3442
Paula

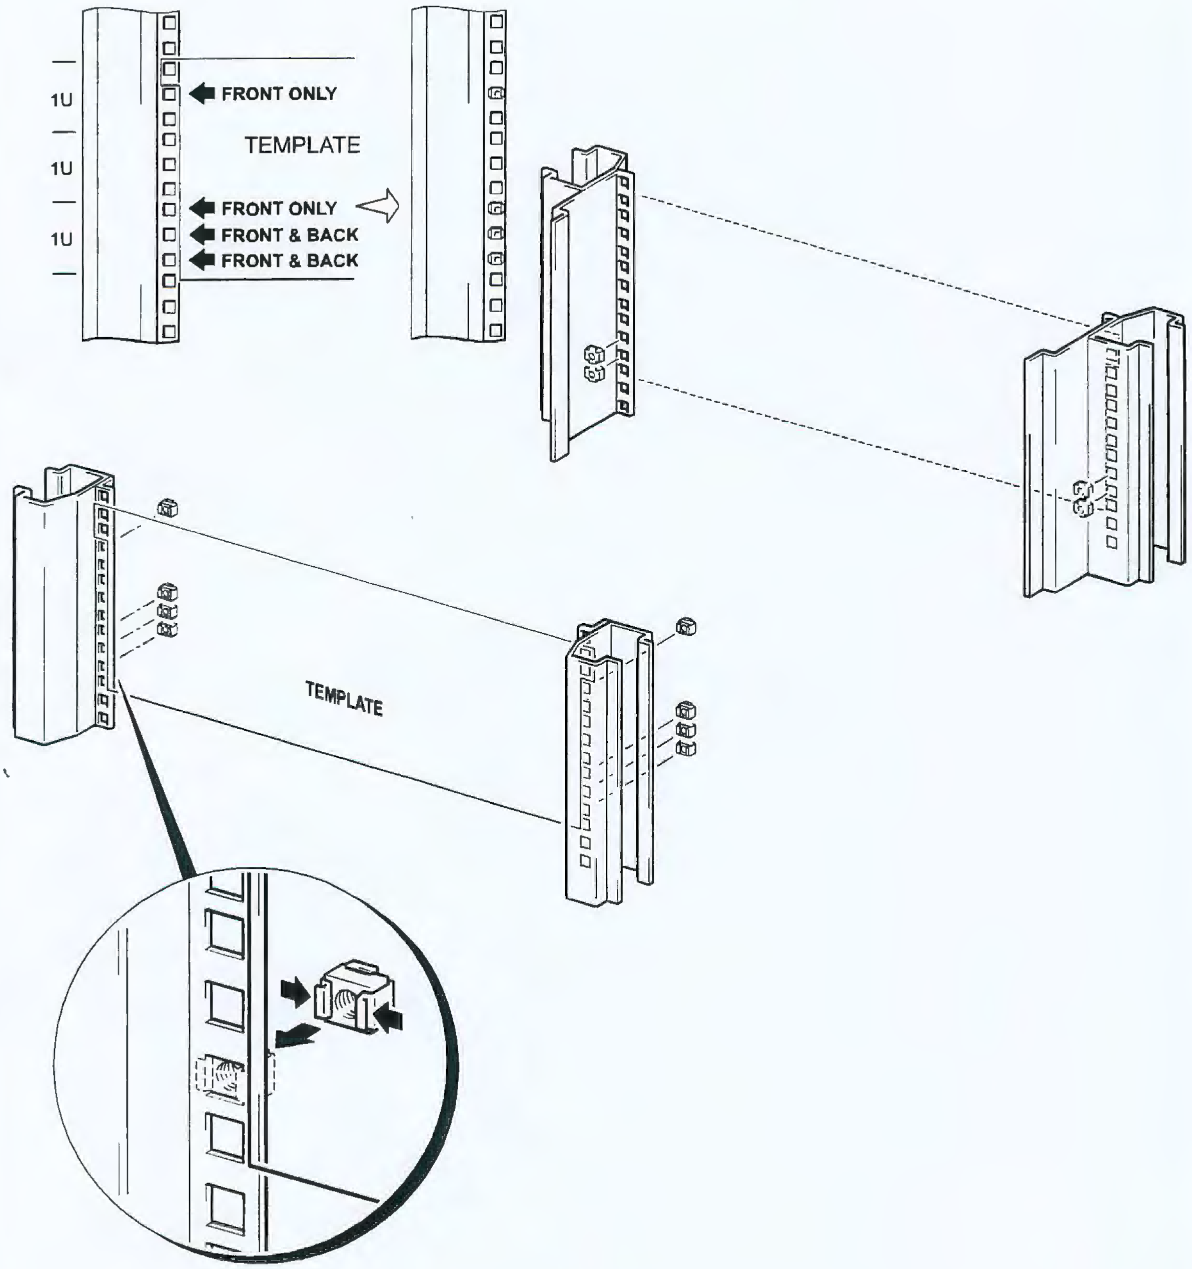

Figure 2b: Install M6 cage nuts in non-HP racks

Step 2b: Install M6 cage nuts in non-HP racks

Note The HP Surestore Tape Array 5300 can be installed in non-HP manufactured racks, such as Compaq and IBM, which have square profile cutouts in the vertical mounting rack column. The square profile cutouts accept M6 mounting cage nuts.

1. Locate the installation template supplied with your HP Surestore Tape Array 5300.
2. Align the template with the front columns of the rack so that the '#' markings on the template align with the required EIA unit[1] numbers on the rack column. If EIA numbers are not given, carefully count the holes in the cabinet columns from the base of the cabinet, and align the template arrows to the holes chosen. Make sure the template is in the position you require for installing the HP Surestore Tape Array 5300.

   Note The square profile holes in the rack's front columns are not uniformly spaced.

3. Place the first eight cage nuts, as shown in Figure 2b, at the locations indicated by the black arrows marked either "FRONT ONLY" or "FRONT & BACK" on the template (four cage nuts to each side). This is done by squeezing the sides of the cage nut and inserting it into the square hole from the inside of the rack column. The cage nuts are spring loaded. If necessary, use the blade of a screwdriver to snap it into place.
4. Note which numbered EIA units were used on the front columns, or count up an equal number of holes from the base.
5. Repeat paragraph 2 and 3 for each of the rear columns using the lower arrow locations indicated on the template marked "FRONT & BACK" (two cage nuts to each side).

[1] Vertical space within a rack is measured by industry-standard EIA units. (1 EIA unit = 1.75 inches or 44.5 millimeters.) Equipment height is also specified in EIA units. System configuration is made easier by counting EIA from the base of the rack. EIA units are numbered on the vertical columns of many (but not all) rack systems.

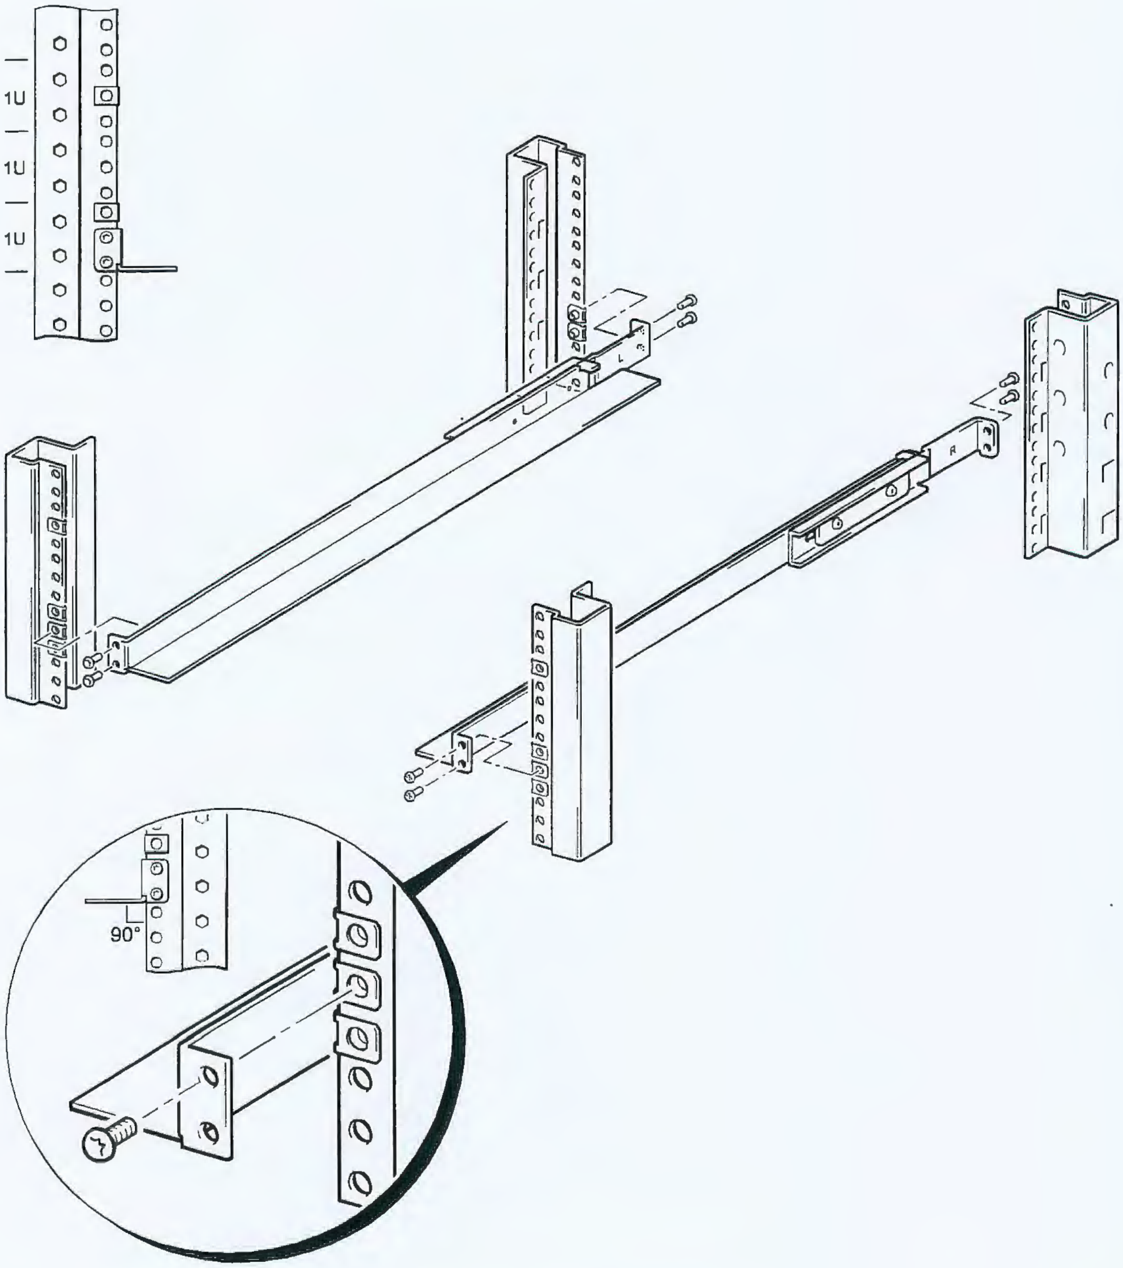

Figure 3a: Install support rails in HP racks

18.339a
Paula

Step 3a: Install support rails in HP racks

1 Refer to Figure 3a. Position a support rail so that the flange at the rear of the rail locates over the rear rack column and the two holes in the rail flange align with the two clip nuts. It is much easier if this is done by two persons.

2 Insert the M5 screws. Partially finger-tighten the screws only at this stage.

3 Position the rail so that the flange at the front of the rail locates over the front rack column and the two holes in the rail flanges align with the bottom two clip nuts.

4 Use the T25 Torx® wrench to tighten the M5 screws. Ensure the rail is secured so that the support flange is at 90 degrees to the vertical rack column.

5 Tighten the screws on the rear rack columns.

6 Repeat paragraphs 1 through 5 for the opposite side of the rack.

18.338
Paula

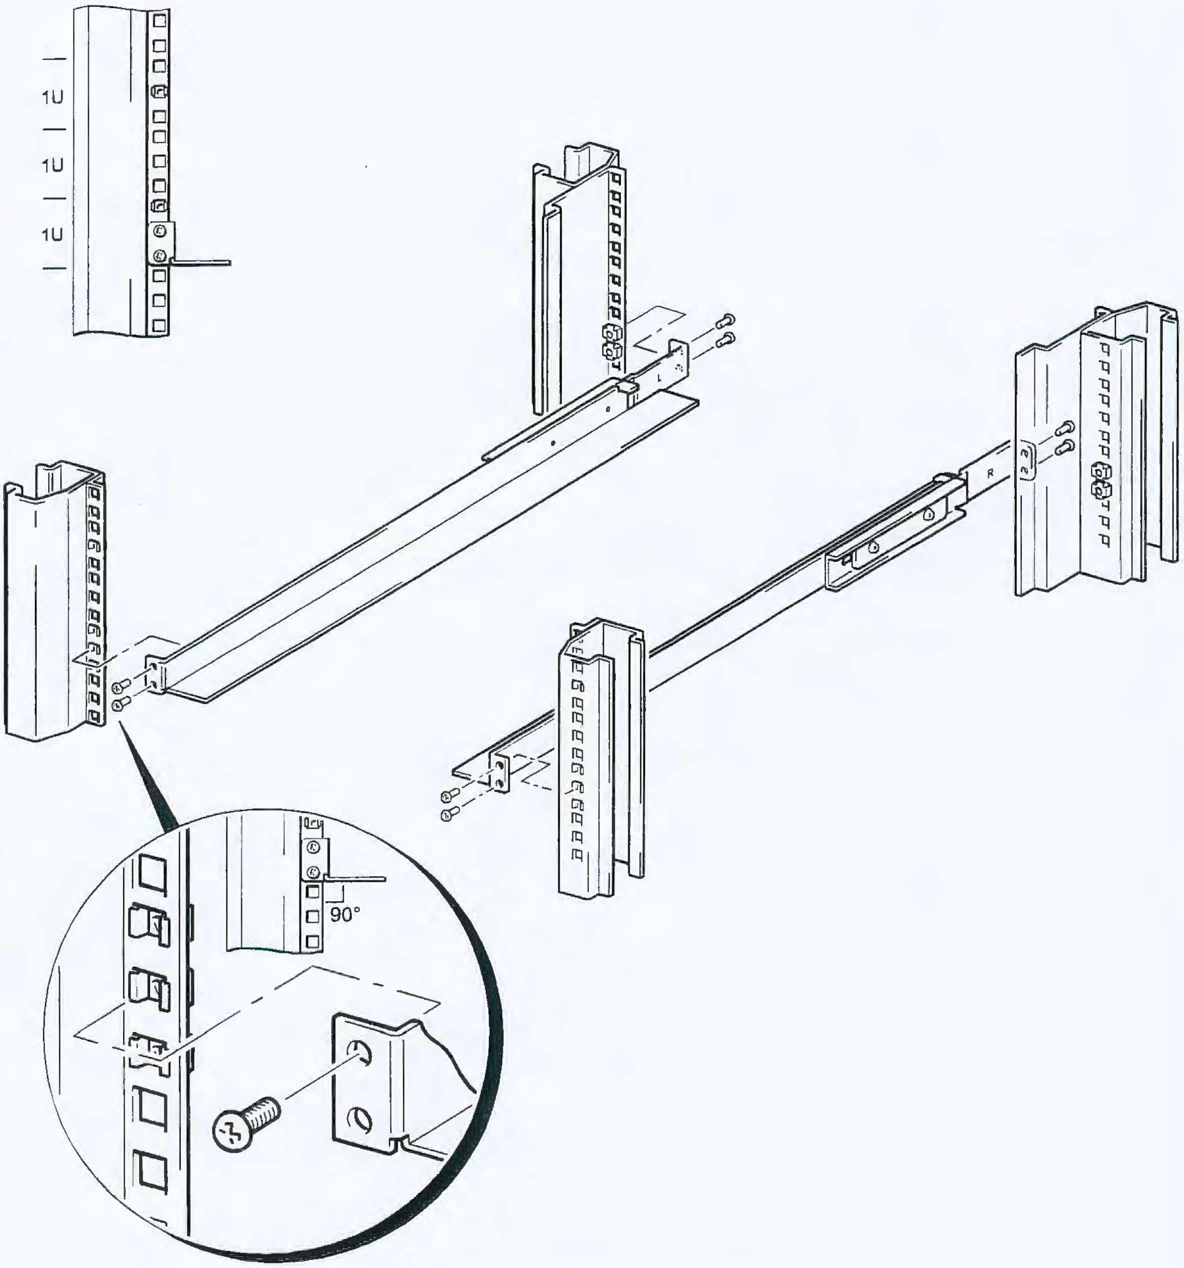

Figure 3b: Install support rails in non-HP racks

18.557
Paula

Step 3b: Install support rails in non-HP racks

1. Refer to Figure 3b. Position a support rail so that the flange at the rear of the rail locates over the rear rack column and the two holes in the rail flange align with the two cage nuts. It is much easier if this is done by two persons.
2. Insert the M6 screws. Partially finger-tighten the screws only at this stage.
3. Position the rail so that the flange at the front of the rail locates over the front rack column and the two holes in the rail flange align with the bottom two cage nuts.
4. Use a Pozidriv® or large Philips screwdriver to tighten the M6 screws. Ensure the rail is secured so that the support flange is at 90 degrees to the vertical rack column.
5. Tighten the screws on the rear rack columns.
6. Repeat paragraphs 1 through 5 for the opposite side of the rack.

18.336
Paula

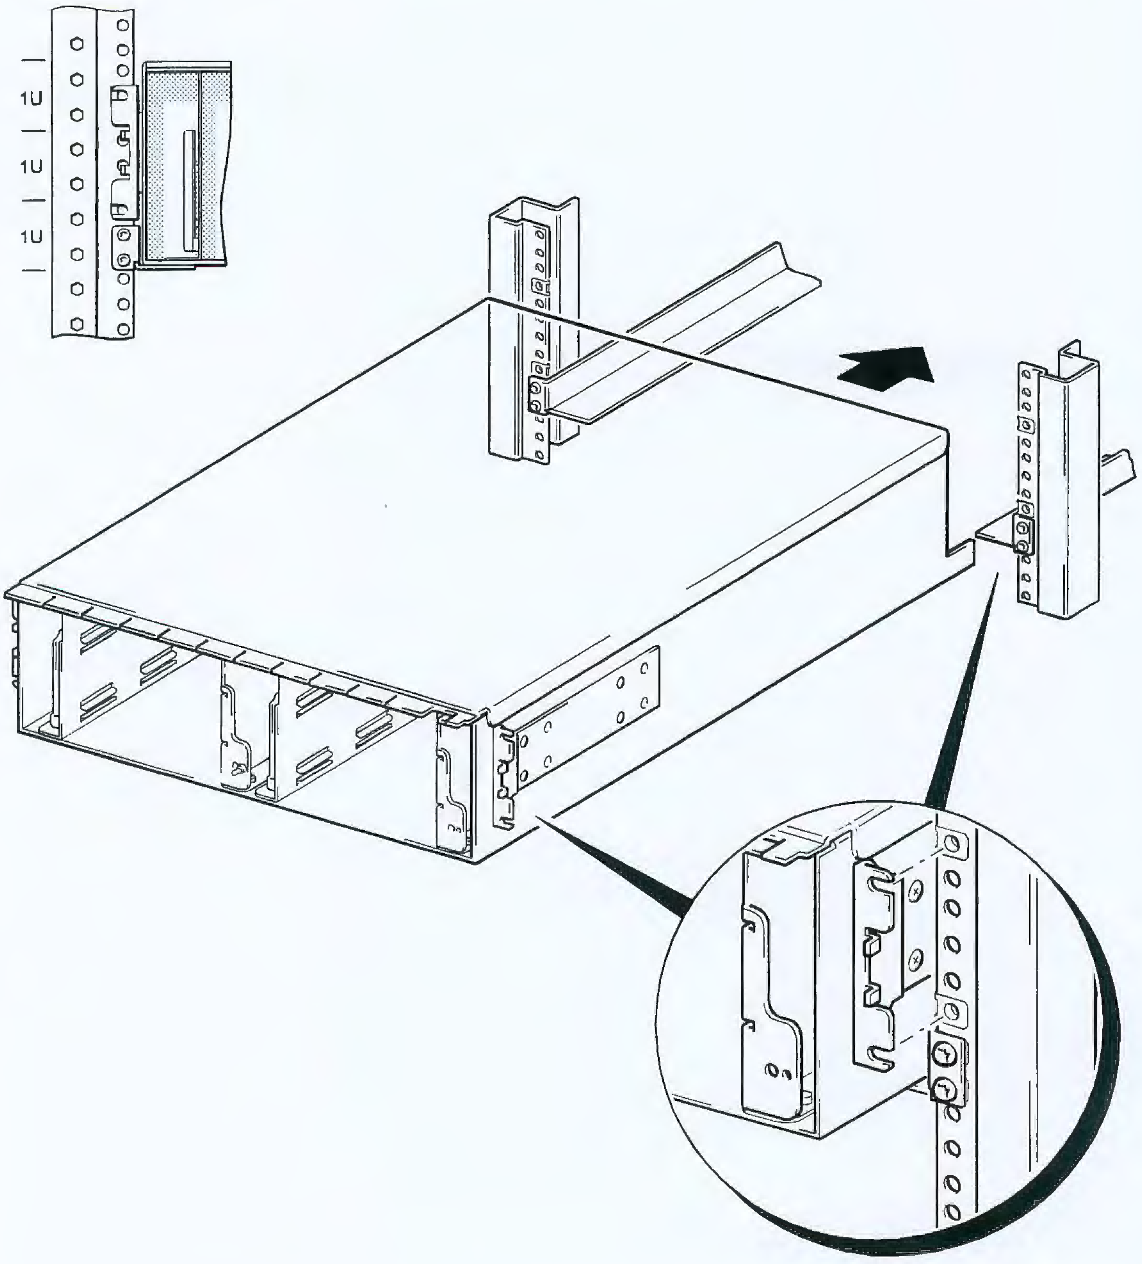

Figure 4a: Install the tape array enclosure in HP racks

18.335
Paula

Step 4a: Install the tape array enclosure in HP racks

Warning The HP Surestore Tape Array 5300 is heavy (13 kg or 28.7 lbs). It is recommended that two people should install it. Prior to installing your tape array in a rack, extend the rack's anti-tip stabilizer and utilize any anti-tip ballast.

To minimize the weight and to prevent personal injury, do not install tape drive modules into the HP Surestore Tape Array 5300 until the tape array is installed in the rack.

1. With an assistant, lift the tape array into the rack and rest the back end of the array on the front of the support rails.
2. Push the tape array back into the rack along the rails until the extended base plate sections of the unit slot under the tabs located on the rails. Push fully in until the flanges of the side mounting brackets are flush with the front rack columns.

   The two mounting slots on the side mounting bracket's flange should now line up with the upper and lower clip nuts located on the front column, as shown in Figure 4a.

18.334
Paula

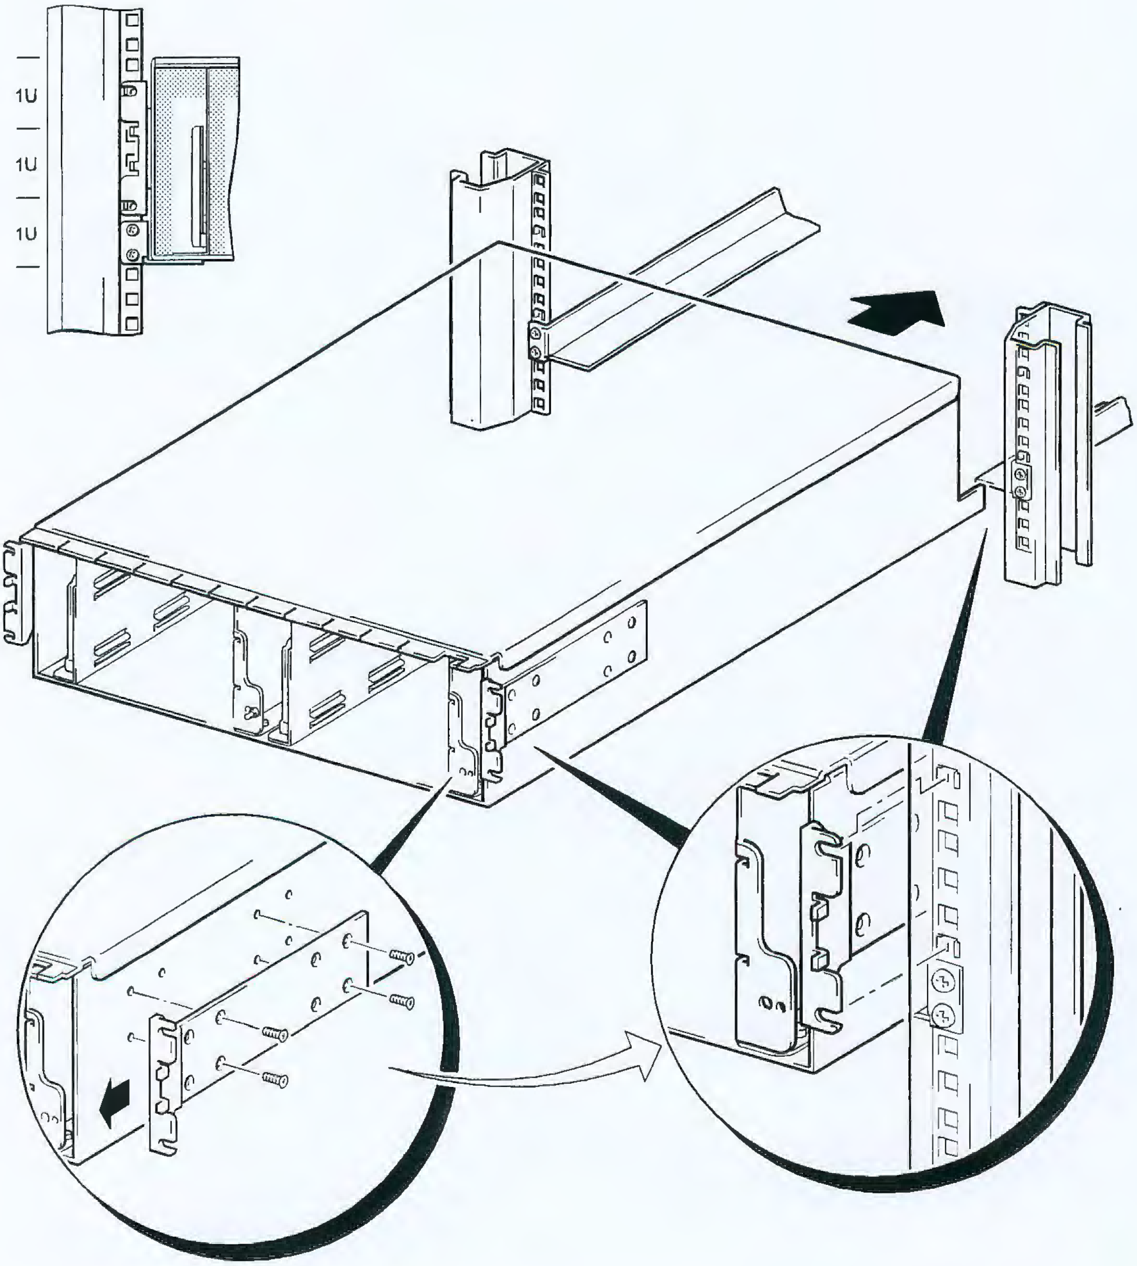

Figure 4b: Install the tape array enclosure in non-HP racks

18 333
Paula

Step 4b: Install the tape array enclosure in non-HP racks

Warning The HP Surestore Tape Array 5300 is heavy (13 kg or 28.7 lbs). It is recommended that two people should install it. Prior to installing your tape array in a rack, extend the rack's anti-tip stabilizer and utilize any anti-tip ballast.

To minimize the weight and to prevent personal injury, do not install tape drive modules into the HP Surestore Tape Array 5300 until the tape array is installed in the rack.

Adjusting the position of the side mounting brackets

When installing into some non-HP racks, for example Compaq, the position of the side-mounting brackets have to be adjusted. This is because the front columns of these racks are positioned further back from the front face of the rack.

1. Remove the Philips head screws on the side-mounting brackets on the side of the tape array.
2. Slide the mounting bracket forward to align the second set of holes in the mounting brackets with the holes in the tape array's chassis.
3. Secure with screws.

Installing the HP Surestore Tape Array 5300

1. With an assistant, lift the tape array into the rack and rest the back end of the array on the front of the support rails.
2. Push the tape array back into the rack along the rails until the extended base plate sections of the unit slot under the tabs located on the rails. Push fully in until the flanges of the side mounting brackets are flush with the front rack columns.

   The two mounting slots on the side mounting bracket's flange should now line up with the upper and lower cage nuts located on the front column, as shown in Figure 4b.

18332
Paula

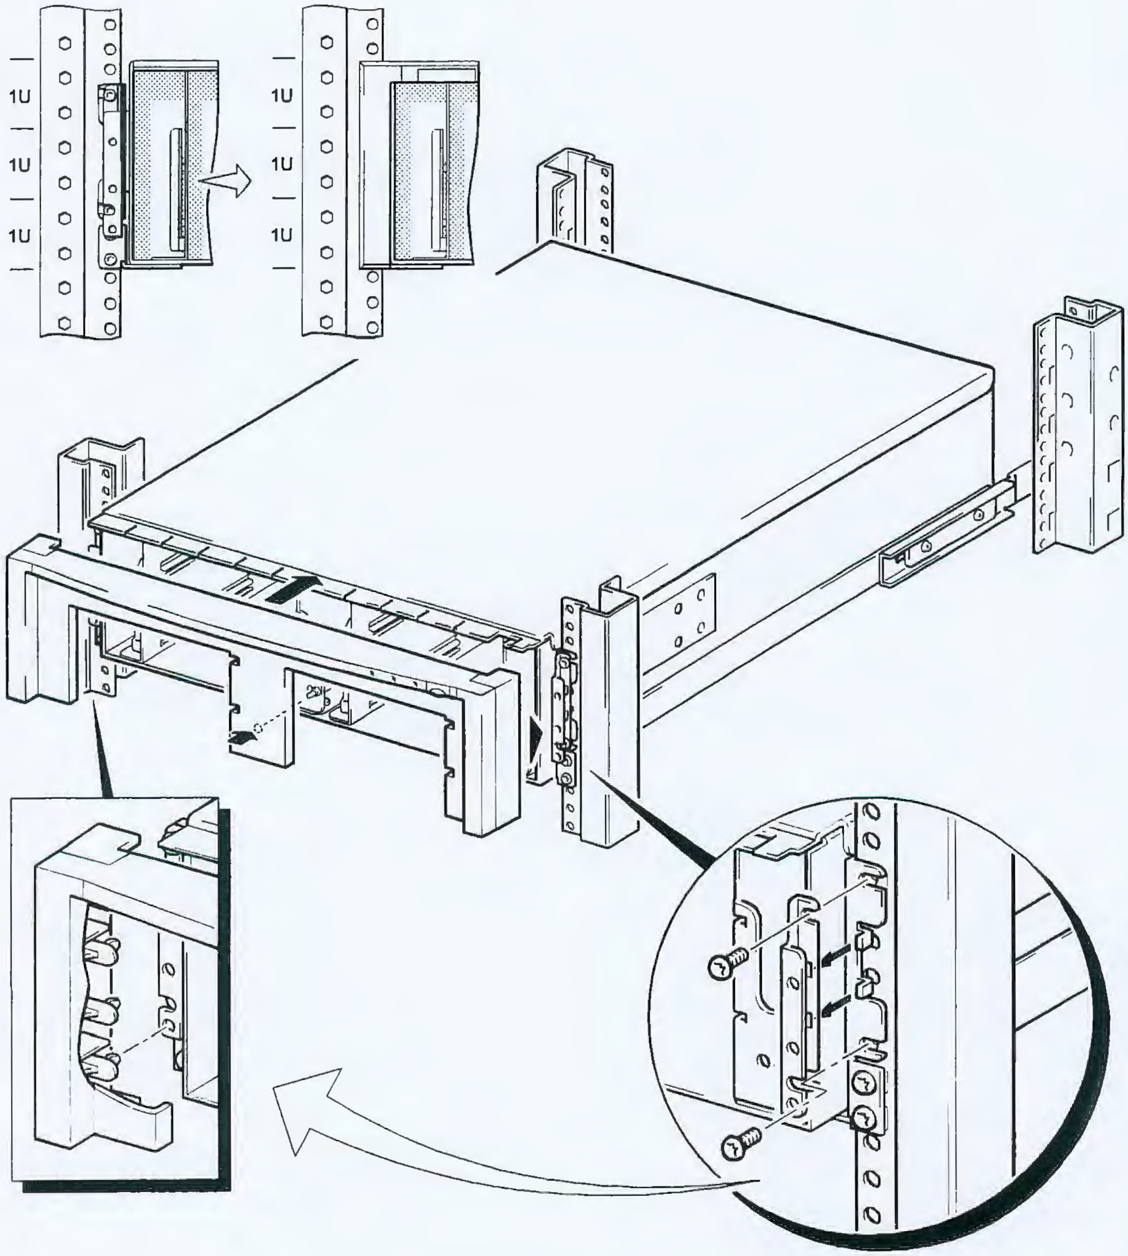

Figure 5a: Install the front bezel on HP racks

Step 5a: Install the front bezel on HP racks

1 Locate the two bezel brackets. The bezel bracket to be fitted to the right hand side of the rack, when the rack is viewed from the front, is stamped with the letter "R". Likewise, the bezel bracket to be fitted to the left hand side of the rack, when the rack is viewed from the front, is stamped with the letter "L".

2 Turn a bezel bracket so that the stamped letter is correctly orientated and is facing the front, and the indents face outward towards the side of the rack as shown on Figure 5a.

3 Attach the bracket to the flange of one of the side mounting brackets.

The two rectangular holes on the bezel bracket hang on the two teeth of the side mounting brackets, as shown in Figure 5a.

4 Insert two M5 mounting screws through the slotted holes on the rear flange of the bezel bracket, so that the screws pass through the aligned slotted holes in the side mounting bracket and locate into clip nuts installed on the front rack column.

5 Repeat paragraphs 2 through 4 with the other bezel bracket.

Caution Take care not to damage the LEDs on the top right hand side of the enclosure when fitting the front bexel.

6 Align the lower three mushroom snaps on both bezel endcaps with the holes on the bezel brackets (the top mushroom snap on each side is not used). Ensuring that the centre locating lug is also aligned and that the top edge of the bexel slides over the top front edge of the enclosure, snap the bexel into place. (No screws are required.)

18.330
Paula

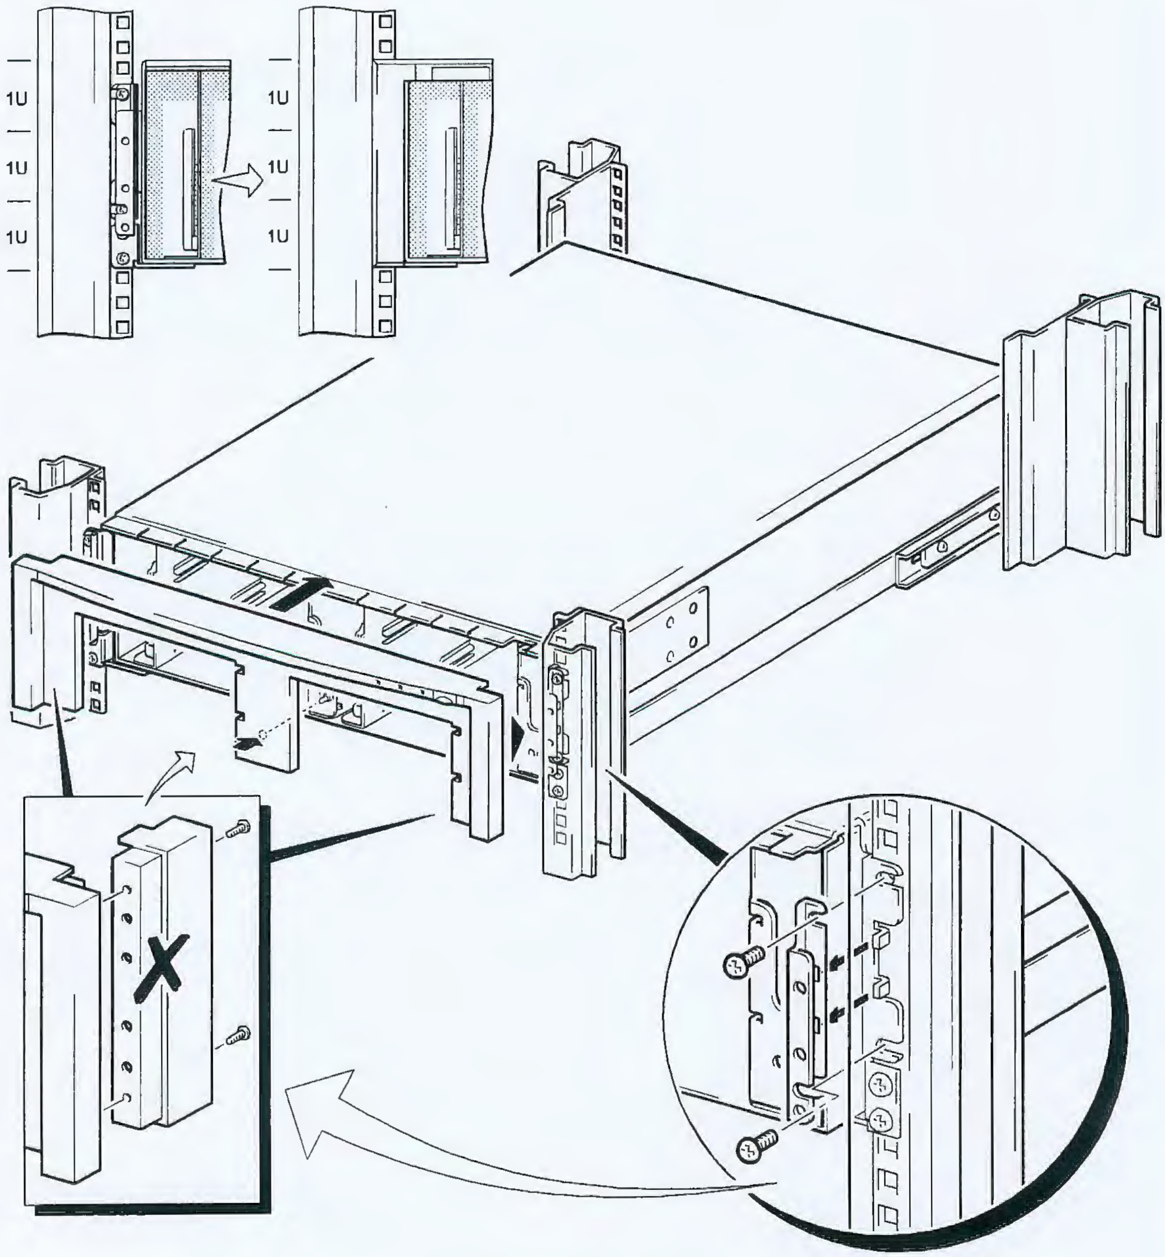

Figure 5b: Install the front bezel on non-HP racks

18 32a
Paula

Step 5b: Install the front bezel on non-HP racks

1 Locate the two bezel brackets. The bezel bracket to be fitted to the right hand side of the rack, when the rack is viewed from the front, is stamped with the letter "R". Likewise, the bezel bracket to be fitted to the left hand side of the rack, when the rack is viewed from the front, is stamped with the letter "L".

2 Turn a bezel bracket so that the stamped letter is correctly orientated and is facing the front, and the indents face outward towards the side of the rack as shown on Figure 5b.

3 Attach the bracket to the flange of one of the side mounting brackets.

The two rectangular holes on the bezel bracket hang on the two teeth of the side mounting brackets, as shown in Figure 5b.

4 Insert two M6 mounting screws through the slotted holes on the rear flange of the bezel bracket, so that the screws pass through the aligned slotted holes in the side mounting bracket and locate into cage nuts installed on the front rack column.

5 Repeat paragraphs 2 through 4 with the other bezel bracket.

6 Remove the bezel endcaps by removing two screws and unsnapping them from the bezel.

Caution Take care not to damage the LEDs on the top right hand side of the enclosure when fitting the front bexel.

7 Align the lower three mushroom snaps on both ends of the bezel with the holes on the bezel brackets (the top mushroom snap on each side is not used). Ensuring that the centre locating lug is also aligned and that the top edge of the bexel slides over the top front edge of the enclosure, snap the bexel into place. (No screws are required.)

18.328
Paula

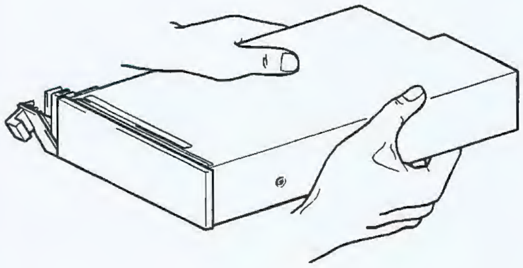

Figure 6a: Handling a half-height tape drive module

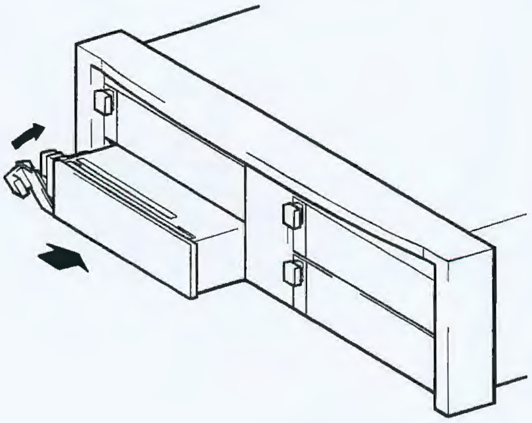

Figure 6b: Installing a half-height tape drive module

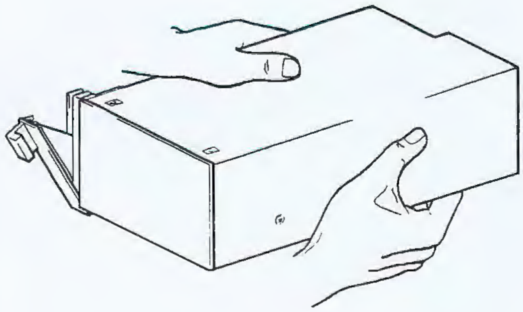

Figure 6c: Handling a full-height tape drive module

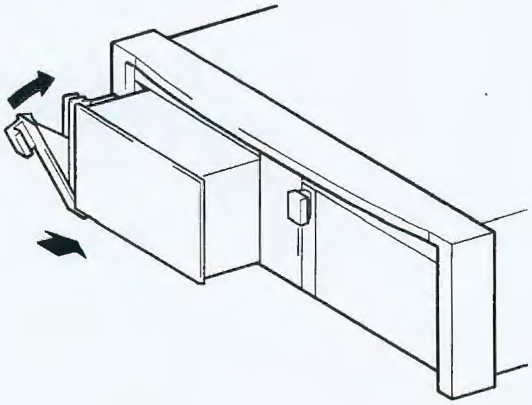

Figure 6d: Installing a full-height tape drive module

18.327
Paula

Step 6: Install the tape drive module

The tape drive modules for the HP Surestore Tape Array 5300 are offline hot-swappable. Tape modules can be installed into bays before or after the tape array is powered up. The HP Surestore Tape Array 5300 can accept up to four half-height tape drives or two full-height tape drives.

Installing the tape drive module

Warning To minimize the weight and prevent personal injury you should install the HP Surestore Tape Array 5300 into the rack cabinet (as described on page 15 or page 17) before you install the tape drive modules.

1 Before installing a tape drive module, make sure that the extractor lever is in the fully open position (pulled out).

Caution When handling a tape drive module, hold the module by the side carrier framework as shown on Figures 6a and 6c. To avoid possible damage, do not hold the body of the module and do not push or pull on the tape drive module's front panel.

2 Slide the tape drive module into one of the open bays, ensuring that the guide rails on both sides of the module locate into the slots on the partition wall as shown on Figures 6b and 6d. Carefully but firmly ensure the tape drive module is pushed fully in so that the rear connectors are securely seated. The extractor level will close slightly.

> **Note** Correct alignment of SCSI modules in the tape array is essential to ensure a secured connection is made between the interface of the module and the interface of the array.

3 Push on the extractor lever to complete the insertion and to lock the tape drive module in place.

Caution It is very important to ensure blanking plates are fitted if less than the full compliment of modules are installed. Failure to do so may affect tape drive performance.

4 If there are less than four half-height (or two full-height) modules installed, insert a blanking plate into the empty bays. This is essential to ensure adequate airflow is maintained.

18.320
Paula

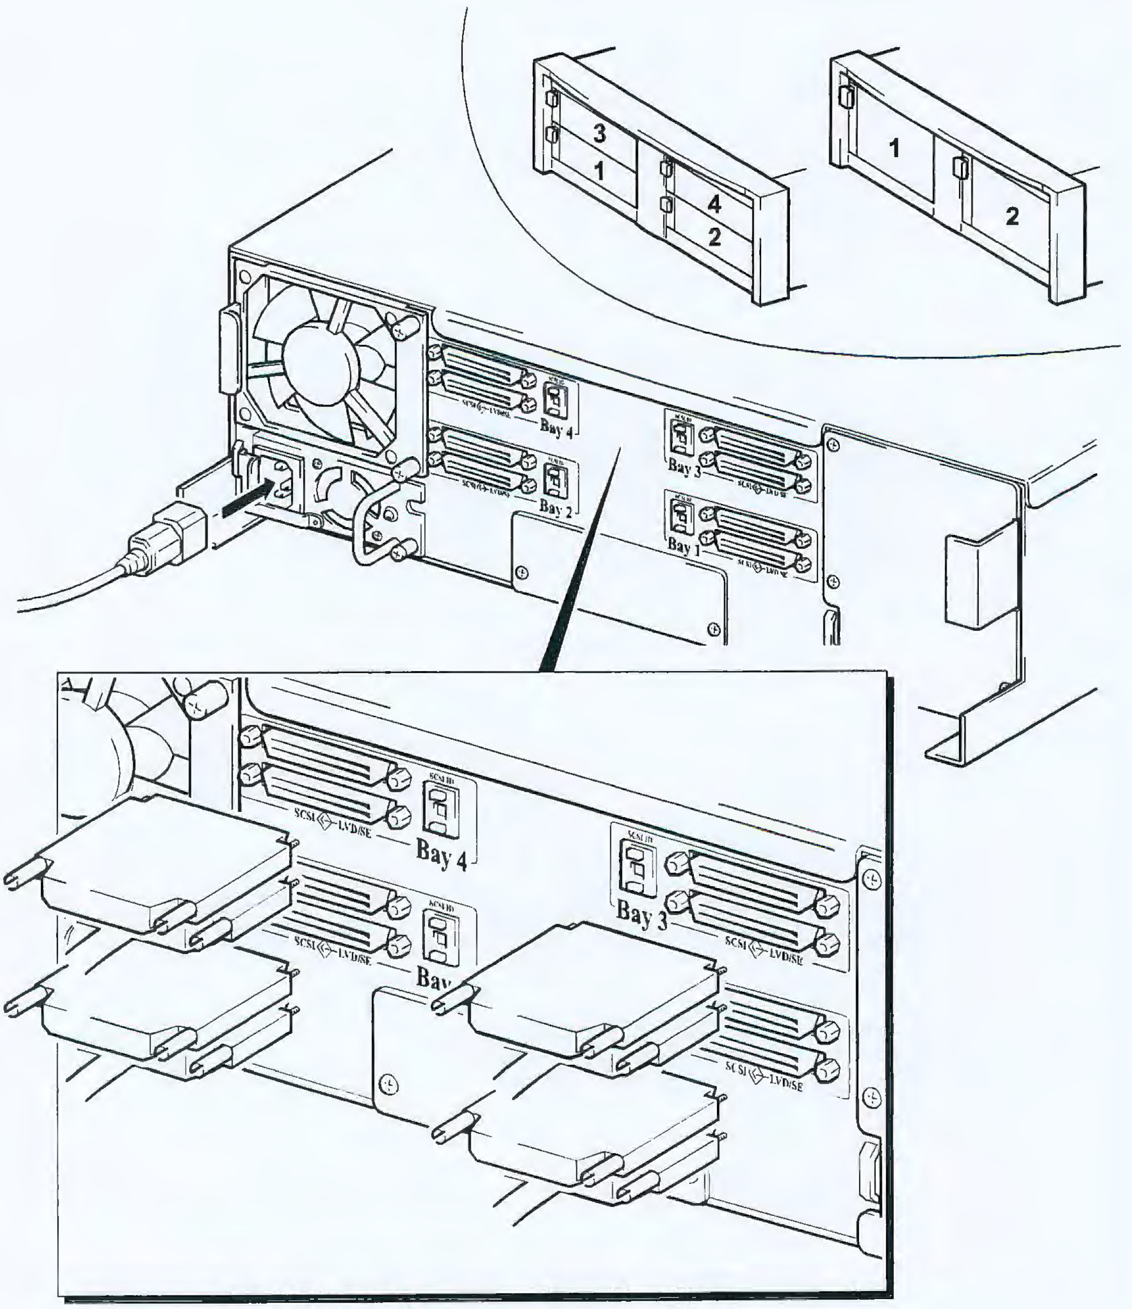

Figure 7: Power, SCSI, and terminator connections

18.325 Paula

Step 7: Connect the SCSI cables, terminator and power cord

1 There are two SCSI connectors on the rear panel of each drive module bay (marked bays 1 to 4). The layout of the four pairs of connectors align with the position of the four half-height tape drive modules in the HP Surestore Tape Array 5300 when viewed from the back. For example, the SCSI connectors to the bottom right, when viewed from the back, (bay 1) correspond to the tape drive module installed in the bottom left bay, when viewed from the front. Where full-height drives are installed, only the lower SCSI connectors (bays 1 and 2) should be used.

The standard configuration is a direct one-to-one SCSI connection between a tape drive module and a host server, as shown in Figure 7.

Connect one end of an appropriate LVD/SE SCSI cable (of the correct length) to one of the SCSI connectors for the bay selected on the rear panel of the HP Surestore Tape Array 5300.

SCSI cables must be ordered separately. To select a cable appropriate for your application, refer to your local HP sales office or view the cable options listed on the HP support web site http://www.hp.com/support/tapearray.

2 Connect the other end of the SCSI cable to the host server.

3 Install an appropriate SCSI LVD/SE terminator to the other SCSI connector for the bay selected on the rear panel of the tape array.

SCSI terminators must be ordered separately. To select a terminator appropriate for your application, refer to your local HP sales office or view the terminator options listed on the HP support web site http://www.hp.com/support/tapearray.

4 Repeat paragraphs 1 through 3 for the remaining tape drive modules.

If other devices are installed on the same SCSI bus, the last device on the bus should be terminated. A tape module should not be connected to the same bus as disk devices or any other non-LVD SCSI tape peripherals.

Note It is also possible to daisy-chain two drive modules together. Contact your local HP sales office, for information on how to order appropriate SCSI cables and terminators. For more information, see "Advice on daisy-chaining" on page 31.

5 Install the power cord to the socket on the power supply unit (lower left side of the rear panel).

18.324
Paula

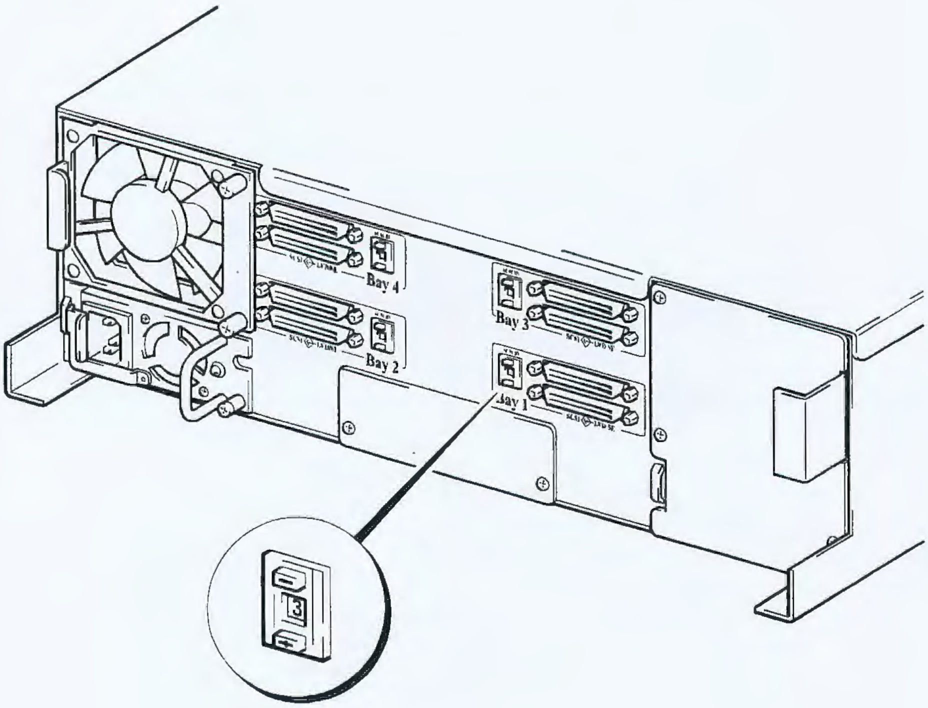

Figure 8: Rear panel view

18.323
Paulo

Step 8: Set the SCSI ID

The rear panel consists of a removable fan, a removable power supply, a power connector, four SCSI In connectors, four SCSI Out connectors, and four SCSI ID switches. See Figure 8.

To set the drive's SCSI ID you need to set the appropriate switch on the panel:

1. Decide which SCSI ID to use and determine whether you need to change the SCSI ID from the default. SCSI ID 7 is normally reserved for the SCSI controller.

Note Each device on the same SCSI bus should have a unique SCSI ID. When daisy-chaining two tape drive modules together or installing other devices onto the bus, ensure that a different number is chosen for each device on the same bus.

It is not necessary to have different numbers where the tape drive modules are configured in direct one-to-one connection with the host servers and no other devices are installed on the bus.

You can install HP Library and Tape Tools from the HP Surestore Tape CD-ROM supplied with your tape drive modules or from our world wide web site at http://www.hp.com/support/tapetools. Use HP Library and Tape Tools and run "Install Check" to check your computer's current SCSI configuration. (This is only available on Windows NT4, Windows 2000, and Novell NetWare operating systems.)

2. Change the tape array's four SCSI ID switches, if necessary.

Use a ball-point pen or similar pointed object to press the indented + or - buttons above and below the SCSI ID number until the required value is displayed.

Guidelines for SCSI Tape Drive Modules

- Each tape drive module that shares the same bus must be assigned a unique SCSI ID before the unit is powered up.
- The last device on the bus must be terminated using a low voltage differential (LVD) terminator.
- A diagnostic application, such as HP Library and Tape Tools, can be used to help troubleshoot detection issues of SCSI modules.

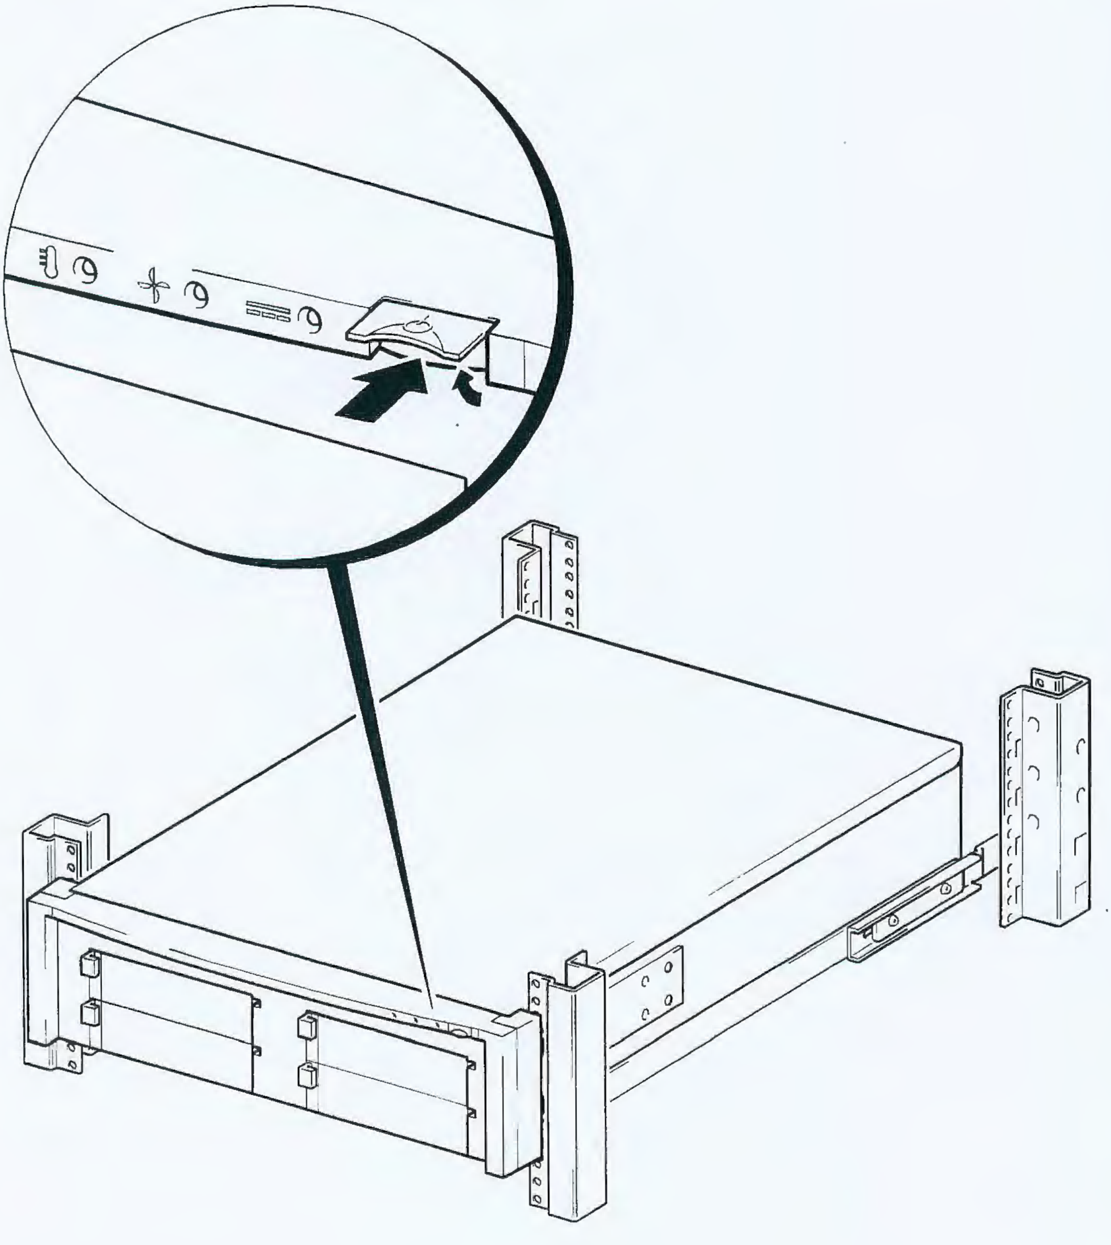

Figure 9: Front bezel with LED indicators

18 321
Laula

Your HP Surestore Tape Array 5300

Switching on the unit

The HP Surestore Tape Array 5300 main power switch is located on the front bezel. Lift the hinged flap and depress the switch to power on the unit.

Front panel LEDs

There are three LEDs on the front bezel. They indicate the temperature, fan and power supply status of the tape array. These LEDs are shown in Figure 9 and are described in the table below:

Indicator	LED color - solid green	LED color - flashing red
Temperature LED	Temperature within safe operating limit for enclosure.	Temperature above safe operating limit for enclosure.
Enclosure fan	Fan present and operating correctly.	Enclosure fan has failed.
Power supply unit	PSU is present and has good output.	PSU present, but its output is not functioning correctly. If the PSU has completely failed, all LEDs will be off.

Where the high availability option kit is installed as described on page 35, so that second fan and power supply unit is added to the enclosure, then the LEDs functionality is as described below:

Indicator	LED color - solid green	LED color - flashing red
Enclosure fans	Both fans present and operating correctly.	One or both enclosure fans have failed.
Power supply units	Both PSUs present and have good output.	Both PSUs present, but either the output at one or both PSUs is not functioning correctly, or one PSU has completely failed.

Rear panel LEDs

The power supply unit has an integral green LED, visible from the rear panel. When lit, the LED indicates that power is on.

Note Each tape drive module also has LEDs on the front panel to monitor the functions of the tape drive itself. Refer to the User's Guide shipped with the tape drive module to understand the status of these LEDs. (The LEDs on the front bezel of the HP Surestore Tape Array 5300 apply to the temperature, fan, and power functions of the enclosure only.)

18.320
Paula

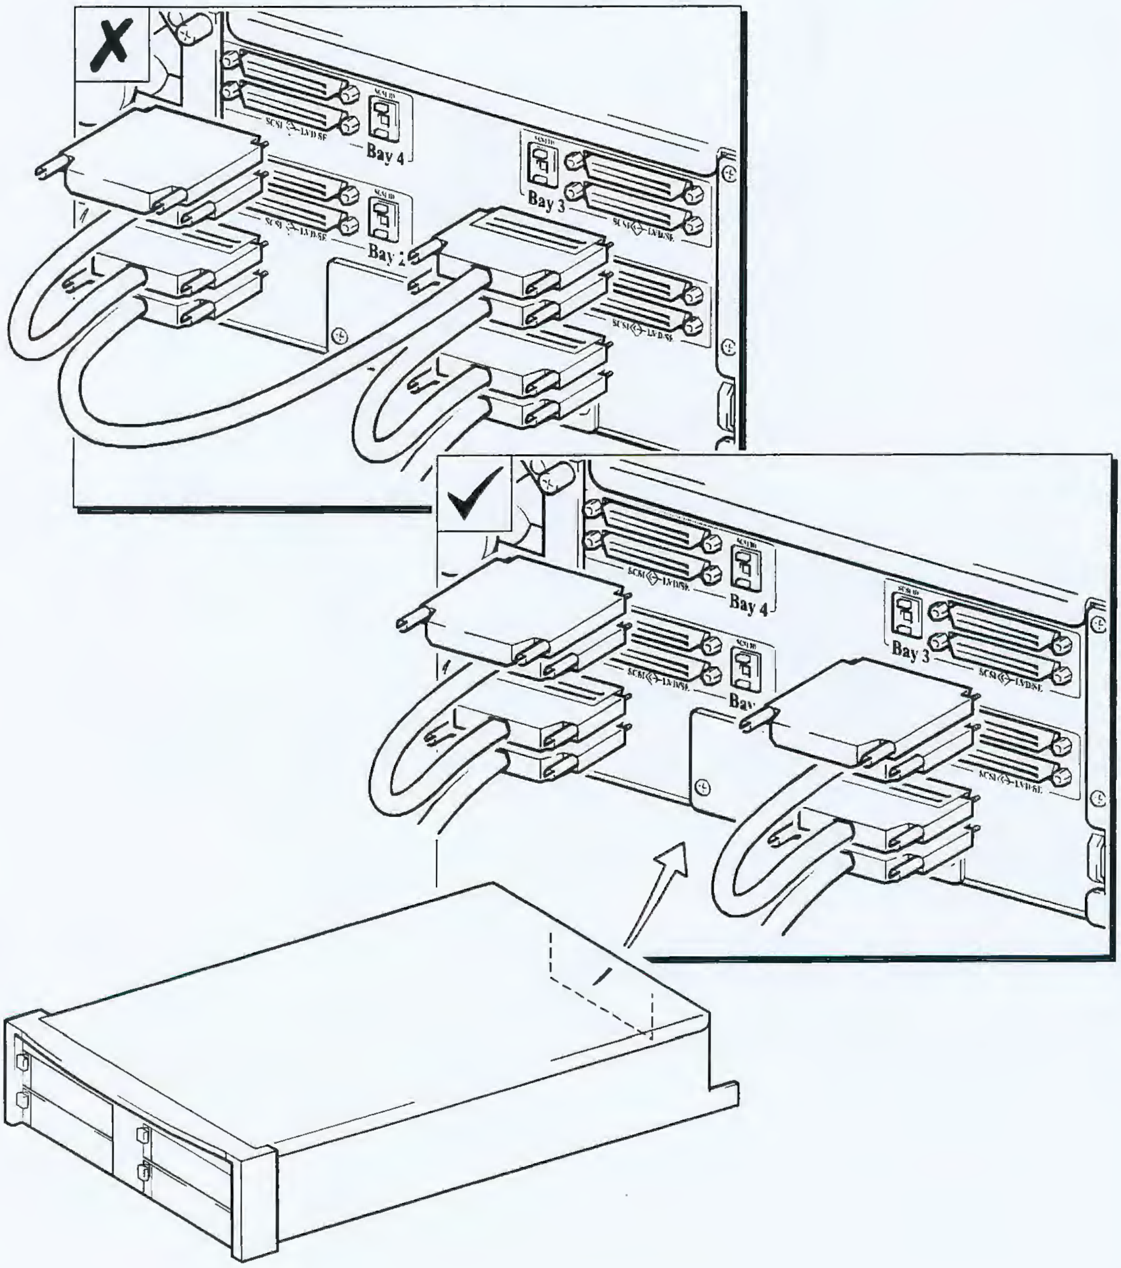

Figure 10: Daisy-chaining two drive modules

18.319
Paula

Advice on daisy-chaining

HP recommends that tape drive modules, as supplied for the HP Surestore Tape Array 5300, should not be placed on a SCSI bus with more than one other module ("daisy-chained"), as illustrated in Figure 10.

Daisy-chaining more than two devices together will not usually result in the drives failing but may degrade their individual performance with respect to transfer rate. This is due to the extra overhead of traffic on the SCSI bus. If more than two drives are run from a single server or workstation, additional SCSI adapters should be used for these drives.

To daisy-chain two tape drive modules:

1. Connect one end of an appropriate LVD/SE SCSI cable to one of the rear panel SCSI connectors of the first module in the daisy-chain.

   SCSI cables must be ordered separately. To select a cable appropriate for your application, refer to your local HP sales office or view the cable options listed on the HP support web site http://www.hp.com/support/tapearray.

2. Connect the other end of the cable to the host server.

3. Take an appropriate short length SCSI cable:
   - Connect one end to the other SCSI connector of the first module.
   - Connect the other end to one of the SCSI connectors of the second module.

4. Install an appropriate SCSI LVD/SE terminator to the other SCSI connector of the second module.

   SCSI terminators must be ordered separately. To select a terminator appropriate for your application, refer to your local HP sales office or view the terminator options listed on the HP Support web site http://www.hp.com/support/tapearray.

5. Ensure that the SCSI IDs are unique on each of the two daisy-chained modules.

18.318
Paula

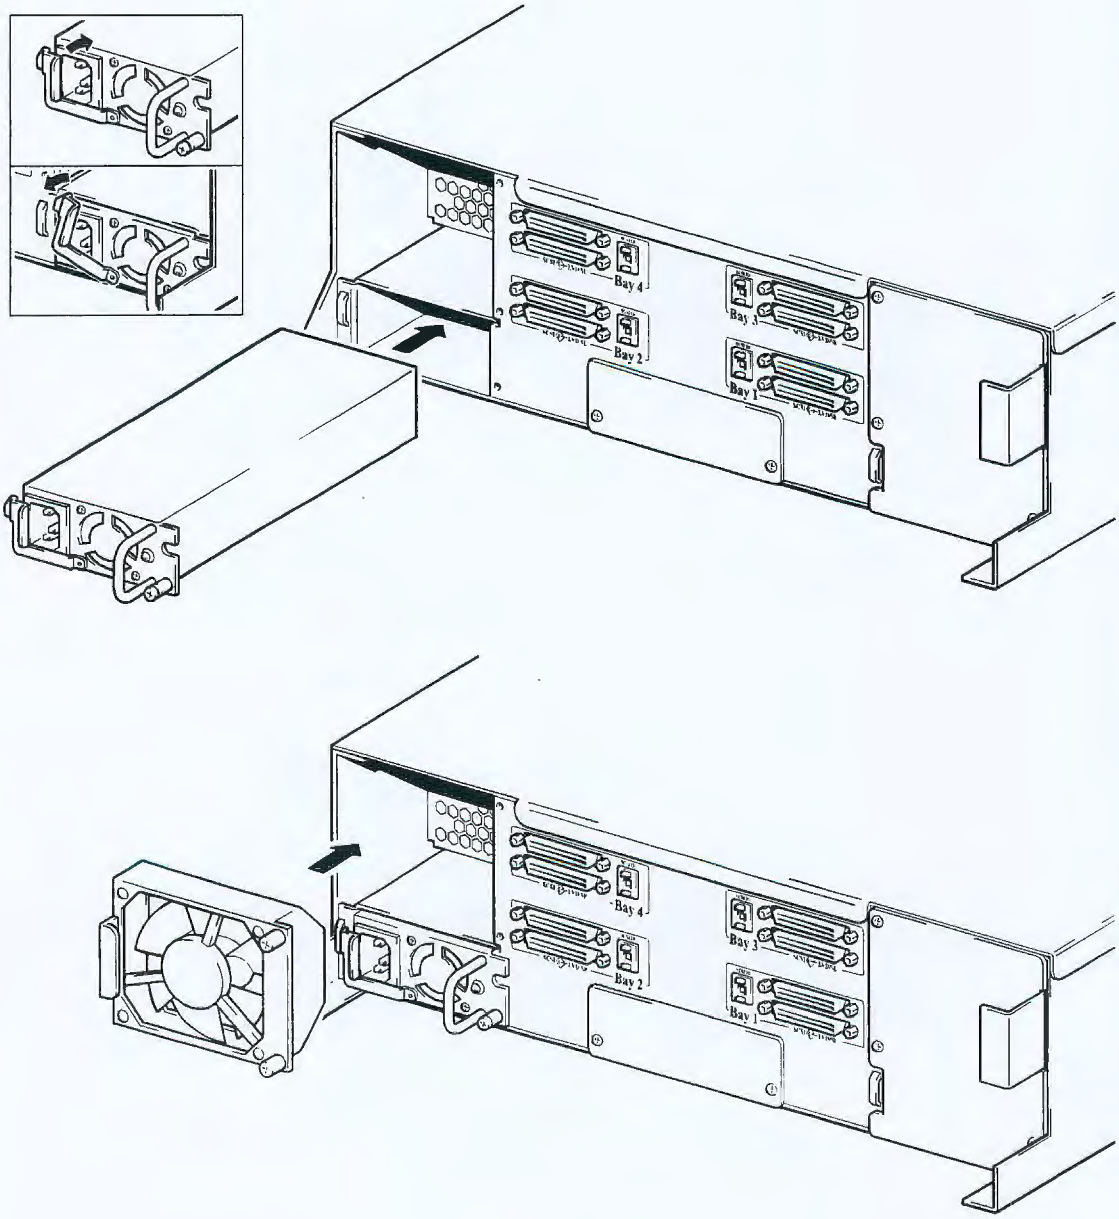

Figure 11: Replace fans and power supply

18.317
Paula

Replacing a tape drive module, fan, or power supply

The tape drive modules are *offline* hot-swappable[1]. The fan and power supply are installed at the factory. Tape modules can be installed into bays before or after the HP Surestore Tape Array 5300 is powered up.

Replacing a tape drive

1 Before installing a tape drive module, make sure that the extractor lever is in the fully open position (pulled out).

2 Slide the tape drive module into one of the open bays until the rear connectors are firmly seated. The extractor level will close slightly.

3 Push on the extractor lever to complete the insertion and to lock the tape drive in place.

4 If there are less than four half-height (or two full-height) modules installed, insert a blanking plate into the empty bays.

Replacing a fan

1 Remove the old fan by unscrewing the two thumb screws. Pull the fan assembly out of its bay.

2 Align the replacement fan unit as shown in Figure 11. Slide the fan unit into the opening.

3 Tighten both thumb screws to complete the insertion and to lock the fan unit in place.

Replacing a power supply

1 Unplug the power cord.

2 Unscrew the single securing screw using a correct-size Philips screwdriver. Rotate the locking mechanism clockwise, sufficient to clear the retaining lug, and pull the power supply unit out of its bay, using the handle supplied.

3 Align the new power supply unit so that the label is uppermost. Insert the power supply unit, which must not have a power cable connected, and slide it into the opening, as shown in Figure 11. Rotate the locking mechanism clockwise, sufficient to clear the retaining lug, and push the unit fully home.

4 Rotate the locking mechanism counter-clockwise to engage with the retaining lug and lock the power supply unit in place.

5 Tighten the single securing screw.

6 Plug in the power cord.

[1] Offline hotswap is defined as the ability to plug and unplug drive modules with the power still connected, but not while data transfers are taking place on other drives on the same SCSI bus.

18.316
Paula

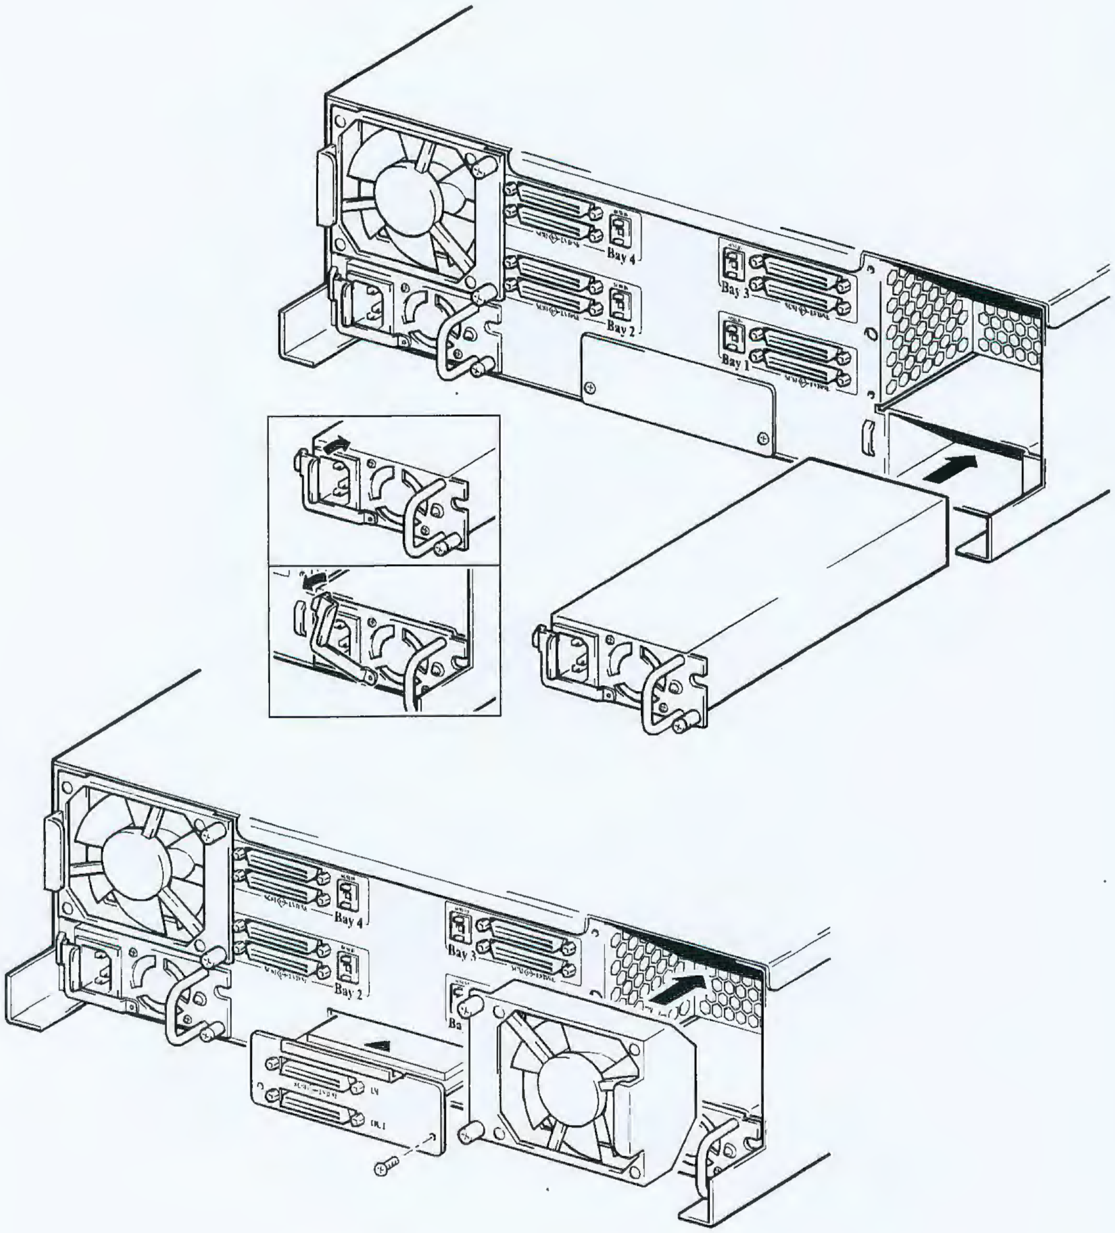

Figure 12: Install upgrade kits

18.315
Paula

Installing high availability/management card upgrade kits

The HP Surestore Tape Array 5300 is shipped with a single fan and power supply unit. An upgrade kit is available to fit an additional fan and power supply unit. This provides high-availability, hot-swap functionality whereby should one unit suffer component failure, the other unit will maintain correct operation. A further upgrade kit allows the addition of a SCSI management card. This gives remote management and diagnostic functionality and provides reporting outputs, via a SCSI connection, for over temperature, fan failure and power supply unit failure.

To power down the HP Tape Array

Ensure that all tape drive devices have finished any backup or restore activities and that no data transfers are taking place between the HP Surestore Tape Array 5300 and its Host device(s).

To install an update kit, first power down the tape array as follows:

1. Lift the power button safety cover on the front bezel and depress the power button.
2. Remove the main cable connection from the power supply on the lower left hand side of the rear panel. This is essential - the tape array's electronics always have some standby power applied unless the power cord is detached. Check that the HP Surestore Tape Array 5300 is fully powered down - all LEDs on the front panel should be unlit.

To install the management card upgrade kit

1. Locate the slot for the management card. The slot is positioned centrally on the HP Surestore Tape Array 5300's rear panel. Remove the blanking plate by undoing the two screws.
2. Observing normal anti-static precautions, remove the management card from its protective packaging and carefully and firmly insert the card into the back panel slot. Ensure the sides of the card engage with the guide rails provided.
3. Ensure the rear connectors are firmly seated, then secure the card into place with the two screws.
4. The management card is now ready for use in conjunction with HP Library and Tape Tools or appropriate Host SAF-TE management software. Refer to you SAF-TE software documentation for details of how to install and operate SAF-TE management in a Host system connected to the HP Surestore Tape Array 5300.

To install the high availability upgrade kit

1. Remove the blanking plate on the right hand side of the tape array's rear panel and install the second fan and power supply unit into their respective slots, as detailed on page 33.
2. Connect mains power cables to both power supply units, and restart the tape array by pressing the power on/off button ON at the front bezel. Check that all three LEDs are GREEN. If all three LEDs are not shown Green, refer to the troubleshooting section on page 36.

Troubleshooting the HP Surestore Tape Array 5300

The first step in problem-solving is establishing whether the problem lies with the HP Surestore Tape Array 5300, its connections to the host computer, or with the removable tape drive modules.

Warning The power button does not power off the unit completely although the power supply and fan(s) are off. The HP Surestore Tape Array 5300 electronics always have some standby power applied unless the power cord(s) is detached.

If the problems relate to poor performance, refer to your backup application manual.

If none of this advice helps you solve the problem, call for technical support. The HP technical support numbers are available on the web site: //www.hp.com/support. They are also listed on the HP Surestore CD-ROM supplied with the tape drive modules.

hysical Installation

Possible reason	**Recommended action**
The side support rails do not fit the cabinet.	• Adjust the length of the rails to fit the cabinet. • Verify the compatibility of the rack used. • Check the template for the correct assembly procedure.
The fixing screws used to attach the side support rails to the cabinet are of the wrong type.	Check the screws are the recommended type for the cabinet. There are two types supplied: M5 screws that locate into clip nuts for HP racks (round profile holes in the cabinet columns) and M6 screws that locate into cage nuts for non-HP racks (square profile holes in the cabinet columns).
The enclosure cannot be pushed fully in.	• Verify the tabs on the rails are not damaged. • Check the template for the correct assembly procedure. • Check the rails are at 90° to the cabinet columns and are fully-tightened up.
The slots in the side mounting brackets do not align with the holes in the cabinet columns.	Check the fittings on the cabinet column with the template for the correct position.
The front bexel does not snap-fit onto the bexel brackets.	• Check for correct orientation of the bezel-mounting bracket on the main enclosures. • Check that the front bexel's central locating lug is correctly aligned and that the top edge of the bexel fits over the top edge of the enclosure.
The front bexel protrudes too far out of the enclosure.	If installing in a Compaq cabinet, check that the side mounting brackets have been repositioned so that the second set of holes on the brackets align with the holes in the tape array's chassis.

18.313
Paula

Installing the tape drive modules

Possible reason	Recommended action
The tape drive modules do not insert properly.	• Check that the extractor lever is in the fully open position before inserting the module. • Check that the guide rails on both sides of the module are correctly located in the slots in the drive bay's partition walls. • Check the module for damage or for any object preventing insertion. • Check that the drive bay is clear of any obstruction. • Check that the connections on both the module and the bay are not damaged.
The blanking panels do not fit into unpopulated drive bays.	Check that the blanking panels are not damaged.
The SCSI cables or SCSI terminators do not fit onto the SCSI connections on the rear panel of the tape array.	• Check that the cables and/or terminators are the correct type and size. • Check the cables, terminators and rear panel connectors for damage.

The tape drive modules are not responding

Possible reason	Recommended action
The tape drive module(s) is not functioning.	• Check that the power cable(s) is firmly connected. • Check that the power supply is operating. • Ensure that the module is fully-seated into the rear connectors of the tape array.
The tape drive module(s) is not seen by the host controller.	• Check that the SCSI host adapter has been correctly installed in the host system. • Check that all SCSI buses are terminated correctly. • Make sure that all SCSI cables are connected firmly. • Remove the SCSI cables and check for damaged contacts or pins. Replace or reinstall the cables. • Remove and reseat the tape drive modules. • Remove and change the position of the modules in the drive bays.
There is a communication problem on the SCSI bus.	• Check the SCSI cabling. • Refer to the SCSI topic of the tape drive module's User's Guide (on the HP Surestore Tape CD-ROM). • Run HP Library and Tape Tools.

18.312
Paula

Problems with the LEDs

Possible reason	Recommended action
Both the power supply unit's LEDs and the front bexel LEDs do not light.	• Check that the power cable(s) is firmly connected and the power is switched on. • Cycle the power to the tape array. • Check the main fuse. • Replace the power supply unit and call for service if the LEDs still do not light.
The power supply unit's LED is on, but the front bexel LEDs do not light.	• Check that the front bexel is properly assembled. • Check that the power-on button on the front bexel is pushed in to the full extent of its travel.
The LEDs on the tape drive module(s) do not light.	• Check the module is fully seated into the rear of the tape array. • Check the connectors on the tape drive module for signs of physical damage. • Refer to the documentation supplied with the tape drive module for details on the correct operation of the module's front panel LEDs.
The Temperature LED is flashing red.	• Check that the fan(s) is working • Check that blanking panels are in place in any unpopulated drive bays. • Check that the tape array is operating within its environmental specifications. • Check that adjacent devices in the cabinet are not overheating. • Check that there are no obstructions blocking the flow of air at the rear or front of the unit.
The Fan LED is flashing red.	• Check that the fan(s) is properly located. • Check that the fan(s) is rotating. If it is not rotating, remove the fan and try repositioning it in the opposite bay. If it is rotating, remove it and check for obstructions. Clean the fan before replacing.
The Power Supply Unit is flashing red.	• If there is only one PSU, and the LED on the PSU is lit, remove the PSU and try repositioning it in the opposite bay. • If there is a second PSU present, remove the left hand unit and power cycle the enclosure. Swap the PSU within the same unit and re-power cycle the unit.
The tape drive module's LED show a fault condition	• Cycle the power for the module. • Refer to the module's User's Guide.

18.311
Paula

Product Specifications

HP Surestore Tape Array 5300 specifications

SCSI Buses	4 individual buses (daisy-chain capable); LVD (Low Voltage Differential)
SCSI Connectors	Eight 68-pin SCSI connectors
SCSI Bus bandwidth	80 MB/s
Rack space used	3U
Power and fan	Redundant and hot-swappable
SCSI Device types supported	Synchronous or asynchronous
Supported NOS	Refer to: http://www.hp.com/go/connect
Supported software	Refer to: http://www.hp.com/go/connect
LED interface	Temperature, power supply, and fan
Height	3U = 133.35 mm (5.25 inches)
Width of enclosure	444.5 mm (17.5 inches)
Width of enclosure plus bexel	480 mm (18.9 inches)
Depth of enclosure	710.5 mm (28 inches)
Depth of enclosure plus bexel	740.5 mm (29.2 inches)
Weight empty of tape drives but with 1 fan and 1 PSU installed	13 kg (28.7 lbs)
Weight fully-populated with half-height tape drives and with front bexel in place	23kg (50.72 lbs) typical - exact weight will vary with type of tape drive modules installed.
Weight fully-populated with full-height tape drives and with front bexel in place	21 kg (46.31 lbs) typical - exact weight will vary with type of tape drive modules installed.

Operating and non-operating ranges

Temperature	Operating: Non-operating	5°C to 40°C -40°C to 70°C
Humidity	Operating: Non-operating	20% – 80% non-condensing 5% – 95% non-condensing
Vibration	Operating random: Survival random: Power off	0.21 grms, 5-500 Hz 2.09 grms, 5-500 Hz
Altitude	Operating: Non-operating	0m to 3.1 km (0 to 10,000 ft) 0m to 4.6 km (0 to 15,000 ft)

These figures apply to the tape rack enclosure only. Refer to the documentation supplied with the tape drive modules for operating parameters for the modules.

18.310
Paula

Warranty

This warranty gives you specific legal rights. You may also have other rights that vary according to where you are located.

Hewlett-Packard warrants its HP Surestore Tape Array 5300 against defects in materials and workmanship for a period of three years, either from the date of delivery or, where the purchase price includes installation by Hewlett-Packard, from the date of installation. During the warranty period, Hewlett-Packard will provide a next-day response, onsite service. This guarantees that an HP engineer will respond to a service call either on the day of the call or on the following working day and will repair or replace defective products onsite.

Should Hewlett-Packard be unable to repair or replace the product within a reasonable period of time, a refund of the purchase price may be given upon return of the product.

Exclusions

The warranty on your HP Surestore Tape Array 5300 does not apply to defects resulting from:

- Improper or inadequate maintenance by the customer
- Customer-supplied software or interfaces
- Operation outside the environmental specifications for the product
- Use of unsupported media
- Improper site preparation and maintenance

Limitations

Any implied warranty of merchantability or fitness is limited to the 3-year duration of this written warranty. Some states or provinces do not allow limitations on how long an implied warranty lasts, so limitations or exclusions may not apply to you.

Limited Warranty Statement for USA, Canada, Australia, and New Zealand

HP warrants to you, the end-user customer, that HP hardware, accessories, and supplies will be free from defects in materials and workmanship after the date of purchase for the period specified above. If HP receives notice of such defects during the warranty period, HP will, at its option, either repair or replace products that prove to be defective. Replacement products may be either new or like-new.

HP warrants to you that HP software will not fail to execute its programming instructions after the date of purchase, for the period specified above, due to defects in material and workmanship when properly installed and used. If HP receives notice of such defects during the warranty period, HP will replace software media that does not execute its programming instructions due to such defects.

HP does not warrant that the operation of HP products will be uninterrupted or error free. If HP is unable, within a reasonable time, to repair or replace any product to a condition as warranted, you will be entitled to a refund of the purchase price upon prompt return of the product.

HP products may contain remanufactured parts equivalent to new in performance or may have been subject to incidental use.

18.309
Paula

Warranty does not apply to defects resulting from (a) improper or inadequate maintenance or calibration, (b) software, interfacing, parts or supplies not supplied by HP, (c) unauthorized modification or misuse, (d) operation outside of the published environmental specifications for the product, or (e) improper site preparation or maintenance.

HP MAKES NO OTHER EXPRESS WARRANTY OR CONDITION WHETHER WRITTEN OR ORAL. TO THE EXTENT ALLOWED BY LOCAL LAW, ANY IMPLIED WARRANTY OR CONDITION OF MERCHANTABILITY, SATISFACTORY QUALITY OR FITNESS FOR A PARTICULAR PURPOSE IS SET FORTH ABOVE.

Some countries, states, or provinces do not allow limitations on the duration of an implied warranty, so the above limitation or exclusion might not apply to you. This warranty gives you specific legal rights and you might also have other rights that vary from country to country, state or state, or province to province.

TO THE EXTENT ALLOWED BY LOCAL LAW, THE REMEDIES IN THIS WARRANTY STATEMENT ARE YOUR SOLE AND EXCLUSIVE REMEDIES. EXCEPT AS INDICATED ABOVE, IN NO EVENT WILL HP OR ITS SUPPLIERS BE LIABLE FOR LOSS OF DATA OR FOR DIRECT, SPECIAL, INCIDENTAL, CONSEQUENTIAL (INCLUDING LOST PROFIT OR DATA), OR OTHER DAMAGE, WHETHER BASED IN CONTRACT, TORT, OR OTHERWISE.

Some countries, states, or provinces do not allow the exclusion or limitation of incidental or consequential damages, so the above limitation or exclusion may not apply to you.

FOR CONSUMER TRANSACTIONS IN AUSTRALIA AND NEW ZEALAND: THE WARRANTY TERMS CONTAINED IN THIS STATEMENT, EXCEPT TO THE EXTENT LAWFULLY PERMITTED, DO NOT EXCLUDE, RESTRICT, OR MODIFY AND ARE IN ADDITION TO THE MANDATORY STATUTORY RIGHTS APPLICABLE TO THE SALE OF THIS PRODUCT TO YOU.

Information in this document is subject to change without notice.

18.308
Paula

ANEXO UNIDADE DE BACKUP ROBOTIZADO
PARTE R

18.307
Paula

MODELS

HP Ultrium 400 GB Data Cartridge

C7972A, Worldwide

Feature List:

- 200 GB native capacity at 2 1 compression
- Unique red cartridge shell to assist with identification in mixed media vaults
- Superior 'smart grabber mechanism and mechanical interlock to prevent the leader pin from being pulled inside the tape housing (a key weakness of DLT media). Sensors detect proper connection and prevent leader loss that would ruin the tape.
- Simplified tape path to reduce wear and tear
- Lowest media cost per GB at launch of *any* tape technology
- LTO cartridge memory in media to improve access time and provide enhanced media monitoring.
- Improved cartridge robustness and durability

OVERVIEW

Backed by HP's exhaustive media qualification process, the C7972A HP Ultrium 400GB data cartridge supports the launch of the HP StorageWorks Ultrium 460 Tape Drive Building on the class-leadership of the Linear Tape Open standard, HP Ultrium 2 cartridges provide exceptional reliability with best-in-class high performance; the new format minimizes network interruption by being capable of protecting nearly 216 GB/hour of data (88% faster than SDLT 320) and storing up to 400GB on one piece of media (double the capacity of Ultrium 1 tapes).

Customer Benefits Include:

- high-capacity: supports capacities up to 400 GB
- high-performance: with HP StorageWorks Ultrium 460 tape drive delivers up to 60 MB/sec compressed transfer rate
- reliable: combines the best of proven technologies to offer users improvements in capacity and performance.
- open standard: for use in all HP and non-HP LTO Ultrium 2 products
- easy to use: possesses a distinctive red cartridge shell

KEY FEATURES

- Tested to extremes in HP's media laboratory, to a level of specification unique to Hewlett-Packard. Many of the test procedures required for HP Brand qualification (e g load/unioad, shoeshine, drop testing and thermal aging) are not required for the Ultrium logo.
- 200 GB (native), 400 GB (2 1)
- Supports transfer speeds of 30 MB/s (native), 60 MB/s (2 1)
- Unique red cartridge shell
- Compatible with all HP and non-HP LTO Ultrium tape products
- LTO Cartridge Memory
- Superior mechanism, tape path and tape engagement
- Durable cartridge designed to minimize wear and debris (important in automation)

PRODUCT HIGHLIGHTS

Ultrium 2 Tape Technology

Ultrium 2 tape drive technology uses a linear format of 512 tracks, writing eight (8) tracks simultaneously The data is written in a serpentine pattern, the tape reverses direction after each set of eight tracks is written.

Ultrium formats are open standards This means that data written on any HP StorageWorks Ultrium 460 tape drive for ProLiants can be interchanged directly with Ultrium 2 tape drives from other vendors

18306
Paula

PRODUCT HIGHLIGHTS *(continued)*

High Reliability

HP ensures the highest level of quality with media specifications that far exceed industry standards.

- HP has 20 custom-built testing chambers that are in use 24 hours per day, 365 days per year.
- It is doubtful if any other media supplier carries out such exhaustive qualification of the drive and media as they are used in the field. This is because no media supplier is exposed to support of hardware in the field and has a vested interest in reducing media-induced hardware issues.
- In 2001, 170,000 different media tests were performed, accounting for 1.3 million test hours.
- HP testing includes procedures like drop testing, load/unload, environmental stress tests and archival simulation that are not required to gain the Ultrium logo.

LTO Ultrium 100/200 GB tape technology uses a "best of breed" technology approach, taking the best features from other tape technologies and combining them into a single new technology without the need for design compromises to accommodate legacy compatibility requirements. Designed for the demanding environment of large-scale libraries, Ultrium 400 GB media is certified for 1 million passes or 100 full back ups and has a 30 year archival storage life.

Compression

LTO-DC is an enhanced version of Advanced Lossless Data Compression (ALDC) hardware compression. Compression is automatically turned "off" if the compressed data would cause an overall expansion in data size.

Compatibility Testing

HP Ultrium 400 GB data cartridges are fully supported and compatible with HP StorageWorks Ultrium tape products. Because HP Ultrium media is Ultrium 2 logo compliant, it may be used with any other non-HP device that bears the Ultrium 2 logo.

Backward Read Compatible

HP Ultrium 400 GB media cannot be read or used to write new data in an Ultrium 1 drive or automation product.

HP Branded Media Specification

HP believes that its test program for HP Ultrium media is the most thorough and comprehensive in the industry. In order to carry the HP Brand, designated cartridges must satisfy an exhaustive battery of additional procedures that relate directly to how the product is used in real life situations when real data and real businesses are at stake. Some of these procedures – e.g. "five corner" environmental interchange, load/unload for automation, ageing simulation and drop testing are not found in the standard logo test. Over 1 million test hours on media and drives per year ensures that HP Ultrium media will always offer maximum reliability even in the most extreme conditions.

Why Buy LTO?

If you have not yet standardized on a large-tape format then LTO Ultrium tape technology is the preferred choice offering

- Highest reliability through no compromise design and innovative features like Data Rate Matching.
- High capacity to meet the challenges of data deluge
- High performance - even on slower hosts
- An industry standard - Linear Tape Open format
- Investment protection - with a four generation roadmap with plans to deliver an 800 GB (native) tape drive by the year 2006

Service and Support

Includes a limited lifetime warranty, fully supported by a worldwide network of resellers and service providers and toll-free 7 x 24 technical phone support during the warranty period

Warranty

Limited lifetime warranty

18.205
Paula

QuickSpecs

hp Ultrium 400GB Data Cartridge

SPECIFICATIONS

Dimensions (H x W x D)	**Shipping**	
	HP Ultrium 400 GB Data Cartridge (single pack)	1.8 x 4.45 x 4.37 in/2.1 x 11.3 x 11.1 cm
Weight	**Shipping**	
	HP Ultrium 400 GB Data Cartridge (single pack)	0.63 lb/285.2 g
UPC Code	HP Ultrium 400 GB Data Cartridge (single pack)	8 08736-39584 7
Capacity	200 GB native/400 GB compressed (2:1)	
	512 tracks with a track pitch of 20.17 um	
Tape Format	Recording Method	Linear 1/2" tape, 8 track parallel recording.
	Recording Format	Ultrium Generation 2
Media Type	Ultrium 2 media – 609m, 1/2" tape, Metal Particle (MP++) formulation with 4K Cartridge Memory	
Media Durability	1,000,000 passes on any area of tape, equates to over 20,000 end-to-end passes/100 full tape backups	
Basefilm Type	6 um Metal particle PEN	
Tape Thickness	8.9um ± 0.3um	
Magnetic Layer Coercivity	2,350 Oe	
Tape Life	30 years archival storage	
Environmental Details	Operating Temperature	50 to 113 ° F/10 to 45 ° C
(Media should not be used outside of these parameters)	Day to Day Storage Temperature	60 to 90 ° F/16 to 32 ° C
	Long Term Storage Temperature – Ambient	41 to 73 ° F/5 to 23 ° C
	Maximum Wet Bulb Temperature	79 ° F/26 ° C
	Non Condensing Relative Humidity	10 to 80%

CONFIGURATION INFORMATION

RELATED OPTIONS

Media	HP Ultrium 400 GB* data cartridge ** (recommended)	C7972A
	HP Ultrium 200 GB* data cartridge** (*assumes 2:1 compression ratio)	C7971A
	HP Ultrium Universal cleaning cartridge**	C7978A

** Customers and partners can order additional LTO media (part no.: C7971A and C7972A) and LTO cleaning cartridges (part no.: C7978A) as HP third-party products through traditional Major Account Direct, Partner Direct, and Compaq Direct processes This is the same method that HP printers are ordered via Compaq direct systems today.

Hewlett-Packard shall not be liable for technical or editorial errors or omissions contained herein The information is provided "as is" without warranty of any kind and is subject to change without notice The warranties for Hewlett-Packard products are set forth in the express limited warranty statements accompanying such products Nothing herein should be construed as constituting an additional warranty

18.304
Paula

ANEXO UNIDADE DE BACKUP ROBOTIZADO
PARTE S

18.303
Paula

QuickSpecs

HP StorageWorks E2400-160 FC Interface Controller

Overview

Product Description

As storage capacity doubles every year businesses struggle to keep up with storage demands. Contributing factors to the growth in storage include applications becoming more complex, increasing file sizes, and an increasing number of government regulations regarding the storage of information. In addition to this, companies need to operate 24 x 7 to not be left behind. To keep up with these demands, they look to service providers to guide them in their decision making when creating a data protection strategy. Key components of this strategy should be reliability, lower total cost of ownership, investment protection, and ease of use.

The StorageWorks E2400-160 Fibre Channel Interface Controller offers all of the above. The E2400-160 is a key component in a complete data protection solution. It is an embedded Fibre Channel-to-SCSI controller with two, user configurable, 1Gb or 2 Gb Fibre Channel ports and four SCSI (LVD) ports and allows multiple host servers to communicate with a SCSI tape device over a Fibre Channel link. It is targeted towards customers whose backup window is impacting mission critical business operations and need to reduce or remove the impact of the backup window.

To install the E2400-160 a card cage is required which is installed into the back of the library. The card cage kit contains one card cage and a two slot adapter panel to accommodate both E2400 and E2400-160 controllers in the same card cage.

The embedded interface controllers are for customers who want to save on rack space by implementing their Fibre Channel to SCSI interface controller inside the library. By connecting tape libraries to the SAN backup window, performance can increase up to five times faster therefore reducing the downtime of business operations.

The E2400-160 ships with an accessory kit, which includes documentation, SCSI cables, cable clips, and serial cable. Additional Fibre Channel cable kits will need to be purchased. A SC-LC for 1 Gb Fibre Channel connections or LC-LC for 2 Gb Fibre Channel connections.

Benefits at a Glance

- Embedded Fibre Channel interface controller
- 1 Gb or 2 Gb Fibre Channel support
- Web-based management
- LVD SCSI Ultra3 support
- Supports 2 Gb/1 Gb FC on same interface controller
- Provides connectivity for one Ultrium or /SDLT/ drive per SCSI bus

Tape Library Connection Chart

ESL9198 Libraries
ESL9322 Libraries
ESL9326 Libraries
ESL9595 Libraries

Key Features

- Embedded Fibre Channel interface – saving rack space
- 1 Gb and 2 Gb Fibre Channel support with the ability to support either on the same interface controller
- LVD SCSI capability (VHDCI connects), Narrow/Wide Fast/Ultra3
- Additional Fibre Channel Cables are required: a 1 Gb Fibre Channel cable OR a 2 Gb Fibre Channel cable
- Minimum Configuration:
  One interface controller card, supports two FC connects and four SCSI Ports/four tape drives
- Maximum Configurations:
  Four interface controller cards, supports eight FC connects and sixteen SCSI Ports/16 tape drives

hp
invent

18.302
Paula

Models

E2400-160 Fibre Channel Models	Product numbers associated with the E2400-160 Fibre Channel Interface Controller	
	ESL9000 E2400-160 FC Interface Controller Includes: a single 2FC x 4LVD SCSI interface controller card, four VHDCI-WIDE SCSI cables and a User Guide CD. One controller is necessary for every four tape drives up to a maximum of four controllers for 16 drives. This SKU does not include a card cage. A card cage will need to be installed first prior to installation of the FC interface Controller. An installation CarePaq will be quoted with this SKU.	330839-B21
	ESL9000 Card Cage Kit The ESL9000 card cage is a six slot cPCI cage that accommodates interface cards. The SKU includes one card cage, two slot adapter panels, an installation guide, and User Guide CD. Only the card cage kit is required per library frame. Fibre Channel Interface Controller Cards must be ordered separately.	330838-B21
	ESL9000 Card Cage Slot Adapter Required only for customers upgrading an existing card cage to include E2400-160 controllers. If a card cage is already installed and E2400-160's are being added, the two slot filler panel will enable both an E2400 and E2400-160 FC Interface Controllers to exist in the same card cage. Includes a two slot adapter panels and four slot covers. HP Service Installation is required and only one slot adapter kit is necessary per card cage.	330838-B22

Interface Controller Cables	2m SW LC/SC FC Cable Kit	221691-B21
	5m SW LC/SC FC Cable Kit	221691-B22
	15m SW LC/SC FC Cable Kit	221691-B23
	2m SW LC/LC FC Cable Kit	221692-B21
	5m SW LC/LC FC Cable Kit	221692-B22
	15m SW LC/LC FC Cable Kit	221692-B23

Product Highlights

Connectivity

It is recommended that one LTO or SDLT tape drive be configured per SCSI bus for maximum performance. The robot arm counts as a SCSI device but does not need its own SCSI bus connection because of the low bandwidth requirement of the robot arm

Manageability

Configure and access device status, statistics and mapping with HTTP, SNMP, serial port visibility, Telnet, or FTP

Web-based management tool – a web application management tool for configuring and managing the interface controller. The graphical user interface allows endusers to configure the interface controller by "pointing and clicking" on graphical representations interface controllers. The manager can be accessed from any computer with a web browser (as long as this computer has security access privileges to access the interface controller).

LUN Mapping/Masking – allows multiple servers to selectively share storage resources while protecting data integrity via unique device map to host assignments.

SCSI Controller Command (SCC) – provides a means for the host server operating system to determine the number of FC LUNs supported by the NSR.

Configurations Supported

For complete supported configurations please visit the HP StorageWorks Enterprise Backup Solutions website at: http //www.hp.com/go/ebs

18.300
Paula

Warranty, Service and Support, and Care Pack Information

Warranty

One year parts exchange with next business day response

Additional warranty protection can be purchased for all library products

Software Product Services

- Stand-alone telephone support
- Rights to a new license
- Media and documentation updates

Hardware Product Services

- Installation services
- On-site maintenance (includes warranty support)
- Response time upgrades during the warranty period
- Post-warranty coverage

For additional hardware installation and maintenance information please refer to the URLs listed below:
http://h18005.www1.hp.com/services/CarePaq/us/hardware/cp_storage.html
http://h18005.www1.hp.com/services/CarePaq/us/install/cp_storage.html

Warranty Upgrade Options

- Response – Upgrade onsite response from next business day to same day 4-hours
- Coverage – Extend hours of coverage from 5 days x 9 hours to 7 days x 24 hours
- Duration – Select duration of coverage for a period of 1, 3 or 5 years

For additional CarePaq (hardware and software) information, as well as orderable part numbers, please refer to:
http://h18005.www1.hp.com/services/carepaq/us/

Care Packs

Installation Services:

Additional warranty protection and/or Hewlett-Packard Installation packages can be purchased for the ESL library products:

Description	Option	Product Number
Tape Library Implementation into SAN (HW install included)	5DQ	HA114A1
Tape Library Installation (SCSI direct attached)	5DQ	HA113A1

Warranty Upgrade Services:

HW-1yr next day onsite	HA101A1
HW-1yr 8x5 onsite	HA103A1
HW-1yr 24x7 onsite	HA104A1
HW-3yr next day onsite	HA101A3
HW-3yr 8x5 onsite	HA103A3
HW-3yr 24x7 onsite	HA104A3

CarePaqs

For the most up-to-date list of CarePaq hardware installation and maintenance information please refer to the URL listed below.
http://www.compaq.com/services/carepaq/index.html

NOTE. Certain restrictions and exclusions apply Consult the Hewlett-Packard Customer Support Center for details

- CarePaq and Care Pack are defined as an upgrade to the product warranty attribute, available for a specific duration and hours of coverage.
- CarePaq and Care Pack is not available for less than the products warranty duration.
- CarePaq and Care Pack are available for sale anytime during the warranty period for most products, but the commencement date will be the same as the Warranty Start Date (delivery date to enduser customer) Proof of purchase may be required
- CarePaq and Care Pack services are prepaid.

18 249
Paula

Technical Specifications

HP StorageWorks Interface controller Specifications

Dimensions (HxWxD)	E2400-160 FC Interface Controller	3.74 x 6.193 in/16.19 x 15.7302 cm
	ESL9000 Card Cage	23.0 x 16.5 x 13.25 in/58.42 x 41.91 x 33.655 cm
	ESL9000 Card Cage Adapter	12.5x 9.5x 2.25 in/31.8cm x 25.1cm x 5.7cm
Weight (lb/kg)	E2400-160 FC Interface Controller	9.6 lb/4.4kg
	ESL9000 Card Cage	17.0 lb/7.711 kg
	ESL9000 Card Cage Adapter	9.6 lb/ 4.4 kg
UPC Codes	E2400-160 FC Interface Controller	808736892773
	ESL9000 Card Cage	808736892759
	ESL9000 Card Cage Adapter	808736968515
Control Indicators	Chassis LED	Power (Green=good)
	Fibre Channel LEDs	Link, Activity
	Ethernet Port	Link, Activity
	SCSI Buses	SCSI Activity LEDs, 1 per channel
Connectivity	Ports	Chassis: RJ 11 Serial port
		Fibre Channel: 22 Fibre Channel per card
		LVD SCSI: 4 LVD SCSI ports per card (VHDCI connects)
		One (1) 10/100 TX Ethernet port
	Topology	Switched Fabric
	Protocols	FCP (Ultra2), PLDA
Temperature Range	Operational	50° to 104° F/5° to 40° C
Relative Humidity	Operational	5% to 85%
	Storage	5% to 85%
Compatibility	http://www.hp.com/go/ebs	

The information contained herein is subject to change without notice.

The only warranties for HP products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. HP shall not be liable for technical or editorial errors or omissions contained herein.

Doc.
000233

CPL/AC

PREGÃO
050/2003

LOCAÇÃO DE EQUIPAMENTOS DE INFORMÁTICA INCLUINDO ASSISTÊNCIA TÉCNICA E TREINAMENTO

COBRA TECNOLOGIA – MANUAL VOLUME 2

2003
PASTA 30

COBRA TECNOLOGIA
Integradora de Soluções em TI e Serviços

ANEXO UNIDADE DE BACKUP ROBOTIZADO
PARTE D/2

13.298
Paula

18.247
Paula

4 Managing Media

In This Chapter

This chapter gives detailed information on how to manage your media, including:

“Overview of Data Protector Media Management” on page 99

“Creating a Media Pool” on page 102

“Adding Media to a Media Pool” on page 107

“Formatting Media” on page 108 and “Importing Media” on page 113

“Appending Backups to Media” on page 117

“Using a Pre-Allocation List of Media for Backup” on page 119

“Selecting Media for Backup” on page 120

“Setting Data Protection for Media” on page 122

“Recycling Media” on page 123

“Moving Media to Another Pool” on page 124

“Exporting Media from Data Protector” on page 125

“Modifying Media Locations” on page 126 and “Modifying Media Descriptions” on page 127

“Verifying Data on a Medium” on page 128

“Scanning Media in a Device” on page 129

“Checking the Condition of a Medium” on page 131

“Searching for and Selecting a Medium” on page 135

“Entering a Medium into a Device” on page 136 and “Ejecting a Medium from a Device” on page 137

“Vaulting Media” on page 140

“Copying Media” on page 143

“Detection of Write-Protected Media” on page 147

“Using Different Media Format Types” on page 148

“Modifying Views in the Media Management Window” on page 149

Overview of Data Protector Media Management

Data Protector provides a powerful media managing functionality that allows simple and efficient management of a large number of media.

NOTE Data Protector recognizes and uses different format types to write data to media. For limitations incurred, refer to "Using Different Media Format Types" on page 148.

- Grouping media into logical groups called media pools, which allow you to manage large sets of media without having to worry about each individual medium.
- Data Protector keeps track of all media and the status of each medium, including data protection expiration time, availability of media for backup, and a catalog of what has been backed up to each medium.
- Fully automated operation. If Data Protector has control of enough media in the library devices, the media management functionality allows backups to run without the need for an operator to handle the media.
- Automated media rotation policies, so that you do not have to enforce policies manually.
- The ability to explicitly define which media and which devices you want to use for a certain backup.
- Optimized media management for specific device types, such as standalone, magazine, library devices, and large silo devices.
- Automatic recognition of Data Protector media and other popular tape formats.
- Recognition and support of barcodes on large library and silo devices with barcode support.
- Recognition, tracking, viewing, and handling of media used by Data Protector in large library and silo devices.

18224
Fauk

- The ability to store information about media in a central place and share this information among several Data Protector cells.
- Support for **media vaulting**, also known as **archiving** or **off-site storage**.
- Creation of additional copies of media. Media copying can be either manually started or automated.

Information about the media used is stored in the IDB.

For more information on media management, see the *HP OpenView Storage Data Protector Concepts Guide*.

Figure 4-1 indicates the relationship among the components, backup specification, devices, and media pools. The pool is used during a backup session. A default pool is part of the device definition. However, a different pool can be specified in the backup specification.

Figure 4-1 **How the Media Pool Relates to Other Components**

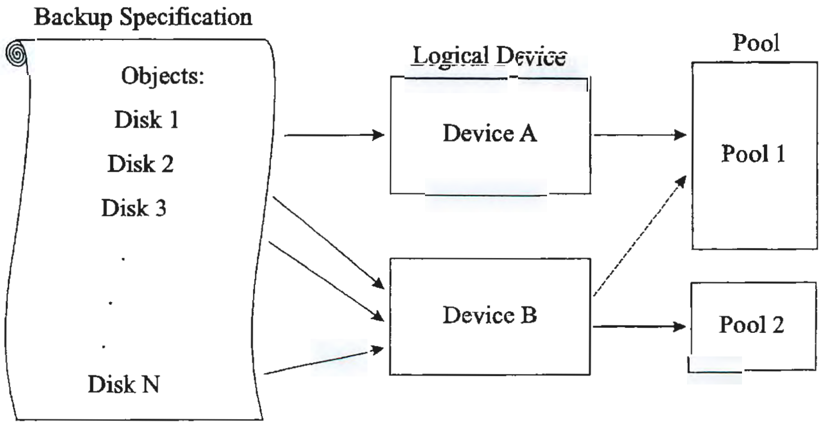

Media Life Cycle

A typical media life cycle consists of the following steps:

1. Preparing media for backup. This includes formatting media for use with Data Protector and assigning media to a media pool. The media pool is used to track these media. See the following topics for detailed information:

Paula
12293

"Creating a Media Pool" on page 102.

"Adding Media to a Media Pool" on page 107.

2. Using media for backups. This includes how the media are selected for a backup, what media condition factors are checked (for example, the number of overwrites), how new backups are appended to the media, and when data on the media can be overwritten.
3. Vaulting media to a safe place (vault).
4. Recycling media once data on the media is not needed anymore. These media can then be reused.
5. Retiring Media. Once the medium has expired (according to its maximum usage criteria), it is marked as Poor and no longer used by Data Protector. See "Factors Influencing the Condition of Media" on page 132 for more information.

Details are explained in the following sections.

18.242
Paula

Creating a Media Pool

What Is a Media Pool?

A media pool represents a set of media of the same type (for example DLT) used for backup, with the same usage policy and properties. For example, you may have one media pool for regular backup, one for archive backup, and one for each department.

What Is a Free Pool?

A free pool is an auxiliary source of media of the same type (for example, DLT) for use when all free media in a regular pool run out. This helps to avoid failed backups due to unavailable media.

Media are moved between regular and free pools in two events:

- Allocation. Media are moved from a free pool to a regular pool.
- Deallocation. Media are moved from a regular pool to a free pool. You can specify in the GUI whether deallocation is performed automatically.

Protected (allocated, used) media belong to a specific regular pool (such as a SAP pool), while free Data Protector media can be automatically moved to a free pool. This free pool is later used for allocation of free media to a specific regular pool during backup, when needed.

See the *HP OpenView Storage Data Protector Concepts Guide* for more information on media pools.

Default Media Pools

Data Protector provides default media pools for each media type that you can use in your initial configuration, for example: Default_DDS.

If you do not want to create a media pool at this time and to use the default media pools instead, go to “Adding Media to a Media Pool” on page 107 for instructions.

How to Create a Media Pool

Create a new media pool in the `Devices & Media` context using the `Add Media Pool` wizard. For detailed steps, refer to the online Help index keyword “adding media pools”.

What’s Next?

The next step is to add media that you want to use for backup to the media pool. See “Adding Media to a Media Pool” on page 107 for instructions.

18.291
Paula

Properties of a Media Pool

This section describes the properties of a media pool. You specify them when you are configuring the media pool. Some of the properties can be modified later.

Pool Name

A media pool name identifies a media pool. It can be up to 32 characters long, including spaces. You should assign a meaningful name that will help you identify the media pool later, for example, your department name.

Description

A description is optional and helps you to identify the media pool. It can contain any characters and can be up to 80 characters long.

Media Type

Data Protector shows you a list of available media types for your configuration.

You can select among DDS, DLT, ExaByte, AIT, QIC, T3480/T4890/T9490, T9840, T9940, T3590, SD-3, Tape, Optical (which stands for magneto-optical media), File, LTO-Ultrium and SuperDLT.

Once you select the media type, Data Protector calculates the available space on the media for that media pool. This calculation is based on the selected media type.

Media Allocation Policy

The media allocation policy defines the order in which media are accessed within a media pool, so that media wear out evenly.

For more information on how Data Protector selects media for backup, see “Selecting Media for Backup” on page 120.

Strict Directs Data Protector to require a specific medium. The medium has to be already formatted for use with Data Protector. If this policy is used, Data Protector does not format media. This allocation policy should be used with library devices to prevent accidental overwrite of non-Data Protector media in the library and where even usage of media has priority.

Loose Directs Data Protector to accept any suitable medium in the pool except a medium in *poor* condition or a protected medium. This option is combined with the `Allocate unformatted media first` option.

18.290
Paula

If InitOnLoosePolicy is set to 1 (by default, it is set to 0) media that are unrecognized by Data Protector (new media) are automatically formatted. This policy is preferred if you want unattended backup to succeed, as it maximizes the number of media Data Protector can choose from.

Unformatted media first This is a modification of the Loose policy. If selected, this policy directs Data Protector to give preference to unknown media, even if unprotected media are available in the library. This is recommended if Data Protector is the only application using the library and you want to have even usage of all media.

Use free pool Directs Data Protector to search in the free pool for suitable media in addition to the regular pool. By default, this option is OFF.

See "Selecting Media for Backup" on page 120 for detailed information.

See the *HP OpenView Storage Data Protector Concepts Guide* for more information on free pools.

Media Usage Policy

Media usage policy controls how new backups are added to already used media.

Appendable A backup session starts writing data to the space remaining on the last medium used in the previous backup session. Subsequent media needed in this session are written from the beginning of the tape, hence only unprotected or new tapes can be used. Data may be appended from any backup specification to any other backup specification. Appending media conserves media space but can add complexity to a restore operation, because one medium can contain data from several backup sessions.

Non-Appendable A backup session writes data beginning at the first position on the first available medium for backup.

Appendable on incrementals only The first medium used in a backup session is appended to only if an incremental backup is performed. If several appendable media are available in the pool, the least recently written to

18 239
Paula

medium is used first. If additional media are needed during the same backup session, they must be free and not contain any protected backups. This media usage policy will create media which will contain a full backup, followed by any number of incremental backups.

NOTE If you use the append functionality and the backup requires more than one medium, only the first medium used can contain backed up data from a previous session. Subsequently, Data Protector will use empty or unprotected media only.

See "Appending Backups to Media" on page 117 and "Selecting Media for Backup" on page 120 for more information.

Magazine Support Magazine support allows you to use a set of media configured as magazines. A backup device used with these media must have support for magazines, such as the HP 12000e.

You can set this option when you configure a new media pool.

See the following sections for more information:

- "Configuring Magazine Devices" on page 34 for instructions on how to configure a magazine device.
- "Formatting Media" on page 108 for instructions on how to format a full magazine or a single medium in the magazine.
- "Importing Media" on page 113 for instructions on how to import a full magazine or a single medium.

Media Condition Factors Media condition factors define the status of the media, thus determining how long media can be reliably used for backup. If a pool uses the free pool option, the media condition factors are inherited from the free pool. Data Protector calculates the status of media in use via media condition factors. The two media condition factors you can select are:

Medium valid for The age of a medium is calculated as the number of months that have elapsed since it was formatted. Once a medium is older than the threshold number of months, it is marked as poor. The default threshold is 36 months.

18288
Paula

Maximum number of overwrites The usage of a medium is defined as the number of overwrites from the beginning of the medium. Once the medium has more than the threshold number of overwrites, it is marked as poor. The default threshold is 250 overwrites, except for DDS tapes, for which it is 100 overwrites.

For more information on how media condition factors are calculated, see "Changing How Media Condition Is Calculated" on page 134.

18.287

Adding Media to a Media Pool

Once you have created a media pool, you have to add the media that you want to use for backup to this media pool.

How to Add the Unused Media

To add unused media to the media pool, see “Formatting Media” on page 108. If your media allocation policy for the media pool is set to `loose`, formatting media as a separate step is not required. If `InitOnLoosePolicy` is set to 1 (by default, it is set to 0), the media are formatted before the backup session in which they are used. See “Media Usage Policy” on page 104 for more information.

How to Add Used Media

To import previously used Data Protector media without overwriting them, see “Importing Media” on page 113.

To add used non-Data Protector media to the media pool, you have to reformat them. See “Formatting Media” on page 108.

For more information on how Data Protector handles media used by other applications, see “Recognizing Other Data Formats” on page 111.

Labeling Media

Data Protector labels each medium with a unique media label and medium ID. Both are stored in the IDB and allow Data Protector to manage the medium. The medium ID is assigned by Data Protector. The media label is a combination of the user-defined description and the barcode of the medium (if the medium has a barcode and the `Barcode Reader Support` option is enabled). For example, [CW8279]Default DLT_1 is a media label with the Default DLT_1 description and the CW8279 barcode.

In the Data Protector GUI, you can sort media by media label. You do this by clicking the `Media label` field in the Results Area.

What’s Next?

Once you have added media to the media pool, you can select data that you want to back up. Refer to Chapter 5, “Backup,” on page 151 for instructions.

13286
Paula

Formatting Media

What Is Formatting Media?

Formatting media prepares them for use with Data Protector by saving the information about the media (media IDs, description and location) in the IDB, and also writes this information on the medium itself (medium header). When you format media, you also specify to which media pool the media belong.

NOTE

Certain media management operations (initialize, scan, enter, and eject) can be performed during backup or restore. Data Protector uses “advisory locking” to ensure that a backup or restore will not fail if the device cannot be locked. Other operations (backup, restore, import, and copy) require device locking to proceed with backup or restore.

When to Format Media

You need to format media before the media can be used for backup. If the media are not formatted before backup and the `Loose` media allocation policy is defined for the media pool, and the global variable `InitOnLoosePolicy` is set to 1 (default is 0) Data Protector automatically formats new media when they are selected for backup. In this case, the media are labelled with default values.

Non-Data Protector media must be formatted before backup.

Recognition of Other Formats

Data Protector recognizes common media formats, if the medium was already in use. See “Recognizing Other Data Formats” on page 111 for detailed information.

Formatting with Padding Blocks

You can extend the size of the medium header and fill it up with incompressible data, padding blocks. This becomes useful when creating media copies. The padding blocks are not copied to the target medium. This way you make sure that the target medium does not reach the end of the tape before the source medium.

Tape padding is disabled by default. To enable it, set the `OB2BLKPADDING_n` variable in the `omnirc` file on the system with the backup device connected. For more information, see “Using Omnirc Options” on page 525.

18285
Paula

How to Format Media

To format media, browse for the specific device, media pool, or library slot in the `Devices & Media` context, right-click it and click `Format`. For detailed steps, refer to the online Help index keyword "formatting media".

If you use library devices, you can select multiple slots using the `Ctrl` key and format several media in a single step. For detailed steps, refer to the online Help index keyword "formatting media in library devices".

TIP

To format media used by other applications, use the `Force Operation` option. Data Protector protected media cannot be re-formatted using this option. You have to first remove the protection. See "Recycling Media" on page 123 for more information.

NOTE

When selecting the `Medium Size` option, choose between `Default` and `Specify MB`. If you have chosen the `Default` medium size, the estimated and not the real size of the media is shown. Be aware that the total media size is set for non-compressed media. Hardware compression of the device may double the space on the media. The correct media size is shown when the media are full.

Cartridge Memory Data Initialization

When using Cartridge Memory enabled LTO drive(s) with Cartridge Memory enabled media, Cartridge Memory data is formatted automatically at the time the medium is formatted.

Cartridge Memory Reformat

To synchronize header segment information in the Cartridge Memory with header segment information on the medium, use the `Cartridge Memory Reformat` action. The information is then updated in the IDB. You can reformat the Cartridge Memory for a specified slot or Data Protector medium. Refer to Figure 4-2.

Figure 4-2 **Cartridge Memory Reformat for Specific Slot**

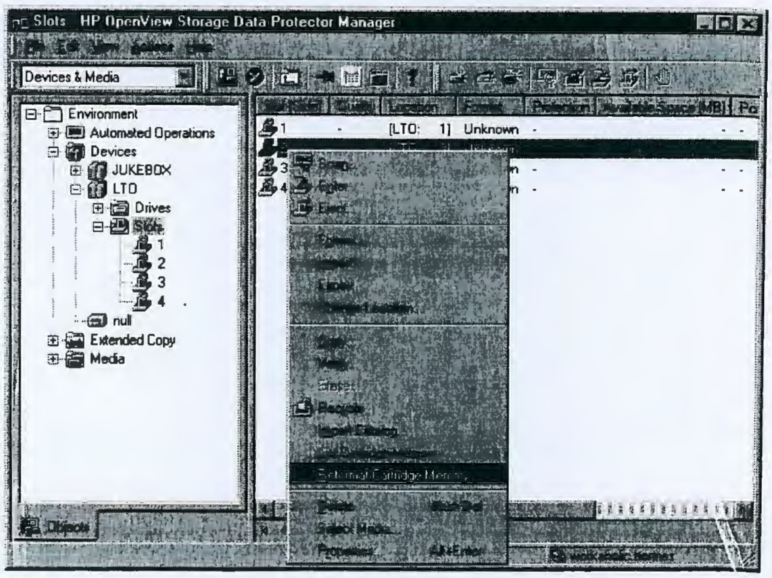

What's Next?

Once you have formatted your media, you may use the media for backup. See Chapter 5, "Backup," on page 151 for more information on how to configure backups.

Formatting Media in a Magazine

If you are using a device with magazine support, Data Protector allows you to format all media or a single medium in the magazine.

How to Format a Full Magazine

To format a full magazine, browse for the media pool used for the device, right-click it and click `Format Magazine`. For detailed steps, refer to the online Help index keyword "formatting media in magazines".

How to Format a Single Medium in a Magazine

To format a single medium in a magazine, browse for the media pool used for the device, right-click it and click `Format`. For detailed steps, refer to the online Help index keyword "formatting a single medium in magazines".

TIP

To format media used by other applications, use the `Force Operation` option. Data Protector protected media cannot be re-formatted using this option. You have to first remove the protection. See "Recycling Media" on page 123 for more information.

What's Next?

Once you have formatted your media, you may use these media for backup. See Chapter 5, "Backup," on page 151 for more information on how to configure backups.

Recognizing Other Data Formats

Recognized Formats

To prevent accidental overwrite of data already written to the media, Data Protector recognizes a number of different tape formats:

Table 4-1 Data Protector Media Format Categories

Media Format	Data Protector Behavior
unknown or new	`Loose` Policy: formatted and used for backup only if the global variable `InitOnLoosePolicy` is set to 1 `Strict` Policy: not used for backup
media written with compression, now used without compression	
media written without compression, now used with compression	
foreign Data Protector (from another cell)	not used for backup unless imported or formatted with the `Force Operation` option
tar, cpio, OmniStorage, OmniBack I, ANSI label, filesystem	not used for backup unless formatted with the `Force Operation` option
Data Protector unprotected media	used for backup
Data Protector protected media	used for appending backups

15.252
Paula

NOTE Do not rely on Data Protector to recognize other media types, as recognition depends on the platforms you use.

Cartridge Memory Enabled Recognition

With Cartridge Memory enabled LTO drive(s) used with Cartridge memory enabled media, Cartridge Memory provides the attributes for giving specific ownership information. Data Protector uses this ability to recognize media under ownership of other applications.

NOTE If you try to read from a medium that was written using hardware compression with a device that does not support hardware compression, Data Protector cannot recognize the medium and read the data. Therefore, the medium will be treated as unknown or new.

18.281

Importing Media

Importing media adds media already used by Data Protector to a media pool, without losing the data on the media. Media used by Data Protector are media that were formatted by Data Protector, but exported from the Data Protector cell.

Importing a medium writes detailed information about backed up data on the medium to the IDB, so that you can later browse it for a restore.

Use media import when moving your media between Data Protector cells.

This operation is not available for media in free pools.

NOTE Attribute information such as object or media size will not be reconstructed during import. Thus the size of the imported objects will be shown as 0 KB.

Importing can take a considerable amount of time, depending on the device and media used.

IMPORTANT Import all media used in one backup session at once. If you add only some media from the backup session, you will not be able to restore data spanning to other media.

How to Import Media

To import media, browse for the specific device, media pool or library slot in the `Devices & Media` context, right-click it and click `Import`. For detailed steps, refer to the online Help index keyword “importing media”.

If you use library devices, you can select multiple slots using the `Ctrl` key and import several media in a single step. Refer to Figure 4-3. For detailed steps, refer to the online Help index keyword “importing media in library devices”.

18.280
Paula

Figure 4-3 **Import Multiple Media**

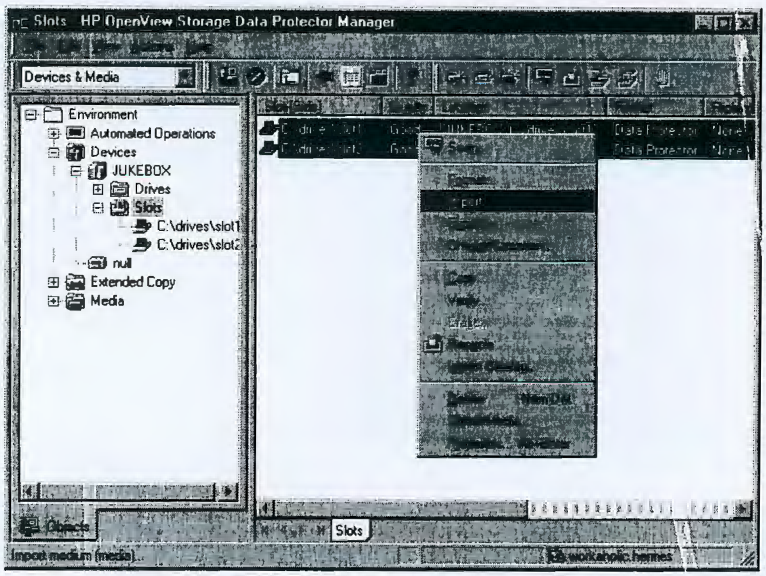

Importing the Catalog from Media

Importing the catalog from a medium writes the information about file versions into the IDB, enabling you to browse files and directories for restore.

Use `Import Catalog` if the catalog protection for a particular object has expired and you can no longer browse its files and directories.

This operation is not available for media in free pools.

How to Import the Catalog from Media

To import the catalog from a medium, browse for the specific medium, device or library slot in the `Devices & Media` context, right-click it and click `Import Catalog`. Refer to Figure 4-4. For detailed steps, refer to the online Help index keyword “importing catalogs from media”.

18.279
Paula

Figure 4-4 **Import Catalog**

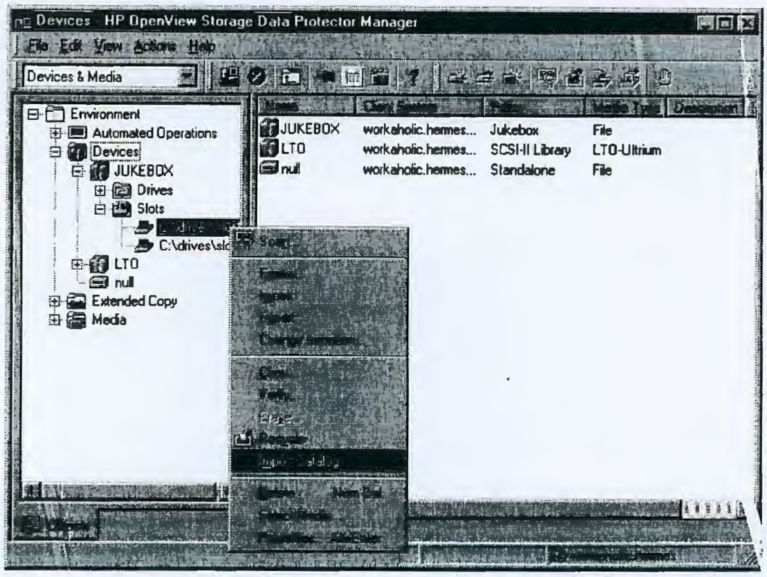
Devices - HP OpenView Storage Data Protector Manager
File Edit View Actions Help
Devices & Media
Environment
Automated Operations
Devices
JUKEBOX
Drives
Slots
C:\drives\sl
LTO
null
Extended Copy
Media
JUKEBOX workaholic.hermes... Jukebox File
LTO workaholic.hermes... SCSI-II Library LTO-Ultrium
null workaholic.hermes... Standalone File
Import Catalog

Importing Media in a Magazine Device

If you use a device with magazine support, Data Protector allows you to import all media or a single medium into the magazine.

Prerequisite

The media pool for the magazine device must be configured with the `Magazine Support` option enabled.

How to Import All Media

To import all media in a magazine device, expand the media pool used for that device in the `Devices & Media` context, right-click the `Magazines` item and then click `Import Magazine`. Refer to Figure 4-5. For detailed steps, refer to the online Help index keyword "importing media in magazines".

18248
Paula

Figure 4-5 **Import Magazine**

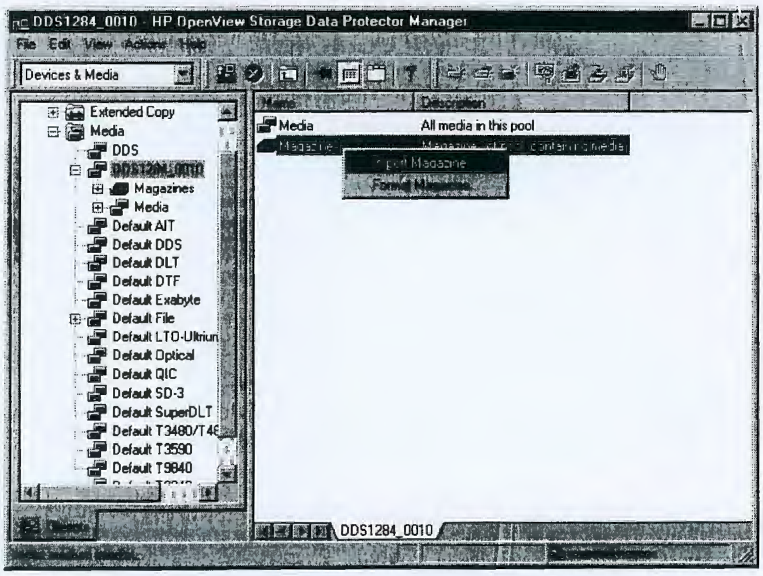

How to Import a Single Medium into a Magazine

To import a single medium into a magazine device, expand the media pool used for that device in the `Devices & Media` context, select the specific magazine, right-click the `Media` item and then click `Import`. For detailed steps, refer to the online Help index keyword “importing a single medium in a magazine”.

What’s Next?

Once you have imported the media, you may use these media for backup. See Chapter 5, “Backup,” on page 151 for more information on how to configure backups.

15.277
Paula

Appending Backups to Media

Data Protector allows you to add new backups to media which already contain backups. This method conserves media space.

Limitation Backups cannot be appended on media used in Travan devices.

The appendable media usage policy can be selected when configuring a media pool. Appendable media contain some currently protected objects; the media must be in good condition and must not be full.

If several devices are used with load balancing, the appendable concept applies on a per device basis, that is, each device uses an appendable medium (if available) as the first medium in a backup session. The backup sessions appending data on the same medium do not have to use the same backup specification.

Two alternatives of appendable media usage policies are available:

- **Appendable:** The first medium used in a backup session uses the space remaining on the medium from the previous backup session. If several appendable media are available in the pool, the least recently used medium is used first. If additional media are needed during the same backup session, they must be free and not contain any protected backups. For this media usage policy, the type of backup (full or incremental backup) can be mixed in any order on the media.
- **Appendable on incrementals only:** The first medium used in a backup session is appended to only if an incremental backup is performed. If several appendable media are available in the pool, the least recently used medium is used first. If additional media are needed during the same backup session, they must be free and not contain any protected backups. This media usage policy will create media which will contain a full backup, followed by any number of incremental backups.

TIP If you want to create tapes which contain only one full backup and the incremental backups related to the same client, configure Data Protector as follows:

- Configure one pool per client with the media usage policy `Appendable on Incrementals only`.

18.376

- Link a different pool to each client in the backup specification, or create a separate backup specification per client.

This is a method to create media containing restore chains. Be aware that occasionally media will be created which contain incremental backups only.

See “Media Usage Policy” on page 104 for a description of media usage policy options like `Appendable`.

See “Selecting Media for Backup” on page 120 for more information on how the media usage policy influences how media are selected for backup.

To modify the settings later, open the properties for the media pool.

13.275
Paula

Using a Pre-Allocation List of Media for Backup

You can specify the order in which media from a media pool will be used for backup. This order is called a **pre-allocation list**. You specify the pre-allocation list when configuring a backup. The purpose of a pre-allocation list is to control which media will be used for a backup session. You have to match the pre-allocation list with the available media before each backup.

Depending on the allocation policy of the media pool, Data Protector behaves in two different ways:

- If the pre-allocation list is used in combination with the `Strict` media allocation policy, Data Protector expects the media in a backup device to be available in that order. If the media are not available, Data Protector issues a mount request. If the media mentioned in the pre-allocation list are loaded in a SCSI-II exchanger, Data Protector handles the media sequence automatically.
- If the pre-allocation list is used in combination with the Loose media allocation policy, media in the pre-allocation list are used first. If the media are not available, any suitable media in the library are used.

Preallocating Media for Backup

On how to preallocate media for backup, refer to the online Help index keyword "preallocating media".

18.274
Paula

Selecting Media for Backup

Data Protector media management automatically selects the most appropriate media for backup. This section explains various factors that influence how media are selected for backup.

Media Allocation Policy

You can influence how media are selected for backup using the **media allocation** policy. You can specify a `Loose` policy where any suitable media are used for backup, or a `Strict` policy where specific media have to be available in a predefined order.

See "Media Allocation Policy" on page 103 for more information.

Pre-Allocating Media

You can specify the order in which media from a media pool will be used for backup. This order is called a **pre-allocation list**. For more information, see "Using a Pre-Allocation List of Media for Backup" on page 119.

Media Condition

The condition of the media also influences which media are selected for backup. For example, media in good condition are used for backup before media in fair condition. Media in poor condition are not used for backup.

CAUTION

Media that are marked as fair will only be used if there are no protected objects on the media. Otherwise, a mount request is issued, and data might be lost before backup completes.

See "Factors Influencing the Condition of Media" on page 132 for more information.

Media Usage

The media usage policy also influences which media are selected for backup. See "Media Usage Policy" on page 104 and "Appending Backups to Media" on page 117 for a detailed description.

Media Selection

This section describes the criteria Data Protector uses to select media for backup.

18273
Paula

Media in poor condition are not used for backup. Media in fair condition are used only if no media in good condition are available. Media in good condition are sorted to use the one with the least number of overwrites first.

Media are always selected first from the specified pool and (optionally) from the free pool.

Table 4-2 **How Media Are Selected for Backup**

Allocation Policy	**Allocate Unformatted Media First**	**Data Protector Selection Order**
Loose	OFF	1. Pre-allocation list (if specified) 2. Appendable (as set in usage policy) 3. Unprotected Data Protector media 4. Unformatted media 5. Fair media
Loose	ON	1. Pre-allocation list (if specified) 2. Appendable (as set in usage policy) 3. Unformatted media 4. Unprotected Data Protector media 5. Fair media
Strict	(Not applicable)	1. Pre-allocation list (if specified) 2. Appendable (as set in usage policy) 3. Unprotected Data Protector media 4. Fair media

18.272
Paula

Setting Data Protection for Media

Data Protector keeps track of data on every medium used. When configuring a backup, you can protect your data from being overwritten by newer backups for a specified time. This protection is on a session basis: if data from several sessions is on the same media, the longest protection defines protection of the media. See “Data Protection: Specifying How Long Data Is Kept on the Media” on page 228 for detailed information.

You can also re-use the media by removing their protection. See “Recycling Media” on page 123 for more information.

18.371
Paula

Recycling Media

Data Protector keeps track of data on every medium used. When configuring a backup, you protect your data from being overwritten by newer backups for a specified time. See Chapter 5, "Backup," on page 151 for detailed information.

Keep in mind that on all media there may be data from several backup sessions. Each session can contain data from several backup objects (file systems).

Recycling removes the data protection from all backed up data on the medium, thus allowing Data Protector to overwrite it during one of the next backups. Recycling does not actually change the data on the medium, it only tells Data Protector that this data is not protected anymore. This option is not available for media in free pools.

For instructions on how to change the protection of a specific session or an object, see Chapter 9, "Managing the Data Protector Internal Database," on page 381.

How to Recycle Media

In the `Devices & Media` context, browse for a medium, right-click it and click `Recycle`. For detailed steps, refer to the online Help index keyword "recycling media".

Moving Media to Another Pool

Data Protector lets you move a medium from one media pool to another media pool of the same media type.

You need this feature if you want to reorganize the backups and rearrange the purpose of each pool. It is also useful when you want to use the medium in a device which is the default device of another media pool.

How to Move Media to Another Pool

In the `Devices & Media` context, browse for a medium, right-click it and click `Move to Pool`. For detailed steps, refer to the online Help index keyword "moving media".

Moving Media Using a Free Pool

When using a free pool, media are moved in two instances:

- When media are selected (allocated) for backup, they are moved from a free pool to a regular pool.
- When the media protection has expired, media are moved from a regular pool to a free pool.

This behavior depends on the free pool options selected.

For further information see "Creating a Media Pool" on page 102.

15.269
Paula

Exporting Media from Data Protector

What Is Exporting Media?

Exporting (removing) a medium removes the information about the medium and its contents from the IDB. Data Protector no longer recognizes that this medium exists. The medium and the data it contains remain unchanged. You can import the medium later, thus re-reading the information about data on the medium back to the IDB. See "Importing Media" on page 113 for instructions.

When to Export Media

If you want to move media to another cell, you have to export the media from one cell and import them to another.

Media that contain protected data cannot be removed. You have to recycle the media first. See "Recycling Media" on page 123 for instructions.

TIP

Export all the media from a backup session. If a backup session spans several media and you do not remove all of them, you will not be able to restore data; Data Protector still recognizes that data exists on the media, but the media will not be available anymore.

How to Export Media

In the `Devices & Media` context, browse for a medium, right-click it and click `Export`. For detailed steps, refer to the online Help index keyword "exporting media".

What's Next?

See "Adding Media to a Media Pool" on page 107 if you want to add media to another pool or move them to another cell.

See "Importing Media" on page 113 if you want to import media into another cell.

17.268
Paula

Modifying Media Locations

What Is a Location?

The media location helps you to physically locate the media. You enter the location when you format the media. The initial location information is written on the media and to the IDB.

You should modify the location whenever you move media to a different place, such as to off-site storage, for example, "Shelf 4-Box 3". The revised location information is only written to the IDB.

Data Protector allows you to create a list of pre-defined locations to simplify vaulting and archiving (also known as off-site storage). See "Vaulting Media" on page 140 for more information.

NOTE

When you modify a location, Data Protector modifies the location in the IDB and not on the medium itself.

If you export and import media again, the location information in the IDB is replaced with the location stored on the media.

TIP

You can modify the location of multiple media at the same time. This is useful for vaulting (archiving) purposes. See "Vaulting Media" on page 140.

How to Modify Media Location

Modify media location in the `General` property page for the medium. For detailed steps, refer to the online Help index keyword "modifying media location".

Modifying Media Descriptions

What Is a Description?

The media description helps you identify media. You can define a media description when you format new media. The initial description is written on the media and to the IDB.

If media were auto-formatted during backup, you may want to change the automatically-created description to something better suited to your needs. The revised description information will only be written to the IDB.

NOTE

When you modify a media description, Data Protector modifies the description in the IDB and *not* on the medium itself.

Therefore, if you export and import media that have note been updated, the description in the IDB is replaced with the description from the media.

Media Label

The media label is composed of the user-defined description and the barcode of the medium (if the medium has a barcode and the Barcode Reader Support option is enabled). For example, [CW8279]Default DLT_1 is the media label with the Default DLT_1 description and the CW8279 barcode. If the media description is changed, the descriptive part of the media label is changed too, but the barcode part remains the same.

How to Modify a Media Description

Modify a media description in the General property page for the medium. For detailed steps, refer to the online Help index keyword "modifying, media descriptions".

Using Cartridge Memory

With Cartridge Memory enabled LTO drive(s) used with Cartridge Memory enabled media, you can also update the medium description on the medium Cartridge Memory. This way, the description is not lost when you export or import the medium (it will be retrieved from the Cartridge Memory).

18266

Verifying Data on a Medium

What Is Verifying?

Verifying a medium shows whether the data on the medium is valid. It also updates the information about the medium in the IDB, such as medium condition.

Data Protector performs the following:

- Checks the Data Protector headers with information about the medium (medium ID, description, and location.)
- Reads all blocks on the medium.
- If the CRC (Cyclic Redundancy Check) option was used while writing to the medium, Data Protector recalculates the CRC and compares it to the one stored on the medium.

  If the CRC option was not used, and the verify operation passed, this means that all the data on the medium has been read. The medium did not cause a read error, so the hardware status of the tape is at the very least acceptable. This level of check can be viewed as partial.

  Additionally, if the CRC option was used, the backup data itself is consistent within each block. This level of check has a high level of reliability.

NOTE

Depending on the backup devices and media you use, this task can take a considerable amount of time to complete.

When to Verify Media

If errors were reported during backup, you can verify the medium to check whether the backup is usable.

How to Verify Data on a Medium

In the `Devices & Media` context, browse for a medium, right-click it, and click `Verify`. For detailed steps, refer to the online Help index keyword "verifying media".

17.265
Paulo

Scanning Media in a Device

What Is Scanning? You scan a device to update Data Protector information about the media in the device or library.

- In a standalone device, you scan a medium in a drive.
- In a library device, you scan media in the selected slots.
- With Cartridge Memory enabled drives, Data Protector can check the library inventory very quickly.

When to Scan the Device You have to scan the device when you change the location of media (enter, eject) manually without using the Data Protector commands. This creates inconsistencies with the information in the IDB, because Data Protector cannot track the actual location of the media.

Scanning loads media from all the selected slots into a drive, checks the format of media, displays the media header information, and updates the information about the repository in the IDB.

NOTE Depending on the number of selected slots, scanning may take a considerable amount of time. Data Protector has to load a medium from each slot into a drive and read the medium header with information about the medium.

How to Scan Media in a Device Scan media in a device by selecting the device and clicking `Scan` from the Actions menu. For detailed steps, refer to the online Help index keyword “scanning backup devices”.

If you are using a library device, you can scan several media in a single action. However, you can only use one drive. For detailed steps, refer to the online Help index keyword “scanning drives in library devices”.

Barcode Scan To scan a library with barcode support, use the `Barcode Scan` option. Data Protector only checks the barcode on the medium and updates the information in the IDB.

18.264
Paula

List Cartridge Memory

List Cartridge Memory, available for Cartridge Memory enabled drives, is equivalent to a standard Data Protector scan, with the difference that the information is retrieved from Cartridge Memory instead of tape. It does not require the loading or unloading of tape and is faster.

However, using this method to synchronize the repository with the IDB is not recommended. Use the standard scan instead. You can perform a Cartridge Memory list for a specific slot. Refer to Figure 4-6.

Figure 4-6 **List Cartridge Memory for Specific Slots**

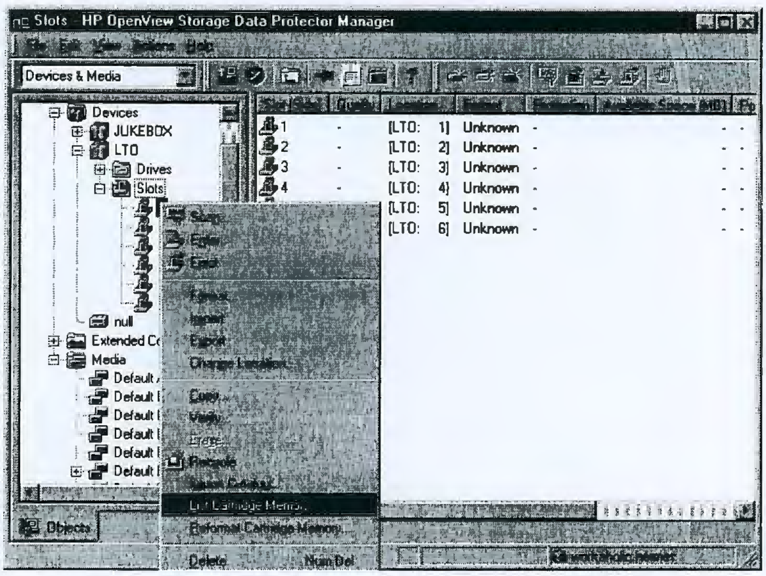

1	-	[LTO: 1]	Unknown -
2	-	[LTO: 2]	Unknown -
3	-	[LTO: 3]	Unknown -
4	-	[LTO: 4]	Unknown -
		[LTO: 5]	Unknown -
		[LTO: 6]	Unknown -

NOTE Certain media management operations (such as initialize, scan, enter, and eject) can be performed during backup or restore. Therefore, Data Protector uses "advisory locking" to ensure that backup or restore will not fail if the device cannot be locked. Other operations (such as backup, restore, import, and copy) require device locking to proceed with backup or restore.

Checking the Condition of a Medium

Data Protector allows you to view information about the usage and condition of a medium. The condition of the medium affects the ability to write to the medium and read the data contained on it. This helps you determine when the medium has to be replaced. See "Factors Influencing the Condition of Media" on page 132 for a description of when to change your media.

Use the `Info` property page of a medium to view information about the medium quality (condition). Refer to Figure 4-7.

Figure 4-7 **Information on Media**

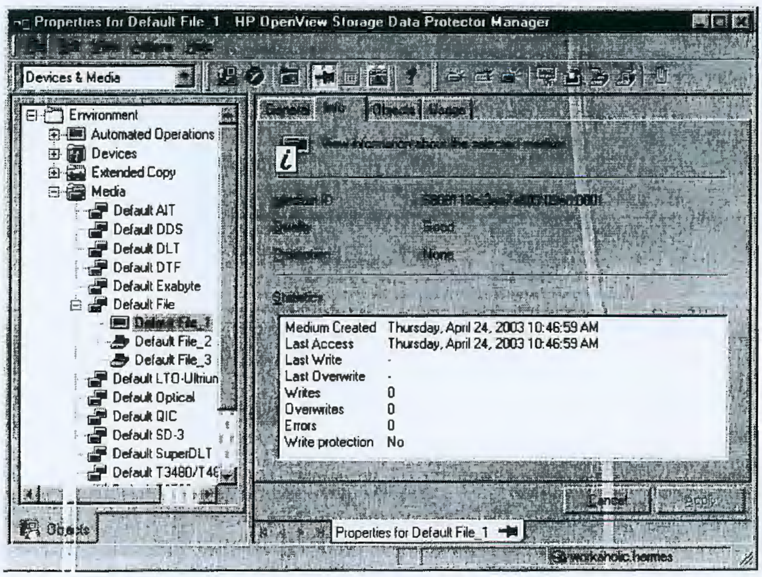

Quality: Good
Protection: None

```
Medium Created    Thursday, April 24, 2003 10:46:59 AM
Last Access       Thursday, April 24, 2003 10:46:59 AM
Last Write        -
Last Overwrite    -
Writes            0
Overwrites        0
Errors            0
Write protection  No
```

Selection of Backup Media

Media condition influences how media are selected for backup. Media in good condition are selected before media in fair condition. Media in poor condition are never selected. See "Selecting Media for Backup" on page 120 for details.

13.263
Paula

Cartridge Memory List

To view some additional information about Cartridge Memory enabled media, you can use the `Cartridge Memory List` feature. This lets you view the contents of Cartridge Memory for the medium.

Information stored in the IDB is *not* related to the information stored in the Cartridge Memory.

Factors Influencing the Condition of Media

Data Protector uses **media condition factors** to calculate the condition of the media. The condition of the media in a media pool determines the condition of the media pool. For example, as soon as one medium in a pool is `poor`, the whole media pool is `poor`. When media that are in poor condition are removed from the pool, the pool status reverts to either `fair` or `good` status.

The condition of a media pool indicates the reliability of that media pool for backups. For example, a backup to old or worn media is more likely to have read/write errors.

Media Condition Factors

The two media condition factors you can select are:

Medium valid for. The age of a medium is calculated as the number of months that have elapsed since the medium was formatted. Once a medium is older than the threshold number of months, it is marked as `poor`. The default threshold is 36 months.

Maximum number of overwrites. The usage of a medium is defined as the number of overwrites at the beginning of the medium. Once the medium has more than the threshold number of overwrites, it is marked as `poor`. The default threshold is 250 overwrites, except for DDS, which is set up with a default of 100 overwrites.

Device Error and Media Condition

If a device fails during backup, the media used for backup in this device are marked as `poor`. This prevents future errors if the problem was caused by the bad media.

If this error was due to a dirty drive, clean the drive and verify the medium to reset its condition.

It is recommended that you investigate if media marked `poor` appear in a pool. You can use `Verify` to get more information on each medium's condition. It is not recommended to simply recycle the medium.

Statuses of Media and Media Pools

Media or media pools can have three statuses, based on the media condition factors:

Good. Less than 80% of the threshold for age or usage.

Fair. 81 to 100% of the threshold for age or usage.

Poor. Exceeds 100% of the threshold for age or usage, or read/write errors have occurred on this medium.

See below for information on how to change the media condition factors.

Changing How Media Condition Is Calculated

When you add a medium to a media pool, you can define the media condition factors that are used to calculate the condition of the medium.

IMPORTANT For Data Protector to accurately calculate the condition of the media, use new media when adding media to the media pool.

Change the media condition factors using the `Condition` property page for the media pool. These condition factors are set for the entire media pool.

Figure 4-8 **The Media Condition Property Page**

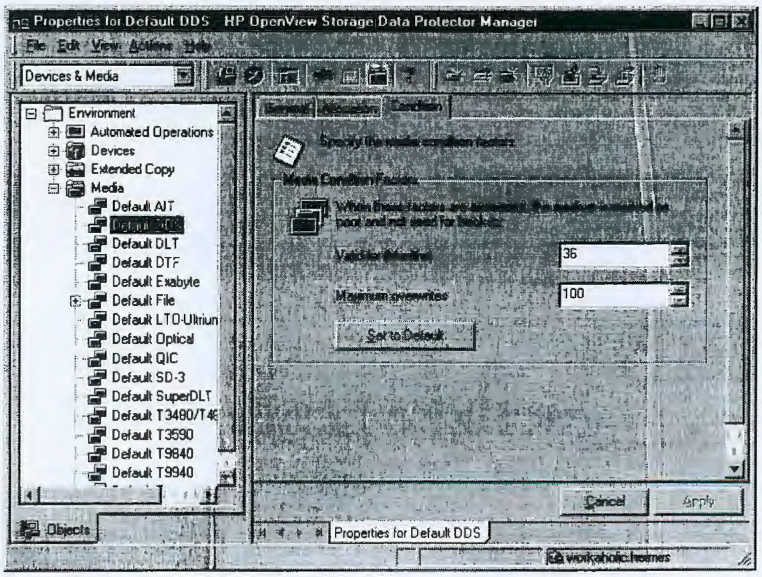

Properties for Default DDS - HP OpenView Storage Data Protector Manager

Specify the media condition factors

Media Condition Factors

When these factors are exceeded, the medium is marked as poor and not used for backup.

Valid for (Months) 36

Maximum overwrites 100

Searching for and Selecting a Medium

Use this function to locate and select specific media without having to browse through the entire list of media.

Media selection is especially useful for vaulting purposes, for example, selecting all media older than 14 days and moving them to a vault. See "Vaulting Media" on page 140 for more information

How to Search for and Select Media

In the `Devices & Media` context, browse for a media pool or a library device, right-click it, and click `Select Media`. For detailed steps, refer to the online Help index keyword "searching for media".

18.258
Paula

Entering a Medium into a Device

Data Protector allows you to physically enter media into a library device. You can select the slot that you want to use. Entering and ejecting media does not affect the media pool to which they belong.

IMPORTANT

It is recommended that you use Data Protector to handle the media in the device. This keeps the information about the media in the IDB up to date. If you enter media into the device manually using the device's controls, the information in the IDB is not consistent, and you have to scan the device to update this information. See "Scanning Media in a Device" on page 129 for instructions.

TIP

You can enter multiple media into a device in a single action. See the instructions below.

How to Enter Media into a Device

1. In the `Data Protector Manager`, switch to the `Devices & Media` context.
2. In the `Scoping Pane`, click `Devices`. The list of configured devices is displayed in the `Results Area`.
3. In the list of configured devices, click the name of the library, then expand it to display the `Drives` and `Slots` items.
4. Click `Slots` to display the list of slots.
5. Right-click the slot (or multiple slots) where you want to enter the media, and then click `Enter Medium`.

   A session starts that will prompt you to insert additional media into the device as needed.

What's Next?

If you want to add media to a media pool, see "Adding Media to a Media Pool" on page 107 for more information.

Ejecting a Medium from a Device

Data Protector allows you to physically eject media from the device. When used with library devices, media are moved to the specified slot. You can select the slot that you want to use.

IMPORTANT It is recommended that you use Data Protector to handle the media in the device. This keeps the information about the media in the IDB up to date. If you eject media from the device manually using the device's controls, the information in the IDB is not consistent, and you have to scan the device to update this information. See "Scanning Media in a Device" on page 129 for instructions.

Bulk Eject of Media You can eject multiple media from a library in a single action. Data Protector instructs you to remove media from a mail slot when the mail slot becomes full, to free up space for other media selected for ejection.

Predefined Eject of Media Some operations include the possibility of ejecting the media automatically when the session finishes. For example, when you copy media, you can specify whether the media will be ejected after the session.

When media cannot be ejected because the mail slot is full, Data Protector retries the operation until the mail slot becomes free or until the predefined time limit expires. During this retry, the robotics are accessible to other sessions.

During the eject execution, none of the specified media can be used by other sessions.

Limitation On Novell NetWare, Bulk Eject functionality is not supported.

How to Eject Media In the `Devices & Media` context, eject media by right-clicking a medium/slot (or multiple media/slots) and then clicking `Eject`. For detailed steps, refer to the online Help index keyword "ejecting media".

13.256

TIP Ejecting of media can be scheduled. Refer to "Scheduled Eject of Media" on page 138 for details.

What's Next? If you want to put media in a vault, see "Vaulting Media" on page 140 for more information.

Scheduled Eject of Media

Data Protector allows you to schedule the ejection of specific media through the reporting mechanism. The scheduled ejection of media is linked to a specific report made using the external send method. This method enables you to send the report to a user-definable external script, which can then parse the report and execute the ejection of media (using `omnimm -eject` command).

Prerequisite A program or script must be created on the Cell Manager to perform the ejection, and any applicable interpreters must also be installed on the Cell Manager. A Perl script is used in this example.

Overview You can set up and schedule a report group so that it creates a report and sends it as an input to a script. Such a report group should be set up so that it lists the media you want to eject (for example, the List of Media Report) by specifying the report parameters, so that the report contains only the media you want to eject. When the Report Group is started (as the result of a schedule or as triggered by a notification, for example the End of Session notification), Data Protector starts the script with the report result as an input for the script. The script then parses the report and performs the ejection of the specified media by using the Data Protector `omnimm` CLI command.

Notification on Mail Slots Full By default, the Event Log Viewer will notify you if you need to remove media from mail slots in order to continue the eject operation. This situation will arise when there are more media to be ejected than there are empty mail slots in a library. Refer to Chapter 7, "Monitoring, Reporting, Notifications, and the Event Log," on page 307 for more information on Data Protector notifications.

13255
Paula

If media are not removed from the mail slots after a default time span, and there are still media to be ejected, the `omnimm` command aborts the operation. You can change the default time span in the `.omnirc` file. Refer to “Using Omnirc Options” on page 525.

For an example of configuring scheduled ejection of media, refer to Appendix, “Example of Scheduled Eject of Media,” on page A-14.

Vaulting Media

What Is Vaulting?

Vaulting is a process of moving media with important information to a safe place where they are kept for a specified period of time. The safe place for media is often called a **vault**. This is also known as off-site storage.

Vaulting and Data Protector

Data Protector supports vaulting on various levels:

- Allows setting up of data protection and catalog protection policies.
- Allows easy selection and ejection of media from the library.
- The media location function tells you the physical location where the media are stored.
- A report shows media used for backup within a specified time frame.
- A report shows which backup specifications have used specified media during the backup.
- A report shows media stored at a specific location with data protection expiring at a specific time.
- Displays a list of media needed for a restore and the physical locations where the media are stored.
- Allows filtering of media from the media view based on specific criteria, such as time written to the media or media with expired protection.

Implementing Vaulting

How you implement vaulting depends on your company's backup strategy and policies for handling data and media. Generally, it consists of the following steps:

1. Specify the desired data protection and catalog protection policies when configuring the backup of data.
2. Configure a vault in Data Protector. Essentially, this means specifying a name for the vault that you will use for the media, such as Vault_1.
3. After a backup is done, copy the media, if desired. You can use either manually started or automated media copying. For more details, refer to "Copying Media" on page 143.

18.253
Paula

4. Select the media that you want to store in the vault, change the location of the media, eject the media, and store them in the vault.
5. Select the media that you want to remove from the vault, such as media with expired data protection. You can get a list of such media using the List of Media report. For how to generate this report, see "Running Individual Reports" on page 338.
6. Enter the media into the library, scan them, and then change the location field.
7. Establish the appropriate media maintenance policy for media in the vault.

Configuring Vaults

Data Protector allows you to create a list of pre-defined vault locations that you often use. This simplifies entering locations when you move media to the vault.

In the `Devices & Media` context, click `Locations` from the `Edit` menu. For detailed steps, refer to the online Help index keyword "configuring lists of vaults".

Moving Media to a Vault

Depending on your company's policies, you can move the original media to a vault directly, or you can create copies and move the copies.

Moving media to a vault consists of two steps:

1. Select media that you want to move and change the location for the media. See "Modifying Media Locations" on page 126.
2. Eject the media from the device and move them to the vault. See "Ejecting a Medium from a Device" on page 137.

Restoring from Media in a Vault

Restoring media from a vault is no different from restoring from any other media. Depending on how your data and catalog protection policies are defined, you may need to take some additional steps:

1. Identify the media needed for restore.

18252
Paula

2. Take the media from a vault, enter the media in the library, and scan them.
3. If the catalog protection for the media is still valid, restore data by selecting what you want to restore, using the Data Protector user interface.

   If the catalog protection for the media has expired, Data Protector may not have detailed information about the backed up data. You can restore by manually specifying the files or directories that you want to restore, or use the `List from media` functionality.

TIP

To re-read the detailed information about files and directories from the media once the catalog protection has expired, export the media and import them back, specifying that you want to read the detail catalog data. Now you will be able to browse files and directories in the Data Protector user interface again.

15 251
Paula

Copying Media

What Is Media Copying?

Data Protector enables copying of backed up data to a second set of media. You can move either the copies or the original media to a safe place for archiving/vaulting purposes, and keep the other set of media on site for restore purposes. For how to configure Data Protector for vaulting, see "Vaulting Media" on page 140. For more information on vaulting, refer to the *HP OpenView Storage Data Protector Concepts Guide*.

Besides manually started media copying, Data Protector also offers automated media copying. For more information, see "Automated Media Copying" on page 145.

How to Copy Media

In the `Devices & Media` context, browse for a medium, right-click it and click `Copy`. For detailed steps, refer to the online Help index keyword "copying media".

You need two devices with the same media type, one as a **source medium**, one as a **target medium**. A source medium is the medium being copied, while a target medium is the medium to which data is copied.

You can specify the protection period for the target medium, during which the data on the medium cannot be overwritten. The default protection is the same as for the original. Other options are Permanent and Until (specified date). A medium is protected until the end of the longest protection period of one of the objects on the medium.

You need to start the copying of each medium separately, as only one medium can be copied in a copying session. The copy operation is not available for media in free pools.

What Is the Result?

The result of copying media is that you have two sets of media with the same data, the original media and the copies.

After the source medium has been copied, Data Protector marks it as non-appendable to prevent appending new backups. (This would result in the original being different from its copy.) The copy is also marked as non-appendable.

You can make multiple copies of the original media. You cannot, however, make copies of copies, also known as second generation copies.

NOTE

When copying media, it is possible that the target medium reaches the end of the tape before the source medium. This may happen if the source medium was written in streaming mode and you make a copy on a busy system or through a loaded network, which can create blank space where the tape has stopped and started again. You can prevent this by enabling tape padding when you format media. See "Formatting Media" on page 108.

Moving Copies

Typically, you want to move the copies of the media to a safe place. See "Vaulting Media" on page 140 and "Ejecting a Medium from a Device" on page 137 for more information.

Exporting Copies

Exporting a medium removes all information regarding this medium from the IDB. If you export the original medium, but one or more copies of the medium exist, one of the copies becomes the original.

If you try to import the removed copy, but the original media are not in the IDB, you have to import these media using the `force` option. See "Importing Media" on page 113 for instructions.

Restoring from a Copy

When you restore data, Data Protector prefers restoring from the original media. However, if the original media are not available, but a copy is available, the copy will be used for the restore.

If neither the original nor a copy is available in the device during restore, Data Protector issues a mount request, displaying both the original and the copy as the media required for restore. You can use any one of these.

If you perform a restore using a standalone device, you can choose to restore from the copy rather than from the original. To do this, insert the copy in the device that will be used for the restore, or select the device containing the copy. However, if you perform a restore using a library device and the original is in the library, Data Protector will use it for the restore.

For detailed instructions on how to restore data from the media archive, see "Vaulting Media" on page 140.

Automated Media Copying

What Is Automated Media Copying?

Automated media copying is an automated process that creates copies of the media containing backups.

Data Protector offers two types of automated media copying: **post-backup media copying** and **scheduled media copying**.

What Is Post-Backup Media Copying?

Post-backup media copying takes place after the completion of a backup session. It copies all media used in that particular session.

Configuring Post-Backup Media Copying

In the `Devices & Media` context, right-click `Automated Operations` and click `Add Post-Backup Media Operation`. For detailed steps, refer to the online Help index keyword "post-backup media copying".

What Is Scheduled Media Copying?

Scheduled media copying takes place at a user-defined time. Media used in different backup specifications can be copied in a single session. You create an automated media copy specification to define which media will be copied.

Configuring Scheduled Media Copying

In the `Devices & Media` context, right-click `Automated Operations` and click `Add Scheduled Media Operation`. For detailed steps, refer to the online Help index keyword "scheduled media copying".

You can configure scheduled media copying to run on specific dates at specific times, or to run periodically. You can reset, disable, or enable a schedule, and disable or enable automated media copying on holidays. For details, refer to the online Help index keyword "automated media copying".

Limitations

- You cannot use standalone devices for automated media copying; only library devices can be used.
- The source medium and the target medium must be of the same type.
- You cannot copy NDMP media.

How Does Automated Media Copying Operate?

First you create an automated media copy specification. When the automated media copy session begins, Data Protector generates a list of media, referred to as **source media**, based on the parameters specified in the automated media copy specification. For each source medium, a

18.248
Paula

target medium is selected to which the data will be copied. The target media are selected from the same media pool as the source media, from a free pool, or from the blank media in a library.

Selection and Use of Devices

For each source medium, Data Protector selects a pair of devices from the devices that you specified in the automated media copy specification. The automated media copy functionality provides its own balancing. Data Protector tries to make optimum use of the available devices by using as many devices as possible and selecting local devices if they are available.

Devices are locked at the beginning of the session. The devices that are not available at that time cannot be used in the session, as device locking after the beginning of the session is not possible. Note that at least a pair of devices must be available for each media type for the entire session to complete successfully. If the minimum number of devices necessary for the session cannot be locked, the session fails.

If a media error occurs, the device with errors will be avoided within that automated media copy session. However, if there are no other devices available, it will be reused.

Destination Pool of the Copies

The source medium defines the destination pool of the target medium. This means that the copied media will belong to the same pool as the original media.

Data Protection of the Copies

The default protection period for the copy is the same as the protection for the original. You can set a different protection period when creating or modifying the automated media copy specification.

Mount and Cleanme Request Handling

The automated media copy functionality does not handle mount or cleanme requests. If a mount request is received, the media pair concerned is aborted, but the session continues. You can manually copy the media that were not copied after the automated media copy session finishes.

For examples of use, refer to the *HP OpenView Storage Data Protector Concepts Guide*.

18.2141
Paula

Detection of Write-Protected Media

Data Protector can detect and handle media that has been mechanically protected by setting the write protection switch on.

NOTE It is recommended not to use write-protected media with Data Protector.

The following operations can detect and handle write-protected media:

- Read-only operations, such as: list, scan, and verify.

  Read-only operations detect the write-protected media and proceed without any warnings.

- Write operations, such as: initialize, erase, and backup.

  Write operations detect the write-protected media and either abort the session or skip the write-protected media. Backup sessions treat write-protected media as unusable media and behave according to the media allocation policy. If the allocation policy is `strict`, a mount request is issued. If the allocation policy is `loose`, the medium is skipped.

The detection of a write-protected medium and all changes to the write-protection state of the medium are logged to the `media.log` file.

Using Different Media Format Types

Data Protector recognizes and uses two different format types to write data to media:

- Data Protector (for backup devices that are under direct Data Protector control)
- NDMP (for backup devices that are connected to NDMP servers)

Both format types use different Data Protector Media Agent components to communicate with backup devices.

Limitations

Take into account the following limitations, when using different media format types:

- Media that are written by one format type will be recognized as blank or as foreign in a backup device that uses a different format type.
- You cannot back up objects using different format types on the same medium.
- You cannot have two different Data Protector Media Agent components installed on the same system.
- It is strongly recommended that you use different media pools for different media format types.

Modifying Views in the Media Management Window

You can customize the information you see about the media in the Media Management window. This enables you to always see the information you need.

To customize your view, do the following:

1. Open the global options file.

   On the UNIX Cell Manager:

   `/etc/opt/omni/options/global`

   On the Windows Cell Manager:

   `<Data_Protector_home>\config\options\Global`

2. Customize the attributes that are to be displayed in the library or media management view by specifying the corresponding token strings.

18.244
Bauka

ANEXO UNIDADE DE BACKUP ROBOTIZADO
PARTE D/3

18.243
Paula

18 [illegible]
[illegible]

5 Backup

In This Chapter

This chapter explains how to back up your data. It also describes some advanced Data Protector features.

"Configuring a Backup" on page 153

"Backing Up UNIX Systems" on page 161

"Backing Up Windows Systems" on page 168

"Backing Up Novell NetWare Systems" on page 194

"Backing Up OpenVMS Systems" on page 201

"Backing Up in a Direct Backup Environment" on page 204

"Scheduling Unattended Backups" on page 207

"Selecting a Backup Type: Full or Incremental" on page 213

"Using Backup Templates" on page 216

"Groups of Backup Specifications" on page 222

"Using Backup Options" on page 225

"Pre- and Post-Exec Commands" on page 250

"Managing Failed Backups" on page 263

For information on how to back up database applications such as Oracle, SAP R/3, MS Exchange, MS SQL, Informix, IBM DB2 UDB or Sybase, refer to the *HP OpenView Storage Data Protector Integration Guide*.

For information on how to back up the Data Protector internal database (IDB), see "Configuring the Database Backup" on page 398.

For information on how to install and configure Data Protector management applications, see Chapter 13, "Integrations with Other Applications," on page 611.

NOTE

Backup devices (such as tape drives) are subject to specific Data Protector licenses. See the *HP OpenView Storage Data Protector Installation and Licensing Guide* for details.

Configuring a Backup

A backup is a process that creates a copy of system data on backup media. This copy is stored and kept for future use in case the original is destroyed or corrupted.

Prerequisites

- You need to have a Disk Agent installed on every system that is to be backed up, unless you use NFS (on UNIX) or Network Share Backup (on Windows) for backing up these systems.
- You need to have at least one backup device configured in the Data Protector cell.
- You need to have media prepared for your backup.
- You need to have appropriate user rights for performing a backup.

Backup Configuration

Configuring a backup consists of the following steps:

1. Selecting what to back up - the data sources on the Disk Agent clients.
2. Selecting where to back up to - the backup devices connected to the Media Agent clients.
3. Selecting how to back up - backup options.
4. Optionally, you can schedule an unattended backup.

You specify these options when creating a **backup specification**. Refer to "Creating a Backup Specification" on page 154.

At a specified time, Data Protector starts the backup session based on the backup specification. A **backup object** is any data selected for a backup, such as a disk, a file, a directory, a database, or a part of the database. During the backup session, Data Protector reads the objects, transfers data through the network, and writes them to the media residing in the devices.

The backup specification defines the devices to be used and, optionally, the media pools. If no media pool is specified, the default media pool, which is a part of the device specification, is used.

A backup specification can be as simple as backing up one disk to a standalone DDS drive, or as complex as specifying a backup for 40 large servers to a tape library with 8 drives.

18.231
Paula

A **backup session** is based on the backup specification, and can be started interactively. During the backup session, Data Protector reads the backup objects, transfers their data through the network, and writes them to the media residing in the devices.

Figure 5-1 **Backup Session**

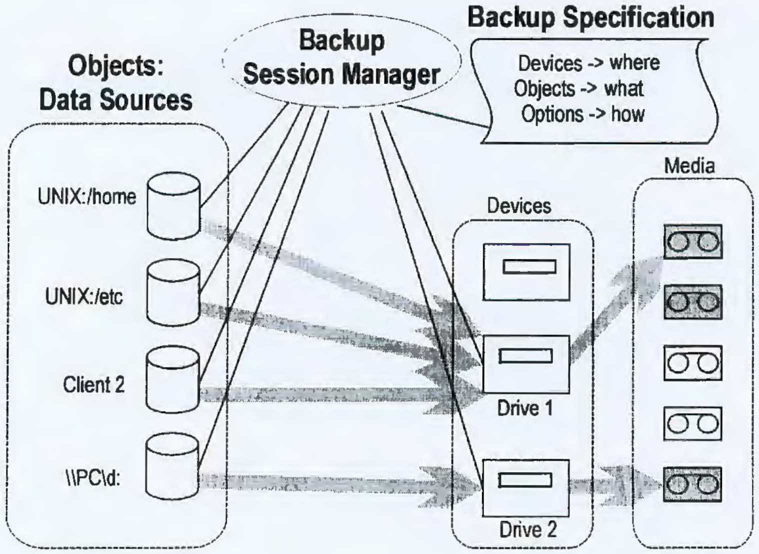

Creating a Backup Specification

You can configure a backup specification using the Data Protector user interface. A backup specification defines the client systems, drives, directories, and files to be backed up, the devices or drives to be used, the backup options for all objects in the specification, and the days and times that you want backups to be performed.

You can create multiple backup specifications by copying an existing specification and then modifying one of the copies.

Data Protector provides default options that are suitable for most cases. To customize the behavior, use Data Protector backup options.

Keep the following key points in mind when you run a backup session:

Key Points

- The backup type (full or incremental) is the same for the whole backup session. All data in a group is backed up using the same backup type.
- A backup object can be added to multiple backup specifications. For example, you may have one backup specification for full backups, one for incremental backups, one for a departmental backup, and one for the archive backup. You can give a description for each object. It is important that you choose the description carefully, because this lets you differentiate among various backups from the same filesystem.
- Objects or clients can be grouped into one backup specification if the media and the backups are managed in the same way, or if media are put into a vault.
- If many backup specifications exist or are planned, you should structure them in groups of backup specifications. If the groups are structured along common option settings (how to back up), then you can apply the backup templates efficiently.
- The Data Protector GUI can display a limited number of backup specifications. The number of backup specifications depends on the size of their parameters (name, group, ownership information and information about whether the backup specification is load balanced or not). This size should not exceed 80 Kb.

Example of Creating a Backup Specification

The following example shows how to create a backup specification for a filesystem and how to start the backup interactively.

1. In the `HP OpenView Storage Data Protector Manager` window, switch to the `Backup` context.
2. In the Scoping Pane, expand `Backup`, and then double-click `Backup Specifications`.
3. In the Results Area, right-click `Filesystem`, and then click `Add Backup`. The `Create New Backup` dialog box appears.
4. In the Create New Backup dialog box, select the `Blank Filesystem Backup` template, and then click `OK` to start the Backup wizard. See Figure 5-2 on page 156.

13.237
Paula

Figure 5-2

Create New Backup Dialog Box

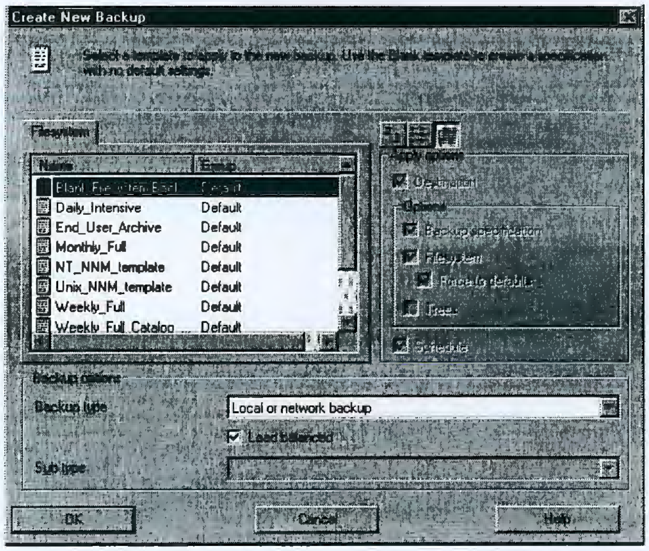

Create New Backup

Select a template to apply to the new backup. Use the Blank template to create a specification with no default settings.

Filesystem

Name	Group
Blank Filesystem Backup	Default
Daily_Intensive	Default
End_User_Archive	Default
Monthly_Full	Default
NT_NNM_template	Default
Unix_NNM_template	Default
Weekly_Full	Default
Weekly_Full_Catalog ...	Default

Apply options
☑ Destination
Options
☑ Backup specification
☑ Filesystem
☑ Force to defaults
□ Trees
☑ Schedule

Backup options
Backup type: Local or network backup
☑ Load balanced
Sub type:

5. Select what you want to back up. Figure 5-3 on page 157 shows data sources selected for backup. Click Next to proceed.

Figure 5-3 **Source Page of the Backup Wizard**

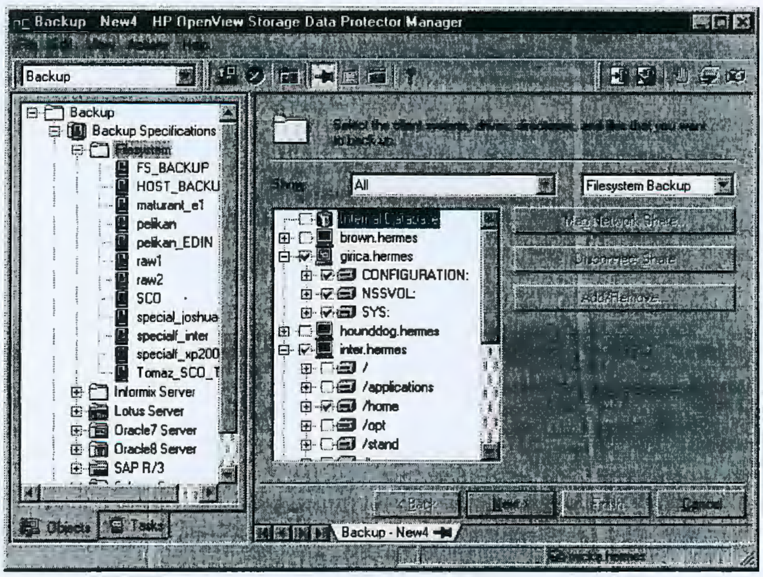

Select the client systems, drives, directories, and files that you want to back up.

6. Select the device(s) that will be used to back up your data. See Figure 5-4 on page 158. Click Next to proceed.

Figure 5-4 **Device Page of the Backup Wizard**

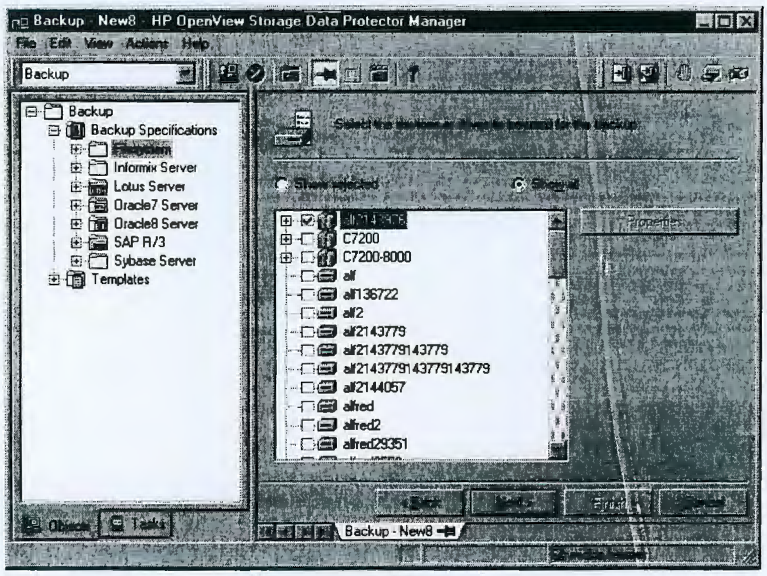

7. Select backup options. See "Using Backup Options" on page 225 for details. Click `Next`.

8. In the `Schedule` page, you can schedule the backup. See "Scheduling Unattended Backups" on page 207 for more information. Click `Next`.

9. In the `Backup Object Summary` page, you can review the backup options. See Figure 5-5 on page 159. Click `Next`.

15:2?4
Paula

Figure 5-5 **Backup Object Summary Page**

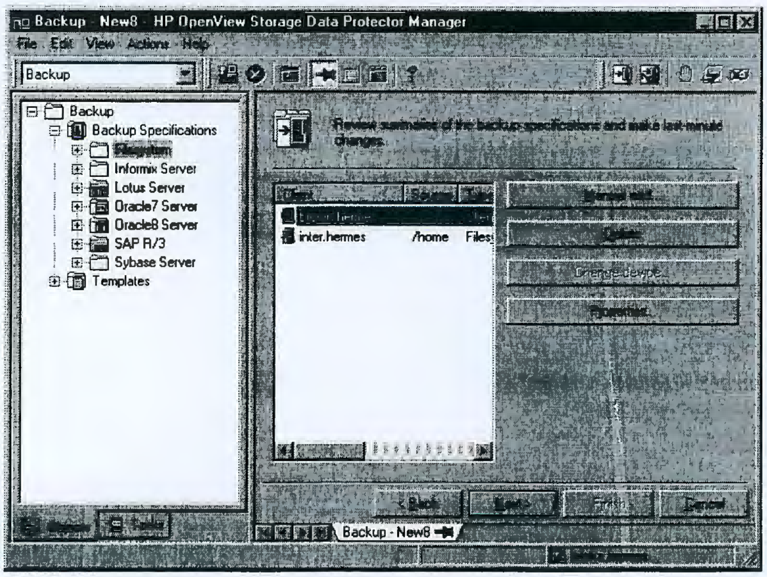

Review summaries of the backup specifications and make last-minute changes.

10. In the final page of the Backup wizard, you can save the backup specification, start the interactive backup, or preview the backup. See Figure 5-6 on page 160.

    It is recommended to save the backup specification so that you can schedule or modify it later.

15 233
Paula

Figure 5-6 **Final Page of the Backup Wizard**

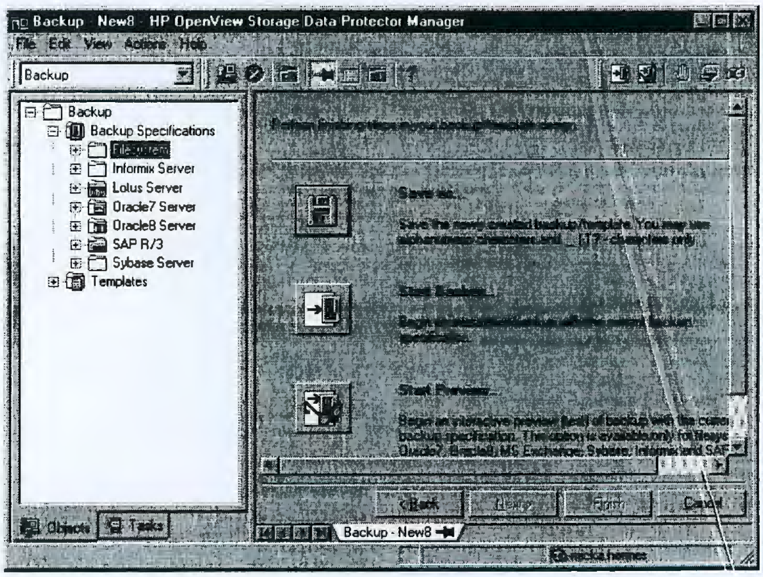

Save the newly created backup/template. You may use alphanumeric characters and _ [17 characters only]

Begin an interactive preview (test) of backup with the current backup specification. This option is available only for filesys Oracle7, Oracle8, MS Exchange, Sybase, Informix and SAP

11. Click Start Backup to run the backup interactively. The Start Backup dialog box appears.

NOTE

During a backup, you may be prompted to add more media to continue your backup. This is called a mount request. See “Responding to Mount Requests” on page 310 for more detailed information.

Backing Up UNIX Systems

You can install a Disk Agent on every UNIX system in order to back it up. Alternatively, you may use the Network Filesystem (NFS) to back up data from systems that do not have a Disk Agent.

See “Backing Up Disks Using NFS” on page 164 for details.

See the *HP OpenView Storage Data Protector Installation and Licensing Guide* or online Help for instructions on how to install a Disk Agent.

See the *HP OpenView Storage Data Protector Software Release Notes* for a complete list of supported platforms.

Backing Up UNIX Filesystems

Limitations

The maximum size of the files you can back up depends on operating system and filesystem limitations. Data Protector has no file size limitations on the following UNIX systems: HP-UX, Solaris, AIX, IRIX, and Linux. On other UNIX systems Data Protector backs up files of up to 2 GB.

Data Protector backs up the directory structure, regular files, and special files. Special files are character device files, block device files, UNIX domain sockets, FIFO files, HP-UX network special files, and XENIX specially-named files.

Softlinks and mountpoints are not followed, and are backed up as softlinks and ordinary empty directories, respectively.

If there are multiple hardlinks referencing the same file, the file is backed up only once. You can change this by setting the `Backup hardlinks as files` option, as explained in “List of Data Protector Backup Options” on page 236.

All file attributes, including file permissions, access times, and Access Control Lists (ACLs) on HP-UX and AIX are backed up together with the files. The time of the last access to each file is saved before reading the file and then returned to the original value after the file is backed up. This behavior can be changed by setting the `Do not preserve access time attributes` option, as explained in “Using Backup Options” on page 225.

Network share backup is not supported.

18.231
Paula

Data Protector provides a sophisticated mechanism for incremental backups. To determine which files have changed, the Data Protector Disk Agent checks when each was last modified. This method keeps Data Protector from detecting moved files, as moving the file does not change the modification time.

NOTE

During a backup session, each file being backed up is opened and read. Therefore, the access time of the file is changed after the backup. Unless the `Do not preserve access time attributes` backup option is set, the access time attribute is set to its original value. OFF is the default value. If this option is set, moved files on UNIX clients are included in the incremental backup, because detection is based on the `inode` modification time.

Selecting Specific Files or Directories

For each filesystem, you can restrict the backup to specific directory trees. For each directory tree you can:

- Exclude any sub-tree or file
- Back up files that match a specific wildcard pattern
- Skip files that match a specific wildcard pattern

Some files are permanently in use, for example, by database applications. These files should be excluded from ordinary filesystem backup and should be backed up in a special way. This is also true for the IDB itself.

Therefore, exclude the IDB directories `/var/opt/omni/db` and `/etc/opt/omni` on UNIX Cell Managers from standard filesystem backups to ensure the consistency of data.

For detailed information on how to back up the IDB, see “Configuring the Database Backup” on page 398.

You should also exclude temporary directories.

How to Back Up UNIX Files

Back up UNIX files using the procedure described in “Example of Creating a Backup Specification” on page 155.

See also “Using Backup Options” on page 225 for information on using and structuring your backup options.

Backing Up Clients Using Disk Discovery

How Are Disks Discovered?

If you specify a client backup with **disk discovery**, Data Protector contacts the client at backup time and finds all filesystems on the disks that are attached to that system. Only mounted disks are identified using the `mount` command. Then Data Protector backs up each filesystem identified as a regular filesystem, except for NFS, CD mounted filesystems, and removable volumes. The description for each filesystem object is generated and the filesystem mountpoint is appended to the description of the client backup.

When to Use Disk Discovery

This backup type is recommended under the following conditions:

- If you back up workstations with relatively small disks that are frequently mounted or unmounted.
- If you would like to back up the data following a mountpoint into one directory, regardless of how many filesystems are mounted. For example, `/home/data`, where `/home/data/disk1 and /home/data/newdisk/disk2` can be mounted or unmounted frequently and independently of each other.

You can use disk discovery by specifying the client as a data source. If another disk is mounted later, it will be included in the backup.

In contrast to a filesystem backup, where you have to specify any newly added disk or mounted filesystem that is not yet specified in the backup specification, this is unnecessary if you use disk discovery.

To create a backup specification that will define a disk discovery backup, follow the procedure described in "Example of Creating a Backup Specification" on page 155.

Once you get to the `Source` property page of the Backup wizard, click the check box next to the client. This selects the entire client to be backed up, as shown in Figure 5-7.

18.229
Paula

Figure 5-7 **Selecting an Entire Client to Be Backed Up**

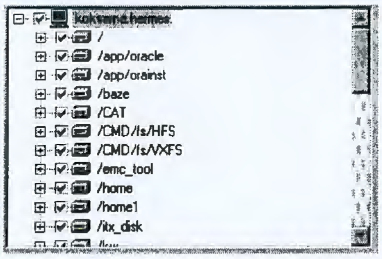

☑ kokvena.hermes
- ☑ /
- ☑ /app/oracle
- ☑ /app/orainst
- ☑ /baze
- ☑ /CAT
- ☑ /CMD/ls/HFS
- ☑ /CMD/ls/VXFS
- ☑ /emc_tool
- ☑ /home
- ☑ /home1
- ☑ /itx_disk

NOTE

Selecting all of the client's drives is not the same as selecting the check box next to the client name, which is the procedure for a Disk Discovery backup.

When you perform a client backup, all the files and directories that belong to the root (/) mountpoint are automatically backed up. Therefore, you cannot exclude the root in the backup specification. If you want to exclude the root, perform a filesystem backup.

To check the configured backup type, see the `Backup Object Summary` property page. Under the `Type` label, you will see `Client System` if you have configured a Disk Discovery backup and `Filesystem` if only the drives have been selected.

Also see "Using Backup Options" on page 225 for information on structuring your backup specifications.

Backing Up Disks Using NFS

What Is NFS?

NFS (Network Filesystem) is a communication protocol that allows a computer to access files over a network as though they were on its local disks.

Figure 5-8 shows a typical configuration. You want to back up the filesystem `/home/special` from system `wiz`, which is not part of the Data Protector cell and has no Data Protector software installed. But the filesystem is mounted as `/nfs/wiz/home/special` on a Data Protector client `bserv`.

18.228
Paula

To back up this filesystem using NFS, follow the same procedure as if you were backing up any other filesystem on `bserv`, except that you have to manually type `/nfs/wiz/home/special` as a mountpoint. Only local filesystems can be browsed.

Figure 5-8 **NFS Environment**

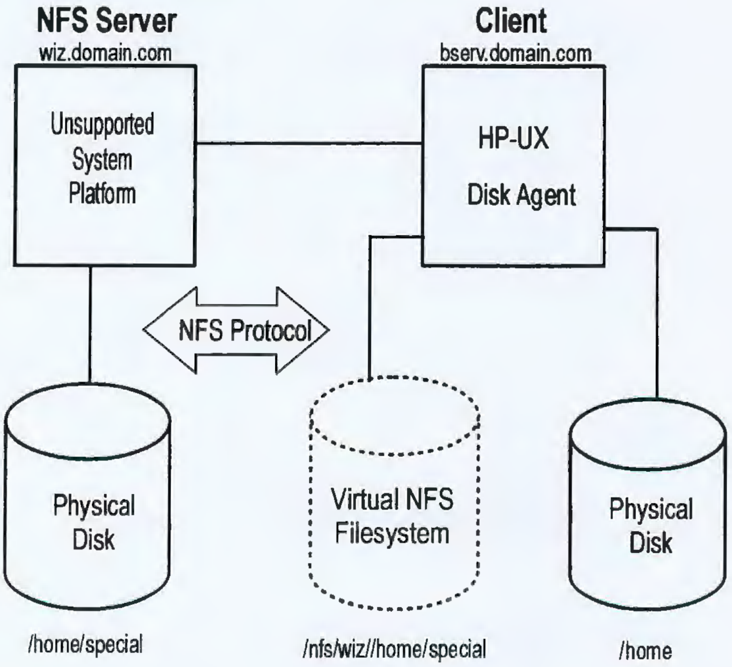

Limitations

- You can use NFS only if backing up files on HP-UX clients. You can not back up soft links, or character and device files.
- ACL attributes are not preserved. NFS does not support ACLs on remote files. Individual manual entries specify the behavior of various system calls, library calls, and commands. When transferring a file with optional entries over the network or manipulating a remote file, the optional entries may be unexpectedly deleted.

13.227
Paula

NOTE It is recommended to have root permission on mounted NFS filesystems.

When to Use NFS Backup

Use NFS backup in either of the following situations:

- A system to be backed up is not a part of the Data Protector cell.
- You want to back up system platforms that are not supported by Data Protector.

To back up a filesystem using NFS, follow the procedure described in "Example of Creating a Backup Specification" on page 155 until you get to the `Backup Object Summary` page of the wizard. Proceed as follows:

1. In the `Backup Object Summary` page, click `Manual Add`.
2. Click the `UNIX Filesystem` button, and then click `Next`.
3. In the `General Selection` page, select a client and manually add the mount point in the `Mountpoint` text box. See online Help for details.

Backing Up UNIX Disks as Disk Image Objects

What Is a Disk Image Backup?

A **disk image backup** is a high-speed backup of disks, disk partitions, or logical volumes without tracking the file and directory structure stored on these data sources. Data Protector stores the disk image structure at the character level.

When to Use a Disk Image Backup

Use a disk image backup in any of the following situations:

- You have lots of small files and a high backup speed is required.
- A full disk backup is needed, for example, for disaster recovery or before a major software update.
- A direct disk-to-disk connection is not possible and you want to duplicate a filesystem to another disk. The latter must be identical to the original disk.

Where to Find Rawdisk Sections

On the HP-UX and Solaris systems, the rawdisk sections are usually listed in the `/dev/rdsk` directory. On HP-UX, raw logical volumes can be found in `/dev/vg<XX>`. The first letter of the new logical volume must be `r`, for instance `/dev/vg01/rlvol2`.

13226
Paula

IMPORTANT Unmount a disk before a disk image backup and mount it later. You can use `pre-` and `post-` `exec` commands for this purpose. See Appendix, "Examples of Pre-Exec and Post-Exec Commands for UNIX," on page A-20.

To back up a disk image object, follow the procedure described in "Example of Creating a Backup Specification" on page 155 until you get to the `Backup Object Summary` page of the wizard. Proceed as follows:

1. In the `Backup Object Summary` page, click `Manual Add`.
2. Click the `Disk image object` button, and then click `Next`.
3. In the `General Selection` page, select a client and manually add the mount point in the `Mountpoint` text box. See online Help for details.

18.225
Paula

Backing Up Windows Systems

Prerequisites

You have to install a Disk Agent on at least one Windows computer in the Data Protector cell. This computer then becomes a Disk Agent client.

Files that do not reside on Disk Agent clients can be backed up if they share their disks with Disk Agent clients. It is better to install a Disk Agent on every Windows system that you want to back up.

See "Backing Up Windows Shared Disks" on page 185 for details.

See the *HP OpenView Storage Data Protector Installation and Licensing Guide* or online Help for instructions on how to install a Disk Agent.

See the *HP OpenView Storage Data Protector Software Release Notes* for a complete list of supported system platforms.

Limitations

- Files of up to 128 GB can be backed up on NTFS. Moved files cannot be detected during an incremental backup.
- To run a VSS filesystem backup, your system must have at least one NTFS filesystem.

Backing Up Filesystems (Logical Disk Drives)

Selecting Backup Objects

Select a file, a directory, or a logical disk drive for backup in the Backup wizard.

See "Example of Creating a Backup Specification" on page 155 and "Using Backup Options" on page 225 for details.

What Is Backed Up?

A filesystem backup of a disk drive involves reading the directory structure and the contents of the files on the selected disk drive. The following data is also backed up along with the data in the file:

- Full Unicode filenames
- FAT16, FAT32, VFAT, and NTFS attributes

  Once a file is backed up, its archive attribute is cleared. You can change this behavior by setting the `Do not use archive attribute` option among the `Advanced` filesystem backup options in the backup specification. See online Help for details.

18 2221
Paula

- NTFS alternate data streams
- NTFS security data

NOTE The sharing properties of a folder are not stored in the filesystem and are not backed up within filesystem backup. Information about shares is stored in the registry and is backed up and restored within CONFIGURATION backup object.

What Is Not Backed Up?

In the backup specification, you can specify the files to be excluded from or skipped by the backup. The list of these files is also known as a **private exclusion list.**

See "Object Options" on page 239 and online Help for more information on how to exclude or skip files and directories.

In addition to the private exclusion list, Data Protector by default excludes the following:

- The *<Data_Protector_home>*\log and *<Data_Protector_home>*\tmp directories from a Windows client or Cell Manager backup.
- The *<Data_Protector_home>*\db40 directory from a Windows Cell Manager backup.

  For example, the *<Data_Protector_home>*\db40 directory is excluded from the Cell Manager backup even though it was selected in the backup specification. This is because the *<Data_Protector_home>*\db40 directory contains the IDB, which must be backed up in a special way to ensure data consistency. See "Configuring the Database Backup" on page 398 for details.

The skipped file is the `Pagefile.sys` system file. Before starting a backup, Data Protector reads the list of excluded and skipped files from the following Registry keys:

```
HKEY_LOCAL_MACHINE\SOFTWARE\Hewlett-Packard\OpenView
\\OmniBack II\Agents\FileSystem\Exclude
```

```
HKEY_LOCAL_MACHINE\SOFTWARE\Hewlett-Packard\OpenView
\\OmniBack II\Agents\FileSystem\Skip
```

19.223
Paula

NTFS 3.x Filesystem Features

The NTFS 3.x filesystem has introduced new file attributes and concepts, which can be summarized as follows:

- The NTFS 3.x filesystem supports **reparse points.** The **volume mount points**, **Single Instance Storage (SIS)**, and **directory junctions** are based on the reparse point concept. See "Glossary" for details.
- The NTFS 3.x filesystem supports **sparse files** as an efficient way of reducing the amount of allocated disk space.
- The NTFS 3.x filesystem supports the **Object IDs** that are backed up by Data Protector along with other alternate data streams.
- Some of the NTFS 3.x filesystem-specific features are controlled by system services that maintain their own data records. These data structures are backed up as a part of CONFIGURATION.

  See "Backing Up CONFIGURATION" on page 173 and "Backing Up the Windows 2000/XP/Server 2003 Services" on page 179 for details.
- The Microsoft-encrypted NTFS 3.x files are backed up and restored encrypted, but their contents can only be properly viewed when they are decrypted. Refer to the *HP OpenView Storage Data Protector Software Release Notes* for details about related limitations.

VSS Filesystem Backup

Volume Shadow Copy service (VSS) is implemented on the Windows Server 2003 operating system. This service provides an additional Windows filesystem backup, where the level of data integrity is slightly increased compared to traditional backup of active volume.

To prepare for creation of the shadow copy, all I/O activity is stopped by the VSS mechanism. When the shadow copy is created, Data Protector starts its normal backup procedure, except that the source volume is replaced by the newly created shadow copy. If the shadow copy creation fails, Data Protector can proceed with the normal filesystem backup, if the `Allow Fallback` option was specified in the backup specification.

During the VSS filesystem backup the consistency of data is improved in comparison with the non-VSS filesystem backup. VSS allows you to create shadow copy backups of volumes and exact point-in-time copies of files, including all open files. In this way, the files changed during the backup are copied correctly.

18.322
Paula

The advantages of the VSS filesystem backup are the following:

- A computer can be backed up while applications and services are running. Therefore, the applications can continue writing data to the volume during a backup.
- Open files are no longer skipped during the backup process, because they appear closed on the shadow copy volume at the time of the shadow copy creation.
- Backups can be performed at any time without locking out users.
- There is little or no impact on performance of the application system during the backup process.

For VSS filesystem backup related options, refer to "Using Backup Options" on page 225. Also refer to the *HP OpenView Storage Data Protector Concepts Guide* for details on the VSS concepts.

Reparse Points

Basically, reparse points are plain filesystem objects with a unique tag attached, known as a reparse point ID. The NTFS 3.x directories or files can contain a reparse point, which typically imitates the contents by directing to data from another location.

When Data Protector encounters reparse points, the reparse point IDs are not followed by default, what is also known as backing up raw reparse points. This affects the way you configure your backups:

✓ If you configure a backup using Disk Delivery, all data will be backed up once.

✓ If you back up filesystems or drives containing reparse points, ensure that the data pointed to by a reparse point gets backed up. For example, the Windows 2000/XP/Server 2003 **directory junctions** reparse points are not followed, so the junctions have to be backed up separately. SIS reparse points are exceptions.

The **Single Instance Storage (SIS)** service regularly checks the files on a disk. If the service detects several identical files, it replaces them with the reparse points and stores the data into a common repository. In this way, the disk space usage is reduced.

Reparse points let you mount logical volumes as disk drives. Data Protector treats the mounted volumes as though they were ordinary drives, so that they are visible as selectable objects for backup.

18.2021
Paula

Sparse Files

Sparse files contain many zero data sets as opposed to, for example, compressed files. At backup time, Data Protector automatically skips zero-parts, so that the media space on the backup device is allocated for non-zero parts only.

UNIX and Windows sparse files are not compatible.

Manual Definition of Multiple Disk Agents

If you want to back up one mount point through multiple Disk Agents (DA), you have to specify each object separately using the `Manual add` functionality. Give a new description to each object and use the `Trees/Exclude` option in the `Manual add` wizard to specify the path for an object. Refer to Figure 5-9.

In addition, consider the following:

- You have to manually define the data area split, taking care to avoid overlapping the same data.
- If more than one DA is concurrently accessing the same mount point, which is defined as one disk, the data transfer speed will drop. This can be different when using disk arrays.

18.220
Paula

Figure 5-9 **Specifying Objects Using Manual Add**

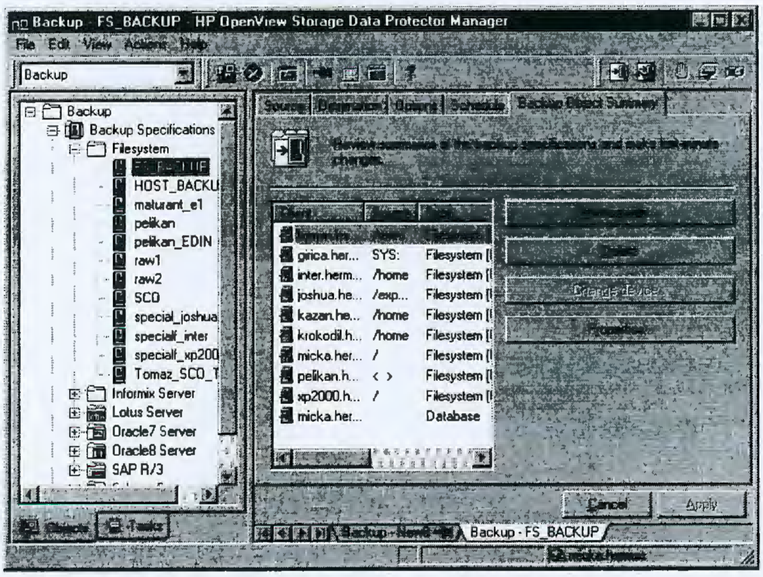

girica.her...	SYS:	Filesystem [
inter.herm...	/home	Filesystem [
joshua.he...	/exp...	Filesystem [
kazan.he...	/home	Filesystem [
krokodil.h...	/home	Filesystem [
micka.her...	/	Filesystem [
pelikan.h...	< >	Filesystem [
xp2000.h...	/	Filesystem [
micka.her...		Database

For detailed steps, refer to the online Help index keyword “concurrency”.

Backing Up CONFIGURATION

The Data Protector `CONFIGURATION` object is a set of data structures maintained by the Windows operating system that are not treated as a part of a filesystem backup when you, for example, select logical drives such as `C:` or `D:` for the backup.

Windows NT CONFIGURATION

`CONFIGURATION` consists of the following objects:

- `EventLog`
- `Profiles`
- `Registry`
- `SystemRecoveryData`
- `EISA Utility Partition`
- `WINS`, `DHCP` (on the Windows NT TCP/IP protocol servers)

18.219

Figure 5-10 **Windows NT CONFIGURATION**

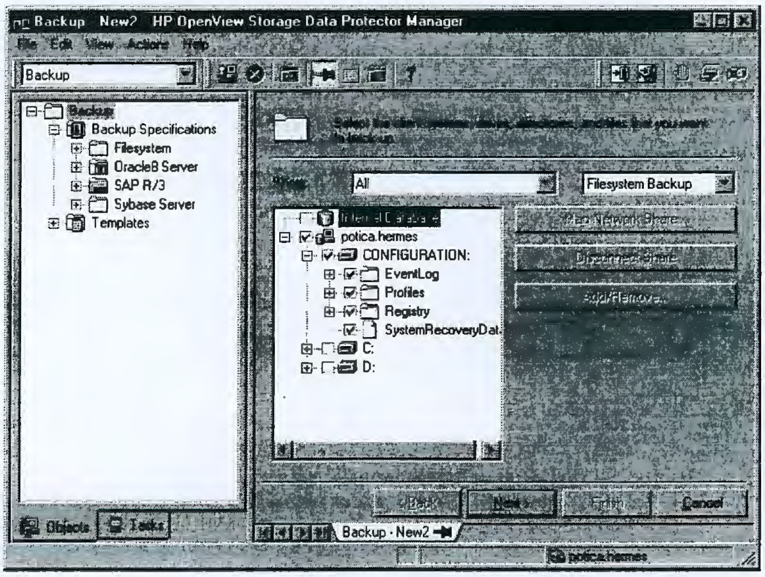

TIP The `SystemRecoveryData` file is only needed for Windows disaster recovery. If a backup version is not used for disaster recovery, clear the `SystemRecoveryData` check boxes when backing up `CONFIGURATION`.

Windows 2000/XP/Server 2003 CONFIGURATION

The items listed at “Windows NT CONFIGURATION” on page 173 also belong to the Windows 2000/XP/Server 2003 `CONFIGURATION`. The following Windows 2000/XP/Server 2003-specific parts are also part of `CONFIGURATION`:

- `QuotaInformation`, `RemovableStorageManagementDatabase`, and `FileReplicationService`.
- The **System State** services

  See “Backing Up the Windows 2000/XP/Server 2003 System State” on page 176.
- `DNSServerDatabase`

18.218

See “Backing Up WINS, DHCP, and DNS” on page 178.

- `SysVol`

  `SysVol` is a shared directory that stores the server copy of the domain's public files, which are replicated among all domain controllers in the domain.

- `IIS`

  Microsoft Internet Information Server is a network file and application server that supports multiple protocols. Primarily, IIS transmits information in Hypertext Markup Language (HTML) pages by using the Hypertext Transport Protocol (HTTP).

Figure 5-11 **Windows 2000/XP/Server 2003 CONFIGURATION**

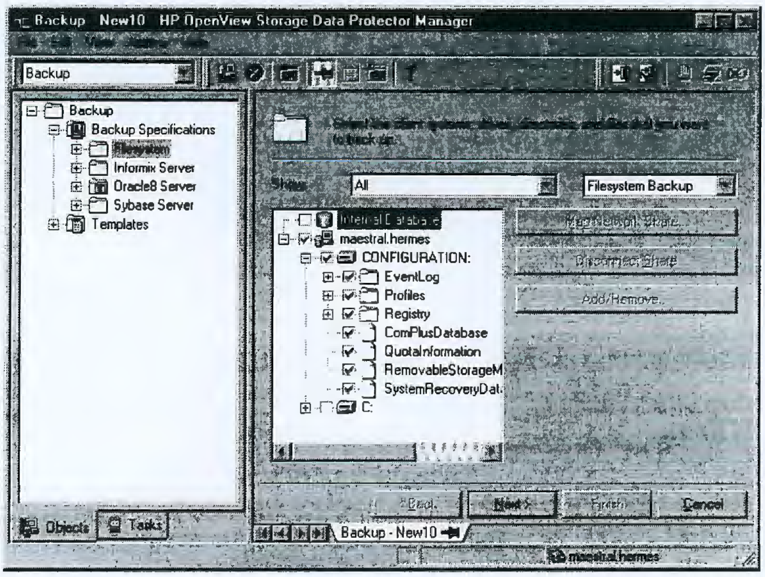

`CONFIGURATION` varies among Windows NT Workstation, Windows NT Server, Windows 2000/XP Professional/XP 64-bit edition, Windows 2000 Server and Windows Server 2003 systems.

18217
Paula

Backing Up CONFIGURATION

Only one CONFIGURATION backup can run on a system at the time. You have to expand a client and select its CONFIGURATION in the Backup wizard.

See "Example of Creating a Backup Specification" on page 155, Figure 5-10 and Figure 5-11.

Backing Up the Windows 2000/XP/Server 2003 System State

The Windows System State consists of several elements related to various aspects of Windows. They are structured under their respective Windows backup object. The Windows System State includes the following:

- Registry and ComPlusDatabase
- The following boot files: Ntldr.exe, Ntdetect.com and boot.ini
- The System Volume Information directory, which keeps data accessed by the System File Protection (SFP) service.

Provided that the services were installed and configured, the System State data of a Windows Server system also includes:

- ActiveDirectoryService
- CertificateServer
- TerminalServiceDatabase

See "Example of Creating a Backup Specification" on page 155 for a detailed backup procedure. Figure 5-12 shows how to select System State in the Backup wizard.

18216
Paula

Figure 5-12 **System State on Windows 2000/XP/Server 2003**

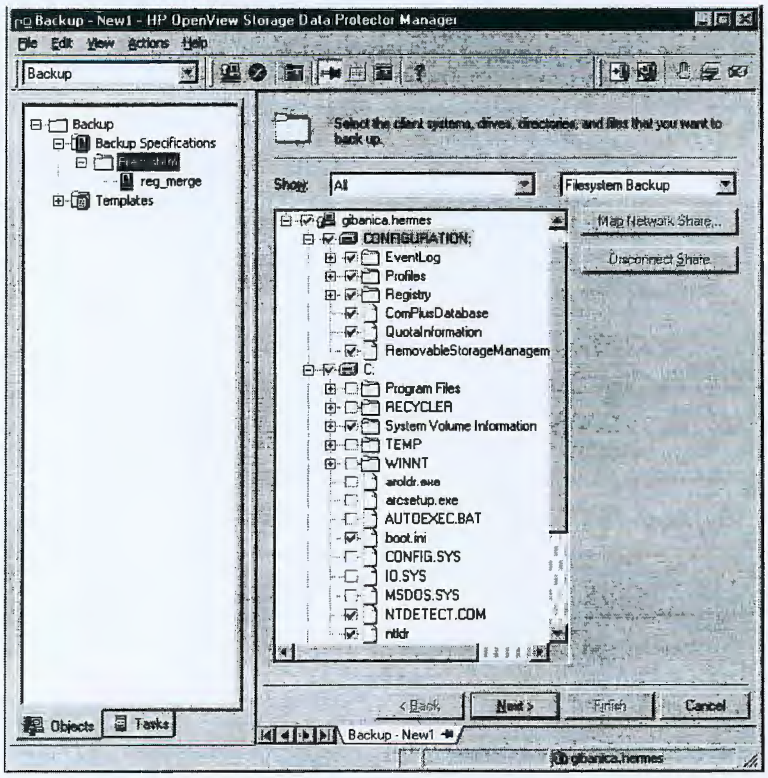

Select the client systems, drives, directories, and files that you want to back up.

Backing Up the Windows Registry

The database repository of information containing the Windows system configuration is known as the `Registry`. The Windows `Registry` is important for the system operation, and must be backed up regularly.

The `Registry` can be backed up as a part of `CONFIGURATION`, or separately by selecting the `Registry` folder as shown in Figure 5-13.

18.215
Paula

Figure 5-13 Backing Up the Windows Registry

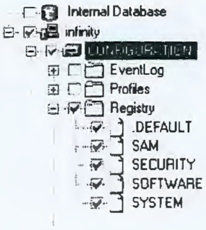

```
□ Internal Database
⊟ ☑ infinity
  ⊟ ☑ CONFIGURATION
    ⊞ □ EventLog
    ⊞ □ Profiles
    ⊟ ☑ Registry
        ☑ .DEFAULT
        ☑ SAM
        ☑ SECURITY
        ☑ SOFTWARE
        ☑ SYSTEM
```

Backing Up WINS, DHCP, and DNS

WINS, DHCP, DNS Servers

In TCP/IP networks, the following services can be configured and run on Windows servers:

- **WINS Server**

  This service, also known as Windows Internet Name Service, is a dynamic replicated database service that can register and resolve NetBIOS names to IP addresses used on a TCP/IP network.

  To back up this database, select `WINS` in the Backup wizard.

- **DHCP Server**

  This service provides dynamic IP address assignment and network configuration for Dynamic Host Configuration Protocol (DHCP) clients.

  To back up this database, select `DHCP` in the Backup wizard.

- **DNS Server**

  This service runs on a Domain Name System server and maintains its own database. A DNS Server answers queries and updates requests for DNS names.

  To back up this database, select `DNSServerDatabase` in the Backup wizard.

18.214 Paula

Backing Up the Windows 2000/XP/Server 2003 Services

Backing up the Windows 2000/XP/Server 2003 services means backing up the data structures used by these services. A particular database is exported (dumped) into a file, which is then backed up. The Windows 2000/XP/Server 2003 services are always backed up if `CONFIGURATION` was selected in the Backup wizard.

NOTE

A Windows 2000/XP/Server 2003 service has to be up and running so that Data Protector can detect it and show it as a selectable item in the Backup wizard. If a service is not running at backup time, the corresponding backup object will fail. See "Managing Failed Backups" on page 263 for more information.

To back up a specific service, you can select the corresponding folder under the `CONFIGURATION` backup object.

Figure 5-14 **Backing Up Windows 2000/XP/Server 2003 Services**

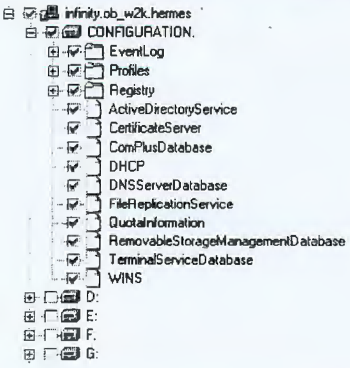

```
⊟ ☑ infinity.ob_w2k.hermes
  ⊟ ☑ CONFIGURATION.
    ⊞ ☑ EventLog
    ⊞ ☑ Profiles
    ⊞ ☑ Registry
      ☑ ActiveDirectoryService
      ☑ CertificateServer
      ☑ ComPlusDatabase
      ☑ DHCP
      ☑ DNSServerDatabase
      ☑ FileReplicationService
      ☑ QuotaInformation
      ☑ RemovableStorageManagementDatabase
      ☑ TerminalServiceDatabase
      ☑ WINS
  ⊞ □ D:
  ⊞ □ E:
  ⊞ □ F.
  ⊞ □ G:
```

See also "Example of Creating a Backup Specification" on page 155 for a step-by-step procedure.

Data Protector can detect and back up the following Windows 2000/XP/Server 2003 services:

- COM+ Event System

1.8213
Paula

This service provides automatic distribution of events to subscribing COM+ components. To back up this database, select the `ComPlusDatabase` in the Backup wizard.

- Removable Storage

  This service manages removable media, drives, and libraries. To back up this database, select `RemovableStorageManagementDatabase` in the Backup wizard.

IMPORTANT You can back up the Removable Storage database, but this service is not used for Data Protector media management. The native robotics driver used with robotics media changers has to be disabled before a device is configured by Data Protector.

Refer to the *HP OpenView Storage Data Protector Installation and Licensing Guide* for details.

- Active Directory Service

  Active Directory Service is the Windows 2000 directory service that enables you to manage data structures distributed over a network. For example, Active Directory Service stores information about user accounts, passwords, phone numbers, profiles, and installed services. It provides the methods for storing directory data and making this data available to network users and administrators.

  To back up the Active Directory data structures that are stored on the local system, select `ActiveDirectoryService` in the Backup wizard.

- Terminal Services

  These services provide a multi-session environment that allows client systems to access a virtual Windows 2000/XP/Server 2003 desktop session and Windows-based programs running on the server.

  To back up this database, select `TerminalServiceDatabase` in the Backup wizard.

- Certificate Services

  These services issue, revoke, and manage certificates employed in public key-based cryptography technologies. To back up this database, select `CertificateServer` in the Backup wizard.

For example, if you use Active Directory to publish Certificate Revocation Lists (CLRs), back up the Active Directory services along with the Certificate Services.

- Remote Storage Service

  Remote Storage Service (RSS) is used to automatically move infrequently accessed files from local to remote storage. Remote files are recalled automatically when the file is opened. Although RSS databases are part of System State data, you must back them up manually. Refer to "Backing Up a Remote Storage Service Database" on page 181.

- System File Protection Service

  System File Protection (SFP) service scans and verifies the versions of all protected system files after you restart your computer. If the SFP service discovers that a protected file has been overwritten, it retrieves the correct version of the file and then replaces the incorrect file. Data Protector enables you to back up and then restore protected files without overwriting them. The protected files can be backed up using the `Move Busy Files` option in a standard filesystem backup procedure.

- DNS, DHCP, and WINS

  See "Backing Up WINS, DHCP, and DNS" on page 178.

Backing Up the DFS

Data Protector backs up the Windows 2000/XP/Server 2003 Distributed File System (DFS) as part of one of the following:

- Windows 2000/XP/Server 2003 Registry, if the DFS is configured in a standalone mode.
- Windows 2000/XP/Server 2003 Active Directory, if the DFS is configured in a domain mode.

Backing Up a Remote Storage Service Database

Data Protector allows you to back up the Remote Storage Server (RSS) database by following the standard filesystem backup procedure. The RSS databases must be backed up offline. You can stop and restart the Remote Storage Service using `pre-` and `post-exec` scripts, or you can perform this manually before and after the backup. Use the following commands:

18.211
Paula

```
net stop/start "Remote Storage Engine"
net stop/start "Remote Storage File"
```

The RSS databases are located in the following directories:

```
<%SystemRoot%>\System32\RemoteStorage
<%SystemRoot%>\System32\NtmsData
```

Backing Up Windows User Profiles, Event Logs, and User Disk Quotas

User Profiles

A `User Profile` contains information about a user configuration. This includes the profile components, such as desktop settings, screen colors, and network connections. When a user logs on, the user profile is loaded and the Windows environment is set accordingly.

The user profile data resides in the following directory:

- `<%SystemRoot%>\Profiles` on Windows NT
- `\Documents and Settings` on Windows 2000/XP/Server 2003

These directories contain all user profiles that are configured on the system and backed up by Data Protector. If a system is configured for multiple users, a separate user profile belongs to each defined user. For example, the `All Users` and `Default User` profiles contain the profile components common to all defined users and those assigned to a newly created user.

Data Protector reads the location of the profiles from the following Registry keys:

```
HKEY_USERS\.DEFAULT\Software\Microsoft\Windows\\
CurrentVersion\Explorer\Shell Folders
```

where information about common profile components resides.

```
HKEY_USERS\.DEFAULT\Software\Microsoft\Windows\\
CurrentVersion\Explorer\User Shell Folders
```

NOTE

If you back up `CONFIGURATION` and the whole Windows system partition as a filesystem, the `Profiles` are backed up twice; as part of a filesystem backup and as part of `CONFIGURATION`. To avoid this, exclude the profile data (see above for location) from the filesystem backup.

18.210
Paula

See also “Warnings When Backing Up System Disks” on page 263.

Event Logs

Event logs are files where the Windows operating system saves information about events, such as starting and stopping services or the logging on and logging off of a user.

User Disk Quotas

User Disk Quotas enable enhanced tracking and control over disk space usage on Windows 2000.

Data Protector backs up user disk quotas on the whole system and for all configured users at a time.

Event Logs, User Profiles, and User Disk Quotas are always backed up if `CONFIGURATION` was selected in the Backup wizard.

See Figure 5-11, “Windows 2000/XP/Server 2003 CONFIGURATION.”, and refer to “Example of Creating a Backup Specification” on page 155 for a step-by-step procedure.

Backing Up Windows Clients Using Disk Discovery

You can use disk discovery by specifying the client as a data source. If another disk is added later, it will be included in the backup.

How Are Disks Discovered?

If you specify a client backup with disk discovery, Data Protector contacts the client and discovers all logical disk drives that belong to physical disks on the client, except for CDs and removable drives. Then it backs up the `CONFIGURATION` folder and each discovered logical drive as a regular filesystem. The description text of each filesystem object will be generated by appending the drive letter in square brackets to the description of the `Client Backup`.

When to Use Disk Discovery

This backup type is recommended under the following circumstances:

- When backing up systems with relatively small disks
- When performing a whole system backup to prepare for disaster recovery
- When the number of disks connected to the system varies.

For a client backup with disk discovery, it is not possible to select only specific directory trees, because this implies a single logical drive backup. It is, however, possible to exclude any directory from the backup.

13.209
Paula

How to Perform a Backup

To perform a Windows client backup, you have to create a backup specification as described in “Example of Creating a Backup Specification” on page 155.

In the `Source` property page of the Backup wizard, select the check box next to the client name to obtain the disk discovery functionality. Then, follow the wizard.

NOTE

Selecting all of the client’s drives is not the same as selecting the check box next to the client name, which is the procedure for a Disk Discovery backup.

To check the configured backup type, see the `Backup Object Summary` property page. Under the `Type` label, you will see `Client System` if you have configured a Disk Discovery backup or `Filesystem` if only the disks have been selected.

Figure 5-15 **Selecting the Client**

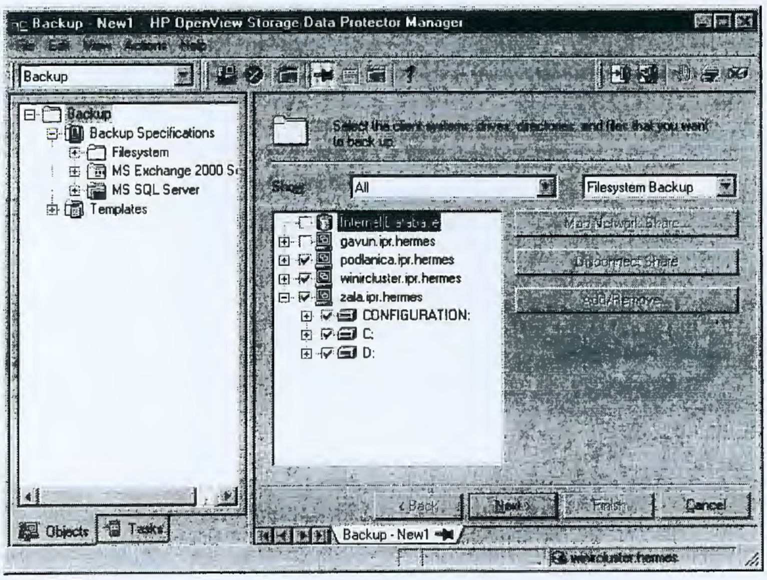

Select the client systems, drives, directories, and files that you want to back up.

18208
Paula

See “Using Backup Options” on page 225 for information on using and structuring your backup specifications.

Backing Up Windows Shared Disks

Data Protector allows you to back up data on Windows shared disks. You have to use a regular Disk Agent client, which can then be used to back up other remote systems via shared disks. Then you can configure a backup specification.

NOTE

Backup using the shared disk method is a workaround for backing up systems which cannot be backed up otherwise. It is better not to use it as the main backup approach.

When to Use Shared Disks Backup

Use shared disks backup in either of the following situations:

- The remote system does not belong to the Data Protector cell and does not have the Data Protector Disk Agent installed.
- The platform to be backed up is not directly supported by Data Protector, for example, Windows 3.11.

TIP

To reduce the network load, the Disk Agent client should be the Media Agent client as well. Otherwise, data is transferred over the network twice.

You can use one Windows client to manage backups and restores involving shared disks or other remote systems. Backup performance may be reduced if you start too many backups at a time, since one Disk Agent is started for each backed up disk. In this case, you should configure additional Disk Agent clients to increase the backup speed.

Limitation

Backing up writers that store their data on network shared volumes using the VSS functionality is not supported.

18.207
Paula

IMPORTANT

The Disk Agent client must have the Inet service configured using an account with access to the shared disks. This must be a specific user account, not the system account. See “Setting the User Account for the Data Protector Inet Service” on page 187 for more information on how to use the appropriate logon account.

Once you have set the user account for the Inet service, you can back up the shared disks as though they were residing on the local system.

How to Perform a Windows Shared Disks Backup

1. In the Data Protector Manager, switch to the Backup context.
2. Expand the Backups item, and then double-click Backup Specifications.
3. Right-click Filesystem, and then Add Backup.
4. In the Create New Backup dialog box, select one of the available templates, and then click OK to open the wizard.
5. On the first page of the wizard, in the drop-down list, select Network Share Backup.

IMPORTANT

You have to map the shared drives using the Backup wizard.
If you are using the GUI on a UNIX system, it is not possible for the system to confirm the existence of a Windows shared drive, or to browse it. Therefore, when specifying a Windows shared drive/directory, you must confirm yourself that it is available and correctly specified, or the backup will fail.

6. Click Map Network Share. The Browse Network Shares dialog box opens.
7. In the Client System drop-down list, select the client with the Disk Agent that will be used to back up the remote system.
8. Select the shared disk. It appears in the Share Name text box.

IMPORTANT

Share names containing spaces are not supported.

9. Enter the required information. See online Help for details.

Setting the User Account for the Data Protector Inet Service

The following procedure describes how to change the user account used by the `Data Protector Inet` service to access disks that belong to remote computers. This account must have permission to access both the local client and the remote shared disks. It must be a specific user account, not the system account.

Windows NT

Proceed as follows to change the user account on a Windows NT Disk Agent client:

1. On Windows NT, click the `Start` button, then `Settings`, and `Control Panel`.

   In the `Control Panel` window, double-click `Services`. The `Services` dialog box appears.

2. Scroll down the list of services, select `Data Protector Inet`, then click `Stop`.
3. Click `Startup`. The `Service` dialog box appears.
4. In the `Log On As` area, select the `This Account` radio button.
5. Enter or browse for the account that has the correct permission to access the shared disks that you want to back up.
6. Enter and confirm the password.

Paula

Figure 5-16 **Inet Logon Option on Windows NT**

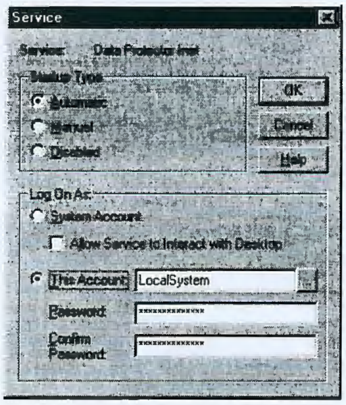

Service

Service: Data Protector Inet

Startup Type
● Automatic
○ Manual
○ Disabled

Log On As:
○ System Account
□ Allow Service to Interact with Desktop
● This Account: LocalSystem
Password: **************
Confirm Password: **************

7. Click OK to return to the Services dialog box.
8. Ensure that Data Protector Inet is still selected, and then click Start.
9. Click Close to confirm and exit this dialog.

   Data Protector is now able to access all disks to which the account you entered has access.

Windows 2000/XP/Server 2003

Proceed as follows to change the user account on a Windows 2000/XP/Server 2003 Disk Agent client:

1. In the Control Panel, click Administrative Tools, and then double-click Services.
2. Scroll down the list of services and select Data Protector Inet.
3. Under the General property page, click Stop. Then select the Log On tab.

Paula

Figure 5-17 **Inet General Property Page on Windows 2000/XP/Server 2003**

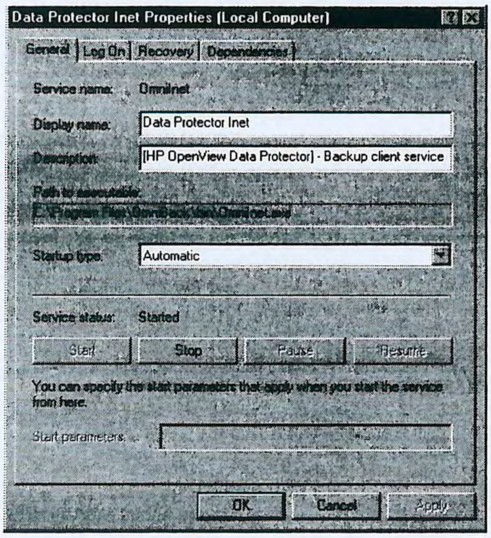

Service name: Omniinet

Display name: Data Protector Inet

Description: [HP OpenView Data Protector] - Backup client service

Path to executable:
C:\Program Files\OmniBack\bin\Omniinet.exe

Startup type: Automatic

Service status: Started

You can specify the start parameters that apply when you start the service from here.

Start parameters:

4. In the `Log On As` area, select the `This Account` button.
5. Enter or browse for the account that has the correct permission to access the shared disks you want to back up.
6. Enter the password, then confirm it.

18 203
Paula

Figure 5-18 **Inet Logon option on Windows 2000/XP/Server 2003**

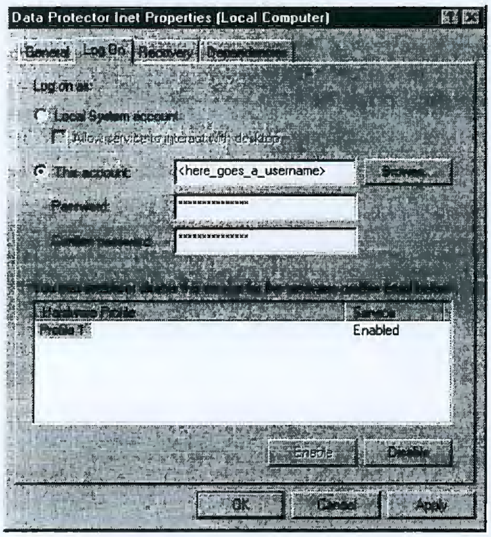

Data Protector Inet Properties (Local Computer)

Log on as:

○ Local System account

☐ Allow service to interact with desktop

● This account: <here_goes_a_username>

Password: ***************

Confirm password: ***************

Hardware Profile	Service
Profile 1	Enabled

7. Click `Apply` to apply the changes and then restart the service by clicking the `Start` button in the `General` property page.

Backing Up Windows Disks as Disk Image Objects

What Is a Disk Image Backup?

A disk image backup is a high-speed backup of disks, disk partitions, or logical volumes without tracking the file and directory structure stored on these data sources.

When to Use a Disk Image Backup

Use a disk image backup in the following situations:

- You have lots of small files and a high backup speed is required.
- A full disk backup is needed, for example, for disaster recovery or before a major software update.
- A direct disk-to-disk connection is not possible and you want to duplicate a filesystem to another disk. The latter must be identical to the original disk.

Paula

How to Specify a Disk Image Section

You can specify a disk image section in two ways. In case of a zero downtime backup (snapshot or split mirror), you must use the second way.

- `\\.\<drive_letter>`, for example: `\\.\E:`
- `\\.\PHYSICALDRIVE#`,

  where # is the current number of the disk you want to back up.

  For example: `\\.\PHYSICALDRIVE3`

Where to Find a Disk Number (Physical Drive Number)

On Windows 2000/XP/Server 2003 systems, you can find the current numbers of your disks (as well as the drive letters) by clicking `Control Panel`, `Administrative Tools`, `Computer Management`, `Storage`, `Disk Management`.

On Windows NT systems, you can find the current numbers of your disks (as well as the drive letters) by clicking `Start`, `Programs`, `Administrative Tools`, `Disk Administrator`.

15201
Paula

Figure 5-19 **The Numbers Representing Disks (Physical Drive Number) on Windows 2000 System**

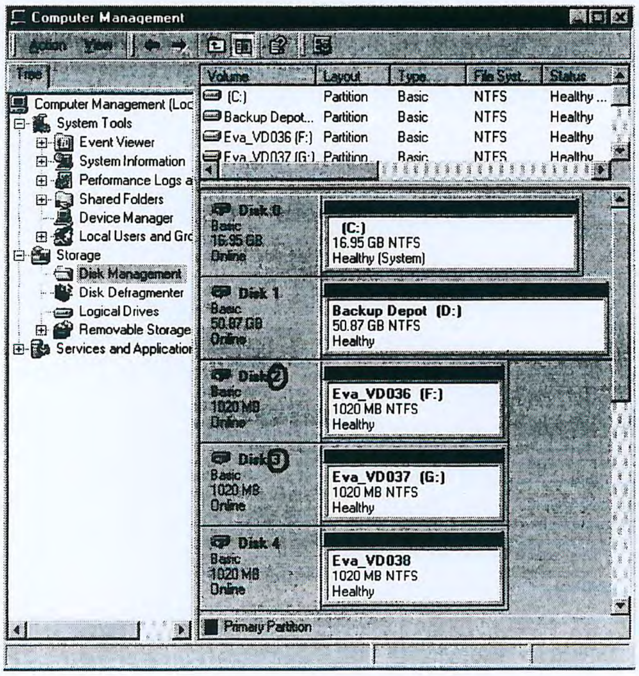

Volume	Layout	Type	File Syst...	Status
(C:)	Partition	Basic	NTFS	Healthy ...
Backup Depot...	Partition	Basic	NTFS	Healthy
Eva_VD036 (F:)	Partition	Basic	NTFS	Healthy
Eva_VD037 (G:)	Partition	Basic	NTFS	Healthy

Disk 0
Basic
16.95 GB
Online

(C:)
16.95 GB NTFS
Healthy (System)

Disk 1
Basic
50.87 GB
Online

Backup Depot (D:)
50.87 GB NTFS
Healthy

Disk 2
Basic
1020 MB
Online

Eva_VD036 (F:)
1020 MB NTFS
Healthy

Disk 3
Basic
1020 MB
Online

Eva_VD037 (G:)
1020 MB NTFS
Healthy

Disk 4
Basic
1020 MB
Online

Eva_VD038
1020 MB NTFS
Healthy

NOTE The numbers representing disks can change if the system is rebooted.

Limitation A disk image backup fails if a file on the target system is open, since Data Protector cannot lock the file.

How to Perform a Disk Image Backup To perform a disk image backup, use the `Manual add` function from the `Backup Object Summary` page. For detailed steps, refer to the online Help index keyword “backing up, disk images”.

18200 Paula

Figure 5-20 The Manual Add Functionality

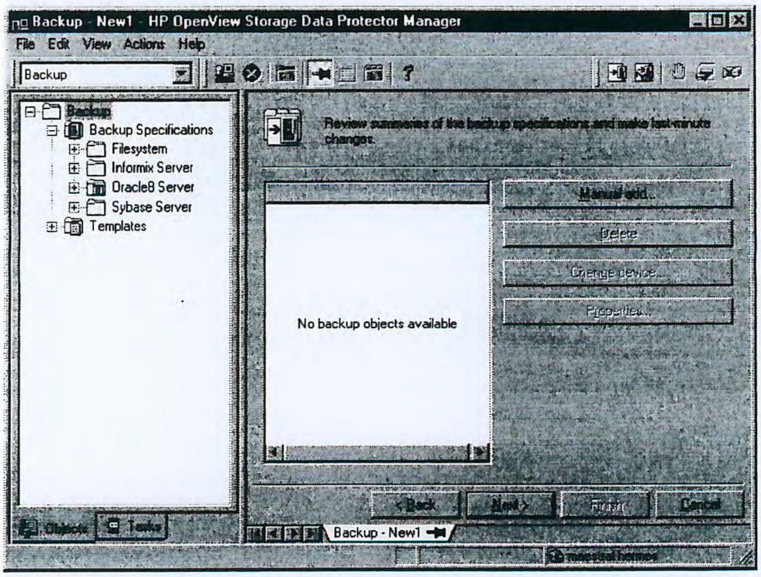

Review summaries of the backup specifications and make last-minute changes.

No backup objects available

18199
Paula

Backing Up Novell NetWare Systems

This section describes how to back up Novell NetWare filesystems and NetWare Directory Services (NDS).

Backing Up Novell NetWare Filesystems (Volumes)

Prerequisites

To back up data on a Novell NetWare system, install the Novell NetWare Disk Agent on the Novell NetWare system. See the *HP OpenView Storage Data Protector Installation and Licensing Guide* for instructions.

To use backup devices connected to a Novell NetWare system with Data Protector, install the Novell NetWare Media Agent on the Novell NetWare system. See the *HP OpenView Storage Data Protector Installation and Licensing Guide* for instructions.

How to Back Up a Novell NetWare System

To back up Novell Netware filesystems, proceed as follows:

1. In the `HP OpenView Storage Data Protector Manager`, switch to the `Backup` context.
2. Expand `Backups`, right-click `Backup Specifications`, and then click `Add Backup`.
3. In the `Create New Backup` dialog box, select one of the available templates, and then click `OK` to open the wizard.
4. In the drop-down list, select `Filesystem Backup`.
5. Expand the client whose data you want to back up.
6. Select the backup objects. Follow the wizard to select a backup device.
7. In the next wizard page, click the `Advanced Filesystem Options` tab to open the `Filesystem Options` dialog box. Click the `NetWare Options` tab to set backup options. Refer to "Object Options" on page 239 for a description of the options.
8. Follow the wizard and then save and start your backup.

What Is Backed Up?

The directory structure and the files are backed up as well as the following filesystem information:

- Four Name Space information: DOS, Mac, NFS, Long

18198

- Trustee information
- Inherited right mask
- File and directory attributes
- Time attributes (creation date/time, modification date/time, last accessed date/time, last modified date/time, last archived date)
- Owner
- Owning name space
- Search mode
- Volume or directory space restrictions. To back up volume restrictions, select the whole volume object for backup.

Server Specific Information is backed up separately as a part of CONFIGURATION mount point.

After backing up each file, the file's archive flag is cleared and the archive time is set.

The Novell NetWare filesystem provides file compression transparently to NLMs and clients. By default, Data Protector backs up and consequently restores such files in their compressed format. Thus, they will be restored to Novell NetWare platforms. To restore Novell NetWare compressed files to non-Novell NetWare platforms, use the `Uncompress NetWare compressed files` option at a backup time.

What Is Not Backed Up?

Files that are opened for shared access with the `Denied read` option enabled cannot be backed up by Data Protector. You can set the `Number of retries` option to increase the probability of the file being backed up. This option is only useful if the applications operate in such a mode that they use a certain file and then release it after a certain time.

- System files that are in Queue directories are not backed up.
- All files that belong to NDS are skipped. You can back up NDS separately.
- Extended attributes (which can be installed as a NetWare addition) are not backed up.

Limitations

The following features are unavailable for NetWare backups:

- `Pre-exec` and `post-exec` options

18197
Paula

- The Compress option
- The omit_deleted_files option (restore option)

Files of up to 4 GB are backed up on NetWare 4.X.

NOTE

To allow users to run backups on the Novell NetWare system, grant them the Backup as Root user right. See Chapter 3, “Configuring Users and User Groups,” on page 81 for details on how to change user rights.

Data Protector cannot back up **moved files** during incremental backup sessions.

What Is Included in an Incremental Backup?

In order to determine which files have changed, the Data Protector Disk Agent checks the last modification time for each file. This method prevents Data Protector from detecting moved files, as moving the file does not change the modification time.

See the *HP OpenView Storage Data Protector Concepts Guide* for details about incremental backups.

Selecting Specific Files or Directories

For each filesystem, you can restrict the backup to specific directory trees. For each directory tree you can do the following:

- Exclude any sub-tree or file
- Back up files that match a specific wildcard pattern
- Skip files that match a specific wildcard pattern

Backing Up CONFIGURATION

Data Protector enables you to back up a special data structure known as CONFIGURATION, which consists of the following components, as shown in Figure 5-21 (NetWare 4.x, 5.x and 6.0).

CONFIGURATION Components

- Server Specific Info
- Schema
- Root

To back up the CONFIGURATION item or part of it, follow the procedure “How to Back Up a Novell NetWare System” on page 194, selecting the appropriate item in the Source page of the backup wizard.

Figure 5-21 **Backing Up NetWare 4.x and NetWare 5.x Configuration**

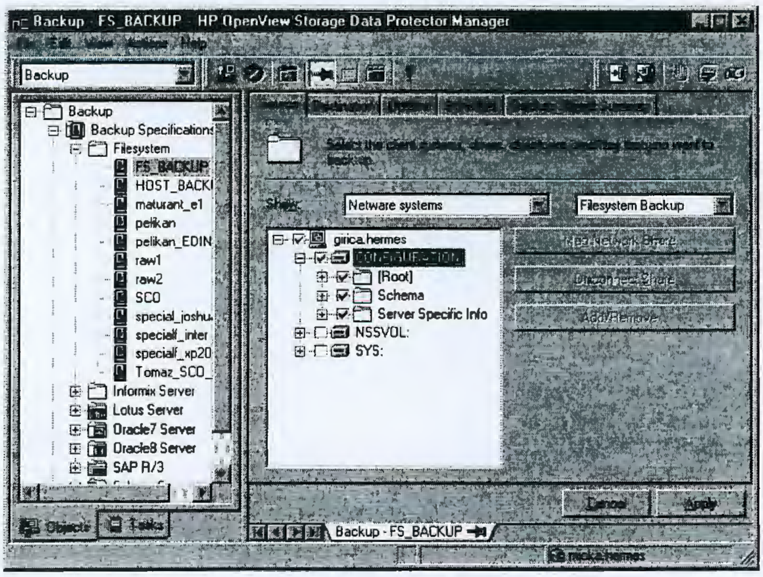

Backup - FS_BACKUP - HP OpenView Storage Data Protector Manager

Select the client systems, drives, directories and files you want to back up.

Show: Netware systems | Filesystem Backup

13135
Paula

Client Backup with Disk Discovery

You can discover disks (volumes) on NetWare just as you can for UNIX or Windows systems.

How Are Disks Discovered?

If you specify a client backup with disk discovery, Data Protector first contacts the client and discovers all volumes that belong to the client. Then it backs up the CONFIGURATION item and each discovered volume as a regular filesystem. The description text for each filesystem object is generated by appending the volume name, in square brackets, to the description of the client backup.

For client backup with disk discovery, it is not possible to select only specific directory trees, because this implies single volume backup. But it is possible to exclude any directory from the backup.

How to Perform a NetWare Client Backup

1. In the Data Protector Manager, switch to the Backup context.
2. Expand the Backup item, and then double-click Backup Specifications.
3. In the Results Area, right-click Filesystem, and then click Add Backup.
4. In the Create New Backup dialog box, select one of the available templates.
5. Click OK to open the wizard.
6. Click the check box next to the client. This selects the entire client to be backed up, similar to what is shown in Figure 5-15.

See "Using Backup Options" on page 225 for information on using and structuring your backup specifications.

18/04
Paula

Backing Up NetWare Directory Services (NDS)

Data Protector backs up NDS using Novell NetWare Storage Management Services (SMS). Data Protector backs up and restores all extensions to the NDS Schema.

NOTE Incremental backup of the NDS database is not possible. A full backup of the NDS database is always performed.

To successfully back up NDS, follow the instructions in the *HP OpenView Storage Data Protector Installation and Licensing Guide* and ensure that:

- `TSANDS.NLM` is loaded
- `HPLOGIN.NLM` is loaded and access information is given to Data Protector

Backing Up NDS Back up NDS as you would a UNIX or Windows filesystem, except that the mountpoint has to be the `CONFIGURATION` item.

Adding NDS Objects to a Backup Specification

Data Protector offers advanced functionality to back up only a part of NDS. However, unless you understand why some parts can be excluded, it is advisable to back up everything.

Each object in the NDS tree has its own fully distinguished name. For example, leaf object `CN=Admin`, which resides in the container object `O=HSL`, has its fully distinguished name as seen by the SMS (`TSANDS.NLM`):

```
.CN=Admin.O=HSL.[Root]
```

Data Protector uses the fully distinguished name to build the tree structure of the NDS as follows:

- The fully distinguished name is reversed.
- The dot-symbol (`.`) separator is replaced with the slash-symbol separator (`/`).

For example, the fully distinguished name

`.CN=Admin.O=HSL.`[Root]

18.193 Paula

has its counterpart used by Data Protector, containing forward slashes, which are used for Windows as well:

```
/[Root]/O=HSL/CN=Admin
```

Except for this naming rule, the Data Protector backup specification syntax is the same as for Novell NetWare or UNIX filesystem objects.

NOTE NDS objects (container and leaf objects) are represented and backed up as directories. These objects can be skipped using the `skip` option or backed up using the `only` option. Data Protector views the [Root] object as a non-containment object, so the [Root] object cannot be excluded.

The Mountpoint Configuration File TSANDS.CFG

For the best protection of your NDS data, you should perform a full directory backup of the NDS Schema and all containers in the tree starting with the [Root] object. However, there are situations where you might prefer to begin backing up NDS from a container other than the [Root] object, but a configured user does not have sufficient rights to browse through to the starting container's context.

To facilitate backing up portions of the NDS tree, Novell has provided a text file, `SYS:SYSTEM\TSA\TSANDS.CFG` file, that allows you to specify the names of containers where you want backups to begin. This file is located on the server where `TSANDS.NLM` is loaded.

To begin your NDS backup from the HSL container, create a `TSANDS.CFG` file containing the line:

```
.O=HSL.[Root]
```

An additional mountpoint becomes available to the backup configuration.

15.192
Paula

Backing Up OpenVMS Systems

This section describes how to back up OpenVMS filesystems.

Backing Up OpenVMS Filesystems

Prerequisites

To back up data on a OpenVMS system, install the OpenVMS Disk Agent on the OpenVMS system. See the *HP OpenView Storage Data Protector Installation and Licensing Guide* for instructions.

To use backup devices connected to an OpenVMS system with Data Protector, install the OpenVMS Media Agent on the OpenVMS system. See the *HP OpenView Storage Data Protector Installation and Licensing Guide* for instructions.

How to Back Up an OpenVMS System

To back up an OpenVMS filesystem, proceed as follows:

1. In the `HP OpenView Storage Data Protector Manager`, switch to the `Backup` context.
2. Expand `Backups`, right-click `Backup Specifications`, and then click `Add Backup`.
3. In the `Create New Backup` dialog box, select one of the available templates, and then click `OK` to open the wizard.
4. In the drop-down list, select `Filesystem Backup`.
5. Expand the client whose data you want to back up.
6. Select the backup objects. Follow the wizard to select a backup device.
7. Select backup options. See "Using Backup Options" on page 225 for details.
8. Follow the wizard and then save and start your backup.

What Is Backed Up?

The directory structure and the files are backed up, together with the following filesystem information:

- File and directory attributes
- ACL (Access Control List)

13.191
Paula

Files can be backed up from mounted FILES-11 ODS-2 or ODS-5 volumes only.

Limitations

- Any file specifications that are entered into the GUI or passed to the CLI must be in UNIX style syntax:

  ```
  /disk/directory1/directory2/filename.ext.n
  ```

  - The string should begin with a slash, followed by the disk, directories, and filename, separated by slashes.
  - Do not place a colon after the disk name.
  - A period should be used before the version number instead of a semi-colon.
  - File specifications for OpenVMS files are case insensitive.

  For example:

  An OpenVMS file specification of:

  ```
  $1$DGA100:[USERS.DOE]LOGIN.COM;1
  ```

  must be specified in the form:

  ```
  /$1$DGA100/Users/Doe/Login.Com.1
  ```

- There is no implicit version number. You always have to specify a version number. Only file versions selected for the backup will be backed up. If you wish to include all versions of the file, select them all in the GUI window, or, using the CLI, include the file specifications under the `Only (-only)` option, including wildcards for the version number, as follows

  ```
  /DKA1/dir1/filename.txt.*
  ```

- If the `Do not preserve access time attributes (-touch)` option is enabled during a backup, the last accessed date will be updated with the current date and time on ODS-5 disks. On ODS-2 disks, this option has no effect, and all the dates remain unchanged.
- Rawdisk backups are not available on OpenVMS. There is no equivalent to a "BACKUP/PHYSICAL".
- The `Backup POSIX hard links as files (-hlink)`, `Software compression (-compress)`, and `Encode (-encode)` options are not available on OpenVMS.

18190
Paula

Files with multiple directory entries are only backed up once using the primary path name. The secondary path entries are saved as soft links. During a restore, these extra path entries will also be restored.

There is no support for an equivalent to BACKUP/IMAGE. To make a restored copy of an OpenVMS system disk bootable, the OpenVMS WRITEBOOT utility has to be used to write a boot block on to the restored disk.

- Files being backed up are always locked regardless of whether the `Lock files during backup (-lock)` option is enabled or disabled. With the `-lock` option enabled any file opened for write is not backed up. With the `-lock` option disabled any open file is backed up as well.
- The default device and directory for pre- and post-exec command procedures is `/omni$root/bin`. To place the command procedure anywhere else the file specification must contain the device and directory path in UNIX style format: For example: `/SYS$MANAGER/DP_SAVE1.COM`
- When specifying wildcards for `Skip (-skip)` or `Only (-only)` filters, use '*' for multiple characters and '?' for single characters.

18.189
Paula

Backing Up in a Direct Backup Environment

This section provides the steps for the configuration of a direct backup backup specification. Please refer to *HP OpenView Storage Data Protector Concepts Guide* for a complete information on direct backup concepts.

Prerequisites

- The application and backup systems must be configured for split mirror or snapshot backup, depending on the disk array used. Refer to the *HP OpenView Storage Data Protector HP StorageWorks Disk Array XP Integration Guide* or to the *HP OpenView Storage Data Protector EVA / VA / MSA Integration Guide*.
- If backing up an Oracle8i server on the application system, the application system must be configured for the Oracle8i split mirror or snapshot backup, depending on the disk array used. Refer to the *HP OpenView Storage Data Protector HP StorageWorks Disk Array XP Integration Guide* or to the *HP OpenView Storage Data Protector EVA / VA / MSA Integration Guide*.
- The XCopy engine must be configured in the same SAN zone as the source (mirror disk connected to the backup system) and the destination (backup device connected to a Fibre Channel bridge). In other words, the XCopy engine must have SAN access to both the mirror disk connected to the backup system and to backup device connected to a Fibre Channel bridge.
- You need to have HP StorageWorks Disk Array XP agent or HP StorageWorks Virtual Array agent installed on every system that is to be backed up (application system). Refer to *HP OpenView Storage Data Protector Installation and Licensing Guide*.
- You need to have Media Agent and HP StorageWorks Disk Array XP agent or HP StorageWorks Virtual Array agent installed on every system that controls a backup device (backup system). Refer to *HP OpenView Storage Data Protector Installation and Licensing Guide*.
- You need to have at least one backup device configured for direct backup in the Data Protector cell. Refer to "Configuring Devices for Direct Backup" on page 38.
- You need to have prepared media for your backup. Refer to Chapter 4, "Managing Media," on page 97.

18.188
Paula

- You need to have appropriate user rights for performing a backup. Refer to Chapter 3, “Configuring Users and User Groups,” on page 81.

Limitations

- The systems in the direct backup environment must be HP-UX 11.0.
- The `min` and `max` options for the `load balancing` option are ignored for direct backup. All devices selected in the backup specification are load balanced, if load balancing is used. Consequently, it is not possible to set the order in which the selected devices are used using the Order devices functionality.
- The pre-exec and post-exec options for backup objects are not possible for raw logical volumes direct backup. They are possible for Oracle8i direct backup.
- The backup device must be either attached to an external FC bridge with the XCopy engine, or must have the FC bridge with the XCopy engine embedded internally.
- Backup and restore of striped logical volumes are not supported.
- The `CRC check` option is ignored with direct backup.
- The disk agent `Concurrency` option is ignored with direct backup.
- The `Block size` option is FC bridge dependent.
- The `Segment size` and `Disk agent buffers` options are ignored with direct backup.

Restore

The data backed up in a direct backup environment can be:

- Restored from a backup medium over the LAN directly to the application system following the Data Protector rawdisk or Oracle8 restore procedure. Refer to “Restoring Disk Images” on page 273 (rawdisk restore) or *HP OpenView Storage Data Protector Integration Guide* (Oracle8 restore).
- Restored using the Data Protector instant recovery functionality. Refer to the *HP OpenView Storage Data Protector HP StorageWorks Disk Array XP Integration Guide* or to the *HP OpenView Storage Data Protector EVA / VA / MSA Integration Guide*.

Backup Specification Configuration Procedure

A direct backup backup specification can be configured for the following objects:

18.187
Paula

- rawdisks
- Oracle8i databases (online)
- Oracle8i databases (offline)

Backing Up Rawdisks

Please refer to the online Help index keyword “configuring direct backup specifications” for detailed information on how to configure a rawdisk direct backup specification.

Backing Up Oracle8i Online and Offline

Please refer to the online Help index keyword “configuring direct backup specifications” for detailed information on how to configure an Oracle8i online or offline direct backup specification.

Starting Direct Backup Using the CLI

After a direct backup specification has been configured, you can start the direct backup session using the GUI as described in the previous section, or using the CLI as follows:

- for rawdisks

  ```
  omnib -datalist <Name>
  ```

- for Oracle8i Online and Offline

  ```
  omnib -oracle8_list <Name>
  ```

where *<Name>* is the name of the direct backup specification.

18.186
Paula

Scheduling Unattended Backups

Data Protector allows you to configure unattended backups by scheduling backups of your systems at specific times.

The configuration and your scheduling policies can significantly influence the effectiveness and performance of your backup.

Key Points

- To simplify scheduling, Data Protector provides backup specifications for group clients. All clients configured in one backup specification are backed up at the same time in a single backup session.
- Be sure to have sufficient media and devices to run unattended backups smoothly. See Chapter 7, “Monitoring, Reporting, Notifications, and the Event Log,” on page 307 for details on monitoring running sessions and setting up e-mail or other notifications for a mount request.
- When the scheduled backup is started, Data Protector tries to allocate all the needed resources, such as licenses, devices, and access to IDB. If one of the needed resources is not available, the session is marked as queued. Data Protector will try to find the needed resources for the queued session once every minute until the time-out period is reached.

  Once Data Protector finds the resources, one of the queued sessions is started. The queued sessions are not started in the order they are displayed.
- To prevent Cell Manager overload, a maximum of up to five backup sessions can be started at the same time. If more are scheduled at the same time, the sessions are queued.
- For each individual or periodic scheduled backup, you can specify the following options: `Backup type` (full or incremental), `Network load`, and `Backup protection`. With split mirror or snapshot backup, in the case of ZDB disk or ZDB disk/tape backups (instant recovery enabled), you specify the `Split mirror/snapshot backup` option. For split mirror and snapshot backups, the backup type is ignored (it is set to full).

18.185
Paulo

- Each backup specification can be scheduled multiple times with different option values. Within one backup specification, you can schedule both ZDB disk and ZDB disk/tape backups, and specify a different data protection period for each individual or periodic scheduled backup.
- Data and catalog protection settings determine the period that data is kept on a medium (data protection) and in IDB (catalog protection). See "Data Protection: Specifying How Long Data Is Kept on the Media" on page 228 and "Catalog Protection: How Long Info Is Kept in the Database" on page 230 for details.
- When applying a backup template, the schedule settings of the template override the schedule settings of the backup specification. After applying the template, you can still modify the backup specification and set a different schedule.

NOTE

You can schedule backups up to a year in advance. Periodic backups do not have a defined time limit. Weekly periodic backups can be configured only if the time between two recurring backups is at most four weeks.

Handling Scheduling Conflicts

When scheduling periodic backups, it can happen that the chosen backup start time is already occupied by another scheduled backup in the same backup specification. In that case, Data Protector prompts you that there are scheduling conflicts, and asks if you wish to continue. If you click `Yes`, the new schedule will be applied where possible (on the days when the time slot is still free). If you click `No`, the new schedule will be discarded.

Planning Your Scheduling Policies

See the *HP OpenView Storage Data Protector Concepts Guide* for answers to questions such as:

- How do I plan a scheduling policy for my environment?
- How does the amount of data influence my scheduling policy?
- How long will the backup take?
- How many media do I need for the backup?
- How do I plan for a disaster recovery?

13.184
Laura

Starting Backups on Specific Dates

Data Protector allows you to define the date and time when you want your unattended backup to start. You usually want to back up on specific dates when configuring exceptions to your regular periodic backups, for example, if you want to back up some data before a specific event.

How to Configure Backups on Specific Dates

To configure a backup on a specific date, you can create a new backup specification, or modify an existing one. For detailed steps, refer to the online Help index keyword "scheduling backups on specific dates and times".

Starting Periodic Backups

Periodic backups are based on a time period after a specific date. For example, you may configure periodic backups so that a full backup is done on Sunday at 3 a.m. and repeated every two days. The next full backup would be at 3 a.m. the following Tuesday. Periodic backups simplify backup configuration for regularly scheduled backups.

Data Protector provides predefined backup schedules to simplify the configuration.

Predefined Backup Schedules

The predefined backup schedules provided can be used to simplify your configuration. You can modify the schedules later. Schedule types include those described in the following sections:

Daily intensive

Data Protector runs a full backup at midnight and two additional incremental backups at 12:00 (noon) and 18:00 (6 p.m.) every day. This backup type is intended for database transaction servers and other environments with intensive backup requirements.

Daily full

Data Protector runs a full backup every day at 21:00 (9 p.m.). This is intended for backups of single workstations or servers.

Weekly full

Data Protector runs a full backup every Friday and Incr1 backups every day from Monday to Friday at 21:00 (9 p.m.). This is intended for small environments.

Fortnightly full

Data Protector runs a full backup every second Friday. Between these backups, Data Protector runs Incr1 backups every Monday to Thursday, all at 21:00 (9 p.m.).

Monthly full

Data Protector runs a full backup on the first of every month, an Incr1 backup every week, and an incremental backup every other day. This is intended for relatively static environments.

How to Use a Predefined Schedule

To configure a backup using a predefined schedule, you can create a new backup specification, or modify an existing one. For detailed steps, refer to the online Help index keyword “scheduling periodic backups”.

Configuring a Recurring Backup

You can schedule a backup so that it starts at a specific time and date on a set schedule. For example, you could schedule a full backup to take place every Friday at 21:00 (9 p.m.) for the next six months.

How to Configure a Recurring Backup

To configure a recurring backup, you can create a new backup specification, or modify an existing one. For detailed steps, refer to the online Help index keyword “scheduling periodic backups”.

Editing Your Backup Schedule

Clearing a Schedule

To eliminate a schedule that you have already set up, click `Reset` in the `Schedule` property page.

When you clear a schedule, you clear all the schedule settings of a specified mode for the current year.

Undoing the Clear

To undo the schedule clearing, click `Undo` in the `Schedule` property page.

Changing the Start Date

To change the start date, follow the procedure for setting up a backup for a specific date. See “Starting Backups on Specific Dates” on page 209.

Disabling a Schedule

18182
Laula

To disable a backup schedule, select the `Disable Schedule` option in the `Schedule` property page. The backup will not be performed until you deselect this option.

Disabling backup schedules does not influence currently running backup sessions.

Skipping Backups During Holidays

By default, Data Protector runs backups on holidays.

If you do not wish to run your backups on holidays, set the `Holidays` option to ON in the `Schedule` page of the Backup wizard. You can identify holidays from the `Holidays` file or as red dates on the `Schedule Calendar`.

To set different holidays, edit the `Holidays` file, located in the following directory:

- on UNIX: `/etc/opt/omni/Holidays`
- on Windows: *`<Data_Protector_home>`*`\Config\Holidays`

How you configure your scheduling policies strongly influences the effectiveness and performance of your backup. For example, if the date January 1 is registered as a holiday, Data Protector will not back up on that date. If you have scheduled a full backup for January 1st and an incremental for January 2nd, Data Protector will skip running the full backup on January 1st but will run the incremental backup scheduled for January 2nd. The incremental backup will be based on the last full backup.

Configuring Backup Options When Scheduling Backups

When scheduling a backup, you can set further options. These options are only valid for scheduled backups and not for those started interactively. Data protection that is specified in the `Schedule Backup` dialog overrides protection settings anywhere else in the backup specification.

How to Set Schedule Backup Options

You can set schedule backup options when creating a new backup specification, or when modifying an existing one. For detailed steps, refer to the online Help index keyword “setting schedule backup options”.

18 181
Paula

Running Consecutive Backups

You can start a backup after one is finished. For example, you can start a backup of an Oracle database after a filesystem backup is finished. For detailed steps, refer to the online Help index keyword "running consecutive backups".

For details on `pre-` and `post-exec` scripts on UNIX systems, refer to the Appendix, "Examples of Pre-Exec and Post-Exec Commands for UNIX," on page A-20.

18.180
Paula

Selecting a Backup Type: Full or Incremental

To save time and media during a backup, you can combine full and incremental backups. For example, you can create a second-level incremental backup based on a previous first-level incremental backup, a third-level incremental backup based on a previous second-level incremental backup, and so on.

The backup type (full or incremental) applies to the entire backup specification and only to filesystem objects.

The backup type is ignored for zero downtime backup sessions (split mirror or snapshot backup). It is set to full.

To combine full and incremental backups, make sure that the backup object has exactly the same:

- client name
- drive/mountpoint
- description

  The description can be set for the whole backup specification or for a specific object. Refer to “Backup Specification Options” on page 236 and “Object Options” on page 239.
- owner

  Backup ownership can be set for the whole backup specification. Refer to “Ownership: Who Will Be Able to Restore?” on page 235.

Backup Types

- Full backup

  A full backup consists of all backup objects, even if they have been backed up before. The first backup of an object is always a full backup. Any subsequent backup will be completed as full if no protected full backups with the same ownership are available at the backup time.
- Incr backup

  This backup type is based on any previous, still protected backup chain, either a full or an incremental backup. An incremental backup includes only the files that have changed since the last still protected

18179
Paula

backup. Even if the previous backup was an incremental (**Incr** or **Incr1**, **Incr2**, ...) backup, the subsequent incremental backup includes only those files that changed in the interim.

- Incr1 backup

  This backup type refers to the most recent still protected full backup with the same ownership. It does not depend on any previous incremental backups. The files that have changed since the most recent still protected full backup are included in the backup.

- Incr2 backup

  This backup type refers to the most recent still protected full backup, provided that there is no Incr1 done afterwards. If there are several Incr1 backups available, it refers to the most recent one. All files that have changed since the reference backup was done are backed up.

- Incr1-9 backup

  The description above explains the concept of incremental levels, which can be extended up to Incr9.

Table 5-1 shows the relative referencing of backup runs with various backup types. See the text following the table for a full explanation.

Table 5-1 **Relative Referencing of Backup Runs**

1	Full	<----	Incr1				
2	Full	<----	<----	<----	Incr2		
3	Full	<----	Incr1	<----	Incr2		
4	Full	<----	Incr				
5	Full	<----	Incr1	<----	Incr		
6	Full	<----	Incr1	<----	Incr2	<----	Incr
7	Full	<----	Incr1	<----	Incr	<----	Incr
8	Full	<----	Incr1	<----	Incr3		
9	Full	<----	Incr1	<----	Incr2	<----	Incr3
10	Full	<----	<----	<----	Incr2	<----	Incr3
11	Full	<----	<----	<----	<----	<----	Incr3

13.175
Paula

How to Read Table 5-1

- The rows in Table 5-1 are independent of each other and show different situations.
- The age of the backups increases from right to left, so that the far left is the oldest and the far right is the most recent backup.
- The full and IncrX represent still-protected objects of the same owner. Any existing IncrX that is not protected can be used for restore, but is not considered for referencing on subsequent backup runs.

Examples:

- In the second row, there is a full, still protected backup and an Incr2 is running. There is no Incr1, so the backup is executed as an Incr1.
- In the fifth row, there is a full backup, an Incr1 and another incremental is running. Data Protector references the currently running backup to the previous incremental, that is Incr1.
- In the eighth row, the Incr3 is executed as Incr2, and in the eleventh row, the Incr3 is executed as Incr1.

How to Select the Backup Type

If you perform an interactive backup, you are prompted to select the backup type. When scheduling a backup, you specify the backup type in the `Schedule Backup` dialog. You can, for example, create a schedule that runs the same backup specification as full on Saturday and as Incr1 on all working days.

Backup Type and the Restore Process

Keep in mind that full backups enable a simple and efficient restore, but require many media that can hold multiple versions of the entire backed up data. The time required to complete a backup is rather long. Incremental backups require fewer media resources, but have a more complex restore algorithm. Compare the following two examples:

1. `full ; Incr ; Incr ; Incr ; Incr  (-> time)`

   This example requires a shorter backup time and the media space required is lower. The restore process is more complex; many media need to be accessed, and the required time is longer if you want to restore to the state of the last Incr.

2. `full ; Incr1 ; Incr1 ; Incr1 ; Incr1  (-> time)`

   This example requires more time for backup and the media space consumption is a bit higher if compared to the first example. The restore process is simple; few media are needed, and the time spent on performing a restore is shorter than in the first example.

18.177

Using Backup Templates

Overview

Data Protector backup templates are a powerful tool that can help you simplify your backup configuration. A template has a set of clearly specified options for a backup specification, which you can use as a base for creating and modifying backup specifications. Data Protector enables you to apply a group of options offered by the template.

A template can be used in two ways:

- It can be used to create a new backup specification.
- It can be applied to existing backup specifications to modify these specifications.

Backup templates are created and modified similarly to backup specifications, except that objects and the backup application configuration are not selected within the backup template.

Data Protector Default Backup Templates

Data Protector offers you default templates for different types of data (Filesystem, Oracle8/SAP, and so on) to configure a filesystem or an application backup. The templates provide typical settings, which can be used as a basis for your backup specifications.

Blank Backup Templates

In blank backup templates, such as `Blank Filesystem Backup`, `Blank Informix Backup`, and so on, there are no objects or devices selected. Backup specification options and object options have Data Protector default values, and there is no backup schedule. You can separately select the `Load balanced` option, enabling Data Protector to automatically balance the usage of devices selected for the backup.

Options Offered by Templates

When using a backup template for creating or modifying a backup specification, you can select or deselect options offered by the template.

18.176
Paula

Figure 5-22 **Options Offered by Templates**

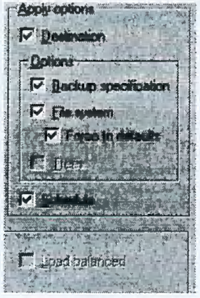

Apply options
☑ Destination
Options
☑ Backup specification
☑ File system
☑ Force to defaults
☐ Trees
☑ Schedule

☐ Load balanced

Destination Backup device settings specified in your template apply to your backup specification.

Backup specification Backup specification options specified in the template apply to your backup specification.

Filesystem Filesystem options specified in the template apply to all filesystem objects of your backup specification.

Force to defaults Filesystem object options specified in the template apply to all filesystem objects of your backup specification. These are the options in the `Backup Object Summary` page.

Trees Trees options specified in the template apply to your backup specification.

Schedule Schedule settings specified in the template apply to your backup specification.

Once you have applied the template options, you can still modify your backup specification and change any setting.

For more information on these options, refer to "Using Backup Options" on page 225.

Load balanced This option directs how data is distributed to devices.

See "Load Balancing: Balancing the Usage of Backup Devices" on page 232 for details about the `Load balancing` options.

18.145
Paula

Using a Backup Template When Creating a New Backup Specification

When creating a new backup specification, Data Protector offers you a set of backup templates, either default templates or templates you have created. Select an appropriate template, or optionally, select or deselect some groups of options, and then proceed with the Backup wizard.

To create a backup specification without predefined settings, select `Blank Filesystem Backup`.

Figure 5-23 **Using Templates When Creating New Backup Specifications**

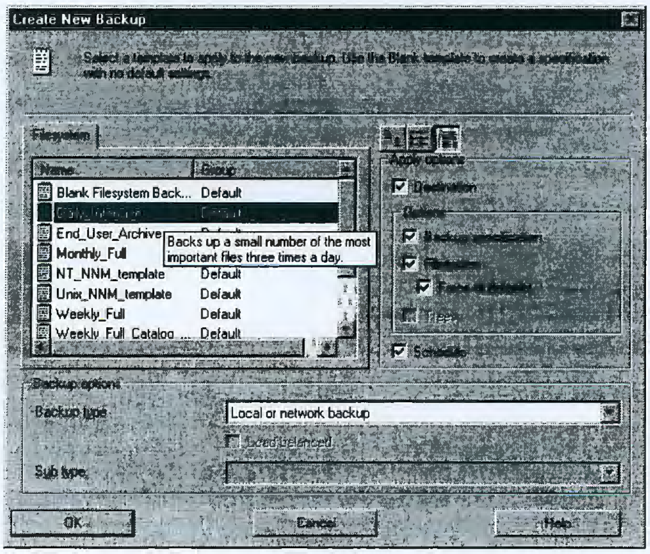

Create New Backup

Select a template to apply to the new backup. Use the Blank template to create a specification with no default settings.

Name	Group
Blank Filesystem Back...	Default
[illegible]	Default
End_User_Archive	Default
Monthly_Full	
NT_NNM_template	Default
Unix_NNM_template	Default
Weekly_Full	Default
Weekly_Full_Catalog ...	Default

Backs up a small number of the most important files three times a day.

Apply options
☑ Destination
Options
☑ Backup specification
☑ Filesystem
☑ Force all defaults
□ Trees
☑ Schedule

Backup options
Backup type: Local or network backup
□ Load balanced
Sub type:

Applying a Backup Template

Data Protector allows you to apply a backup template to saved backup specifications. When applying a template to backup specifications, you can select which option groups should be applied. Refer to "Options Offered by Templates" on page 216.

The result of applying an option group is that all related options in this group are set to the state specified in the template.

18.174
Paula

To apply a template to backup specifications, right-click the backup specification and click `Apply Template`. The `Apply Template` window appears, in which you apply the desired options. For detailed steps, refer to the online Help index keyword “applying backup templates”.

Figure 5-24 **The Apply Template Dialog Box**

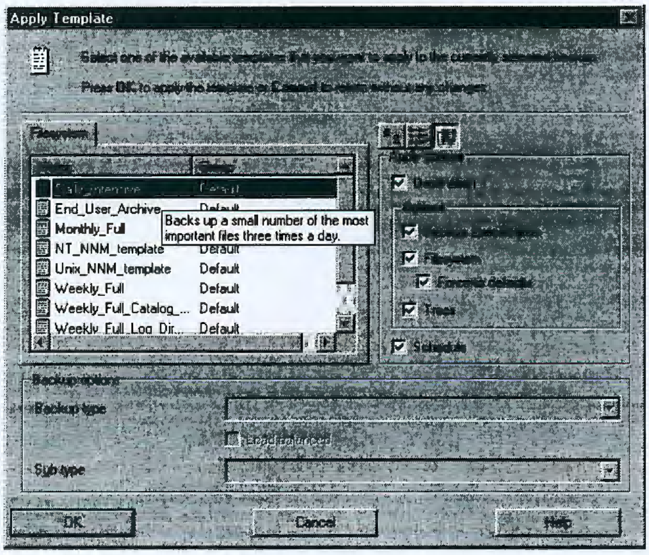
Apply Template

Select one of the available templates that you want to apply to the currently selected backup specs.

Press OK to apply the template or Cancel to return without any changes.

Name	Group
Daily_Intensive	Default
End_User_Archive	Default
Monthly_Full	
NT_NNM_template	Default
Unix_NNM_template	Default
Weekly_Full	Default
Weekly_Full_Catalog_...	Default
Weekly_Full_Log_Dir...	Default

Backs up a small number of the most important files three times a day.

Integration Backup Specification

To apply a template to an integration backup specification, the backup specification you would like to apply should not be opened in the Results Area. If you first click on the backup specification to open it, and then try to apply the template to this backup specification, the `Apply Template` option will not be available.

IMPORTANT If you select the `Force to defaults` option, the options specified in your template apply to all filesystem objects of your backup specification for which you changed options in the `Backup Object Summary` page.

18173
Paula

Creating a New Template

You can create new templates and use them for creating or modifying backup specifications.

To create a new template, use the Backup context. For detailed steps, refer to the online Help index keyword “creating backup templates”.

For more information on individual options, refer to “Using Backup Options” on page 225.

Figure 5-25 **Creating a New Template**

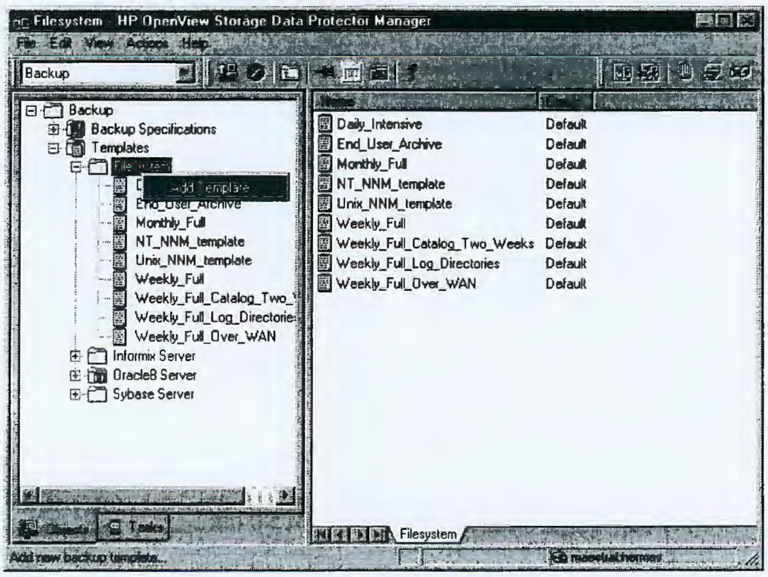

Daily_Intensive	Default
End_User_Archive	Default
Monthly_Full	Default
NT_NNM_template	Default
Unix_NNM_template	Default
Weekly_Full	Default
Weekly_Full_Catalog_Two_Weeks	Default
Weekly_Full_Log_Directories	Default
Weekly_Full_Over_WAN	Default

Add new backup template...

Modifying an Existing Template

You can modify Data Protector default templates, as well as templates that you have created.

To modify an existing template, open the properties of the template. For detailed steps, refer to the online Help index keyword “modifying backup templates”.

18.172
Paula

For more information on individual options, refer to "Using Backup Options" on page 225.

18.171
Paula

Groups of Backup Specifications

Data Protector offers you the ability to organize backup specifications into different groups. The purpose of grouping is to organize the specifications of multiple backups.

For example, backup specifications for *"Corporation X"* can be classified into three different groups:

- **USER_FILES**: This group contains backup specifications that perform weekly, full backups for all users in each of the ten departments. This is the main production backup.
- **SERVERS_DR**: This group contains backup specifications for the company's servers to prepare for disaster recovery. Each time a new server is installed or an existing server is upgraded, a new backup specification is created and added to this group.
- **END_USER**: This group is used to save backup specifications that are made as a result of an end-user request. For example, end users who want to free up some disk space have to archive their own hard disk first.

See Figure 5-26 on page 223.

Such a configuration can result in many backup specifications, often as many as 50, which are hard to manage if they are viewed together. Grouping the backup specifications into meaningful groups can facilitate finding and maintaining single backup specifications. This allows you to apply common options settings from a template to the entire group.

For example, if you want to change the list of devices to all backup specifications in the group, you can selectively apply the device settings of a template.

18 170
Paula

Figure 5-26 **Example of Backup Specification Groups**

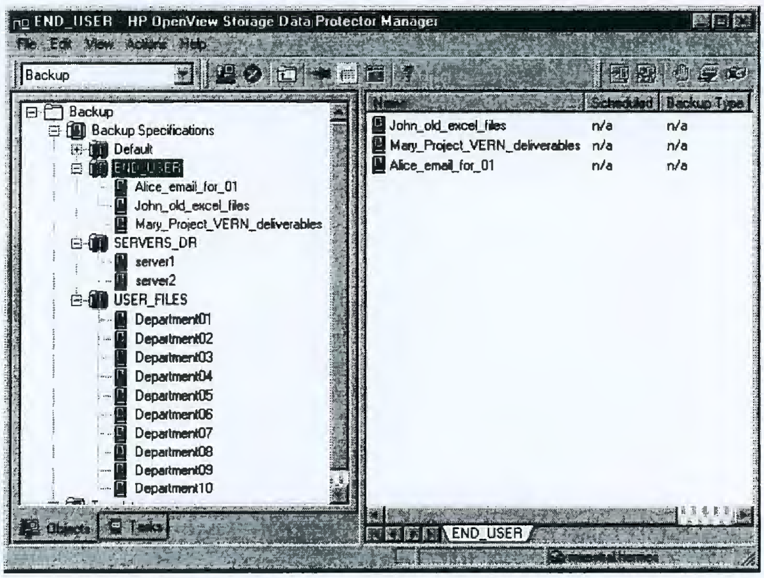
END_USER - HP OpenView Storage Data Protector Manager
File Edit View Actions Help
Backup
Objects
Tasks
END_USER

- Backup
  - Backup Specifications
    - Default
    - END_USER
      - Alice_email_for_01
      - John_old_excel_files
      - Mary_Project_VERN_deliverables
    - SERVERS_DR
      - server1
      - server2
    - USER_FILES
      - Department01
      - Department02
      - Department03
      - Department04
      - Department05
      - Department06
      - Department07
      - Department08
      - Department09
      - Department10

Name	Scheduled	Backup Type
John_old_excel_files	n/a	n/a
Mary_Project_VERN_deliverables	n/a	n/a
Alice_email_for_01	n/a	n/a

How to View and Create Groups

The following procedure describes how to view the available backup groups and how to create a new one:

1. In the `Data Protector Manager`, switch to the `Backup` context.
2. In the `View` menu, click `By Group`. The list of available backup groups appears under the `Backup Specifications` item. Clicking a group lists the backup specifications within that group.
3. Right-click the `Backup Specifications` item, and then click `Add Group`. The `Add New Group` dialog box appears.
4. In the `Name` text box, enter a name for your new group, and then click `OK`. Your new group will appear under the `Backup Specifications` item.

How to Save a Backup Specification in a Group

While saving the backup specification, you are also adding it to a group of backup specifications. If you do not specify the name of the group, a backup specification will be added to the `Default` group.

18.169
Paula

How to Delete a Group

Before deleting a group, you have to empty it first. One way of doing this is to move the backup specifications into other groups. See online Help for details.

13.168
Paula

Using Backup Options

Data Protector offers a comprehensive set of backup options to help you fine-tune your backups. All options have default values that are appropriate for most cases.

The availability of backup options depends on the type of data being backed up. For example, not all backup options available for a filesystem backup are available for a disk image backup. Common and specific application options for Exchange, SQL, and so on, are described in the *HP OpenView Storage Data Protector Integration Guide*.

Additionally, the `User defined variables` function lets you specify a variable name and its value for flexible operation on some platforms and integrations, for example, for backing up MPE platforms.

The backup options can be grouped as follows:

- Backup specification options, such as **Load balancing, Ownership**, and **pre-** and **post-exec** options for the whole backup specification.
- Object options specifying how different backup objects, such as filesystems or disk images, are backed up.

  It is important to understand that object options can be set on two levels. First, you can set the *default object options* for all filesystems and for all disk image objects in the backup specification separately. Then you can set them differently *for a specific object*. These settings will override the defaults. For example, to compress data from all clients except for one with a slow CPU, set `compression` to `ON` when setting filesystem options. Then, select the slow client and set `compression` to `OFF` for this client.
- Device options define the behavior of backup devices. If you do not set the device options, the values are read from the device definition.
- Schedule options define the backup type, network load, and data protection for each individual or periodic scheduled backup. With split mirror or snapshot backup, in the case of ZDB disk or ZDB disk/tape backups (instant recovery enabled), you specify also the `Split mirror/snapshot backup` option.

18.167
Paula

For split mirror and snapshot backups, the backup type is ignored (it is set to full). Data protection that is specified in the `Schedule Backup` dialog overrides protection settings anywhere else in the backup specification.

See Figure 5-27 for a graphic scheme of how some of these options work. You can use backup templates to apply the same group of options to a number of backup specifications. Applying a template changes the backup specification according to the template. If you later change the template, you have to apply it again if you want the changes to take effect.

You can selectively apply schedule, device, and object options and the private exclusion list.

See also "Using Backup Templates" on page 216 for details about the backup templates.

1816
Paula

Figure 5-27 **Backup Options**

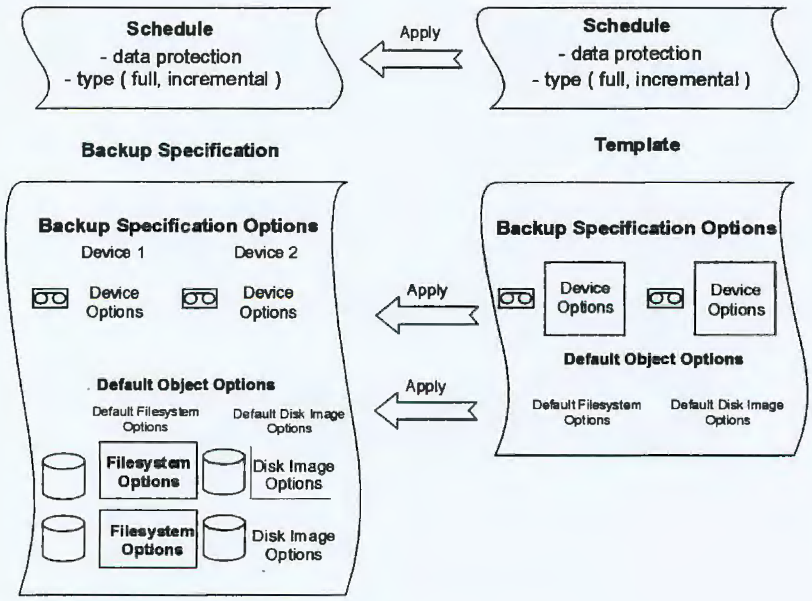

Most Frequently Used Backup Options

This section describes the options that are most likely to be modified according to specific backup policies. These are the following:

- "Data Protection: Specifying How Long Data Is Kept on the Media" on page 228
- "Catalog Protection: How Long Info Is Kept in the Database" on page 230
- "Logging: Changing Details About Data Stored in the Database" on page 231
- "Load Balancing: Balancing the Usage of Backup Devices" on page 232
- "Ownership: Who Will Be Able to Restore?" on page 235

Data Protection: Specifying How Long Data Is Kept on the Media

Configuring protection policies is extremely important for the safety of your data and for successful management of your environment. See the *HP OpenView Storage Data Protector Concepts Guide* for more detailed information on how to define these policies.

Based on your company data protection policies, you have to specify how long your backed up data is kept on the medium. For example, you may decide that data is out of date after three weeks and can be overwritten during a subsequent backup.

NOTE

Due to operating system limitations, the latest protection date that can be set is Jan 18th, 2038.

You can specify data protection in different places. Different combinations are available, depending on whether you are running an interactive backup, starting a saved backup specification, or scheduling a backup. The default value is `Permanent`.

- Interactive backups

  When configuring an interactive backup, you can change the default data protection for the entire backup. See Figure 5-28 on page 229. Additionally, you can specify different data protection periods for individual backup objects. The protection that is specified on the backup object level overrides the default protection setting. See Figure 5-29 on page 230.

- Backups using a saved backup specification

  When starting saved backups using the GUI, the data protection is applied as described for interactive backups.

  When starting saved backups using the CLI, you can also specify data protection. This will override all data protection settings in the backup specification.

- Scheduled backups

18/04
[illegible]

You can specify a different period of protection for each individual or periodic scheduled backup. The data protection specified in the `Schedule Backup` dialog overrides all other data protection settings in the backup specification. If you leave the default protection, data protection is applied as described for interactive backups.

On how to specify data protection, refer to the online Help index keyword "specifying data protection".

NOTE If you apply a backup template to an existing backup specification and select the `Filesystem` and/or `Schedule` options, the protection settings from the template will replace the previous data protection settings in the respective parts of the backup specification. For more information, refer to "Options Offered by Templates" on page 216.

Figure 5-28 **Backup Options: Protection**

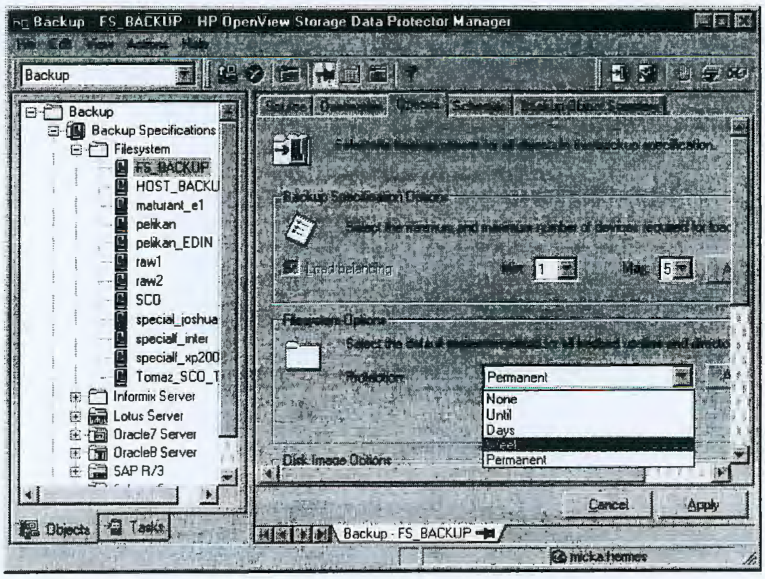

18163
faulr

Figure 5-29 **Backup Object Properties - Options: Protection**

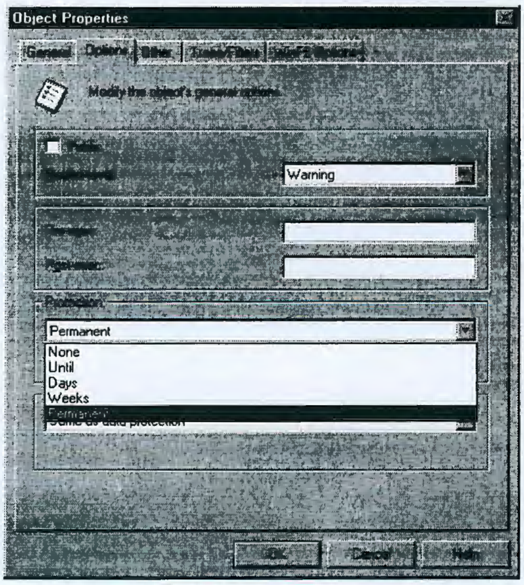

Catalog Protection: How Long Info Is Kept in the Database

Besides the `Protection` option, which controls how long data is protected on media, you can set the `Catalog Protection` option, which controls the time for which information about backed up files and directories is kept in IDB. Catalog protection and data protection can be set independently. Catalog protection has no effect if the log level is `Log None`.

The default value for catalog protection is `Same as data protection`. This means that you can browse and select files or directories as long as the media are available for restore.

NOTE If data protection expires, the catalog protection is cancelled. That is, when the data protection ends and a medium is overwritten, the catalogs for the objects are removed regardless of the catalog protection.

18.163 Paula

Even when catalog protection expires, you are still able to restore, but you must specify filenames manually.

Be aware that catalog protection, together with logging level, has a very big impact on the growth of the IDB. Therefore, it is very important to define a catalog protection policy appropriate to your environment. Refer to the IDB section in the *HP OpenView Storage Data Protector Concepts Guide* for more information on catalog protection and usage recommendations.

NOTE Due to operating system limitations, the latest protection date that can be set is Jan 18th, 2038.

Logging: Changing Details About Data Stored in the Database

The logging level determines the volume of detail on files and directories written to the IDB during backup. Note that you can restore your data regardless of the logging level used during backup. Data Protector provides the following four logging levels:

Table 5-2

Log All	This is the default logging level. All detailed information about backed up files and directories (names, versions, and attributes) is logged to the IDB. You can browse directories and files before restoring and in addition look at file attributes. Data Protector can fast position on the tape when restoring a specific file or directory.
Log Files	When this logging level is selected, detailed information about backed up files and directories (names and versions) is logged to the IDB. You can browse directories and files before restoring, and Data Protector can fast position on the tape when restoring a specific file or directory. The information does not occupy much space, since not all file details (file attributes) are logged to the database.

13.161
Paula

Table 5-2

Log Directories	When this logging level is selected, all detailed information about backed up directories (names, versions, and attributes) is logged to the IDB. You can browse only directories before restoring. However, during the restore Data Protector still performs fast positioning because a file is located on the tape near the directory where it actually resides. This option is suitable for filesystems with many auto-generated files, such as news and mail systems.
No Log	When this logging level is selected, no information about backed up files and directories is logged to the IDB. You will not be able to search and browse files and directories before restoring.

The different logging level settings influence the IDB growth, backup speed, and the convenience of browsing data for restore.

To be able to select the most appropriate logging level setting, it is important to understand the consequences. Refer to the *HP OpenView Storage Data Protector Concepts Guide* for more information on logging level and usage recommendations.

Load Balancing: Balancing the Usage of Backup Devices

What Is Load Balancing?

By default, Data Protector automatically balances the usage of backup devices specified for backup. This is also called load balancing, and it ensures equal usage of the devices. When you run backup with the `Load Balancing` option, Data Protector uses devices in the order they are specified in the load balanced backup specification.

NOTE

If you disable the `Load Balancing` option, you have to select the backup device which is used to back up each object in the backup specification. If a device becomes unavailable, then the objects that should be backed up to the device will not be backed up.

See the *HP OpenView Storage Data Protector Concepts Guide* for more information on load balancing.

13.160
Paula

When to Use Load Balancing

It is recommended that you use the `Load Balancing` option when you want to back up a large number of objects to a number of available devices, and you would like Data Protector to keep all the devices busy all of the time. You should use `Load Balancing` to minimize the impact of unavailable devices on the backup. A device may become unavailable because it:

- failed during a backup
- stopped during a backup
- is in use by another session
- cannot be started at all

When Not to Use Load Balancing

Deselecting the `Load Balancing` option is recommended when

- you want to back up a small number of objects
- objects are backed up on simple devices, such as DDS
- you want to manually select the devices to which objects will be backed up
- you want to know on which medium/media objects will be backed up

How Are the Parameters Used?

The `Load Balancing` option has `MIN` and `MAX` parameters:

`MIN` specifies the minimum number of backup devices out of the list of devices in the backup specification that can be used immediately. It means that they are not used by some other backup session and that you have enough licenses.

`MAX` specifies the maximum number of devices that are used at the same time, even if there are more devices defined in the backup specification. The rest are used if needed.

For example, assume that there are four devices in the backup specification and `MIN` and `MAX` are both configured at two. The backup session will queue until any of those two devices can be used. If any of them fail, one of the two devices in reserve will be used.

How Are Objects Assigned to an Available Device?

The first device from the list of devices is started. The number of selected objects for a device is defined by its `concurrency`. The next device is started and objects are selected until there are no more objects in the list or the maximum number of devices are running.

Objects to be backed up are assigned according to the following criteria:

18.154
Paula

- Objects that reside on the client connected to the backup device have a higher priority.
- Objects are selected so that the number of Disk Agents per client is kept as low as possible.

The size of objects does not play a role in assigning an object to a device.

If a device becomes unavailable, the following happens:

- All objects backed up to the device before the failure time are actually backed up.
- All objects that are being backed up to the device at failure time are aborted.
- All objects pending to be backed up to the device will be backed up to some other available device specified in the backup specification, if the maximum number of devices has not been used.

Example

For example, assume that there are 100 objects configured for backup to four devices with concurrency set to three and with load balancing parameters `MIN` and `MAX` both configured at two. If at least two devices are available, the session will start with three objects being backed up in parallel to each of the first two available devices. The other 94 objects will be pending and will not be assigned to a particular device at that time.

Once a backup of a particular object is done, the next pending object is started and assigned to the device that has less than three concurrent objects being backed up. Load balancing ensures that the two devices are running in parallel as long as there are still pending objects to be backed up. If a device fails during backup, one of the two devices in reserve is used. The objects that were being backed up to the failed device are aborted, while the next three pending objects are assigned to the new device. This means that each failure of a device can cause a maximum of three objects to be aborted, provided that other devices are available for the backup session to continue.

The following rules should be considered when applying device options from a template:

- If the load balancing option is not selected in the template, the device options are not used with the backup specification.
- If the load balancing option is selected in both the template and the backup specification, the device options are applied.

18.158
Paula

- If load balancing is only selected in the template, the device options are applied only if the backup specification has no devices.

For more information on failed backups, refer to "Managing Failed Backups" on page 263.

Ownership: Who Will Be Able to Restore?

Who Is a Backup Session Owner?

A user who starts an interactive backup is the session owner. If a user starts an existing backup specification without modifying it, the backup session is not considered interactive.

If a modified backup specification is started by a user, the user is the owner unless the following conditions apply:

- The user has the `Switch Session Ownership` user right.
- The backup session owner is explicitly defined in the backup specification, where the username, group or domain name, and the system name are specified. In that case, the backup session owner is the user specified in the backup specification.

If a backup is scheduled on a UNIX Cell Manager, the session owner is `root:sys`, unless the above conditions apply.

If a backup is scheduled on a Windows Cell Manager, the session owner is the user specified at installation time, unless the above conditions apply.

Who Can Restore a Private Object?

The following users can restore a private object:

- Members of the `Admin` and `Operator` user group.
- The backup session owner who has the `Start Restore` user right. Other user rights may be required, such as `Restore to Another Client`.
- Users who have the `See Private Objects` user right.

Why Change the Backup Owner?

Sometimes, you may want to change the backup owner. For example, if the administrator configures and schedules a backup specification, operators are allowed to run it, but they cannot modify or save it. If the `Private` backup option is set for all objects, the operators are not able to restore anything, but can still manage backups and restart failed sessions.

Changing the owner works only for saved backup specifications. If the backup configuration is changed and not saved, the backup is treated as an interactive backup and the owner is not changed. This could result in a different kind of backup than expected. For example, if you interactively start an incremental backup and you are not the owner of the full backup, you will get another full backup instead of an incremental one.

List of Data Protector Backup Options

This section describes three sets of backup options. The options are ordered alphabetically within each set.

Backup Specification Options

Setting Options for a Backup Specification

1. Select the backup specification whose options you want to set.
2. Click the `Options` tab.
3. Under `Backup Specification Options`, click `Advanced`. The `Backup Options` window appears.
4. Select the options for `General`, `Other`, and `Clustering`. The `EMC` and the `StorageWorks` tabs appear only if you have the respective devices connected and configured.

   Ignore `Clustering` if you do not have the `MC/Service Guard` or the `Microsoft Cluster Server` installed and configured.
5. Click `OK` to confirm and exit the Backup Options window. Refer to online Help for details.

Figure 5-30 **Backup Specification Options - General and Other**

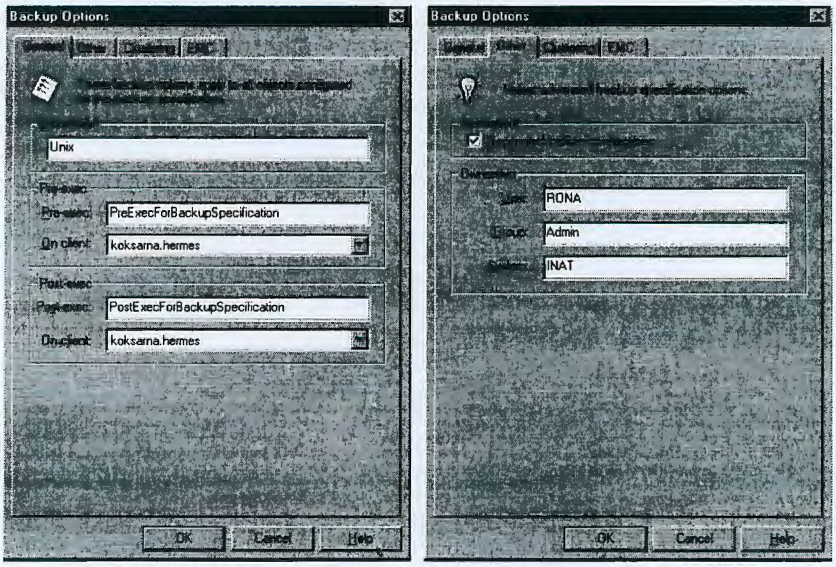

Backup Options

Unix

Pre-exec

Pre-exec: PreExecForBackupSpecification

On client: koksarna.hermes

Post-exec

Post-exec: PostExecForBackupSpecification

On client: koksarna.hermes

Backup Options

☑

Ownership

RONA

Group: Admin

INAT

Available Backup Specification Options

Description

You can type in any text to describe the purpose or contents of the backup specification. This text has no effect on the backup session.

Load Balancing

If this option is set, Data Protector dynamically assigns backup objects to available devices. This means that devices are evenly used, and if one fails, a backup continues on other available devices. If it is not set, the backup objects are backed up to devices assigned to them in the exact order specified.

The default value is ON.

See "Load Balancing: Balancing the Usage of Backup Devices" on page 232 for more information.

Ownership

The session owner is the user who started the interactive backup, unless the owner is specified in the backup specification. Otherwise, the owner is:

13.155

- root on UNIX Cell Managers
- the user specified at installation time on Windows Cell Managers

The default value is not specified.

See “Ownership: Who Will Be Able to Restore?” on page 235 for more information.

You can change the session owner by using the following steps:

1. In the HP OpenView Storage Data Protector Manager, switch to the Backup context.
2. Double-click Backup Specifications, then right-click on the backup specification you want to modify.
3. Choose Properties, Options, then, under Backup Specification Options, choose Advanced. Choose the Other tab.
4. Modify session ownership as necessary. Use uppercase on Windows systems.

NOTE Make sure to specify the information as it was specified when the user was configured.

Pre-Exec

The command specified in this field is run on a specified client before any object is backed up. If the client is not defined, the command is run on the Cell Manager.

See “Pre- and Post- Exec Commands on Windows Systems” on page 251 for details of specifying pre-exec commands on Windows.

See “Pre- and Post- Exec Commands on UNIX Systems” on page 257 for details of specifying pre-exec commands on UNIX.

See Appendix, “Examples of Pre-Exec and Post-Exec Commands for UNIX,” on page A-20 for some sample scripts on UNIX.

The default value is not specified.

Post-Exec

The command specified in this field is run on a specified client after all objects have been backed up. If the client is not defined, the command runs on the Cell Manager.

See “Pre- and Post- Exec Commands on Windows Systems” on page 251 for details of specifying post-exec commands on Windows.

See “Pre- and Post- Exec Commands on UNIX Systems” on page 257 for details of specifying pre-exec commands on UNIX.

See Appendix, “Examples of Pre-Exec and Post-Exec Commands for UNIX,” on page A-20 for some sample scripts on UNIX.

Default value is `not specified`.

Reconnect Broken Connections

When this option is set, Data Protector reconnects the Backup Session Manager and Disk Agents or Media Agents in the event of short-term network problems. Otherwise, the session is aborted.

This setting is useful if you have the Cell Manager on one LAN and Disk Agents or Media Agents on another. Assuming that the connection between these two LANs is unreliable (WAN connections), Data Protector tries to reconnect for 1200 seconds. This can be set in the `omnirc` variable `OB2RECONNECT_RETRY`.

The default value is `OFF`.

Object Options

Setting the Filesystem Options

1. Select the backup specification and from the `Options` property page, under `Filesystem Options`, click `Advanced`.
2. Select the options to be set from the `Options`, `Other`, `WinFS Options`, or `Netware Options` tabs.

NOTE

On the `Options` tab, if specifying Pre- and/or Post- exec command names, you may or may not have to specify the full paths for the commands.

See “Pre- and Post- Exec Commands on Windows Systems” on page 251 for details of specifying pre-exec commands on Windows.

See “Pre- and Post- Exec Commands on UNIX Systems” on page 257 for details of specifying pre-exec commands on UNIX.

3. Click `OK` to confirm and exit this dialog box.

   See online Help for specific help on each option.

18.153

Figure 5-31 **Filesystem Options**

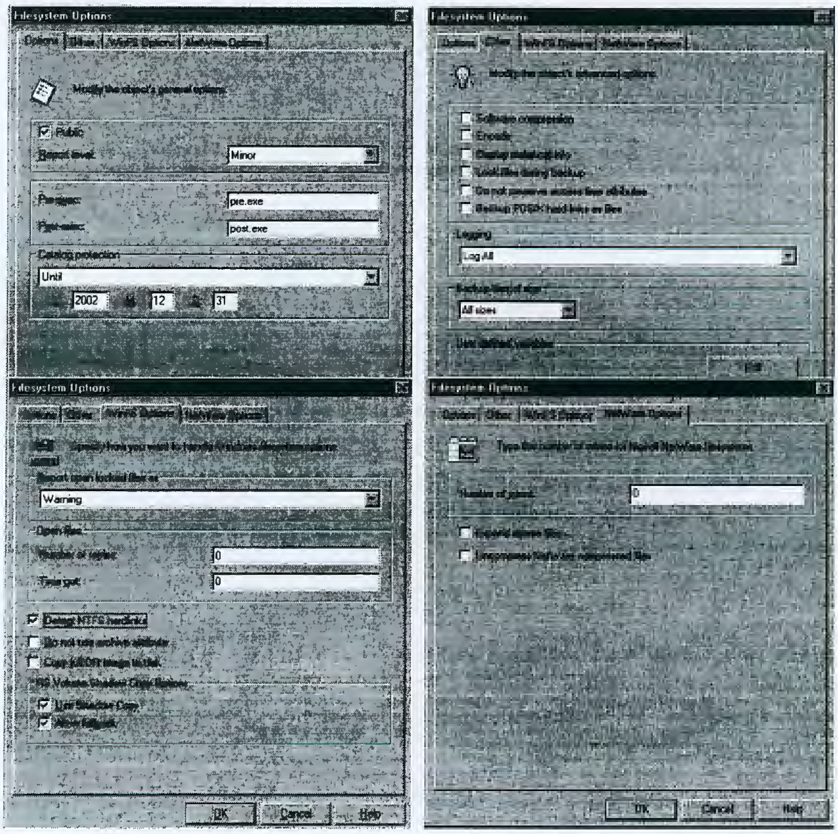

Filesystem Options

Options | Other | WinFS Options | NetWare Options

Modify the object's general options.

☑ Public

Report level: Minor

Pre-exec: pre.exe

Post-exec: post.exe

Catalog protection

Until

2002 12 31

Filesystem Options

Options | Other | WinFS Options | NetWare Options

Modify the object's advanced options.

□ Software compression
□ Encode
□ Display statistical info
□ Lock files during backup
□ Do not preserve access time attributes
□ Backup POSIX hard links as files

Logging

Log All

Backup files of size

All sizes

User defined variables

Edit

Filesystem Options

Options | Other | WinFS Options | NetWare Options

Specify how you want to handle Windows filesystem options.

Report open locked files as

Warning

Open files

Number of retries: 0

Time-out: 0

☑ Detect NTFS hardlinks
□ Do not use archive attribute
□ Copy full DR image to disk

MS Volume Shadow Copy Options

☑ Use Shadow Copy
☑ Allow fallback

OK | Cancel | Help

Filesystem Options

Options | Other | WinFS Options | NetWare Options

Type the number of values for Novell NetWare filesystems.

Number of values: 0

□ Expand sparse files
□ Uncompress NetWare compressed files

OK | Cancel | Help

Setting the Disk Image Options

1. Select the backup specification.
2. Select the `Options` property page.
3. Under `Disk Image Options`, click `Advanced`.
4. Click either the `Options` or the `Other` tab, and specify the options as desired. For a description of each option, click `Help` in the dialog box.

NOTE On the `Options` tab, if specifying Pre- and/or Post- exec command names, you may or may not have to specify the full paths for the commands.

See “Pre- and Post- Exec Commands on Windows Systems” on page 251 for details of specifying pre-exec commands on Windows.

See “Pre- and Post- Exec Commands on UNIX Systems” on page 257 for details of specifying pre-exec commands on UNIX.

Figure 5-32 **Disk Image Options**

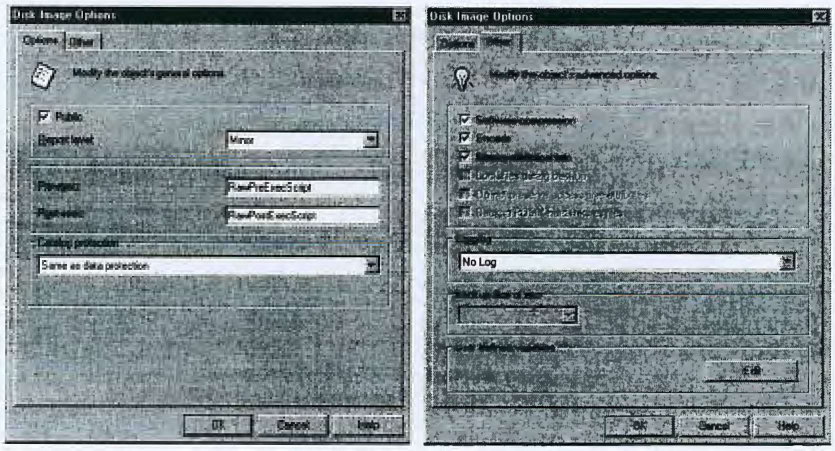

Disk Image Options

☑ Public

Report level: Minor

RawPreExecScript

RawPostExecScript

Same as data protection

Disk Image Options

No Log

Options
Other
Modify the object's general options
Modify the object's advanced options
Edit
OK
Cancel
Help

5. Click OK to confirm and exit the dialog box.

Setting the Object Specific Options

1. Select the backup specification whose options you want to set.
2. Select the Backup Object Summary property page.
3. Right-click the backup object, then select Properties. The contents of the Object Properties window depend on the type of backup object you selected. This can be a UNIX filesystem, a Windows filesystem, or a UNIX disk image.

   The Object Properties window for a **Windows filesystem** contains the General, Options, Other, Trees/Filters, and the WinFS Options tabs. Options, Other, and WinFS Options are the same as shown in Figure 5-31, while General and Trees/Filters are shown in Figure 5-33.

Figure 5-33 **Object Properties - General and Trees/Filters**

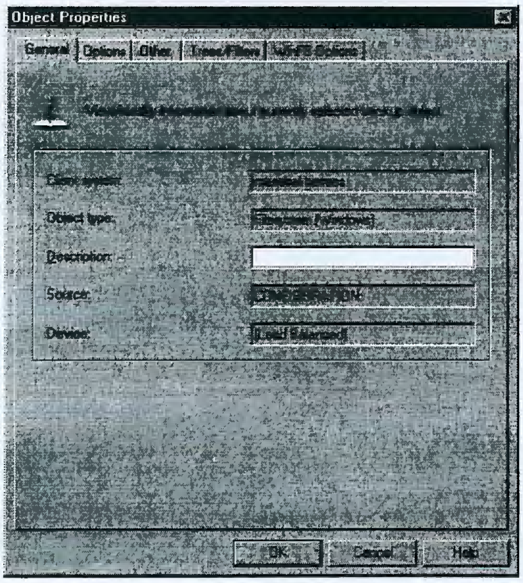

Client system: [illegible]

Object type: [illegible]

Description:

Source: [illegible]

Device: [illegible]

The Object Properties window for a **UNIX filesystem** contains the General, Options, Other, and the Trees/Filters tabs. Options and Other are the same as shown in Figure 5-31, while General and Trees/Filters are the same as in Figure 5-33, except that Object type is described as Filesystem [UNIX].

The Object Properties window for a **disk image** object contains the General, Options, Other, and the Disk Image Options tabs. Options and Other are the same as for the backup specification shown in Figure 5-32, while General and Disk Image Options are shown in Figure 5-34.

13150
Paula

Figure 5-34 **Object Properties - General and Disk Image Options**

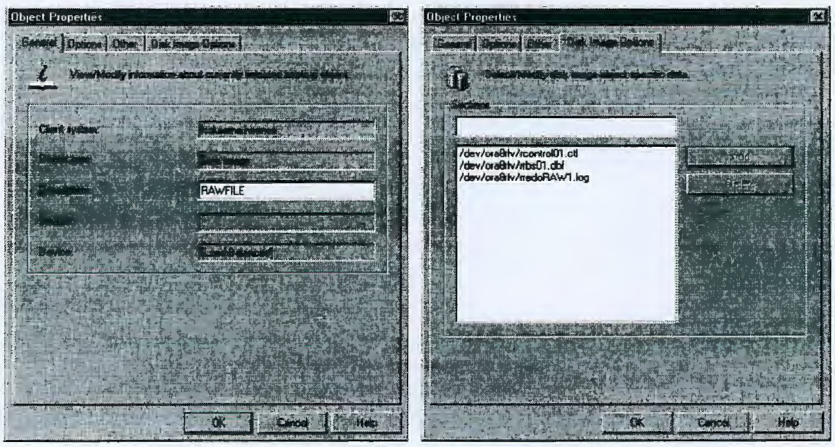

RAWFILE

/dev/ora8itv/rcontrol01.ctl
/dev/ora8itv/rtbs01.dbf
/dev/ora8itv/redoRAW1.log

4. Set the options and click `OK` to confirm the selection. See below for details on a particular option.

Allow Fallback (Windows-Specific Options)

If the `Use Shadow Copy` option is specified, but the shadow copy creation on the system where the VSS filesystem backup is running fails, the backup also fails by default. However, you can avoid backup failure by specifying the `Fallback to legacy` option. In this case, the backup will continue as a non-VSS backup.

Backup Files of Size

Use this option to specify the size of backed up files. You can back up `All Files` (default), `Files Bigger` Than, `Smaller` Than, or files within a specified size range in kilobytes.

Backup POSIX Hard Links as Files

This option is only relevant for UNIX filesystems.

A hard link is a directory entry that actually points to a physical file. If this option is not set, Data Protector traverses the directory trees twice. In the first traverse, a table of all hard links that point to the same files

18.129
Paula

is made. In the second traverse, only one hard link is backed up with the file contents, while all the others are backed up as hard links. The first traverse also allows Data Protector to estimate the size of the backup.

If set, Data Protector backs up the entire file contents for each hard link. Data Protector traverses the filesystem tree only once, thus significantly speeding up the backup process.

Use this option when there are no hard links in your directory. When this option is set, Data Protector cannot estimate the size of the backup or display the percentage of the backup finished.

The default value is `OFF`.

Catalog Protection

The default `Catalog Protection` value is **Same as data protection**. It can be changed by specifying the **None**, **Until**, **Days**, and **Weeks** values.

Refer to “Catalog Protection: How Long Info Is Kept in the Database” on page 230 for more information.

Do Not Preserve Access Time Attributes

When this option is not set, the access time attributes remain as they were before the backup. (They are reset to their original values after each file is backed up.) When this option is set, the access time values are set to the moment of backup.

See also “Backing Up UNIX Systems” on page 161.

The default value is `OFF`.

Do Not Use Archive Attribute (Windows-Specific Option)

Data Protector clears the archive attribute after each backup (after the file has been read). If you have other applications that make specific use of this attribute, you should use this option.

The default value is `OFF`.

Detect NTFS Hardlinks (Windows-Specific Option)

This option is similar to `Backup hardlinks as files` except that it is only valid for NTFS and the default value is `OFF`, meaning that hardlinks are backed up as ordinary files. The default value is `OFF` because the NTFS hardlinks are not often used and setting this option decreases backup performance.

Encode

Open Systems and public networking make data security in large enterprises essential. Data Protector lets you encode file and disk image data so that it becomes unreadable. Data is encoded before it is transferred over the network and written to the media. Data Protector uses a fixed, built-in algorithm for this purpose.

The default value is `OFF`.

Expand Sparse Files (Novell Netware Specific)

By default, Data Protector backs up Novell NetWare sparse files in their compressed format. Though this approach speeds up the backup process, it makes it impossible to restore the Novell NetWare sparse files to any non-NetWare platform. When this option is selected, Novell NetWare sparse files are expanded before being backed up. Files backed up in this form can be restored to non-NetWare platforms.

Lock Files During Backup

If set, files are locked while being backed up, preventing them from being modified during the backup. Mandatory locking is used.

The default value is `OFF`.

With the option set for the NetWare Storage Management Service (SMS) integration, the Disk Agent attempts to open files in the Deny Write mode. If this attempt fails, the file is locked. If the file cannot be locked, then it will not be backed up.

Logging

The default logging level is **Log All**. It can be changed to **No Log**, **Log Directories**, or **Log Files**.

For more information on each logging level, see "Logging: Changing Details About Data Stored in the Database" on page 231.

Number of Retries (Novell NetWare Specific)

The number that you enter in the **Number of Retries** text box is the number of Data Protector's attempts to back up a file. If a backup cannot be made within this number of retries, Data Protector issues an error message. If you use applications that open and release files, you can use this option to increase the probability that the files are backed up.

The default value is 1.

Open Files (Windows-Specific Option)

This option controls what Data Protector does when it encounters open Windows files. If the **Number of retries** value is specified, this number defines how many times Data Protector tries to back up an open or busy file. The **Time out** value is the amount of time in seconds during which Data Protector waits before retrying to back up an open or busy file.

Protection (Data Protection)

This option enables you to set the protection level for backed up data. In this way, you prevent the backup media from being overwritten for the specified period. The `Protection` values are **None**, **Until**, **Days**, **Weeks**, and **Permanent**.

The default value is **Permanent**.

Public/Private

This option lets you set the access rights for restoring data that you back up. If a filesystem is backed up with the **Private** setting, it can be restored only by you or users who are part of the Data Protector `Admin` group.

Setting the value to **Public** lets anyone with the `Start Restore` user right restore the data.

The default value is **Private**.

Report Level

This option defines the level of errors that are reported during a backup session. Setting a level means errors of this level and higher are reported. You can choose from **Warning**, **Minor**, **Major**, and **Critical** report level.

For example, when the value `Minor` is set, all errors graded as **Minor**, **Major**, and **Critical** are reported in the `Messages` field. Messages keyed as **Normal** always appear in the `Messages` field. The default value is **Warning**.

NOTE The number of messages per backup system stored in the IDB is limited to 3000.

Report Open Locked Files As (Windows-Specific Option)

This option sets the report level for files that are opened and locked at the time Data Protector attempts to back them up. Data Protector reports such files as per the regard to the **Report Level** setting. The default value is **Warning**.

Software Compression

Data Protector can compress data on a Disk Agent client before sending it to a Media Agent client. This feature is also known as software compression. Select `Software compression` in the `Other` property page of the `Object Properties` window to enable software compression. In this way, you reduce traffic over the network, as well as number of media needed and thus improve overall backup performance. Depending on the data type, compression ranges from 30% to 70% and is based on the Lempel-Ziv 4.3 compression algorithm, which is compatible with the standard UNIX `compress` utility. The progress indication on the monitor is not accurate if this option is used.

The default value is `OFF`.

NOTE

Most modern backup devices provide built-in hardware compression that can be set when you create a device file or SCSI address in the device configuration procedure. Do not use software and hardware compression at the same time, since double compression decreases performance without giving better compression results. See *HP OpenView Storage Data Protector Installation and Licensing Guide* for details on how to enable hardware compression.

HP Ultrium LTO devices do not let you disable automatic hardware compression. Keep the default software compression value (`OFF`) when you configure an HP Ultrium LTO drive with Data Protector.

Uncompress NetWare Compressed Files (Novell Netware Specific)

By default, Data Protector backs up Novell NetWare compressed files in their compressed format. Though this approach speeds up the backup process, it makes it impossible to restore the Novell NetWare compressed files to any non-NetWare platforms. When this option is selected, Novell NetWare compressed files are uncompressed before being backed up. Files backed up in this form can be restored to non-NetWare platforms.

18.145
Paula

Use Shadow Copy (Windows-Specific Option)

When performing filesystem backup on Windows Sever 2003 systems, Data Protector uses MS Volume Shadow Copy service (VSS) for coordinating the point-in-time backup. VSS allows you to create shadow copy backups of volumes and exact point-in-time copies of files, including all open files. This means that the VSS mechanism commits all pending I/O operations and holds incoming writing requests during the preparation of a shadow copy volume. In this way all files on the filesystem are closed and unlocked during the shadow copy creation.

User Defined Backup Variables

Set user defined backup variables (a variable name and its value) to enable flexible operation on some platforms and integrations with Data Protector. For detailed steps, refer to the online Help index keyword "setting user definable backup variables".

The list of variables and their values that are configurable with Data Protector is dynamic and comes with Data Protector patches.

Figure 5-35 **User Defined Variables**

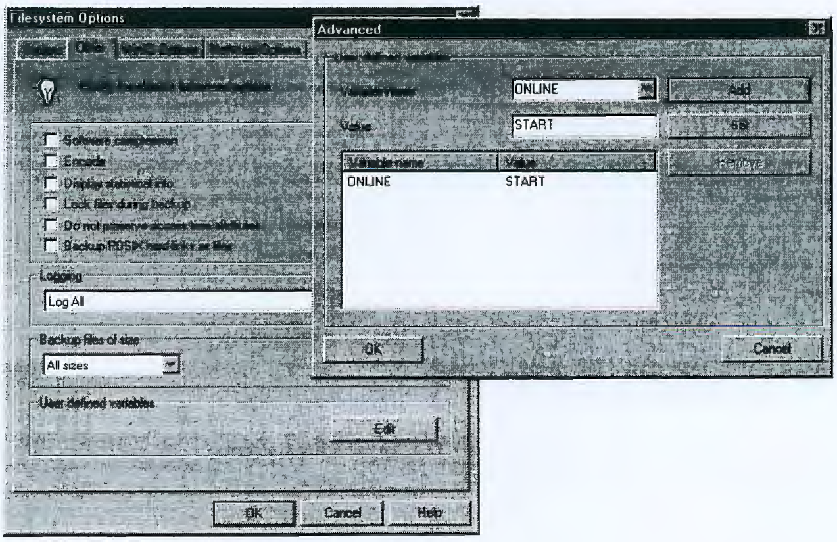

Filesystem Options

□ Software compression
□ Encode
□ Display statistical info
□ Lock files during backup
□ Do not preserve access time attributes
□ Backup POSIX hard links as files

Logging: Log All

Backup files of size: All sizes

User defined variables

Advanced

Variable name: ONLINE

Value: START

Variable name	Value
ONLINE	START

18.144
Paula

Device Backup Options

You can select the options listed below for each backup device in use. None of the settings are needed, because CRC Check, Concurrency, and Media Pool use the default values that are set when the device is configured. The Prealloc List value is specified along with the media pool settings.

CRC Check

Set this option to have Data Protector calculate the CRC (Cyclic Redundancy Check) when a backup runs. CRC is an enhanced checksum function that lets you later confirm using the Verify option whether or not data has been written correctly to the medium.

The default value is OFF.

Concurrency

Concurrency allows more than one Disk Agent to write to one backup device. Data Protector can then keep the devices streaming if data can be accepted faster than a Disk Agent can send it. The maximum concurrency value is 32.

Data Protector provides default values for all supported devices.

Media Pool

This option selects the media pool with the media you will use for a backup. If not defined, a default pool, which is a part of device specification, is used.

Prealloc List

The **Prealloc List** is a subset of media in the media pool used for a backup. It specifies the order in which the media will be used. When using the **Prealloc List** and the Strict media allocation policy with the backup device, Data Protector expects the sequence of the media in the device to correspond with that specified in the **Prealloc List**. If the media are not available in this sequence, Data Protector issues a mount request. If no media are specified in this list, then the Data Protector allocation procedure is used to allocate media.

18.243
Paula

Pre- and Post-Exec Commands

Before a backup or restore session begins, an additional action is sometimes necessary. For example, you may want to check the number of files to back up, stop some transaction processing, or shut down a database. Such actions are performed using `pre-` and `post-exec` commands. `Pre-` and `post-exec` commands are not supplied by Data Protector. Depending on your needs, you have to write your own executables to perform the required actions.

For backup, `pre-` and `post-exec` commands can be configured on two levels:

Backup Specification

The `pre-exec` command is executed before the backup session starts. The `post-exec` command is executed when the backup session stops. You specify these commands as backup options for the entire backup specification. By default, `pre-` and `post-exec` commands for the session are executed on the Cell Manager, but you can choose another system.

Specific Backup Object

The `pre-exec` command for a specific backup object starts before the object is backed up. The `post-exec` command for the backup object is executed after the object is backed up. You specify these commands as backup options that apply for all objects, or for individual objects. `Pre-` and `post-exec` commands for the object are executed on the system where the Disk Agent that backs up the object is running.

`Pre-` and `post-exec` commands are run in the following order:

1. The `pre-exec` command for the entire backup specification starts and completes.
2. For each object in the backup specification:
    a. The `pre-exec` starts and completes.
    b. The object is backed up.

18.142
Paula

c. The post-exec (for each object in the backup specification) starts and completes.

3. The post-exec command for the entire backup specification starts and completes.

Figure 5-36 **Pre- and Post-Exec Control Flow**

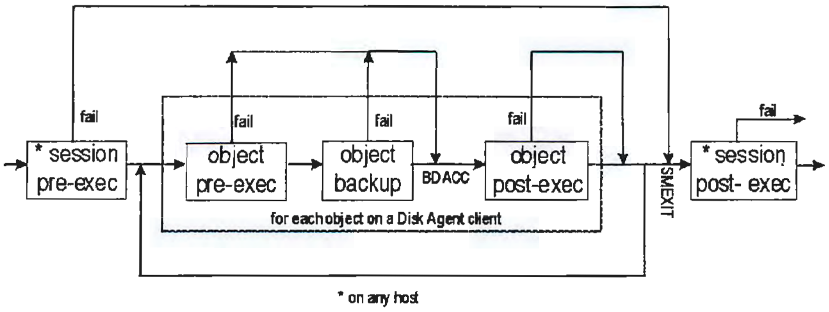

Pre- and Post- Exec Commands on Windows Systems

This section describes how to implement pre- and post-exec commands on Windows Cell Managers and clients.

How to Write the Commands

Pre- and post-exec commands can be written as executables or batch files. All the commands that run within the batch file must return an exit code 0 to signify success or greater than 0 to signify a failure.

Carefully follow the implementation guidelines provided in this section.

Pre- and Post-Exec Commands for a Backup Specification

Pre- and post-exec commands for a backup session are started before and after the session. These commands are usually executed on the Cell Manager, but you can choose another system.

Where to Locate the Commands

Pre- and post-exec scripts executed on the Cell Manager are started under the Data Protector Inet Service account (by default, Local System account). They can be located in any directory.

For scripts located in the *<Data_Protector_home>*\bin directory and executed on the Cell Manager, you can specify only the filename. For scripts located in the *<Data_Protector_home>*\bin directory and executed on the system other than Cell Manager, you must specify only the filename. For scripts that are not located in the `<Data_Protector_home>\bin` directory, you must specify a full path name of the script.

How to Specify the Filename or Pathname

In the backup specification, click the `Options` tab. Under `Backup Specification Options`, click `Advanced`. Write the filename or pathname in the `Pre-exec` and/or `Post-exec` text box.

When entering a full pathname, if your directory names are longer than 8 characters, write the pathname either in quotes or in the short 8.3 MS-DOS compatible form.

IMPORTANT

If you use quotes ("") to specify a pathname, do not use the combination of backslash and quotes (`\"`). If you need to use a trailing backslash at the end of the pathname, use the double backslash (`\\`).

Environment Variables

The following environment variables are set by Data Protector, and can be used only in `pre-` and `post-exec` scripts for a backup specification on the Cell Manager:

DATALIST The name of a backup specification.

MODE Backup operation type, such as full, incremental, incremental1, incremental2.

OWNER Owner of the session.

The contents of this variable are in the same format as in the database (case-sensitive):
`<user>.<group>@<hostname>` for UNIX
`<DOMAIN>\<user>@<hostname>` for Windows

PREVIEW The value is 1 if a preview is running and 0 if a backup is running. Use this variable to modify your commands so that they are executed only during a backup and not during a preview. By default, the `pre-` or `post-exec` commands are not executed for preview. You can enable them by setting the global option `ExecScriptOnPreview`.

15840
Paula

RESTARTED Set to 1 if this is a restarted backup session, otherwise set to 0. The `post-exec` can use this variable to prevent an additional restart in the case that **SMEXIT** equals 0.

SESSIONID Is used to identify a finished session and is recorded in the database. You cannot use this to preview a session (use SESSIONKEY).

SESSIONKEY Is used to identify a running session. You may, for example, abort a backup session before it is started if something is wrong.

SMEXIT The exit code of the Session Manager is the same as the exit code of the `omnib` command. You can only use this variable with the `post-exec` command. Agents can refer to Disk Agents, Media Agents, Application Agents, Symmetrix Agents, and so on.

Table 5-3 **SMEXIT VALUES**

Value	**Description**
0	All files were successfully backed up.
10	All Agents completed successfully, but not all files were saved.
11	One or more Agents failed or there was a database error.
12	None of the Agents completed the operation.
13	Session was aborted.

Key Points

- ✓ The `pre-` and `post-exec` commands for a backup specification have to be executables or batch files. It is important to specify a filename extension on Windows.
- ✓ The `pre-` and `post-exec` commands can be located in any directory on the system where the Disk Agent is running. If they are located in a directory other than `<Data_Protector_home>\bin` a full pathname must be specified.

15 139
Paula

- ✓ The execution of pre- and post-exec commands is implemented using the Windows pipe mechanism. All processes started in the pre- or post-exec functions must finish before processing continues.
- ✓ A pre- or post-exec command must return a non-negative value upon successful completion.
- ✓ If a pre-exec command fails (returns a value less than 0), the status of the backup session is set to Failed and the session is aborted.
- ✓ If a post-exec command fails (returns a value less than 0), the backup session status is set to Completed with errors.
- ✓ The pre- and post-exec commands for a backup specification are by default NOT executed during a preview of the backup. This behavior is defined by the ExecScriptOnPreview variable in the global options file. See "Global Options File" on page 523 for details on how to modify these values.
- ✓ Pre- and post-exec commands are handled in the same way as commands entered at the DOS prompt. Therefore, special characters, such as the pipe (|) and the redirect symbols (>, <) are not allowed.
- ✓ While pre- or post-exec commands are running, the backup session cannot be aborted.
- ✓ The pre- and post-exec commands run in the background mode. Therefore, do not use any commands that require user interaction.
- ✓ Standard output of the pre- and post-exec commands is written to the IDB as messages and shown on the monitor screen of the Data Protector GUI.
- ✓ You can disable a session's pre- and post-exec command execution on the Cell Manager by setting SmDisableScript global option to 1.
- ✓ You can disable remote session pre- and post-exec command execution on any client by adding OB2REXECOFF=1 into the omnirc file on the specific client.
- ✓ You can secure the client by specifying which Cell Managers are allowed to access the client. Only permitted Cell Managers will be able to execute pre- and post-exec commands on the client. For more information on securing a client, refer to the *HP OpenView Storage Data Protector Installation and Licensing Guide*.

13.138
Paula

Pre- and Post-Exec Commands for a Specific Backup Object

`Pre-` and `post-exec` commands for an object are executed before and after the backup of the object, respectively. You can specify these commands for all objects in a backup specification, or for each individual object. When backing up integrations, for example Oracle, the database is considered as an object, so the commands are executed before and after the database backup. These commands are executed on the system where the Disk Agent is running.

Where to Locate the Commands

`Pre-` and `post-exec` scripts for a backup object are started under the Data Protector Inet Service account (by default, Local System account) and can be located in any directory, except for host backup object. For host backups they must reside in the *<Data_Protector_home>*`\bin` directory on the system where the Disk Agent is running. For scripts located in the *<Data_Protector_home>*`\bin` directory you specify only the filename, otherwise a full path name of the script must be specified.

How to Specify the Filename or Pathname

To apply `pre-` and `post-exec` commands to all objects in the backup specification, click the `Options` tab in the backup specification. Under `Filesystem Options` (`Disk Image Options` in a saved backup specification for disk image backup), click `Advanced`.

To apply `pre-` and `post-exec` commands to individual objects only, click the `Backup Object Summary` tab in the backup specification. Right-click an object and click `Properties`. In the `Object Properties` dialog box, click the `Options` tab.

To apply `pre-` and `post-exec` commands to an integration object, click the `Options` tab in the backup specification. Under `Application Specific Options`, click `Advanced`.

Write the filename or pathname in the `Pre-exec` and/or `Post-exec` text box.

When entering a full pathname, if your directory names are longer than 8 characters, write the pathname either in quotes or in the short 8.3 MS-DOS compatible form.

Environment Variables

BDACC

The Disk Agent sets its exit code (0 is successful) to the **BDACC** environment variable. This variable can be checked in the `post- exec` command, thus making the `post-exec` command dependent upon successful termination of the Disk Agent.

18.134
Paula

NOTE

If you perform a host backup, the pre-exec script is started once, before the first filesystem backup for the particular system, while the post-exec script is started after the backup. In this case, **BDACC** cannot be exported because the variable is related to a single filesystem object, not to a whole client.

Key Points

- ✓ The pre- and post-exec commands for a backup object have to be executable or batch files. It is important to specify the filename extension on Windows.
- ✓ The pre- and post-exec commands can be located in any directory on the system where the Disk Agent is running except for host backups. If they are located in a directory other than *<Data_Protector_home>*\bin a full pathname must be specified.
- ✓ If a pre-exec command fails (returns a non-zero value), the backup of this object is aborted. The status of the object is set to aborted and the backup Disk Agent stops processing. No backup of the object exists.
- ✓ If a post-exec command fails (returns a non-zero value), the backup object status is set to aborted. The backup of the object exists and data can be restored.
- ✓ The pre- and post-exec commands are handled in the same way as commands entered at the DOS prompt. Therefore, special batch characters such as the pipe (|) and the redirect symbols (>, <) are not allowed.
- ✓ While pre- or post-exec commands are running, the backup session cannot be aborted.
- ✓ The pre- and post-exec processes run in the background mode. Therefore, do not use any commands that require user interaction.
- ✓ Standard output of the pre- and post-exec commands is written to the IDB as messages and shown on the monitor screen of the Data Protector GUI.
- ✓ The pre- and post-exec scripts have to send some output at least every 15 minutes by default, or the sessions waiting for the scripts are aborted. You can change this time interval by modifying the ScriptOutputTimeout variable in the global options file.

- ✓ Time-out is provided. If no message is received within the specified time-out in seconds, the session is aborted.
- ✓ You can disable a `pre-` and `post-exec` script by adding the line `OB2OEXECOFF=1` in the `omnirc` file on any client.
- ✓ You can secure the client by specifying which Cell Managers are allowed to access the client. Only permitted Cell Managers will be able to execute `pre-` and `post-exec` commands on the client. For more information on securing a client, refer to the *HP OpenView Storage Data Protector Installation and Licensing Guide*.

Pre- and Post- Exec Commands on UNIX Systems

This section describes how to implement `pre-` and `post-exec` commands on UNIX Cell Managers and clients.

How to Write the Commands

`Pre-` and `post-exec` commands can be written as shell scripts.

See Appendix, "Examples of Pre-Exec and Post-Exec Commands for UNIX," on page A-20.

Pre- and Post-Exec Commands for a Backup Specification

`Pre-` and `post-exec` commands for a backup session are started before and after the backup session, respectively. These commands are usually executed on the Cell Manager, but you can choose another system as well.

Where to Locate the Commands

`Pre-` and `post-exec` commands for backup specifications on UNIX systems are started by the backup session owner, unless the backup session owner has the `Back up as root` permission and the commands are then started under root.

On a UNIX Cell Manager, the exec commands for backup specifications can reside in any directory, but the full pathname must be specified when defining the backup specification.

On a remote UNIX client the exec commands for backup specifications must be located as follows:

- Solaris 7/8, HP-UX: `/opt/omni/lbin`
- Solaris 2.6, other UNIX systems: `/usr/omni/bin`

In this case, the command filename can be specified without the path.

18.135
Paula

How to Specify the Filename or Pathname?

For information on how to specify the commands, refer to the online Help index keyword "pre- and post-exec commands for backup specifications".

Environment Variables

The following environment variables are exported, and can be used in `pre-` and `post-exec` scripts for a backup specification session on any host:

DATALIST The name of the backup specification.

MODE Backup operation type, such as full, incremental, incremental1, and so on.

OWNER Owner of the session.

The content of this variable is in the same format as in the database (case-sensitive):
<user>.<group>@<hostname> for UNIX and
<DOMAIN>\<user>@<hostname> for Windows NT.

PREVIEW Set to 1, if the preview is running. Set to 0, if a backup is running. Use this variable to modify your commands so that they are executed only during a backup and not during a preview. By default, `pre-` and `post-exec` commands are not executed for preview. You can enable this with global option `ExecScriptOnPreview`.

RESTARTED Set to 1 if this is a restarted Backup session, otherwise set to 0. The `post-exec` can use this variable to prevent an additional restart if **SMEXIT** equals 0.

SESSIONID Is used to identify a finished session and is recorded in the database. You cannot use this to preview a session (use SESSIONKEY).

SESSIONKEY Is used to identify a running session. You may, for example, abort a backup session before it is started if something is wrong.

SMEXIT The exit code of the Session Manager is the same as the exit code of the `omnib` command. You can only use this variable with the `post-exec` command. Agents can refer to Disk Agents, Media Agents, Application Agents, and Symmetrix Agents. Refer to Table 5-3 on page 253 for details on SMEXIT values.

18134
Paula

Key Points

Check the following before configuring pre- and post- exec commands for a backup specification on a local or remote host:

✓ If a pre-exec command fails (returns a non-zero value), the backup status of the session is set to failed and the session is aborted.

✓ If a post-exec command fails (returns a non-zero value), the backup of the session is set to completed with errors.

✓ The pre- and post-exec commands for a backup specification are by default NOT executed during a preview of a backup. This behavior is defined by the ExecScriptOnPreview variable in the global options file. See "Global Options File" on page 523 for details.

✓ While the pre- or post-exec commands are running, the backup session cannot be aborted.

✓ The pre- and post-exec processes operate in the background mode. Therefore, do not use any interactive commands for pre- and post-exec processing.

✓ The pre- and post-exec scripts have to send some output at least every 15 minutes by default, or the sessions waiting for the scripts are aborted. You can change this time interval by modifying the ScriptOutputTimeout variable in the global options file.

✓ Time-out is provided. If no message is received within the specified time-out in seconds, the session is aborted.

✓ If there is no executable script on the host or if the path of the script is wrong, Data Protector displays an error message that the script failed and the session is aborted.

✓ If a command writes any text to stdout, this text is sent to the Session Manager and written to the database. A stderr is redirected to /dev/null. You can redirect it to stdout to get error messages logged to the database.

NOTE

A pre- or post-exec script may hang because it did not close all file descriptors before forking a new process. If the new process runs in the background and does not exit, such as, for example, the database server process (dbstart), the scripts hang. You can use the detach command. The source of the detach command is provided in the detach.c file, but is officially unsupported. For example:

18.133
Paula

```
/opt/omni/lbin/utilns/detach pre_script [arguments...]
```

- You can disable a session's `pre-` and `post-exec command` execution on the Cell Manager by setting the `SmDisableScript` global option to 1.
- You can disable remote session `pre-` and `post-exec` command execution on any client by adding `OB2REXECOFF=1` into the `omnirc` file on the specific client.
- You can secure the client by specifying which Cell Managers are allowed to access the client. Only permitted Cell Managers will be able to execute `pre-` and `post-exec` commands on the client. For more information on securing a client, refer to the *HP OpenView Storage Data Protector Installation and Licensing Guide*.

Pre- and Post-Exec Commands for a Specific Backup Object

`Pre-` and `post-exec` commands for an object are executed before and after the backup of the object, respectively. You can specify these commands for all objects in a backup specification, or for each individual object. When backing up integrations, for example Oracle, the database is considered as an object, so the commands are executed before and after the database backup. These commands are executed on the system where the Disk Agent is running.

Where to Locate the Commands

`Pre-` and `post-exec` commands for backup objects on UNIX systems are started by the backup session owner, unless the backup session owner has the `Back up as root` permission and the commands are then started under root.

On UNIX Cell Managers and clients, the `exec` commands for backup objects can reside in any directory, and the full pathname must be specified when defining the backup specification. However, for host backups the `exec` commands must reside in the `/opt/omni/lbin` directory on HP-UX or in `/usr/omni/bin` on other UNIX systems. In this case, only the filename can be specified. The commands have to be located on the system where the Disk Agent is running.

How to Specify the Filename or Pathname

For information on how to specify the commands, refer to the online Help index keyword "pre- and post-exec commands for backup objects".

15132
Paula

Environment Variables

The following environment variables are exported, and can be used in the pre- and post-exec scripts for an object on the system where the Disk Agent is running:

BDACC

The Disk Agent sets its exit code (0 is successful) to the BDACC environment variable. This variable can be checked in the post-exec script, thus making the post-exec command dependent on the successful termination of the Disk Agent.

NOTE

If you perform a host backup, the pre-exec script is started once, before the first filesystem backup for the particular system, while the post-exec script is started after the backup. In this case, BDCACC cannot be exported because the variable is related to a single filesystem object, not to a whole client.

Key Points

Check the following key points before configuring the pre- and post-exec commands:

- ✓ The pre- and post-exec commands for an object are executed during the preview of a backup. Therefore, you may want to preview your backup first and then add the pre- and post-exec commands, or check the PREVIEW environment variable in your scripts.
- ✓ If a pre-exec command for an object fails (returns a non-zero value), the backup status of the object is set to `Aborted` and the Disk Agent stops processing. No backup of the object exists.
- ✓ If a post-exec command fails (returns a non-zero value), the backup status of the object is set to `Aborted`. A backup of the object exists and data can be restored.
- ✓ The pre- and post-exec commands should send some output to the Disk Agent at least every 120 minutes by default, or the backup of the object is aborted. This time period can be changed by modifying the `SmDaIdleTimeout` variable in the global options file.
- ✓ Pre- and post-exec commands are handled in the same way as commands entered at the shell prompt. Special shell characters, such as the pipe (|) and the redirect symbols (>,<) are not allowed.
- ✓ While the pre- and post-exec commands are running, the backup session cannot be aborted.

18.131
Paula

- ✓ The pre- and post-exec processes operate in background mode. Therefore, do not use any interactive commands for the pre- and post-exec processing.
- ✓ If a command writes any text to stdout, this text is received by the Disk Agent, sent to the Session Manager, and written to the database. A stderr is redirected to /dev/null. You can redirect it to stdout to get error messages logged to the database.
- ✓ The pre- and post-exec commands for an object have to be located on the client where the Disk Agent is running.
- ✓ The pre- and post-exec commands must be executable and specified with the full pathname.
- ✓ You can disable pre- and post-exec scripts by adding the line OB2OEXECOFF=1 into the omnirc file on any client.
- ✓ You can secure the client by specifying which Cell Managers are allowed to access the client. Only permitted Cell Managers will be able to execute pre- and post-exec commands on the client. For more information on securing a client, refer to the *HP OpenView Storage Data Protector Installation and Licensing Guide*.

18.130
Paula

Managing Failed Backups

During a backup, some systems may not be available because they were shut down, there were some networking problems, or similar occurrences. This results in some systems not being backed up entirely.

Setup Notification

Data Protector lets you configure a notification so that you are informed about unexpected events, such as a mount request or a device error during a backup session. You can choose among the methods that most suit your needs, for example e-mail, or a broadcast message to your Windows display.

See Chapter 7, "Monitoring, Reporting, Notifications, and the Event Log," on page 307 for details.

Checking Failed Backups

One of the most important aspects of managing backups is the regular checking of the backup status. Data Protector provides a comprehensive reporting functionality that allows you to view reports on the backup status. See "Monitoring Sessions" on page 309 for details on the reporting functionality.

Warnings When Backing Up System Disks

Data Protector issues warnings when backing up the system disk on Windows systems. This is because certain files on the system disk are always busy and cannot be opened by any application, including the Disk Agent. The contents of these files can only be backed up as a part of `CONFIGURATION`.

When these files are accessed by a filesystem backup, such as when the whole system disk is backed up, Data Protector fails to open them and reports warnings or errors, depending on the backup options. See "Using Backup Options" on page 225.

While this behavior is correct from the filesystem backup point of view, it can create a manageability problem. Due to the large number of warnings that are always reported, it is likely that a failure of another file may be overlooked.

These specific files can only be backed up through a `CONFIGURATION` backup. Knowing this, you can exclude them from a filesystem backup to avoid warnings.

18.129
Paula

The following example is a list of files that cannot be opened on an active Windows NT 4.0 system with the Windows NT software installed on the C: drive:

```
<%SystemRoot%>\system32\config\default
<%SystemRoot%>\system32\config\default.LOG
<%SystemRoot%>\system32\config\SAM
<%SystemRoot%>\system32\config\SAM.LOG
<%SystemRoot%>\system32\config\SECURITY
<%SystemRoot%>\system32\config\SECURITY.LOG
<%SystemRoot%>\system32\config\software
<%SystemRoot%>\system32\config\software.LOG
<%SystemRoot%>\system32\config\system
<%SystemRoot%>\system32\config\SYSTEM.ALT
```

For each user who is logged on, the following files also cannot be opened:

```
<%SystemRoot%>\Profiles\<user>\NTUSER.DAT
<%SystemRoot%>\Profiles\<user>\ntuser.dat.LOG
```

IMPORTANT When performing a filesystem backup of a system disk, the previously listed files are not backed up. Excluding them only solves the problem of managing the session reports. You should perform a CONFIGURATION backup to back up the contents of these specific files.

When backing up an inactive system disk (for example in a dual-boot situation) the previously listed files are not a part of the currently active CONFIGURATION. These files can be backed up in a filesystem backup, and should not be excluded.

Preventing Backup Failure

Data Protector provides a set of features that improve backup robustness, thus lessening the chance that a backup could fail.

18.128
Lauh

If a backup of an object fails to start, Data Protector tries to back up this object again at the end of the backup session. If it fails again, the object is not backed up, and the status of the object and the session is set to `Failed`. A backup is repeated when it is scheduled. If some objects finish properly, the session status is completed with failures.

Clients that are not up and running when they are scheduled to be backed up are retried after the rest of the objects are completed. Before the first failed object is retried, the backup session is suspended for 30 seconds. This waiting time can be changed using the `WaitBeforeRetry` global option. See "Global Options File" on page 523 for information on how to change global options.

IMPORTANT If you have an infrequent backup schedule, this may result in a period of time when there is no recent backup of your data.

NOTE Data Protector always needs one full backup of data. If no protected full backup is available, a full backup will be done next time, even though an incremental backup was scheduled. To avoid this, run a full backup of the failed system interactively before you schedule a backup.

For details on full and incremental backup behavior, see the *HP OpenView Storage Data Protector Concepts Guide*.

When you configure a backup, you can use the `Reconnect Broken Connection` option. When this option is set, Data Protector reconnects the Backup Session Manager and Disk or Media Agents in the case of short-term network problems during a backup session. This often happens on unreliable LAN networks.

Enabling Wake ONLAN Support

If you have any machines that support remote power-up (**Wake ONLAN**), you can use the Data Protector Wake ONLAN support. When a Backup Session Manager fails to connect to a client that is configured to use Wake ONLAN support, it sends a wake-up request according to the Wake ONLAN protocol, and retries connecting to the client. This allows the full use of the power-saving features of desktop systems, which would otherwise interfere with the backup process.

18127
Paula

NOTE You can enable Wake ONLAN support for computers equipped with a Wake ONLAN-compatible LAN interface, such as the HP NightDIRECTOR series. The Wake ONLAN (WOL) option is available in the BIOS setup.

When you install a Disk Agent on a Windows client and add it to a cell, the client's Mac address is automatically discovered. You can also manually change the Mac address in the same section where you enable the Wake ONLAN (WOL) option, as shown below.

Use the following steps to enable Wake ONLAN support for the Windows client:

1. In the `Data Protector Manager`, switch to the `Clients` context.
2. In the Scoping Pane, right-click the client whose WOL option you want to enable, and then click `Properties`.
3. Click the `Advanced` tab.
4. Under the `Magic Packet` section, select the `Enable Magic Packet` check box, and then click `Apply`.

Restarting Failed Backups

Data Protector provides a simple way of restarting the backup of failed objects only. This can be done as follows:

1. In the `Data Protector Manager`, switch to the `Internal Database` context.
2. Under `Internal Database`, expand the `Sessions` item.
3. In the `Results Area`, search for your backup.

   You can sort your sessions using the buttons on the top of each of the columns.
4. Right-click the failed session, and then select `Restart Session`.

   A dialog box appears asking you to confirm that you want to restart the session. Click `Yes`.

15.126
Paula

ANEXO UNIDADE DE BACKUP ROBOTIZADO
PARTE D/4

18125
Paula

6 Restore

18.124
Paula

In This Chapter

This chapter describes restore topics, such as how to restore specific data and how to use restore options to achieve a desired restore behavior.

"Restoring Your Data" on page 269

"Restoring UNIX Systems" on page 276

"Restoring Windows Systems" on page 277

"Restoring Novell Netware Filesystems" on page 287

"Restoring OpenVMS Filesystems" on page 291

"Restore Options" on page 294

"Restore Techniques" on page 299

For information on how to restore database applications such as Oracle, SAP R/3, MS Exchange, MS SQL, Informix, IBM DB2 UDB or Sybase, refer to the *HP OpenView Storage Data Protector Integration Guide*.

For information on how to restore the IDB, refer to Chapter 9, "Managing the Data Protector Internal Database," on page 381 and "Recovering the IDB" on page 417.

18123
Paula

Restoring Your Data

A restore is a process that recreates the original data from a backup copy on a disk. This process consists of the preparation and actual restore of the data, and optionally some post-restore actions that make the data ready for use.

Data Protector includes an internal database (IDB) that keeps track of data, including what files from which system are kept on a particular medium. The IDB provides fast and convenient access to the data to be restored.

Data Protector offers you some special restore features:

- The ability to restore on different levels: session, client, object, directory, specific file, or specific file version
- The option to specify an alternative location to restore your data
- Cross-platform restore
- Parallel restore of multiple objects from a session, on a client, or in a cell

Depending on the platform, the way you specify these features and available options can vary.

Standard Restore Procedure

Prerequisite

In order to perform a restore, you need to have the appropriate user rights. These rights are defined according to the user group.

What You Need to Do to Perform a Restore

As part of the standard restore procedure, you need to do the following:

- Select the data to be restored
- Find the media needed
- Start the restore session

Other Settings

Other settings are already predefined according to the backup process, but can be modified. If you want to change these predefined settings, you can specify the following:

- The backup version you want to restore

18.122
Paula

- The location you want to restore data to
- The device to restore from
- How to handle file conflicts with existing files
- Restore options, such as locking files during restore

For detailed steps of standard restore tasks, refer to the online Help index keyword “standard restore procedure”.

Selecting Your Data for Restore

The Data Protector `Restore` context offers two possible ways of browsing objects for restore:

- `Restore Objects` with a list of backed up objects classified by client systems in the cell and by different data types, such as Filesystem, Disk Image, Internal Database, and so on.
- `Restore Sessions` with a list of filesystem sessions with all objects backed up in these sessions. You can choose to view only sessions from the last year, last month, or last week. By default, all filesystem sessions are listed. You cannot perform restore of the online database integrations from a specific backup session.

You can select either one object to perform a single restore, or multiple objects to perform a parallel restore. For more information on parallel restore, refer to “Restoring Files in Parallel” on page 300.

You can also specify a `Search Interval` and browse only objects backed up within a specific timeframe.

Data Protector offers the `Restore by Query` task, which searches for files and directories you want to restore and restores them. Refer to “Restoring by Query” on page 302.

18121
Bauke

Figure 6-1 **Selecting Data for Restore**

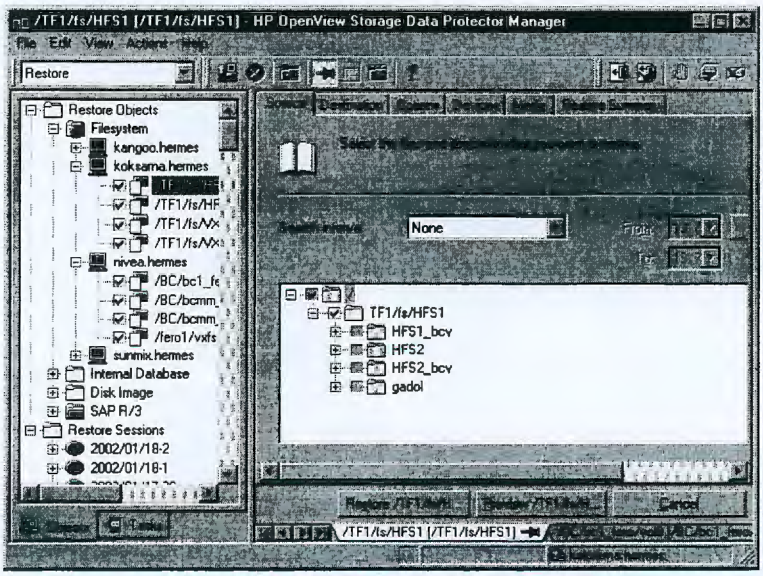

Selecting a Backup Version

When selecting data that you want to restore, the last backup version is selected by default. This means that only directories and/or files from the last backup session are selected for restore. Directories and files in the same tree structure that have not been backed up in the same backup session are shaded.

If you want to restore the data from any other backup session, browse for the file or directory that you want to restore, right-click it, and click `Restore Version`.

In the `Version` tab, click "..." to get additional information about the backup versions. The "..." button is available only if the backup was performed using a logging level that logs attributes.

Handling File Conflicts

In the `Description` property page of your restore, you can specify how to handle conflicts between the version currently on the disk and the backup version of a file. `File Conflict Handling` offers you three

18.120
fauh

possible options: `Keep most recent`, `No overwrite`, and `Overwrite`. For more information on these options, refer to “Restore Options” on page 294.

Specifying Restore Location

By default, Data Protector restores data to the same client and directory from which it was backed up. You can change these default settings in the `Destination` property page by specifying where to restore your data to:

- With the appropriate user rights, you can restore to another client.
- You can restore to another directory.

This specification can be set on a per-object basis.

Additionally, Data Protector offers the `Restore As/Into` option for specifying a different location for individual files and directories from the same backup object. This specification can be set on a per-object basis or for the individual files.

For more information on specifying restore location, refer to “Restoring Files to Different Paths” on page 299.

Setting Restore Options

Set restore options in the `Options` property page of your restore. These are available according to the type of data being restored. For example, not all restore options available for a filesystem restore are available for a disk image restore. For more information on restore options, refer to “Restore Options” on page 294.

Restoring Under Another Device

By default, the device used for restore is the same device as the one the backup was made to. You can restore your data from any device configured in the same Data Protector cell. To specify a new device, click the `Change` button in the `Devices` property page of your restore. The new device will be used for this session only.

NOTE

With *some* database integrations, you can set the changed device as a default restore device for *all* Data Protector integration restore sessions (regardless of the type of integration), by clicking the `Save as default` button.

Finding Needed Media

To get a list of the media on which your data is stored, go to the `Media` property page after you select data for restore.

You can also find the media needed for the restore by clicking the `Needed Media` button in the `Start Session/Preview Session` dialog box. This dialog box appears when you start or preview the restore.

Previewing and Starting a Restore

Ensure that the media are loaded properly before starting the restore. Otherwise, the media will not be detected.

If restoring objects selected in the `Source` property page of your restore, use the `Start` or `Preview` buttons.

If also restoring objects selected in the Scoping Pane, click `Start Restore` or `Preview Restore` from the `Actions` menu.

Aborting a Restore

Aborting a restore session stops the restore. Data processed before the session was aborted is restored to the specified location.

To abort a restore session, click `Abort` in the `Actions` menu.

You can also abort restore sessions in the Data Protector `Monitor` context.

Restoring Disk Images

A disk image restore is a sector-by-sector restore of a disk image backup. Data Protector restores a complete image of the disk that was backed up (as a disk image) at a certain point in time. This method is particularly fast. It is available for Windows and UNIX systems.

18.118
Paula

Prerequisites

You need to meet the following prerequisites in order to perform a disk image restore:

- The disk must have been backed up using the disk image backup.
- To restore a disk image on a disk other than the disk from which you backed it up, the new disk must be of the same size or larger.
- On UNIX systems, unmount the disk before a disk image restore and then mount it back afterwards.

Limitation

On Windows systems, disk image restore fails if a file or section is in use.

Procedure

To restore a disk image backup, expand the disk image object under the `Restore` context as shown in Figure 6-2 on page 274, and then use the standard restore procedure. Refer to "Standard Restore Procedure" on page 269.

Figure 6-2 **Disk Image Objects**

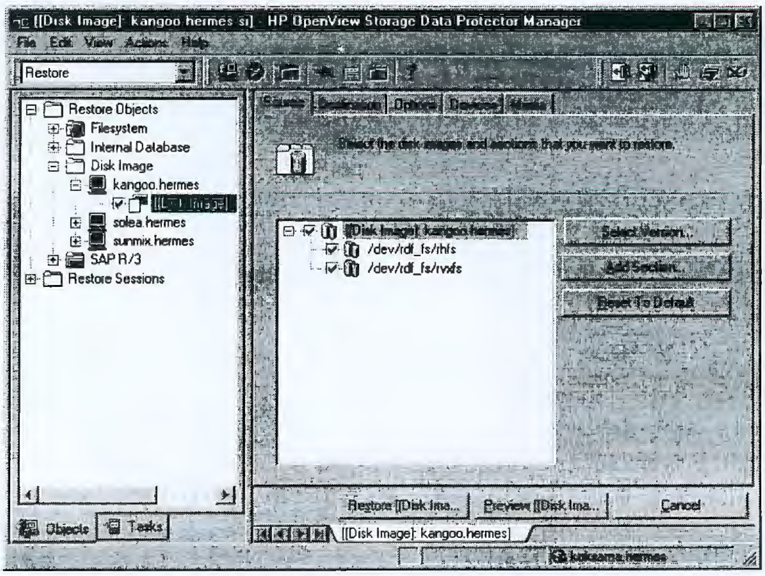

Select the disk images and sections that you want to restore.

18.117
Paulo

Restoring Your Data to a Shared Disk

Data Protector allows you to restore UNIX and Windows data to a Windows shared disk, even if the data was not originally backed up from the shared disk. The Data Protector user account and its Inet service must have permission to access the remote computer and permission on the Disk Agent client. Refer to "Setting the User Account for the Data Protector Inet Service" on page 187 for more information on how to use the appropriate logon account.

Here are some cases in which one would restore a UNIX or Windows filesystem to a Windows shared disk:

- If the system is not part of the Data Protector cell and does not have the Data Protector Disk Agent installed.
- If you are restoring to platforms not directly supported by Data Protector, such as Windows for Workgroups or Windows 3.1 systems.
- If you want to make the data available from several systems.

NOTE When you restore your data to a different filesystem type than it was backed up from (UNIX to Windows), filesystem-specific attributes may be lost.

How to Restore to a Shared Disk

In the `Destination` property page of your restore, you can specify the target client and a Windows shared disk as a new location for the data you want to restore. For detailed steps, refer to the online Help index keyword "shared disks, restoring to".

18.116
Paula

Restoring UNIX Systems

What Is Restored?

When restoring files to the original location from which the backup was performed, Data Protector restores the files, including file attributes.

System-specific data, such as ACL (Access Control List) on UNIX, is restored only on the same filesystem type and operating system from which the backup was made.

Restoring Regular UNIX Files

Use the standard restore procedure to restore UNIX files and directories. Refer to “Standard Restore Procedure” on page 269.

Restoring VxFS

When restoring VxFS data backed up to a temporary directory, use the `Restore As` option and restore it to the desired location. Refer to “Restoring Files to Different Paths” on page 299 for information on how to use the `Restore As` option.

Restoring OmniStorage Backups

Beside restoring backed up data into an OmniStorage controlled file system (MFS), Data Protector A.05.10 offers the possibility to restore OmniStorage filesystem data backed up with OmniBack II or Data Protector, using a normal filesystem restore on HP-UX 11.x. In this case, the “migration attributes” of OmniStorage, like migration policies, will be lost.

OmniStorage files can be restored to any filesystem on HP-UX, but in order to retain the VxFS specific file attributes it is recommended that the target filesystem is of JFS type with a VxFS3 or later layout.

Restoring Disk Images

Refer to “Restoring Disk Images” on page 273.

Restoring to a Shared Disk

Refer to “Restoring Your Data to a Shared Disk” on page 275.

18.115
Paula

Restoring Windows Systems

What Is Restored? When restoring a Windows filesystem, Data Protector restores the data within the files and directories, as well as Windows-specific information about the files and directories.

Consider the filesystem restore limitations when restoring to a different filesystem from the one where the backup was performed. See "Filesystem Limitations" on page 278.

The following Windows-specific information is restored:

- Full Unicode filenames
- FAT16, FAT32, VFAT, and NTFS attributes

  Once a file has been backed up, its archive attribute is cleared. You can change this behavior by setting the `Do not use archive attribute` option in the Advanced filesystem backup options in the backup specification.
- NTFS alternate data streams

  For example, Object IDs on Windows 2000 are backed up as sets of alternate data streams.
- NTFS security data

Additionally, the following applies on Windows systems, using NTFS 3.x:

- The NTFS filesystem supports reparse points.

  The volume mount points, Single Instance Storage (SIS), and directory junctions are based on the reparse point concept. These reparse points are selected like any other filesystem object.
- The NTFS filesystem supports sparse files as an efficient way of reducing the amount of allocated disk space.

  These files are backed up as sparse to save tape space. Sparse files are backed up and restored as sparse to the NTFS 3.x filesystem only.
- Some of the NTFS filesystem specific features are controlled by the system services, which maintain their own data records. These data structures are backed up as a part of `CONFIGURATION`.
- Encrypted files.

18.114
Paula

Filesystem Limitations

You can select a different target filesystem from the one where the backup was performed. This functionality has limitations that should be taken into consideration. See Table 6-1.

Table 6-1 **Windows Filesystem Restore Limitations**

	TO						
FROM	FAT32	FAT16	CDFS	UDF	NTFS 1.1[a]	NTFS 3.0[b]	NTFS 3.1[c]
FAT32	FC	FC	N/A	N/A	FC	FC	FC
FAT16	FC	FC	N/A	N/A	FC	FC	FC
CDFS	FC	FC	N/A	N/A	FC	FC	FC
UDF	FC	FC	N/A	N/A	FC	FC	FC
NTFS 1.1[a]	*	*	N/A	N/A	FC	FC	FC
NTFS 3.0[b]	***	***	N/A	N/A	**	FC	FC
NTFS 3.1[c]	***	***	N/A	N/A	**	FC	FC

How to Read This Table

a Also called NTFS 4.0. It is used by Windows NT.

b Also called NTFS 5.0. It is used by Windows 2000.

c Also called NTFS 5.1. It is used by Windows XP/Server 2003.

FC Full Compatibility, meaning that the file attributes are entirely preserved.

* Files are restored without security information and alternate data streams.

** Reparse points, sparse files and encrypted files are not restored.

*** Combines * and **.

18.113
Paula

Table 6-1 shows that NTFS 3.x filesystem objects can only be adequately restored to the NTFS 3.x filesystem. The filesystem-specific attributes and alternate data streams are lost when restoring into a different or older filesystem version.

- A Windows 2000/XP/Server 2003 reparse point, such as a directory junction or a volume mountpoint, can only be restored to an NTFS 3.x filesystem. UNIX reparse points cannot be restored to an NTFS 3.x filesystem.

NOTE When you restore an NTFS 3.x filesystem that contains SIS reparse points, a full disk condition may occur. This happens if the original file is restored into multiple target files, which can take up more space than available.

- Sparse files are restored as sparse to the NTFS 3.x filesystem only.
- User Disk Quotas cannot be restored using Data Protector.
- If a user attempts to restore a sparse file to a non-NTFS 3.x filesystem, Data Protector will issue a warning. A sparse file restored to a filesystem other than NTFS 3.x will not include zero sections.
- Microsoft encrypted NTFS 3.x files can only be restored to the NTFS 3.x filesystem, because other filesystem drivers cannot decrypt them.

Restoring Regular Windows Files and Directories

Use the standard restore procedure to restore Windows files and directories. Refer to "Standard Restore Procedure" on page 269.

Restoring Shared Disks

Objects that were backed up as shared disks are associated with the Disk Agent client that was used to back them up. If the environment has not changed, you can restore the shared disk as you would a local Windows filesystem. By default, the same Disk Agent client that was used to back up the shared disk is used to restore the data to the original location.

For information on how to choose and configure the Disk Agent client that restores the shared disks, refer to "Backing Up Windows Shared Disks" on page 185.

For information on restoring a UNIX or Windows filesystem to a shared disk, refer to "Restoring Your Data to a Shared Disk" on page 275.

18.112
Paula

Restoring Disk Images

Refer to “Restoring Disk Images” on page 273.

Restoring the Windows CONFIGURATION

To restore the Windows CONFIGURATION, select the CONFIGURATION object and follow the standard restore procedure. See Figure 6-3.

Figure 6-3 **Restoring Windows CONFIGURATION**

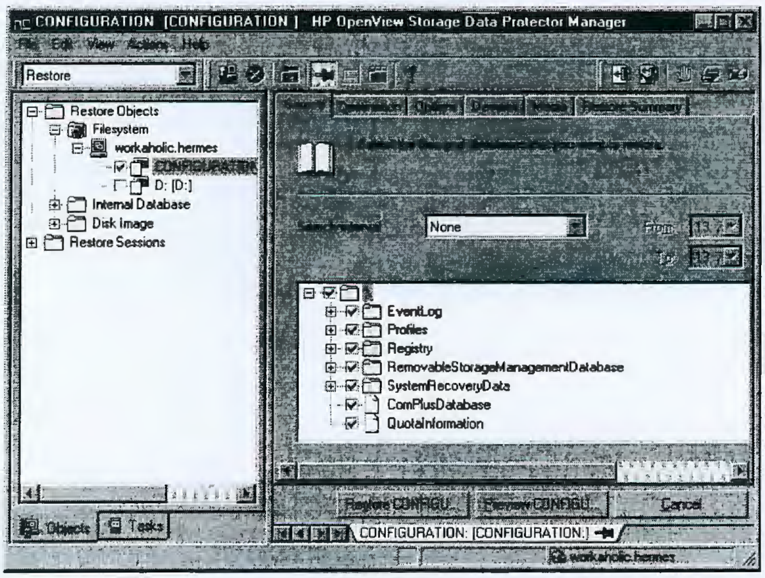

Prerequisites

The CONFIGURATION consists of data structures that influence system operation. Therefore, the system must be prepared for such a restore. The prerequisites depend on the contents of the CONFIGURATION item and the Windows operating system version. Refer to “Backing Up CONFIGURATION” on page 173. They can be summarized as follows:

- User profiles that are currently being used cannot be restored. The login account has to be changed or the relevant service has to be stopped.

  Refer to “Restoring Windows User Profiles and Event Logs” on page 285 for details.

18.111
Paula

- You have to boot the system in the Active Directory restore mode to restore the Active Directory.

  Refer to "Restoring Windows 2000/XP/Server 2003 Services" on page 283 for details.

When the whole `CONFIGURATION` is restored, restart the system to read the restored data in the `Registry`. Refer to "Restoring the Windows Registry" on page 282 for details.

Restoring the SysVol

You can perform a restore of SysVol directory in one of three modes:

- `Nonauthoritative` restore

  If at least one domain controller in the domain is available and working, files are restored to their original location. The restored data is not propagated to other domain controllers.

- `Authoritative` restore

  Perform an `authoritative` restore if critical SysVol data is deleted from the local domain controller and the deletion is propagated to other domain controllers.

- `Primary` restore

  If all domain controllers in the domain are lost and you want to rebuild a domain controller from backup, the FRS is informed that you are restoring primary files, and files are restored to their original location.

Restoring the Windows 2000/XP/Server 2003 System State

Prerequisites

If you use Active Directory, which is always a part of the System State, you have to start the system in the Active Directory restore mode.

Refer to "Restoring Windows 2000/XP/Server 2003 Services" on page 283 for details on Active Directory modes.

You restore the System State by selecting the following objects in the Restore wizard:

18.110
Paula

1. System State objects that belong to CONFIGURATION. Refer to “Backing Up the Windows 2000/XP/Server 2003 System State” on page 176 for a list of these objects.
2. The SystemVolumeInformation folder and the boot files. These are located on the system drive.

Figure 6-4 **Selecting System State Items**

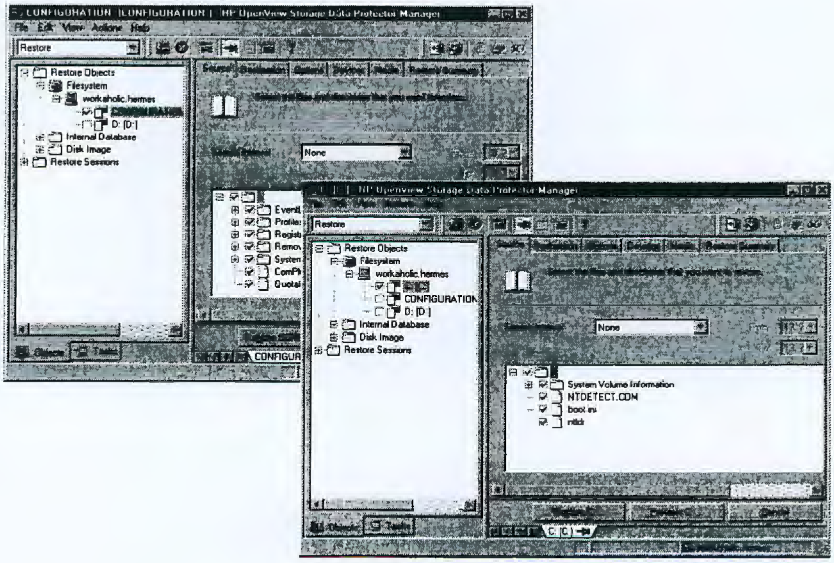

NOTE From the Data Protector point of view, the System State consists of ordinary filesystem objects and CONFIGURATION objects. As opposed to selecting objects in the Backup wizard, different objects for restore are selected in separate Restore wizards.

Once the restore session is completed, restart the system.

Restoring the Windows Registry

To restore the Windows Registry, expand the CONFIGURATION item and select only the Registry item.

13 109
Paula

Once the restore session is completed, restart the system.

NOTE

If you select the whole Windows 2000/XP/Server 2003 `Registry` for a restore, some of the Registry keys are not restored, and others are treated in a special way during a restore. This is because certain keys are being used by the operating system. You can find them under the following `Registry` key:

```
\HKEY_LOCAL_MACHINE\SYSTEM\CurentControlSet\Control\BackupRe
store\KeysNotToRestore
```

Restoring Windows 2000/XP/Server 2003 Services

To restore Windows 2000/XP/Server 2003 services, expand `CONFIGURATION` and select the service you want to restore.

Prerequisites

The following information that belongs to Windows 2000/XP/Server 2003 services can be selected under `CONFIGURATION`:

- `COMPlusDatabase`
- `FileReplicationService`
- `RemovableStorageManagementDatabase`
- `ActiveDirectoryService`
- `TerminalServiceDatabase`
- `CertificateServer`
- `DHCP`, `WINS`, and `DNSServerDatabase`

For a detailed explanation of these terms, refer to “Glossary”.

The list below describes specifics related to restoring a particular Windows 2000/XP/Server 2003 service.

Active Directory Restore

If you want to restore the Active Directory service, restart the system using the `Directory Services Restore Mode` start-up option.

When the system is started in the `Directory Services Restore Mode`, the domain user accounts cannot be used. Configure the `Data Protector Inet` and the `crs` service (for a Cell Manager) to log on using the local system account and then restart the services. Refer to "Setting the User Account for the Data Protector Inet Service" on page 187 for more details.

Select `Active Directory`, and set a replication mode by choosing among the `Windows 2000 specific options`: `Primary`, `Nonauthoritative`, `Authoritative`. For information on these options, refer to "Active Directory Specific Options" on page 297.

NOTE

To perform an `Authoritative` restore, you also need to run `ntdsutil.exe` after the restore session has finished. For example, to perform a typical `authoritative` restore, at a command prompt enter `ntdsutil`, then `authoritative restore`, then `restore database`. Restart the server and wait for replication to take place.

TIP

You can also create a post-exec command to perform the additional action needed for the Active Directory authoritative restore. For example, to perform an authoritative restore of an entire directory, use the following line:

```
ntdsutil "popups off" "authoritative restore" "restore
database" quit quit
```

Figure 6-5 **Active Directory Restore Modes**

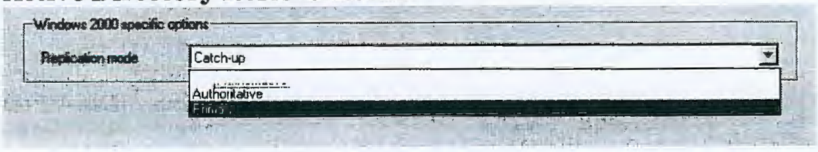

Windows 2000 specific options

Replication mode: Catch-up

Certificate Services Restore

Certificate Server Services are restored offline. You have stop them before you can start a restore. `Authoritative` is the only possible replication mode.

Once the restore has finished, restart the system.

18.107
Paula

Remote Storage Service Restore

Although the RSS databases are part of System State data, you restore them manually. The RSS database must be restored offline. You can provide pre- and post-exec scripts to stop and restart the service, or you can stop and restart it manually before and after the restore, respectively.

Select the following directories for restore:

- `<%SystemRoot%>\System32\RemoteStorage`
- `<%SystemRoot%>\System32\NtmsData`

Restoring DFS

Data Protector restores the configuration of the Windows 2000/XP/Server 2003 Distributed File System (DFS) as part of one of the following:

- Windows 2000/XP/Server 2003 Registry, if the DFS is configured in a standalone mode.
- Windows 2000/XP/Server 2003 Active Directory, if the DFS is configured in a domain mode.

Restoring Windows User Profiles and Event Logs

To restore the Windows User Profiles and Event Logs, expand the `CONFIGURATION` object and select the items you want to restore.

User Profiles

Data Protector will not restore any files that are currently accessed. You have to log off the system and stop all the services that are running under the user account whose profiles you want to restore.

The restore session can be started from another system or by logging on the restore target system as a different user.

Deleted User Profiles

A user profile can only be restored when its location is already defined on the system. Individual files of existing user profiles or deleted profiles can be still restored as long as they exist among the system's profiles. Otherwise, you need to recreate them before restoring the files. Proceed as follows:

1. Log on as the user whose profile you want to restore in order to create a default user profile.

18.100
Paula

2. To keep the restored files unmerged, you can delete the files in the newly created profile before running a restore session.
3. Log off and start the restore session by logging on as a different user or by using another system.

The system may assign a different name to the user. In this case, use the `Restore As` option to restore the files to the newly assigned location.

When the restore has finished, restart the system.

User Disk Quotas

User Disk Quotas cannot be restored using Data Protector. The backed up information can be restored using Microsoft utilities.

Restoring Windows TCP/ IP Services

WINS, DHCP, DNS Servers

On a Windows Server that runs a Microsoft TCP/IP protocol and is configured as a **WINS Server**, a **DHCP Server,** or a **DNS Server,** you can restore the services that manage network communication.

To restore Windows TCP/IP services, expand the `CONFIGURATION` item and select `WNS`, `DHCP`, or `DNSServerDatabase`.

Each of these services is automatically stopped before the restore.

When the restore has finished, restart the system.

18 105
fauk

Restoring Novell Netware Filesystems

Use the standard restore procedure to restore Novell NetWare filesystems. Refer to “Standard Restore Procedure” on page 269.

Restoring Namespace Information and Volume Space Restrictions

To restore only volume space restrictions, specify the `Volume space restrictions only restore` option in the `Destination` page. The object selected for the restore must be a volume.

Data Protector restores Novell NetWare volume namespace information during a regular filesystem restore session. Namespace information is restored on a per-file/directory basis for the following Name Spaces: DOS, Mac, NFS, OS/2.

To restore files or directories, note the following:

- Backed up namespace information will be successfully restored only if the same Name Spaces are installed on the volume where you are attempting to restore the data.
- DOS namespace exists on each installed Novell NetWare volume and is always restored.
- A Mac's resource fork can only be restored to a volume that has the Mac namespace installed.
- Specific namespace information depends on the existence of NDS objects, such as user and group IDs in NFS namespace.
- After restoring the Queue objects, manually create a queue directory in the `SYS:SYSTEM` directory with the proper name *<queue_ID>*`.qdr`. Use the appropriate utility (`NWADMIN.EXE` or `SYSCON.EXE`) to retrieve the *<queue_ID>*.qdr from the NDS.
- NSS volumes on Novell NetWare 5.0 or later support files larger than 4 GB. You cannot restore any of these large Novell NetWare files to non-NetWare platforms.
- You cannot restore Novell NetWare sparse files that have been backed up in their native compressed format to non-NetWare platforms.

18.104
Paula

- You cannot restore Novell NetWare files that have been backed up in their compressed format to non-NetWare platforms.

Restoring File Ownerships and Trustees

Data Protector restores owner and trustee information on a per-file/directory basis. The owner and trustees of the file or directory are restored correctly if the relevant objects exist in the NDS database (Novell NetWare 4.X).

At restore time, select `Trustee only restore` and the appropriate Trustee Conflict Handling option in the `Destination` page of the `Restore` context.

Restoring the Novell NetWare CONFIGURATION

Data Protector enables you to restore the special data structure known as `CONFIGURATION`, which consists of the following components:

CONFIGURATION Components

- Server Specific Info
- Schema
- Root

To restore a component of the `CONFIGURATION` item, select the `CONFIGURATION` object and follow the standard restore procedure. Refer to Figure 6-6.

18.103
Paula

Figure 6-6 **Restoring the NetWare Configuration**

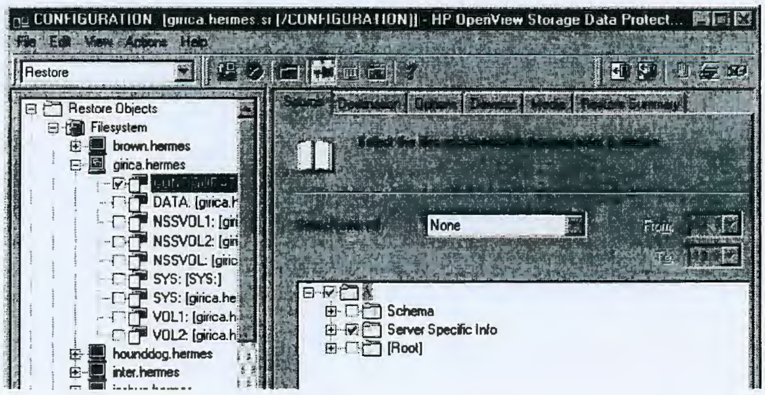

Restoring Novell NDS

Prerequisites

The prerequisites for performing a successful restore are the same as for a backup of the NDS database. Data Protector restores NDS objects in the same way as Novell NetWare filesystem data, except in the following cases:

- NDS objects cannot be restored to other Novell NetWare volumes.
- Container and leaf objects (treated as directories by Data Protector) cannot be restored into other container objects or as other container objects.

Restoring the NDS does not affect the current partitioning and replication in the NDS tree. If partitions and replicas exist when NDS information is restored, those partitions and replicas are fully utilized. If partition information does not exist at restore time, the entire tree structure is placed in one partition.

NOTE

Data Protector does not restore the NDS partitions and replica information. Partitions and replicas have to be manually reestablished.

13 102
Paula

For Novell NDS restore you can specify how to handle conflicts between the version currently on the disk and the backup version of a file. `File Conflict Handling` offers you three possible options: `Keep most recent`, `No overwrite`, and `Overwrite`. For more information on these options, refer to "Restore Options" on page 294.

Restoring NDS Schema and NDS Objects

Data Protector allows single NDS object restore. Within a Data Protector restore session, it is possible to:

- Restore the trees of the NDS using the `-trees` option
- Exclude a subtree of the NDS using the `-exclude` option
- Skip NDS objects using the `-skip` option
- Overwrite existing NDS objects using the `-overwrite` option

Troubleshooting

Sometimes an NDS restore session is completed successfully but some of the objects are not correctly restored and are marked as unknown. This happens when the NDS container object is deleted from NDS after the backup session. To solve this problem, restore this object again using the `-overwrite` option.

18.101
falta

Restoring OpenVMS Filesystems

Use the standard restore procedure to restore OpenVMS filesystems. Refer to "Standard Restore Procedure" on page 269.

What is Restored?

The directory structure and the files are restored, together with the following filesystem information:

- The directory and file attributes.
- ACL (Access Control List) if available (see Limitations below).
- Secondary file entries.

  Files with multiple directory entries are backed up once using the primary path name. Secondary path entries are saved as soft links. During a restore, these extra path entries are restored. Refer to "Limitations" in "Backup Specification Configuration Procedure" on page 205.

Files can be restored to mounted FILES-11 ODS-2 or ODS-5 volumes only.

Limitations

- For files and directories saved on any other operating system platform not all file attributes are restored and no ACL will be restored in this case.
- Directories that are created during a restore but have not been included in a save will get the attributes of the first file restored in the directory unless disabled by the `-no_protection` option.
- Any file specifications that are entered into the GUI or passed to the CLI must be in UNIX style syntax:

  ```
  /disk/directory1/directory2/filename.ext.n
  ```

  - The string should begin with a slash, followed by the disk, directories, and filename, separated by slashes.
  - Do not place a colon after the disk name.
  - A period should be used before the version number instead of a semi-colon.

18.100
Paula

— File specifications for OpenVMS files are case insensitive.

For example:

An OpenVMS file specification of:

```
$1$DGA100:[USERS.DOE]LOGIN.COM;1
```

must be specified in the form:

```
/$1$DGA100/Users/Doe/Login.Com.1
```

- There is no implicit version number. You always have to specify a version number. Only file versions selected for the backup will be restored. If you wish to include all versions of the file, select them all in the GUI window, or, using the CLI, include the file specifications under the `Only (-only)` option, including wildcards for the version number, as follows

```
/DKA1/dir1/filename.txt.*
```

- If you restore to a location other that the original location, only the disk device and starting directory are changed. The original directory path is added to the destination path to form the new restore location.
- If the `Restore Time Attributes (-notouch)` option is disabled during a restore, the last accessed date will be updated with the current date and time on ODS-5 disks. On ODS-2 disks, the original dates will be set on the files.
- A file saved as a soft link will be restored using the equivalent of a "DCL SET FILE/ENTER" command. No data will be restored in this case. The soft link entered points to the primary path/filename for this file from the time the file was saved. If the primary path/filename does not exist or was not restored, the creation of the soft link will fail.

  To make a restored copy of an OpenVMS system disk bootable, the OpenVMS WRITEBOOT utility has to be used to write a boot block after the disk has been restored.
- The `Move Busy Files (-move)` and `Restore Sparse Files (-sparse)` options are not available on OpenVMS.
- Files backed up from an ODS-5 disk on an OpenVMS system that have extended filesystem names (i.e. upper and lower case letters, Unicode characters, etc) may not be restored to an ODS-2 disk.
- Files being restored are always locked regardless of whether the `Lock`

18.C99
Paula

`Files during Restore (-lock)` option is enabled or disabled.

- The default device and directory for pre- and post-exec command procedures is `/omni$root/bin`. To place the command procedure anywhere else the file specification must contain the device and directory path in UNIX style format: For example: `/SYS$MANAGER/DP_SAVE1.COM`
- If the `Restore Protection Attributes (-no_protection)` option is disabled, the files are created with the default owner, protection and ACL.
- When specifying wildcards for `Skip (-skip)` or `Only (-only)` filters, use '*' for multiple characters and '?' for single characters.

18.095
Paula

Restore Options

Data Protector offers a set of comprehensive restore options that allow fine-tuning of a restore. All these options have default values which are appropriate in most cases.

Restore options depend on the data being restored. For example, restore options for a filesystem are different from those for a disk image restore.

List of Restore Options

The following list of restore options can be set for a particular object. They apply to all the data restored from the backed up object.

General Restore Options

Target Client By default, you restore to the same client system from which the data was backed up. You can select another system in your cell from the drop-down list. The Disk Agent is started on the selected client system and the data is restored there. You need to have the `Restore to other clients` user right to be able to restore to another client system.

Omit Deleted Files This option removes files that were deleted between a full and an incremental backup. It recreates the state of your disk or directory as it was at the time when the last incremental was run. It does not apply to files that were created after the incremental backup. By default, this option is disabled.

WARNING

If, between the full and incremental backup, a user has created files with the same name as those that have been deleted, the newly created files are also deleted.

When using the `Restore As` functionality, a file on the new location will be deleted if it was removed from the original location between the full and incremental backup and its modification time is older than the time of the last incremental backup.

18C97
Paula

The time on the Cell Manager and clients must be synchronized for the `Omit Deleted Files` option to function properly.

Move Busy Files This option is relevant if a file on the disk is being used by an application when a restore wants to replace this file. The option is used with the `Keep most recent` or `Overwrite` options. By default, this option is disabled.

On UNIX systems, Data Protector moves the busy file *`filename`* to #*`filename`* (adds a hash in front of the filename). The application will keep using the busy file until it closes the file. Subsequently, the restored file is used.

On Windows systems, the file is restored as *`filename.001`*. All applications keep using the old file. When the system is rebooted, the old file is replaced with the restored file.

List Restored Data When this option is enabled, Data Protector displays the names of the files and directories in the monitor window as the objects are being restored. By default, this option is disabled.

Display Statistical Information When this option is enabled, Data Protector reports statistical information (such as size and performance) for each object that is restored. You can view the information in the monitor window. By default, this option is disabled.

Omit Unrequired Incrementals This option enables repositioning within a medium when restoring individual files of a specific object. The Media Agent restores a specific item, repositions itself directly on the next requested item, and continues the restore. This improves restore performance when restoring multiple single files. Note that several Disk Agents may be started per object. Disable this option if you intend to restore empty directories. By default, this option is enabled.

18 caé
Paula

Restore Sparse Files This option restores sparse files in their original compressed form. This is important because sparse files can consume additional disk space unless they are restored in their original form. By default, this option is disabled.

This option applies to UNIX sparse files only. Windows sparse files are always restored as sparse.

Lock Files During Restore This option denies access to files during the restore. By default, this option is disabled.

Restore Time Attributes This option preserves the time attribute values of each restored file. When this option is disabled, Data Protector sets the time attributes of the restored objects to the current date and time. By default, this option is enabled.

Restore Protection Attributes This option preserves the original protection attributes of each restored file. If this option is disabled, Data Protector applies the protection attributes of the current restore session. By default, this option is enabled.

On Windows systems, this option applies to file attributes only. Security information is always restored, even when this option is disabled.

Pre- and Post-Exec Commands

For general information on `pre-` and `post-exec` commands, refer to "Pre- and Post-Exec Commands" on page 250. For examples of these commands on UNIX, refer to "Examples of Pre-Exec and Post-Exec Commands for UNIX" on page A-20. Note that `pre-` and `post-exec` commands are executed before and after the restore of each object, and not the entire restore session.

Pre-Exec This option allows you to enter a command to be executed before the restore of each object is initiated. This command must return success for Data Protector to proceed with the restore. The `pre-exec` command is executed on the client system where the Disk Agent is running. On how to specify the command, refer to online Help.

15.095
Paula

Post-Exec This option allows you to enter a command to be executed after the restore of each object is completed. The `post-exec` command is executed on the client system where the Disk Agent is running. On how to specify the command, refer to online Help.

File Conflict Handling Options

Keep Most Recent If this option is selected, the most recent versions of files are kept. If a file on the disk is newer than the backed up version, the file is not restored. If a file on the disk is older than the backed up version, the file is overwritten with the newer version from the backup. By default, this option is enabled.

No Overwrite If this option is selected, files that exist on the disk are preserved. This means that they are not overwritten by other versions of these files from the backup. Only non-existing files are restored from the backup. By default, this option is disabled.

Overwrite If this option is selected, existing files on the disk are replaced with files from the backup. By default, this option is disabled.

Active Directory Specific Options

Authoritative The Active Directory database is *not* updated after the restore, and the restored data overwrites the existing data in the target destination. An authoritative restore can only be performed by running `ntdsutil.exe` from the command prompt after the restore session has finished.

Nonauthoritative The `Nonauthoritative` replication mode is the default option. The Active Directory database is updated after the restore using standard replication techniques.

Prımary The `Primary` replication mode allows you to keep the NT Directory Service online, and is used when you restore `FileReplicationService` along with the Active Directory service. This option must be used when all replication partners for a replicated share have been lost. With regard to the Certificate Server and the Active Directory Server, `Primary` is the same as `Authoritative`.

18.094
Paula

User Defined Restore Variables

You can use variables (a variable name and its value) for flexible operations on some platforms and integrations with Data Protector. For detailed steps, refer to the online Help index keyword “setting user definable restore variables”.

Figure 6-7 **User Defined Restore Variables**

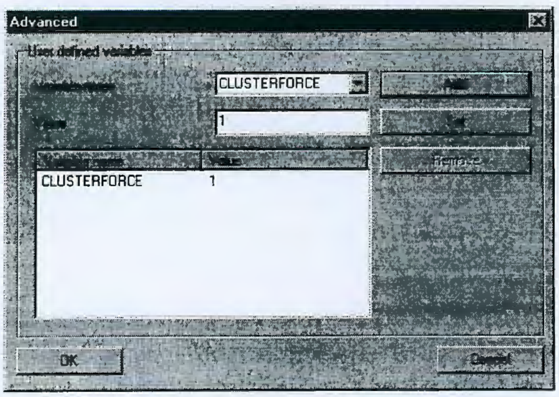

CLUSTERFORCE
1

CLUSTERFORCE	1

18.093
Paula

Restore Techniques

The following restore techniques apply to the UNIX and Windows platforms.

Restoring Files to Different Paths

By default, Data Protector restores data to the same client and directory from which it was backed up. You can restore your data to a different client system and directory. For individual files and directories, you can specify a different path and different name.

Different Location for an Object

In the `Destination` page of your restore, you can specify a different restore location for an object selected for restore:

- With appropriate user rights you can restore to a different client system by selecting the client system in the `Target client` drop-down list. By default, Data Protector restores the object using the same directory structure. For example, if the object was originally backed up from the `C:\temp` directory on system A, it will restore the data to the `C:\temp` directory on system B.
- You can restore to a different directory by selecting the `Restore to new location` option, and then entering or browsing for a new path in the text box. The original path is appended to the new one. For example, if data was backed up from the `C:\sound\songs` directory and you enter `\users\bing` as the new path, the data is restored to the `C:\users\bing\sound\songs` directory.

For detailed steps, refer to the online Help index keyword "location options for restore".

Different Location for Individual Files

The individual location specified under the `Restore As/Into` option overrides the default destination specified in the `Destination` property page.

You can restore individual files and directories to different paths and under a different name using the `Restore As/Into` option available from the `Source` property page of your restore.

18.09.2
Paula

This capability is available for the initially selected tree node (directory) and for tree nodes that are not hierarchically dependent on any already selected tree nodes. A selected tree node is indicated by a blue check mark, and a dependent tree node is indicated by a black check mark.

`Restore Into` appends the source path to the new one entered under `Location`. For example, if the `colors.mp3` file was backed up from the `C:\sound\songs` directory and you enter `\users\bing` as the new path, the file is restored to the `C:\users\bing\sound\songs` directory.

`Restore As` replaces the source path with the one entered under `Location`. The destination path can be a new directory or an existing one. You can rename the files and directories as you restore them. For example, if the `colors.mp3` file was backed up from the `C:\sound\songs` directory and you enter `\users\bing\colors.mp3` as the new path, the file is restored to the `C:\users\bing` directory.

CAUTION

Consider the risk of deleting data with the `Overwrite` option enabled when:

- Specifying restore under a name that already exists
- Entering an existing path without specifying the file or directory name

For example, when you enter the new path `\users\bing` in the `Location` text box when restoring the file `colors.mp3`, but you do not enter the name of the file, then the `colors.mp3` file will be restored as `bing`. What used to be the `bing` directory is deleted and substituted with the restored file.

Restoring Files in Parallel

What Is Parallel Restore?

Parallel restore lets you restore files to multiple disks at the same time, assuming that the disks have been backed up to the same device using a concurrency higher than 1. This improves the speed of the restore. This behavior is complementary to a parallel backup, where files from multiple disks are backed up concurrently to the same device.

18.091
Paula

Figure 6-8 **Restoring Files in Parallel**

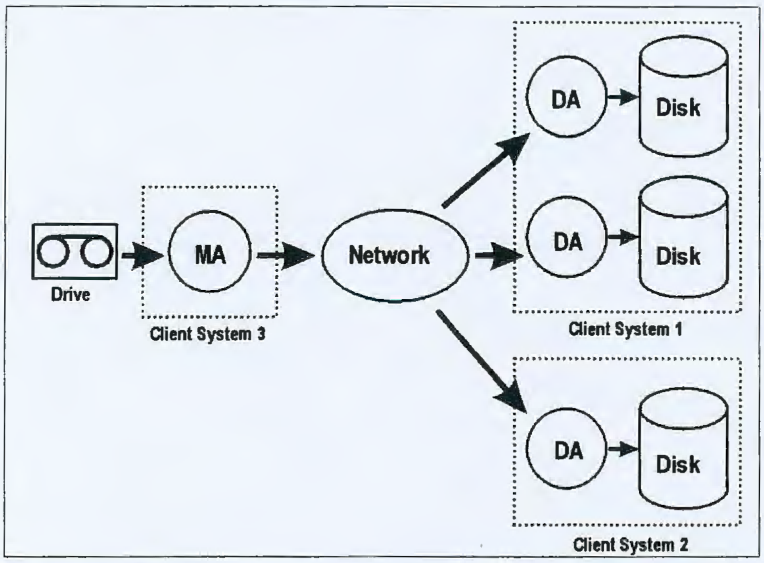

The figure shows an example of restoring files in parallel from one medium. Each object uses a different DA.

How to Run a Parallel Restore

Select the data that you want to restore to different disks and start the restore. Data Protector asks you if you want to perform a parallel or single restore. Choosing parallel restore enables multiple Data Protector Disk Agents to run in parallel. Refer to See "Selecting Your Data for Restore" on page 270.

Viewing Files Not in the IDB

Data Protector allows you to view and restore data directly from backup media even though the information about this data is no longer in the IDB.

When to Restore Directly from Media

The following must apply in this case:

- You have removed information about backed up data or media from the IDB.

18.090
Paula

- The catalog protection has expired. Refer to "Most Frequently Used Backup Options" on page 227 for more information about data and catalog protection.
- The media are not from the same Data Protector cell and, as such, are not recognized in the IDB of the cell. In this case, you need to import it first.

Prerequisite

A large amount of memory on the Cell Manager is required. The amount of memory needed can be estimated by using the following formula: *number_of_files* multiplied by 200 bytes.

Limitations

- You cannot list database application objects from the media.
- Files that span several media cannot be restored directly fro media. All media needed to restore the file have to be imported, and then the file can be restored using the `List From Database` option.

How to Restore Directly from Media

To restore data directly from media, click `List From Media` in the `Actions` menu of your restore context, and follow the `Restore from media` wizard. For detailed steps, refer to the online Help index keyword "restoring directly from media".

Restoring Files in Use

Data Protector allows you to back up and restore files, such as databases or word processing documents, that are in use (opened) by other applications.

Data Protector provides restore options that allow you to specify the behavior when files being restored are in use by setting the **Lock Files During Restore** and the **Move Busy Files** options. Refer to "Restore Options" on page 294.

Restoring by Query

Data Protector provides the `Restore by Query` task, which searches for files and directories you want to restore and restores them.

Figure 6-9 **Restore by Query**

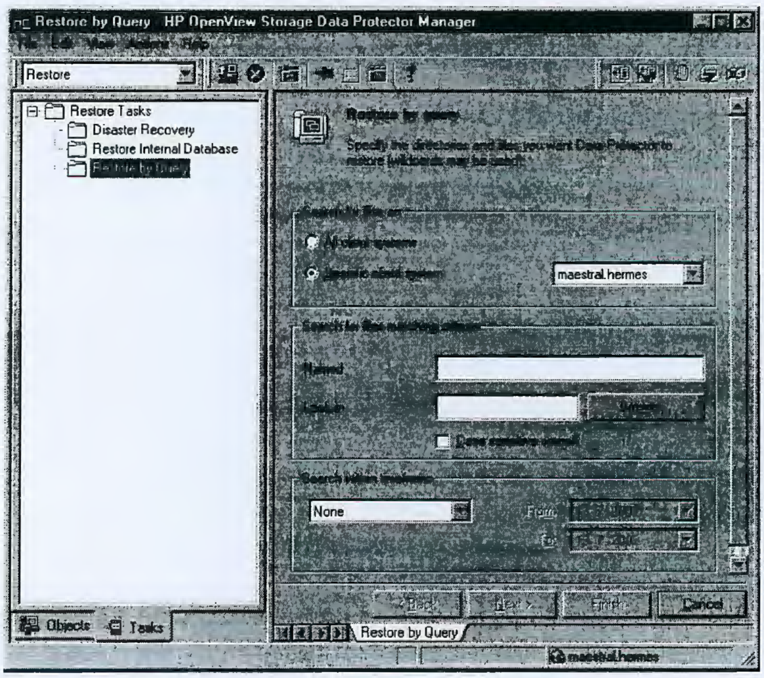

Restore by Query HP OpenView Storage Data Protector Manager

To search for a file or a directory, you need to know at least a part of its name. The `Restore by Query` task enables you to search for files and directories backed up from a specific client system in a specified timeframe, or for files and directories with specific criteria (using wildcards: *, ?).

When to Restore by Query

You may want to use this task in the following cases:

- You do not know the full path where a file or directory that you want to restore is located.
- You do not know on which system (object) the file or directory you are looking for is located.

18088
Paula

How to Restore by Query

Start the `Restore by Query` task from the `Restore` context of the `Data Protector Manager`. Use the `Tasks` navigation tab. See Figure 6-9. For detailed steps, refer to the online Help index keyword “restore by query”.

Skipping Files for Restore

Data Protector allows you to skip certain files during restore. By using wildcards (`*` or `?`), you can skip files matching specific criteria. For example, entering `*.exe` skips the files that end in `.exe`.

How to Skip Files for Restore

In the `Source` property page of your restore, select the tree node to be restored and right-click it to open its properties. In the `Skip` property page, specify the criteria to match the files to be skipped. For detailed steps, refer to the online Help index keyword “skipping files”.

Figure 6-10 **Skipping Files for Restore**

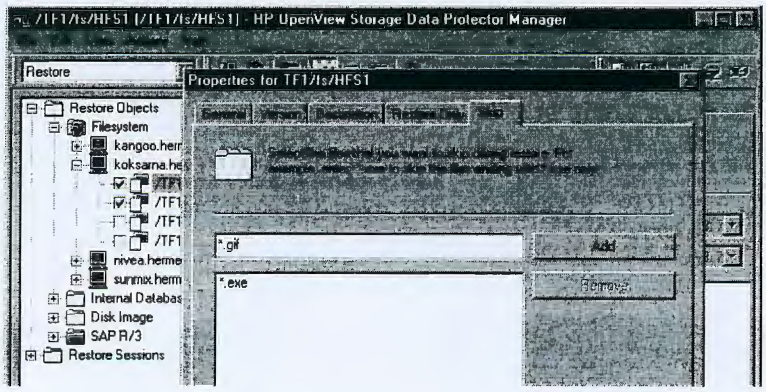

*.gif

*.exe

Selecting Only Specific Files (Matching) for Restore

Data Protector allows you to restore only specific files. By using wildcards (`*` or `?`), you can restore files matching specific criteria. For example, entering `*.exe` restores only the files that end in `.exe`.

How to Match Files for Restore

In the `Source` property page of your restore, select the tree node to be restored and right-click it to open its properties. In the `Restore Only` property page, specify the criteria to match the files to be restored. For detailed steps, refer to the online Help index keyword “selecting only specific files for restore”.

18087
Paula

Figure 6-11 **Matching Files for Restore**

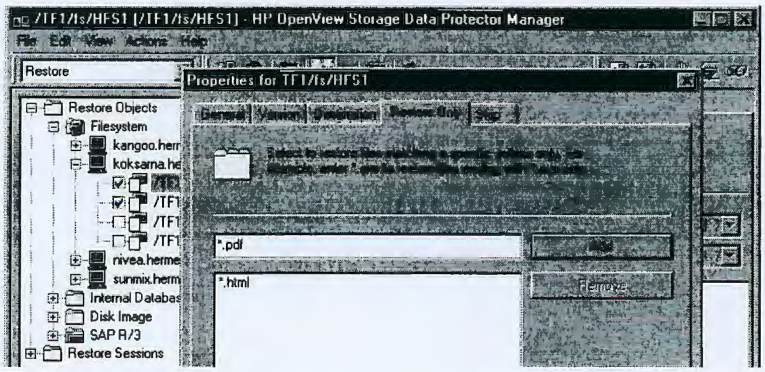

*.pdf

*.html

Restoring Files and Directories Manually

You need to restore a file or a directory manually when you can no longer browse for the file or directory. This happens when the catalog protection for your data has expired, or when backup was done using the `No log` option.

Prerequisite

To add a file or a directory manually, you need to know the exact path and the name of the file or the directory. The file and path names are case sensitive.

How to Add Files and Directories Manually

In the `Restore Summary` page of your restore, write the exact path and name of the file or the directory, and then click `Add`. For detailed steps, refer to the online Help index keyword “manually restoring files or directories”.

18056
Paula

18.035
Paulo

7 Monitoring, Reporting, Notifications, and the Event Log

18.084
Paula

In This Chapter

This chapter consists of the following sections:

"Monitoring Sessions" on page 309

"Monitoring Several Cells Simultaneously" on page 314

"Data Protector Reporting" on page 315

"Configuring Reports Using the Data Protector GUI" on page 335

"Running Reports and Report Groups Using the Command-Line Interface" on page 339

"Data Protector Notifications" on page 342

"Configuring Reports and Notifications on the Web" on page 353

"Data Protector Event Log" on page 356

You can monitor several cells at the same time using the Manager-of-Managers functionality. See Chapter 8, "Manager-of-Managers Environment," on page 359 for more information.

If you do not have access to the Data Protector user interface, you can still view reports and set notifications using your Web browser. See "Configuring Reports and Notifications on the Web" on page 353 for information on how to do this.

18.083
Paula

Monitoring Sessions

Data Protector allows you to manage running sessions and to respond to mount requests. You can view the status of sessions, their type, owner, session ID, and start time, as well as the names of the corresponding backup specifications.

When you run an interactive backup, restore, or media management session, a monitor window opens, showing the objects and backup devices used, and the messages generated during the session. Note that even if the user interface is closed, the session continues.

You can change the level of reported messages during a backup or restore session by changing the Report level option when configuring a backup specification or when starting a restore session.

NOTE

Only the Data Protector users in the `Admin` group and those granted the `Monitor` user rights are given access to the Data Protector monitoring functionality.

Viewing Currently Running Sessions

Currently running sessions are available in the Data Protector GUI, in the `Monitor` context. As soon as the sessions finish or are aborted, they disappear from the current view. To learn how to view finished sessions, see the next section.

NOTE

You can view currently running sessions only if the pre-exec script has finished. If no sessions appear in the window, there are no running sessions or they are still in the pre-exec stage.

Use the following steps to monitor currently running sessions:

1. In the `Data Protector Manager`, switch to the `Monitor` context.

13.082
Paula

The progress and status of current sessions appear in the Results Area. You can sort the sessions by clicking the column headings in the Results Area.

TIP To view the details of a running session, double click the session. The detailed monitor view of the session appears.

Viewing Finished Sessions

As soon as a session is finished or is aborted, it is no longer displayed in the Data Protector GUI, `Monitor` context. The finished session is moved to the `Internal Database` context.

Use the following steps to view finished sessions:

1. In the Data Protector GUI, switch to the `Internal Database` context.

   If you are running the Manager-of-Managers, select `Monitor` in the Context List, and then select a Cell Manager of your choice. From the Tools menu, select `Database Administration` to open a new Data Protector GUI with the `Internal Database` context selected.

2. Expand `Sessions` to display all the sessions stored in the IDB.

   The sessions are sorted by date.

   To view details on a specific session, double-click the session.

Responding to Mount Requests

Data Protector issues a mount request in the following cases:

- The end of the currently used medium has been reached and Data Protector needs a free medium.
- A mail slot is open. In this case, shut the mail slot.

You respond to a mount request to confirm that the needed medium is in a device. Use the following procedure to respond to the mount request while monitoring the session:

1. In the Context List, select `Monitor`.

18.081
Paula

2. Insert the needed medium into the device. If you have a library device, it is not necessary to use the slot requested by the mount request.
3. In the Results Area, double-click the session with the mount request status to display details about the session.
4. Select the device with the mount request status.
5. In the `Actions` menu, click `Confirm Mount Request`. The status of the session and device changes to `Running`.

TIP You can also right-click the device with the mount request status and select `Confirm Mount Request`.

Restarting Failed Backups

During backup, some systems may not be available because they were shut down, there were some networking problems, and so on. This results in some systems not being backed up or being backed up just partially - some objects failed.

This section gives you detailed instructions on how to restart failed backup sessions. For more information on how to manage failed backups, see "Managing Failed Backups" on page 263.

You cannot restart failed sessions that are the result of an unsaved backup specification.

After you have resolved the related problems, restart a failed session, as follows:

1. In the `Data Protector Manager`, switch to the `Internal Database` context.

   If you are running the Manager-of-Managers, select `Clients` in the Context List, and then expand `Enterprise Clients`. Select a Cell Manager with the failed backup. From the `Tools` menu, select `Database Administration` to open a new Data Protector window with the `Internal Database` context displayed.
2. Under the `Internal Database` item, expand the `Sessions` item.
3. In the Results Area, search for your backup.

18.080
Paula

You can sort your sessions using the buttons on the top of each of the columns.

4. Right-click on your failed session, and then select `Restart Failed Object`.
5. Click `Yes` to confirm.

Aborting Running Sessions

You can abort a session if you want to stop a backup, restore, or media management operation. A backup copy or restored data will exist only for data that was backed up or restored before you aborted the session.

1. In the Context List, click `Monitor`. The progress and status of current sessions appear in the Results Area.

   If you are running the Manager-of-Managers, expand the `Enterprise Monitor` in the Scoping Pane, and then select the Cell Manager you want to monitor. The progress and status of current sessions appear in the Results Area.
2. Click the column headings to sort the sessions.
3. Right-click the session that you wish to abort and select `Abort`.

If you abort a backup session while it is still determining the sizes of the disks that you have selected for the backup, it does not abort immediately. The backup is aborted once the size determination (treewalk) is completed.

TIP If you started a backup, restore, or media management session interactively, you can also abort the session in the Data Protector `Backup`, `Restore`, or `Devices & Media` context respectively.

Changing the Amount of Messages Shown

You can change the level of reported messages for backup and restore sessions by changing the Backup and Restore options.

See “Using Backup Options” on page 225 for information on which backup options affect your displayed messages.

18.079
Paula

See “Restore Options” on page 294 for information on which restore options affect your displayed messages.

18.078
Paula

Monitoring Several Cells Simultaneously

You can monitor several cells at the same time using the Manager-of-Managers functionality.

See Chapter 8, "Manager-of-Managers Environment," on page 359 for more information.

18077
Paula

Data Protector Reporting

What Is Reporting?

Data Protector reports provide various information on your backup environment. For example, you can check the status of the last backup, check which systems in your network are not configured for backup, check the status of devices, and more.

Data Protector reporting represents a powerful, customizable, and flexible tool for managing and planning your backup environment.

You can configure reports and report groups using the Data Protector GUI or any Web browser with Java support.

NOTE

Only the Data Protector users in the Admin group and those granted the `Reporting, notifications and event log` user rights are given access to the Data Protector reporting functionality.

Prerequisite

The Data Protector user under whose account the CRS service is running should not be removed. This user is configured by default at installation time. On a Windows Cell Manager, this is the user under whose account the installation was performed. On a UNIX Cell Manager, this is the root user of the Cell Manager.

Report Groups

You can gather various reports in a report group, which can be scheduled, started interactively, or triggered by a notification.

Starting Reports

Reports can be started using the Data Protector GUI, the Data Protector command-line interface, the Data Protector Web reporting interface, the Data Protector scheduler, a notification event, or a post-exec script that includes a Data Protector command-line interface command.

Reports on Multiple Cells

Reporting is also available for a multiple cell configuration when you use the Manager-of-Managers functionality.

Report Parameters

Reports can be customized by configuring optional input parameters (optional selections). Some input parameters allow multiple selections.

18.076
Paula

If no optional input parameters (optional selections) are specified when configuring a report, a default value is set, which is *<all>* in case of objects and *<no time limit>* in case of time frames.

To configure a report or report group, you need to provide the following information:

- name of the report
- type of report
- send method
- recipient(s)
- format

All other input parameters (selections) depend on the type of the report.

Report Formats

Output of the reports is provided in various formats and optionally displays input parameters (selections), too. Refer to “Report Formats” on page 329.

Report Send Methods

Reports can be sent using various methods. Refer to “Report Send Methods” on page 331.

18 075
Paula

Report Types

Data Protector provides various types of reports, as shown in Table 7-1:

Table 7-1

Backup Specifications	Provides information on backups, such as average size of backed up objects, schedule of backups, filesystems not configured for backup, and so on.
Configuration	Provides information on the configuration of the Data Protector cell, on devices not configured for backup, on systems not configured for backup, and so on.
IDB	Provides information on the size of the IDB and on the results of the database purge sessions.
Pools and Media	Provides information on media pools and used media.
Sessions in Time Frame	Provides information on backup sessions that have run in a specified period of time.
Single Session	Provides detailed information on a specific session.

Backup Specification Reports

The following table lists the Backup specification reports. Backup specification reports provide information on backups, such as average size of backed up objects, schedule of backups, filesystems not configured for backup, and so on.

18.074
Paula

For supported formats, refer to “Report Formats” on page 329.

Table 7-2 **Backup Specification Reports**

Report and `omnirpt` Option	**Description**	**Required Selections**	**Optional Selections**	**Supported Formats**
Trees in Backup Specification `dl_trees`	Lists all trees in the specified backup specification. It also shows names of drives and the name of a tree.	none	• Backup Specifications • Backup Specification Group	all formats
Objects without Backup `obj_nobackup`	Lists all objects that are part of a backup specification and do not have a valid backup (successfully completed backup, the protection has not yet expired).	none	• Backup Specifications • Backup Specification Group • Number of Days[1]	all formats
Object’s Latest Backup `obj_lastbackup`	Lists all objects for each specified backup specification, together with the last full and the last incremental backup time.	none	• Backup Specifications • Backup Specification Group • Number of Days[1]	all formats
Average Object Size `obj_avesize`	Displays the average size of an object in the specified backup specification. It displays the size of the full and the incremental backup of the object.	none	• Backup Specifications • Backup Specification Group • Number of Days[1]	all formats

15073
Paula

Table 7-2 **Backup Specification Reports**

Report and `omnirpt` Option	**Description**	**Required Selections**	**Optional Selections**	**Supported Formats**
Not Configured Filesystems `fs_not_conf`	Lists all disks (filesystems) that are not configured in any of the selected backup specifications.	none	• Backup Specifications • Backup Specification Group	all formats
Backup Specification Information `dl_info`	Shows the following information: backup specification name, type, group, owner, and pre & post exec commands for all specified backup specifications.	none	• Backup Specifications • Backup Specification Group	all formats
Backup Specification Schedule `dl_sched`	Lists the next backup time for each specified backup specification.	none	• Backup Specifications • Backup Specification Group	all formats

1. Counted from the moment of starting the report backwards.

18.072
Paula

Configuration Reports

The following table lists the Configuration reports. Configuration reports provide information on the configuration of the Data Protector cell, devices, systems not configured for backup, and so on. For supported formats, refer to "Report Formats" on page 329.

Table 7-3 **Configuration Reports**

Report and `omnirpt` Option	Description	Required Selections	Optional Selections	Supported Formats
Cell Information `cell_info`	Lists Data Protector cell related information (number of clients, backup specifications, media management server, licensing server).	none	none	all formats
Configured Clients Not Used by Data Protector `hosts_unused`	Lists all configured clients that are not used for backup and do not have any device configured.	none	none	all formats
Configured Devices Not Used by Data Protector `dev_unused`	Lists configured devices that are not used for backup at all.	none	none	all formats
Look up Schedule `lookup_sched`	Lists backup specifications that are scheduled to start in the next specified number of days.	Number of Days	none	all formats
Clients Not Configured for Data Protector `hosts_not_conf`	Lists clients in the selected domains that are not part of the current cell.	Network Ranges	none	all formats
Licensing `licensing`	Lists all licenses with their total and available amount.	none	none	all formats

18.071
Paula

Table 7-3 **Configuration Reports**

Report and omnirpt Option	Description	Required Selections	Optional Selections	Supported Formats
Client Backup `host`	Lists information about the specified clients such as: filesystems not configured, all objects, and all objects with a valid backup. Reports also list times and average sizes.	Host Name	none	all formats

IDB Reports

The following table lists the IDB reports. IDB reports provide information on the size of the IDB and on the results of the database purge sessions. For supported formats, refer to "Report Formats" on page 329.

Table 7-4 **IDB Reports**

Report and omnirpt Option	Description	Required Selections	Optional Selections	Supported Formats
IDB Size `db_size`	Provides a table that contains information about the Media Management DB, Catalog DB, DB extension files, statistics for DC binary files, SMBF, and SIBF and low DB disk space.	none	none	all formats
IDB Purge `db_purge`	Lists all purged sessions together with the following information: start time, end time, duration, inactivity time, and number of the file name records and the amount of Mb read.	none	none	all formats

18.070
Paula

Table 7-4 **IDB Reports**

Report and `omnirpt` Option	Description	Required Selections	Optional Selections	Supported Formats
Purge Preview `db_purge_preview`	Lists the following information: overall number of filenames in database (in thousands), estimated number of obsolete filenames in database (in thousands) and estimated duration of database purge (in seconds).	none	none	all formats

18.069
Paula

Table 7-4 **IDB Reports**

Report and omnirpt **Option**	**Description**	**Required Selections**	**Optional Selections**	**Supported Formats**
System Dynamics `db_system`	Lists for each Data Protector client in the cell: the number of filenames (in thousands) in the IDB, the number of active filenames (in thousands) in the IDB, the IDB filenames growing ratio (new filenames per day), the number of deleted filenames in the IDB per day, active growth per year, and a dynamics indicator (medium/high/low/critical). The filenames that are not active are filenames of the backed up files in the IDB that have no associated file versions in the IDB. The active growth per year is calculated in two ways: If there is no Data Protector database purge session recorded in the Data Protector database, the active growth per year is calculated on the basis of data in last 11 days and then extrapolated to one year. If there is a Data Protector database purge session recorded in the Data Protector database, the active growth per year is calculated on the basis of data in the time span since the last Data Protector database purge session and then extrapolated to one year.	none	none	all formats

18.065
Paula

Pools and Media Reports

The following table lists the Pools and Media reports. Pools and media reports provide information on media pools and used media. For supported formats, refer to "Report Formats" on page 329.

Table 7-5 **Pools and Media Reports**

Report and omnirpt Option	Description	Required Selections	Optional Selections	Supported Formats
Extended List of Media `media_list_extended`	Lists all media matching the specified search criteria. For each medium, the following information is provided: medium ID, medium label, medium location, medium condition, medium protection, used and total space (MB), time when medium was last accessed, media pool and media type, and the backup specifications that have used this medium during the backup.	none	• Description • Locations • Poolnames • Media Types (DDS, DLT and so forth) • Condition • Expiration[1] • Timeframe[2] • Library Devices	all formats
List of Pools `pool_list`	Lists all pools matching the specified search criteria. For each pool the following information is provided: pool name, description, media type, total number of media, number of full and appendable media containing protected data, number of free media containing no protected data, number of poor, fair and good media.	none	• Pool Names • Locations • Media Types (DDS, DLT, and so forth) • Library Devices • Timeframe[2]	all formats

18.067
Paula

Table 7-5 **Pools and Media Reports**

Report and `omnirpt` Option	Description	Required Selections	Optional Selections	Supported Formats
Media Statistics `media_statistics`	Reports the statistics on the media matching the search criteria. The following information is provided: number of media; number of scratch media; number of protected, good, fair and poor media; number of appendable media; and total, used, and free space on media.	none	• Description • Locations • Poolnames • Media Types (DDS, DLT and so forth) • Status • Expiration[1] • Timeframe[2] • Library Devices	all formats
List of Media `media_list`	Lists all media matching the specified search criteria. For each medium, the following information is provided: medium ID, medium label, medium location, medium condition, medium protection, used and total space (MB), time when medium was last accessed, and media pool and media type.	none	• Description • Locations • Poolnames • Media Types (DDS, DLT and so forth) • Condition • Expiration[1] • Time frame[2] • Library Devices	all formats

1. The following are possible:

Don't care / Unprotected / Protected; the last with the following suboptions:

18066
Paula

Number of remaining days in which the data protection will expire, counted from the moment of starting the report / Never

2. Timeframe in which the medium was used for a backup.

Relative time: the first parameter sets the starting point of the timeframe (number of hours counted from the moment of starting the report backwards), the second parameter sets the end point of the timeframe (number of hours counted from the starting point).

Absolute time: the first parameter sets the starting point of the timeframe (date), the second parameter sets the end point of the timeframe (date).

Sessions in Timeframe Reports

The following table lists the Data Protector Sessions in Timeframe reports. Sessions in Timeframe reports provide information on backup sessions that have run in a specific period of time. For supported formats, refer to "Report Formats" on page 329.

Table 7-6 **Sessions in Timeframe Reports**

Report and `omnirpt` Option	Description	Required Selections	Optional Selections	Supported Formats
List of Backup Sessions `list_sessions`	Lists all sessions in the specified timeframe.	TimeFrame[1]	• Backup Specifications • Backup Specification Group	all formats
Session Flow `session_flow`	Graphically presents the duration of each session for the specified timeframe. A flow chart of the backup sessions matching the search criteria is shown.	TimeFrame[1]	• Backup Specifications • Backup Specification Group	HTML

18.065
Paula

Table 7-6 **Sessions in Timeframe Reports**

Report and `omnirpt` Option	Description	Required Selections	Optional Selections	Supported Formats
Device Flow `device_flow`	Graphically presents the usage of each medium. A flow chart of the backup sessions matching the search criteria is shown.	TimeFrame[1]	• Backup Specifications • Backup Specification Group	HTML
Used Media `used_media`	Lists media that have been used during the backup sessions in the specified timeframe, together with their statistics.	TimeFrame[1]	• Backup Specifications • Backup Specification Group	all formats
Client Statistics `host_statistics`	Lists clients and their backup status statistics. Only the clients that match the search criteria are listed.	TimeFrame[1]	• Backup Specifications • Backup Specification Group • Hostnames	all formats
Backup Statistics `backup_ statistics`	Shows statistics about backup status in the selected timeframe.	TimeFrame[1]	• Backup Specifications • Backup Specification Group	all formats
Backup Errors `backup_errors`	Displays a list of messages that occurred during backup. The messages are grouped by client.	TimeFrame[1]	• Backup Specifications • Backup Specification Group • Hostnames • Message Level	all formats

18004
Paula

Table 7-6 **Sessions in Timeframe Reports**

Report and `omnirpt` **Option**	**Description**	**Required Selections**	**Optional Selections**	**Supported Formats**
Extended Report on Used Media `used_media _extended`	Provides extended information about all media that were used in the selected session.	TimeFrame[1]	• Backup Specifications • Backup Specification Group	all formats

[1] Timeframe in which the medium was used for a backup.

Relative time: the first parameter sets the starting point of the timeframe (number of hours counted from the moment of starting the report backwards), the second parameter sets the end point of the timeframe (number of hours counted from the starting point).

Absolute time: the first parameter sets the starting point of the timeframe (date), the second parameter sets the end point of the timeframe (date).

Single Session Report

The following table lists the Data Protector Single Session Reports. For supported formats, refer to "Report Formats" on page 329.

Table 7-7 **Single Session Reports**

Report and `omnirpt` **Option**	**Description**	**Required Selections**	**Optional Selections**	**Supported Formats**
Single Session `single_session`	Displays all relevant information about a single Data Protector backup session.	Session ID	Message Level	all formats
Session Objects `session_objects`	Lists all backup objects and their statistics that took part in a selected session.	Session ID	none	all formats

18063
Paula

Table 7-7 **Single Session Reports**

Report and omnirpt Option	**Description**	**Required Selections**	**Optional Selections**	**Supported Formats**
Session per Client `session_hosts`	Provides information about each client that took part in the selected session. Using the `Generate multiple reports` option, this report can be split into smaller reports, one for each client.	Session ID	Message Level	all formats
Session Devices `session_devices`	Provides information about all devices that were used in the selected session.	Session ID	none	all formats
Session Media `session_media`	Provides information about all media that were used in the selected session.	Session ID	none	all formats

Report Formats

Data Protector reports can be produced in various formats.

If you view each report individually, the report is displayed in the Data Protector Manager and you do not have to choose the report format.

If you group reports into report groups so that you can send reports on a specific event or schedule the reports, you also need to specify the format and the recipients of each report.

The following is a list of report formats:

ASCII A report is generated as plain text.

HTML A report is generated in HTML format. This format is useful for viewing using a Web browser. For example, you can check if your systems have been backed up by clicking a link and viewing the report on the intranet.

18.062
Paula

IMPORTANT

When sending an HTML report on a Windows Cell Manager using the email send method, how the report is displayed will depend on the email client used to open it. Many email clients will display the report as plain ASCII text. To ensure the report displays correctly as HTML, open it in a Web browser.

Short	A report is generated as plain text, but in a short, summary form, showing the most important information. This is the suggested format for broadcast messages.
Tab	A report is generated with fields separated with tabs.

TIP

The Tab format is useful to import the reports into some other applications or scripts for further analysis, such as Microsoft Excel.

The following command creates a list of media used in the last 24 hours in a Microsoft Excel spreadsheet:

```
omnirpt -report used_media -timeframe 24 24 -log
used_media.xls -tab
```

18.061
Paula

Report Send Methods

Report Send Methods

Reports can be sent using various methods:

- Email send method
- Broadcast message send method
- SNMP send method
- External send method
- Log to file send method

The following sections describe specifics of each method.

Email Send Method

The email send method allows you to send or receive an email with the output of the report.

IMPORTANT

When sending an HTML report on a Windows Cell Manager using the e-mail send method, how the report is displayed will depend on the email client used to open it. Many email clients will display the report as plain ASCII text. To ensure the report displays correctly as HTML, open it in a Web browser.

To send e-mail reports from a Windows system with Microsoft Exchange, create a Data Protector Exchange profile called `OmniBack` on this system (usually the Data Protector Cell Manager).

Creating a New Microsoft Exchange Profile

To create a Data Protector user profile on the system that will be sending the email messages, follow the steps below:

1. In the Windows `Control Panel`, click the `Mail and Fax` icon.
2. In the `Properties` dialog box, click `Show Profiles`. The `Mail and Fax` dialog box appears.
3. Click `Add` to start the `Microsoft Exchange Setup` wizard.
4. Select `Use the following information services`.

18.060
Paula

5. In the information services list, select `Microsoft Exchange Server`.
6. Click `Next` to display the `Profiles` page.
7. In the `Profile Name` text box, enter `OmniBack`, and then click `Next`. The `Microsoft Exchange Server` page appears.
8. In the `Microsoft Exchange Server` text box, enter the name of the server.
9. In the `Mailbox` text box, enter the name that you want to appear in email messages. This is usually Data Protector or the administrator's name.
10. The remaining information is optional. Follow the on-screen instructions, and then click `Finish` to complete the wizard. You will then have a new MS Exchange profile for Data Protector.

Broadcast Message Send Method

The broadcast message send method allows you to send a broadcast message with the output of the report to specified systems.

Broadcast messages can be sent to Windows systems only, by specifying the system to which the broadcast message should be sent. Broadcast messages are limited in length, so the `short` format is preferred. The reports are limited to 1000 characters.

Log to File Send Method

The log to file send method allows you to post a log file with the output of the report to a specified file.

The log file is posted to the Cell Manager system. Specify the name of the file to which you want to post the report. The file will be overwritten if it exists.

SNMP Send Method

SNMP send method allows you to send an SNMP trap with the output of the report. The SNMP trap can be further processed by applications using SNMP traps.

18059
Paula

NOTE

On a UNIX Cell Manager, SNMP traps are sent to the systems configured in the notification.

On a Windows Cell Manager, SNMP traps are sent to the systems configured in the Windows SNMP traps configuration.

Windows NT

To configure Windows NT SNMP traps, proceed as follows:

1. On the Cell Manager, click `Control Panel`, `Network`, and then `Services`. If there is no SNMP Service entry in the list of network services under the Services tab, perform step 2. If there is, skip step 2.
2. Click `Add` and select `SNMP Service`. Click `OK`. Insert the Windows NT installation CD, or provide an alternative path to the requested files. Click `Continue`.
3. Select `SNMP Service` and click `Properties`. Select the `Traps` tab and enter `public` in the `Community Name` drop-down list. Click `Add` next to the `Community Name` drop-down list.
4. Click `Add` under `Trap Destinations` text box and enter the hostname of the VPO Management Server. Click `Add`. Repeat this step to add any number of VPO Management Servers.
5. Click `OK`. Click `Close`.
6. Start `omnisnmp`.

Windows 2000/XP/Server 2003

To configure Windows 2000/XP/Server 2003 SNMP traps, proceed as follows:

1. On the Cell Manager, click `Settings`, and then `Network and Dial-up Connections`.
2. In the `Advanced` menu, select `Optional Networking Components` to start the wizard.
3. In the wizard, select `Management and Monitoring` tools and click `Next`.
4. Follow the wizard to install the `Management` and `Monitoring` tools.
5. Open `Control Panel`, `Administrative Tools` and then `Services`.

18.058
Paula

6. Right-click SNMP Service and select Properties.

   a. Select the Traps tab and enter public in the Community name text box and the hostname of the VPO Management Server in the Trap Destinations text box.

   b. Select the Security tab. Under Accepted community names, select the community public, click Edit and set Community rights to READ CREATE.

   c. Confirm your settings.

7. Start omnisnmp.

External Send Method

The external send method allows you to process the output of the report in your own script. The script receives the output as standard input (STDIN). The recommended format for script processing is the tab format.

The script, which is located on the Cell Manager system, must reside in the /opt/omni/lbin (UNIX systems) or *<Data_Protector_home>*\bin (Windows systems) directory. You need to provide only the name of the script, not the entire path.

TIP You can use this delivery method to perform a scheduled eject of the specified media. Refer to “Scheduled Eject of Media” on page 138.

18.057
Paula

Configuring Reports Using the Data Protector GUI

This section describes how to configure Data Protector reports using the Data Protector GUI.

NOTE

To display the input parameters (selections) in the output of a report, select the `Show selection criteria in report` option in the `Report Wizard`. The `Show selection criteria in report` is not available for the reports that have no required or optional input parameters (selections). The output of the report displays only required parameters and optional parameters with the changed default values.

Configuring Report Groups and Adding Reports

Report Groups

You can start Data Protector reports individually (interactively) or you can group them into report groups and then start the report group. You can add individual reports to an already configured report group.

Using the Data Protector GUI, a report group allows you to:

- Start all the reports at once (interactively).
- Schedule the group to start the reports at a specified time.
- Start the group when triggered by a notification.

Examples

These are some examples of the use of reports:

- A backup operator wants to receive an email with the status of the backup performed on the previous night.
- Administrators of specific departments want to receive a broadcast message with information on the backup of the systems they are responsible for.
- A full report with tab delimited data is posted as a log file and is used by an application that records backup statistics.

18.056
Paula

Administrators can configure a report group and add a separate report for each of the requirements. They can schedule the report group to be executed early enough in the morning, so that all recipients receive the reports before coming to work.

NOTE The Mount Request Report and Device Error Report can only be used in a report group and are not available as interactive reports.

To configure a report group, do the following:

1. In the `Data Protector Manager`, switch to the `Reporting` context.
2. Click the `Objects` tab below the Scoping Pane to switch to the Objects view.
3. Right-click `Reports` and then select `Add Report Group`. The `Add Report Group` wizard appears.

   Follow the wizard. You will go through the following steps:

   a. Name the report group.

   b. Optionally schedule when the group should be started. For more information on how to use the Scheduler, see "Scheduling Unattended Backups" on page 207.

   c. Choose and configure a report for the group. For each report, you must configure a format used to deliver a report, recipients for each report, and a send method. See "Report Formats" on page 329 for more information on report formats. See "Data Protector Notifications" on page 342 for more information on various send methods.

IMPORTANT When sending an HTML report on a Windows Cell Manager using the email send method, how the report is displayed will depend on the email client used to open it. Many email clients will display the report as plain ASCII text. To ensure the report displays correctly as HTML, open it in a Web browser.

18.055
Paula

NOTE To trigger a report group by a notification, you first need to configure a report group and then configure the notification to use the `Use Report Group` send method.

4. The report group is created and displayed in the Scoping Pane.
5. To add multiple reports to the group, right-click the group and then select `Add Report`.

18.054
Paula

Running Reports and Report Groups Using the Data Protector GUI

Data Protector reports can be run individually, or they can be grouped into report groups and then run.

Running Individual Reports

To run each report individually, do the following:

1. In the `Data Protector Manager`, switch to the `Reporting` context.
2. Click the `Tasks` tab below the Scoping Pane to switch to the tasks context. Browse the provided reports and select the one that you want.
3. Follow the `Report Wizard` to configure and run the report.

Running Report Groups

To run a configured report group, do the following:

1. In the `Data Protector Manager`, switch to the `Reporting` context.
2. In the Scoping Pane, browse for and right-click the report group you want to run and then click `Start`.
3. Click `Yes` to confirm.

18.053
Paula

Running Reports and Report Groups Using the Command-Line Interface

Data Protector reports can be generated using the command-line interface. The command-line interface allows you to include Data Protector reports in some other configuration scripts you are using. It allows you to generate individual reports, run report groups, and define report formats and send methods.

The omnirpt command is used to generate reports. For a detailed description of the command, see the omnirpt man page.

Here are some examples of omnirpt usage:

```
omnirpt -rptgroup <ReportGroup>
```

Runs the report group named *<ReportGroup>*.

NOTE You first need to configure a report group using the Data Protector GUI or Web reporting interface before running it using the Data Protector command-line interface.

```
omnirpt -report host -host <Hostname> -html
```

This generates a Client Backup Report for system *<System_Name>* in the HTML format.

IMPORTANT When sending an HTML report on a Windows Cell Manager using the email send method, how the report is displayed will depend on the email client used to open it. Many email clients will display the report as plain ASCII text. To ensure the report displays correctly as HTML, open it in a web browser.

Example 1 The following command creates a Session Flow Report for the last 24 hours and logs it to the file in HTML format, as shown in Figure 7-1 on page 340:

18.052
Paula

```
omnirpt -report session_flow -timeframe 24 24 -log
session_flow.html -html
```

Figure 7-1 **A Session Flow Report**

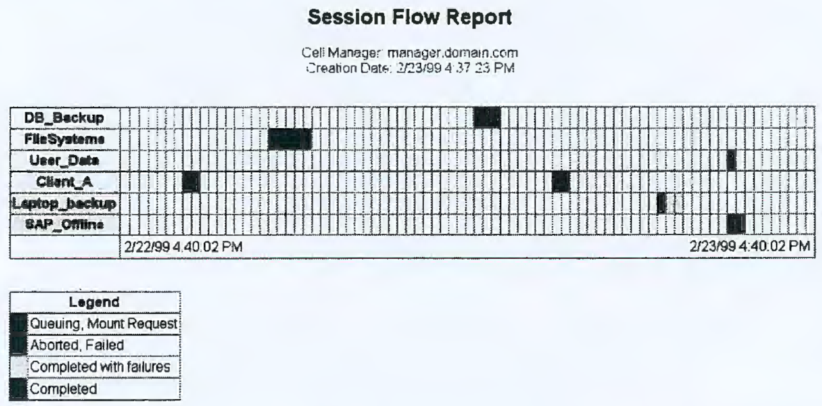

Session Flow Report

Cell Manager: manager.domain.com
Creation Date: 2/23/99 4:37:23 PM

Legend
- ■ Queuing, Mount Request
- ■ Aborted, Failed
- □ Completed with failures
- ■ Completed

Example 2 The following command creates a Media Statistics Report on media in poor condition and logs it in the file in the ASCII format, as shown in Figure 7-2 on page 340:

```
omnirpt -report media_statistics -status poor -log
media_statistics.txt -ascii
```

Figure 7-2 **A Media Statistics Report**

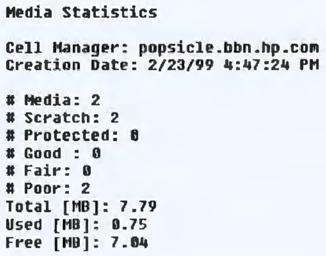

```
Media Statistics

Cell Manager: popsicle.bbn.hp.com
Creation Date: 2/23/99 4:47:24 PM

# Media: 2
# Scratch: 2
# Protected: 0
# Good : 0
# Fair: 0
# Poor: 2
Total [MB]: 7.79
Used [MB]: 0.75
Free [MB]: 7.04
```

Example 3 The following command creates a Device Flow Report for the last 24 hours and sends it via email in HTML format, as shown in Figure 7-3 on page 341:

18.05!
Paula

```
omnirpt -report device_flow -timeframe 24 24 -email
ulmo@outersea.ea -html
```

Figure 7-3 **A Device Flow Report**

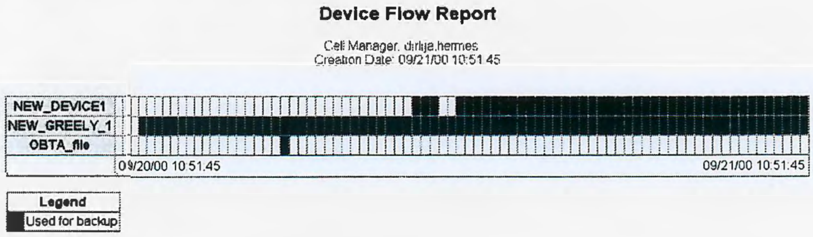

Device Flow Report

Cell Manager. dirlija.hermes
Creation Date: 09/21/00 10:51 45

18.050
Paula

Data Protector Notifications

What Are Notifications?

The Data Protector notification functionality allows you to receive notifications when specific events occur. For example, when a backup session is completed, you can receive an email with a status of the session.

You can set up a notification so that it triggers a report. For more information about Data Protector reports, refer to "Data Protector Reporting" on page 315.

NOTE

Only the Data Protector users in the Admin group and those granted the `Reporting, notifications and event log` user rights are given access to Data Protector notification functionality.

Configuring Notifications

Notifications can be configured using the Data Protector user interface or any Web browser with Java support.

Notifications can be customized by configuring input parameters.

All notifications have the following common input parameters:

- Name (a name for the notification)
- Message Level (the default value depends on the notification and is listed for each notification in the table below)
- Send Method (the default value is Data Protector Event Log)

Notification Types

There are two main types of notifications:

- Notifications that are triggered when an event occurs:
  - ✓ Alarm
  - ✓ Backup Error
  - ✓ Database Corrupted
  - ✓ Device Error

18049
Paula

- ✓ End of Session
- ✓ Mail Slots Full
- ✓ Mount Request

- Notifications that are scheduled and started by the Data Protector checking and maintenance mechanism:
  - ✓ Database Purge Needed
  - ✓ Database Space Low
  - ✓ Health Check Failed
  - ✓ License Will Expire
  - ✓ Not Enough Free Media
  - ✓ Unexpected Events
  - ✓ User Check Failed

For more information on the Data Protector checking and maintenance mechanism, refer to "Data Protector Checking and Maintenance Mechanism" on page 605.

Table 7-8 **Data Protector Notifications**

Name	Optional Input Parameters	Default Message Level and Optional Input Parameter Default Values	Message Displayed
Database Corrupted	none	• Critical	Corruption in the *<DB_part>* part of the internal database has been detected *<error_message>*
Backup Error	Single Message Level (<Any>/Warning/Minor/Major/Critical/ Normal - only the Data Protector messages of the specified level of messages and above trigger this notification)	• Major • Major	Backup session *<session_ID>* of the backup specification *<backup_spec>* has errors: *<number_of_errors>*

18.048
Paula

Table 7-8 **Data Protector Notifications**

Name	**Optional Input Parameters**	**Default Message Level and Optional Input Parameter Default Values**	**Message Displayed**
Unexpected Events	Number of Events (threshold value for the number of events in the Data Protector Event Log that triggers this notification)	• Warning • 20	Data Protector Event log increased for `<Number of Events>` unexpected events in last day
Health Check Failed	none	Critical	Health check message: `<healthcheck_command>` failed, check HealthCheck.log file.
User Check Failed	Command Path	• Major • none	User check failed with exit code `<error_code>:<error_description>`
End of Session	• Datalist • Session Status	• Warning • All • Completed with Errors	Session `<session_ID>` of backup specifications `<Datalist>` completed with overall status `<Session Status>`
Device Error	Device	• Critical • `<Any>`	Error on device `<Device>` occurred
Database Space Low	• Maximum Size of filenames.dat [MB] • Disk Free for Internal Database [MB] • DCBF Size Limit [MB]	• Major • 250 MB • 50 MB • 250 MB	Internal database is running out of space

18.047
Paula

Table 7-8 **Data Protector Notifications**

Name	Optional Input Parameters	Default Message Level and Optional Input Parameter Default Values	Message Displayed
Database Purge Needed	• Days Last Purge [days]) • Num. Estimated Filenames [mio] • Estimated Time Purge [min] • Num. Filenames [mio]	• Warning • 180 days • 6 million • 120 minutes • 100 million	Filename purge should be run for Internal Database
Mount request	Device	• Warning • *<Any>*	Mount request on device *<Device>*
Not Enough Free Media	• Media Pool • Number of Free Media (threshold value for the lowest number of free media that triggers this notification)	• Warning • *<Any>* • 2	Media pool *<Media Pool>* contains only *<number_of_media>* free media
Mail Slots Full	• Device	• Warning • *<Any>*	All mail slots of library *<Device>* are full. Please remove them immediately
License Will Expire	License expires in days	• Warning • 10	The first license will expire in *<License expires in days>* days
Alarm	none	• Warning	Alarm: *<Alarm_message>*

18 046
Paula

Explanation of Some Notifications

Alarm The Alarm notification is used to display critical Data Protector messages triggered by Data Protector internal conditions.

Database Purge Needed By default, once per day Data Protector will check the Database Purge Needed condition as a part of Data Protector checking and maintenance mechanism and trigger the notification if:

- For any Data Protector client in the cell, the number of days since the last IDB filename purge is larger than the <*Days Last Purge [days]*> input parameter and at least one of the following two conditions is true:
  - ✓ The number of filename records likely to be purged is larger than the <*Num. Estimated Filenames [mio]*> input parameter.
  - ✓ It is estimated that more than <*Estimated Time Purge [sec]*> seconds will be needed to finish the purge.
- The number of filenames in the IDB is larger than the <*Num. Filenames [mio]*> input parameter.

For more information on the Data Protector checking and maintenance mechanism, refer to “Data Protector Checking and Maintenance Mechanism” on page 605.

Database Space Low By default, once per day Data Protector will check the Database Space Low condition, and will trigger notification if the allocated space for CDB extension files is running low, if any of the disks containing the IDB are running out of space, or if the allocated space for all DC directories is running low. In other words, the notification will be triggered if any of the following is true:

- The difference between the maximum size of *all* CDB extension files, (the sum of all CDB extension files maximum sizes) and the current size of all CDB extension files drops below the <*Maximum Size of filenames.dat [MB]*> input parameter.
- The free disk space on *any* of the disks containing the IDB drops below the <*Disk Free for Internal Database [Mb]*> input parameter.
- The difference between the maximum size of *all* DC directories and the current size of all DC directories drops below the <*DCBF Size Limit [MB]*> input parameter.

18045
Paula

For more information on the Data Protector checking and maintenance mechanism, refer to "Data Protector Checking and Maintenance Mechanism" on page 605.

Health Check Failed As a part of the Data Protector checking and maintenance mechanism, Data Protector will by default once per day start the Health Check, which starts the `omnihealthcheck` command and triggers the notification if the `omnihealthcheck` command fails. For more information on the `omnihealthcheck` command, refer to the `omnihealthcheck` man page. The `omnihealthcheck` command checks:

- whether the Data Protector services (`rds`, `crs`, `mmd`, `omnitrig`, and `OmniInet`) are active
- whether the Media Management database is consistent
- whether at least one backup of the IDB exists

The exit code of the command is 0 (OK) only if all three checks completed successfully (exit code for every check was 0). Exit values other than 0 indicate that one of the checks failed. For more information on exit codes, refer to the `omnihealthcheck` man page.

User Check Failed By default, once per day Data Protector will start the User Check, which executes the script/command specified as the *`<script/command pathname>`* input parameter. Create the command/script in the `/opt/omni/lbin` (HP-UX and Solaris) or *`<Data_Protector_home>`*`\bin` (Windows) directory of the application system. Enter the filename here. The notification is triggered if the script/command exits with the return value other than 0.

For more information on the User Check Failed notification, refer to "The User Check Failed Notification" on page 606.

End of Session The End of Session notification is triggered when a Data Protector session for the backup specification(s) specified by the *`<Datalist>`* input parameter ends with the status specified by the *`<Session Status>`* input parameter. The default value is Completed with Errors.

Notification Send Methods

Notifications can be sent using various methods:

- Email send method

13044
Paula

- Broadcast message send method
- SNMP send method
- External send method
- Log to file send method
- Use Report Group send method
- Data Protector Event Log send method

NOTE By default, all notifications are configured to be sent to the Data Protector Event Log. In order to send an additional notification using some other send method, an additional notification has to be configured.

Email Send Method

Email notifications allow you to receive an email with desired information when a specified event occurs.

TIP To send email notifications from a Windows system with Microsoft Exchange, create a Data Protector Exchange profile called `OmniBack` on the Data Protector Cell Manager. On UNIX systems, no additional configuration is needed.

Creating a New Microsoft Exchange Profile

To create a Data Protector user profile on the system that will be sending the email messages, follow the steps below:

1. In the Windows `Control Panel`, click the `Mail and Fax` icon.
2. In the `Properties` dialog box, click `Show Profiles`. The `Mail and Fax` dialog box appears.
3. Click `Add` to start the `Microsoft Exchange Setup` wizard.
4. Select `Use the following information services`.
5. In the information services list, select `Microsoft Exchange Server`.
6. Click `Next` to display the `Profiles` page.

18043
Paula

7. In the `Profile Name` text box, enter `OmniBack`, and then click `Next`. The `Microsoft Exchange Server` page appears.
8. In the `Microsoft Exchange Server` text box, enter the name of the server.
9. In the `Mailbox` text box, enter the name that you want to appear in email messages. This is usually Data Protector or the administrator's name.
10. The remaining information is optional. Follow the on-screen instructions, and then click `Finish` to complete the wizard. You will then have a new MS Exchange profile for Data Protector.

Broadcast Message Send Method

Broadcast message notifications allow you to send a broadcast message to systems when a specified event occurs.

Broadcast messages can be sent to Windows systems only, by specifying the system to which the broadcast message should be sent. Broadcast messages are limited in length, so the `short` format is preferred. The reports are limited to 1000 characters.

Log to File Send Method

Log to file notifications allow you to post a log file with desired information when a specified event occurs.

The log file is posted to the Cell Manager system. Specify the name of the file to which you want to post the report.

SNMP Send Method

SNMP traps notifications allow you to send an SNMP trap with desired information when a specified event occurs. The SNMP trap can be further processed by applications using SNMP traps.

NOTE On a UNIX Cell Manager, SNMP traps are sent to the systems configured in the notification.

On a Windows Cell Manager, SNMP traps are sent to the systems configured in the Windows SNMP traps configuration.

18.042
Paula

Windows NT

To configure Windows NT SNMP traps, proceed as follows:

1. On the Cell Manager, open `Control Panel`, `Network`, `Services`. If there is no SNMP Service entry in the list of network services under the Services tab, perform step 2. If there is, skip step 2.
2. Click `Add` and select `SNMP Service`. Click `OK`. Insert the Windows NT installation CD, or provide an alternative path to the requested files. Click `Continue`.
3. Select `SNMP Service` and click `Properties`. Select the `Traps` tab and enter `public` in the `Community Name` drop-down list. Click `Add` next to the `Community Name` drop-down list.
4. Click `Add` under `Trap Destinations` text box and enter the hostname of the VPO Management Server. Click `Add`. Repeat this step to add any number of VPO Management Servers.
5. Click `OK`. Click `Close`.
6. Start `omnisnmp`.

Windows 2000/XP/Server 2003

To configure Windows 2000/XP/Server 2003 SNMP traps, proceed as follows:

1. On the Cell Manager, open `Settings`, `Network and Dial-up Connections`.
2. In the `Advanced` menu, select `Optional Networking Components` to start the wizard.
3. In the wizard, select `Management and Monitoring` tools and click `Next`.
4. Follow the wizard to install the `Management` and `Monitoring` tools.
5. Open `Control Panel`, `Administrative Tools`, `Services`.
6. Right-click `SNMP Service` and select `Properties`.
   a. Select the `Traps` tab and enter `public` in the `Community name` text box and the hostname of the VPO Management Server in the `Trap Destinations` text box.
   b. Select the `Security` tab. Under `Accepted community names`, select the community `public`, click `Edit` and set `Community rights` to `READ CREATE`.
   c. Confirm your settings.

18041
Paula

7. Start omnisnmp.

External Send Method

External script notification allows you to process the output of the report in your own script. The script receives the output as standard input (STDIN). The recommended format for script processing is the tab format.

The script, which is located on the Cell Manager, must reside in the /opt/omni/lbin (HP-UX and Solaris systems) or *<Data_Protector_home>*\bin (Windows systems) directories. You need to provide only the name of the script, not the whole path.

TIP You can use this delivery method to perform a scheduled eject of the specified media. Refer to "Scheduled Eject of Media" on page 138.

Use Report Group Send Method

Report group notification allows you to start a report group when a specified event occurs. See "Configuring Reports Using the Data Protector GUI" on page 335 for more information on report groups.

Data Protector Event Log Send Method

By default, all notifications are sent to the Data Protector Event Log. The Data Protector Event Log is accessible only for Data Protector users in the Admin group and to Data Protector users that are granted the Reporting, notifications and event log user rights. You can view or delete all events in the Data Protector Event Log. Refer to "Data Protector Event Log" on page 356.

Configuring Notifications

To configure a notification, do the following:

1. In the Data Protector Manager, switch to the Reporting context.
2. Click the Objects tab below the Scoping Pane to switch to the Objects view.
3. Right-click Notifications and then select Add Notification. The Add Notification wizard appears. Follow the wizard.

18.040
Paula

TIP

To trigger a report group by a notification, configure a report group and then configure the notification to use the Use Report Group send method.

4. The notification is created and displayed in the Scoping Pane.

13034
Paula

Configuring Reports and Notifications on the Web

You can use your Web browser to view Data Protector reports and notifications.

Using the web reporting and notifications interface, you can view, configure, and start Data Protector reports and notifications from any system on your network. You can configure reports and notifications that are delivered using various reporting methods and formats.

All reporting and notifications functionality accessible using the Data Protector GUI is also accessible using Data Protector web reporting and notifications. See below for the limitations.

When you install the Data Protector Cell Manager, the web reporting user (called Java) is automatically created. By default, no password is needed to use the Data Protector web reporting and notifications. By configuring a Web user password you restrict the access to the Data Protector web reporting and notifications functionality.

Limitations

The following is a list of Data Protector web reporting and notifications interface limitations:

- You cannot edit, view, or delete the saved reports using the web reporting and notifications interface.
- You cannot start a report group using the web reporting and notifications interface.
- Whenever multiple input parameters (selections) are to be *typed* in the web reporting and notifications interface, every parameter (selection) has to be enclosed in double quotes if it contains spaces.

To use Data Protector web reporting and notifications, do the following:

1. Have a system with a configured and running web server. Data Protector works with all popular web servers.
2. Copy Data Protector Java programs to the web server. The system does not have to be a Data Protector client. The steps are described below.

18038 Paula

3. Optionally, configure a password to limit access to Web reports. The steps are described below.

Copying Data Protector Java Programs to the Web Server

To allow access to Data Protector Web reporting and notifications interface from a browser from any system, copy Data Protector Java reporting programs to the web server.

From the system with the Data Protector user interface installed, copy the following directory with all subdirectories:

- On Windows: `<Data_Protector_home>\java`
- On UNIX:`/opt/omni/java`

Access the `\bin\WebReporting.html` (Windows systems) or the `/bin/webreporting.html` (UNIX systems) file from the copied `java` folder in a browser to display the Data Protector reporting. Make this file available to the users of the web reporting in the full URL form. For example, you can put a link to this file from your intranet site.

Restricting Access to Web Reporting

When you install the Data Protector Cell Manager, the web reporting and notifications user (called `java`) is automatically created. By default, no password is needed to use the Data Protector web reporting and notifications. By configuring a web user password, you restrict the access to Data Protector web reporting and notifications functionality. Any user using web reporting and notifications will have to provide this password to browse the Data Protector reports on the web.

To change the password for the Data Protector web reporting and notifications interface, do the following:

1. In the `Data Protector Manager`, switch to the `Users` context.
2. Choose `Action, Set Web User Password`. A dialog box appears, where you change the password.

   Any user using web reporting and notifications will have to provide this password to browse the Data Protector reports on the web.

18034

Generating the Reports

To generate reports using the Data Protector Web reporting and notifications interface, you have to access this interface. The actual steps depend on your configuration. Once you are logged on the Cell Manager, you can generate various types of reports. See "Data Protector Reporting" on page 315 for more information on report types.

To view a report, click the report and provide the needed information.

When the report is displayed, you can print the report or save it. When you save the report, you can also add this report to an existing or a new report group. See the next section for more information.

Configuring Notifications

To configure notifications using the Data Protector Web reporting and notifications interface, you have to access this interface. The actual steps depend on your configuration. Once you are logged on the Cell Manager, you can configure notifications. See "Data Protector Notifications" on page 342 for more information on notifications.

To configure a notification, select `Notifications` and click `Add Notification`. Provide the needed information and save the notification.

Configuring Report Groups

Report Groups

See "Configuring Report Groups and Adding Reports" on page 335 for more information on report groups.

In the web reporting and notifications interface, you can create a new report group when you save the report:

1. Choose the report you want to generate.
2. Enter the needed information.
3. Once the report is displayed, click `Save`. Enter the report name and a new or an existing report group to which you want to add the report.

Data Protector Event Log

The Data Protector Event Log represents a centralized event management mechanism, dealing with specific events that occurred during the Data Protector operation. The events are logged in the `<Data_Protector_home>\log\Ob2EventLog.txt` (Windows systems) or in the `/var/opt/omni/log/Ob2EventLog.txt` (HP-UX and Solaris systems) file on the Cell Manager. Viewing the Data Protector Event Log using the Data Protector GUI helps you troubleshoot possible problems.

The events are logged by the notifications functionality. Refer to "Data Protector Notifications" on page 342 for more information on notifications.

NOTE Only the Data Protector users in the Admin group and those granted the `Reporting, notifications and event log` user rights are given access to Data Protector Event Log functionality.

Event Log To access the Event Log, select the `Reporting` context in the Data Protector GUI and expand `Reporting`. Select `Event Log` to display events.

NOTE The Data Protector Event Log is not refreshed automatically. If you want to view new messages, refresh it manually by pressing `F5`.

Deleting Event Log Contents Right-click `Event Log` and select `Empty Event Log`. This will delete all entries in the `Event Log`.

NOTE Deleting the Event Log contents will not delete the `<Data_Protector_home>\log\Ob2EventLog.txt` (Windows systems) or the `/var/opt/omni/log/Ob2EventLog.txt` (HP-UX and Solaris systems) file.

18.035
Paula

When the Data Protector graphical user interface is started by a user, if there are new notifications that have not been seen by this user in the Data Protector Event Log, the following message is displayed:

Figure 7-4 **The Event Log Message**

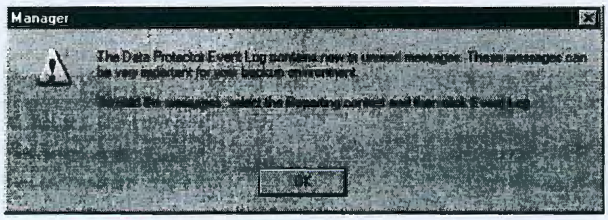

The Data Protector Event Log contains new or unread messages. These messages can be very important for your backup environment.

To read the messages, select the Reporting context and then click Event Log.

18035
Paula

8 Manager-of-Managers Environment

In This Chapter

This chapter shows you how to configure and use the Data Protector Manager-of-Managers, which is used to control an enterprise backup environment. It consists of the following sections:

"Manager-of-Managers" on page 361

"Configuring the Manager-of-Managers" on page 362

"Centralized Media Management Database (CMMDB)" on page 366

"Configuring a Centralized Media Management Database" on page 368

"Centralized Licensing" on page 372

"Working with a MoM Environment" on page 377

"Restoring, Monitoring, and Reporting in an Enterprise Environment" on page 380

NOTE

MoM is subject to specific Data Protector licenses. See the *HP OpenView Storage Data Protector Installation and Licensing Guide* for details.

18031
Paula

Manager-of-Managers

The Data Protector Manager-of-Managers (MoM) allows administrators to centrally manage a large environment consisting of several Data Protector cells, also known as MoM clients, from a single point. Refer to the *HP OpenView Storage Data Protector Concepts Guide* for further details about the enterprise environment.

NOTE Each MoM client and the MoM Manager need to run the same version of Data Protector.

The Data Protector MoM is flexible enough to expand the backup environment as the enterprise grows. It provides the following features:

Centralized management of all tasks

Enables configuration, management, and control over the enterprise environment from a single point. This includes configuring backup, media management, restoring, and monitoring; and reporting about the status of the whole backup environment.

Centralized Media Management Database

Optionally, all the cells in the environment can share a common, central database to manage devices and media within the enterprise. The Centralized Media Management Database (CMMDB) enables you to share high-end devices between cells. This means that any device in a cell using the CMMDB is available to all cells using the CMMDB.

Centralized licensing

Data Protector enables you to configure centralized licensing for the whole MoM environment. All Data Protector licenses are installed and kept on the MoM Manager and can be allocated to specific cells as needed.

18.030
Paula

Configuring the Manager-of-Managers

To configure the MoM environment, you need to do the following:

- Set up the MoM Manager. See “Setting Up MoM Manager” on page 363.
- Import Data Protector cells into MoM environment. See “Importing Data Protector Cells” on page 363.
- Create a Data Protector user in the Admin user group on every cell in the environment (MoM administrator). See “Adding a MoM Administrator” on page 364.
- Restart Data Protector services. See “Restarting Data Protector Services” on page 364.

Optionally, you can also:

- Configure the Centralized Media Management Database. See “Configuring a Centralized Media Management Database” on page 368.
- Configure centralized licensing. See “Centralized Licensing” on page 372.
- Distribute the MoM configuration. See “Distributing the MoM Configuration” on page 378.

Prerequisites

Choose the system you will configure as your MoM Manager. Follow the guidelines below:

- The MoM Manager system should be highly reliable.
- The system has to already be a Data Protector Cell Manager with the software installed. See the *HP OpenView Storage Data Protector Installation and Licensing Guide* for more information on how to configure the Data Protector Cell Manager system.

Install the required licenses on the MoM cell and every prospective MoM client cell.

18 029
Paula

Setting Up MoM Manager

To set up an enterprise environment, configure one of your Cell Managers as a MoM Manager.

1. In the `Data Protector Manager`, click `Clients` in the Context List.
2. In the `Actions` menu, click `Configure CM as Manager-of-Managers Server`.
3. Stop and restart Data Protector services. Refer to the section "Restarting Data Protector Services" on page 364.

TIP On Windows, you can also use the `Control Panel` to stop the Data Protector services. See "Setting the User Account for the Data Protector Inet Service" on page 187 for details.

4. Run the MoM graphical user interface:
   - On Windows: from the `Start` menu select `Manager-of-Managers` in the `HP OpenView Storage Data Protector` program group.
   - On UNIX: run the `/opt/omni/bin/xomnimom` command.

Importing Data Protector Cells

Once you have configured the MoM Manager, you can start adding (importing) the Data Protector cells to the MoM environment. To import a Data Protector cell to the MoM environment, proceed as follows:

1. In the `Data Protector Manager-of-Managers`, click `Clients` in the Context List.
2. Right-click `Enterprise Clients`, and then click `Import Cell Manager`.

IMPORTANT In order to import a Cell Manager into the MoM as an Enterprise Client, you must be a member of the admin user group on that Cell Manager. If you are not, the import will fail.

3. Enter, or browse for, the name of the Cell Manager that you want to import, and then click `Finish`. The selected Cell Manager is now a part of your MoM environment.

NOTE If you are adding a Cell Manager installed on a cluster to the MoM cell, ensure that you enter its **virtual server name**.

Adding a MoM Administrator

A MoM administrator can perform administration tasks in all cells in the enterprise environment.

You need to have a certain user that is in the Admin user group on every Cell Manager in the MoM environment. For example, you may have a user called *MoM_Admin*. This user will be the MoM administrator.

1. Using the Data Protector Manager, connect to each Cell Manager in the MoM environment as an `Admin` user.
2. Add the user that will be the MoM Administrator to the Data Protector `Admin` user group.

On how to add users, see "Adding or Deleting a User" on page 90.

Restarting Data Protector Services

When you have configured the MoM environment, you will be notified to restart the Data Protector services.

If the Windows Service Control Manager is used to start and stop services on the Cell Manager, only the current and previous copies of the database log are kept. Using the `omnisv -stop` and the `omnisv -start` commands will save all previous database logs.

1. Stop all Data Protector services by entering the following command:
   - on Windows: `<`*`Data_Protector_home`*`>\bin\omnisv -stop`
   - on UNIX: `/opt/omni/sbin/omnisv -stop`

MC/ServiceGuard If the Cell Manager is configured on MC/SG, stop the Data Protector package using the following command:

18027
Paula

```
cmhaltpkg <pkg_name>
```

where *<pkg_name>* is the name of the Data Protector cluster package.

2. Restart the Data Protector services by entering the following command:
   - on Windows: `<Data_Protector_home>\bin\omnisv -start`
   - on UNIX: `/opt/omni/sbin/omnisv -start`

MC/ServiceGuard

If the Cell Manager is configured on MC/SG, restart the Data Protector package using the following command:

```
cmrunpkg -n <node_name> <pkg_name>
```

55.026
Paula

Centralized Media Management Database (CMMDB)

The IDB is an embedded database that keeps information about backup, restore, and media management sessions, devices, and media. It consists of five parts that are located on the Cell Manager.

- MMDB - Media Management Database
- CDB - Catalog Database
- DCBF - Detail Catalog Binary Files
- SMBF - Session Messages Binary Files
- SIBF - Serverless Integrations Binary Files

In a typical cell-oriented environment, all parts are located on the Cell Manager system and each keeps information on devices, media, and backup information for that cell. For security reasons, it is impossible to access and use this data from another Data Protector cell. Therefore, media and devices used in that cell cannot be accessed and used in some other cell without moving them to that cell.

In larger multi-cell environments with high-end backup devices, you may want to share these devices and media among several cells. This can be achieved by having one centralized MMDB database for all the cells and keeping an individual CDB for each cell. This allows media and device sharing while preserving the security capabilities of the multi-cell structure.

With the CMMDB, media are owned by the Data Protector cell that performed the first backup on the media. The media owner is displayed in the media view. While media are protected, only backups from that cell can be appended on the media. Therefore, media can only be owned by one cell at a time. Once the protection expires, the media become available to other cells again.

NOTE A backup anywhere in the enterprise environment will not run if the cell running the backup does not have access to the CMMDB. For example, this happens if a network failure occurs between the cell and the MoM cell.

Figure 8-1 The Central Media Management Database

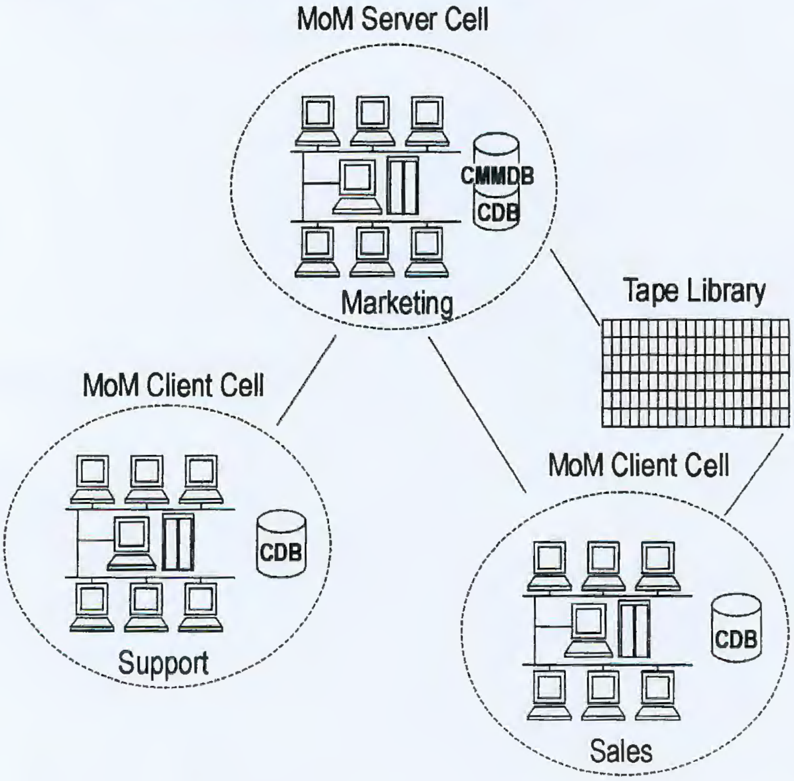

18.024
Paula

Configuring a Centralized Media Management Database

It is not required to set up a Centralized Media Management Database (CMMDB). If you do not set up a CMMDB, Data Protector will work in a multi-cell environment, but each cell will have its own IDB. See "Centralized Media Management Database (CMMDB)" on page 366 for more information on this functionality.

This section describes how to configure a Centralized Media Management Database for the whole multi-cell environment. If it is needed, this process will merge the local Media Management Database into the CMMDB. You can decide for each cell if it will use the CMMDB or its own local MMDB.

IMPORTANT The CMMDB has a major effect on licensing. Immediately after the MMDB is changed from local to remote, all the licenses associated with libraries and devices are taken (validated) from the MoM Manager and can be removed from client cells.

When the CMMDB is used, it does not have to reside on the MoM Manager system. The CMMDB can reside on any Cell Manager in the MoM environment. The Cell Manager on which the CMMDB is located is specified in the file `mmdb_server` in the following directory:

- On Windows: `<Data_Protector_home>\config\cell`
- On UNIX: `/etc/opt/omni/cell`

Each medium with protected data on it has information showing which cell currently owns the data. Once this protection has expired, any cell can reuse the medium. If a tape has been initialized by one cell, any other cell can use it, as long as it does not have any protected data on it. If a tape is loaded in a library and not yet initialized, any cell can initialize it, assuming that there is a `loose` media allocation policy and no other tapes are available.

The media allocation rules apply in exactly the same way to shared tapes, except that appendable media can only be appended by the cell that owns it.

On the MoM, add one cell at a time to the CMMDB.

Prerequisites

- Data Protector Cell Managers in all cells have to have the same version of Data Protector installed and running.
- Check that there are no backup, restore, or media management sessions running on any of the cells to be added to the multi-cell environment.

How to Configure the CMMDB

To configure the CMMDB in the MoM environment, two phases are required:

- Configuration of the CMMDB on the MoM Manager. See "Configuring the CMMDB on the MoM Manager" on page 369.
- Configuration of the CMMDB on the client cell. See "Configuring the CMMDB on the Client Cell" on page 370.

NOTE

Once you have configured the CMMDB and start using it, it is not possible to split it back into local MMDBs. It is not recommended to recover the old state of a MMDB. Instead, you should create a new MMDB from scratch.

Configuring the CMMDB on the MoM Manager

Log on to the MoM Manager and perform the following steps:

1. Copy the following directory to a temporary location for safety reasons:
   - On Windows: `<Data_Protector_home>\db40\datafiles\mmdb`
   - On UNIX: `/var/opt/omni/db40/datafiles/mmdb`
2. Run the following command to merge the local MMDB into the CMMDB:
   - On Windows: `<Data_Protector_home>\bin\omnidbutil -mergemmdb <Cell_Server_Hostname>`
   - On UNIX: `/opt/omni/sbin/omnidbutil -mergemmdb <Cell_Server_Hostname>`

18.022
Paula

TIP If you are configuring a new cell, (and you do not yet have devices and media configured) there is no need to merge the database. You only want to merge cells with the CMMDB that already have devices and media configured.

3. Run the following command to synchronize the local CDB:
   - On Windows: `<Data_Protector_home>\bin\omnidbutil -cdbsync <Cell_Server_Hostname>`
   - On UNIX: `/opt/omni/sbin/omnidbutil -cdbsync <Cell_Server_Hostname>`
4. On the MoM Server, edit the duplicated names of media pools and devices (in the user interface). The duplicated names have a "_N" appended to their name, where N represents a number. This always happens to default pools if they exist on both cells. In this case, manually change the backup specifications that use these devices to use the new device names. It would be a good idea to add a line to the media pool's description to say from which cell the pool has come.

Repeat the steps 2 to 4 for all client cells that you want to add to the CMMDB.

Configuring the CMMDB on the Client Cell

On each MOM client cell, perform the following:

1. Log on to the Cell Manager of the client cell as `Administrator` or `root`.
2. Create the file containing the name of the MMDB Server (fully qualified):
   - On Windows: `<Data_Protector_home>\config\cell\mmdb_server`
     Save the file as Unicode.
   - On UNIX: `/etc/opt/omni/cell/mmdb_server`
3. Stop and restart the Data Protector services. See "Restarting Data Protector Services" on page 364.
4. Update configuration files by running the following command:

18.001

- On Windows: `<Data_Protector_home>\bin\omnicc -update_mom_server`
- On UNIX: `/opt/omni/bin/omnicc -update_mom_server`

13020
Paula

Centralized Licensing

It is not required to set up centralized licensing. Individual licenses can be installed on each Cell Manager. Without centralized licensing, these individual licenses are restricted to the cell on which they are installed, and all licensing administration tasks have to be performed locally.

NOTE

If you have clusters configured in the MoM cell, make sure you identify a cluster client with its virtual hostname.

Why Use Centralized Licensing?

Data Protector allows you to configure centralized licensing for the whole MoM environment. All licenses are installed and kept on the MoM Manager system and can be allocated to specific cells as needed.

Centralized licensing simplifies license management. Licensing administration is performed by the MoM administrator for all cells in the MoM environment. This also includes the distribution and moving of the licenses.

When licenses are installed locally on the Cell Managers, they cannot be moved among the cells without the approval of the *HP Password Delivery Center*. See the *HP OpenView Storage Data Protector Installation and Licensing Guide* for instructions on how to move licenses.

Setting Up Centralized Licensing

Prerequisite

If you are consolidating existing Data Protector cells into a MoM environment, send a request to *HP Password Delivery Center* to move the licenses from the existing Cell Managers to the new MoM Manager.

Configuring Centralized Licensing

1. Log on to the MoM Manager and create the `licdistrib.dat` file:

   On Windows:
   `<Data_Protector_home>\config\cell\licdistrib.dat`

   On UNIX: `/etc/opt/omni/cell/licdistrib.dat`

2. Log on to each client Cell Manager in the MoM environment and create the `lic_server` file with the name of the MoM Manager:

On Windows: *<Data_Protector_home>*`\config\cell\lic_server`

On UNIX: `/etc/opt/omni/cell/lic_server`

3. Stop and restart Data Protector services on each Cell Manager where you made the changes. See "Restarting Data Protector Services" on page 364.
4. In the `Data Protector Manager-of-Managers`, click `Clients` in the Context List.
5. In the Scoping Pane, right-click the Cell Manager that has the licensing information you want to change, and then click `Configure Licensing` to open the wizard. The types and numbers of licenses available to your selected Cell Manager are displayed.

   The **USED** column shows the number of licenses assigned to that particular Cell Manager. Increasing the number in this column will correspondingly decrease the number of available licenses, and vice-versa.

   The **AVAILABLE** column shows the number of licenses available to the entire enterprise. This is the number of licenses not taken by any cell within the enterprise environment.

   The **TOTAL** column shows the total number of licenses both used and available in the entire enterprise.
6. Click the `Remote` option to change the licensing from local to remote. Note that **USED** column is changed into **ALLOCATED**.
7. Modify the license configuration. Note that only **ALLOCATED** column is available during the modification process.

Releasing Licenses

To release (give up) a license type, thus increasing the number available, reduce its corresponding number in the **ALLOCATED** column.

Assigning Licenses

To assign a license type, increase its corresponding number in the **ALLOCATED** column by double-clicking it.

8. Click `Finish` to apply the configuration.
9. Repeat the steps for all Cell Managers for which you want to set up the centralized licensing.

18.015
Paula

NOTE

Data Protector checks the license configuration with the MoM Manager every hour. The licensing status is kept for 72 hours. In case of a communication problem, after this 72 hour period, local licenses are used.

Moving Licenses in the MoM Environment

If you have not configured centralized licensing, you cannot move licenses between cells without the approval of the *HP Password Delivery Center*. This is, however, possible in the MoM environment with configured centralized licensing, where the MoM administrator allocates licenses as needed.

In the example below, assume that the clients from one cell were moved to another. This resulted in the need to move the licenses.

Enterprise Environment Before the Reorganization

Assume that two Cell Managers, Aztec and Mayan, are configured in the enterprise environment with centralized licensing. Aztec is an HP-UX Cell Manager with a Cell Manager for UNIX - Single Drive license. There is also an NDMP server connected in the cell that requires an NDMP Server Backup Extension license. Mayan is also an HP-UX Cell Manager with one Cell Manager for UNIX - Single Drive license.

Reorganization of the Enterprise Environment

The Aztec cell needs to be reorganized, with most of the clients and the NDMP server being transferred to the Mayan cell. Mayan now needs the NDMP Server Backup Extension license. Follow the procedure described below to move the license:

1. In the `Data Protector Manager-of-Managers`, click `Clients` in the Context List.
2. Right-click the Aztec Cell Manager and then click `Configure Licensing`. The types and numbers of licenses available to the Aztec Cell Manager are displayed. Remove the NDMP Server Backup Extension license.
3. Click `Finish` to apply the configuration.
4. Right-click the Mayan Cell Manager and then click `Configure Licensing`. Add the NDMP Server Backup Extension license.
5. Click `Finish` to apply the configuration.

Enterprise Environment After the Reorganization

The Aztec Cell Manager now has one Cell Manager for UNIX - Single Drive license and the Mayan Cell Manager has a Cell Manager for UNIX - Single Drive license and an NDMP Server Backup Extension license for the NDMP server.

For more information on Data Protector licensing policies, see the *HP OpenView Storage Data Protector Installation and Licensing Guide*.

18.016
Paula

Deactivating Centralized Licensing

Centralized licensing can be deactivated and changed back to local licensing.

Deactivation Procedure

1. In the Data Protector Manager-of-Managers, click Clients in the Context List.
2. In the Scoping Pane, right-click the Cell Manager for which you want to deactivate centralized licensing, then click Configure Licensing to open the wizard. The types and numbers of licenses available to your selected Cell Manager are displayed.
3. Click the Local option to change licensing from remote to local.
4. Click Finish to apply the configuration.
5. Repeat the steps for all Cell Managers for which you want to deactivate centralized licensing.
6. Log on to the MoM Manager and mount the following directory:

   On Windows systems *<Data_Protector_home>*\config\cell

   On UNIX systems /etc/opt/omni/cell
7. Rename the licdistrib.dat file, for example, to licdistrib.old.

The changes will take effect after you stop and restart Data Protector services on the MoM Manager and each Cell Manager where you made the changes. See “Restarting Data Protector Services” on page 364.

18.05
Paula

Working with a MoM Environment

The Manager-of-Managers interface enables you to configure, manage, and control an enterprise backup environment from a single point.

In the MoM user interface, you can import and export cells, move clients among cells, and distribute the MoM configuration to other cells in the environment.

Other tasks are performed on the MoM Manager in the same way as if you were a local administrator. Follow the standard procedure to configure backup and restore, manage devices and media for a specific cell, configure Data Protector users and user groups, add clients, monitor running sessions and the status of the backup environment, and configure reporting and notifications.

Importing and Exporting Data Protector Cells

Importing a cell into a MoM environment allows it to be centrally managed using the MoM Manager. Exporting a cell will remove it from the enterprise environment.

NOTE

Cluster clients identify themselves to the MoM Manager with their virtual server names. If you import or export a cluster in a MoM environment, use only its virtual server name.

Importing a Cell Manager

1. In the `Data Protector Manager-of-Managers`, click `Clients` in the Context List.
2. Right-click `Enterprise Clients`, and then click `Import Cell Manager`.
3. Select a Cell Manager you want to import and click `Finish`.

Exporting a Cell Manager

1. In the `Data Protector Manager-of-Managers`, click `Clients` in the Context List.
2. In the `Scoping Pane`, right-click the Cell Manager you want to export, and then click `Export Cell Manager`.
3. Confirm your choice.

Moving Client Systems Among Cells

Data Protector allows you to move systems among cells. During the process, Data Protector:

- Checks whether the system to be moved is configured in any backup specification and leads you through the steps to reconfigure the backup of this system in the new cell.
- Checks whether there are any devices configured on the system and leads you through the steps to move devices to another system.
- Checks whether there are media used in the devices on this system and leads you through the steps to move media.

Moving Clients

1. In the `Data Protector Manager-of-Managers`, click `Clients` in the Context List.
2. Expand the Cell Manager that has the system that you want to move to another cell.
3. Right-click the client system and then click `Move Client System to Other Cell` to open the wizard.
4. Select the target Cell Manager and click `Finish` to move the client.

Distributing the MoM Configuration

Data Protector allows you to create a common user class specification, holidays file settings, global options file settings, and vaulting on all Cell Managers in a MoM environment.

How to Distribute the MoM Configuration

To distribute the `MoM` configuration, follow these steps:

1. In the `Data Protector Manager-of-Managers` click `Clients` in the Context List, right-click `Enterprise Clients`, and then click `Distribute Configuration`.
2. In the `Distribute Configuration` dialog box, select the type of configuration and the Cell Managers to which you want to distribute the selected configuration.
3. Click `Finish` to distribute the configuration.

18.013
Paula

Configuring Data Protector Users

You can add users or user groups to a MoM environment as you would for a single Cell Manager. This procedure updates all Cell Managers with the new users. See Chapter 3, "Configuring Users and User Groups," on page 81 for more information about users and user groups.

To configure Data Protector users or user groups, follow these steps:

1. In the `Data Protector Manager-of-Managers`, click `Users` in the Context List.
2. Select a Cell Manager to which you want to add users.
3. In the `Edit` menu, click `Add` and select `Users` if you want to add a new user, or `User Group` if you want to add a new user group.
4. Enter the required information and click `Finish`.

Managing Devices and Media for a Specific Cell

You can configure devices and media for specific devices and media anywhere within your enterprise environment. To do so, follow these steps:

1. In the `Data Protector Manager-of-Managers`, click `Clients` in the Context List.
2. Select the cell that has the devices or media that you want to manage.
3. In the `Tools` menu, click `Device & Media Administration`. In the `Device and Media` context, configure devices and media as if you were a local administrator.

Restoring, Monitoring, and Reporting in an Enterprise Environment

Restoring data in an enterprise environment is the same as restoring data in a single cell environment.

Select data from the appropriate source and restore as described in Chapter 6, “Restore,” on page 267.

Data Protector allows you to monitor currently running or previously run sessions for any cell in the enterprise environment. When you use Web Reporting, you can also get reports on the entire enterprise environment using the MULTICELL item in the Scoping Pane.

See Chapter 7, “Monitoring, Reporting, Notifications, and the Event Log,” on page 307 for more information on how to use these features in an enterprise environment.

ANEXO UNIDADE DE BACKUP ROBOTIZADO
PARTE D/5

18/01/
Paula

COBRA Tecnologia S.A.
Estrada dos Bandeirantes 7966
CEP 22783-110 Rio de Janeiro RJ
Tel. 21 2442-8800
www.cobra.com.br

18010
Paula

9 Managing the Data Protector Internal Database

Paula

In This Chapter

This chapter provides information about the Data Protector internal database (IDB) and tasks related to managing the database. It is organized as follows:

"About the Data Protector Internal Database" on page 383

"The IDB Architecture" on page 384

"Configuring the IDB" on page 388

"Maintaining the IDB" on page 402

"Restoring the IDB" on page 414

"Recovering the IDB" on page 417

About the Data Protector Internal Database

What Is the Data Protector Internal Database (IDB)?

The Data Protector internal database (IDB) is an embedded database, located on the Cell Manager, which keeps information regarding what data is backed up; on which media it resides; the result of backup, restore, and media management sessions; and what devices and libraries are configured.

Why Is the IDB Used?

There are three key reasons for using the IDB:

- Fast and convenient restore

  The information stored in the IDB enables you to browse the files and directories to be restored. You can quickly find the media required for a restore and therefore make the restore much faster.

- Backup management

  The information stored in the IDB enables you to verify the result of backup sessions.

- Media management

  The information stored in the IDB enables you to allocate media during backup, track media management operations and media attributes, group media in different media pools, and track media location in tape libraries.

How to Manage the IDB

One of the important steps in setting up your Data Protector backup environment is to configure the IDB. Once the IDB is configured as described in "Configuring the IDB" on page 388, you will be notified if you need to perform any of the IDB maintenance tasks.

The IDB maintenance tasks, and the cases when they need to be performed, are described in "Maintaining the IDB" on page 402.

If you receive error messages, refer to "Troubleshooting the IDB" on page 592 and "Recovering the IDB" on page 417.

For information on IDB limitations, refer to the *HP OpenView Storage Data Protector Software Release Notes*.

18.004
Paula

The IDB Architecture

The IDB consists of the following parts:

- MMDB (Media Management Database)
- CDB (Catalog Database)
- DCBF (Detail Catalog Binary Files)
- SMBF (Session Messages Binary Files)
- SIBF (Serverless Integrations Binary Files).

18.006

Figure 9-1 **IDB Architecture**

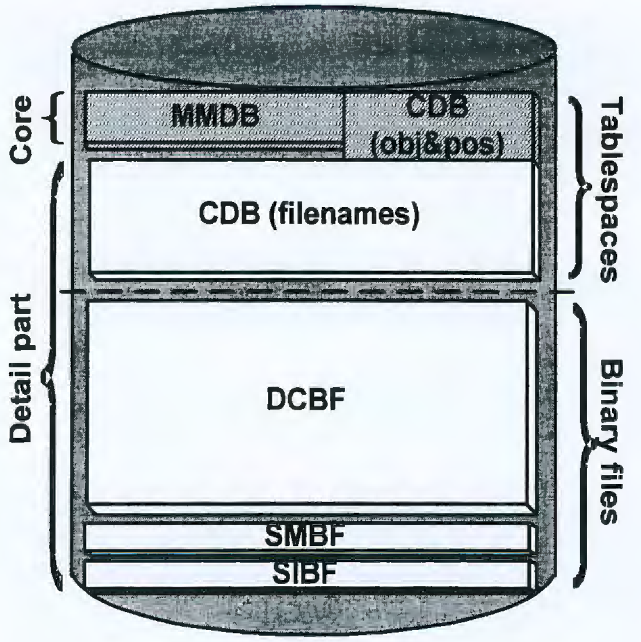

Each of the IDB parts stores specific Data Protector information (records), influences the IDB size and growth in different ways, and is located in a separate directory on the Cell Manager.

MMDB The Media Management Database stores information about the following:

- Configured devices, libraries, library drives, and slots.
- Data Protector media used for backup.
- Configured media pools and media magazines.

CDB The Catalog Database stores information about the following:

- Backup sessions and restore sessions. This is the copy of the information sent to the Data Protector Monitor window.
- Backed up objects and their versions.
- Pathnames of backed up files (filenames) together with client system names. Filenames are stored only once per client system. The filenames created between backups are added to the CDB.
- Positions of backed up objects on media. For each backed up object, Data Protector stores information about the media and data segments used for the backup.

DCBF — The Detail Catalog Binary Files part stores file version information. This is information about backed up files, such as file size, modification time, attributes/protection, and so on.

One DC (Detail Catalog) binary file is created for each Data Protector medium used for backup. When the medium is overwritten, the old binary file is removed and a new one is created.

SMBF — The Session Messages Binary Files part stores session messages generated during backup sessions and restore sessions. One binary file is created per session. The files are grouped by year and month.

SIBF — The Serverless Integrations Binary Files part stores raw NDMP restore data. This data is necessary for restore of NDMP objects.

The MMDB and CDB parts are implemented using an embedded database consisting of tablespaces. This database is controlled by the rds database server process. All changes to the MMDB and CDB are updated using transaction logs. CDB (objects and positions) and MMDB present the core part of IDB.

The DCBF, SMBF, and SIBF parts of the IDB consist of binary files. Updates are direct (no transactions).

In the Manager-of-Managers (MoM) environment, the MMDB can be moved to a central system to create the Central Media Management Database (CMMDB).

18.2004
Louiza

For additional information on each of the IDB parts, refer to the *HP OpenView Storage Data Protector Concepts Guide.*

18.003
Laura

Configuring the IDB

The IDB configuration helps to manage the following:

- the size of the IDB and available disk space
- the location of the IDB directories
- transaction log usage
- the IDB backup necessary in case of IDB corruption or a disaster
- configuration of the IDB reports and notifications

Once the IDB is configured, it should be maintained only when you are notified about the need.

General Procedure This is the general procedure for IDB configuration:

1. Allocate disk space for future needs.
   Refer to "Allocating Disk Space for Future Use" on page 388.
2. Prepare for the IDB recovery.
   Refer to "Preparing for IDB Recovery" on page 390.
3. Set the appropriate reports and notifications about the IDB.
   Refer to "Configuring the Database Reports and Notifications" on page 400.

Allocating Disk Space for Future Use

Over time, the IDB can occupy a considerable amount of disk space on the Cell Manager. You need to plan in advance and consider the allocation of the disk space for future IDB needs.

Prerequisites

- You need to understand the key factors influencing the IDB growth, such as number of files, file dynamics, environment growth, and so on. Refer to the *HP OpenView Storage Data Protector Concepts Guide* for additional information.
- You need to set logging level and catalog protection polices according to your environment requirements and available disk space. To get this information, together with the usage recommendations for logging level and catalog protection settings, refer to the *HP OpenView Storage Data Protector Concepts Guide.*

- You need to estimate future IDB size (disk space necessary for future IDB needs). Refer to the *HP OpenView Storage Data Protector Concepts Guide* for the IDB size estimation.

How Much Disk Space Is Needed?

The disk space needed to accommodate the IDB varies significantly as a function of many configuration aspects and policies used in defining and operating backups.

The following simplified scenario of an environment requires about 900 MB of disk space for the IDB after 3 months, with very little growth afterwards:

- 100 systems to be backed up (10,000 files each; without mail-servers)
- 350 GB total data volume
- filesystem backups with typical dynamics of 3% of new files per month
- one full backup and four incremental backups per week
- logging level set to Log all (to allow convenient browsing of filenames before restore). This is the most demanding logging option.
- catalog protection setting of three months for the full backups and two weeks for the incremental backups.

Note that large configurations or long catalog protection periods in the IDB can require more than 20 GB for the IDB.

A detailed estimation can be performed using the IDB Capacity Planning Tool located on the Cell Manager:

- On UNIX: `/opt/omni/doc/C/IDB_capacity_planning.xls`
- On Windows: `<Data_Protector_home>\docs\IDB_capacity_planning.xls`

What to Plan for in Advance

Typically the IDB grows rapidly in the beginning, until the catalog retention periods have been reached. After that, the growth of the IDB is mainly determined by the dynamics of systems that have a large percentage of new files per month and the growth of the environment itself (new systems to be backed up).

It is important to understand the various IDB growth functions:

- The filenames part of the IDB is proportional to the total number of filenames in the cell (but not the data volume and the number of backups). Typically the filename growth is moderate, with the exception of some mail servers or other systems with a large amount of automatically generated files.
- The file versions part of the IDB grows with the number of backups, the number of files in the cell, and the duration of the catalog protection.
- Using the IDB transaction log files requires additional disk space. Size prediction is not simple. Dominating factors influencing the size are the number of new filenames being backed up and the total backup activities (or weeks, if scheduled backups are the main operation) between IDB backups.

Preparing for IDB Recovery

You need to make advance preparations in order to be able to recover the IDB at any point in time. The IDB recovery restores information stored in the IDB and is essential for the restore of backed up data in case the Cell Manager crashes.

Prepare for IDB recovery by:

- Considering recommendations for optimizing robustness. Refer to “Robustness Considerations” on page 390.
- Relocating IDB directories. Refer to “The IDB Directories” on page 391.
- Enabling of transaction logs. Refer to “Enabling Transaction Logs” on page 396.
- Configuring the IDB backup and backing it up regularly. Refer to “Configuring the Database Backup” on page 398.

Robustness Considerations

This section outlines some aspects and recommendations you should consider to optimize robustness and reliability of the IDB.

- The core part of the IDB, which contains CDB (objects & positions) and MMDB, is essential for the operation of Data Protector.

38.000
Paula

- The DCBF and SMBF parts of the IDB are not required for basic operation of Data Protector, such as backup and restore. However, if they are not present, restore becomes less convenient (no filename browsing) and the session messages are lost.
- If the IDB recovery file and the IDB transaction logs are lost, normal operation would not be affected, but IDB restore would be considerably more difficult, and replaying the IDB data generated since the last IDB backup would not be possible. Instead, the used media would need to be reimported.

Recommendations to Optimize Robustness

- Ensure that the IDB recovery file and the transaction logs do not reside on the same physical disk as the core part of the IDB.

  This is to ensure a fast and simple restore of the IDB in case the physical disk A crashes. It also for the replay of the transactions that happened since the last IDB backup. Refer to Figure 9-2.
- Relocating the DCBF, SMBF, and SIBF parts to a disk other than the one that holds the core part of the IDB is also recommended, but less important. If this is done, the load on disk A is reduced significantly and IDB space management is easier, because these parts are usually the largest part of the IDB.

TIP Following the recommendations to optimize robustness will also increase performance, allowing for more backup activities on the Cell Manager system.

The IDB Directories

The IDB is located on the Cell Manager. In order to improve space management, you may want to relocate some IDB directories.

Limitations

- On Windows NT 4.0 systems, it is not possible to change the location of the IDB directories.
- The IDB files can be located only on locally attached disks (not using NFS or on shared disks).

17.999
Paula

- If the IDB is installed in a cluster, it must be installed on disks in the cluster group (Microsoft Cluster Server) or cluster package (MC/ServiceGuard).

Table 9-1 **Location of IDB Directories on Windows**

IDB	Location on Windows
Tablespaces (CDB and MMDB)	`<Data_Protector_home>\db40\datafiles`
Binary files (DCBF, SMBF, SIBF)	• `<Data_Protector_home>\db40\dcbf` • `<Data_Protector_home>\db40\msg` • `<Data_Protector_home>\db40\meta`
Transaction logs	`<Data_Protector_home>\db40\logfiles\syslog`
IDB recovery file	`<Data_Protector_home>\db40\logfiles\rlog`

Table 9-2 **Location of IDB Directories on UNIX**

IDB	Location on UNIX
Tablespaces (CDB and MMDB)	`/var/opt/omni/db40/datafiles`
Binary files (DCBF, SMBF, SIBF)	• `/var/opt/omni/db40/dcbf` • `/var/opt/omni/db40/msg` • `/var/opt/omni/db40/meta`
Transaction logs	`/var/opt/omni/db40/logfiles/syslog`
IDB recovery file	`/var/opt/omni/db40/logfiles/rlog`

17998

Figure 9-2 **Recommended Location of IDB Directories**

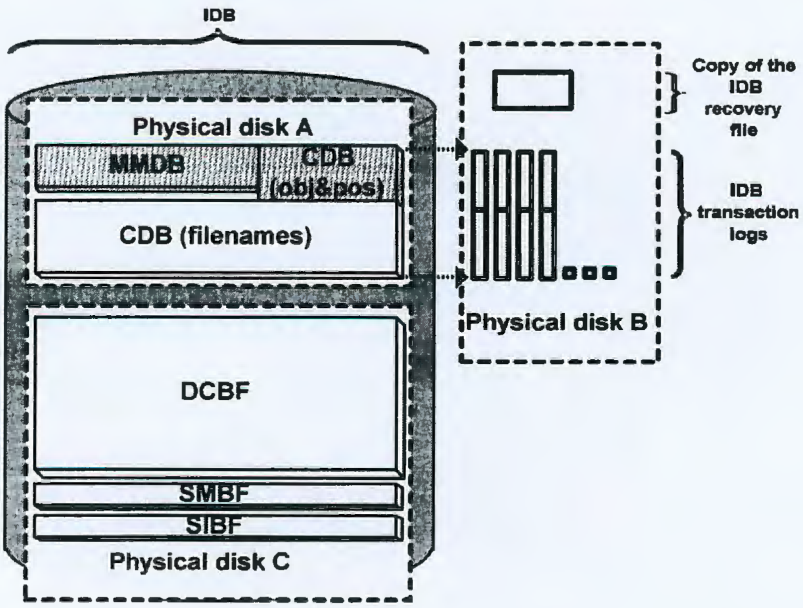

Relocating the IDB Directories

You can change the location of any of the following IDB directories:

the

- `datafiles` directory, containing CDB (objects, positions, and filenames) and MMDB parts of the IDB
- the `logfiles` directory, containing transaction logs and the IDB recovery file
- the `dcbf` directory, containing the DCBF part of the IDB
- the `msg` directory, containing the SMBF part of the IDB
- the `meta` directory, containing the SIBF part of the IDB

You can also modify the directory path for the `dcbf` directory (using the Data Protector user interface) and for the `msg` and `meta` directories (using the global options file).

17991
Paula

NOTE On UNIX, you can use symbolic links to relocate the directories, but the links are not allowed beneath the `/var/opt/omni/db40/datafiles` directory.

Follow the described below to relocate the IDB directories:

1. Stop all backups and other Data Protector activities and run the `omnisv -stop` command to stop the Data Protector services:
   - On Windows: `<Data_Protector_home>\bin\omnisv -stop`
   - On UNIX: `/opt/omni/sbin/omnisv -stop`

   If the IDB is installed on MC/ServiceGuard, run the `cmhaltpkg` *<pkg_name>* command on the active node to stop the Data Protector package, where *<pkg_name>* is the name of the Data Protector cluster package.
2. Rename the *<IDB_dir>* directory that you want to move to *<IDB_dir>*`.save`. For example, to relocate the transaction logs and the IDB recovery file, rename `<Data_Protector_home>\db40\logfiles` to `<Data_Protector_home>\db40\logfiles.save` (on Windows), or `/var/opt/omni/db40/logfiles` to `/var/opt/omni/db40/logfiles.save` (on UNIX).
3. Create a new empty directory with the same relative path, for example `<Data_Protector_home>\db40\logfiles` on Windows systems, or `/var/opt/omni/db40/logfiles` on UNIX systems.
4. On Windows, add a new disk or mount a new volume at an NTFS folder as `<Data_Protector_home>\db40\<IDB_dir>`. For example, mount it as `<Data_Protector_home>\db40\logfiles`.

   On UNIX, add a new disk or create a new logical volume and mount it as `/var/opt/omni/db40/<IDB_dir>`. For example, mount it as `/var/opt/omni/db40/logfiles`.
5. Copy the contents of *<IDB_dir>*`.save` into *<IDB_dir>* on the new disk or new volume.
6. Run the `omnisv -start` command to start the Data Protector services:
   - On Windows: `<Data_Protector_home>\bin\omnisv -start`

- On UNIX: `/opt/omni/sbin/omnisv -start`

  If the IDB is installed on MC/ServiceGuard, run the `cmrunpkg <pkg_name>` command on the active node to start the Data Protector package.

Creating an Additional Copy of the IDB Recovery File

Creating an additional copy of the IDB recovery file prevents you from losing important data for IDB recovery.

Use the following steps to make another copy of the IDB recovery file:

1. Stop all backups and other Data Protector activities and run the `omnisv -stop` command to stop the Data Protector services.

   - On Windows: `<Data_Protector_home>\bin\omnisv -stop`
   - On UNIX: `/opt/omni/sbin/omnisv -stop`

   If the IDB is installed on MC/ServiceGuard, run the `cmhaltpkg <pkg_name>` command on the active node to stop the Data Protector package, where `<pkg_name>` is the name of the Data Protector cluster package.

   If the IDB is installed on Microsoft Cluster Server, take the `OBVS_VELOCIS` cluster group offline using the Cluster Administrator utility on the active node.

2. Edit the global options file by setting the value for the `RecoveryIndexDir` variable: specify an additional location where Data Protector makes a copy of the IDB recovery file, `obrindex.dat`. It is recommended to specify a different physical disk.

3. Run the `omnisv -start` command (on UNIX, located in the `/opt/omni/sbin` directory) to start the Data Protector services.

   - On Windows: `<Data_Protector_home>\bin\omnisv -start`
   - On UNIX: `/opt/omni/sbin/omnisv -start`

   If the IDB is installed on MC/ServiceGuard, run the `cmrunpkg <pkg_name>` command on the active node to start the Data Protector package.

   If the IDB is installed on Microsoft Cluster Server, bring the `OBVS_VELOCIS` and `OBVS_MCRS` cluster groups online using the Cluster Administrator utility.

Creating or Relocating DC Directories

Creating a DC Directory

Create a DC directory using the `Database` context in the `Data Protector Manager`. See Figure 9-3. For detailed steps, refer to the online Help index keyword “creating DC directories”.

Relocating a DC Directory

To change the location of a DC directory, proceed as follows:

1. Create a new DC directory on a new location, using the Data Protector user interface. See Figure 9-3.
2. Verify that the new DC directory has been created and has enough disk space.
3. Move DC binary files from the source DC directory to the new DC directory.
4. Run the `omnidbutil -remap_dcdir` command to update the pathnames of DC binary files.
5. Remove the old DC directory from the list of configured DC directories.

Figure 9-3 **Creating a DC Directory**

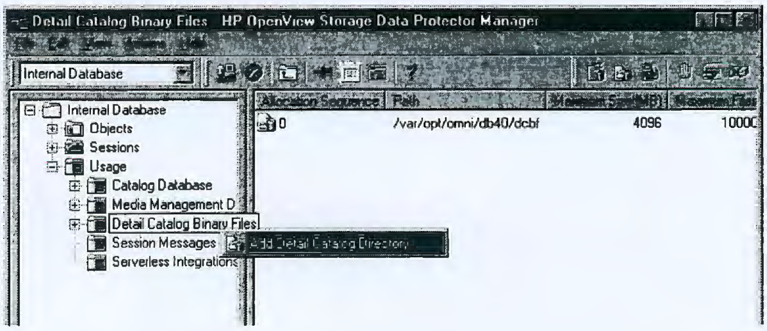

Allocation Sequence	Path	Maximum Size(MB)	Maximum Files
0	/var/opt/omni/db40/dcbf	4096	1000C

Enabling Transaction Logs

Transaction logs used by the MMDB and CDB parts of the IDB are created in the following directory:

- On Windows: `<Data_Protector_home>\db40\logfiles\syslog`
- On UNIX: `/var/opt/omni/db40/logfiles/syslog`

17.904
Paula

By default, transaction logging is disabled. If enabled, transaction logs from the latest IDB backup are kept until the next backup. If a transaction log file reaches 2 MB, a new one is created. An IDB backup removes all existing transaction logs, except for the currently active one, and starts to create new ones.

Why Enable Transaction Logs?

In order to perform the most convenient IDB recovery method, **guided autorecovery**, with replaying logs, you need to have available the transaction log files created after the last IDB backup.

Disk Space Considerations

The disk space used for the transaction logs depends on the amount of backups done between two IDB backups. If the filenames are already in the IDB, the amount is fairly small and the reserved space of 100 MB should be enough for most cases. If new filenames are backed up, the disk space usage is considerable (estimation is 200 bytes per filename). It is recommended to enable transaction logs *after* the first full backup of the environment (when all filenames are stored in the IDB).

How to Enable the Transaction Logs

To enable transaction logs, proceed as follows:

1. Stop all backups and other Data Protector activities and run the `omnisv -stop` command to stop Data Protector services:
   - On Windows: `<Data_Protector_home>\bin\omnisv -stop`
   - On UNIX: `/opt/omni/sbin/omnisv -stop`
2. Ensure that there is enough disk space in the following directory:
   - On Windows: `<Data_Protector_home>\db40\logfiles\syslog`
   - On UNIX: `/var/opt/omni/db40/logfiles/syslog`
3. Edit the `velocis.ini` file and set the value of the Archiving parameter to 1.
   - On Windows: `<Data_Protector_home>\db40\datafiles\catalog\velocis.ini`
   - On UNIX: `/var/opt/omni/db40/datafiles/catalog/velocis.ini`
4. Start the Data Protector services using the `omnisv -start` command:

17.993
Paula

- On Windows: `<Data_Protector_home>\bin\omnisv -start`
- On UNIX: `/opt/omni/sbin/omnisv -start`

Configuring the Database Backup

An essential part in the IDB configuration is to configure the backup of the IDB itself. Once the IDB backup is performed regularly, the most important preparation for recovery in case of a disaster is done. The IDB recovery is essential for restore of other backed up data in the event that the Cell Manager crashes.

How to Configure the IDB Backup

Configure the IDB backup like any standard backup, but be sure to select the `Internal Database` object and specify the object options in the Backup Object Summary page of the IDB backup specification. For detailed steps, see the Data Protector online Help index keyword "configuring IDB backups".

Figure 9-4 **Selecting the Internal Database Object**

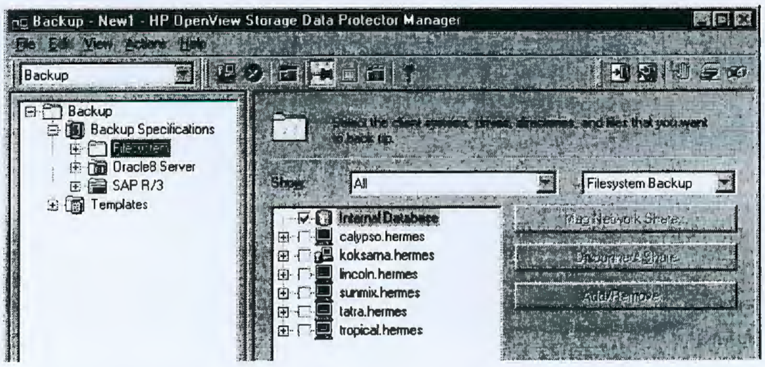

Recommended IDB Backup Configuration

We recommend the following when configuring the IDB backup:

- Create a separate backup specification for the IDB. This simplifies scheduling and restoring in case of a disk crash. To create an IDB backup specification, follow the standard backup procedure, but select the `Internal Database` object.
- Schedule the IDB backup to be performed once per day. This ensures that you always have an almost up-to-date backup of the IDB.

17.992
Paula

- Perform the IDB backup using a separate media pool on separate media, on a specific device. Make sure you know which media you use for the IDB backup. You can configure a `Session Media Report` to be informed about the media used for the backup. This greatly simplifies eventual restore. If possible, use a device locally connected to the Cell Manager. Refer to "Data Protector Reporting" on page 315.
- Set data protection and catalog protection to a few days only. Set these options such that you have at least the last two IDB backup versions protected.
- Always have the `Check Internal database` option enabled (default). See Figure 9-5.
- Do not overwrite the previous IDB backup with the new one (keeping several copies is suggested).

What Happens During the IDB Backup

During the IDB backup, Data Protector does the following:

- Checks the consistency of the IDB, thus preventing the backing up and later restoring a corrupted IDB. For this check to happen, you need to have the `Check Internal database` option enabled (default).

  The check operation takes approximately 1.5 hours for a 10 GB database with a `fnames.dat` file size of 1 GB.
- Backs up the IDB online (while the IDB is in use). Therefore, other backup or restore sessions can run while the IDB backup runs. But, if possible, back up the IDB when no other backup and restore activities are in progress.
- Backs up all Data Protector configuration data, including the data on devices, backup specifications, and schedules. This simplifies recovery in case of a disaster.

NOTE Only one IDB backup can run at a time.

Disabling the Automatic Check Before Backup

By default, Data Protector automatically checks the consistency of the IDB before the database is backed up. You can enable or disable the automatic consistency check. It is strongly recommended that you keep the automatic IDB check enabled.

In environments where the Cell Manager is used heavily and the time needed to perform the check of the IDB creates a problem, you may need to disable the `Check Internal database` option. In such cases, consider the following suggestions:

- Schedule the IDB backup with the IDB check option enabled to be performed when the automatic check activity is acceptable.
- Schedule the daily IDB backup with the IDB check option disabled.
- Keep at least the most recent checked IDB backup.

For detailed steps, refer to the online Help index keyword "disabling automatic IDB checks".

Figure 9-5 **The Check Internal database Option (enabled by default)**

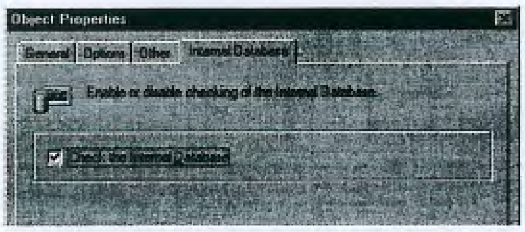

Enable or disable checking of the Internal Database.

☑ Check the Internal Database

Configuring the Database Reports and Notifications

Configure the IDB reports and notifications so that you are notified if you need to perform IDB maintenance tasks such as purging the IDB, extending the size of the IDB, and so forth.

IDB Reports The following list presents the IDB reports:

IDB Purge Preview Report Lists the number of filenames per client, the estimated number of obsolete filenames per client, and the estimated duration of the filename purge session per client.

Report on System Dynamics Reports on the dynamics of the growth of filenames on a particular client.

IDB Purge Report Lists the filenames that have been removed from the IDB.

17.990
Paula

IDB Size Report Lists the sizes of the individual parts of the IDB.

There are also other Data Protector reports to be considered. For example, the `List of Sessions` report shows the number of files backed up in one session. Refer to "Data Protector Reporting" on page 315 for more information.

IDB Notifications

The following list presents IDB notifications:

Low Database Space Informs you if the IDB is running out of space.

Database Purge Needed Informs you if you need to run the filename purge of the IDB.

Database Corrupted Informs you if any kind of IDB corruption is detected.

For detailed information on each report and notification, refer to "Report Types" on page 317.

Procedure for Configuring IDB Reports and Notifications

Configure the IDB reports and notifications using the `Reporting` context in the `Data Protector Manager`. For detailed steps, refer to the online Help index keywords "configuring IDB reports" and "configuring IDB notifications".

What's Next?

Once you have configured the IDB reports and notifications, you have completed the last step in IDB configuration. If you need to perform any IDB maintenance task, you will be notified by Data Protector. Now, you can continue to set up your environment.

Maintaining the IDB

Once you have configured the IDB, you need to perform IDB maintenance tasks in the following cases:

- the IDB is running out of space

  If configured, the `Low Database Space` notification informs you about this.

- the IDB needs a file version purge

  With the OmniBack II A.03.50/A.03.51 and earlier databases, the level of the granularity of the purge was at the object level. Now, the granularity of the purge is the complete medium. This means that the catalog protection for all object versions on the medium must expire before the file version records are purged. Then, the related medium binary file containing the detail catalog is removed. This purges many file versions in a very short time. This happens automatically on a daily basis. Obsolete sessions and messages are also purged automatically.

- the IDB needs a filenames purge

  Purging filenames was a regular maintenance task in OmniBack II A.03.50/A.03.51 and earlier versions. With the OmniBack II A.04.10 release, the frequency is reduced to once per year in an environment that can generate 100,000 obsolete filenames per day. You will be notified automatically if the filenames purge is needed. Filenames purge can be executed selectively on a per host basis. The operation must run exclusively, so no backups can run concurrently. This purge takes more time to execute than in the previous version of Data Protector.

- the dynamics of the client system are high or critical

  If configured, the `System Dynamics` report informs you about this.

- you want to move the IDB to a different Cell Manager

- you want to check the size of the IDB

  The `Database Size` report informs you of the size of the IDB.

- the IDB does not work properly (might be corrupted) and you want to check its consistency

17955
Paula

The `Database Corrupted` notification informs you about IDB corruption.

Refer to Table 9-3 for information on which of the maintenance tasks you can perform in which cases.

Table 9-3 **IDB Maintenance Tasks**

Situation	Which Task Can You Perform?	Reference
The IDB is running out of space	• Extend the size of the IDB • Purge the IDB filenames • Reduce the growth of the IDB • Reduce the current size of the IDB	• "Extending the Database Size" on page 408 • "Purging Obsolete Filenames" on page 408 • "Reducing the IDB Growth" on page 405 • "Reducing the IDB Size" on page 406
Obsolete filenames in the IDB	• Purge IDB filenames	• "Purging Obsolete Filenames" on page 408
The dynamics of a client system are high or critical	• Reduce the growth of the IDB • Extend the size of the IDB	• "Reducing the IDB Growth" on page 405 • "Extending the Database Size" on page 408
You want to check the size of the IDB	• Check the size of the IDB	• "Checking the Database Size" on page 410
The IDB does not work properly (might be corrupted)	• Check the consistency of the IDB	• "Checking the Consistency of the Database" on page 411
You want to move the IDB to a different Cell Manager	• Move the IDB to a different Cell Manager on the same platform	• "Moving the Database to a Different Cell Manager" on page 412

17.486
Paula

Reducing the IDB Growth

You can reduce the growth of the IDB by reducing the logging level and catalog protection settings of your backup specifications. These actions do not influence the current size of the IDB, but they do influence its future growth.

The effect of reducing the logging level is a reduction in browse comfort at restore time.

The effect of reducing the catalog protection is that browsing is not possible for some restores (namely of those backups that have exceeded the catalog protection).

Refer to the *HP OpenView Storage Data Protector Concepts Guide* for information on key factors and tunable parameters for IDB growth and performance, as well as for usage recommendations.

How to Reduce the IDB Growth

Modify the backup specifications by changing the logging level and catalog protection settings using the `Data Protector Backup` context in the `Data Protector Manager`. See Figure 9-6. For detailed steps, refer to the online Help index keyword “reducing IDB growth”.

By reducing the logging level settings for a backup specification, you reduce the amount of data (files/directories) that will be stored in the IDB (`Log all -> Log files -> Log directories -> No log`).

By reducing the catalog protection, you reduce the protection for the (restore browse) information in the IDB only. The information is still stored on media.

Figure 9-6 **Changing Logging Level and Catalog Protection Settings**

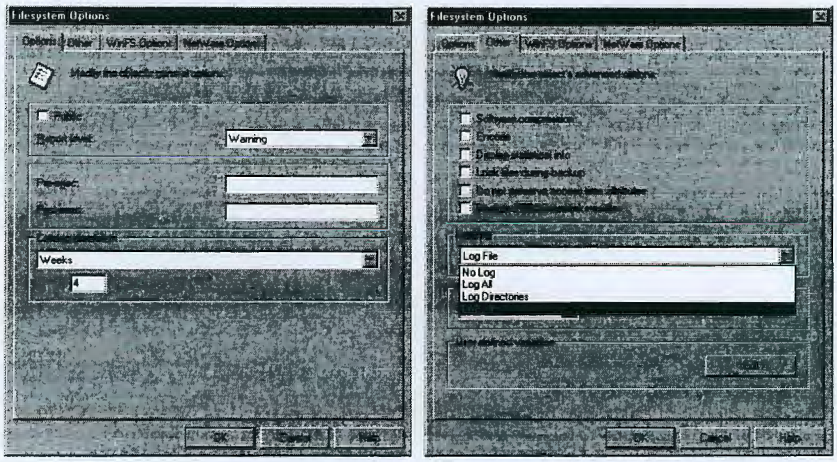

Reducing the IDB Size

You can reduce the IDB size by changing the catalog protection settings for a complete backup session (all objects in the session) or for specific objects only.

The effect of reducing the catalog protection is that browsing is not possible for some restores (namely of those backups that have exceeded the catalog protection).

This action does not influence the future growth of the IDB.

When Does the Change Take Effect?

The change takes effect:

- If catalog protection is removed from all objects on the medium.
- Once per day (by default, at noon) when Data Protector automatically removes obsolete data from the IDB. The time can be specified in the `DailyMaintenanceTime` global options variable, using the twenty-four hour clock notation. Refer to "Global Options File" on page 523.

  You can start the purge immediately by running the `omnidbutil -purge -dcbf` command. Refer to the `omnidbutil` man page for information on removing other obsolete items from the IDB.

By changing the catalog protection, you change protection in the IDB only. The information is still stored on media. Therefore, if you export media and import it back, Data Protector rereads information about catalog protection from the media.

How to Reduce the IDB Size

Change the logging level and catalog protection settings using the `Internal Database` context in the `Data Protector Manager`. See Figure 9-7 and Figure 9-8. For detailed steps, refer to the online Help index keyword "reducing IDB current size".

Figure 9-7 **Changing Catalog Protection for a Session**

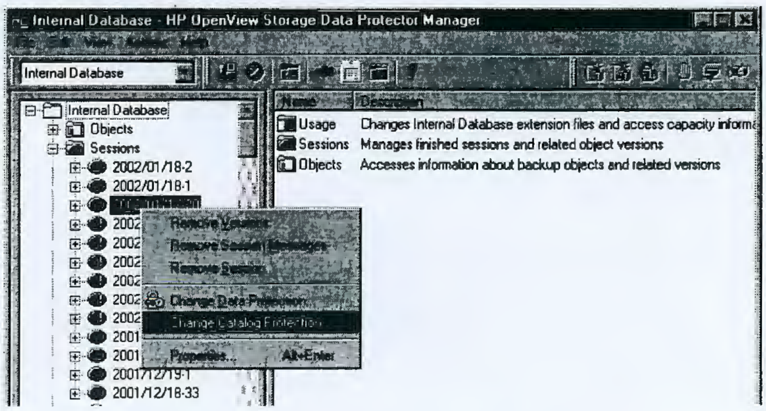

Name	Description
Usage	Changes Internal Database extension files and access capacity informa
Sessions	Manages finished sessions and related object versions
Objects	Accesses information about backup objects and related versions

Figure 9-8 **Changing Catalog Protection for an Object**

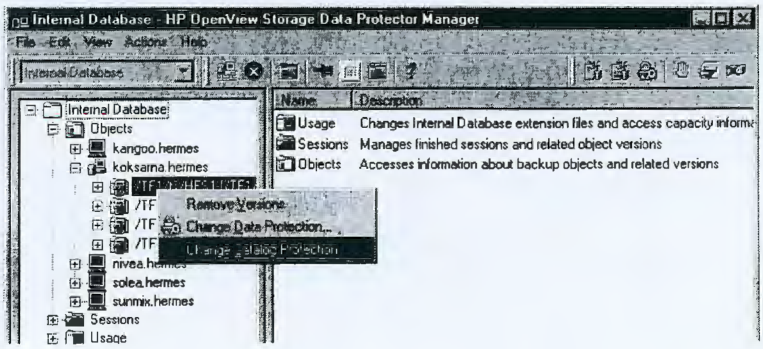

Name	Description
Usage	Changes Internal Database extension files and access capacity informa
Sessions	Manages finished sessions and related object versions
Objects	Accesses information about backup objects and related versions

17983
Paula

Purging Obsolete Filenames

During the purge process, Data Protector automatically checks for and purges obsolete filenames from the IDB to free up space for new information. A filename becomes obsolete when there are no file versions for the filename in the IDB.

Use the `Internal Database Purge Preview Report` and `Internal Database Purge Report` to get more information about the purge. Refer to "Configuring the Database Reports and Notifications" on page 400.

How to Purge Obsolete IDB Filenames

Purge the IDB when no other backups are running on the Cell Manager. Run the following command:

```
omnidbutil -purge -filenames
```

You can limit the purge to one or more clients by running the following command:

```
omnidbutil -purge -filenames <host_1 ... host_n>
```

Data Protector skips purging filenames on the clients that have fewer than 1,000,000 obsolete filenames. In order to purge filenames on these clients as well, use the `-force` subcommand.

Extending the Database Size

It is required to extend the IDB size for the following reasons:

- The space for the filenames is consumed and another `fnames.dat` file is needed.
- More disk space is needed for the detail part of the IDB (file versions and attributes)

You can extend the size of the IDB in either of two ways:

- By creating new DC (Detail Catalog) directories and, possibly, locating them on different disks.
- By creating additional `fnames.dat` files.

Creating New DC Directories

You create a new DC directory using the `Internal Database` context in the `Data Protector Manager`. See Figure 9-3 on page 396. For detailed steps, refer to the online Help index keyword "creating DC directories".

17983
Paula

Creating New fnames.dat Files

What Are `fnames.dat` Files?

The `fnames.dat` files contain information on the names of backed up files. Typically, these files occupy about 20% of the IDB. The default size of a `fnames.dat` file is 2 GB; the maximum size is 32 GB.

How to Create `fnames.dat` Files

You add new `fnames.dat` files using the `Internal Database` context in the `Data Protector Manager`. See Figure 9-9. For detailed steps, refer to the online Help index keyword "creating fnames.dat files".

On Windows Cell Managers, it is recommended that the extension files are created on the same logical disk as the IDB.

The IDB extension files are backed up as a part of the IDB backup and are restored using the IDB recovery.

Figure 9-9 **Creating a New `fnames.dat` File**

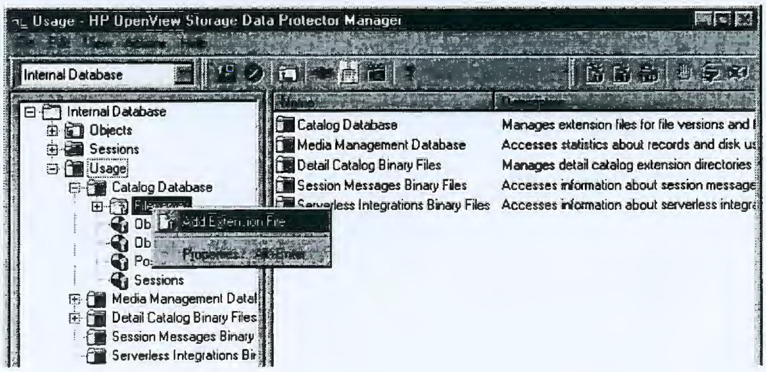

Catalog Database	Manages extension files for file versions and f
Media Management Database	Accesses statistics about records and disk us
Detail Catalog Binary Files	Manages detail catalog extension directories
Session Messages Binary Files	Accesses information about session message
Serverless Integrations Binary Files	Accesses information about serverless integra

17951
Paula

Checking the Database Size

You can check the current size of the IDB parts using the Data Protector GUI.

Also, if configured, the Database Size Report and Low Database Size notifications report on the IDB size.

How to Check IDB Size

Check the size of the IDB parts, CDB, MMDB, DCBF, SMBF, and SIBF using the `Internal Database` context in the `Data Protector Manager`. See Figure 9-10. For detailed steps, refer to the online Help index keyword “checking, IDB size”.

Figure 9-10 **Checking the Size of the `fnames.dat` File (CDB Part)**

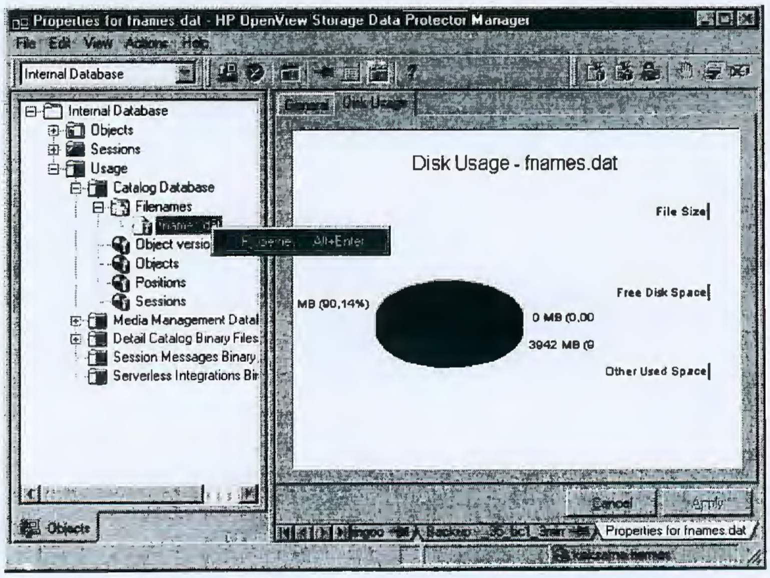

17.985
Paula

Checking the Consistency of the Database

Data Protector by default checks the consistency of the IDB before the IDB is backed up. This is extremely important for recovering the IDB and backed up data in case of a disaster.

Additionally, you can manually perform the following IDB checks:

Check of the core part of the IDB Checks the MMDB (Media Management Database) and CDB (Catalog Database) parts without information about filenames. It takes approximately 5-10 minutes for a medium size IDB. To perform it, run the `omnidbcheck -core` command.

Filenames check Checks IDB information about filenames. It takes approximately one hour for a medium size IDB. To perform it, run the `omnidbcheck -filename` command.

Simple check of the DCBF part Checks if the DC binary files exist and what their size is. It takes approximately 10-30 seconds for a medium size IDB. To perform it, run the `omnidbcheck -bf` command.

Complete check of the DCBF part Checks the consistency of media positions and the DC binary files. It takes approximately 10 minutes for each GB of the DCBF part. To perform it, run the `omnidbcheck -dc` command.

Check of the SMBF part Checks for the presence of session messages binary files. It takes approximately 5-10 minutes. To perform it, run the `omnidbcheck -smbf` command.

Check of the SIBF part Checks the consistency of object versions and Serverless Integrations Binary Files. It takes approximately 10 minutes for each GB of the SIBF part. To perform it, run the `omnidbcheck -sibf` command.

Quick check Checks the core part (MMDB and CDB), filenames, and the DCBF part. It takes approximately two and a half hours for a medium size IDB. To perform it, run the `omnidbcheck -quick` command.

17979
Paula

Extended check Checks the critical part (MMDB and CDB), filenames, the DCBF part, and the DC part. To perform it, run the `omnidbcheck -extended` command.

If you run into problems using the IDB, refer to the troubleshooting section "Troubleshooting the IDB" on page 592 and "Recovering the IDB" on page 417.

Moving the Database to a Different Cell Manager

You can move the IDB to a different Cell Manager that runs on the same operating system by following the steps below:

1. Stop all Data Protector services on the source and target systems using the `omnisv -stop` command:

   - On Windows: `<Data_Protector_home>\bin\omnisv -stop`
   - On UNIX: `/opt/omni/sbin/omnisv -stop`

   If the IDB is installed on MC/ServiceGuard, run the `cmhaltpkg <pkg_name>` command on the active node to stop the Data Protector package, where *<pkg_name>* is the name of the Data Protector cluster package.

   If the IDB is installed on Microsoft Cluster Server, take the `OBVS_VELOCIS` cluster group offline using the Cluster Administrator utility on the active node.

2. Copy the following IDB files to the target system:

   - Tablespaces to the same relative pathname:

     On Windows systems:
     `<Data_Protector_home>\db40\datafiles` to
     `<Data_Protector_home>\db40\datafiles`

     On UNIX systems: `/var/opt/omni/db40/datafiles` to
     `/var/opt/omni/db40/datafiles`

   - Extension files to the same full pathname as they were on the source system. You can get a list of the files by using the `omnidbutil -extendinfo` command.
   - SMBF files to the same relative pathname:

17978
Paula

On Windows systems: *<Data_Protector_home>*\db40\msg to *<Data_Protector_home>*\db40\msg

On UNIX systems: /var/opt/omni/db40/msg to /var/opt/omni/db40/msg

- SIBF files to the same relative pathname:

  On Windows systems: *<Data_Protector_home>*\db40\meta to *<Data_Protector_home>*\db40\meta

  On UNIX systems: /var/opt/omni/db40/meta to /var/opt/omni/db40/meta

- DC directories to the same or other locations. You can get the list of DC directories using the omnidbutil -list_dcdir command.

3. Start Data Protector services on the target system using the omnisv -start command:

   - On Windows: *<Data_Protector_home>*\bin\omnisv -start
   - On UNIX: /opt/omni/sbin/omnisv -start

   If the IDB is installed on MC/ServiceGuard, run the cmrunpkg *<pkg_name>* command on the active node to start the Data Protector package.

   If the IDB is installed on Microsoft Cluster Server, bring the OBVS_VELOCIS and OBVS_MCRS cluster groups online using the Cluster Administrator utility.

4. Run the omnidbutil -change_cell_name command.
5. Relocate DC directories on the target system.
6. Run the omnidbutil -remap_dcdir command for Data Protector to refresh the new locations of the DC directories.

Restoring the IDB

If you have backed up the IDB using the standard procedure, you can restore it using the methods described in this section.

For a detailed description of how to handle the IDB recovery in case of a disaster, refer to “Recovering the IDB” on page 417.

Restoring the IDB consists of two phases:

1. Restoring the IDB to a temporary location.

IMPORTANT This step is necessary because the IDB is in use during the restore. If you try to restore the IDB to the original location, you will corrupt the IDB.

2. Moving the IDB to the original location.

Ensure that you have enough disk space before you begin.

Restoring the IDB to a Temporary Directory

To restore the IDB files to a temporary location, proceed as follows:

1. In the `Data Protector Manager`, switch to the `Restore` context.
2. Expand the `Internal Database` item.
3. Expand the client system with the IDB backup and then click the database object to open the `Source` property page.
4. In the `Source` property page, select the IDB directories that you want to restore. By default, the last backup version is selected for restore. If you want to restore any other version, right-click the selected directory and click `Restore Version`. From the `Backup version` drop-down list, select the backup version that you want to be restored. Click `OK`.
5. In the `Destination` property page select `Restore to new location` option and select the temporary directory for IDB files (for example, the `temp` directory).

1776
Paula

NOTE You should not select the *<Data_Protector_home>* directory, as this directory is the original location of the IDB.

If you want to restore to a different system, specify the new Cell Manager's name.

6. Click `Restore`.

Moving the IDB to the Original Location

After you have restored the IDB to a temporary location, you need to move the IDB directories to their original location. Proceed as follows:

On an UNIX Cell Manager

1. Stop all running Data Protector sessions and close the Data Protector GUI. This prevents access to the IDB.
2. Stop all Data Protector processes by running:

   ```
   /opt/omni/sbin/omnisv -stop
   ```

3. Move the existing IDB directories:

   `/var/opt/omni/db40` and `/etc/opt/omni`

   This prevents merging of old and new files.
4. Copy the IDB directories from the temporary directory to the original directories

   `/var/opt/omni/` and `/etc/opt/omni`

   If your extension files were located on some other directory, be sure to copy them to the original disk and directory as well.
5. Restart the Data Protector processes by running:

   ```
   /opt/omni/sbin/omnisv -start
   ```

On a Windows Cell Manager

1. Stop all running Data Protector sessions and close the Data Protector GUI. This prevents access to the IDB.
2. Stop all Data Protector services by running:

   ```
   <Data_Protector_home>\bin\omnisv -stop
   ```

3. Move the existing IDB directories (`db40` and `config`) from the *`<Data_Protector_home>`* directory. This prevents merging of old and new files.

4. Copy the IDB directories from the temporary directory to the original directory *`<Data_Protector_home>`*.

   If your extension files were located on some other directory, be sure to copy them to the original disk and directory as well.

5. Restart the Data Protector services by running:

   ```
   <Data_Protector_home>\bin\omnisv -start
   ```

TIP

You can check the consistency of the IDB after the restore. See "Checking the Consistency of the Database" on page 411 for more information.

17974
Paula

Recovering the IDB

When Is Recovery Needed?

IDB recovery is needed if all or some of the IDB files are not available or are corrupted.

There are three levels of IDB issues, each with its own techniques for repair:

- Troubleshoot the IDB problems that are caused by OS configuration issues, such as not mounted filesystems, nameservice problems, and so on. Refer to the troubleshooting section "Troubleshooting the IDB" on page 592.
- Omit or remove non-core parts (binary files or filenames part) of the IDB that contain problems. This is possible if the identified level of IDB corruption is minor or major (meaning the corruption is not in the core part of the IDB).
- Perform a complete recovery. This consists of restoring the IDB and updating information that has been modified since the last IDB backup. This is a must if the identified level of IDB corruption is critical (meaning the corruption is in the core part).

Complete Recovery

Complete recovery consists of two phases:

1. IDB restore, which gets the IDB to the last (available) consistent state.
2. Updating the IDB from the last consistent state up to the last moment when the IDB was still operational.

Depending on how well you prepared for IDB recovery before problems occurred (availability of IDB recovery file, IDB backup, original device and transaction logs), the recovery procedure can differ. If all these are available, you can use a very convenient IDB recovery method, guided autorecovery.

Overview of IDB Recovery Methods

Several recovery methods are available for recovering the IDB. Depending on the identified level of corruption, your requirements, and the availability of the IDB recovery file and the original device and transaction logs, the recovery procedure can differ.

17.973
Paula

The Most Convenient Complete Recovery

When the complete IDB is missing or the core part is corrupted, the corruption level is critical. If the IDB recovery file and the original device used for the IDB backup are available, you can perform the Guided Autorecovery (IDB Restore and Replay Logs). Refer to "Performing Guided Autorecovery" on page 421. Otherwise, follow one of the methods given under "More Recovery Methods" on page 418.

The guided autorecovery method guides you through restoring the IDB and replaying transaction logs. If transaction logs are not available, you can still update the IDB by importing all media since the last IDB backup.

Omitting (Removing) Corrupted IDB Parts

If the identified level of corruption is major or minor (corruption is not in the core part), you can consider omitting (removing) the missing or corrupted parts of the IDB or perform the complete IDB recovery instead.

When the filename tablespace is corrupted, the corruption level is major. Refer to "Handling Major Database Corruption in the Filenames Part" on page 423.

When the DC binary files are missing or corrupted, the corruption level is minor. Refer to "Handling Minor Database Corruption in the DCBF Part" on page 422.

More Recovery Methods

These recovery procedures are adapted to specific situations. They assume that you want to recover the complete IDB, but for some reason you cannot perform the guided autorecovery method. The recovery consists of restoring the IDB and updating the IDB.

Table 9-4 **Restoring the IDB**

Current situation	Remark	Recovery Procedure
The IDB recovery file is available but the original device used for the IDB backup has changed.	The method is essentially the same as the guided autorecovery method, but less guided, and more complex and time consuming.	"Recovering the IDB Using IDB Recovery File and Changed Device" on page 424.

1792
Paula.

Table 9-4 **Restoring the IDB**

Current situation	**Remark**	**Recovery Procedure**
The IDB recovery file is not available.	The method is essentially the same as the guided autorecovery method, but less guided, and more complex and time consuming.	"Recovering the IDB Without the IDB Recovery File" on page 426.
You want to recover the IDB from a specific IDB backup (not the latest one).	This method does not provide the latest state of the IDB as a result.	"Recovering the IDB from a Specific IDB Session" on page 428.
You want to recover to a different disk layout.	This method is equivalent to disaster recovery from a Data Protector configuration where you lost the IDB transaction logs, the IDB recovery file, and the `media.log` file. It is far more complex than the guided autorecovery and does not provide the latest state of the IDB as a result.	"Recovering the IDB to a Different Disk Layout" on page 431.

If the transaction logs are available, the recovery procedures in Table 9-4 guide you through replaying the IDB transaction logs. Refer to "Replaying IDB Transaction Logs" on page 430.

If the transaction logs are not available, you can update the IDB by importing media. Refer to "Updating the IDB by Importing Media" on page 433.

Identifying the Level of Database Corruption

IDB Corruption Levels

There are three levels of IDB corruption: critical, major, and minor. The level depends on the part of the IDB where the corruption occurs.

17.971
Paula

You can use the IDB consistency check to determine which part of the IDB is corrupted. Depending on the level of corruption, the IDB recovery procedure differs.

Figure 9-11 **IDB Corruption Levels**

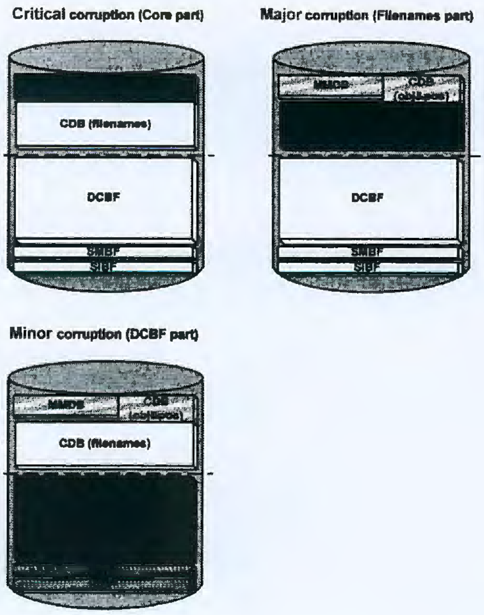
Critical corruption (Core part)
CDB (filenames)
DCBF
SMBF
SIBF
Major corruption (Filenames part)
MMDB
CDB (objects)
DCBF
SMBF
SIBF
Minor corruption (DCBF part)
MMDB
CDB (objects)
CDB (filenames)

How to Identify the Corruption Level

Identify the level of IDB corruption using the `omnidbcheck -extended` command:

NOTE

The extended check may take several hours. To avoid an extended period of system downtime, you can run subparts of the `omnidbcheck` command instead. For example, run the `omnidbcheck -core` to determine whether the core part of the IDB is corrupted.

After identifying the level of corruption, perform the appropriate recovery procedure. Refer to “Overview of IDB Recovery Methods” on page 417.

Performing Guided Autorecovery

Guided autorecovery is the most convenient IDB recovery method. You can perform it if the IDB recovery file and the original device used for the IDB backup together with the IDB backup medium are available.

This method guides you through restoring the IDB and replaying transaction logs since the last IDB backup. If the transaction logs are not available, you can still update the IDB since the last IDB backup by importing media.

Transaction replay updates the core part of the IDB. Binary files are not updated and changes to binary files are lost.

The following are not available for the backups that were running from the last IDB backup before the IDB corruption:

- Session messages
- Browsing of file versions (restores of complete objects are possible). Import the catalog on the media used by the backups to recover the changes.
- SIBF updates. Export and import the media used by the backups to recover the changes.

Prerequisites

Ensure the following before performing guided autorecovery:

- Mount a disk of the same size as before the disaster on the same directories as at the IDB backup time (on Windows systems, the same drive letters must be assigned). If this cannot be ensured, follow the procedure for recovering the IDB to a different disk/volume layout. You can use the `-preview` option of the `omnidbrestore` command to see where the files will be restored.
- Verify that Data Protector is installed on the Cell Manager and the system where a device is attached (preferably, the device used for the IDB backup).
- If the IDB is installed on MC/ServiceGuard, run the `cmhaltpkg` *<pkg_name>* command on the active node before performing the guided autorecovery, to stop the Data Protector package. When the guided autorecovery has finished, run the `cmrunpkg` *<pkg_name>* command on the active node to start the Data Protector package, where *<pkg_name>* is the name of the Data Protector cluster package.

17969
Paula

- If the IDB is installed on Microsoft Cluster Server, take the `OBVS_VELOCIS` cluster group offline using the Cluster Administrator utility on the active node before performing the guided autorecovery. When the guided autorecovery has finished, bring the `OBVS_VELOCIS` and `OBVS_MCRS` cluster groups online using the Cluster Administrator utility.

Recovery Procedure

To recover the IDB, run the `omnidbrestore -autorecover` command.

The command reads the IDB recovery file and if IDB backups are logged to the file, it stops the services and starts restore of the IDB back in place. All the options are generated automatically using data from the IDB recovery file.

Once the restore is complete, the `omnidbrestore` checks if transaction logs are available to be replayed. If logs are available, you are asked to confirm the replay of the logs. If this step is cancelled or transaction logs are not available, output describes how to update the IDB since the last IDB backup by:

- importing media
- finding the transaction logs and replaying them later

Once you replay logs or import media to update the IDB, the full IDB should be successfully recovered.

Handling Minor Database Corruption in the DCBF Part

If you detect that the IDB corruption is of minor severity, it means that some DC binary files are missing or corrupted. If this is the case, there is no need for complete IDB recovery. You can easily recreate the binary files by importing catalog from media. Choose the recovery procedure depending on the corruption type.

Recovering if DC Binary Files Are Missing

DC binary files are organized so that one binary file exists for each medium. If some DC binary files are missing, media positions of some media point to the non-existent files. An error message is displayed when browsing the relevant filesystems. Proceed as follows:

1. From the `omnidbcheck -bf` output, identify the Medium ID of the missing binary file. Run the `omnimm -media_info` *<Medium>* command to get other attributes of the medium, such as medium

label and media pool.

2. Run the `omnidbutil -fixmpos` command to establish consistency between media positions (mpos) and binary files.
3. Import the catalog from the media to recreate the binary files. Refer to "Importing the Catalog from Media" on page 114.

Recovering if DC Binary Files Are Corrupted

If some DC binary files are corrupted, you can remove the DC binary files and recreate them. The only effect of removing the files is that some media positions point to the non-existent binary files, and thus an error message is displayed when browsing the relevant filesystems. Proceed as follows:

1. From the `omnidbcheck -dc` output, identify the Medium ID of the corrupted DC binary file. Run the `omnimm -media_info` *`<Medium>`* command to get other attributes of the medium, such as medium label and media pool.
2. Identify the DC binary file for the affected medium. DC binary files are named: *`<Medium>`*`_`*`<TimeStamp>`*`.dat` (in the *`<Medium>`*, and colons ":" are replaced with underscores "_").
3. Remove the corrupted DC binary files.
4. Run the `omnidbutil -fixmpos` command to establish consistency between media positions (mpos) and binary files.
5. Import the catalog from the media to recreate the binary files. Refer to "Importing the Catalog from Media" on page 114.

Handling Major Database Corruption in the Filenames Part

If you detect that the corruption is of major severity, which means that a filename tablespace is corrupted, you can remove the detail catalogs (filenames and DC binary files) instead of recovering the whole IDB.

The procedure is fast and results in an IDB without detail catalogs (as though all backups were done with the `No log` option). The IDB is still fully operational in terms of all backups, restores, and media management operations, except that browsing is not possible (information about backed up data should be read from media).

Since all detail catalogs are lost, this method of recovery is only applicable if:

17967
Paula

- The catalogs created by subsequent backups are good enough.
- There is no IDB backup available.

Recovery Procedure

Proceed as follows:

1. Run the command:

   ```
   omnidbutil -writedb -no_detail -cdb <Directory> -mmdb
   <Directory>
   ```

   to write the IDB without detail catalogs to ASCII files.

2. Run the command:

   ```
   omnidbutil -readdb -cdb <Directory> -mmdb <Directory>
   ```

   to read the IDB from the ASCII files.

The operation lasts approximately 5-20 minutes.

After the detail catalogs are removed, all DC binary files can be deleted, although the DC directories are still registered. Subsequent backups will store the file versions in the DC binary files.

Recovering the IDB Using IDB Recovery File and Changed Device

Use this procedure to recover the IDB if the IDB recovery file (`obrindex.dat`) is available but the original device used for the IDB backup is different from the one to be used for recovery, or the medium is located in a different slot.

Prerequisites

Ensure the following before performing the database recovery:

- Mount a disk of the same size as before the disaster on the same directories as at the IDB backup time (on Windows systems, the same drive letters must be assigned). If this cannot be ensured, follow the procedure for recovering the IDB to a different disk/volume layout. You can use the `-preview` option of the `omnidbrestore` command to see where the files will be restored.
- If possible, move the `media.log` file from the previous installation to a safe place. It will provide you with the information about the media used since the last IDB backup. This is very helpful for updating the IDB if transaction logs are not available.

17 966
Paula

- Verify that Data Protector is installed on the Cell Manager and the system where a device is attached (preferably, the device used for the IDB backup).
- If the IDB is installed on MC/ServiceGuard, run the `cmhaltpkg` *`<pkg_name>`* command on the active node before performing the guided autorecovery, to stop the Data Protector package. When the guided autorecovery has finished, run the `cmrunpkg` *`<pkg_name>`* command on the active node to start the Data Protector package, where *`<pkg_name>`* is the name of the Data Protector cluster package.
- If the IDB is installed on Microsoft Cluster Server, take the `OBVS_VELOCIS` cluster group offline using the Cluster Administrator utility on the active node before performing the guided autorecovery. When the guided autorecovery has finished, bring the `OBVS_VELOCIS` and `OBVS_MCRS` cluster groups online using the Cluster Administrator utility.

Recovery Procedure

1. Run the following command to create a text file with the restore job options:

```
omnidbrestore -logview -autorecover -skiprestore -save
C:\TEMP\restjob.txt
```

IMPORTANT

The specified `-logview` command lists first transaction logs, next to the session IDs. Remember the first transaction log for the session you want to restore, because you will need it in order to update the IDB after the restore. For example, from the output `2001/02/09-2 AAAAAAH`, you would remember the first transaction log `AAAAAAH` in order to restore the `2001/02/09-2` session.

The created `restjob.txt` file has the information on original devices and on slots in which media were originally located (at IDB backup time).

For example, if the IDB backup was done on a DDS drive with the SCSI address `scsi0:0:0:0`, a file like this is created:

```
-name LDEV

-policy 1

-type 1
```

```
-dev scsi0:0:0:0
-mahost goedl.hermes
-maid 0100007f:3a486bd7:0410:0001
-position 3:0
-daid 977824764
```

2. Modify the `restjob.txt` file to specify the current device or the slot in which the media are currently located.

   For example, if the DDS drive that had the SCSI address `scsi0:0:0:0` at backup time has the SCSI address `scsi0:0:1:0` at restore time, the `restjob.txt` file should be modified accordingly:

   ```
   -name LDEV
   -policy 1
   -type 1
   -dev scsi0:0:1:0
   -mahost cm.dom.com
   -maid 0100007f:3a486bd7:0410:0001
   -position 3:0
   -daid 977824764
   ```

3. Run the restore with the `omnidbrestore -read C:\TEMP\restjob.txt` command.

The command guides you through restoring the IDB and replaying transaction logs since the last IDB backup.

If the transaction logs are not available, you can still update the IDB by importing all media used since the last IDB backup. In this case, refer to "Updating the IDB by Importing Media" on page 433.

Recovering the IDB Without the IDB Recovery File

Use this procedure to recover the IDB if the IDB recovery file (`obrindex.dat`) is not available.

Prerequisites

Ensure the following before performing the database recovery:

- Mount a disk of the same size as before the disaster on the same directories as at the IDB backup time (on Windows systems, the same drive letters must be assigned). If this cannot be ensured, follow the procedure for recovering the IDB to a different disk/volume layout. You can use the `-preview` option of the `omnidbrestore` command to see where the files will be restored.

- If possible, move the `media.log` file from the previous installation to a safe place. It will provide you with the information about the media used since the last IDB backup. This is very helpful for updating the IDB if transaction logs are not available.
- Verify that Data Protector is installed on the Cell Manager and the system where a device is attached (preferably, the device used for the IDB backup).
- If the IDB is installed on MC/ServiceGuard, run the `cmhaltpkg` *`<pkg_name>`* command on the active node before performing the guided autorecovery, to stop the Data Protector package. When the guided autorecovery has finished, run the `cmrunpkg` *`<pkg_name>`* command on the active node to start the Data Protector package, where *`<pkg_name>`* is the name of the Data Protector cluster package.
- If the IDB is installed on Microsoft Cluster Server, take the `OBVS_VELOCIS` cluster group offline using the Cluster Administrator utility on the active node before performing the guided autorecovery. When the guided autorecovery has finished, bring the `OBVS_VELOCIS` and `OBVS_MCRS` cluster groups online using the Cluster Administrator utility.

Recovery Procedure

1. Configure the device using the `Data Protector Manager`.
2. Find the medium with the latest IDB backup.
3. Insert the medium into the device and use the following command to display the contents of the medium:

   ```
   omnimlist -dev <LogicalDevice>
   ```

   The information you need for the IDB restore is the Medium ID and Disk Agent ID for the backup session you want to restore.
4. Use the following command to display the information on the device configuration:

   ```
   omnidownload -dev <LogicalDevice>
   ```

   The information you need for the IDB restore is the following:

17023
Paula

- Mahost (Media Agent host)
- Policy (number)

  A policy number can be obtained using the following translation: 1 for Standalone devices, 3 for Stacker devices, 10 for SCSI-II Libraries, and 5 for Jukebox devices.
- Media type (number)

  A media type number can be obtained using the following translation: 1 for DDS, 3 for ExaByte, 10 for DLT, or 7 for File.
- SCSI address
- Robotics SCSI address (only if using Exchanger library devices)

5. Run the `omnidbrestore` command using the obtained information:

   ```
   omnidbrestore -policy <log. device policy> -type <log.
   device_type> [-ioctl <RoboticsDevice>] -dev <PhysicalDevice>
   -mahost <DeviceHostname> -maid <mediumID> -daid <DAID>
   ```

   For example, you would use the following command to restore the IDB from a backup session with the medium ID 0100007f:3a486bd7:0410:0001 and the Disk Agent ID 977824764, performed using a standalone device of the type DLT, connected to the system `cm.dot.com` and with the SCSI address `scsi0:1:2:0`:

   ```
   omnidbrestore -policy 1 -type 10 -dev scsi0:1:2:0 -mahost
   cm.dom.com -maid 0100007f:3a486bd7:0410:0001 -daid  977824764
   ```

The command guides you through restoring the IDB and replaying transaction logs since the last IDB backup.

If the transaction logs are not available, you can still update the IDB by importing all media used since the last IDB backup. In this case, refer to "Updating the IDB by Importing Media" on page 433.

Recovering the IDB from a Specific IDB Session

Use this procedure to recover the IDB from a backup other than the latest one if the IDB recovery file (`obrindex.dat`) is available.

Prerequisites

Ensure the following before performing the database recovery:

- Mount a disk of the same size as before the disaster on the same directories as at the IDB backup time (on Windows systems, the same drive letters must be assigned). If this cannot be ensured, follow the

17962

procedure for recovering the IDB to a different disk/volume layout. You can use the -preview option of the omnidbrestore command to see where the files will be restored.

- If possible, move the media.log file from the previous installation to a safe place. It will provide you with the information about the media used since the last IDB backup. This is very helpful for updating the IDB if transaction logs are not available.
- Verify that Data Protector is installed on the Cell Manager and the system where a device is attached (preferably, the device used for the IDB backup).
- If the IDB is installed on MC/ServiceGuard, run the cmhaltpkg *<pkg_name>* command on the active node before performing the guided autorecovery, to stop the Data Protector package. When the guided autorecovery has finished, run the cmrunpkg *<pkg_name>* command on the active node to start the Data Protector package, where *<pkg_name>* is the name of the Data Protector cluster package.
- If the IDB is installed on Microsoft Cluster Server, take the OBVS_VELOCIS cluster group offline using the Cluster Administrator utility on the active node before performing the guided autorecovery. When the guided autorecovery has finished, bring the OBVS_VELOCIS and OBVS_MCRS cluster groups online using the Cluster Administrator utility.

Recovery Procedure

1. Check all backups using the following command:

   ```
   omnidbrestore -autorecover -logview -skiprestore
   ```

2. Choose the backup session you want to restore from and perform the restore by running the omnidbrestore -autorecover -session *<sessionID>* command.

   For example, if you choose to restore from the backup session 2000/12/26-1 and the original device used for the IDB backup exists, run:

   ```
   omnidbrestore -autorecover -session 2000/12/26-1
   ```

The command guides you through restoring the IDB and replaying transaction logs since the last IDB backup. If the transaction logs are not available, you can still update the IDB by importing all media used since the last IDB backup. In this case, refer to “Updating the IDB by Importing Media” on page 433.

17461
Paula

Replaying IDB Transaction Logs

In a successful `omnidbrestore -autorecover`, transaction logs are already replayed. Use this procedure only if you need to retry replaying of transaction logs or you postponed it before.

Replaying transaction logs after the IDB restore is completed recovers the IDB to the same state as before the crash, except that binary files are not updated and changes to binary files are lost.

The following are not available for the backups that were running from the last IDB backup until the IDB corruption:

- Session messages.
- Browsing of file versions (restores of complete objects are possible). Perform the import catalog on the media used by the backups, to recover the changes.
- SIBF updates. Export and import the media used by the backups to recover the changes.

Limitation

Replay of the transaction logs can only be done if archiving of the transaction logs is enabled. (The archiving parameter in the `velocis.ini` file must be set to 1.)

Prerequisites

- Transaction logs must be available. For more information on transaction logs, refer to “Preparing for IDB Recovery” on page 390. You can verify that the transaction logs are available by listing the directory: `/db40/logfiles/syslog`

  If transaction logs are not available, refer to “Updating the IDB by Importing Media” on page 433.
- If the IDB is installed on MC/ServiceGuard, run the `cmhaltpkg` *`<pkg_name>`* command on the active node before running the `omnidbrestore` command in the procedure below, to stop the Data Protector package. Before running the `omnidbcheck` command in the procedure below, run the `cmrunpkg` *`<pkg_name>`* command on the active node to start the Data Protector package, where *`<pkg_name>`* is the name of the Data Protector cluster package.
- If the IDB is installed on Microsoft Cluster Server, take the `OBVS_VELOCIS` cluster group offline using the Cluster Administrator utility on the active node before running the omnidbrestore command in the procedure below. Before running the omnidbcheck command in

17.960
Laura

the procedure below, bring the `OBVS_VELOCIS` and `OBVS_MCRS` cluster groups online using the Cluster Administrator utility.

How to Replay Transaction Logs

Proceed as follows:

1. Run the following command to replay the transaction logs:

   ```
   omnidbrestore -replay_only -firstlog
   <FirstTransactionLog>
   ```

   where *<first_trans_log>* is the first transaction log that was created just after the IDB backup was started.

   At the end of the `omnidbrestore -autorecover` output, Data Protector displays the exact command you should use to replay the transaction logs, giving you the name of the first transaction log.

   For example, the command could be:

   ```
   omnidbrestore -replay_only -firstlog AAAAAC
   ```

   where `AAAAAC` is the first transaction log created after the IDB backup was started.

2. Run the `omnidbcheck` command.

This completes the recovery procedure.

Recovering the IDB to a Different Disk Layout

You can restore the IDB to a disk of a different size than before the disaster, and to different directories than at the backup time.

Prerequisites

Ensure the following before recovering the IDB to a different disk layout:

- If possible, store the `media.log` file from the previous installation to a safe place. It will provide you with information about the media used since the last IDB backup.
- Verify that Data Protector is installed on the Cell Manager and the system where a device is attached. (Preferably, this device was used for the IDB backup.)
- Import the media with the IDB backup.

Recovery Procedure

After you meet the prerequisites, proceed as follows to recover the IDB:

1. In the Data Protector Manager, browse the Internal Database backup object and select it for restore. Refer to "Selecting Your Data for Restore" on page 270.
2. For the db40/datafiles directory, use the Restore As/Into option to specify a restore location other than the default one. Refer to "Restoring Files to Different Paths" on page 299.
   You may want to restore the Detail Catalog and Session Messages Binary Files to a different restore location. In this case, also use the Restore As/Into option.
3. Start the IDB restore. Refer to "Previewing and Starting a Restore" on page 273.
4. Move the db40/datafiles directory back in place and start the Data Protector services using the omnisv -start command.
   - On Windows: *<Data_Protector_home>*\bin\omnisv -start
   - On UNIX: /opt/omni/sbin/omnisv -start

   If the IDB is installed on MC/ServiceGuard, run the cmrunpkg *<pkg_name>* command on the active node to start the Data Protector package, where *<pkg_name>* is the name of the Data Protector cluster package.

   If the IDB is installed on Microsoft Cluster Server, bring the OBVS_VELOCIS and OBVS_MCRS cluster groups online using the Cluster Administrator utility.
5. If you restored the Detail Catalog and Session Messages Binary Files to a different restore location, you need to do the following:
   a. Create a new DC directory and remove the old one. Refer to "Creating a DC Directory" on page 396.
   b. Run the omnidbutil -remap_dcdir command to update the pathnames of DC binary files.
6. Verify that you have all files back by running the omnidbcheck command.

What's Next?

After you have restored the IDB, you need to update the IDB by importing media if the media.log file is available. Refer to "Updating the IDB by Importing Media" on page 433.

Updating the IDB by Importing Media

To successfully complete the IDB recovery, you need to update the IDB changes after the IDB is restored.

If transaction logs are not available, update the changes by importing all media since the last IDB backup. Do this once the IDB restore has finished.

To verify that transaction logs are available, or to update the changes using transaction logs, refer to "Replaying IDB Transaction Logs" on page 430.

To update the changes by importing media, proceed as follows:

1. Start the Data Protector processes and services using the `omnisv -start` command:
   - On Windows: `<Data_Protector_home>\bin\omnisv -start`
   - On UNIX: `/opt/omni/sbin/omnisv -start`
2. Increase the session counter to 200 using the following command:

   ```
   omnidbutil -set_session_counter 200
   ```

   If necessary, you can now start with backups.
3. Export and import the media with the last IDB backup. This creates consistent information about the last IDB backup.
4. Import (export if already in IDB) the media used between the last IDB backup and the time of the IDB recovery. See the `/var/opt/omni/log/media.log` (on UNIX systems) or `<Data_Protector_home>\log\media.log` (on Windows systems) file for a list of media.
5. Run the `omnidbcheck` command.

The complete IDB should be successfully recovered.

NOTE If recovering an IDB that encompasses a CMMDB or a remote MMDB to a different disk layout, you need to run the `omnidbutil -cdbsync` command after updating the IDB.

17957
Paula

17.956
Paula

ANEXO UNIDADE DE BACKUP ROBOTIZADO
PARTE D/6

COBRA Tecnologia S.A.
Estrada dos Bandeirantes 7966
CEP 22783-110 Rio de Janeiro RJ
Tel. 21 2442-8800
www.cobra.com.br

17955
Paula

10 Disaster Recovery

In This Chapter

This chapter provides an overview of disaster recovery on Windows UNIX clients and Cell Managers. The following sections are included:

"Introduction" on page 437

"Preparing for a Disaster Recovery" on page 443

"Assisted Manual Disaster Recovery of a Windows System" on page 450

"Disk Delivery Disaster Recovery of a Windows Client" on page 459

"Enhanced Automated Disaster Recovery of a Windows System" on page 463

"One Button Disaster Recovery of a Windows System" on page 472

"Automated System Recovery" on page 480

"Restoring the Data Protector Cell Manager Specifics" on page 487

"Advanced Recovery Tasks" on page 490

"Manual Disaster Recovery of an HP-UX Client" on page 498

"Disk Delivery Disaster Recovery of an UNIX Client" on page 507

"Manual Disaster Recovery of an UNIX Cell Manager" on page 512

"Troubleshooting Disaster Recovery on Windows" on page 514

17053
Paula

Introduction

This section explains the basic terms used in the Disaster Recovery chapter. For overview and concepts of the available disaster recovery methods as well as table outlining the possible combinations of disaster recovery methods and operating system, please see the Disaster Recovery section in the *HP OpenView Storage Data Protector Concepts Guide*.

Table 10-1 **Supported Disaster Recovery Methods and Operating Systems**

	Cell Manager	**Client**
Windows NT/2000	• “Assisted Manual Disaster Recovery of a Windows System” on page 450 • “Enhanced Automated Disaster Recovery of a Windows System” on page 463 • “One Button Disaster Recovery of a Windows System” on page 472	• “Assisted Manual Disaster Recovery of a Windows System” on page 450 • “Disk Delivery Disaster Recovery of a Windows Client” on page 459 • “Enhanced Automated Disaster Recovery of a Windows System” on page 463 • “One Button Disaster Recovery of a Windows System” on page 472

17952
Paula

Table 10-1 **Supported Disaster Recovery Methods and Operating Systems**

	Cell Manager	**Client**
32-bit Windows XP[a]/Server 2003	• "Assisted Manual Disaster Recovery of a Windows System" on page 450 • "Automated System Recovery" on page 480	• "Assisted Manual Disaster Recovery of a Windows System" on page 450 • "Disk Delivery Disaster Recovery of a Windows Client" on page 459 • "Automated System Recovery" on page 480
64-bit Windows XP/Server 2003		• "Assisted Manual Disaster Recovery of a Windows System" on page 450 • "Automated System Recovery" on page 480
HP UX 11.x	• "Manual Disaster Recovery of an UNIX Cell Manager" on page 512	• "Manual Disaster Recovery of an HP-UX Client" on page 498 • "Disk Delivery Disaster Recovery of an UNIX Client" on page 507
Solaris 7/8	• "Manual Disaster Recovery of an UNIX Cell Manager" on page 512	• "Disk Delivery Disaster Recovery of an UNIX Client" on page 507

17951 Paula

Table 10-1 **Supported Disaster Recovery Methods and Operating Systems**

	Cell Manager	Client
Tru64/AIX		• "Disk Delivery Disaster Recovery of an UNIX Client" on page 507

a. ASR is not available on Windows XP Home Edition, therefore it is not supported.

What Is a Computer Disaster?

A **computer disaster** refers to any event that renders a computer system unbootable, whether due to human error, hardware or software failure, virus, natural disaster, etc. In these cases it is most likely that the boot or system partition of the system is not available and the environment needs to be recovered before the standard restore operation can begin. This includes repartitioning and/or reformatting the boot partition and recovery of the operating system with all the configuration information that defines the environment. *This has to be completed in order to recover other user data.*

What Is an Original System?

Original system refers to the system configuration backed up by Data Protector before a computer disaster hit the system.

What Is a Target System?

Target system refers to the system after the computer disaster has occurred. The target system is typically in a non-bootable state and the goal of Data Protector disaster recovery is to restore this system to the original system configuration. The difference between the crashed and the target system is that the target system has all faulty hardware replaced.

What Are Boot and System Disks/Partitions/ Volumes?

A **boot disk/partition/volume** refers to the disk/partition/volume that contains the files required for the initial step of the boot process, whereas the **system disk/partition/volume** refers to the disk/partition/volume that contains the operating system files.

NOTE

Microsoft defines the boot partition as the partition that contains the operating system files and the system partition as one that contains the files required for the initial step of the boot process.

17450
Paula

What Is a Hosting System?

Hosting system is a working Data Protector client used for Disk Delivery Disaster Recovery with Disk Agent installed.

What Is Auxiliary Disk?

Auxiliary disk is a bootable disk that has a minimal operating system with networking and Data Protector Disk Agent installed. It can be carried around and used to boot the target system in Phase 1 of Disk Delivery Disaster Recovery of UNIX clients.

What Is a Disaster Recovery Operating System (DR OS)?

Disaster recovery operating system (DR OS) is operating system environment where the process of disaster recovery is running. It provides Data Protector a basic runtime environment (disk, network, tape and filesystem access). It has to be installed and configured before the Data Protector disaster recovery can be performed.

DR OS can be either temporary or active. **Temporary DR OS** is used exclusively as a host environment for some other operating system restore along with the target operating system configuration data. It is deleted after the target system is restored to the original system configuration. **Active DR OS** not only hosts the Data Protector disaster recovery process but is also a part of the restored system because it replaces it's own configuration data with the original configuration data.

What Are Critical Volumes?

Critical volumes are volumes required to boot the system and Data Protector files. Regardless of the operating system, these volumes are:

- boot volume
- system volume
- Data Protector executables
- IDB (Cell Manager only)

NOTE

If IDB is located on different volumes than all volumes where IDB resides, are critical.

Apart from the critical volumes stated above, `CONFIGURATION` is also a part of the critical volumes set for Windows systems. Services are backed up as a part of the `CONFIGURATION` backup.

Some items included in the `CONFIGURATION` can be located on volumes other than system, boot, Data Protector or IDB volumes. In this case these volumes are also part of critical volumes set:

17.449

- user profiles volume
- Certificate Server database volume on Windows Server
- Active Directory Service volume on domain controller on Windows Server
- quorum volume on Microsoft Cluster Server.

What is Online Recovery?

Online recovery is performed when Cell Manager is accessible. In this case most of Data Protector functionalities are available (Cell Manager runs the session, restore sessions are logged in the IDB, you can monitor the restore progress using GUI, etc.).

What is Offline Recovery?

Offline recovery is performed if the Cell Manager is not accessible (for example, due to network problems, Cell Manager has experienced a disaster, online recovery has failed, etc.). Only standalone and SCSI-II Library devices can be used for offline recovery. Note that recovery of Cell Manager is always offline.

What is Local/Remote Recovery?

Remote recovery is performed if all Media Agent hosts specified in SRD file are accessible. If any of them fails, disaster recovery process fails over to local mode. This means that the target system is searched for locally attached devices. If only one device is found, it is automatically used. Otherwise Data Protector prompts you to select the device which will be used for restore. Note that offline OBDR is always local.

Disaster is always serious, however the following factors can exacerbate the situation:

- The system has to be returned to online status as quickly and efficiently as possible.
- Administrators are not familiar with the required steps to perform the disaster recovery procedure.
- The available personnel to perform the recovery may only have fundamental system knowledge.

Disaster recovery is a complex task that involves extensive planning and preparation before execution. You have to have a well-defined, step-by-step process in place to prepare for, and recover from, disastrous situations.

17945
Paula

The Recovery Process

The disaster recovery process consists of 4 phases with the *Phase 0* (preparation) being the prerequisite for a successful disaster recovery. In *Phase 1*, DR OS is installed and configured, which usually includes repartitioning and reformatting of the boot partition, since the boot or system partition of the system are not always available and the environment needs to be recovered before normal restore operations can resume. Operating system with all the configuration information that defines the environment with Data Protector (as it was) is restored in *Phase 2*. Only after this step is completed, is the restore of applications and user data possible (*Phase 3*). A well-defined, step-by-step process has to be followed to ensure fast and efficient restore.

17947
Paula

Preparing for a Disaster Recovery

Carefully follow the instructions in this section to prepare for a disaster recovery and to ensure fast and efficient restore. Preparation does not depend on the disaster recovery method, however, it does include developing a detailed disaster recovery plan, performing consistent and relevant backups and updating the SRD file on Windows.

Planning

Developing a detailed disaster recovery plan has a major impact on the success of a disaster recovery. To deploy disaster recovery in a large environment with many different systems, proceed as follows:

1. **Plan**

   Planning must be prepared by IT administration and should include the following:

   - Determine the systems that need to be recovered as well as the time and level of recovery. Critical systems are all systems required for network to function properly (DNS servers, domain controllers, gateways, etc.), Cell Managers and Media Agent clients.
   - Determine a recovery method to be used (impacts the required preparations).
   - Determine a method to obtain the required information at recovery time, such as the media that holds the IDB, location of updated SRD file and location and labels of Cell Manager backup media.
   - Create a step-by-step detailed checklist to guide you through the process.
   - Create and execute a test plan to confirm that the recovery will actually work.

2. **Prepare for recovery**

   Depending on the recovery method to be used, the preparation should include:

17946

On UNIX systems:

- Creation of tools, such as the auxiliary disk with the minimum operating system, network resources, and the Data Protector Disk Agent installed.
- Creation of pre-execution scripts, which collect the storage structure and other client-specific preparations.

On Windows systems:

- Updating **System Recovery Data (SRD)** and storing it to a safe place. You should restrict access to SRD files due to security reasons.

On all systems:

- Performing regular and consistent backups.

3. **Perform recovery procedures**

   Follow the procedures and checklists you have tested to recover the crashed system.

Consistent and Relevant Backup

In the case of a disaster, the target system should be put back into the state it was at the time of the last valid known backup. Additionally, the system should function as it had functioned just before the last valid backup performance.

NOTE

On UNIX systems, some daemons or processes are active as soon as the system finishes booting, for various reasons (HP-UX example: License server at run level-2). Such an early process may even read the data into memory and write a “dirty flag” into some file while it runs. A backup taken at the standard operating stage (the standard run level-4) cannot be expected to yield a problem-free restart of such an application. To follow the example, the license server, if started after such a pseudo recovery, will realize that the data read from the file is inconsistent and will refuse to run the service as expected.

Paula

On Windows, while the system is up and running, many system files cannot be replaced because the system keeps them locked. For example, the user profiles that are currently being used cannot be restored. The login account has to be changed or the relevant service has to be stopped.

Data consistency of an application can be violated depending on what is active on the system when the backup runs, thereby causing re-start and execution issues after recovery.

How to Create a Consistent and Relevant Backup?

- ✓ Ideally, you would perform a backup with the relevant partition(s) set off-line, which is usually not possible.
- ✓ Examine the activity on the system during the backup. Only operating system related processes and database services which are backed up online can remain active during the backup execution.
- ✓ None of the low-level (UNIX) or background-level (Windows) application specific services should be running.

Updating the System Recovery Data (SRD)

What Is SRD?

System recovery data (SRD) is a Unicode text file that contains information required for the configuration and restore of the Windows target system. A SRD file is generated when `CONFIGURATION` backup is performed on a Windows client and then stored in *`<Data_Protector_home>`*`\Config\dr\srd` (Windows Cell Manager) or in `/etc/opt/omni/dr/srd/` (UNIX Cell Manager).

IMPORTANT When IDB is not available, information about objects and media is stored only in SRD file.

The SRD filename on the Cell Manager is identical to the hostname of the computer where it was generated - for example `computer.company.com`.

After the `CONFIGURATION` backup, the SRD contains only system information required for installation of the DR OS. In order to perform a disaster recovery, additional information about backup objects and

17 ayy
Paula

corresponding media must be added to the SRD. The SRD can be updated only on a Windows client. The name of the updated SRD file is `recovery.srd`.

How to Update SRD?

There are three different methods possible for updating the SRD file:

- Update SRD File Wizard
- `omnisrdupdate` command as a standalone utility
- `omnisrdupdate` command as a backup session post-exec script

Using SRD Update Wizard

To update the SRD file using the Update SRD File Wizard:, proceed as follows:

1. In the `Data Protector Manager` switch to the `Restore` context and then click the `Tasks` Navigation tab.
2. In the `Scoping Pane` of the `Tasks` Navigation tab, check the `Disaster Recovery`.
3. In the `Results Area`, check the `SRD File Update` option button, select the client and click `Next`.
4. For each of the critical objects, select an object version and click `Next`.
5. Type the destination directory where the updated SRD file is to be placed and click `Finish`.

IMPORTANT

Because the SRD file is saved on the Cell Manager system, it is not accessible if the Cell Manager fails. As a result, you need an additional copy of the Cell Manager's SRD which should be stored in a vault.
In addition to the Cell Manager, you should save the updated SRD file to several secure locations as a part of the disaster recovery preparation policy. See "Preparation" on page 451.

Using omnisrdupdate

It is also possible to update the SRD file using the `omnisrdupdate` command as a standalone command. The `omnisrdupdate` command is located in the `<Data_Protector_home>\bin` directory.

17.943
Paula

`Omnisrdupdate` requires a `session_ID` to update an existing SRD file with backup object information belonging to the given session. Using this value, `omnisrdupdate` will update the SRD file with the backup object information which belongs to the passed `session_ID` value. After the SRD is updated it will be saved back on the Cell Manager.

This procedure will only succeed if all critical backup objects (as specified in the SRD file) were actually backed up during the specified session. To view which objects are considered as critical for the SRD update, open the SRD file in a text editor and find the objects section. All critical objects for the SRD update are listed there. Note that the database is represented as "/".

Here is an example of an objects section of the SRD file:

```
-section objects

-objcount 3

-object /C -objtype 6 -objpurpose 283

-endobject /C

-object / -objtype 3 -objpurpose 32

-endobject /

-object /CONFIGURATION -objtype 6 -objpurpose 4

-endobject /CONFIGURATION

-endsection objects
```

In this case, there are 3 critical objects: `/C`, / (database) and `/CONFIGURATION`.

TIP

To obtain the session ID, execute the `omnidb` command with the option `-session`. To obtain the latest session ID, at the command prompt type **`omnidb -session -latest`**.

The updated SRD file should be kept in a safe place so that it is not lost in the case of disaster. To locate where the updated SRD file will be saved, use the `-location` option with the `omnisrdupdate` command. There can be more than one `-location` parameters specified (including network shares on which you have write permission), each of which will receive an updated copy of the SRD file. See "Preparation" on page 451.

17442
Paula

To determine for which hostname the SRD file from the Cell Manager should be updated, use the option `-host` with the command `omnisrdupdate`. If you don't specify the hostname, the local host is assumed. SRD file on the Cell Manager is not updated.

Example

To update the SRD file with the backup object information which belongs to a session `2002/05/02-5` for the client with the hostname `computer.company.com` and to store an updated copy of the SRD file on the floppy disk and in the `SRDfiles` share on computer with the hostname `computer2`, type **`omnisrdupdate -session 2002/05/02-5 -host computer.company.com -location a: -location \\computer2\SRDfiles`**

Make sure that you have the write permission on that share.

Using a Post-Exec Script

Another method to update the SRD is using the `omnisrdupdate` command as a backup post-exec script. To do so you have to either modify an existing backup specification or create a new one. Perform the following steps to modify a backup specification so that the SRD file is updated with information about backed up objects when the backup session stops:

1. In the `Backup` context, expand the `Backup Specifications` item and then `Filesystem`.
2. Select the backup specification that you would like to modify (it must include all backup objects marked as critical in the SRD file, otherwise the update will fail. It is recommended to perform the client backup with disk discovery) and click `Options` in the `Results Area`.
3. Click the `Advanced` button under the `Backup Specification Options`.
4. Type **`omnisrdupdate.exe`** in the `post-exec` text box.
5. In the `On client` drop down list, select the client on which this post-exec script will be executed and confirm with OK. This should be the client that was marked for backup on the source page.

When `omnisrdupdate` command is executed as a post-exec utility, the session ID is obtained automatically from the environment and the user is not required to specify the session ID.

All other options can be specified the same way as with the standalone utility (`-location <`*`path`*`>, -host <`*`name`*`>`).

17411
Paula

IMPORTANT You should restrict access to SRD files due to security reasons.

17.940
Paula

Assisted Manual Disaster Recovery of a Windows System

The following sections explain how to prepare and execute an Assisted Manual Disaster Recovery on Windows systems. For details on supported operating systems, refer to the *HP OpenView Storage Data Protector Software Release Notes*.

Assisted Manual Disaster Recovery is an elementary method that consists of the following steps:

1. Installing the Windows NT operating system temporarily (temporary DR OS) or installing other Windows systems to its original location (active DR OS). This includes the creation and formatting of the boot and system partition, needed for the Windows installation.
2. Creating and formatting additional partitions as they existed on the crashed system, including original drive letter assignments.
3. Executing the Data Protector `drstart.exe` command, which will install a temporary Data Protector suite and start the restore of the system critical volumes.
4. Booting the system and deleting the Windows NT temporary installation.
5. Recovering the vendor-specific partition, if it existed before the disaster.

NOTE The preparation and recovery procedure are different for the recovery of a Data Protector client and of a Data Protector Cell Manager. The differences are marked in the text.

Note that Windows provide additional possibilities to recover a system before deciding on a disaster recovery. This can be done by booting the system in the safe mode or from the recovery floppy disks and trying to resolve problems. Another option is to start the computer using the last known good configuration.

17939
Paula

Requirements

- The partitions have to be the same size or larger than the partitions on the failed disk. This way the information stored on the failed disk can be restored to the new one. Also, the type of filesystem and compression attributes of the volumes must match (FAT, NTFS).
- The hardware configuration of the target system must be the same as of the original system. This includes SCSI BIOS settings (sector remapping).

Limitation

- Internet Information Server (IIS) Database, Terminal Services Database and Certificate Server Database are not restored automatically during Phase 2. They can be restored on the target system using the standard Data Protector restore procedure.
- Boot and system partition on Windows NT 4.0 must be physically bellow the first 7,8 GB due to the operating system limitations. Refer to MSDN Q224526 for more information and workaround for the problem.

Preparation

To prepare for a successful disaster recovery, you should follow the instructions related to the general preparation procedure together with the specific method requirements. Advance preparation is essential to perform the disaster recovery fast and efficiently. You should also give special attention to the disaster recovery preparation of the Cell Manager.

WARNING **It is too late to prepare for a disaster recovery once a disaster has occurred.**

See also “Preparing for a Disaster Recovery” on page 443, in addition to completing the steps listed in this section. To recover from a disaster quickly and efficiently, consider the following steps and prepare your environment accordingly:

1. You need a Windows bootable installation CD-ROM to enable your system to start from the CD-ROM. If you do not have a bootable CD-ROM, use the standard procedure for booting the computer from diskettes.
2. Ensure that you have drivers for the system you want to recover. You may need to install some drivers, such as network, HBA and SCSI drivers during Windows Setup.
3. To recover the crashed system, you need the following information about the system before the disaster (stored also in the SRD file):
   - If DHCP was not used before the disaster, the TCP/IP properties (IP address, Default gateway, Subnet mask and DNS order)
   - Client properties (Hostname)
4. Ensure that the following is true:
   - You should have a successful full client backup. See "Backing Up Filesystems (Logical Disk Drives)" on page 168 and "Backing Up CONFIGURATION" on page 173.
   - You should have a SRD file updated with information about backed up objects in the chosen successful backup session. See "Updating the System Recovery Data (SRD)" on page 445.
   - In the case of a Cell Manager recovery, you need a successful IDB backup of the Cell Manager. Refer to "Preparing for IDB Recovery" on page 390 for more information on how to perform a IDB backup.
   - The disk with the boot partition requires free disk space that is needed for:
     - ✓ On Windows NT: the Data Protector disaster recovery installation (15 MB) and temporary DR OS installation (150 MB). Additionally, you need as much free space, as required for the restore of the original system. If you had applied the Compress Drive on the original partition, its size will be doubled when restored.
     - ✓ On other Windows systems: the Data Protector disaster recovery installation (15 MB) and active DR OS installation. Additionally, you also need as much free space, as required for the restore of the original system.
5. Copy the contents of *`<Data_Protector_home>`*`\Depot\DRSetup` or

17937
Paula

`\i386\tools\DRSetup` (located on Data Protector installation medium) for 32 bit Windows Client or Cell Manager on two floppy disks (**drsetup diskettes**) or `<Data_Protector_home>\Depot\DRSetup64` or `\i386\tools\DRSetup64` (Data Protector installation medium) for 64 bit Windows systems on three floppy disks. In case of a disaster, save the updated SRD file of the crashed client to the first floppy disk (disk1). Only one set of drsetup diskettes is required per site for all Windows systems, but you must always copy an updated SRD file of the crashed client on the first floppy disk. If multiple SRD files are found, Data Protector will ask you to select the appropriate version.

6. In order to re-create disk partitions to their initial state prior to the crash, record the following information for each partition (it will be needed during the recovery process):

   - partitions length and order
   - drive letters assigned to the partitions
   - partitions filesystem type

   This information is stored in the SRD file. The `-type` option in the diskinfo section of the SRD file shows the partition filesystem type for a particular partition:

Table 10-2 **How to Determine the Filesystem Type from the SRD File**

Type number	Filesystem
1	Fat12
4 and 6	Fat32
5 and 15	Extended partition
7	NTFS
11 and 12	Fat32
18	EISA
66	LDM partition

17.930

The table on the next page is an example of the preparation for the disaster recovery. Note that data in the table belongs to a specific system and cannot be used on any other system. Refer to the Appendix A, "Windows Manual Disaster Recovery Preparation Template," on page A-49 for an empty template which can be used when preparing for the Assisted Manual Disaster Recovery.

17935
Fauk

Table 10-3

client properties	computer name	ANDES
	hostname	andes.company.com
drivers		hpn.sys, hpncin.dll
Windows Service Pack		Windows NT SP6
TCP/IP properties	IP address	3.55.61.61
	default gateway	10.17.250.250
	subnet mask	255.255.0.0
	DNS order	11.17.3.108, 11.17.100.100
medium label / barcode number		"andes - disaster recovery" / [000577]
partition information and order	1st disk label	
	1st partition length	31 MB
	1st drive letter	
	1st filesystem	EISA
	2nd disk label	BOOT
	2nd partition length	1419 MB
	2nd drive letter	C:
	2nd filesystem	NTFS/HPFS
	3rd disk label	
	3rd partition length	
	3rd drive letter	
	3rd filesystem	

17.034

Recovery

Follow the procedure below to recover a Windows system using Assisted Manual Disaster Recovery. If you are performing advanced recovery tasks (such as disaster recovery of a Cell Manager or IIS), see also "Advanced Recovery Tasks" on page 490.

1. Install the Windows system from the CD-ROM and install additional drivers if needed. The Windows operating system has to be installed on the same partition as prior to the disaster. Do not install the Internet Information Server (IIS) during the installation of the system. Refer to "Restoring Internet Information Server (IIS) Specifics" on page 496 for more details.

IMPORTANT If Windows has been installed using the Windows unattended setup, use the same script now to install Windows to ensure that the `<$SystemRoot$>` and `\Documents and Settings` folders are installed to the same position.

2. When the Windows Partition Setup screen appears, proceed as follows:

   - If an vendor-specific partition (e.g. EISA Utility Partition) existed on the system before the crash, create (if it does not exist due to the crash) and format a "dummy" FAT partition using the EUP information gathered from the SRD file. The EUP will be later on recovered to the space occupied by the "dummy" partition. Create and format a boot partition immediately after the "dummy" partition. To do this, you need the data as described in "Preparation" on page 451.
   - If an EUP did not exist on the system before the crash, create (if the boot partition does not exist due to the crash) and format the boot partition as it existed on the disk before the crash. To do this, you need the data as described in "Preparation" on page 451.

Windows NT When the Windows NT setup prompts you for the Windows NT installation directory, specify *any new directory on the boot partition* that is not the location where the original Windows NT installation resided (for example, `DPWINNT`). A new directory has to be specified because temporary DR OS is used to recover Windows NT.

17933
Paula

Other Windows Systems

If you are recovering a Windows system other than Windows NT, install Windows into its original location, i.e. the same drive letter and directory as in the original system before the disaster. This information is stored in the SRD file.

NOTE

During the installation, do not add the system to the previous location where the Windows domain resided, but add the system to a workgroup instead.

3. Install TCP/IP protocol. If DHCP was not used before the disaster, configure the TCP/IP protocol as prior to the disaster by providing the following information: hostname of the crashed client, its IP address, default gateway, subnet mask and DNS server. Make sure that the field labeled `Primary DNS suffix of this computer` contains your domain name

WARNING

By default, Windows 2000/XP/Server 2003 install the Dynamic Host Configuration Protocol (DHCP) during the Windows 2000/XP/Server 2003 setup.

4. Create a temporary disaster recovery account in the Administrators group. See "Adding or Deleting a User" on page 90. Note that the account must not had existed on the system before the disaster and that it will be removed at a later time during this procedure.
5. Log off and log in to the system using the newly created account.
6. If you are recovering a Windows NT system, install SP4 or later. No service packs are required for a successful disaster recovery of other Windows systems.
7. Execute the `drstart.exe` command from the *`<Data_Protector_home>`*`\Depot\drsetup\Disk1` (Windows Cell Manager) or `\i386\tools\drsetup\Disk1` (Data Protector installation medium) directories.
   If you have prepared the drsetup diskettes (see "Preparation" on page 451), you can also execute the `drstart.exe` command from the first diskette.
8. `Drstart.exe` first scans the current working directory, floppy and

17932
Pauk

CD drives for the location of disaster recovery setup files (`Dr1.cab` and `omnicab.ini`). If the required files are found, the `drstart` utility installs the disaster recovery files in the *`<%SystemRoot%>\system32\OB2DR`* directory. Otherwise enter their path in the `DR Installation Source` text box or browse for the files.

9. If the `recovery.srd` file is saved in the same directory as `dr1.cab` and `omnicab.ini` files, then `drstart.exe` copies `recovery.srd` file to the `<%SystemRoot%>\system32\OB2DR\bin` directory and the `omnidr` utility is started automatically. Otherwise, you can enter the location of SRD file (`recovery.srd`) in the `SRD Path` field or browse for the file and click `Next`.

   If multiple SRD files are found on the floppy disk, Data Protector will ask you to select an appropriate version of the SRD file.

   After `omnidr` successfully finishes, all critical objects required for a proper boot of the system are restored.

10. Reboot the computer, log on and verify that the restored applications are running.

11. If you are recovering a Cell Manager, perform the procedure described in "Restoring the Data Protector Cell Manager Specifics" on page 487

12. Use Data Protector to restore user and application data.

The temporary DR OS will be deleted after the first login except in the following cases:

- You have interrupted the Disaster Recovery Wizard during the 10 seconds pause after it has found the DR installation and SRD file on the backup medium, and have selected the `Use Debugs` option.
- You have manually started the `omnidr` command with the `no_reset` or `debug` options.
- Disaster recovery fails.

17.931
Paula

Disk Delivery Disaster Recovery of a Windows Client

To perform the Disk Delivery Disaster Recovery, use a working Data Protector client (Data Protector disaster recovery host) to create the new disk while connected to this client. The administrator has to ensure before the disaster that enough data is collected to correctly format and partition the disk. However, Data Protector automatically stores the relevant information as part of the configuration backup.

The recovered partitions are:

- the boot partition
- the system partition
- the partitions containing Data Protector

Any remaining partitions can be recovered by using the standard Data Protector recovery procedure.

For details on supported operating systems, refer to the *HP OpenView Storage Data Protector Software Release Notes*.

TIP

This method is specially useful with hot swap hard disk drives, because you can disconnect a hard disk drive from a system and connect a new one while the power is still on and the system is operating.

Requirements

- The partitions have to be the same size or larger than the partitions on the failed disk. This way the information stored on the failed disk can be restored to the new one. Also, the type of filesystem format has to match (FAT, NTFS).
- The system on which the disk is created and the system in which the disk is used have to use the same sector mapping/addressing (SCSI BIOS enabled/disabled; EIDE: both systems have to use the same addressing mode: LBA, ECHS, CHS).

Fl 930
Paula

Limitations

- Disk Delivery Disaster Recovery is not supported for Microsoft Cluster Server.
- RAID is not supported. This includes software RAIDs (fault-tolerant volumes and dynamic disks).
- Internet Information Server (IIS) Database, Terminal Services Database and Certificate Server Database are not restored automatically during Phase 2. They can be restored on the target system using the standard Data Protector restore procedure.
- Boot and system partition on Windows NT 4.0 must be physically bellow the first 7,8 GB due to the operating system limitations. Refer to MSDN Q224526.

Preparation

Complete a few steps in order to prepare for disaster recovery. Read and follow the section, "Preparing for a Disaster Recovery," in addition to completing the steps listed in this section.

IMPORTANT You have to prepare for disaster recovery *before* a disaster occurs.

In order to recover from a disaster quickly, efficiently and effectively, you need the following:

- The last valid known full backup of the client that you want to recover.
- A new hard disk to replace your crashed disk.
- A Data Protector hosting system, which has to be of the same operating system as the crashed client and must have the same hardware I/O path required to connect the new disk.

In order to re-create disk partitions to their initial state prior to the crash, record the following information for each partition (it will be needed during the recovery process):

- partitions length and order
- drive letters assigned to the partitions

- partitions filesystem type

You can refer to Table 10-3 on page 455 as an example of the preparation for the Disk Delivery disaster recovery. Refer to the Appendix A, "Windows Manual Disaster Recovery Preparation Template," on page A-49 for an empty template which can be used when preparing for the Disaster Recovery.

Recovery

This section provides the procedure for recovering your Windows client using the Disk Delivery method. See also "Advanced Recovery Tasks" on page 490

With the Disk Delivery method on Windows, use a Data Protector disaster recovery host (DR host) to restore the last valid known full backup of your crashed disk to a new hard disk connected to the client. Then replace your crashed disk on the faulty system with this new hard disk.

Disk Delivery Disaster Recovery Procedure

The actual Disk Delivery Disaster Recovery procedure consists of the following steps:

1. Connect the new disk to a DR host.
2. Reboot the DR host to recognize the new disk.
3. Use Data Protector GUI on disaster recovery host and switch to the `Restore` context and click the `Tasks` tab. Select the `Disaster Recovery` item in the Scoping Pane, select the client from the drop down list and check the `Disaster recovery with disk delivery` in the Results Area.
4. For each of the critical objects, select an object version that will be restored and click `Next`.
5. If partitioning has not already been done, partition the new disk using the Disk Administrator. Use the partition information you have gathered as part of the preparation for Disk Delivery disaster recovery.
6. When partitioning the system, you have to assign partitions in the same order as prior to the time that the full backup was performed. This simplifies drive letter reassignment after the restore and prevents a possibility of failure at system restart because of an inappropriate path to the system partition in the `boot.ini` file.

17928
Paula

IMPORTANT You have to assign drive letters for Windows 2000/XP/Server 2003 mountpoints. In this case you must have enough unassigned drive letter available in order to be able to assign a drive letter for each mount point.

7. Perform all necessary drive letter mappings by right clicking on the original drive letter. This is necessary because drive letters on hosting and original system can be different.
8. Press `Finish`.
9. Remove the new disk from the DR host, and then connect it to the target system.
10. Power on the target system.
11. Use the standard Data Protector restore procedure to restore user and application data. This completes the recovery of the client.

Disk Delivery can also be a valuable method in case one of disks in a multi boot system has crashed, and the user can still boot at least one configuration.

NOTE Data Protector does not restore volume-compression flag after recovery. All files, that were compressed at backup time, will be restored as compressed but you will have to manually set volume compression if you want any new files created to be compressed as well.

17207
Paula

Enhanced Automated Disaster Recovery of a Windows System

Enhanced Automated Disaster Recovery (EADR) is a fully automated Data Protector recovery method for Windows clients and Cell Manager, where user intervention is reduced to minimum. For details on supported operating systems, refer to the *HP OpenView Storage Data Protector Software Release Notes*.

The EADR procedure for Windows platforms collects all relevant environment data automatically at backup time. During configuration backup, data required for temporary DR OS setup and configuration is packed in a single large **DR OS image file** and stored on the backup tape (and optionally on Cell Manager) for each backed up client in the cell.

In addition to this image file, a Phase 1 startup information (stored in the **P1S** file), required for correct formatting and partitioning of the disk is stored on the Cell Manager. When a disaster occurs, EADR Wizard is used to restore the DR OS image from the backup medium (if it has not been saved on the Cell Manager during the full backup) and convert it to a **disaster recovery CD ISO image**. CD ISO image can then be burned on a CD using any burning tool and used to boot the target system.

Data Protector then automatically installs and configures DR OS, formats and partitions the disks and finally recovers the original system with Data Protector as it was at the time of backup.

IMPORTANT You have to perform a new backup and prepare a new DR CD after each hardware, software or configuration change. This also applies to any network configuration changes, such as change of IP address or DNS server.

The recovered volumes are:

- the boot partition
- the system partition
- the partitions containing Data Protector

Any remaining partitions can be recovered using the standard Data Protector recovery procedure.

The following sections explain the limitations, preparation, and recovery that pertains to EADR of the Windows clients. See also "Advanced Recovery Tasks" on page 490.

Before selecting this method of disaster recovery, consider the following requirements and limitations:

Requirements

- The Data Protector Automatic Disaster Recovery component must be installed on clients for which you want to enable recovery using this method and on the system, where the DR CD ISO image will be prepared. See *HP OpenView Storage Data Protector Installation and Licensing Guide*.
- The hardware configuration of the target system must be the same as of the original system. This includes SCSI BIOS settings (sector remapping).
- Replacement disks have to be attached to the same host bus adapter on the same bus.
- Boot partition has to be larger than 100 MB or disaster recovery will fail.
- An additional 200 MB of free disk space is required on the boot partition at backup time. If this disk space is not available, the disaster recovery fails. If you had applied the Compress Drive on the original partition, you must have 400 MB free.
- All drivers required for boot must be installed under `<%SystemRoot%>` folder.
- Network must be available when you boot the system in Safe Mode with Networking or in Directory Services Restore Mode (Domain Controller only), but you must do the backup of the system after it was booted with normal boot process.
- The system's BIOS must support bootable CD extensions as defined in the El-Torito standard and read/write access to hard disk drive using LBA addressing via INT13h function XXh. The BIOS options can either be checked in the user's manuals of the system or by inspecting the system setup before the boot.

17.2005
Paula

- When backing up the client, the default 64 kB block size should be used to write to the device if you plan to perform an offline restore. This is the only default block size available on Windows when performing disaster recovery. To verify that the default 64 kB block size is set, choose `Advanced` in the `Properties` box, as shown in Figure 10-1:

Figure 10-1 **Verifying the Default Block Size**

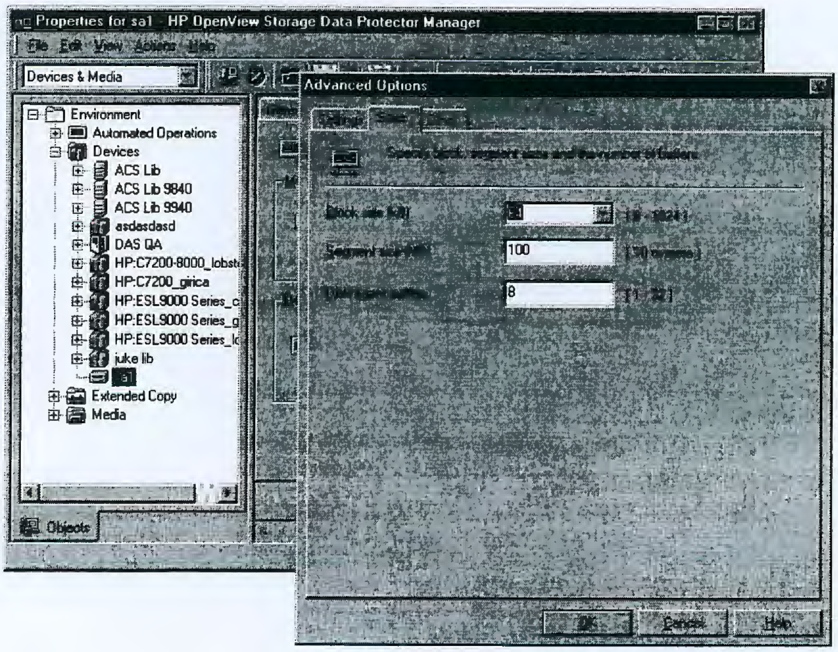

Block size (kB) [illegible]
Segment size (MB) 100
[illegible] 8

Limitations

General

- The disaster recovery CD for a Windows 2000 client or Cell Manager should be created on a Windows 2000 system.
- Multiboot systems that do not use Microsoft's boot loader are not supported.
- Internet Information Server (IIS), Terminal Services Database and

17.924
Paula

Certificate Server Database are not restored automatically during Phase 2. They can be restored on the target system using the standard Data Protector restore procedure.

Disk and Partition Configuration

- With fault-tolerant disk drives on the OS level for Windows NT systems, mirror set is supported while stripe and volume set are not. Dynamic disks are not supported on Windows 2000 (including mirror set upgraded from Windows NT).
- New disk has to be the same size or bigger than the crashed disk. If it is larger than the original disk, the difference will remain unallocated.
- Boot and system partition on Windows NT 4.0 must be physically bellow the first 7,8 GB due to the operating system limitations. Refer to MSDN Q224526.
- Only vendor specific partitions of type 0x12 (including EISA) and 0xFE are supported for Enhanced Automated Disaster Recovery.

Preparation

Complete the steps described in the section “Preparing for a Disaster Recovery” on page 443 in order to prepare for disaster recovery in addition to completing the steps listed in this section. See also “Advanced Recovery Tasks” on page 490.

IMPORTANT You have to prepare for disaster recovery *before* a disaster occurs.

Prerequisite

Full client backup (including the configuration) is prerequisite for successful EADR. See “Backing Up Filesystems (Logical Disk Drives)” on page 168 and “Backing Up CONFIGURATION” on page 173.

DR Image File

Data required for temporary DR OS installation and configuration (**DR image**) is packed in a single large file and stored on the backup medium and optionally on the Cell Manager during a full client backup. If you want to save the full disaster recovery image file to the Cell Manager for all clients in the backup specification, perform the following steps:

1. In the Context List, select `Backup`.

17.923
Paula

2. In the `Scoping pane`, expand the `Backup Specifications` and then Filesystem.
3. Select the backup specification you will use for a full client backup (create it if you have not created it already).
4. In the Results Area, click `Options`.
5. Under `Filesystem Options` click `Advanced`.
6. Click the `WinFS Options` and check the `Copy full DR image to disk` check box.

Figure 10-2 **WinFS Options Tab**

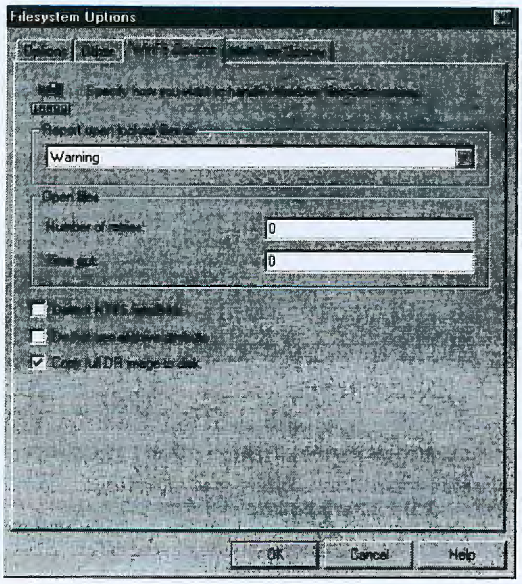

Report open locked files as: Warning
Open files
Number of retries: 0
Time out: 0
□ [illegible]
□ [illegible]
☑ Copy full DR image to disk

If you want to copy the DR image files only for particular clients in the backup specification, perform the following steps:

1. In the Context List, select `Backup`.
2. In the `Scoping pane`, expand the `Backup Specifications` and then `Filesystem`.
3. Select the backup specification you will use for a full client backup

17922
Paula

(create it if you have not created it already).

4. In the `Results Area`, click `Backup Object Summary`.
5. Select the client for which you would like to store its DR image file onto the Cell Manager and click `Properties`.
6. Click the `WinFS Options` and check the `Copy full DR image to disk` check box.

If the disaster recovery image is saved to the Cell Manager during backup, it is stored into `<Data_Protector_home>\Config\dr\p1s` (Windows Cell Manager) or into `/etc/opt/omni/dr/p1s` (UNIX Cell Manager) with the name `<client name>.img`. This is useful if you are going to prepare a disaster recovery CD ISO image on the Cell Manager, because it is much faster to obtain DR image from disk than from the backup medium.

TIP

If you do not have enough free disk space in the destination directory, you can create a link to another volume on UNIX or create a mount point on Windows.

Phase 1 Startup File (P1S)

In addition to the DR image file, a **Phase 1 Startup file (P1S)** is created during full backup. It is saved on backup medium and on the Cell Manager into `<Data_Protector_home>\Config\dr\p1s` directory (Windows Cell Manager) or into `/etc/opt/omni/dr/p1s` directory (UNIX Cell Manager) with the filename equal to the hostname (for example, `computer.company.com`). It is a Unicode UTF-8 encoded file that contains information on how to format and partition all disks installed in the system, whereas the updated SRD file contains only system information and data about backup objects and corresponding media.

After a disaster occurs, you can use the EADR Wizard to merge DR image, SRD and P1S files with disaster recovery installation into a **disaster recovery CD ISO image**, which can be burned on a CD using any CD burning tool. This **disaster recovery CD** can then be used to perform automated disaster recovery. Note that disaster recovery CD has to be prepared in advance for the Cell Manager. Additional steps are required if you are preparing disaster recovery CD of a Microsoft Cluster node. See "Restoring the Microsoft Cluster Server Specifics" on page 490.

17.921
Paula

IMPORTANT

It is recommended to restrict access to backup media, DR images, SRD files and disaster recovery CDs due to security reasons.

Preparing DR CD ISO Image

To prepare a DR CD ISO image, perform the following steps:

1. In the Context List, select `Restore`.
2. Click the `Tasks` navigation tab and select `Disaster Recovery` in the Scoping Pane.
3. From the drop down list in the Results Area, select the client you would like to recover.
4. Click `Enhanced Automated Disaster Recovery` and then `Next`.
5. For each critical object select an appropriate object version and click `Next`.
6. If you have saved the DR image file on the Cell Manager, specify its location, otherwise click `Restore from backup medium`. Click `Next`.
7. Select the destination directory where you want to place the ISO CD image (`recovery.iso`) and click `Finish` to create the ISO CD image.

WARNING

If you place a new ISO CD image to a location where a `recovery.iso` is already located, the old ISO CD image will be overwritten by the new one without a warning.

8. Burn the disaster recovery ISO CD image on a CD using any CD burning tool.

IMPORTANT

You have to perform a new backup and prepare a new DR CD after each hardware, software or configuration change. This also applies to any network configuration changes, such as change of IP address or DNS server.

17920
Paula

Recovery

You need the following to successfully perform a disaster recovery on the crashed system:

- A new hard disk to replace your crashed disk.
- A successful full client backup of the client that you want to recover.
- The Data Protector disaster recovery CD.

EADR of a Windows Client

The following is a step-by-step procedure for performing EADR of a Windows system:

1. Boot from the disaster recovery CD of the original system.
2. Press **F12** when the following message is displayed: `To start recovery of the machine <HOSTNAME> press F12.`
3. Select the scope of recovery and press **Enter**. There are 5 different scopes of recovery:
   - `No recovery`: Disaster recovery is not performed and the computer is rebooted.
   - `Default Recovery`: Critical volumes are recovered. All other disks are not partitioned and formatted and are ready for Phase 3.
   - `Minimal Recovery`: Only system and boot disks are recovered (available for EADR and OBDR only).
   - `Full Recovery`: (For future releases).
   - `Full with Shared Volumes`: Available for MSCS only. This option should be used if all nodes in the MSCS have crashed and you are performing Enhanced Automated Disaster Recovery of the first node. It will recover all volumes in the Restore Set including cluster shared volumes that were locked by the backed-up node at backup time.

     If at least one node is up and the MSCS service is running, than shared volumes will not be restored because the node keeps them locked. In this case, you should use `Default Recovery`.
4. After you have selected the scope of the recovery, Data Protector starts setting up the DR OS directly to the hard disk. You can monitor the progress and, when the DR OS is set up, the system reboots.
5. Wait for 10 seconds when prompted `To start recovery of the`

17.919
Paulo

machine <HOSTNAME> press F12, to boot from the hard disk and not from the CD.

6. Data Protector will then reestablish the previous storage structure within the selected scope of recovery and restore all critical volumes. The temporary DR OS will be deleted after the first login, except in the following cases:
   - Minimal Recovery is selected.
   - You have interrupted the Disaster Recovery Wizard during the 10 seconds pause after it has found the DR installation and SRD file on the backup medium, and have selected the Use Debugs option.
   - You have manually started the omnidr command with the no_reset or debug options.
   - Disaster recovery fails.
7. Additional steps are required if you are recovering a Cell Manager or performing advanced recovery tasks. See "Advanced Recovery Tasks" on page 490 for more information.
8. Restore user and application data using the standard Data Protector restore procedure.

NOTE Data Protector does not restore the volume-compression flag after recovery. All files that were compressed at backup time will be restored as compressed, but you will have to manually set the volume compression if you want any new files created to be compressed as well.

One Button Disaster Recovery of a Windows System

One Button Disaster Recovery (OBDR)is a fully automated Data Protector recovery method for Windows clients and Cell Manager, where user intervention is reduced to minimum. For details on supported operating systems, refer to the *HP OpenView Storage Data Protector Software Release Notes*.

OBDR collects all relevant environment data automatically at backup time. During backup, data required for temporary DR OS setup and configuration is packed in a single large OBDR image file and stored on the backup tape. When a disaster occurs, OBDR device (backup device, capable of emulating CD-ROM) is used to boot the target system directly from the tape which contains the OBDR image file with disaster recovery information.

Data Protector then installs and configures the disaster recovery operating system (DR OS), formats and partitions the disks and finally restores the original operating system with Data Protector as it was at the time of backup.

IMPORTANT You have to perform a new backup after each hardware, software or configuration change. This also applies to any network configuration changes, such as change of IP address or DNS server.

The recovered volumes are:

- the boot partition
- the system partition
- the partitions containing Data Protector

Any remaining partitions can be recovered using the standard Data Protector recovery procedure.

The following sections explain the requirements, limitations, preparation and recovery pertaining to One Button Disaster Recovery on Windows systems. See also “Advanced Recovery Tasks” on page 490.

17.917
Paula

Requirements

- Data Protector Automatic Disaster Recovery and User Interface components must be installed on the systems for which you want to enable recovery using this method. See *HP OpenView Storage Data Protector Installation and Licensing Guide*.
- It is essential to have an OBDR capable computer configuration: the system's BIOS must support bootable CD extensions as defined in the El-Torito standard and read/write access to hard disk drive using LBA addressing via INT13h function XXh. The OBDR device must conform to the same standard when emulating the CD-ROM. The BIOS options can either be checked in the user's manuals of the system or by inspecting the system setup before the boot.

  For more information about supported systems, devices and media, please refer to the HP StorageWorks Tape Hardware Compatibility Table on the World Wide Web:

  http://www.openview.hp.com/products/datapro/spec_0001.html. Also see the *HP OpenView Storage Data Protector Software Release Notes*.
- The hardware configuration of the target system must be the same as of the original system. This includes SCSI BIOS settings (sector remapping).
- Replacement disks have to be attached to the same host bus adapter on the same bus.
- An additional 200 MB of free disk space is required on the boot partition at backup time. If this disk space is not available, the disaster recovery fails. If you had applied the Compress Drive on the original partition, you must have 400 MB free.
- All drivers, required for boot must be installed under the `<%SystemRoot%>` folder.
- Network must be available when you boot the system in Safe Mode with Networking or in Directory Services Restore Mode (Domain Controller only), but you must do the backup of the system after it was booted with normal boot process.
- A media pool with a `Non-appendable` media usage policy and `Loose` media allocation policy has to be created for the OBDR capable device. Only the media from such pool can be used for disaster recovery.

17916
Paula

- When backing up the client, the default 64 kB block size should be used to write to the device if you plan to perform an offline restore. This is the only default block size available on Windows when performing disaster recovery. To verify that the default 64 kB block size is set, choose `Advanced` in the `Properties` box, as shown in Figure 10-3:

Figure 10-3 **Verifying the Default Block Size**

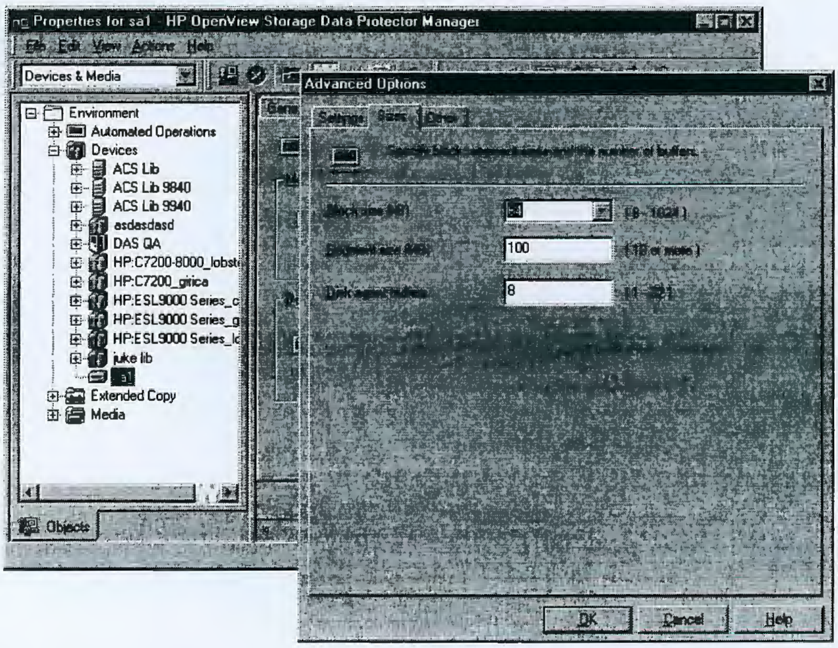

Limitations

General

- Multiboot systems that do not use Microsoft's boot loader are not supported.
- Internet Information Server (IIS) Database, Terminal Services Database and Certificate Server Database are not restored automatically during Phase 2. They can be restored on the target system using the standard Data Protector restore procedure.

17.915

- One Button Disaster Recovery backup session can only be performed for one selected client or Cell Manager on the same OBDR device at a time. This has to be done on a single, locally attached OBDR capable device.

Disk and Partition Configuration

- With fault-tolerant disk drives on the OS level for Windows NT systems, mirror set is supported while stripe and volume set are not. Dynamic disks are not supported on Windows 2000 (including mirror set upgraded from Windows NT).
- New disk has to be the same size or bigger than the crashed disk. If it is larger than the original disk, the difference will remain unallocated.
- Boot and system partition on Windows NT 4.0 must be physically bellow the first 7,8 GB due to the operating system limitations.
- Only vendor specific partitions of type 0x12 (including EISA) and 0xFE are supported for OBDR.

Preparation

Complete the steps described in the section "Preparing for a Disaster Recovery" on page 443 in order to prepare for disaster recovery in addition to completing the steps listed in this section. See also "Advanced Recovery Tasks" on page 490.

IMPORTANT You have to prepare for disaster recovery *before* a disaster occurs.

Create a media pool for DDS or LTO media with `Non-appendable` media usage policy (to ensure that this will be the only backup on tape) and `Loose` media allocation policy (because the tape is formatted during OBDR backup). In addition, the media pool must be selected as a default media pool for the OBDR device. Refer to "Creating a Media Pool" on page 102 for more information. Only media from such pool can be used for OBDR.

OBDR Backup

Use the following steps to perform OBDR backup locally on the system, for which you want to enable recovery using OBDR:

1. In the Context List, select `Backup`.

17.914
Paula

2. Click Tasks navigation tab and check One Button Disaster Recovery Wizard in the Scoping Pane.
3. From the drop-down list in the Results Area, select the client for which you would like to perform OBDR backup and click Next.
4. All critical objects are already selected (including the IDB in case of the Cell Manager OBDR backup) and can not be deselected. Manually select any other partitions you want to keep, because during the recovery procedure, Data Protector deletes all partitions from your system. Click Next.
5. Select the locally attached OBDR device you are going to use for backup and click Next.
6. Select backup options. See "Using Backup Options" on page 225 for details.
7. Click Next to proceed to the Scheduler page, which can be used to schedule the backup. See "Scheduling Unattended Backups" on page 207 for more information.
8. Click Next to display the Backup Object Summary page, in which you can review the backup options.

NOTE

In the Summary page, you cannot change a previously selected backup device or the order in which the backup specifications follow one another (move up and move down functionalities are not available). Only OBDR non-essential backup objects can be deleted as well as general object properties can be viewed.

However, a backup object's description can be changed.

9. In the final page of the Backup wizard, you can save the backup specification, start the interactive backup, or preview the backup.

   It is recommended to save the backup specification so that you can schedule or modify it later.

Modifying an OBDR Backup Specification

Once a backup specification is saved, you can edit it. Right-click the backup specification and select Properties. You are offered to treat the modified backup specification as a standard Data Protector backup specification or as an OBDR backup specification. Save it as an OBDR backup specification to keep it in the original One Button Disaster Recovery format. If saved as a standard backup specification, it is not usable for OBDR purposes.

10. Click `Start Backup` to run the backup interactively. The `Start Backup` dialog box appears. Click `OK` to start the backup.

A bootable image file of the system, containing all information required for installation and configuration of temporary DR OS, will be written at the beginning of the tape to make it bootable.

IMPORTANT You have to perform a new backup and prepare a bootable backup medium after each hardware, software or configuration change. This also applies to any network configuration changes, such as change of IP address or DNS server.

IMPORTANT It is recommended to restrict access to backup media due to security reasons.

Recovery

You need the following to successfully perform a disaster recovery on the crashed system:

- A new hard disk to replace your crashed disk (if needed).
- A bootable backup medium with all critical objects of the client that you want to recover.
- An OBDR device connected locally to the target system.

OBDR Procedure The following is a step-by-step procedure for performing a One Button Disaster Recovery of a Windows system:

1. Insert the tape containing the image file and your backed up data into an OBDR device.
2. Shut down the target system and power off the tape device.
3. Power on the target system and while it is being initialized, press the eject button to the tape device and power it on. For details see the

17.912
Paula

device documentation.

4. In the screen that appears, select the scope of recovery and press **Enter**. There are 5 different scopes of recovery:

   - `No recovery`: Disaster recovery is not performed and the computer is rebooted.
   - `Default Recovery`: Critical volumes are recovered. All other disks are not partitioned and formatted and remain empty and ready for Phase 3.
   - `Minimal Recovery`: Only system and boot disks are recovered (available for EADR and OBDR only).
   - `Full Recovery`: (For future releases).
   - `Full with Shared Volumes`: Available for MSCS only. This option should be used if all nodes in the MSCS have crashed and you are performing One Button Disaster Recovery of the first node. It will recover all volumes in the Restore Set including cluster shared volumes that were locked by the backed-up node at backup time.

TIP To enable automatic restore of all shared disk volumes in the MSCS, move all volumes temporarily to the node, for which you are preparing OBDR boot tape. It is namely impossible to collect enough information to configure disks in Phase 1 for shared disk volumes that are locked by another node at backup.

   If at least one node is up and running than shared volumes will not be restored because the node keeps them locked. In this case, you should use `Default Recovery`.

5. After you have selected the scope of recovery, Data Protector starts setting up the DR OS directly to the hard disk. You can monitor the progress and, when the DR OS is set up, the system reboots.

6. Data Protector will then reestablish the previous storage structure and restore all critical volumes. The temporary DR OS will be deleted after the first login, except in the following cases:

   - `Minimal Recovery` is selected.

1791
Paula

- You have interrupted the Disaster Recovery Wizard during the 10 seconds pause after it has found the DR installation and SRD file on the backup medium, and have selected the `Use Debugs` option.
- You have manually started the `omnidr` command with the `no_reset` or `debug` options.
- Disaster recovery fails.

7. Additional steps are required if you are recovering a Cell Manager or performing advanced recovery tasks. See "Advanced Recovery Tasks" on page 490 for more information.
8. Restore user and application data using the standard Data Protector restore procedure.

NOTE Data Protector does not restore the volume-compression flag after recovery. All files that were compressed at backup time will be restored as compressed, but you will have to manually set the volume compression if you want any new files created to be compressed as well.

17.910
Paula

Automated System Recovery

Automated System Recovery (ASR) is an automated system on Windows systems, which reconfigures a disk to its original state (or resizes the partitions if the new disk is larger than the original disk) in the case of a disaster. This includes disk partitioning and logical volume configuration (file formats, drive letter assignments, volume mountpoints, and volume characteristics). ASR thus enables the Data Protector `drstart.exe` command to install the active DR OS which provides Data Protector disk, network, tape and file system access.

Data Protector then recovers the target system to the original system configuration and finally restores all user data.

For details on supported operating systems, refer to the *HP OpenView Storage Data Protector Software Release Notes*.

IMPORTANT You have to perform a full client backup after each hardware, software or configuration change and to update the ASR diskettes. This also applies to any network configuration changes, such as change of the IP address or DNS server.

IMPORTANT You have to create the ASR set for the Cell Manager in advance, because you will not be able to obtain the ASR archive file after the disaster. ASR sets for other systems can be created using Cell Manager when a disaster occurs.

The recovered volumes are:

- the boot partition
- the system partition
- the partitions containing Data Protector

Any remaining partitions can be recovered using the standard Data Protector recovery procedure.

17.909
Paula

The following sections explain the requirements, limitations, preparation, and recovery pertaining to Automated System Recovery on Windows systems. See also "Advanced Recovery Tasks" on page 490.

Requirements

- Data Protector Automatic Disaster Recovery component must be installed on systems for which you want to enable recovery using ASR. See the *HP OpenView Storage Data Protector Installation and Licensing Guide*.

Hardware Configuration

- The hardware configuration of the target system must be identical to that of the original system, except for hard disk drives, video cards and network interface cards. If you have replaced a network card or a video card, you will have to manually configure it.
- Floppy disk drive must be installed.
- Floppy and CD drives must be connected to IDE or SCSI controllers. External devices such as USB or PCMCIA devices are not supported.

Hard Disk Drives

- The target system must have the same number of physical disks with critical volumes as the original system.
- Replacement disks must be attached to the same host bus adapter on the same bus.
- The storage capacity of each replacement disk on the target system must be bigger than or equal to the capacity of the corresponding disk on the original system. In addition, disk geometry of the replacement disk must be the same as on the replaced disk.
- All disks on the target system must have 512 bytes-per-sector.
- All disks used in ASR must be accessible to the system (hardware RAID must be configured, SCSI disks must be correctly terminated, etc.)
- When backing up the client, the default 64 kB block size should be used to write to the device if you plan to perform an offline restore. This is the only default block size available on Windows when performing disaster recovery. To verify that the default 64 kB block size is set, choose `Advanced` in the `Properties` box, as shown in Figure 10-4:

17908
Paula

Figure 10-4 **Verifying the Default Block Size**

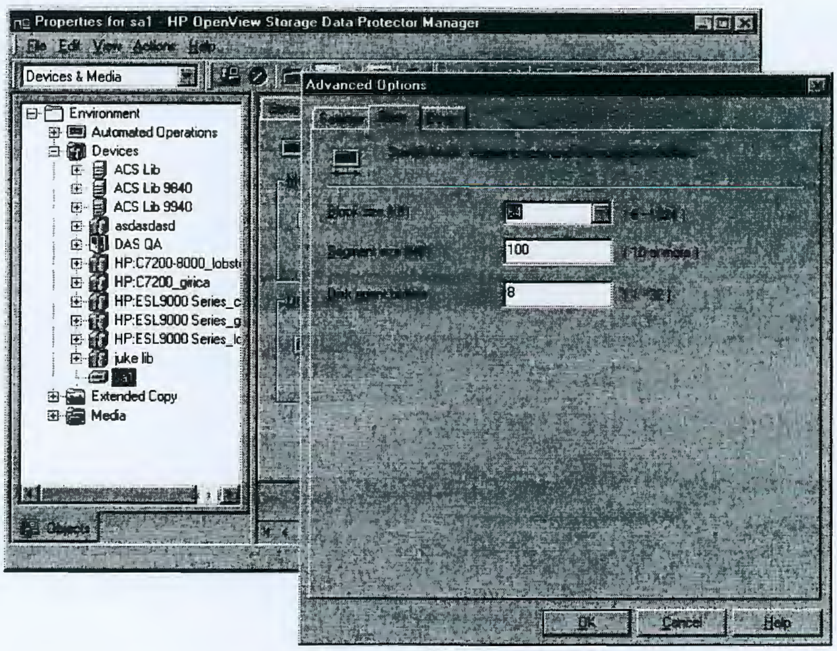

Limitations

- Windows XP Home Edition does not support ASR.
- Multiboot systems that do not use Microsoft's boot loader are not supported.
- Internet Information Server (IIS) Database, Terminal Services Database, and Certificate Server Database are not restored automatically during Phase 2. They can be restored on the target system using the standard Data Protector restore procedure.
- Data stored on vendor specific partitions is not automatically restored during ASR. The partitions will be recreated during the ASR but you will have to restore the data manually using the vendor specific procedure for restoring data. However, you can restore data on EISA utility partition using the standard Data Protector restore procedure.

- Only those local backup devices are supported, that can be installed by Windows during OS installation (no additional drivers are required).

Preparation

Complete the steps described in the section "Preparing for a Disaster Recovery" on page 443 in addition to completing the steps listed in this section. See also "Advanced Recovery Tasks" on page 490 in order to prepare for disaster recovery.

IMPORTANT You have to prepare for disaster recovery *before* a disaster occurs.

Prerequisite

A full client backup (including the configuration) is a prerequisite for successful ASR. See "Backing Up Filesystems (Logical Disk Drives)" on page 168 and "Backing Up CONFIGURATION" on page 173.

After you have performed the full client backup you have to prepare an ASR set. An ASR set is a collection of files stored on two or three diskettes, required for proper reconfiguration of the replacement disk (disk partitioning and logical volume configuration) and automatic recovery of the original system configuration and user data that was backed up during the full client backup. These files are stored as an ASR archive file on the Cell Manager (in `<Data_Protector_home>\Config\dr\asr` on a Windows Cell Manager or in `/etc/opt/omni/dr/asr/` on a UNIX Cell Manager) as well as on the backup medium. ASR archive file is extracted to two diskettes for 32-bit Windows system or three diskettes for 64-bit Windows system after a disaster occurs. You need these diskettes to perform ASR.

NOTE You have to create the ASR set for the Cell Manager in advance, because you will not be able to obtain the ASR archive file after the disaster.

Creation of ASR Set

Perform the following steps to create an ASR set:

1. Perform a full client backup.
2. Insert a diskette in the floppy drive.

17906

3. In the `HP OpenView Storage Data Protector Manager`, switch to the `Restore` context.
4. Click the `Tasks` navigation tab and select `Disaster Recovery` in the Scoping Pane.
5. From the drop down list in the Results Area, select the client for which you would like to create an ASR set.
6. Click `Create Automated System Recovery set` and then click `Next`.

Figure 10-5

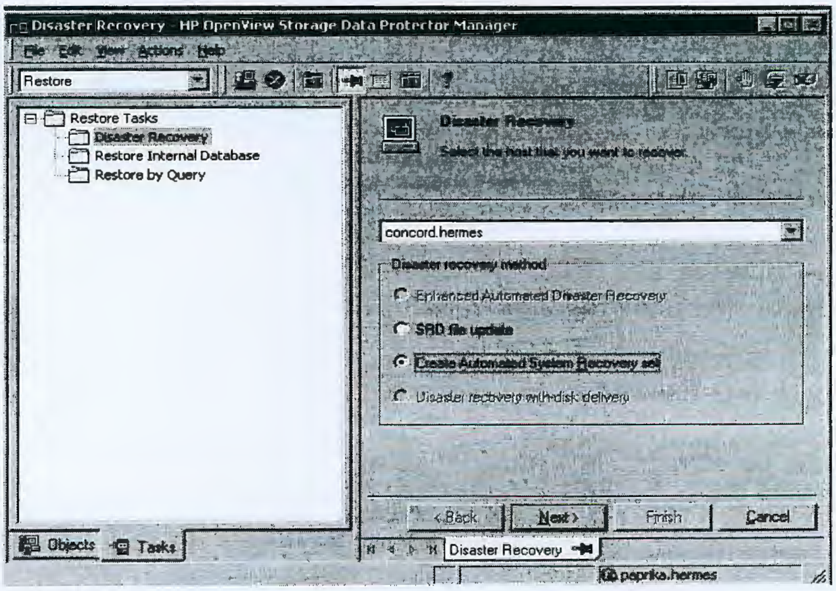

Disaster Recovery - HP OpenView Storage Data Protector Manager

Restore Tasks
- Disaster Recovery
- Restore Internal Database
- Restore by Query

Disaster Recovery

Select the host that you want to recover.

concord.hermes

Disaster recovery method

○ Enhanced Automated Disaster Recovery
○ SRD file update
● Create Automated System Recovery set
○ Disaster recovery with disk delivery

Data Protector will obtain the ASR archive file from the Cell Manager. If it is not saved on the Cell Manager, the Disaster Recovery wizard will offer you to recover it from the backup medium.

7. For each critical object, select the appropriate object version and click `Next`.
8. ASR archive file created during a full client backup is downloaded from the Cell Manager. Select the destination location where you want your ASR archive file extracted and select the `Copy DR`

17905
Paula

`installation` check box to copy DR installation files to the same location. The recommended destination is your floppy drive because you will need these files stored on diskettes (ASR set) to perform ASR.

Data Protector will create two diskettes for a 32 bit Windows system and three diskettes for a 64 bit Windows system. ASR set for the Cell Manager has to be prepared in advance, while you can prepare ASR diskettes for other systems using the Cell Manager when a disaster occurs.

Once the ASR set is created, you have to update only the first diskette (which contains ASR information) after each hardware, software or configuration change. This also applies to any network configuration changes, such as a change of the IP address or DNS server. In order to update the first diskette from the ASR set, repeat the whole procedure, but you do not have to select the `Copy DR installation` check box. This option copies the DR installation files (to a selected destination), which do not need to be updated.

IMPORTANT It is recommended to restrict access to ASR diskettes due to security reasons.

Local Devices

If you are using a locally attached device for ASR, test if it is supported. To do so, perform the following steps:

1. Run `devbra -dev` from the command prompt (from *<Data_Protector_home>*\bin).
2. Rename the `scsitab` file (located in *<Data_Protector_home>*) and run `devbra -dev` from the command prompt again.
3. Compare the both outputs of the `devbra -dev` command. If they are identical, ASR using this device is possible, otherwise copy the `scsitab` file to the first ASR diskette. You have to copy the `scsitab` file only the first time you are preparing the ASR set. You do not have to copy it when you are only updating the ASR set. Refer to the "Support of New Devices" on page 41 for more information.
4. Rename the `scsitab` file back to the original name.

17.904
Paula

Recovery

To successfully perform a disaster recovery of the crashed system, you need the following:

- A new hard disk to replace your crashed disk.
- A successful full client backup of the client that you want to recover.
- Updated ASR set.
- Windows installation medium.

ASR Procedure

The following is a step-by-step procedure for performing ASR:

1. Boot from the Windows installation medium.
2. Press **F2** during the start of the OS setup to enter the ASR mode.
3. Provide the first (updated) diskette from the ASR set.
4. After reboot, Disaster Recovery Wizard pops-up and requires input for the `DR installation source` and `SRD Path`. DR installation and SRD file are both located on the first diskette of the ASR set (`a:\`).
5. Change diskette(s) when prompted.

   Original storage structure will be automatically reestablished and all critical data automatically restored based on the information in the ASR set.
6. Reboot the system when prompted and remove the Windows installation medium and ASR diskette.
7. Restore user and application data using the standard Data Protector restore procedure.

17.903
Paula

Restoring the Data Protector Cell Manager Specifics

This section explains additional steps for particular methods that should be performed when restoring Windows Cell Manager.

Making IDB consistent (all methods)

The procedure described in this section should only be used after you have performed the general disaster recovery procedure.

To make the IDB consistent, you have to import the medium with the last backup so that the information about the backed up objects is imported to the database. In order to do so, you have to perform the following steps:

1. Using the Data Protector GUI, recycle the medium or media with the backup of the partitions that remain to be restored for enabling the medium or media to be imported in the IDB. Refer to "Recycling Media" on page 110 for more information on how to do this. Sometimes it is not possible to recycle a medium since Data Protector keeps it locked. In such a case stop Data Protector processes and delete the `\tmp` directory by running the following commands:
   ```
   <Data_Protector_home>\bin\omnisv -stop
   del <Data_Protector_home>\tmp\*.*
   <Data_Protector_home>\bin\omnisv -start
   ```
2. Using the Data Protector GUI, export the medium or media with the backup of the partitions that remain to be restored. Refer to "Exporting Media from Data Protector" on page 112 for more information on how to do this.
3. Using the Data Protector GUI, import the medium or media with the backup of the partitions that remain to be restored. Refer to "Importing Media" on page 100 for more information on how to do this.

Enhanced Automated Disaster Recovery Specifics

Two additional steps are required in Phase 0 if you are recovering Windows Cell Manager using Enhanced Automated Disaster Recovery:

17.902
Paula

- Disaster recovery CD for the Cell Manager should be prepared in advance.

IMPORTANT You have to perform a new backup and prepare a new DR CD after each hardware, software or configuration change. This also applies to any network changes, such as change of IP address or DNS server.

- In addition to the Cell Manager, you should save the updated SRD file of the Cell Manager on several secure locations as a part of the disaster recovery preparation policy, because the SRD file is the only file in Data Protector where information about objects and media is stored, when IDB is not available. If the SRD file is saved only on the Cell Manager, it is not accessible if the Cell Manager fails. See "Preparation" on page 451.

IMPORTANT It is recommended to restrict access to backup media, DR images, SRD files and disaster recovery CDs.

One Button Disaster Recovery Specifics

Since the IDB is not available if the Cell Manager has crashed, you have to know the location of OBDR bootable medium.

IMPORTANT You have to perform a new OBDR backup and prepare a new bootable medium after each hardware, software or configuration change. This also applies to any network changes, such as change of IP address or DNS server.

IMPORTANT It is recommended to restrict access to backup media.

Automated System Recovery Specifics

An additional step is required in Phase 0 if you are recovering Windows Cell Manager using Automated System Recovery (ASR):

- ASR diskette for the Cell Manager should be prepared in advance.

IMPORTANT You have to perform a new backup and update the ASR diskette after each hardware, software or configuration change. This also applies to any network changes, such as change of IP address or DNS server.

IMPORTANT It is recommended to restrict access to backup media and ASR diskettes.

Advanced Recovery Tasks

This section provides explanation of the steps you will need to take if you want to perform advanced recovery tasks such as restoring Microsoft Cluster Server and Internet Information Server.

Restoring the Microsoft Cluster Server Specifics

This section provides explanation of the steps you will need to take if you want to perform disaster recovery of a Microsoft Cluster Server (MSCS). For concepts and general information please refer to the clustering section in the *HP OpenView Storage Data Protector Concepts Guide* and "Cluster Integrations with Data Protector" on page 613 in the *HP OpenView Storage Data Protector Administrator's Guide*.

Select the disaster recovery method that is appropriate for your cluster and include it in your disaster recovery plan. Consider the limitations and requirements pertaining to disaster recovery methods before deciding which method to use. Execute tests from the test plan.

Possible Scenarios

There are two possible scenarios for disaster recovery of a MSCS:

- at least one of the nodes is up and running
- all nodes in the cluster have experienced a disaster

IMPORTANT MSCS can be recovered using any disaster recovery method except for Disk Delivery Disaster Recovery. All specifics, limitations and requirements pertaining a particular disaster recovery method you are going to use also apply for the disaster recovery of a MSCS. For details on supported operating systems, refer to the *HP OpenView Storage Data Protector Software Release Notes*.

All prerequisites for disaster recovery (i.e. consistent and up-to-date backup, updated SRD file, all faulty hardware replaced...) must be met to recover MSCS.

Consistent backup for MSCS should include:

- all nodes
- administrative virtual server (defined by the administrator)
- if Data Protector is configured as a cluster aware application, then Cell Manager virtual server and IDB should be included in the backup specification.

Disaster Recovery of a Secondary Node

This is the basic scenario for disaster recovery of a MSCS. The following must be true in addition to other prerequisites for disaster recovery:

- at least one of the cluster nodes is functioning properly
- the cluster service is running on that node
- all physical disk resources must be online (i.e. owned by the cluster)
- all normal cluster functionality is available (the cluster administration group is online)
- the Cell Manager is online

In this case, the disaster recovery of a cluster node is the same as the disaster recovery of a Data Protector client. You should follow the instructions for the specific disaster recovery method that you will use to restore the secondary node.

NOTE Only local disks are restored, because all shared disks are online and owned by the working node(s) during recovery and locked.

After the secondary node has been recovered, it will join the cluster after boot.

You can restore the MSCS database after all nodes have been recovered and have joined the cluster to ensure its coherency. The MSCS database is part of the `CONFIGURATION` on Windows. See “Restoring the Windows CONFIGURATION” on page 280.

Disaster Recovery of the Primary Node

In this case all nodes in the MSCS are unavailable and the cluster service is not running.

17.898
Paula

Merging P1S files of all nodes for EADR

Another step is required for EADR after backup has been performed. Information on shared cluster volumes in P1S files for all nodes in the MSCS has to be merged so that P1S file of each node contains information on shared cluster volumes configuration. It is namely impossible to collect enough information to configure the disk in Phase 1 for shared disk volumes that are locked by another node at backup. Merging P1S files is necessary to enable restore of all cluster shared disks if all nodes in the cluster have experienced a disaster. To merge the P1S files of all nodes, execute the `mmerge.cmd` command from the *<Data_Protector_home>*`\bin\drim\bin`:

```
mmerge p1sA_path ... p1sX_path
```

Where `p1sA` is the full path of the first node's P1S file and `p1sX` is the full path of the P1S file of the last node in the MSCS. Merged P1S files will be saved in the same directory as the source P1S files with the `.merged` appended to their filename (for example, `computer.company.com.merged`). Rename the merged P1S files back to the original name (delete the `.merged` extension).

The `mmerge.cmd` command works only on Windows systems with Data Protector Automatic Disaster Recovery component installed. If you are using an UNIX Cell Manager, copy the P1S files to a Windows client which has Automatic Disaster Recovery component installed and merge the files. Rename the merged P1S files back to the original name and copy them back to the Cell Manager.

Example for merging P1S files for MSCS with 2 nodes: `mmerge` *<Data_Protector_home>*`\config\dr\p1s\node1.company.com` *<Data_Protector_home>*`\config\dr\p1s\node2.company.com`. You have to enclose the path in quotes on Windows if the path contains a space character. The merged files will be `node1.company.com.merged` and `node2.company.com.merged`. Rename the files back to their original names (you will have to rename the source P1S files first): node1.company.com and node2.company.com.

You can avoid merging P1S files after backup by moving all shared cluster volumes temporarily to the node which you are going to back up. In this case all required information about all shared cluster volumes can be collected. In this case only that node can be the primary node.

The following must be true in addition to other prerequisites for disaster recovery:

- the primary node must have write access to the quorum disk (the quorum disk must not be locked)

17897
Paula

- the primary node must have write access to all IDB volumes, when recovering the Cell Manager
- all other nodes must be shut down until all physical disk resources are online

In this case, you have to restore the primary node with the quorum disk first. The IDB has to be restored as well if the Cell Manager has been installed in the cluster. Optionally you can restore the MSCS database. After the primary node has been restored, you can restore all remaining nodes.

NOTE The MSCS service uses a hard disk signature written into the MBR of every hard disk to identify physical disks. If the shared cluster disks have been replaced, this means that the disk signatures were changed during Phase 1 of disaster recovery. As a consequence, the Cluster Service will not recognize the replaced disks as valid cluster resources, and cluster groups depending on those resources will fail. See "Restoring Hard Disk Signatures On Windows" on page 495 for more information.

Perform the following steps to restore the primary node:

1. Perform disaster recovery of the primary node (including the quorum disk).
   - Assisted Manual Disaster Recovery: All user and application data on the quorum disk will be restored automatically by the `drstart` command. (`-full_clus` option)
   - EADR and OBDR: When you are asked to select the scope of recovery, select `Full with Shared Volumes` to restore quorum disk.
   - Automated System Recovery: All user and application data on the quorum disk will be automatically restored.

TIP

To enable automatic restore of all shared disk volumes in the MSCS using OBDR method, move all volumes temporarily to the node for which you are preparing OBDR boot tape. It is namely impossible to collect enough information to configure the disk in Phase 1 for shared disk volumes that are locked by another node.

2. Reboot the computer.
3. Restore the cluster database. MSCS database is part of the `CONFIGURATION` on Windows. See “Restoring the Windows CONFIGURATION” on page 280.

NOTE

The MSCS service must be running in order to be able to restore the MSCS database. Therefore it can not be restored automatically during Phase 2 of disaster recovery. However, the cluster database can be restored at the end of Phase 2 using the standard Data Protector restore procedure.

4. Make the IDB consistent if you are recovering a Cell Manager. See “Making IDB consistent (all methods)” on page 487.
5. The quorum and IBD volumes are restored. All other volumes are left intact and are claimed by the recovered primary node if they are not corrupted.

   If they are corrupted you have to:

   a. disable the cluster service and cluster disk driver (the steps required to do so are described in MSDN Q176970)
   b. reboot the system
   c. reestablish the previous storage structure
   d. enable the cluster disk driver and cluster service
   e. reboot the system
   f. restore user and application data
6. Restore the remaining nodes. See “Disaster Recovery of a Secondary Node” on page 491.

17.895
Paula

Restoring Hard Disk Signatures On Windows

The MSCS service uses a hard disk signature written into the MBR of every hard disk to identify physical disks. If the shared cluster disks have been replaced, this means that the disk signatures were changed during Phase 1 of disaster recovery. As a consequence, the Cluster Service will not recognize the replaced disks as valid cluster resources, and cluster groups depending on those resources will fail. This applies only to the restore of the active node, since shared cluster resources are operational as long as at least one of the nodes is up and running and claims ownership of the resources. This problem does not apply to EADR and OBDR critical disks because the original disk signatures of all EADR/OBDR critical disks are automatically recovered. In case you have replaced any other disks, you will have to restore their hard disk signatures as well.

The most critical shared disk is the cluster quorum resource. If it has been replaced, than the original disk signature must be restored, or the cluster service will not start.

During Phase 2, the MSCS Database is restored into the `\TEMP\ClusterDatabase` directory on the system volume. After the system is rebooted, the cluster service will not be running, because the quorum resource will not be identified due to the changed hard disk signature in Phase 1. This can be resolved by running the `clubar` utility (located in the `<`*`Data_Protector_home`*`>\bin\utilns`), which restores the original hard disk signature. After `clubar` successfully finishes, the cluster service is automatically started.

Example

At the command prompt type `clubar r c:\temp\ClusterDatabase force q:` to restore a MSCS Database from `c:\temp\ClusterDatabase`.

For more information on `clubar` usage and syntax, see the `clubar.txt` file located in the `<`*`Data_Protector_home`*`>\bin\utilns`.

If the Data Protector shared disk on the Cell Manager is different from the quorum disk, it has to be restored as well. To restore the signature of the Data Protector shared disk and any other application disk, you should use the `dumpcfg.exe` utility included in the Windows 2000 Resource Kit. For details on using `dumpcfg.exe`, run `dumpcfg /?` or see the Windows 2000 Resource Kit documentation. For more information on the problems with hard disk signatures on Windows 2000, see MSDN article Q280425.

17894
Paula

There is no simple procedure for recovering the disk signatures on Windows NT. MSDN article Q243195 describes a complicated procedure that includes manual modifying of the registry, and should be followed to recover disk signatures. Note that using Registry Editor incorrectly may cause serious problems that may require you to reinstall your operating system. On the other hand, EADR automatically recovers disk signatures of all EADR critical volumes and is therefore the preferred method to be used to restore the MSCS.

You can obtain the original hard disk signatures from the SRD files. The signature is a number following the volume keyword in the SRD file.

Example

```
-volume 5666415943 -number 0 -letter C -offslow 32256
-offshigh 0 -lenlow 320430592 -lenhigh 2 -fttype 4 -ftgroup
0 -ftmember 0

-volume 3927615943 -number 0 -letter Q -offslow 320495104
-offshigh 2 -lenlow 1339236864 -lenhigh 0 -fttype 4 -ftgroup
0 -ftmember 0
```

The number following the -volume keyword is the signature of the hard disk. In this case the SRD file stores information about a local hard disk (with drive letters C) and quorum disk (with drive letter Q). The signature of the quorum disk is stored only in the SRD file of the active node (at backup time), because it keeps the quorum disk locked and thus prevents other nodes from accessing the quorum disk. It is therefore recommended to always back up the whole cluster, because you need the SRD files of all nodes in the cluster, since only all SRD files together include enough information to configure the disk in Phase 1 for shared disk volumes. Note that a hard disk signature stored in the SRD file is represented as a decimal number, whereas dumpcfg requires hexadecimal values.

Restoring Internet Information Server (IIS) Specifics

Internet Information Server (IIS) is not supported for disaster recovery. To perform Assisted Manual Disaster Recovery of an IIS, follow these steps (in addition to the steps required for Assisted Manual disaster recovery):

1. Do not install the IIS during clean installation of the system.
2. Stop or uninstall the `IIS Admin Service`, if it is running.
3. Run the `drstart` command.

4. The IIS Database is restored as a plain file (with the filename `DisasterRecovery`) into the default IIS location (`%SystemRoot%\system32\inetsrv`).
5. After the successful boot, restore the IIS Database using the standard Data Protector restore procedure or IIS Backup/Restore snap-in. Note that this may take quite some time.

Troubleshooting

1. If any of the IIS dependant services (for example, SMTP, NNTP) do not start automatically, try to start them manually.
2. If this fails, stop the `IIS Admin Service` and restore the `%SystemRoot%\system32\inetsrv\MetaBase.bin` file, using the `overwrite` option.

NOTE

`%SystemRoot%\system32\inetsrv` is the default location of `IIS Service`. If you have installed the service into other location, use this location as a destination for restore of `MetaBase.bin` file.

3. Start the `IIS Admin Service` and all dependant services.

Manual Disaster Recovery of an HP-UX Client

This chapter explains the procedure that should be used to recover a HP-UX client from a disaster.

The procedure is based on the Ignite-UX product; an application primary developed for HP-UX system installation and configuration tasks, which offers (in addition to a powerful interface for the system administration) preparation and recovery of the system from a disaster.

While Ignite-UX is focused on the disaster recovery of the target client, Data Protector has to be used to restore the user and application data in order to complete the *Phase 3* of disaster recovery.

This chapter cannot cover the full functionality of Ignite-UX. For detailed information please refer to the “Ignite-UX Administration Guide”.

Concept

Ignite-UX offers 2 different approaches to prepare a system for and recover a system from a disaster:

- Using custom installation medium (**Golden Image**)
- Using system recovery tools (**make_tape_recovery**, **make_net_recovery**)

While the usage of Golden Image is most suitable for IT environments with a large number of basically identical hardware configurations and OS releases, the usage of the system recovery tools supports the creation of recovery archives, which are customized for your individual systems.

Both methods allow the creation of bootable installation media like DDS-Tapes or CD’s. Using these media, the system administrator is able to perform a local disaster recovery directly from the system console of the failed client.

In addition, both methods can also be used to run a network based recovery of the client by assigning the failed client a suitable Golden Image or the previously created “recovery archive”. In this case, the client boots directly from the Ignite Server and runs the installation from the assigned depot, which has to be located on a NFS share on your network.

17.891
Paula

Use Ignite-UX GUI where it is supported.

Using Custom Installation Medium

Overview

Large IT environments often consist of a large number of systems that are based on identical hardware and software. Installation of OS, applications and required patches can be significantly reduced if a complete snapshot of the installed system is used to install other systems. Ignite-UX includes a feature, which allows you to modify parameters like networking or filesystem settings and add software like Data Protector to the image (with Ignite-UX command `make_config`) before you assign such a Golden Image to another system. This feature can thus be used to recover a system from a disaster.

Creating a “Golden Image”

Steps to Create a Golden Image

The following steps explain how to create a Golden Image of a client system on a target system, which will share the image via NFS to your network. In this example, Data Protector client is already installed on the client system and will be included in the “Golden Image” without additional configuration steps.

1. Copy the `/opt/ignite/data/scripts/make_sys_image` file from your Ignite-UX Server into a temporary directory on the client system.
2. Run the following command on the client node to create a compressed image of the client on another system: `make_sys_image -d <directory of the archive> -n <name of the archive>.gz -s <IP address of the target system>`

   This command will create a gzipped file depot in the specified directory on the system defined with the `-d` and `-s` options. Make sure that your HP-UX client has granted a passwordless access to the target system (an entry in the `.rhosts` file with the name of the client system on the target system) otherwise the command will fail.
3. Add the target directory to the `/etc/exports` directory on the target system and export the directory on the target server (`exportfs -av`).
4. On the Configuring Ignite-UX server, copy the archive template file

`core.cfg` to `archive_<name>.cfg`:

```
cp /opt/ignite/data/examples/core.cfg
/var/opt/ignite/data/<OS_Release>/archive_<name>.cfg
```

Example

```
cp /opt/ignite/data/examples/core.cfg
/var/opt/ignite/data/Rel_B.11.11/archive_HPUX11_11_DP50_C
L.cfg
```

5. Check and change the following parameters in the copied configuration file:

   - In the `sw_source` section:

     ```
     load_order = 0

     source_format = archive

     source_type="NET"

     # change_media=FALSE

     post_load_script = "/opt/ignite/data/scripts/os_arch_post_l"

     post_config_script =
     "/opt/ignite/data/scripts/os_arch_post_c"

     nfs_source = "<IP Target System>:<Full Path>"
     ```

   - In the matching `OS archive` section:

     ```
     archive_path = "<archive_name>.gz
     ```

6. Determine the "impacts" entries by running the command `archive_impact` on your image file and copy the output in the same "OS archive" section of your configuration file:

   ```
   /opt/ignite/lbin/archive_impact -t -g
   <archive_name>.gz
   ```

Example

```
/opt/ignite/lbin/archive_impact -t -g
/image/archive_HPUX11_11_DP50_CL.gz


impacts = "/" 506Kb

impacts = "/.root" 32Kb

impacts = "/dev" 12Kb

impacts = "/etc" 26275Kb

impacts = "/opt" 827022Kb

impacts = "/sbin" 35124Kb
```

17.889
Paula

```
impacts = "/stand" 1116Kb
impacts = "/tcadm" 1Kb
impacts = "/usr" 729579Kb
impacts = "/var" 254639Kb
```

7. To make Ignite-UX aware of the new created depot, add an `cfg` entry to the `/var/opt/ignite/INDEX` file with the following layout:

```
cfg "<This_configuration_name>" {
description "<Description of this configuration>"
"/opt/ignite/data/<OS>/config"
"/var/opt/ignite/data/<OS>/ archive_<name>.cfg
}
```

Example

```
cfg "HPUX11_11_DP50_Client" {
description "HPUX 11.i OS incl Patches and DP50 Client"
"/opt/ignite/data/Rel_B.11.11/config"

"/var/opt/ignite/data/Rel_B.11.11/archive_HPUX11_11_DP50_CL.cfg
"

}
```

8. Make sure that one or more IP addresses reserved for booting clients are configured in the `/etc/opt/ignite/instl_boottab` file. The number of IP addresses is equal to the number of parallel booting clients.

After the above described procedure is completed, you have a Golden Image of an HP-UX client (with a specific hardware and software configuration), which can be used to recover any client of a similar layout.

You have to repeat these steps to create a Golden Image for all systems with different hardware and software configuration.

17.888
Paula

NOTE Ignite-UX enables you to create a bootable tape or CD based on the created Golden Image. Please refer to the Ignite-UX Administration Guide for more information.Ignite-UX enables you to create a bootable tape or CD based on the created Golden Image. Please refer to the Ignite-UX Administration Guide for more information.

Recovery

Recovery Using a Golden Image

To recover an HP-UX client by applying the Golden Image, which is located on a NFS share on your network, perform the following steps:

- On the Client System

1. Replace the faulty hardware.
2. Boot the HP-UX client from the Ignite-UX server:
   `boot lan.<IP-address Ignite-UX server>install.`
3. Select `Install HP-UX` when the `Welcome to Ignite-UX` screen appears.
4. Choose `Remote graphical interface running on the Ignite-UX server` from the UI Option screen.
5. Respond to the Network configuration dialogue.
6. The system is now prepared for a remote Ignite-UX Server controlled installation.

- On the Ignite-UX Server

7. Right click the client icon in the Ignite-UX GUI and select `Install Client - New Install.`
8. Select the Golden Image you want to install, check the settings (network, filesystem, time zone,...) and click the `Go!` button.
9. You can check the installation progress by right clicking the client icon and choosing `Client Status...`
10. After the installation has finished, restore additional user and application data using the standard Data Protector restore procedure.

Using System Recovery Tools

Overview

The usage of the system recovery tools, bundled with the Ignite-UX, enables you a fast and easy recovery from a disk failure. The recovery archive of system recovery tools includes only essential HP-UX directories. However, it is possible to include other files and directories (for example, additional volume groups or the Data Protector files and directories) in the archive to speed up the recovery process.

`make_tape_recovery` creates a bootable recovery (installation) tape, customized for your system and enables you unattended disaster recovery by connecting the backup device directly to the target system and booting the target system from the bootable recovery tape. The backup device has to be locally connected to the client during the creation of the archive and recovery of the client.

`make_net_recovery` allows you to create a recovery archive over the network onto the Ignite-UX server or any other specified system. The target system can be recovered across subnets after booting ether from a bootable tape created by the Ignite-UX `make_boot_tape` command or the system boots directly from the Ignite-UX server. Booting directly from the Ignite-UX server can be automated with the Ignite-UX `bootsys` command or interactively specified on the boot console.

Creating Recovery Archives

The easiest way to create a recovery archive of an HP-UX client is to use the Ignite-UX GUI on the Ignite-UX server. All GUI commands can also be executed from the command line. Refer to the “Ignite-UX Administration Guide” for more information.

Prerequisites

Before you are able to prepare your system for disaster, the Ignite-UX fileset has to be installed on the client in order to enable the Ignite-UX server to communicate with the client.

Make sure that the revisions of the Ignite-UX fileset on the Ignite-UX sever and on the client are the same. The simplest way to keep everything consistent is to install Ignite-UX from a depot build on the Ignite-UX server. This depot can be constructed by running the following command on the Ignite-UX server:

```
pkg_rec_depot -f
```

17886
faula

This creates an Ignite-UX depot under `/var/opt/ignite/depots/recovery_cmds`, which can be specified as a source directory by `swinstall` on the client for the Ignite-UX software installation.

After you have installed Ignite-UX on the client node, you can use the GUI on the Ignite-UX server to create recovery archives using `make_net_recovery` or `make_tape_recovery`.

Creating an Archive Using make_tape_recovery

Perform the following steps to create an archive using `make_tape_recovery`:

1. Make sure that a backup device is connected to the HP-UX client.
2. Start the Ignite-UX GUI by executing the following command: `/opt/ignite/bin/ignite &`.
3. Right click the client icon and select `Create Tape Recovery Archive`.
4. Select a tape device, if more than one device is connected to the HP-UX client.
5. Select the volume groups you want to include into the archive.
6. The tape creation process will now begin. Check the status and log file on the Ignite-UX server by right clicking the client icon and selecting `Client Status`.

NOTE Ignite-UX recommends the usage of 90m DDS1 backup tapes to ensure that the tape will work with all DDS with any DDS drive.

Creating an Archive Using make_net_recovery

The procedure for creating a recovery archive using `make_net_recovery` is almost the same as using `make_tape_recovery`. The advantage is that there is no need for a locally attached backup device, as the recovery archive is stored on the Ignite-UX server by default.

1. Start the Ignite-UX GUI by executing the following command: `/opt/ignite/bin/ignite &`
2. Right click the client icon and select `Create Network Recovery Archive`.
3. Select the destination system and directory. Make sure that there is

17.385
Paula

enough space to store the compressed archive.

4. Select the volume groups you want to include into the archive.
5. The archive creation process will now begin. Check the status and log file on the Ignite-UX server by right clicking the icon and selecting `Client Status`.

NOTE Ignite-UX allows you to create bootable archive tape out of the compressed archive file. See the chapter *Create a Bootable Archive Tape via the Network* in the *Ignite-UX Administration Guide*.

Recovery

Recovery From the Backup Tape

To recover a system from a disaster using the bootable tape created by `make_tape_recovery` follow the steps below:

1. Replace the faulty hardware.
2. Make sure that the tape device is locally connected to the crashed HP-UX client and insert the medium with the archive you want to restore.
3. Boot from the prepared recovery tape. To do so, type in `SEARCH` at the boot admin menu to get a list of all available boot devices. Determine which one is the tape drive and type in the boot command: `boot <hardware path>` or `boot P<number>`.
4. The recovery process starts automatically.
5. After the recovery has completed successfully, restore additional user and application data using the standard Data Protector restore procedure.

Recovery From the Network

To recover an HP-UX client from a disaster via the network, follow the instructions on how to perform recovery with a Golden Image. Make sure you have selected the desired archive for the installation.

- On the Client
  1. Replace the faulty hardware.
  2. Boot the HP-UX client from the Ignite-UX server: `boot lan.<IP-address Ignite-UX server> install`

17.884
Paula

3. Select `Install HP-UX` from the `Welcome to Ignite-UX` screen.
4. Choose `Remote graphical interface running on the Ignite-UX server` on the UI Option screen.
5. Respond to the Network configuration dialogue.
6. The system is now prepared for a remote installation controlled from the Ignite-UX Server.

- On the Ignite-UX Server

7. Right click the client icon within the Ignite-UX GUI and select `Install Client - New Install`.
8. Under `Configurations:` select the `Recovery Archive` you want to install, check the settings (network, filesystem, time zone,...) and click the `Go!` button.
9. You can check the installation progress by right clicking the client icon and choosing `Client Status...`
10. After the recovery has completed successfully, restore additional user and application data using the standard Data Protector restore procedure.

17883
Paula

Disk Delivery Disaster Recovery of an UNIX Client

To perform a Disk Delivery Disaster Recovery of a UNIX client, connect a bootable disk that contains a minimal OS installation and Data Protector Disk Agent to the crashed system. The administrator has to ensure (before the disaster) that enough data has been collected to correctly format and partition the disk.

For details on supported operating systems, refer to the *HP OpenView Storage Data Protector Software Release Notes*.

Limitations

- This description does not cover the recovery of a cluster environment. Depending on the configuration of the cluster environment, additional steps and modification to the environment are necessary.
- RAID is not supported.
- Auxiliary disk should be prepared on a system of the same hardware class as the target system.

Preparation

Preparation for this disaster recovery method should be performed on several levels: gathering the information for your backup specification, preparing the disk, preparing your backup specification (pre-exec), and executing the backup. All of these preparatory steps are necessary before executing disaster recovery of the client.

Gathering Information

This section provides a list of items that need to be executed for each target system at backup time, in order to perform successful disaster recovery. If the information is collected as part of a pre-exec command, it is important to document the location of these files in the Disaster Recovery plan so that the information can be found once disaster strikes. Also version administration (there is a collection of the "auxiliary information" per backup) has to be considered.

- If the system that will be backed up has application processes active at low run levels, establish a state of *minimal activity* (modified *init 1 run level*) and enter the single user mode to prevent errors after recovery (see "Consistent and Relevant Backup" on page 444). Consult your operating system documentation for details.

HP-UX Example

1. Move some kill links from `/sbin/rc1.d` to `/sbin/rc0.d` and complement the changes for the boot-up section. The kill links include the basic services that would otherwise be suspended by moving to run level 1, and they are needed for the backup. For an example, see Appendix A, "Disaster Recovery: Move Kill Links on HP-UX 11.x," on page A-25.
2. Ensure that `rpcd` is configured on the system (configure the variable `RPCD=1` within the file `/etc/rc.config.d/dce`).

   This prepares the system so that it enters the state of minimal activity. The state can be characterized as follows:

   — `Init-1 (FS_mounted, hostname_set, date_set, syncer_running)`

   — Network must be running

   — The following processes should also be running: `inetd`, `rpcd`, `swagentd`

Solaris Example

1. Move the `rpc` kill link from `/etc/rc1.d` to `/etc/rc0.d` and complement the change for the boot-up section. The kill links include the basic services that would otherwise be suspended by moving to run level 1, and they are needed for the backup.
2. Ensure that `rpcbind` is configured on the system.

   This prepares the system so that it enters the state of minimal activity. The state can be characterized as follows:

   — `Init 1`

   — Network must be running

   — The following processes should also be running: `inetd`, `rpcbind`.

Tru64

— If the system is powered down, boot up the system and enter the System Reference Manual (SRM) console (the firmware console). Execute the following command from the SRM console to enter the

17.851
Paula

single user mode:

- — `boot -fl s` to boot using already generated vmunix file
- — `boot -fi genvmunix -fl s` to boot into the single user mode with the generic kernel.

— If the system is already powered up and running, change from the current run level to single-user mode by executing the following command: `init s`

AIX

— No action is required, because the `alt_disk_install` command, used to prepare the auxiliary disk, ensures consistent disk image without entering the state of minimal system activity.

Creating an Auxiliary Disk

- If you want to work with the auxiliary boot disk, you have to prepare it. Only one bootable auxiliary disk is required per site and platform. This disk has to contain the operating system and network configuration, and has to be bootable.

Preparing the Backup Specification

- Provide a Pre-exec script that performs the following:

— Collects all the necessary information about the environment and puts it in an available location in case of a disaster recovery. It is suggested to put it onto a different system which can be accessed easily. The information should cover:

- ✓ Physical and logical storage structure of the storage
- ✓ Current logical volume structure (for example, on HP-UX, using `vgcfgbackup` and `vgdisplay -v`)
- ✓ ServiceGuard configuration data, disk-mirroring, striping
- ✓ Filesystems and mountpoints overview `(for example, on HP-UX, using` bdf or copy of `/etc/fstab`)
- ✓ System paging space information, for example, on HP-UX, using the output of the `swapinfo` command
- ✓ I/O-structure overview `(for example, on HP-UX, using ioscan -fun` and `ioscan -fkn`)
- ✓ Client network settings

— An emergency copy of the data can also be put into the backup itself. If done so, the information has to then be extracted prior to the actual recovery.

17830
Paula

- — Consider logging out all users from the system.
- — Shut down all applications, unless the application data gets backed up separately, for example, using online database backup.
- — You may want to restrict network access to the system, so that no one can log on to the system while the backup is running (for example, on HP-UX, overwrite `inetd.sec` and use `inetd -c`).
- — If needed, enter the state of minimal system activity (for example, on HP-UX, use `sbin/init 1`; wait 60; check if run_level 1 is reached). Note that this is a modified "init 1" state.

- Provide a post-exec script that elevates the system to the standard run-level, restarts applications, and so on.
- Setup a backup specification for the client on the Data Protector Cell Manager. It should include all the discs (with disc discovery) and include the pre- and post-exec scripts.

Testing the Procedure

- Execute this backup procedure and repeat it on a regular basis, or at least at every major system configuration change, especially any change in the logical volume structure (for example, using LVM on HP-UX).

Recovery

This section describes how to restore a system to the state when the backup was done. You will need the following to successfully perform a Disk Delivery Disaster Recovery:

- A new hard disk to replace your crashed disk.
- An auxiliary disk containing the relevant operating system and the Data Protector agents.
- A successful full backup of the client that you want to recover.

The following steps need to be performed:

1. Replace the faulty disk with a new disk of comparable size.
2. Attach the auxiliary disk (which contains the relevant operating system and the Data Protector client) to the system and make it the boot device.
3. Boot from the auxiliary operating system.

17879
Paula

4. Reconstruct the logical volume structure if applicable (for example, using LVM on HP-UX). Use the saved data for the non-root volume groups (for example, with `vgcfgrestore` or SAM on HP-UX).
5. Additionally, the root volume group to be restored has to be created on the repaired disk (for example, using `vgimport` on HP-UX). It will not look like a root volume group during the restore process. This is because the OS from the auxiliary disk will be running. For more information on `vgimport`, see its man page.
6. Make the new disk bootable.
7. Reconstruct any other storage structures like mirror, striping, service guard, and so on from the data saved on a secondary storage device during backup.
8. Create the filesystems and mount them as required by the data from the backup; use similar but not the original mountpoint names (like `/etc_restore` for `/etc`, and so on).
9. Remove any files in the mountpoints to be restored, they must be clean.
10. Start the Data Protector GUI and open a connection to the Cell Manager. Import the system with the auxiliary disk into the cell.
11. Select the version from which you want to restore. First list all the required media for the restore and make sure they are available. Restore all the required mountpoints including the (future) root-volume to the system, using the option `Restore As` <*new_mountpoint*>. The root-volume from the backup is restored to the root-volume on the repaired disk. Nothing is restored to the currently-running auxiliary operating system on the auxiliary disk.
12. Shut down the system that was just restored.
13. Disconnect the auxiliary disk from the system.
14. Reboot the system from the new (or repaired) disk.

NOTE Instead of using an auxiliary disk, the new disk can also be temporarily connected to a client that has to have a Disk Agent installed. After being restored, it can be connected to the faulty system and booted.

17.848
Paula

Manual Disaster Recovery of an UNIX Cell Manager

Manual Disaster Recovery is a basic method, that involves recovering the system by reinstalling it in the same way as it was initially installed. In addition, Data Protector is used to then restore all files, including the operating system.

Limitation

For details on supported operating systems, refer to the *HP OpenView Storage Data Protector Software Release Notes*.

This description does not cover the recovery of a cluster environment. Depending on the configuration of the cluster environment, additional steps and modification to the environment are necessary.

Preparation

Perform the same preparatory steps without the steps pertaining to the auxiliary disk, as for Disk Delivery Disaster Recovery of an HP-UX or Solaris client. See "Preparation" on page 507 for reference. In addition to completing those steps, you also have to complete the following:

1. The IDB has to be backed up regularly, ideally in a separate backup specification, scheduled after the backup of the Cell Manager.
2. The IDB and configuration backup must run to a specific device located on the Cell Manager system, to make the administrator aware that the medium in the device contains the most recent version of the IDB.

Recovery

Use the following method to recover your UNIX Cell Manager.

Prerequisites

You will need the following to successfully perform a disaster recovery:

- Media containing the last valid known backup of the root partition of the Cell Manager and IDB.

17.877

- A device connected to the Cell Manager.

Procedure

The following steps need to be performed to recover a Cell Manager:

1. Replace the crashed disk.
2. Boot your system from the installation media of your operating system.
3. Reinstall the operating system. Refer to your system administrator's manual for instructions. During the installation, using the data gathered during the preparation phase (pre-exec script), re-create and configure the physical and logical storage structure of the storage, logical volume structure, filesystem and mountpoints, network settings and other.
4. Reinstall the Data Protector on the Cell Manager.
5. Restore the latest backup of your database and `/etc/opt/omni` to a temporary directory. This simplifies the restore of all other files from media. Note that you cannot restore the database directly. See Chapter 6, "Restoring Data," for instructions. This includes stopping all Data Protector processes with the `/opt/omni/sbin/omnisv -stop` command. This ensures that no files will be in use.
6. Remove the `/etc/opt/omni/` directory and replace it with the `/etc/opt/omni` directory from the temporary area. This re-creates the previous configuration.
7. Start Data Protector processes with the `/opt/omni/sbin/omnisv -start` command.
8. Start the Data Protector user interface and restore all the files used from your backup.
9. Reboot the system.

Your Cell Manager should now be successfully recovered.

17.876
Fauta

Troubleshooting Disaster Recovery on Windows

This section provides explanation of the steps you will need to take if you happen to encounter problems with Manual, Disk Delivery, Enhanced, or One Button Disaster Recovery procedures on Windows systems.

General Troubleshooting

Problem

Problems Logging on to the System After Disaster Recovery Finishes

You may receive the following error message after the system is recovered:

```
The system cannot log you on to this domain, because the
system's computer account in its primary domain is
missing or the password on that account is incorrect.
```

This type of message is usually caused by one of the following reasons:

- After collecting all information for successful disaster recovery (including full backup), you reinstalled Windows and (re)inserted into the offending domain.
- After collecting all information for successful disaster recovery (including full backup), you removed your system from the offending domain and later (re)inserted it into the same or some other domain.

  In cases like this, Windows generates new system security information, which is incompatible with information that is restored during disaster recovery. The solution is the following:

Action

1. Log on to the system locally with an Administrator account.
2. In the `Control Panel`, click `Network` and, using the `Identification` tab, remove the system from its current domain to a temporary workgroup (for example, TEMP). After this is done, reinsert the system into the domain from which it was previously removed. You need a domain administrator's password.
3. After the computer is again in the proper domain, click `OK` in the `Network` window. Windows will force you to reboot the system.

17875 Paula

4. To update this new state with disaster recovery, you should perform all necessary procedures (collecting system data, backup) once more, as described in the "Preparing for a Disaster Recovery" section.

Troubleshooting Assisted Manual Disaster Recovery

Problem

Drstart reports: "Can not copy <filename>"

This error is reported because the `drstart` utility can not copy the specified file. One of the reasons may be that the file is locked by the system. For example, if `drstart` cannot copy `omniinet.exe`, it might be because the `Inet` service is already running. This is not a normal scenario and should not happen after a clean install.

Action

A dialog box will appear asking you whether you would like to proceed with copying the rest of the files. If you click `Yes`, `drstart` will skip the locked file and continue copying other files. This will solve the problem if the file is locked by the system, as the process required for the disaster recovery is already running and therefore the file does not need to be copied.

You can also close the `drstart` utility by clicking the `Abort` button.

Troubleshooting Disk Delivery Disaster Recovery

Problem

"Cannot Find Physical Location of Drives Selected for Disk Delivery"

When using the Disk Delivery method for disaster recovery, it is possible that you will receive the following error: "Cannot find physical location of drives selected for disk delivery." Objects will be restored when creating a partition on the new disk if you select a drive letter that has not been used before. The better solution would be:

Action

Disaster recovery checks disk information before restoring objects. An internal function reads the Registry value Information, which is created by the Disk Administrator. If the Disk Administrator is started several times, the `Information` value becomes corrupted (format is changed during update) - the parsers fail in such cases. If you delete the `HKEY_LOCAL_MACHINE\SYSTEM\DISK Information` key and restart the Disk Administrator, the function will succeed.

17874
Paula

Problem **"No Operating System Found"**

Action After performing disaster recovery, if the final boot of a Windows NT system fails with "`No Operating System Found`", check the `boot.ini` file for information about where the partition information is located. See Step 4 in the section "Recovery" on page 461 for additional information.

Problem **Disk Delivery Disaster Recovery of a Media Agent Client**

If you are performing a Disk Delivery disaster recovery, Data Protector first tries to connect to the original client where the backup device was attached (Media Agent client) in order to use the same device for restore. However, when you are performing Disk Delivery disaster recovery of the crashed Media Agent client where the backup has been made, Data Protector will not be able to connect to it and will proceed with offline restore and search for a local device for the restore. If there is no local device attached, Data Protector will issue a notification that there is no local device attached and will abort the disaster recovery.

Action There are three methods to avoid this:

- Move the media to another pool. This way you assign the media to the new device. Then proceed with Disk Delivery disaster recovery.
- The third method involves preparation prior to the disaster. If you have two Media Agent clients in the cell, you can back up of the first Media Agent client to another and vice versa before the disaster to avoid problems when performing Disk Delivery disaster recovery of a Media Agent client.

Troubleshooting EADR and OBDR

Problem **Automatic DR information could not be collected**

When using EADR or OBDR, it is possible that you will receive the following error: "`Automatic DR information could not be collected. Aborting the collecting of system recovery data`"

Action

- Check if all storage devices are configured correctly. If Device Manager reports a device as "`Unknown Device`", you have to install the proper device drivers before you can perform EADR/OBDR. A similar entry would appear in `autodr.log` (located in

17873
Paula

<Data_Protector_home>\tmp) if improperly configured storage devices are attached to your system:

```
DRIM_WIN_ERROR 13 SetupDiGetDeviceRegistryProperty
```

- There must be enough registry space available. It is recommended to set the maximum registry size to at least twice that of the current registry size. If there is not enough registry space available, a similar entry would appear in the autodr.log:

```
ERROR registry  'Exception while saving registry'
...
WindowsError: [Errno 1450] Insufficient system resources
exist to complete the requested service.
```

If the problem persists, uninstall the Data Protector Automatic Disaster Recovery component (so that at least Manual Disaster Recovery and Disk Delivery Disaster Recovery will work) and contact technical support.

Problem

Some Non-critical Errors Were Detected

When using EADR or OBDR, it is possible that you will receive the following error: "Some non-critical errors were detected during the collecting of Automatic DR data. Please review the Automatic DR log file."

A non-critical error detected during the execution of the Automatic Disaster Recovery module, means that such backup can most likely still be used for disaster-recovery purposes. Possible reasons for non-critical errors are stored in autodr.log (located in *<Data_Protector_home>*\tmp):

Action

- Services or drivers outside of the <%SystemRoot%> folder (for example, virus scanners). Autodr.log would contain a similar error message:

```
ERROR safeboot  'unsupported location' 'intercheck support 06' 2
u'\\??\\D:\\Program Files\\Sophos SWEEP for NT\\icntst06.sys'.
```

You can ignore this error message, as it does not affect the success of disaster recovery.

- Boot disk larger than 7.8 GB on Windows NT. The autodr.log file would

```
contain a similar error message:
INFO  storage  'check_system_partition' 'boot' u'C:'
ERROR storage  'boot' 'volume' u'C:' 'last_cyl' 1105 ' >= 1024.'
```

17.842
Paula

```
'System may not boot after disaster recovery.
```

If the boot and system volumes are both physically located below 7.8 GB, this error message can be ignored. The boot and system partitions on Windows NT 4.0 must namely be physically bellow the first 7.8 GB due to the operating system limitations. Refer to MSDN Q224526.

Problem

Blank Screen During Recovery

Certain system configurations have been encountered where the video display does not work, if Windows is started in safe mode. This error is not related to Data Protector and can occur even if only Windows is installed.

Action

If the screen is blank during disaster recovery, this does not mean that the recovery has failed. You can monitor the progress of disaster recovery on the Cell Manager or use `ping` and `telnet 5555` (or appropriate) commands from another client to see if the target system responds. Other indicators that the recovery is still in progress are that the device is working and that hard disk lights are blinking.

If the target system responds to `ping` and `telnet 5555` commands, but hard disk lights are not blinking and the device is not active, it is possible that the auto logon failed. Press **Enter** to log on using the administrator's account with a blank password.

The display on the restored system will then work just as it did at backup time.

Problem

Network is Not Available During Restore

Action

Ensure that the problem is not with switch, cables, etc. Another possibility is also that the DNS server (as configured at backup time) is offline during the restore. Since the configuration of the DR OS is the same as at backup time, the network will not be available. In this case perform offline restore and change the DNS settings after recovery. You can also edit the registry (`HKey_Local_Machine\SYSTEM\CurrentControlSet\Services\Tcpip\Parameters`) before Phase 2 is started. In this case you have to reboot before Phase 2 for the changes to take effect. After Phase 2 finishes, you can correct the settings before Phase 3 can be started.

WARNING **Editing the registry incorrectly can result in failed disaster recovery.**

Problem **Auto Logon Does Not Work**

Action Sometimes auto logon does not work and you have to manually log on using an administrator's account with a blank password.

Problem **Computer Freezes During EADR**

Action Check if the CD is readable. Do not reuse CD-RWs too many times.

Problem **Cannot Create a CD ISO Image for EADR of MSCS**

Action The quorum disk has to be backed up in order to be able to create an CD ISO Image.

17.870
Paula

ANEXO UNIDADE DE BACKUP ROBOTIZADO
PARTE D/7

17.869
Paula

17.868
Laila

11 Customizing the Data Protector Environment

17.867
Paula

In This Chapter

This chapter describes how you can customize Data Protector to better suit your needs. The chapter consists of the following sections:

"Global Options File" on page 523

"Using Omnirc Options" on page 525

"Firewall Support" on page 528

IMPORTANT For specific information on Data Protector limitations and recommendations, see the *HP OpenView Storage Data Protector Software Release Notes*. For details about adding security to your Data Protector cell, refer to the *HP OpenView Storage Data Protector Installation and Licensing Guide*.

17.866
Paula

Global Options File

Global options affect the entire Data Protector cell, and cover various aspects of Data Protector, such as time-outs and limits. All global options are described in the global options file, which you can edit in order to customize Data Protector. It is located in the /etc/opt/omni/options directory on the UNIX Cell Manager and in the *<Data_Protector_home>*\config\options directory on Windows Cell Manager. The file is named global.

How to Use Global Options

Each option has a hash mark, or pound sign (#), which comments out the option and provides an explanation of the option in the text following the hash mark. For options not described in this guide, refer to the file itself.

To use a global option, uncomment the line that has the option name and set an appropriate value. To uncomment a line, simply remove the '#' mark.

NOTE

Most users should be able to operate Data Protector without changing the global options.

Most Often Used Variables

The following list includes the most often used global variables. See the Global Options file for a complete description.

- MediaView: Changes the fields and their order in the Media Management context.
- MaxBSessions: Changes the default limit of five concurrent backups.
- InitOnLoosePolicy: Allows Data Protector to automatically initialize blank or unknown tapes under a loose media policy.
- MaxMAperSM: Increases the default limit of concurrent devices per backup session. (Maximum device concurrency is 32.)
- DCDirAllocation: Determines the algorithm used for selecting into which dcbf directory a new detail catalog binary file goes. Three algorithms are available: fill in sequence (default), balance size, and

balance number.

- DailyMaintenanceTime: Determines the time after which the daily maintenance tasks can begin, using the twenty-four hour clock notation. By default, this time is set to 12:00 (Noon). For a list of daily maintenance tasks, refer to "Data Protector Checking and Maintenance Mechanism" on page 605.
- DailyCheckTime: Determines the time after which daily check can begin, using the twenty-four hour clock notation. By default, this time is set to 12:30 P.M. If you do not wish to perform a daily check, you can disable it. For a list of daily check tasks, refer to "Data Protector Checking and Maintenance Mechanism" on page 605.

17.864
Paula

Using Omnirc Options

The omnirc options are most useful for troubleshooting or overriding other settings affecting the behavior of the Data Protector client only. However, even advanced users should not use them unless their operating environment demands it. The Disk Agents and Media Agents use the values of these options.

These options are found in the following locations:

Locations

- `/opt/omni/.omnirc` on HP-UX and Solaris clients
- `/usr/omni/.omnirc` on other UNIX clients
- *<Data_Protector_home>*`\omnirc` on Windows clients
- `sys:\usr\omni\omnirc` on Novell NetWare clients

How to Use Omnirc Options?

Installation provides a template for the omnirc file (`.omnirc.TMPL` or `omnirc.TMPL`, depending on the platform). This file is not active. To create an active omnirc file, copy the template file to `omnirc` (or `.omnirc`) and edit it. To use a specific option, uncomment the line (remove the '#' character) and edit the value if necessary.

- When creating the `omnirc` file (either by copying or by using an editor), verify its permissions. On UNIX, permissions will be set according to your umask settings and may be such that some processes may be unable to read the file.

  Set the permissions to 644 manually.

- When changing the `omnirc` file, you have to restart the Data Protector services/daemons on the Data Protector client where you modified the `omnirc` file. This is mandatory for the `crs` daemon on UNIX and recommended for Data Protector CRS and `Inet` services on Windows. Specifically on Windows, restarting is not required when adding or changing entries, only when removing entries (or renaming the file).

Most Often Used Variables

The most often used omnirc variables include:

- **OB2BLKPADDING_n:** This is a set of variables that can be used to specify the number of empty blocks written to the media at the initialization time.

17865
Paula

- **OB2DEVSLEEP:** Changes the sleep time between each retry while loading a device.
- **OB2ENCODE:** Allows a user to always turn on data encoding, regardless how the backup options are set in the backup specification.
- **OB2OEXECOFF:** Allows a user to restrict or disable any object pre- and post-exec scripts defined in backup specifications for a specific client.
- **OB2INCRDIFFTIME** and **OB2CHECKCHANGETIME:**

  The **OB2INCRDIFFTIME** variable specifies a time period (in minutes) imposed after checking the inode change time for incremental backups. The latter one is specified by the **OB2CHECKCHANGETIME** variable.

  It means that the referential time, the time of the previous backup, received from the Session Manager is incremented by the specified period and the **OB2CHECKCHANGETIME** value is checked against incremented time to qualify for backup.

  OB2INCRDIFFTIME is specified in minutes and works only together with **OB2CHECKCHANGETIME**.
- **OB2RECONNECT_ACK:** Defines how long Data Protector should wait for the message of acknowledgment (default 1200 seconds). In other words, if the agent does not get an acknowledgment in OB2RECONNECT_ACK seconds, it will assume that the socket connection is no longer valid.
- **OB2RECONNECT_RETRY:** Defines how long Data Protector should wait before trying to reconnect after a socket connection has been broken (the default is 600 seconds). In other words, the WAN line between the Backup Session Manager and agents cannot be down more than OB2RECONNECT_RETRY seconds.
- **OB2REXECOFF:** Allows a user to disable any remote session pre- and post-exec scripts for a specific client.
- **OB2SHMEM_IPCGLOBAL:** This option should be set to 1 on HP-UX clients that have both the Disk Agent and the Media Agent installed, in case the following error occurs during the backup:

  ```
  Cannot allocate/attach shared memory (IPC Cannot Allocate
  Shared Memory Segment)

  System error: [13] Permission denied) => aborting
  ```

17.862
Paula

- **OB2VXDIRECT:** Enables direct (without cache) reading for Advanced VxFS filesystems, as well as improving performance.
- **OB2PORTRANGE:** This option limits the range of port numbers that Data Protector uses when allocating listen ports dynamically. This option is typically set to enable the administration of a cell through a firewall. Note that the firewall needs to be configured separately and that the specified range does not affect the `Inet` listen port.

Example

```
OB2PORTRANGE=40000-40199
```

This sets the port range to ports from 40000 to 40199.

- **OB2PORTRANGESPEC:** This option allows you to specify a range of port numbers for every binary. This mechanism gives you more control over the ranges and helps to keep their sizes smaller. Note that the firewall needs to be configured separately and that the specified range does not affect the `Inet` listen port.

  For configuration examples, refer to "Firewall Support" on page 528.

Firewall Support

This section describes how to configure Data Protector in an environment where the Data Protector processes communicate across a firewall.

Communication in Data Protector

Data Protector processes communicate using TCP/IP connections. Every Data Protector system accepts connections on port 5555 by default. In addition, some processes dynamically allocate ports on which they accept connections from other Data Protector processes.

To enable Data Protector processes to communicate across a firewall, Data Protector allows you to limit the range of port numbers from which dynamically allocated ports are selected. Port ranges are defined on a per system basis. It is possible to define a port range for all Data Protector processes on a specific system, as well as to define a port range for a specific Data Protector agent only.

Configuration Mechanism

The port allocation behavior can be configured through two `omnirc` variables: `OB2PORTRANGE` and `OB2PORTRANGESPEC`. By default, both variables are not set and ports are assigned dynamically by the operating system.

Limiting the Range of Port Numbers

For All Data Protector Processes

You can limit the port range for all Data Protector processes on a system by using the `OB2PORTRANGE` variable in the `omnirc` file:

```
OB2PORTRANGE=<start_port>-<end_port>
```

Data Protector processes use dynamically allocated ports and select ports from this range. The port range is allocated by taking the first available port, starting with port "start_port". If there is no available port within the specified range, the port allocation fails and the requested operation is not done. Refer to Table 11-1 on page 530 for information on port consumption.

17.860
Paula

NOTE

The OB2PORTRANGE variable only applies to dynamically allocated ports. It does not affect the usage of the default Data Protector port number 5555.

Defining a port range for the Data Protector processes limits the port usage of Data Protector. It does not prevent other applications from allocating ports from this range as well.

For a Specific Data Protector Agent

In many cases it is not required that all Data Protector agents communicate across a firewall. For example, one specific agent can be outside a firewall, while all other components are inside of it. In such environments it is useful to limit the range of port numbers only for the specific agent. This allows you to define a much smaller port range and so reduce the need of open ports through the firewall.

You can limit the port range on a system on which a specific agent runs by using the OB2PORTRANGESPEC variable in the omnirc file:

```
OB2PORTRANGESPEC=<AGENT>:<start_port>-<end_port>;...
```

All agent processes check the OB2PORTRANGESPEC for range restrictions. If there is a range defined for an agent process, all dynamically allocated ports select from this specified range. The port range is allocated by taking the first available port, starting with port "start_port". If there is no available port within the specified range, the port allocation fails and the requested operation is not done. See "Examples of Configuring Data Protector in Firewall Environments" on page 535 for information on how to calculate the required range of port numbers.

17859
Paula

The table below lists all possible Data Protector agent identifiers that can be used in the OB2PORTRANGESPEC variable. Note that agent processes that do not dynamically allocate listen ports are not listed in the following table.

Table 11-1 **Agent Identifiers**

Data Protector Component	Agent Identifier	Description	Port Consumption
Cell Manager	BSM	Backup Session Manager	1 port per concurrently running BSM
	RSM	Restore Session Manager	1 port per concurrently running RSM
	DBSM	Database Session Manager	1 port per concurrently running DBSM
	xSM	Wildcard matching all Session Managers	1[a] + 1 port per concurrently running Session Manager
	MMD	Media Management Daemon	1 port
	CRS	Cell Request Server Service	1 port
Media Agent	BMA-NET	Backup Media Agent[b]	1 port per concurrently running Media Agent
	RMA-NET	Restore Media Agent[b]	1 port per concurrently running Media Agent
	xMA-NET	Wildcard matching all Media Agents[b]	1 port per concurrently running Media Agent

[a] This additional port is required during database operations such as filename purges or database upgrades.

[b] BMA and RMA fork two processes, the main process and a NetIO process. The listen port is allocated by the BMA-NET / RMA-NET process.

NOTE

The OB2PORTRANGESPEC variable only applies to dynamically allocated ports. It does not affect the usage of the default Data Protector port number 5555.

Defining a port range for a specific Data Protector agent process limits the port usage of this agent. It does not prevent other processes (applications or other Data Protector agents) from allocating ports from this range as well.

Using Both Variables Together

If both variables OB2PORTRANGESPEC and OB2PORTRANGE are set, OB2PORTRANGESPEC overrides the settings of OB2PORTRANGE.

For example, the setting

```
OB2PORTRANGESPEC=BMA-NET:18000-18009

OB2PORTRANGE=22000-22499
```

limits the port range used by the Media Agent to port numbers 18000-18009, while all other Data Protector processes use port numbers from the range 22000-22499.

By using both variables it is possible to force a specific agent to use only a dedicated port range (OB2PORTRANGESPEC) and, at the same time, prevent other Data Protector processes from selecting port numbers from this range.

Port Usage in Data Protector

The following section provides two tables that describe the port requirements of the different Data Protector components. Table 11-2 breaks down the different Data Protector components and shows to which other components they can connect. It also defines the destination specification for the firewall rules. Table 11-3 gives the same list of components but shows from which other components they can accept connections. It also determines the source port of the firewall rule.

The following table provides a list of all Data Protector components. The first two columns list the process identifiers and their listen ports. The last two columns list all applicable connecting processes.

Table 11-2

Listening Component		**Connecting Component**	
Process	**Port**	**Process**	**Source Port**
Cell Manager			
`Inet`	5555	Application Agent	N/A[a]
		GUI/CLI	N/A[a]
CRS	Dynamic	Application Agent	N/A[a]
		GUI/CLI	N/A[a]
MMD	Dynamic	xSM	N/A[a]
		CLI (from CM)	N/A[a]
xSM	Dynamic	GUI/CLI	N/A[a]
		xMA[b]	N/A[a]
		xDA[b]	N/A[a]
		Application Agent	N/A[a]
Disk Agent			
`Inet`	5555	xSM	N/A[a]
xDA	Does not accept connections		
Media Agent			
`Inet`	5555	xSM	N/A[a]
xMA	Does not accept connections		
xMA-NET	Dynamic	xDA	N/A[a]
		Application Agent	N/A[a]

17836
Paula

Table 11-2

Listening Component		**Connecting Component**	
Process	**Port**	**Process**	**Source Port**
Application Host			
`Inet`	5555	xSM	N/A[a]
Application Agent	Does not accept connections		

a. The source port of a connection is always assigned by the operating system and cannot be limited to a specific range.

b. Only for backup sessions with the reconnect feature enabled. The Disk Agent and the Media Agent communicate with the Cell Manager using the existing TCP connection. The connection in this column is only established after the original connection is broken.

When writing the firewall configuration rules, the process in the first column must be able to accept new TCP connections (SYN bit set) on the ports defined in the second column, from the process listed in the third column.

In addition, the process listed in the first column must be able to reply to the process in the third column on the existing TCP connection (SYN bit not set).

For example, the `Inet` process on a Media Agent system must be able to accept new TCP connections from the Cell Manager on port 5555. The Media Agent must be able to reply to the Cell Manager using the existing TCP connection. It is not required that the Media Agent is capable of opening a TCP connection.

The following table provides a list of all Data Protector components. The first two columns list all applicable connecting processes, while the last two columns list the process identifiers and their listen ports. Processes that do not initiate connections are not listed (for example, `Inet`).

17855

Table 11-3

Connecting Component		Listening Component	
Process	**Port**	**Process**	**Port**
Cell Manager			
xSM	N/A[a]	xMA[b]	5555
	N/A[a]	xDA[b]	5555
	N/A[a]	Application Agent[b]	5555
	N/A[a]	MMD[c]	Dynamic
User Interface			
GUI/CLI	N/A[a]	`Inet` on CM	5555
	N/A[a]	CRS	Dynamic
	N/A[a]	BSM	Dynamic
	N/A[a]	RSM	Dynamic
	N/A[a]	MSM	Dynamic
	N/A[a]	DBSM	Dynamic
CLI (Cell Manager only)	N/A[a]	MMD	Dynamic
Disk Agent			
xDA	N/A[a]	xMA-NET	Dynamic
	N/A[a]	xSM[d]	Dynamic
Media Agent			
xMA	N/A	xSM[d]	Dynamic
	N/A[a]	UMA[b, e]	5555
Application Agents			

17.854
Paula

Table 11-3

Connecting Component		**Listening Component**	
Process	**Port**	**Process**	**Port**
Application Agent	N/A[a]	`Inet` on CM	5555
	N/A[a]	CRS	Dynamic
	N/A[a]	RSM	Dynamic
	N/A[a]	BSM	Dynamic
	N/A[a]	xMA-NET	Dynamic

a. The source port of a connection is always assigned by the operating system and cannot be limited to a specific range.

b. To be more precise, it is the `Inet` process that accepts the connection on port 5555 and then starts the requested agent process. The agent process inherits the connection.

c. This applies only to the MMD on the system running the CMMDB in a Manager-of-Managers (MoM) environment.

d. Only for backup sessions with the reconnect feature enabled.

e. Connections to the Utility Media Agent (UMA) are only required when sharing a library across several systems.

Examples of Configuring Data Protector in Firewall Environments

The following section provides examples on how to configure Data Protector in four different firewall environments.

Example 1: Disk Agent and Media Agent Installed Outside, Other Components Installed Inside a Firewall

You can configure your backup environment so that the Cell Manager and GUI are in the intranet and some Disk Agents and Media Agents are in the Demilitarized Zone (DMZ):

17.853
Paula

Figure 11-1 **Configuration Diagram**

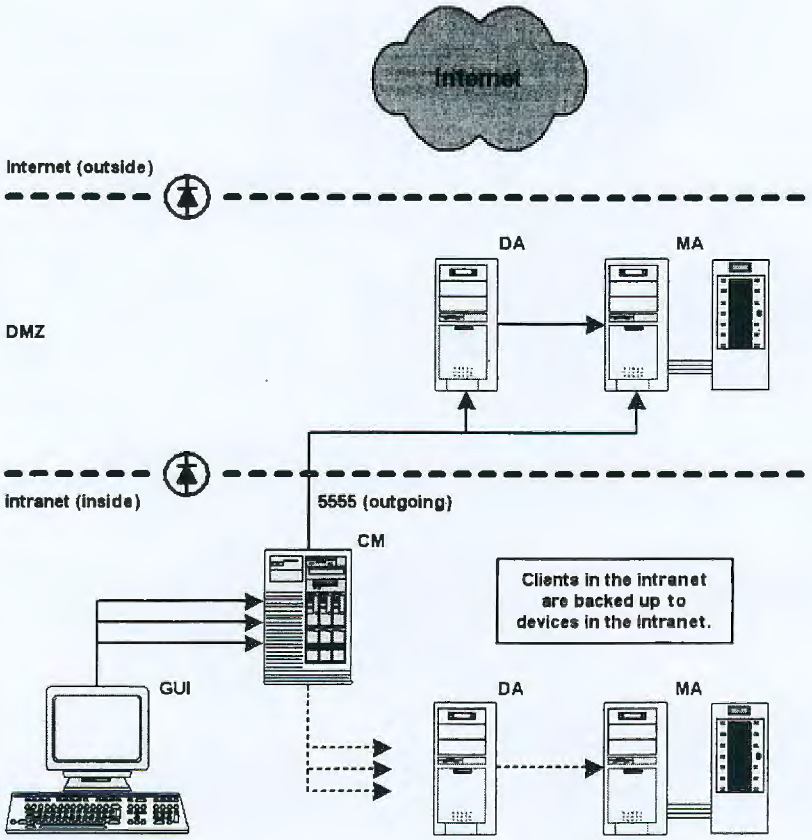

The following two items define the port range settings for this configuration:

1. In order to determine which processes need to communicate across the firewall, see Table 11-2 for the Disk Agent and Media Agent. It shows that the Disk Agent and Media Agent need to accept connections from the Session Manager on port 5555. This leads to the following rules for the firewall:

   ✓ Allow connections from the CM system to port 5555 on the DA system

17.552
Paula

✓ Allow connections from the CM system to port 5555 on the MA system

This table also shows that the Media Agent needs to accept connections from the Disk Agent. However, since these two agents do not communicate through the firewall, you do not need to define a firewall rule for them.

2. See also Table 11-3 for the Disk Agent and Media Agent.

   This table also shows that both agents may connect to the Session Manager and that the Media Agent may need to connect to a utility Media Agent (UMA). However, this only occurs when shared tape libraries are used or the `Reconnect broken connections` option is enabled. See "Backup Specification Options" on page 236 for information on this option.

Port Range Settings

Since all connections that need to go through the firewall connect to the fixed port number 5555, you do not need to define `OB2PORTRANGE` or `OB2PORTRANGESPEC` variables in this environment.

Limitations

- Remote installation of clients across the firewall is not supported. You need to install clients locally in the DMZ.
- This cell can back up clients in the DMZ, as well as clients in the intranet. However, each group of clients must be backed up to devices configured on clients that are on the same side of the firewall.

IMPORTANT

If your firewall does not restrict connections from the intranet to the DMZ, it is possible to back up clients in the intranet to devices configured on clients in the DMZ. However, this is not recommended, as the data backed up in this way becomes more vulnerable.

- If a device in the DMZ has robotics configured on a separate client, this client must also be in the DMZ.
- This setup does not allow the backup of databases or applications using Application Agents on the clients in the DMZ. For details on Application Agents in the DMZ, refer to "Example 4: Application Agent and Media Agent Installed Outside, Other Components Installed Inside a Firewall" on page 543.

17851
Paula

Example 2: Disk Agent Installed Outside, Other Components Installed Inside a Firewall

You can configure your backup environment so that the Cell Manager, Media Agent, and GUI are in the intranet and some Disk Agents are in the DMZ:

Figure 11-2 **Configuration Diagram**

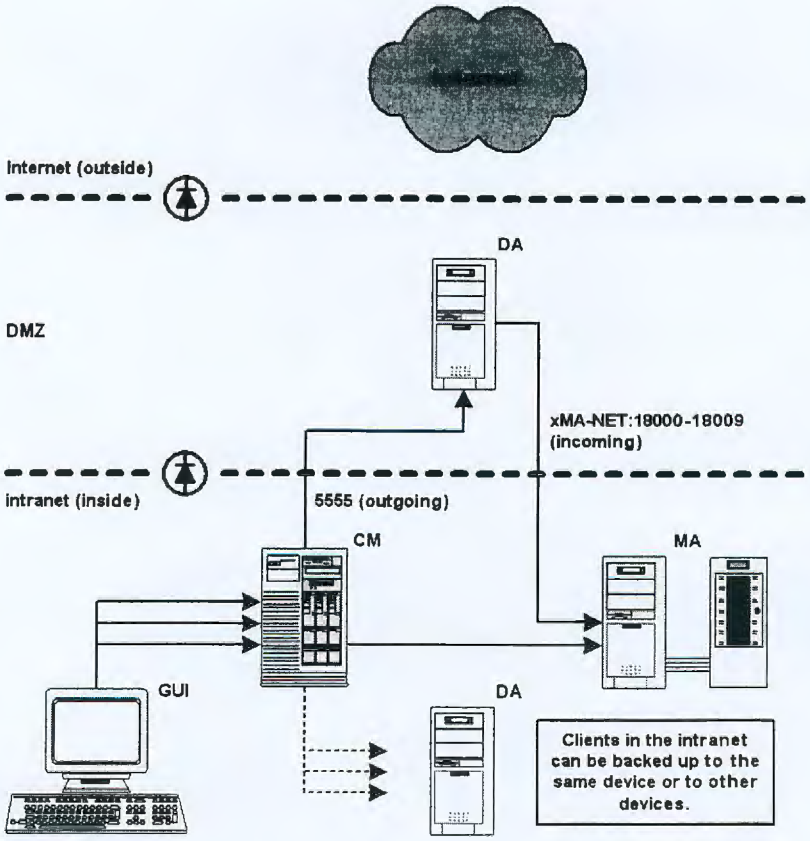

The following three items define the port range settings for this configuration:

17850
Paula

1. In order to determine which processes need to communicate across the firewall, see Table 11-2 (Disk Agent column). It shows that the Disk Agent needs to accept connections from the Session Manager on port 5555. This leads to the following rule for the firewall:

   ✓ Allow connections from the CM system to port 5555 on the DA system

2. See also Table 11-3 for the Disk Agent. It shows that the Disk Agent connects to a dynamically allocated port on the Media Agent. Since you do not want to open the firewall for communication between the Disk and Media Agent in general, you need to limit the range of ports from which the Media Agent can allocate a listen port.

   See Table 11-1 for the port consumption of the Media Agent. The Media Agent requires only one port per running Media Agent. For example, if you have four tape devices connected, you may have four Media Agents running in parallel. This means that you need at least four ports available. However, since other processes may allocate ports from this range as well, you should specify a range of about ten ports on the MA system:

   ```
   OB2PORTRANGESPEC=xMA-NET:18000-18009
   ```

   This leads to the following firewall rule for the communication with the Media Agent:

   ✓ Allow connections from the DA system to port 18000-18009 on the MA system

NOTE This rule allows connections from the DMZ to the intranet, which is a potential security risk.

3. Table 11-3 also shows that the Disk Agent needs to connect to the Session Manager (BSM/RSM) when the `Reconnect broken connections` option is enabled. You can specify a required port range on the CM system analogous to the previous item.

   ```
   OB2PORTRANGESPEC=xSM:20100-20199
   ```

17849
Paula

NOTE All Session Managers allocate ports from this range, not only the one communicating through the firewall.

Limitations

- Remote installation of clients across the firewall is not supported. You need to install clients locally in the DMZ.
- This setup does not allow the backup of databases or applications using Application Agents on the clients in the DMZ. For details on Application Agents in the DMZ, refer to "Example 4: Application Agent and Media Agent Installed Outside, Other Components Installed Inside a Firewall" on page 543.

Example 3: GUI Installed Inside, Other Components Installed Outside a Firewall

You can configure your backup environment so that the entire cell is in the DMZ and only the Graphical User Interface is in the intranet:

17.848

Figure 11-3 **Configuration Diagram**

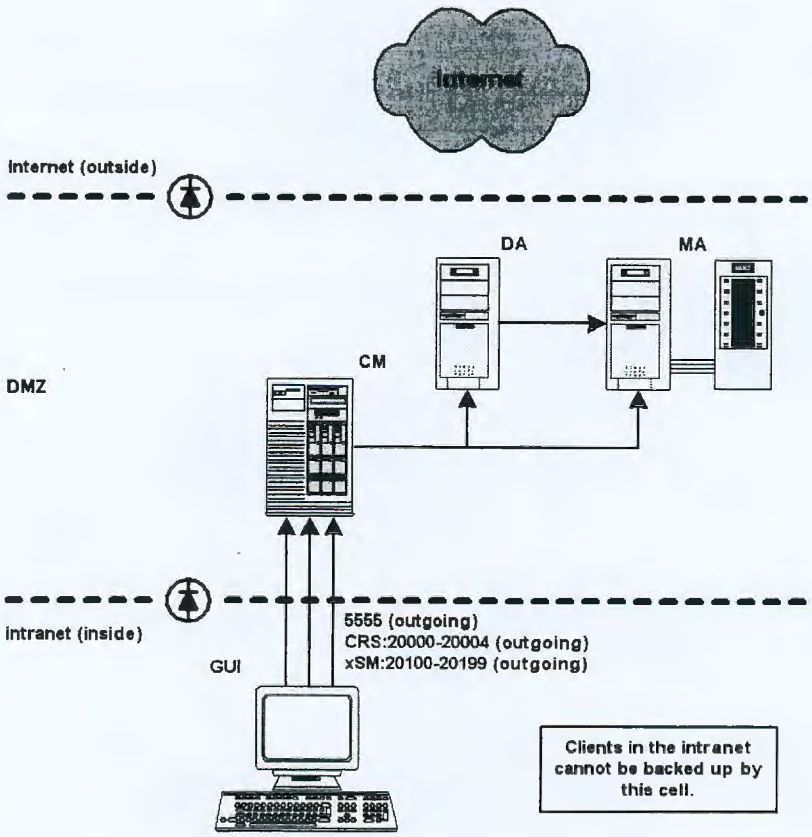
Internet
Internet (outside)
DA
MA
CM
DMZ
intranet (inside)
GUI
5555 (outgoing)
CRS:20000-20004 (outgoing)
xSM:20100-20199 (outgoing)
Clients in the intranet cannot be backed up by this cell.

The following three items define the port range settings for this configuration:

1. Table 11-2 and Table 11-3 show that the GUI does not accept any connections. However, it needs to connect to the following processes on the Cell Manager:

Table 11-4

Process	**Port**
`Inet`	5555
CRS	Dynamic
BSM	Dynamic
RSM	Dynamic
MSM	Dynamic
DBSM	Dynamic

This leads to the following firewall rule for the connection to the `Inet` listen port:

- ✓ Allow connections from the GUI system to port 5555 on the CM system

2. Table 11-1 shows that CRS requires only one port. However, since other processes may allocate ports from this range as well, you should specify a range of about five ports on the CM system. The port range could be defined as follows:

   ```
   OB2PORTRANGESPEC=CRS:20000-20004
   ```

   The resulting firewall rule for the connection to the CRS process is:

   - ✓ Allow connections from the GUI system to ports 20000-20004 on the CM system

3. For the Session Manager, the situation is much more complex. Every Session Manager requires only one port. However, the number of Session Managers (BSM, RSM, MSM, DBSM) heavily depends on the backup environment. The minimum requirement can be estimated with the following formula:

$$NoOfPorts = NoOfConcurrentSessions + NoOfConnectingGUIs$$

17.846
Paula

Port Range Settings on the Cell Manager

For example, if there are 25 backup and five restore sessions running and two GUIs opened, you need to have at least 32 ports available. However, since other processes may allocate ports from this range as well, you should specify a range of about 100 ports on the CM system. The port range could be defined as follows:

```
OB2PORTRANGESPEC=xSM:20100-20199
```

or:

```
OB2PORTRANGESPEC=BSM:20100-20139;RSM:20140-20149;DBSM:201
50-20199
```

Limitations

For this configúration almost all Data Protector functionality is available, including remote installation and online backup of databases and applications.

- This cell cannot be a part of a MoM environment if centralized media management or centralized licensing is used and the MoM cell is inside.
- All backup clients must be in the DMZ. The GUI client cannot be backed up by the Media Agent from the DMZ. The GUI can also be run from a client that is a member of another cell located in the intranet, provided that both cells use the same `Inet` listen port.

Example 4: Application Agent and Media Agent Installed Outside, Other Components Installed Inside a Firewall

You can configure your backup environment so that the Cell Manager and GUI are in the intranet and some Application Agents (SAP R/3, Oracle ...) and Media Agents are in the DMZ:

17.845
Paula

Figure 11-4 **Configuration Diagram**

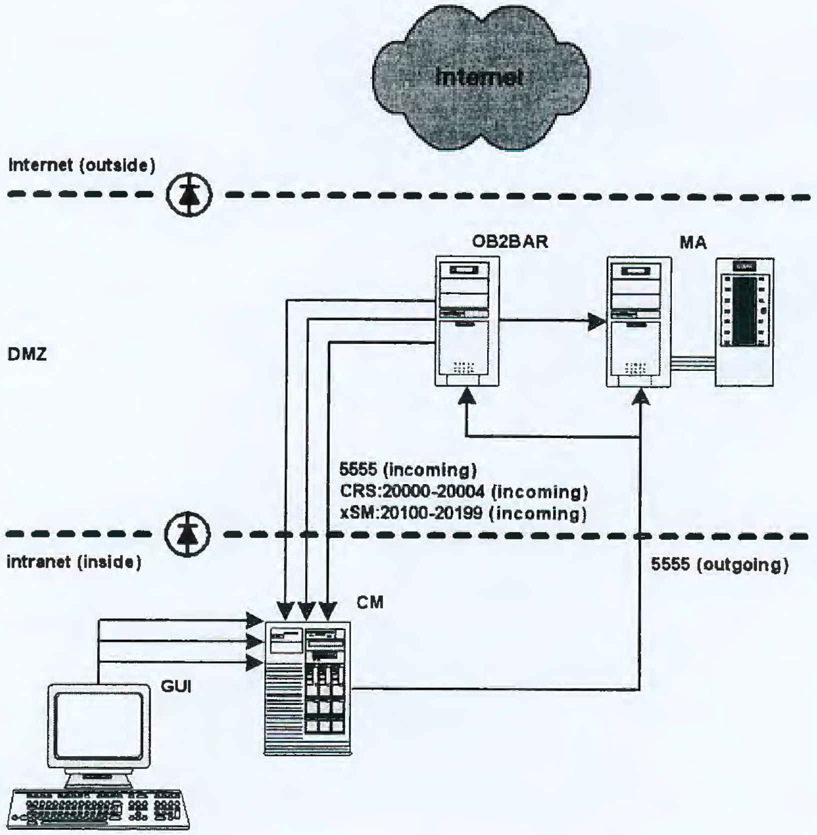

The following three items define the port range settings for this configuration:

1. Table 11-2 shows that Application Agents connects to the following processes on the Cell Manager:

Table 11-5

Process	Port
`Inet`	5555

17.841
Paula

Table 11-5

Process	**Port**
CRS	Dynamic
RSM	Dynamic
BSM	Dynamic
DBSM	Dynamic
xMA-NET	Dynamic

Here, the application Agent connects to the Media Agent. However, this connection does not go through the firewall and so you do not need to specify a port range.

This leads to the following firewall rule for the connection to the `Inet` listen port.

✓ Allow connections from the Application Agent system to port 5555 on the CM system

NOTE This rule allows connections from the DMZ to the intranet, which is a potential security risk.

2. Table 11-1 shows that CRS requires only one port. However, since other processes may allocate ports from this range as well, you should specify a range of about five ports on the CM system. The port range could be defined as follows:

   ```
   OB2PORTRANGESPEC=CRS:20000-20004
   ```

   The resulting firewall rule for the connection to the CRS process is:

   ✓ Allow connections from the Application Agent system to ports 20000-20004 on the CM system

3. For the Backup and Restore Session Manager, the situation is more complex. Every backup and restore session is started by one Session Manager, and every Session Manager requires one port. Additionally, an Application Agent may need to start some DBSMs. For Microsoft Exchange, Microsoft SQL, and Lotus Domino R5 Server integrations,

17.843
Paula

one DBSM will be started. For Oracle and SAP R/3 integrations, "concurrency + 1" DBSMs will be started. The port range for the Session Managers needs to be added to the `OB2PORTRANGESPEC` variable on the CM system:

Port Range Setting on the Cell Manager

```
OB2PORTRANGESPEC=CRS:20000-20004;xSM:20100-20199
```

Therefore, the firewall rule for the connections to the Session Managers is the following:

✓ Allow connections from the Application Agent system to ports 20100-20199 on the CM system

Limitations

- Remote installation of clients across the firewall is not supported. You need to install clients locally in the DMZ.
- This cell can back up clients in the DMZ, as well as clients in the intranet. However, each group of clients must be backed up to devices configured on clients that are on the same side of the firewall.

IMPORTANT If your firewall does not restrict connections from the intranet to the DMZ, it is possible to back up clients in the intranet to devices configured on clients in the DMZ. However, this is not recommended, as the data backed up in this way becomes more vulnerable.

- If a device in the DMZ has robotics configured on a separate client, this client must also be in the DMZ.
- This setup does not allow backup of databases or applications using Application Agents on the clients in the DMZ.

17.342
Paula

12 Troubleshooting

17.841
Paula

In This Chapter

If you have problems with Data Protector, use the suggestions in this chapter to get back on track, including information on:

"Before Calling Your Support Representative" on page 549

"Data Protector Log Files" on page 550

"Debugging" on page 553

"Browsing Troubleshooting Messages" on page 561

"When You Cannot Access Online Troubleshooting" on page 562

"Description of Common Problems" on page 564

"Troubleshooting Networking and Communication" on page 565

"Troubleshooting Data Protector Services and Daemons" on page 569

"Troubleshooting Devices and Media" on page 574

"Troubleshooting Backup and Restore Sessions" on page 580

"Troubleshooting Data Protector Installation" on page 588

"Troubleshooting User Interface Startup" on page 590

"Troubleshooting the IDB" on page 592

"Troubleshooting Data Protector Online Help" on page 603

"Check Whether Data Protector Functions Properly" on page 605

For an overview and hints on the performance aspects of the Data Protector, refer to Appendix , "Performance Considerations," on page 8.

Backup devices (such as tape drives) are subject to specific Data Protector licenses. See the *HP OpenView Storage Data Protector Software Release Notes* for details.

17.840
Paula

Before Calling Your Support Representative

In order to speed up the process of solving your problem, you should prepare before reporting a problem to HP Customer Support Service. See the suggestions below for preliminary steps you can take.

Ensure that:

- You are not running into known limitations that cannot currently be overcome. For specific information on Data Protector limitations and recommendations, as well as on known Data Protector and non-Data Protector problems, see the *HP OpenView Storage Data Protector Software Release Notes*.
- Your problems are not related to third-party hardware and software. If they are, contact the third-party vendor for support.
- You have the latest Data Protector patches installed. Patches can be obtained from the HP OpenView Web site: http://support.openview.hp.com/patches/patch_index.jsp. The list of OS patches is available in the *HP OpenView Storage Data Protector Software Release Notes*.

Collect the following data about the problem you encountered:

- A description of your problem, including the session output (or equivalent output, depending on the type of problem).
- Output from the `get_info` command located in the following directory:
  - On UNIX: `/opt/omni/sbin/utilns`
  - On Windows: `<Data_Protector_home>\bin\utilns`

  The script collects system data from your Data Protector Cell Manager, and configuration data about your Data Protector installation.
- All log files from the Cell Manager and from all clients involved.

Data Protector Log Files

If you encounter problems using the Data Protector application, you can use information in the log files to determine your problem.

Location of Data Protector Log Files

The Data Protector log files are located in the following directories:

- On Windows systems: `<Data_Protector_home>\log`
- On HP-UX and Solaris systems: `/var/opt/omni/log`
- On other UNIX systems: `/usr/omni/log`
- On Novell NetWare systems: `SYS:\USR\OMNI\LOG`

Format of Data Protector Log Files

Most Data Protector log file entries are of the following format:

```
<time_stamp> <process:PID:Thread_ID> <source_file and
branch> <Data Protector_version> <log_entry_message>
```

For example:

```
09/06/00 16:20:04 XOMNI.11561.0 ["/src/lib/ipc/ipc.c
/main/r31_split/10":3414] A.04.10 b325[ipc_receiveDataEx]
buffer 102400 bytes to small to receive data 796226418 bytes
=> ignored
```

17.338

Log Files and Their Contents

The table below describes the information found in Data Protector log files:

Table 12-1 **Data Protector Log Files**

Log File	Description
debug.log	Unexpected conditions are logged to this file. While some can be meaningful to you, it will be used mainly by the support organization.
Ob2EventLog.txt	Data Protector events that occurred during Data Protector operation and all Data Protector notifications are logged into this file. The Event Log represents a centralized Data Protector event depository.
inet.log	Requests made to the `Data Protector Inet` service are logged to this file. It can be useful to check the recent activity of Data Protector on clients.
IS_install.log	This file contains a trace of the remote installation and is located on the Installation Server.
media.log	Each time a medium is used for backup, initialized, or imported, a new entry is made to this log. The `media.log` can be used in IDB recovery to find the tape with the database backup and to find out which media were used after the last backup of the database.
omnisv.log	Contains information on when Data Protector services were stopped and started.

17837
Paula

Table 12-1 **Data Protector Log Files**

Log File	Description
purge.log	Contains traces of the background purge of the IDB.

RDS.log	Contains IDB logs. The file resides on the Cell Manager:
	On Windows:
	`<Data_Protector_home>\db40\datafiles\catalog`
	On UNIX:
	`/var/opt/omni/db40/datafiles/catalog`

sm.log	Contains errors that occurred during backup and restore sessions, such as errors in parsing the backup specifications.
Upgrade.log (UNIX only)	This log is created during the upgrade and contains traces of the upgrade process.
sap.log, oracle8.log, informix.log, sybase.log, db2.log	Application specific logs contain traces of the integration calls between the application and Data Protector. The files are located on application servers and can be used for troubleshooting integrations.

17836
Lauta

Debugging

You should collect debugs only when the support organization requires them to resolve a technical issue. When Data Protector runs in debug mode, it creates debug information that consumes a large amount of disk space. Consult the support organization about the detail level that should be applied and environmental conditions for running Data Protector in the debug mode.

Limiting the Maximum Size of Debugs

Circular Debugging

Data Protector can run in a special debugging mode called circular debugging. In this mode, debugging messages are added until the size of the debug file reaches a preset size (*n*). The counter is then reset and the oldest debugging messages are overwritten. This limits the trace file size, but does not affect the latest records.

When to Use Circular Debugging

Using this mode is recommended only if the problem occurs near the end of the session or if Data Protector aborts or finishes soon after the problem has occurred.

Estimating the Required Disk Space

With circular debugging turned on, an estimate of the maximum required disk space is as follows:

- On Media Agent client(s): `2*`*`n`*` [kB]` for each running MA in a backup or restore.
- On Disk Agent client(s): `2*`*`n`*` [kB]` for each mount point in a backup or restore.
- On the Cell Manager client: `2*`*`n`*` [kB]`.
- On a integration client: `2*`*`n`*` [kB] * `*`parallelism`*.
- For Inet and CRS debugging, the upper limit cannot be reliably determined, because separate debug traces are produced for various actions.

17.835
Paula

Ways of Debugging

You can start Data Protector in the debug mode in different ways and use it to generate debug traces. For more details about debugging options refer to the section "Debug Syntax" on page 555.

IMPORTANT When Data Protector runs in the debug mode, debug information is generated for every action. For example, if you start a backup specification in the debug mode, Disk Agents deliver output on each client backed up in this backup specification.

Debugging Using the Data Protector GUI

To set the options for debugging using the Data Protector GUI, in the `File` menu, click `Preferences`, and then click the `Debug` tab. Specify the debug options and restart the GUI. The GUI will be restarted in the debug mode.

Debugging Using the Trace Configuration File

Another way to set debugging options is to edit the trace configuration file (`/etc/opt/omni/options/trace` on UNIX and *`<Data_Protector_home>`*`\Config\Options\trace` on Windows).

Debugging Using the `OB2OPTS` Variable

Debugging parameters for Data Protector integrations can be set using the `OB2OPTS` environment variable. For more details about the `OB2OPTS` variable refer to the *HP OpenView Storage Data Protector Integration Guide*.

Debugging Scheduled Sessions

To debug scheduled sessions, edit the schedule file (`/etc/opt/omni/schedules` or `/etc/opt/omni/barschedules` on UNIX and *`<Data_Protector_home>`*`\Config\Schedules` or *`<Data_Protector_home>`*`\Config\BarSchedules` on Windows). Debugging parameters must be added in the first line of the file.

17.834
Paula

NOTE

Before you edit the file, make a copy of it, as the changes have to be reverted when debugging is no longer desired.

Example of a Modified Schedule

```
-debug 1-99 sch.txt
-full
-only 2002
-day 14 -month Dec
-at 22:00
```

Debug Syntax

Almost all Data Protector commands can be started with an additional `-debug` parameter that has the following syntax:

`-debug 1-99[,C:<n>] [,T:<s>] <XYZ> [<host>]`

where:

`1-99` is the debug range. The range should always be specified as 1-99 unless instructed otherwise.

`C:<n>` limits the size of debug files to *n* kilobytes. The minimum value is 4 (4kB) and the default value is 1024 (1MB).

`T:<s>` is the timestamp resolution, where the default value is 1, 1000 means the resolution is one millisecond and the value 0 means timestamps are turned off. The timestamp resolution and size limit for circular debugging are supplied as a part of the ranges parameter.

<XYZ> is the debug postfix, for example DBG_01.txt

<host> is the list of hostnames where debugging is turned on.

NOTE

On some platforms (Novell NetWare, MPE), millisecond resolution is not available.

The list of hostnames limits the systems where debugging is turned on during the execution of the Data Protector command. If there are multiple systems on the list, they should be delimited by spaces. The entire list must be within quotation marks, for example: `"host1.company.com host2.company.com"`.

17.833
Paula

Trace File Name

The debug postfix option is used for creating the trace files in the following directory:

- On UNIX systems: `/tmp`
- On Windows systems: *`<Data_Protector_home>`*`\tmp`
- On Novell NetWare systems: `SYS:\USR\OMNI\TMP`

The files are named

`OB2DBG_`*`<did>__<Program>_<Host>_<pid>_<XYZ>`*

where:

`<did>` (debugging ID) is the process ID of the first process that accepts the debugging parameters. This ID is used as an ID for the debugging session. All further processes will use this ID.

`<Program>` is the code name of the Data Protector program writing the trace.

`<Host>` is the name where the trace file is created.

`<pid>` is the process ID.

`<XYZ>` is the postfix as specified in the `-debug` parameter.

Once the backup or restore session ID(*`<sid>`*) is determined, it will be added to the filename:

`OB2DBG_`*`<did>_<sid>_<Program>_<Host>_<pid>_<XYZ>`*

Processes that add the *`<sid>`* are BMA/RMA, xBDA/xRDA, and other processes started by the session, but not by the BSM/RSM itself.

NOTE The session ID is intended to help you identify sets of debug files. Other debug files may belong to the same session and you may have to provide them as well.

trace.log A `trace.log` file is generated on the Cell Manager, containing information where (on which hosts) debug files are generated and which debugging prefixes are being used. Note that this file does not contain a complete list of all generated files.

17.832
Paula

OB2DBGDIR

The default location of trace files can be changed on a per system basis with the omnirc variable OB2DBGDIR. For more details about omnirc variables, refer to “Using Omnirc Options” on page 525.

INET Debug on UNIX

To debug Inet on UNIX systems, edit the /etc/inetd.conf file and change the following line:

```
omni stream tcp nowait root /opt/omni/lbin/inet inet -log
/var/opt/omni/log/inet.log
```

to

```
omni stream tcp nowait root /opt/omni/lbin/inet inet -log
/var/opt/omni/log/inet.log -debug 1-99 SSF
```

NOTE If you enable Inet debugs, all integrations will generate trace log files.

INET Debug on Windows

To debug Data Protector Inet on Windows systems, restart the Data Protector Inet service with the following startup parameters: -debug 1-99 POSTFIX.

NOTE If you enable Inet debugs, all integrations will generate trace log files.

CRS Debug on Windows

In order to debug CRS on Windows, in the Control Panel, go to Services (Windows NT) or Administrative Tools, Services (other Windows systems), then stop the CRS service and restart it with the following startup parameters:

```
-debug 1-99 POSTFIX <Cell_Manager_name>
```

17831
Paula

NOTE

Use the -debug option carefully because execution traces can become quite large.

CRS Debug in the Microsoft Cluster Environment

In the Data Protector shared directory, edit the *<Data_Protector_home>*\Config\options\Trace file. Add the following lines:

```
ranges=1-99,110-500
postfix=DBG
select=obpkg.rc.aus.hp.com
```

From the Cluster Administrator GUI, take the CRS service resource (OBVS_MCRS) offline.

CAUTION

Do not stop the CRS from Control Panel as it will cause the Data Protector package to failover.

Sample Debugging

Follow the procedure described below to collect debug information for problems occurring during backup sessions involving one client and the Cell Manager:

1. Reduce the error environment as much as possible:
   - Create a backup specification that contains just one or a few files or directories.
   - Include only one failing client in the debug run.
2. Create an `info` text file that contains the following:
   - Hardware identification of the Cell Manager, Media Agent, and Disk Agent clients. For example, `HP-9000 T-600 Series; Vectra XA`.
   - The SCSI controller's name, for example,

17.820
Paula

`onboard_type/Adaptec xxx/...` for Windows Media Agent clients.

- The operating system version, for example, `HP-UX 11.00, Windows NT Server 4.0 SP2...`
- Topology information obtained from the `omnicellinfo -cell` command output.
- The output of the `devbra -DEV` command if you have issues with backup devices.

3. Discuss the technical issue with the support organization and request the following information:
    - Debug level (For example, "1-99." This is a command option needed later.).
    - Debug scope (client only, Cell Manager only, all)
4. Delete any files from previous debugging sessions in the following directories:
    - On Windows systems: `<Data_Protector_home>\tmp`
    - On UNIX systems: `/tmp`
    - On Novell NetWare systems: `SYS:\USR\OMNI\TMP`
5. Exit all user interfaces and stop all other backup activities in the cell.
6. In case you need to collect the CRS debugs as well, you need to:
    - Stop the Data Protector services on the Cell Manager.
    - Restart the services in the debug mode.
7. On the Cell Manager run the following command to start the GUI in debug mode:
    - On Windows systems: `manager -debug 1-99 error_run.txt`
    - On UNIX systems: `xomni -debug 1-99 error_run.txt`

    You can define the postfix of the trace file names created by substituting the `error_run` text with your preference.
8. Reproduce the problem using Data Protector.
9. Exit all user interfaces to quit the debug mode.

    If you have to collect CRS debugs as well, you have to stop the Data

17.829
Paula

Protector services on the Cell Manager and restart them without the debug option.

10. Copy the files from any clients involved with the problem.

11. Compress and pack the contents of the tmp directories (_error_run.txt files) on the Cell Manager and clients and the info file. Use WINZIP or TAR.

12. Email the files to the support organization and include in the email information about how you packed and compressed the files.

13. Delete the _error_run.txt files from the tmp directories.

17828
falta

Browsing Troubleshooting Messages

Data Protector provides an interactive online troubleshooting utility, where you can get a detailed explanations of your error messages, including suggestions for correcting problems.

When you receive an error message from Data Protector, the error number is presented as a clickable link. To see detailed information about the error, click the link. The error message dialog appears providing extensive information about the error. Click `Details` to see a detailed description of the error message and the actions you can take to avoid or solve the problem.

Error message dialog consists of the following:

- `Error Message`: Exact message as it appears.
- `Description`: Detailed description of the error message.
- `Action`: Possible actions to take to solve or avoid the problem.

Figure 12-1 **Sample Error Message Dialog**

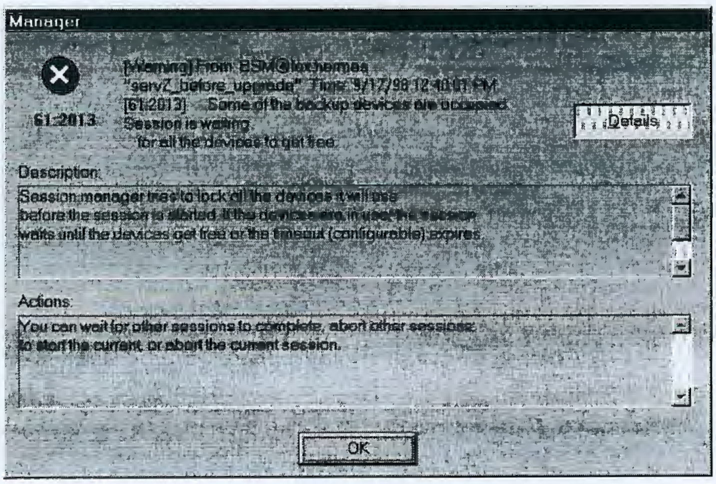

Manager

61:2013

[Warning] From: BSM@lox.hermes
"servZ_before_upgrade" Time: 9/17/98 12:40:07 PM
[61:2013] Some of the backup devices are occupied.
Session is waiting
for all the devices to get free.

Description:

Session manager tries to lock all the devices it will use
before the session is started. If the devices are in use, the session
waits until the devices get free or the timeout (configurable) expires.

Actions:

You can wait for other sessions to complete, abort other sessions
to start the current, or abort the current session.

17.827
Paula

When You Cannot Access Online Troubleshooting

If the user interface cannot be started, you can access the troubleshooting file. This is a text file containing all Data Protector error messages each of which includes the following information:

- MESSAGE: The error message as it appears in Data Protector.
- DESCRIPTION: A detailed or extended information about the error.
- ACTION: Actions you can take to solve or avoid the problem.

The troubleshooting file is only available in the directory where the Cell Manager is installed. It can be found in the following locations:

- On UNIX: `/opt/omni/gui/help/C/Trouble.txt`
- On Windows: `<Data_Protector_home>\help\enu\Trouble.txt`

An example of an error message is shown below:

```
MESSAGE:
[12:5]  Internal error in (\"p\":num) => process aborted
This is an unexpected condition and is likely due to a
combination of circumstances involving both this product
and the operating system.

Report this error to your post-sales Data Protector
Support Representative.

DESCRIPTION:

An internal error occurred. The process was not able to
recover and aborted ungracefully immediately after
reporting this condition.

ACTION:

Before contacting your post-sales Data Protector Support
Representative, please gather as much information as
possible:

* Write down product version and build number.
* Make a note of the circumstances that cause this error.
* Save session output to a file (e.g. session.txt).
```

17.826
Paula

```
* Collect all log files (*.log) in
<Data_Protector_home>/log directories
on all hosts involved in the situation when this error
occurred
(i.e. host running VBDA, host running BMA and host
running BSM).
```

17.825
Paula

Description of Common Problems

If you have problems with Data Protector, find the problem area listed below that most closely matches the problem you are having:

- Networking and Communication, on page 565
- Service Startup, on page 569
- Device Usage, on page 574
- Starting Backup and Restore Sessions, on page 580
- User Interface Startup, on page 590

Certain functionality of Data Protector is subject to particular license requirements. Refer to the *HP OpenView Storage Data Protector Installation and Licensing Guide* for more information on licensing.

17.324
Paula

Troubleshooting Networking and Communication

The section addresses the following networking and communication problems:

- "Hostname Resolution Problems" on page 565
- "Client Fails with "Connection Reset by Peer"" on page 567

Hostname Resolution Problems

Hostname resolution is a very common problem in a Data Protector environment. It means that host A is unable to communicate with host B.

The table below shows Data Protector components and how they should communicate within the Data Protector environment. Communication among hosts means that host A in the table should resolve host B by its fully qualified domain name (FQDN). Resolving a host means that host A can interpret the FQDN and determine its IP address.

Table 12-2 **Data Protector Components Name Resolution**

Host A	Host B
Disk Agent Client Host	Media Agent Client Host
Disk Agent Client Host	Cell Manager Host
Disk Agent Client Host	MoM Server Host
Media Agent Client Host	Disk Agent Client Host
Media Agent Client Host	Cell Manager Host
Media Agent Client Host	MoM Server Host
Cell Manager Host	Media Agent Client Host
Cell Manager Host	Disk Agent Client Host
Cell Manager Host	MoM Server Host
MoM Server Host	Disk Agent Client Host

17823
Paula

Table 12-2 **Data Protector Components Name Resolution**

Host A	Host B
MoM Server Host	Media Agent Client Host
MoM Server Host	Cell Manager Host

DNS Resolution Problem

Test DNS resolution among hosts using the `omnicheck` command. Refer to the "Verifying DNS Connections within Data Protector Cell" section in the *HP OpenView Storage Data Protector Installation and Licensing Guide* and to `omnicheck` man page for more information on how to use the command.

Enter the following command:

```
omnicheck -dns
```

This will check all DNS connections needed for normal Data Protector operating.

Problem

If the response to the `omnicheck` command is:

```
<client_1> connects to <client_2>, but connected system
presents itself as <client_3>
```

The message may occur when the `hosts` file on *client_1* is not correctly configured or the hostname of the *client_2* does not match its DNS name.

If the response to the `omnicheck` command is:

```
<client_1> failed to connect to <client_2>
```

The message may occur when the `hosts` file on *client_1* is not correctly configured or *client_2* is unreachable (for example, disconnected).

Action

Consult your network administrator. Depending on how your environment is configured to perform name resolution, you may need to resolve this problem either in your DNS configuration or by editing the `hosts` file located in the following directories:

- On Windows: `<%SystemRoot%>\System32\drivers\etc`
- On UNIX: `/etc`

17822
Paula

Problem

The response to the `omnicheck` command is:

`<client_1> cannot connect to <client_2>`

This means that the packet has been sent, but not received because of the timeout.

Action

Check for and resolve any network problems on the remote host.

Checking the TCP/IP setup

An important aspect of the TCP/IP configuration process is the setup of a hostname resolution mechanism. Each system in the network must be able to resolve the address of the Cell Manager and all machines with Media Agents and physical media devices. The Cell Manager must be able to resolve the names of all systems in the cell.

Action

Once you have the TCP/IP protocol installed, you can use the `ping` and `ipconfig` utilities to verify the TCP/IP configuration. For detailed steps, refer to the online Help index keyword “checking, TCP/IP setup”.

HOSTS file resolution problem

Action

If you encounter resolution problems when using the `Hosts` file, do the following:

- On Windows: edit the `LMHosts` file in the `<%SystemRoot%>\System32\drivers\etc` directory.
- On UNIX: edit the `/etc/hosts` file.

Client Fails with “Connection Reset by Peer”

On Windows, default configuration parameters of the TCP/IP protocol may cause connections to break. This can be due to a high network or computer usage, unreliable network, and connections between different operating systems.

The connection breaks and the system displays the error: `[10054] Connection reset by peer.`

Action

You can configure the TCP/IP protocol to use 8 instead of the default 5 retransmissions. It is better not to use higher values because each increment doubles the timeout. The setting applies to all network connections, not only to connections used by Data Protector.

17821
Paula

On Windows, apply the change to the Cell Manager first.

If you run the UNIX Cell Manager and the problem persists, apply the change to any problematic Windows clients.

1. Add a new `DWORD` parameter `TcpMaxDataRetransmissions` and set its value to `0x00000008(8)` under the following registry key:

   ```
   HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\Tcpi
   p\Parameters
   ```

   On Windows 98 use:

   ```
   HKEY_LOCAL_MACHINE\System\CurrentControlSet\Services\VxD\
   MSTCP
   ```

   ```
   MaxDataRetries:(DWORD):8
   ```

   Making a mistake in editing the registry can cause your system to become unstable and unusable.

2. Restart the system after making this change.

17820
Paula

Troubleshooting Data Protector Services and Daemons

The Data Protector services and daemons run on the Cell Manager. Run the `omnisv -status` command to check whether services are running.

If the Data Protector services seem to be stopped or have not been installed on the target Data Protector client, ensure first that you don't have a name resolution problem. Refer to "Troubleshooting Networking and Communication" on page 565 for more information.

You can run into the following problems with Data Protector services and daemons:

- "Problems Starting Data Protector Services on Windows" on page 569
- "Problems Starting Data Protector Daemons on UNIX" on page 571

Problems Starting Data Protector Services on Windows

You do not have permission to start the services

The following error message displays:

```
Could not start the <Service_Name> on <System_Name>.

Access is denied.
```

Action

The system administrator should grant you permission to start, stop, and modify services on the system that you administer. You have to log off and log on again on Windows NT system for the changes to take effect or log off and log on as system administrator and then start/stop or modify the services. On other Windows systems you do not have to log off the system in order to start/stop or modify the services. Instead, you can run the `services.msc` (located in the *<%SystemRoot%>*`\system32` directory) as an administrator by right-clicking the file while holding down the **Shift** button and selecting `Run as` from the pop-up menu. Then provide administrator's user name and password.

17819
[illegible]

Changed service account properties

If the service account does not have permission to start the service or if the service account properties (the password, for example) have been changed, you get the following error message:

```
The Data Protector Inet service failed to start due to the
following error:
```

```
The service did not start due to a logon failure.
```

Action

1. Modify the service parameters: in the Windows Control Panel, go to `Services` (Windows NT) or `Administrative Tools, Services` (other Windows systems).
2. If this does not solve the problem, contact your system administrator to set up the account with appropriate permissions. The account should be a member of the `Admin` group and should have the `Log on as a service` user right set.

A specific service has not been found

The location of the service is registered in the `ImagePath` key. If the executable does not exist in the location specified under this key, the following error message is displayed:

```
Could not start the <Service_Name> on <System_Name>. The
system can not find the file specified!
```

Action

1. On the Cell Manager, copy the `<Data_Protector_home>\db40` and `<Data_Protector_home>\config` directories to a safe location before uninstalling Data Protector.
2. Copy the `<Data_Protector_home>\db40` and `<Data_Protector_home>\config` directories back in place.
3. Uninstall the current Data Protector installation either on the client or on the Cell Manager, and then reinstall the software.

This guarantees a clean installation with all the binaries in place.

MMD fails upon starting the CRS service

If the Data Protector CRS service fails to start and `mmd.exe` invokes a *Dr.Watson* diagnosis, this points to a corruption in the database log files.

Action

1. Delete the `mmd.ctx` file on the `<Data_Protector_home>\tmp`

17.813
Paula

directory and the problems should be resolved.

2. Restart the services using the `omnisv -start` command.

RDS does not work on the Windows TSE Cell Manager

Use TCP transport instead of local transport by modifying the `<Data_Protector_home>\db40\datafiles\catalog\velocis.ini` file:

Under `TCP Configuration`, set `Enabled` to `yes`.

Problems Starting Data Protector Daemons on UNIX

The following daemons run on the UNIX Cell Manager:

- Data Protector CRS daemon: `/opt/omni/lbin/crs`
- IDB daemon: `/opt/omni/lbin/rds`
- Data Protector Media Management daemon: `/opt/omni/lbin/mmd`

The `Data Protector Inet` service (`/opt/omni/lbin/inet`) is started by the system inet daemon when an application tries to connect to the Data Protector port, which is by default port number 5555.

Normally, these daemons are started automatically during the system's start-up.

To manually stop, start, and get the status of Data Protector daemons, log on to the Cell Manager as root.

Stopping Daemons

To stop the Data Protector daemons, enter the following command in the `/opt/omni/sbin` directory:

```
omnisv -stop
```

Starting Daemons

To start the Data Protector daemons, enter the following command in the `/opt/omni/sbin` directory:

```
omnisv -start
```

Checking the Status of the Daemons

To check the running status of the Data Protector daemons, enter the following command in the `/opt/omni/sbin` directory:

```
omnisv -status
```

There are several possible reasons why the Data Protector daemon has failed to start:

Raima Velocis server daemon could not be started

```
/opt/omni/sbin/omnisv -start
```

```
Could not start Raima Velocis server daemon.
```

Action

See `/var/opt/omni/db40/datafiles/catalog/RDS.log` for details.

Check that you have all IDB files in the `/var/opt/omni/db40` directory. Compare the list of files in the `/opt/omni/newconfig/var/opt/omni/db40` to the list of files in the `/var/opt/omni/db40` directory. Ensure that these directories are mounted.

Raima Velocis server daemon is apparently not running

If any of the Data Protector commands terminate with following message:

```
[12:1166] Velocis daemon error - the daemon is probably not
running
```

Action

Check if the database server is really not running using following command: `/opt/omni/sbin/omnisv -status`

- If the database server is not running, start it by running: `/opt/omni/sbin/omnisv -start`
- If the database server is running, then it is likely either that the `/var/opt/omni/db40` directory does not exist or some of the files are missing. This can happen if someone has accidently removed the directory or some of the IDB files. Recover the IDB. Refer to "Recovering the IDB" on page 417.

Data Protector Cell Manager daemon could not be started

```
/opt/omni/sbin/omnisv -start
```

```
Could not start the Cell Manager daemon.
```

Action

See `/var/opt/omni/tmp/omni_start.log` for details.

Ensure that the following configuration files exist:

- /etc/opt/omni/options/global
- /etc/opt/omni/options/users/UserList
- /etc/opt/omni/options/ClassSpec

Data Protector Processes

Table 12-3 shows which processes run and where they run while Data Protector is idle, or doing a backup, a restore or a media management session.

Table 12-3 Which Processes Run Where, and When

	Idle	Backup	Restore	Media Management
Windows Cell Manager	rds.exe, crs.exe, omniinet.exe, bsm.exe	rds.exe, mmd.exe, omniinet.exe, mmd.exe	rds.exe, omniinet.exe, mmd.exe, crs.exe, rsm.exe	rds.exe, omniinet.exe, mmd.exe, crs.exe, msm.exe
UNIX Cell Manager	rds, mmd,crs	rds, mmd, crs, bsm	rds, mmd, crs, rsm	rds, mmd, crs, msm
Windows Disk Agent Client	omniinet.exe	omniinet.exe, vbda.exe	omniinet.exe, vrda.exe	omniinet.exe
UNIX Disk Agent Client		vbda	vrda	
Windows Media Agent Client	omniinet.exe	omniinet.exe, bma.exe	omniinet.exe, rma.exe	omniinet.exe, mma.exe
UNIX Media Agent Client		bma	rma	mma

17.815
Paula

Troubleshooting Devices and Media

This section describes solutions to the following problems that can arise while using backup devices:

- "Cannot Access Exchanger Control Device on Windows 2000/XP/Server 2003" on page 574
- "Device Open Problem" on page 575
- "Using Unsupported SCSI Adapters on Windows" on page 575
- "Medium Quality Statistics" on page 575
- "Medium Header Sanity Check" on page 577
- "Cannot Use Devices After Upgrading to Data Protector A.05.10" on page 578
- "Other Common Problems" on page 579

Problems involving device SCSI addresses are explained in detail in Appendix B of the *HP OpenView Storage Data Protector Installation and Licensing Guide*.

Cannot Access Exchanger Control Device on Windows 2000/XP/Server 2003

Data Protector uses the SCSI mini-port driver to control backup drives and libraries. Data Protector may fail to manage devices if other device drivers are loaded on the same system. The error message `Cannot access exchanger control device` appears when device operations such as media formatting or scanning are started.

Action

Run the `<Data_Protector_home>\bin\devbra -dev` command on the system where the devices are located, to list all physical devices that are configured on the system. If any of the SCSI addresses have the `CLAIMED` status value, they are used by another device driver.

Disable the Windows 2000/XP/Server 2003 robotic driver. For detailed steps, refer to the online Help index keyword "robotics drivers".

17.814
Paula

Device Open Problem

The error message `Cannot open device (not owner)` appears when trying to use a DDS device.

Action

Check whether you are using a medium that is incompatible with the Media Recognition System. Media used with DDS drives must comply with the Media Recognition System.

Using Unsupported SCSI Adapters on Windows

System fails due to usage of unsupported SCSI adapters with backup devices.

Typically, the problem occurs when the SCSI device was accessed by more than one Media Agent at the same time or when the length of the transferred data defined by the device's block size was larger than the length supported by the SCSI adapter.

Action

You can change the `Block size` in the `Advanced Backup Options` for the backup specification.

For information on supported SCSI adapters, refer to the *HP OpenView Storage Data Protector Software Release Notes*.

For detailed steps, refer to the online Help index keyword “setting advanced options for devices and media”.

Medium Quality Statistics

This functionality is used to detect any problems with media while they're still in their early stages. Before each medium is ejected from a drive, Data Protector queries the `SCSI log sense` command for medium read and write statistic information. The information is written to the `media.log` file.

The medium quality statistics feature is disabled by default. To enable it, set the following global variable: `Ob2TapeStatistics=1` in the Global Options file.

Global Options file is located:

- on UNIX: `/etc/opt/omni/options` in the
- on Windows: `<`*`Data_Protector_home`*`>\config\option`

If you receive media related errors during write operations, or if the medium is marked as poor, you can check the `media.log` file for media errors statistics. You can do this also when receiving media related errors during read operations.

`Media.log` file contains the following error statistics:

Error statistics	**Explanation**
`errsubdel=`*n*	errors corrected with substantial delays
`errposdel=`*n*	errors corrected with possible delays
`total=`*n*	total number of re-writes
`toterrcorr=`*n*	total number of errors corrected and recovered while writing
`totcorralgproc=`*n*	total number of times correction algorithm processed
`totb=`*n*	total bytes processed (write)
`totuncorrerr=`*n*	total number of uncorrected errors (write)

where *n* stands for number of errors.

If a parameter has the value `-1`, it means that the device does not support this statistic parameter. If all parameters have the value `-1`, it can either mean that during processing the tape quality statistics an error occurred or the device does not support medium quality statistics at all.

Although the tape statistical results are reported in bytes for `total bytes processed`, this is not true for all devices. LTO and DDS devices report data sets and groups, respectively, and not bytes.

Examples

Here are a few examples from the `media.log` file:

- Log sense write report for DLT/SDLT devices - total bytes processed.

  ```
  Media ID from tape= 0fa003bd:3e00dbb4:2310:0001; Medium Label=
  DLT10; Logical drive= dlt1; Errors corrected no delay= 0; Errors
  corrected delay= 0; Total= 13639; Total errors corrected= 13639;
  Total correction algorithm processed= 0; Total bytes processed=
  46774780560; Total uncorrected errors= 0
  ```

  46774780560 bytes of native data after compression were processed (a full DLT8000 tape).

17312
Paula

- Log sense write report for LTO devices - total data sets processed.

  ```
  Media ID from tape=0fa003bd:3e0057e6:05b7:0001; Medium Label=
  ULT2; Logical drive=ultrium1; Errors corrected no delay= 0;
  Errors corrected delay= 0; Total= 0;Total errors corrected= 0;
  Total correction algorithm processed= 0; Total bytes processed=
  47246; Total uncorrected errors= 0
  ```

  One data set is 404352 bytes. To calculate the amount of total bytes processed, use the following formula:

  ```
  47246 data sets * 404352 bytes = 19104014592 bytes after
  compression a full tape).
  ```

- Log sense write report for DDS devices - total groups processed.

  ```
  Media ID from tape= 0fa0049f:3df881e9:41f3:0001; Medium Label=
  Default DDS_5; Logical drive= DDS; Errors corrected no delay=
  -1; Errors corrected delay= -1; Total= -1; Total errors
  corrected= 0; Total correction algorithm processed= 154; Total
  bytes processed= 2244; Total uncorrected errors= 0
  ```

  DDS1/2: One group is 126632 bytes
  DDS3/4: One group is 384296 bytes

  To calculate the amount of total bytes processed, use the following formula:

  ```
  2244 groups * 126632 bytes = 284162208 bytes after compression
  (a 359 MB backup on DDS2).
  ```

  359 MB of data was backed up, resulting in 271 MB of native data on tape.

Medium Header Sanity Check

Data Protector performs a medium header sanity check before a medium is ejected from a drive to validate the medium header.

The medium header sanity check is enabled by default. The global variable can be set by uncommenting the following line in the Global Options file: `Ob2HeaderCheck=1`.

Problem

In case the medium header sanity check detects any header consistency errors on the medium, an error message is displayed and all the objects on the medium are marked as failed.

If the medium header is corrupt, all objects on the affected medium are marked as failed and the medium state is marked as poor.

17.511
Paula

Action

Export the medium from the IDB and restart the failed session using a different medium.

Cannot Use Devices After Upgrading to Data Protector A.05.10

Problem

After upgrading to Data Protector A.05.10, you cannot use devices that were configured as different device types in previous releases. For example, you cannot use 9940 devices that were configured as 9840 devices, or SuperDLT devices that were configured as DLT devices. The following error occurs:

```
[Critical] From: BMA@ukulele.company.com "SDLT"  Time:
2/22/2003 5:12:34 PM
[90:43]  /dev/rmt/1m
Invalid physical device type => aborting
```

Action

Manually reconfigure these devices using the `mchange` command, located on the Cell Manager in the following directories:

- On HP-UX: `/opt/omni/sbin/utilns/HPUX`
- On Solaris: `/opt/omni/sbin/utilns/SOL`
- On Windows: `<Data_Protector_home>\bin\utilns\NT`

Command Syntax

```
mchange -pool PoolName -newtype NewMediaClass
```

where:

PoolName is the name of the media pool with devices that are currently configured and should be reconfigured (for example, Default DLT or Default T9840).

NewMediaClass is the new media type of the devices, for example, T9940 for 9940 devices and SuperDLT for SuperDLT device.

Example

```
mchange -pool "Default DLT" -newtype "SuperDLT"
```

The command changes media types for all media, drives and libraries that use the defined media pool. After you have executed this command for each device you wanted to change, move the media associated with the reconfigured devices from the current media pool to the media pool corresponding to these media.

17810
Paula

For example, move the media associated with the reconfigured 9940 devices to the Default T9940 media pool, and the media associated with the reconfigured SuperDLT devices to the Default SuperDLT media pool. For related procedures, refer to the online Help.

Other Common Problems

Other common problems are hardware-related.

Action

Check the SCSI communication between the system and the device, such as adapters or SCSI cables and their length. Try running an OS-provided command, such as `tar`, to verify that the system and the device are communicating.

17309
Paula

Troubleshooting Backup and Restore Sessions

You may run into the following problems while running or starting backup and restore sessions:

- "Filenames Are Not Displayed Correctly in GUI" on page 580
- "Full Backups Are Performed Instead of Incrementals" on page 580
- "Unexpected Mount Request for a Standalone Device" on page 581
- "Unexpected Mount Request for a Library Device" on page 582
- "Unexpected Mounted Filesystems Detected" on page 583
- "Data Protector Fails to Start a Scheduled Session" on page 584
- "Data Protector Fails to Start an Interactive Session" on page 585
- "Backup Protection Expiration" on page 586
- "Troubleshooting Application Database Restores" on page 586
- "Problems with non-ASCII Characters in Filenames" on page 587

Filenames Are Not Displayed Correctly in GUI

When using the Data Protector GUI on Windows, some filenames belonging to the non-Windows objects can be displayed incorrectly. This happens when different encoding is used.

Action

To view these objects correctly, specify the appropriate encoding in the Data Protector GUI by selecting `Encoding` from the `View` menu, then selecting the appropriate codeset.

Full Backups Are Performed Instead of Incrementals

There are several reasons, outlined below, that Data Protector might run a full backup despite the fact that you specified an incremental backup.

17.808
Laura

No previous full backup

Before performing an incremental backup of an object, Data Protector requires a full backup. Data Protector uses a full backup as a base for comparison to determine which files have changed and consequently need to be included in the incremental backup. If a protected full backup is not available for this comparison, a full backup is performed.

Action Set the protection for the full backup.

The description has changed

An object is defined by the client, disk, and description. If any one of these three change, Data Protector considers it as a new object, even if the client and disk are the same, and Data Protector performs a full backup instead of an incremental.

Action Use the same description for full and incremental backups.

The backup owner is different

If your backups are configured to run as private, the person starting the backup is the owner of the data. For example, if USER_1 performs a full backup and USER_2 tries to start an incremental backup, the incremental backup will be executed as a full backup. This is because the data for USER_1 is private and cannot be used as a base for the USER_2's incremental backup.

Action Configure backup session `Ownership` in the `Advanced Backup Options` for the backup specification. The backup owner should be a user from the `Admin` user group. This will make all backups owned by this user, regardless of who actually starts the backup session.

Unexpected Mount Request for a Standalone Device

There are several situations, described below, that may cause Data Protector to issue a mount request for a standalone device while media are available in the backup device.

The media in the device are in a media pool that has the Non Appendable policy

Even though there is still available space on the media, the media will not be used because of the `Non Appendable` policy of the pool.

17.807
Paula

Action

Modify the media pool policy to `Appendable` to enable the appending of backups to the media until the media are full.

The media in the device are not formatted and the media pool to be used has a Strict policy

If your pool uses a `Strict` media allocation policy, media that are not formatted will not be used for backup. If no formatted media are available, Data Protector issues a mount request.

Action

If you would like Data Protector to automatically format unformatted media, set the media pool policy to `Loose` and change global variable `InitOnLoosePolicy` to 1.

The media in the device are not formatted and the media pool to be used has a Loose policy

If your pool uses a `Loose` media allocation policy, media are not automatically formatted.

Action

If you would like Data Protector to automatically format unformatted media, you need to change global variable `InitOnLoosePolicy` to 1.

The media in the device are formatted but are different from those in the preallocation list

The media in the device are formatted but are different from those in the preallocation list of the backup specification, and the pool specified has a `Strict` policy

If you use a preallocation list of media in combination with the `Strict` media policy, the exact media specified in the preallocation list need to be available in the device when a backup is started. If the exact media are not available, a mount request is issued.

Action

To use media available in the device in combination with the preallocation list, modify the media pool allocation policy to `Loose`.

Unexpected Mount Request for a Library Device

There are several situations, described below, that may cause Data Protector to issue a mount request for a library device while media are available in the library.

17.806
Paula

The media in the library are not formatted and the media pool with the media used for backup has a Strict policy

If your pool uses a `Strict` media allocation policy, unformatted media are not used for backup. If no formatted media are available in the library, Data Protector issues a mount request.

Action

If you would like Data Protector to automatically format unformatted media that are available in the library, set the media pool policy to `Loose`. This can be modified in the media pool `Properties`.

The media in the library are formatted but are different from those in the preallocation list

The media in the library are formatted but are different from those in the preallocation list of the backup specification, and the media pool specified has a `Strict` policy.

If you are using a preallocation list of media in combination with the `Strict` policy and the exact media specified in the preallocation list are not available in the device when backup is started, a mount request is issued.

Action

The exact media specified in the preallocation list need to be available in the device when the backup is started.

To use other media, if available in the device, in combination with the preallocation list, modify the media pool allocation policy to `Loose`.

To use any available media in the device without the preallocation list, remove the preallocation list from the backup specification. Do this by changing backup device options for the backup specification.

Unexpected Mounted Filesystems Detected

When restoring a disk image, you may get a message that the disk image being restored is a mounted file system and will not be restored:

```
Object is a mounted filesystem => not restored.
```

This happens when an application on the disk image leaves some patterns on the disk image. The patterns confuse the system call that verifies whether the eventually mounted filesystem on the disk image is mounted or not, so the system call reports that there is a mounted filesystem on the disk image.

17805
Paula

Action

1. Before you start a restore erase the disk image on the Data Protector client with the disk image being restored by entering the following commands:

```
    prealloc null_file 65536
dd if=null_file of=<device_file>
```

   where *<device_file>* is a device file for the disk image being restored.

2. Start the restore.

Data Protector Fails to Start a Scheduled Session

The scheduled sessions no longer run

The scheduled sessions no longer run since the Data Protector system account, which is supposed to start scheduled sessions, is not in the `Admin` user group on the Cell Manager.

This account is added to the Data Protector `Admin` group on the Cell Manager at installation time. If this is modified and the permission for this account is removed, or if the service account changes, the scheduled sessions no longer run.

Action

Add the Data Protector account to the `Admin` user group on the Cell Manager.

The session fails and Data Protector issues the session status No licenses available.

A backup session is started only after Data Protector has checked the available licenses. Otherwise, the session fails and Data Protector issues the session status `No licenses available`.

Action

Obtain information on available licenses by clicking the `Help` menu and then `About` in the Data Protector Manager.

Request new licenses and apply them to the Data Protector system. See the *HP OpenView Storage Data Protector Installation and Licensing Guide* for licensing details.

17804
Paula

Data Protector Backup sessions are not started at all (UNIX-specific)

Action

Run the `crontab -l` command to check whether the `omnitrig` program is included in the `crontab` file. If the following line does not appear, the `omnitrig` entry was automatically added by Data Protector:

```
0,15,30,45 * * * * /opt/omni/sbin/omnitrig
```

Stop and start the Data Protector daemons by running the `omnisv -stop` and the `omnisv -start` commands in the `/opt/omni/sbin` directory.

Data Protector Fails to Start an Interactive Session

Every time a backup is started, permission to start a backup session is required and checked for the user who is currently running Data Protector. If the user does not have sufficient permission, the session cannot be started.

Action

Check and change the user rights for the particular user. Refer to Chapter 3, “Configuring Users and User Groups,” on page 81.

Poor Backup Performance on Novell NetWare Server

Backup performance on a Novell NetWare Server may be poor. Backup does not run continuously, but intermittently. This is a well-known problem caused by the system `TCPIP.NLM`.

Action

Set the following parameters:

- NW5.1/NW6.0: `SET TCP DELAYED ACKNOWLEDGEMENT = OFF`
- NW5.0: `SET TCP DELAYED ACK = OFF`

This increases backup performance without any secondary effects.

Data Protector Fails to Start Parallel Restore Media Agent on Novell NetWare Clients

Data Protector UNIX session manager sometimes fails to start restore media agents in parallel on Novell NetWare clients with an error message like, for example, `Could not connect to inet` or `Connection`

17803
Lauix

`reset by peer`. It is possible that some parallel restore sessions are completed without errors, while other restore sessions are not even started.

Action A workaround for this problem is to set the `SmMaxAgentStartupRetries` global variable in the Data Protector global options file (located in `/etc/opt/omni/options/global`) to `2` or more (max. 50). This variable specifies the maximum number of retries for the session manager to restart the failed agent before it fails. Refer to "Global Options File" on page 523 for more information about the Data Protector global options file.

Backup Protection Expiration

When scheduling backups, you have set the same protection period for full and incremental backups, which means that incremental backups are protected for the same duration as the relevant full backup. The consequence of this is that your data will actually only be protected until the full backup expires. You cannot restore incremental backups that have been based on expired full backups.

Action Configure the protection for your full backups so that they are protected for longer than your incremental backups.

The time difference between the protection for the full backup and the incremental backup should be the amount of time between the full backup and the last incremental backup before the next full backup. For example, if you run incremental backups Monday through Friday and full backups on Saturday, you should set the protection of the full backup to at least 6 days more than for the incremental backups. This will keep your full backup protected and available until your last incremental backup expires.

Troubleshooting Application Database Restores

A poorly-configured DNS environment could cause problems with database applications. If you try to restore a database and it fails with the message `Cannot connect to target database` or `Cannot create restore set`, the problem is as follows:

When backing up the database on a system, the agent that starts on the system logs the system's name to the database as *<system.company.com>*. The Restore Session Manager wants to restore

17802
Paula

to the `<system_name.company.com>`, but it cannot because it does not know this system as `<system_name.company.com>`, but only as `<system_name>`. The system name cannot be expanded to the full name because the DNS is improperly configured. This situation can also be the other way around, where DNS is configured on the Cell Manager and not on the Application Client.

Action Set up the TCP/IP protocol and configure DNS properly. Refer to Appendix B in the *HP OpenView Storage Data Protector Installation and Licensing Guide* for information.

Problems with non-ASCII Characters in Filenames

In mixed platform environments, there are some limitations regarding filenames containing non-ASCII characters.

Action See Appendix B in the *HP OpenView Storage Data Protector Software Release Notes* for types of problems that occur, for situations in which they occur, and for workarounds.

17801
Paula

Troubleshooting Data Protector Installation

If you run into problems while installing the Data Protector software, check the system's log files on UNIX and setup log files on Windows to determine the problem:

System	Log File
UNIX (local installation)	/var/adm/sw/swinstall.log /var/adm/sw/swagent.log
UNIX (remote installation)	/var/opt/omni/log/IS_install.log
Windows (local installation)	*<System_disk>*:*<Temp>*\OB2_Setup_ui_*<Date>*_*<Time>*.txt
Windows (remote installation)	*<System_disk>*:*<Temp>*\OB2_Setup_exe_*<Date>*_*<Time>*.txt

If the setup log files were not created, run the installation with the -debug option.

Problems with Remote Installation of Windows Clients

When using Data Protector remote installation to update Windows clients, you get the following error:

```
Error starting setup process, err=[1326] Logon failure:
unknown user name or bad password.
```

The problem is that the `Data Protector Inet` service on the remote computer is running under a user account that does not have access to the OmniBack II share on the Installation Server computer. This is most probably a local user.

Action

Change the user name for the `Data Protector Inet` service that can access the OmniBack II share.

17800
Paula

Name Resolution Problems when Installing the Windows Cell Manager

During the installation of the Data Protector Cell Manager on Windows, Data Protector detects and warns you if the DNS or the LMHOSTS file is not set up as required. In addition, Data Protector notifies you if the TCP/IP protocol is not installed on your system.

Name resolution fails when using DNS or LMHOSTS

If the name resolution fails, the `"error expanding hostname"` message is displayed and the installation is aborted.

- If you encounter resolution problems when using DNS, you get a warning message about your current DNS configuration.
- If you encounter resolution problems when using LMHOSTS file, you get a warning message to check your LMHOSTS file configuration.
- If you have not configured either DNS or LMHOSTS, you get a warning message to enable the DNS or the LMHOSTS resolution in the TCP/IP properties dialog.

Action

Check your DNS or LMHOSTS file configuration or activate it. Refer to "Hostname Resolution Problems" on page 565.

The TCP/IP protocol is not installed and configured on your system

If the TCP/IP protocol is not installed and configured on your system, the installation is aborted.

Data Protector uses the TCP/IP protocol for network communications; it must be installed and configured on every client in the cell.

Action

Check the TCP/IP setup. For detailed steps, refer to the online Help index keyword "checking, TCP/IP setup".

Troubleshooting User Interface Startup

Data Protector user interface start-up problems are usually the result of services not running, services not being installed, or problems with network communication.

Inet Is Not Responding on the Cell Manager

The following message appears:

```
Cannot access the system (inet is not responding). The Cell
Manager host is not reachable, is not up and running, or has
no Data Protector software installed and configured on it.
```

Action

If communication between the systems is not the problem, check the installation using `telnet`.

It is possible that some components were not or were improperly installed. Review the steps in the *HP OpenView Storage Data Protector Installation and Licensing Guide*.

If the installation is correct, run the `omnisv -status` command to check whether the services on the Cell Manager are running properly.

No Permissions to Access the Cell Manager

The following message appears:

```
Your Data Protector administrator set your user rights so
that you do not have access to any Data Protector
functionality.

Contact your Data Protector administrator for details.
```

Action

Contact the Data Protector administrator to add you as a user and give you appropriate user rights in the cell. Refer to Chapter 3, “Configuring Users and User Groups,” on page 81.

17.795
Paula

Connection to a Remote System Refused on Windows or Novell NetWare

The response of the `telnet <hostname> 5555` command is `Connection refused`.

Action

If the `Data Protector Inet` service is not running on the remote system, run the `omnisv -start` command to start it.

If Data Protector is not installed on the remote system, install Data Protector on the remote system.

Connection to Windows 98 Clients Fails

Problems connecting to a Data Protector Windows 98 client system can be identified when using the `telnet <hostname> 5555` command or when the following error message appears in the Data Protector message log window: `Cannot connect to inet for getting filesystem list on <hostname>`.

Action

Run the `<Data_Protector_home>\bin\omnii95` command to make sure that the `Data Protector Inet` service is running on the Windows 98 client.

Run the `<Data_Protector_home>\bin\omnii95 -kill` command to stop the `Data Protector Inet` service on a Windows 98 client.

17797
Paula

Troubleshooting the IDB

This section provides troubleshooting for the following problems using the IDB:

- "Problems During the Upgrade of the IDB on Solaris" on page 592
- "Problems While Running the User Interface" on page 595
- "Libraries (Executables) Missing" on page 595
- "Data Files (Directories) Missing" on page 596
- "Temporary Directory Missing" on page 597
- "Problems During Backup and Import" on page 598
- "Performance Problems" on page 599
- "MMDB and CDB Are Not Synchronized" on page 600

Problems During the Upgrade of the IDB on Solaris

Once the IDB upgrade on Solaris is started, the `upgrade.log` file is created in the `/var/opt/omni/log` directory.

The file contains core and detail part upgrade messages, which enable you to see the status of the upgrade. You can see when a session was started and ended, as well as any problems that occurred during the upgrade.

You can also run the `omnidbutil -upgrade_info` command from the command line to display the current status of the IDB upgrade. The possible return values for the command are:

- `No upgrade in progress.`

  Database was initialized, the core upgrade was not started.

- `Upgrade of core part failed.`

  The core upgrade was started, but failed.

- `Upgrade of core part finished.`

  The core upgrade finished successfully, the detail upgrade was not started.

17796
Paula

- `Upgrade of detail part running.`

  The detail upgrade was started and is currently running.

- `Upgrade of detail part finished.`

  The detail upgrade finished.

The Cell Manager crashes during the core upgrade

The following methods can help you to identify the problem:

- The `omnidbutil -upgrade_info` command on the Cell Manager reports `Core upgrade failed.`
- The `upgrade.log` file on the Cell Manager contains the `UCP session started` entry but does not contain either the `UCP session finished` nor the `Session was aborted (Upgrade core part)` entry.

Action

1. Run `/opt/omni/sbin/omnisv -status` to check whether the Data Protector services are running on the Cell Manager. If they are not running, start them using the `/opt/omni/sbin/omnisv - start` command.
2. Run the `omnidbinit` command on the Cell Manager to initialize the new database. The A.03.51 IDB is set to read-only mode and is therefore left intact.
3. Run the `omnidbupgrade -ucp` command to restart the core upgrade.
4. When the core upgrade is finished, continue with the detail part upgrade.

The core upgrade runs out of disk space

The `upgrade.log` file on the Cell Manager contains the `Not enough disk space` or `not enough configure extension/binary files to perform upgrade` entry after the `UCP session started` entry. Refer to the *HP OpenView Storage Data Protector Software Release Notes* for disk space requirements.

Action

1. Free up disk space on the system where the IDB is installed.
2. Run the `omnidbinit` command to initialize the new database. The A.03.51 IDB is set to read-only mode and is therefore left intact.
3. Restart the core upgrade using the `omnidbupgrade -ucp` command.

17.745 Paula

The detail upgrade runs out of memory on the system

- The upgrade.log file on the Cell Manager contains the database network communication error entry.
- Data Protector Event Log contains the Session was aborted. (Upgrade detail part.) entry.
- On UNIX Cell Managers, the Data Protector RDS daemon (/opt/omni/lbin/rds) is not displayed when listing the Data Protector processes using the ps -ef | grep omni command. On Windows Cell Managers, low virtual memory notification may be received in the Windows Event Log and the Data Protector RDS process (rds.exe) may not be listed among processes in Windows Task Manager.

Action

1. Close any applications that do not need to run.
2. Run the omnisv -stop and omnisv -start commands from the /opt/omni/sbin directory to stop and restart the Data Protector services. Detail upgrade is automatically restarted during the Data Protector services restart.
3. If the problem persists, add more RAM to the computer. On UNIX systems, you can also configure a bigger data segment in the operating system. Refer to the *HP OpenView Storage Data Protector Software Release Notes* for a list of installation requirements.

The detail upgrade runs out of disk space on the system

The upgrade.log file on the Cell Manager contains the Not enough disk space or not enough configured extension/binary files to perform upgrade entry after the UDP session started entry. Refer to the *HP OpenView Storage Data Protector Software Release Notes* for disk space requirements.

Action

1. Using system applications, find out how much disk space is still available on the disk where the database resides and run the omnidbupgrade -udpcheck command to see how much disk space is required to run the detail upgrade of the database.
2. Free up space on the disk, where the IDB is installed.
3. Run the omnidbupgrade -udp command to restart the detail upgrade of the database.

17794
Paula

Problems While Running the User Interface

IDB is corrupted

Any of the following messages can be displayed:

- `Database is corrupted.`
- `Interprocess communication problem.`
- `Cannot open Database/File.`
- `Error - Details Unknown.`

Action

Recover the IDB. For more information, refer to "Recovering the IDB" on page 417.

The IDB Session Manager is not running on the Cell Manager

If the IDB Session Manager is not running on the Cell Manager when Data Protector tries to access or use the IDB, the `Interprocess communication problem` error message is displayed.

- On Windows Cell Manager, the Data Protector process `dbsm.exe` is not displayed among the processes in the Windows Task Manager.
- On UNIX Cell Manager, the `/opt/omni/lbin/dbsm` is not displayed when listing the Data Protector processes using the `ps -ef | grep omni` command.

Action

Close and restart the Data Protector GUI.

Libraries (Executables) Missing

On Windows Cell Managers, the following library files should exist in the `<Data_Protector_home>\bin` directory:

- `libob2ecmn.dll, libob2eadm.dll, libob2ecdb.dll, libob2emmdb.dll, _eadm32.dll, _erdm32.dll`

On UNIX Cell Managers, the following library files should exist in the `/opt/omni/lib` directory:

- `libob2ecmn.sl, libob2eadm.sl, libob2ecdb.sl, libob2emmdb.sl, _eadm.sl, _erdm.sl`

17 713 Paula

The RDS service/process cannot be started

If one or several shared library files are missing, the `omnisv -status` command informs you that the RDS service/process is down, while all other services/processes are running.

Action Reinstall Data Protector and reboot your Cell Manager. This will reinstall the shared libraries and restart the RDS service/process.

Data Files (Directories) Missing

The following IDB data files (directories) should exist on the Cell Manager in the following directories:

- On Windows systems: *`<Data_Protector_home>`*`\db40`
- On UNIX systems: `/var/opt/omni`

```
\datafiles\catalog

\datafiles\cdb

\datafiles\mmdb

\dcbf

\logfiles\rlog

\logfiles\syslog

\meta

\msg
```

One or several IDB data files or directories are missing

If one or several IDB data files or directories are missing, the following errors are displayed when Data Protector tries to access or use the IDB:

- `Database network communication error`
- `Cannot open database/file`

Action Reinstall Data Protector and reboot your Cell Manager. This will reinstall the IDB data files and directories.

17792
Paula

Temporary Directory Missing

The following temporary directories should exist on the Cell Manager:

- On Windows: `<Data_Protector_home>\tmp`
- On UNIX: `/var/opt/omni/tmp`

The Data Protector GUI cannot connect to the Cell Manager

When Data Protector GUI tries to connect to the Cell Manager, the following error message is displayed if Data Protector temporary directory is missing:

```
Cannot access the Cell Manager system. (inet is not
responding) The Cell Manager host is not reachable or is not
up and running or has no Data Protector software installed
and configured on it.
```

Action

1. Close the Data Protector GUI.
2. Run the `omnisv -stop` command on the Cell Manager to stop the Data Protector services/processes:
    - On Windows: `<Data_Protector_home>\bin\omnisv -stop`
    - On UNIX: `/opt/omni/sbin/omnisv -stop`
3. On the Cell Manager, manually create the temporary directory:
    - On Windows: `<Data_Protector_home>\tmp`
    - On UNIX: `/var/opt/omni/tmp`
4. Run the `omnisv -start` command to start the services/processes.
    - On Windows: `<Data_Protector_home>\bin\omnisv -start`
    - On UNIX: `/opt/omni/sbin/omnisv -start`
5. Restart the Data Protector GUI.

17.791
Paula

Problems During Backup and Import

The BSM or RSM is terminated during the IDB backup or import session

If the BSM or RSM get terminated during the IDB backup or import session, the following error message is displayed:

```
IPC Read Error System Error: [10054] Connection reset by
peer
```

In the `Internal Database` context, the session status of the IDB backup or import session is still marked as `In progress` but the session is actually not running.

Action

1. Close the Data Protector GUI.
2. Run the `omnidbutil -clear` command to set the status of all sessions that are actually not running but are marked as `In Progress` or `Failed`, to `Failed`.
3. Run the `omnidbutil -show_locked_devs` command to see if any devices and media are locked by Data Protector.
4. If there are, run the `omnidbutil -free_locked_devs` to unlock them.
5. Restart the Data Protector GUI.

The MMD is terminated during the IDB backup or import session

If the media management daemon MMD is terminated during the IDB backup or import session, the following two error messages are displayed:

- `Lost connection to MMD`
- `IPC Read Error System Error: [10054] Connection reset by peer`

Use the following methods to check whether the MMD services/processes are running:

- The `omnisv -status` command informs you that the MMD service/process is down.

17.790
Paulo

- On UNIX, the Data Protector MMD (`/opt/omni/lbin/mmd`) is not displayed when listing the Data Protector processes using the `ps -ef | grep omni` command.

  On Windows, the Data Protector MMD process (`mmd.exe`) is not listed among processes in the Windows Task Manager.

Action

1. Close the Data Protector GUI.
2. Run the `omnisv -stop` command to stop the Data Protector services/processes.
3. Run the `omnisv -start` command to start the Data Protector services/processes.
4. Run the `omnisv -status` command to check if all the services/processes are running.

The DC binary files are corrupted or missing

If the DC binary files are corrupted or missing, the error message `Open of Detail Catalog Binary File failed` is displayed when browsing backed up objects in the `Restore` context.

- The `omnidbcheck -bf` command reports that one or several DC binary files are missing or of incorrect size, or the `omnidbcheck -dc` command reports that one or several DC binary files are corrupted.
- The `debug.log` file on the Cell Manager, located in the `<`*`Data_Protector_home`*`>\log\debug.log` (Windows systems) or in the `/var/opt/omni/log/debug.log` (UNIX systems) contains one or several entries on Data Protector not being able to open a DC binary file.

Action

Recreate DC binary files by importing catalog from media. For more information refer to "Handling Minor Database Corruption in the DCBF Part" on page 422.

Performance Problems

The number of IDB objects and IDB objects' sizes are too large

When browsing object versions and single files for restore, it can take a long time before the information is read from the IDB and displayed.

Action

Set the time interval, which will be used when browsing object versions for restore. You can change this time interval in the `Restore` context when searching for the specific object version you want to restore.

Set the *default* time interval used when browsing object versions for restore.

1. In the Data Protector GUI, click the `File` menu and then click `Preferences`.
2. Click the `Restore` tab and in the `Search interval` drop-down list, select the search interval. Select `Interval` if you want to set an absolute search interval, or `None` if you want all object versions to be listed.
3. Click `OK` to apply the change.

MMDB and CDB Are Not Synchronized

The MMDB and CDB may not be synchronized when the following is true:

- The MMDB and CDB contain information from different periods in time. This may be the result of importing the CDB and the MMDB (the `omnidbutil -readdb` command) from files generated in separate export (the `omnidbutil -writedb` command) sessions.
- In a MoM environment, when the local CDB and CMMDB are not synchronized. This may be the result of the CMMDB restore.

Data Protector reports when an object in the IDB has no medium assigned or when the data protection for a medium is not correctly set.

Action

In a one-cell environment:

- Run the `omnidbutil -cdbsync` *`<Cell_Server_Hostname>`* command in the `/opt/omni/sbin` (UNIX Cell Manager) or in the *`<Data_Protector_home>`*`\bin` (Windows Cell Manager) directory to synchronize the MMDB and CDB.

In a MoM environment:

- Run the `omnidbutil -cdbsync` *`<Cell_Server_Hostname>`* command in the `/opt/omni/sbin` (UNIX Cell Manager) or in the *`<Data_Protector_home>`*`\bin` (Windows Cell Manager) directory with the CMMDB installed (MoM). Run this command for every Cell Manager in the MoM environment by specifying its hostname as the argument.

17.787
Paula

Troubleshooting Reporting and Notifications

If you use Outlook XP or Outlook 98/2000 with the latest security patch installed, you following problem appears: when you add a report to a report group specifying email as a send method, and then try to start a report group, GUI hangs. The same happens if you configure a notification and select the email send method. The cause of the problem is that Outlook requires user interaction before sending an email notification. This feature cannot be disabled since it is a part of the Outlook security policy. To solve this problem, start a report from the CLI:

```
omnirpt -report licensing -email <email_address>
```

When a warning asking whether you allow sending email on your behalf appears, click `Yes` to receive a notification.

For more information on how to customize security settings, refer to *HP OpenView Storage Data Protector Software Release Notes*.

Troubleshooting Data Protector Online Help

Data Protector online Help consists of two parts: Help Topics and the Help Navigator. Help Navigator is context-sensitive help, explaining screens and options in the Data Protector GUI, while Help Topics provide conceptual information, procedure instructions, and examples.

The Help system you use depends on the platform (Windows or UNIX) on which you are running Data Protector. You use HTML Help on Windows systems and WebHelp on UNIX systems.

Troubleshooting Online Help on Windows

When accessing online Help on Windows systems, you can run into the following Help Navigator display problem:

The Help Navigator contents do not change in parallel with the Data Protector windows.

Action

1. If you use `Microsoft HTML Help` mode (default option), ensure that the button shown below is enabled.

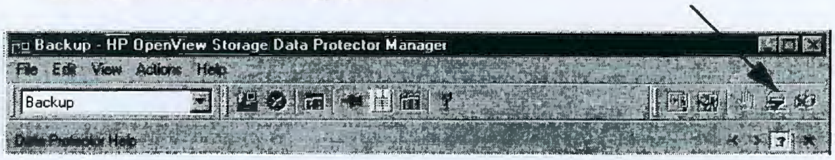
Backup - HP OpenView Storage Data Protector Manager
File Edit View Actions Help
Backup
Data Protector Help

2. If you use `Default HTML Browser` mode (an external HTML browser for displaying the help files) go to `File` menu, click `Preferences` and enable the `Check the box to enable the context-sensitive help navigator` option. Then restart the Help Navigator.

Troubleshooting Online Help on UNIX

If your browser (HTML viewer) is not properly set, you can run into online Help start and display problems. You need to set the browser as follows:

Action

1. In the `File` menu, click `Preferences`. In the drop-down list, select `Netscape`, if your browser is Netscape Navigator. If your browser is

17.785
Paula

not Netscape Navigator, select `Custom`.

IMPORTANT Data Protector supports only Netscape Navigator for online Help viewing.

2. Click `Settings` to open the `HTML Viewer Settings` window.
3. In the `Location of executable script or binary file` text box, enter the location of your browser (for example, `/opt/netscape`).
4. In the `Command to start viewer` text box, enter the command that will start the browser. For Netscape Navigator, enter `netscape $HTML$`.
5. In the `Command to reuse existing viewer window` text box, you can enter a command that will be used to open each HTML file in the same window. If you do not enter the command, each HTML file will be opened in a separate window. For Netscape Navigator this command is `netscape -remote OpenFile($HTML$)`.

17734
Paula

Check Whether Data Protector Functions Properly

The following sections provide an overview of the Data Protector Checking and Maintenance Mechanism and an overview of things to be checked in order to determine whether Data Protector is properly configured in your backup environment.

Data Protector Checking and Maintenance Mechanism

Data Protector provides its own checking and maintenance mechanism, which is performing the following checking and maintenance tasks on a daily basis:

Maintenance Tasks

- Deletes obsolete DC binary files, sessions, and related messages every day at 12:00 (Noon) by default.
- Finds any free (unprotected) media in pools with the `Use free pool` and `Move free media to free pool` options set and deallocates the found free media to a free pool by issuing the following command every day at 12:00 (Noon) by default:

  ```
  omnidbutil -free_pool_update
  ```

  For more information on the `omnidbutil` command, refer to the `omnidbutil` man page. For more information on the above mentioned options, refer to Chapter 4, "Managing Media," on page 97.

Checks

Every day at 12:30 P.M. by default, starts checks for the following Data Protector notifications:

- Database Space Low
- Not Enough Free Media
- Health Check Failed
- User Check Failed
- Unexpected Events
- License Will Expire

17.483
Paula

- Database Purge Needed

For more information on Data Protector notifications, refer to “Data Protector Notifications” on page 342. Any notification that is triggered is by default sent to the Data Protector Event Log. For more information on the Data Protector Event Log, refer to“Data Protector Event Log” on page 356.

The default schedule values for maintenance tasks and checks can be changed by changing the options in the Data Protector global options file. Refer to “Global Options File” on page 523 for more information on global options.

The User Check Failed Notification

The `User Check Failed` notification automates the task of checking whether your backup environment is functioning normally. Note that the definition of “normal” depends on your backup environment (backup policy, network configuration, hardware used, etc.). For an overview of items to be checked in an “average” backup environment, refer to “Overview of Items to Be Checked” on page 607. For more information on Data Protector notifications, refer to “Data Protector Notifications” on page 342.

The `User Check Failed` notification executes the command or script entered as an input parameter to this notification and triggers the notification if the return value of the executed command or of any of the executed commands in the script is other than zero. The command/script should be created in the `/opt/omni/lbin` (on UNIX systems) or `<`*`Data_Protector_home`*`>\bin` (on Windows systems) directory of the application system. The `User Check Failed` notification can be configured to be sent using various send methods (e-mail, broadcast message, SNMP traps, log file, etc.) when it is triggered. It can also be configured to start a Report Group when it is triggered.

Thus, scripts containing checks specified in accordance with your backup environment can be developed and configured in a `User Check Failed` notification. Data Protector, using its maintenance and checking mechanism, then prompts you whenever something goes wrong in your backup environment.

All *configured* `User Check Failed` notifications are by default scheduled to be started every day at 00:00 (Midnight) and are, if triggered, sent to Data Protector Event Log.

17.782
Paula

Overview of Items to Be Checked

In order to ensure that Data Protector is functioning properly and to identify potential problems before they arise, it is recommended that you perform regular checks as described in the following sections.

Using the `User Check Failed` notification, it is possible to automate these checks by developing scripts including these checks. Some of the checks (for example the `omnihealthcheck` and `omnitrig -run_checks` commands) are already automated by the means of Data Protector checking and maintenance mechanism.

Check the Data Protector Cell Manager

1. Run the `omnihealthcheck` command to check the following:

   - whether the Data Protector services (`rds`, `crs`, `mmd`, `omnitrig`, and `OmniInet`) are active
   - whether the Data Protector Media Management Database is consistent
   - whether at least one backup of the IDB exists

   The exit code of the command is 0 (OK) only if all three checks completed successfully (exit code for every check was 0). Exit values other than 0 indicate that one or more of the checks failed.

   For more information on exit codes, refer to `omnihealthcheck` man page.

2. Run the `omnidbcheck -core` command to check the core parts of the IDB.

   The exit code of the command is 0 (OK) only if the check completed successfully. Exit values other than 0 indicate that the check failed.

   For more information on exit codes, refer to `omnihealthcheck` man page.

3. Check the critical parts of IDB using the `omnidbcheck -critical` command. For more information on the `omnidbcheck` command, refer to `omnidbcheck` man page.

   The exit code of the command is 0 (OK) only if the check completed successfully. Exit values other than 0 indicate that the check failed. For more information on exit codes, refer to `omnidbcheck` man page.

17.781
Paula

Check whether backups are configured properly

1. Run the backup preview for crucial backup specifications. Refer to Chapter 5, "Backup," on page 151 for more information on previewing backups. Successfully completed previews prove that:
   - All clients in the backup specification are accessible from the Cell Manager.
   - All files are accessible.
   - The amount of data to be backed up is determined.
   - All backup devices are configured properly.
2. Run the `omnirpt -report dl_sched` command to check whether the backup specifications are scheduled in compliance with your backup policy. For more information on `omnirpt` command, refer to `omnirpt` man page. The command will list all backup specifications and their schedule.

Verify the Data Protector installation

Verify the installation using the Data Protector GUI, `Clients` context, to check whether the Data Protector software components are up and running on the Cell Managers or the clients. Refer to the *HP OpenView Storage Data Protector Installation and Licensing Guide* for information on how to verify the Data Protector installation.

Inspect the Data Protector log files

Inspect the following Data Protector log files and identify possible problems:

- `event.log`
- `debug.log`
- `purge.log`

For more information on Data Protector log files, refer to "Data Protector Log Files" on page 550.

Run the Notifications Checks

Any Data Protector notification that is triggered is sent to Data Protector Event Log by default. You can also run the `omnitrig -run_checks` command to start checks for the following notifications:

17.780
Paula

- ✓ Database Space Low
- ✓ Not Enough Free Media
- ✓ Health Check Failed
- ✓ User Check Failed
- ✓ Unexpected Events
- ✓ License Will Expire
- ✓ Database Purge Needed

For more information on Data Protector notifications, refer to "Data Protector Notifications" on page 342. For more information on Data Protector Event Log, refer to"Data Protector Event Log" on page 356.

Check Other System Resources

Inspect the following operating system log files and identify possible problems:

- On UNIX systems: `/var/adm/syslog/syslog.log`
- On Windows systems: inspect the `Windows Event Viewer` and its `Security`, `System` and `Application` logs.

Check whether IDB System Configuration Backups are Being Made Regularly

Check the Data Protector recovery file, `obrindex.dat`, to make sure that the IDB and configuration files, needed for successful recovery of a system, are created regularly. For more information on `obrindex.dat` file, refer to "Preparing for IDB Recovery" on page 390.

17.779
Paula

DOC - 000233

CPL/AC

PREGÃO
050/2003

LOCAÇÃO DE
EQUIPAMENTOS
DE INFORMÁTICA
INCLUINDO
ASSISTÊNCIA
TÉCNICA E
TREINAMENTO

COBRA
TECNOLOGIA –
MANUAL
VOLUME 4

2003
PASTA 32

19395
E

Table 1 Disk LED Activity Defiinitions 12
Table 2 Disk System Power LED Activity Definitions 13
Table 3 AC Power Requirements 15
Table 4 SCSI Switch Setting Definitions 17
Table 5 SCSI Address Priority 18
Table 6 LED Status Indicators 29
Table 7 Troubleshooting Table 34
Table 8 Products and Options 35
Table 9 Available Upgrade Disk Modules 36
Table 10 Upgrade Products for Hewlett-Packard Systems 36
Table 11 Replaceable Parts 37
Table 12 Supported HP Cables and Terminators (for Enterprise Systems) 37
Table 13 Supported HP Cables and Terminators (for Commercial Systems) 38
Table 14 Physical Dimensions 38
Table 15 AC Power Requirements 38
Table 16 Environmental Specifications 39

19403
E

ANEXO UNIDADE DE BACKUP ROBOTIZADO
PARTE T

19402
E

HP SureStore Disk System 2100

User's Guide

Edition E0301
Order No. A5675-90901
Printed in U.S.A.

invent

19401
E.

FS-3800

Especificações

Velocidade de impressão		25 ppm
Primeira página		9 segundos c/ a máquina pronta e em linha
Tempo de Aquecimento		< 35 segundos ao ligar
Processador		PowerPC 740 200MHz
Ciclo mensal de trabalho		médio de 100.000 páginas por mês (máximo de trabalho de 120.000 páginas)
Memória RAM padrão		16MB
Memória máxima		272MB (incluindo expansão de 128MB x 2 módulos DIMM)
Cartão memória (Compact Flash)		Compact Flash Tipo I (Compact Flash /até de 256 MB)
Micro disco Rígido (Opcional)		Micro Disco Rígido até 1GB (Tipo HDD)
Emulações		PCL6, PostScript Nível 2 e 3***, PJL, Diablo 630,IBM Proprinter X24E, Epson LQ-850, Epson FX, ESC P2, Line Printer, KPDL 2 (Padrão)
Fontes		1 fonte Bitmap, 80 fontes escaláveis PCL e PostScript, 45 fontes códigos de barra, incluindo a Bidimensional PDF-417
Conectividade		I/F paralela bi-direcional (IEEE1284 de alta velocidade), I/F RS-232C/ 422(Opcional) e I/F opcional (Tipo KUIO-LV)
Resolução	Modo fino 1200:	1200 x 1200 dpi
	Modo rápido 1200:	1800 x 600 dpi com KMIR*
	Modo Normal:	600 x 600 dpi, 2400 x 600 dpi com KMIR*
Entrada de papel (padrão)		Bandeja interna de 500 folhas e bandeja multiuso de 100 folhas
Saída de papel (padrão)		250 folhas (Opcional de 300 Folhas)
Tipos de mídia	Bandeja interna	Papel comum/ reciclado/ espesso (60 a 105g/m2)
	Bandeja multiuso	Papel comum/ reciclado/ espesso (60 a 200g/m2), Filme OHP, Etiqueta adesiva, Envelope, Cartão Postal
Tamanho da mídia	Bandeja inter.	A4, A5, B5, Carta, Clientizado (147,3 x 210,8 a 215,9 x 297,2)mm
	Bandeja multiuso	A4, A5, B5, Carta, Ofício, Clientiz. (71 x 147,3 a 215,9 x 297,2)mm
Placas de rede opcionais (NIC´s)		Ethernet 10Base-T/100Base-TX, 10Base2, Token-Ring, Ether-Kyo-FL (Infravermelho), etc.
Opcionais (Entrada de Papel)	PF-60	Alimentador multiformato de 500 folhas (Máx. de 3)
	PC-60LG	Bandeja de 500 folhas de papel ofício (uso como band. interna)
	PF-8E	Alimentador especial de 2.000 folhas (requer acessório PB-60)
	EF-60	Alimentador de Envelopes (Máx. de 70 envelopes)
	DU-60	Duplexador - Impressão frente/verso (papel de 60 a 105g/m2)
Opcionais (Saída de Papel)	SO-60	Alceador / empilhador c/5 caixas de saída (Máx. 200 folhas por caixa)
	HS-8E	Empilhador para 2.000 folhas
	PT-60 / PT-4	Bandeja de saída (face p/ cima) para 250 folhas
Opcionais	**BC-1**	Leitora de código de barras (opcional)
Toner	**TK-60**	Toner monocomponente, 20.000 páginas @ 5% de cobertura**
Emissão de ozônio		0.02mg/m3
Nível de ruído		Imprimindo - 53dB, Em espera - 35dB
Dimensions (A x L x P)		(300 x 345 x 391) mm
Peso		13 Kg.
Compatibilidade:		Windows 95/98/ME/NT/2000/XP; Unix; Linux; Novell

*Modo de refinamento de imagem Kyocera Mita **O toner de inicialização é de 10.000 pág. @ 5% de cobertura ***Exclusivo para a América Latina / Brasil

ECOSYS® O coração das impressoras Ecosys é o nosso cilindro de silício amorfo (a-Si), o qual tem uma durabilidade aproximadamente trinta vezes maior que os cilindros convencionais. A tecnologia " cartridge-free " elimina a necessidade da troca do cilindro e do revelador, tratados como consumíveis nas impressoras convencionais, a cada troca do toner. Por isso, as impressoras Ecosys tem o menor custo total de propriedade (TCO) do mercado.

Economia

O custo de impressão é excepcionalmente baixo, pois o toner é básicamente o único item de consumo.

Ecologia

As impressoras Ecosys não agridem o meio ambiente pois o único "lixo" gerado é a vasilha de toner vazia.

System

PRESCRIBE 2e: linguagem de programação e elaboração de formulários com funções gráficas expandidas para tarefas em todos os sistemas.

Cartão Memória Compact Flash:

Para cartão compacto, até de 256MB. Para armazenar formulários eletrônicos, fonte especiais, logotipos, etc.

Leitora de cód. de barras:

Para recuperar trabalhos de impressão armazenados no HD.

Kyocera Mita Net Viewer

Para gerenciamento e monitoramento das impressoras conectadas em rede.

energy

Como empresa membro da ENERGY STAR, a Kyocera Mita America, Inc. assegura que este modelo encontra-se de acordo com as diretrizes da organização.

TN INDUSTRIAL S.A

Fabricante das Impressoras Kyocera Mita no Brasil

KYOCERA mita

Brasília
Setor de rádio TV norte,
Q 701 lote C Bl. B sala 200
térreo.Centro Empresarial Norte
Distrito Federal.
CEP 70719-000.
Tel: (61)327-1975 ou 327-7008.

Rio de janeiro
(escritório Central)
Rua Buenos Aires, 90 8° andar
Rio de Janeiro RJ Cep: 20.070-020
Tel.: (21) 3231-6446 / Fax: (21) 2242-4785

Fábrica
Rodovia Darly Santos, 4800 Araçás
Vila Velha ES Cep: 29.103-460
Tel.: (27) 3339-8800
Fax: (27) 3349-7088

Suporte Técnico
Tel.: 0800 7034347
e-mail: suporte@kyocera.com.br

*** Exclusivo para a América Latina- TN Industrial

19400
E

Notice

© Hewlett-Packard Company, 2000. All rights reserved.

Hewlett-Packard Company makes no warranty of any kind with regard to this document, including, but not limited to, the implied warranties of merchantability and fitness for a particular purpose. Hewlett-Packard shall not be liable for errors contained herein or for incidental or consequential damages in connection with the furnishing, performance, or use of this material.

This document contains proprietary information, which is protected by copyright. No part of this document may be photocopied, reproduced, or translated into another language without the prior written consent of Hewlett-Packard. The information contained in this document is subject to change without notice.

Warranty

If you have any questions about the warranty for this product, contact your dealer or local Hewlett-Packard sales representative.

Safety Notices

⚠ Warning

Weight exceeds 50 lbs. (22.5 kg.)

Do NOT lift unassisted. Use a lift device or two people.

To protect against personal injury and product damage, do not attempt to lift the product without the assistance of another person or lift device.

Components bearing this symbol may be hot to touch.

Components bearing this symbol are fragile. Handle with care.

Components bearing this symbol are susceptible to damage by static electricity. ESD precautions are required.

Operation

The front door should be closed and locked at all times during the operation of this product except when replacing disks.

This product is intended to be operated in a restricted access area.

Service

Maintenance or repair of the backplane and mezzanine boards must be performed by authorized service-trained personnel.

Format Conventions

	Denotes
WARNING ⚠	**A hazard that can cause personal injury**
CAUTION	A hazard that can cause hardware or software damage
Note	Significant concepts or operating instructions
this font	Text to be typed verbatim: all commands, path names, and file names. Also menu and button selections in GUI contexts
`this font`	Text displayed on the screen

1939
ε

Components of the Disk System 12
Front Panel 12
System Power LED 12
Rear Panel 13
SCSI ID Switch 14
Power Connector 15
Electrical Requirements. 15
AC Site Requirements 15
Installing the Disk System 16
Hardware Requirements 16
Preparing for Installation. 16
Setting the SCSI IDs for the Disk Modules 17
Installing the Disk System 18
Connect the SCSI Cables 18
Connecting the Power Cable 20
Powering On and Off 20
Power On the Disk System 20
Power Off the Disk System 21
Adding Disk Modules 22
Add a Disk Module. 22
Configure the new Disk Module 23
Remove and Replace A Disk Module 23
Removing a Disk Module 23
Replace a Disk Module 24
Remove the Disk System. 25
Setting Up the Hardware Event Monitor (HP-UX Only) 26
Configuration Overview. 26
Event Notification (HP-UX Only) 27
Status LEDs. 29
View Disk Status 30
STM Disk Information: HP-UX. 30
STM Disk Information: MPE/iX 6.5 or 7.0. 32
Isolating Faults 34
Reference Information 35
Product Numbers and Options 35
Upgrade Disk Products Available 36
Specifications. 38
Electrical Specifications. 38
Operating Temperatures 39
Regulatory Statements 39
Safety Certifications. 39
EMC Compliance 39
A. FCC Notice for United States 40
B. FCC Notice for Canada 40

C. Notice for Japan 41
Harmonics Conformance (Japan) 41
D. Notice for Taiwan 41
E. Notice for Korea. 41
F. Notice for Germany 41
Product Web Site 43
Related Documents 43

19397
E

Figure 1. Disk Module LEDs. 12
Figure 2. Disk System Power LED 12
Figure 3. Rear View of the Disk System 13
Figure 4. Power Button 13
Figure 5. SCSI ID Switch Settings 14
Figure 6. AC Power Connector Location. 15
Figure 7. SCSI Switch Location 17
Figure 8. SCSI Port Locations 18
Figure 9. AC Power Connector Location. 20
Figure 10. Power Button Location 21
Figure 11. Installing a disk module 23
Figure 12. Removing the Disk System 25
Figure 13. Sample Hardware Event Notification 28
Figure 14. LED Status Indicators 29
Figure 15. Sample STM Information Log (HP-UX) 31
Figure 16. Sample STM Expert Tool Disk Error Log (MPE/iX 6.5 or 7.0) 33

The HP SureStore Disk System 2100 is a first-generation 1-EIA Unit disk system capable of accomodating up to four disk modules. It is designed to provide inexpensive, high performance rack-optimized online mass storage for entry-level configurations in three distinct markets: (1) HP servers/workstations, (2) open systems JBODs, and (3) open systems IDACs (Internal Disk Array Controller).

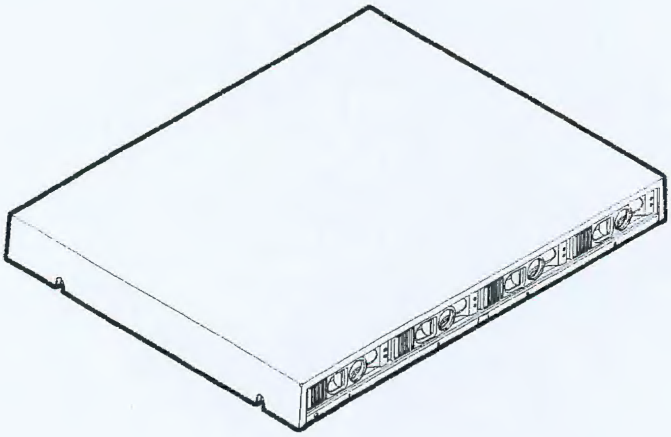

19393
E

Components of the Disk System

Front Panel

The disk modules can be accessed from the front of the disk system. It can accept up to 4 low profile disk drives. If your storage system contains less than 4 disk modules, the remaining empty slots require filler panels. These filler panels (part number A6198-60001) ensure that the proper cooling is maintained within the storage system.

FIGURE 1 Disk Module LEDs

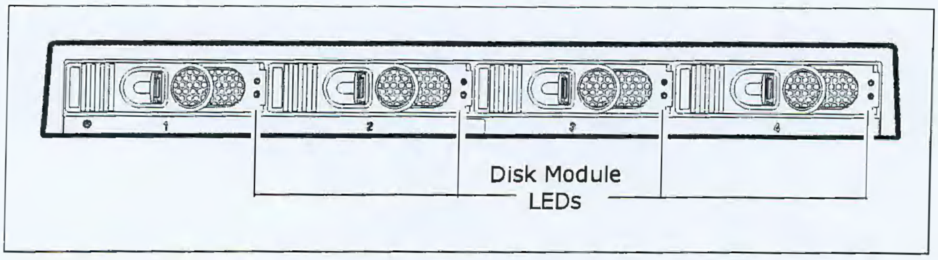

TABLE 1 Disk LED Activity Defiinitions

LED Activity	Indication
Blinking	Disk activity
All on solid for approximately1 second	Bus reset
Off	Disk is idle or off

System Power LED

The System Power LED gives diagnostic information about the disk system operation. Refer to Table 2 on page 10 for more information.

FIGURE 2 Disk System Power LED

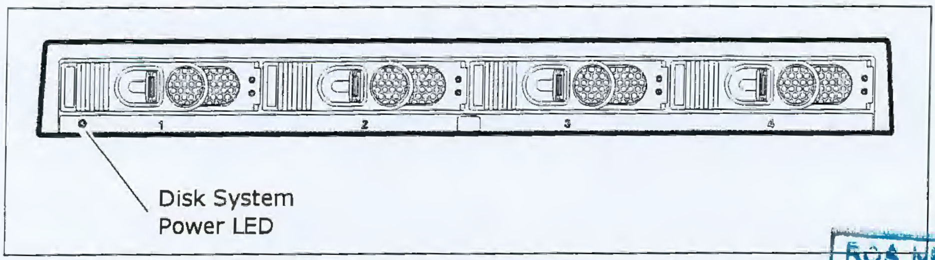

19392

TABLE 2 Disk System Power LED Activity Definitions

LED Activity	Indication
Blinking	Malfunction - either a fan is not operating properly or internal voltage is too low.
On solid	Disk system is operating properly.
Off	Disk system is off.

Rear Panel

FIGURE 3 Rear View of the Disk System

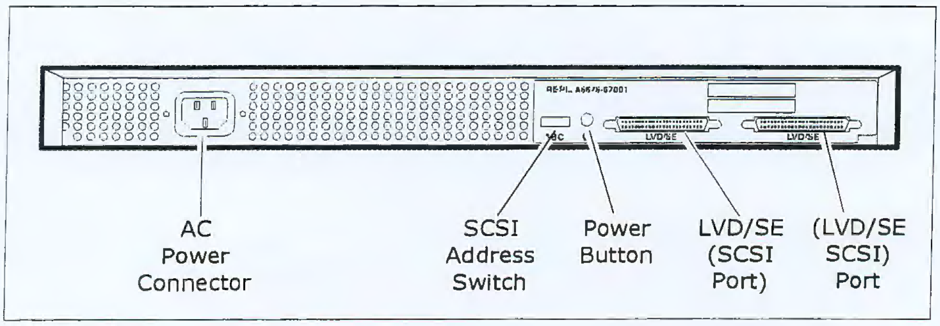

FIGURE 4 Power Button

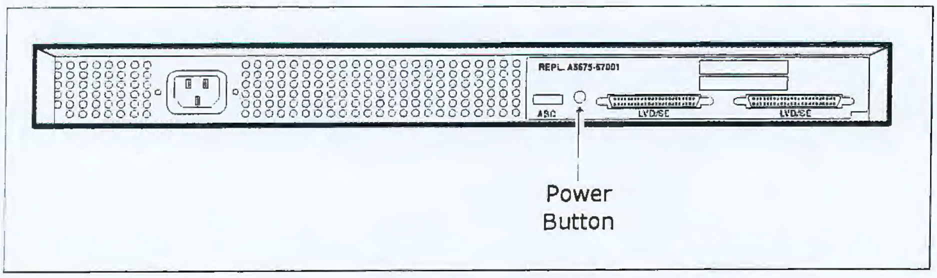

With the power button in the "ON" position, power is supplied to the disk system.

SCSI ID Switch

The SCSI ID switch sets the address for the disk drives.

There are three possible settings for the SCSI address switch: A, B, and C. Each setting assigns a different *set* of addresses to the disk modules installed in the disk system. For SCSI address switch setting A, the addresses for the installed disk modules are (from left to right, looking at the disk system from the front) 0, 2, 4, and 6. For SCSI address switch setting B, the addresses for the installed disk modules are 8, 10, 12, and 14. For SCSI address switch setting C, the addresses for the installed disk modules are 9, 11, 13, and 15. See Figure 5, below.

SCSI address 7 is reserved for use by the host bus adapter.

FIGURE 5 SCSI ID Switch Settings

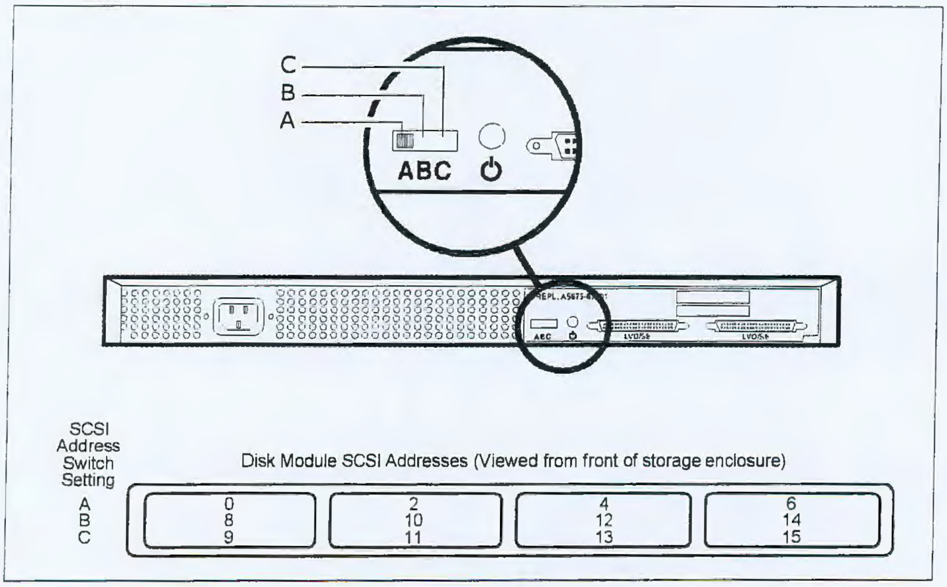

SCSI Address Switch Setting	Disk Module SCSI Addresses (Viewed from front of storage enclosure)			
A	0	2	4	6
B	8	10	12	14
C	9	11	13	15

Note The disk system requires a terminator when it is installed at the end of the SCSI bus.

19390
ε

Power Connector

FIGURE 6 AC Power Connector Location

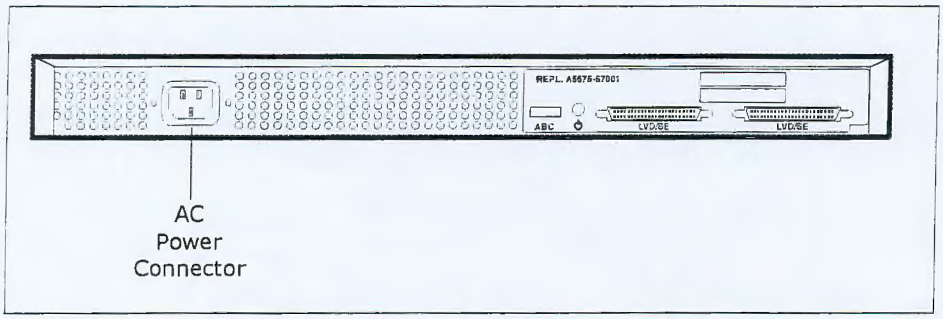

Electrical Requirements

AC Site Requirements

Overcurrent protection devices are required for each cabinet where the disk system is installed. They must be positioned between the power source and the disk system. These protective devices must not trip when exposed to an inrush current of 30 amps lasting 5 ms.

Note These protection devices must meet all applicable electrical safety requirements and be approved for the intended purpose.

TABLE 3 AC Power Requirements

Electrical Element	Requirements
Voltage	100-240 VAC
Frequency	50-60 Hz
Input current	< 1 amp
Maximum Surge Current	30 amps peak

19389
E

Installing the Disk System

Hardware Requirements

The following hardware is included with the disk system:

One power cord comes with each disk system.

Your host computer must have one of the following:

- An on board UltraSCSI port
- An UltraSCSI host adapter board installed in the host system:
  - A5149A Single Port Ultra 2 SCSI Host Bus Adapter (HP Series 9000 only)
  - A5150A Dual Port Ultra 2 SCSI Host Bus Adapter (HP Series 9000 only)
  - D5025A Ultrawide Host Bus Adapter for HP Netserver
  - D5955A Netraid 3si Host Bus Adapter for HP Netservers
  - D9161A Netraid 4M/64 Host Bus Adapter for HP Netservers
  - D9351A Netraid 4M/128 Host Bus Adapter for HP Netservers
  - P3413A Single Port Ultra 3 SCSI Host Bus Adapter for HP Netservers

The following accessories are available for your storage disk system:

- A5149A Single Port Ultra 2 SCSI Host Bus Adapter
- A5150A Dual Port Ultra 2 SCSI Host Bus Adapter
- C2364A High Density 68-pin Terminator for LVD or Single-Ended

For Ultra 3 SCSI performance, the host port or host adapter to which you connect your disk system must be Ultra 3 SCSI capable. If you connect your disk system to a SCSI narrow or wide adapter, the device performs at lower speeds. Also, if you connect the disk system to an UltraSCSI capable host bus adapter using a non-Ultra 3 SCSI cable, the data transfer rate will be lower.

Preparing for Installation

Before the disk system is ready for installation, its SCSI IDs must be set and the host system must be prepared to recognize the newly installed disk system. See your operating system administration manual for configuration procedures for the host bus adapter.

19388 E

Setting the SCSI IDs for the Disk Modules

1 Determine which SCSI addresses are not being used on the host system.

Refer to your system administration manual for information on determining which SCSI addresses are available on the host system.

2 Locate the SCSI switch on the back of the disk system.

FIGURE 7 SCSI Switch Location

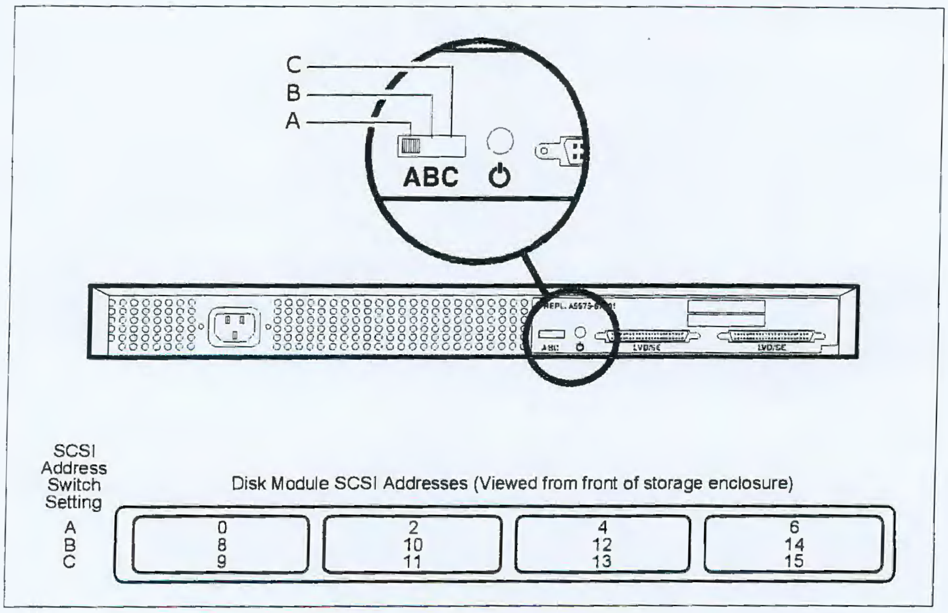
C
B
A
ABC
REPL. A5675-6
LVD/SE
LVD/SE
SCSI Address Switch Setting
A
B
C
Disk Module SCSI Addresses (Viewed from front of storage enclosure)
0
8
9
2
10
11
4
12
13
6
14
15

3 Set the SCSI IDs

Verify that the SCSI IDs you have chosen are available. If you are daisy-chaining these disk systems together, each one's SCSI ID switch must be set for a unique available set of SCSI IDs.

TABLE 4 SCSI Switch Setting Definitions

SCSI Switch Setting	Disk Slot			
	1	2	3	4
A	0	2	4	6
B	8	10	12	14
C	9	11	13	15

19387
E

Keep in mind that the host bus adapter should have the highest SCSI address priority. See Table 5 below.

TABLE 5 SCSI Address Priority

SCSI ID	7	6	5	4	3	2	1	0	15	14	13	12	11	10	9	8
Priority	Highest															Lowest

Installing the Disk System

See the installation instructions enclosed with your rackmount kit.

Connect the SCSI Cables

1 Make sure that the host system has been powered down.

2 Connect one end of an ultra SCSI cable (get it from the shipping box) to the host system. To acheive Ultra 3 SCSI speeds, the host bus adapter, the SCSI cables, and the disks must be able to transmit data at UltraSCSI 3 speeds. For the SCSI port location on your host system, refer to your host system's documentation.

You may connect non-Ultra 3 SCSI cables to the disk system, but the data transmission rates will be lower.

3 Connect the other end of the SCSI cable to one of the LVD/SE ports on the rear of the disk system.

FIGURE 8 SCSI Port Locations

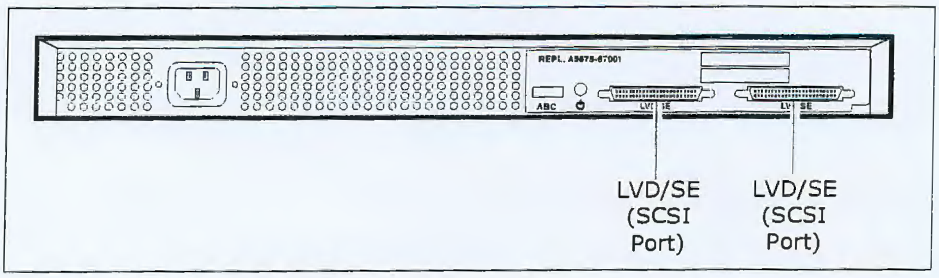

4 Determine if this disk system is at the beginning or end of the SCSI bus.

- If the disk system is at the end of the SCSI bus (you are not going to daisy-chain another device off this one), install the appropriate terminator (part number 5183-2657) on the available SCSI port and go to “Connecting the Power Cable” (see Figure 9 on page 17).
- If the disk system is at the beginning or middle of the SCSI bus (you are going to daisy-chain another device off this one), consider the following:

  No more than 3 of these disk systems can be daisy-chained together.

  Ensure that the SCSI IDs assigned for the disk drives in the second disk system or other peripheral are not already assigned to another device on the SCSI bus. Refer to “Setting the SCSI IDs for the Disk Drives” (see Figure 7 on page 14 for the SCSI IDs assigned for each SCSI switch setting).

  If you connected the SCSI cable to the on-board UltraSCSI port on the host system, verify that the SCSI IDs assigned to the disk drives in the second disk system or additional peripherals are not already assigned to any other peripherals installed in the host system.

  The overall SCSI bus length cannot exceed 12 meters. The SCSI bus length for the disk system internal cables is 1.1 meters. If another type of peripheral is being connected on this SCSI bus, refer to that peripheral’s documentation for its internal SCSI cable length.

5 Connect one end of an ultra SCSI cable (get it from the shipping box) to the available SCSI port on the rear of the first disk system on the SCSI bus.

6 Connect the other end of the ultrascsi cable to the SCSI IN port on the rear of the second disk system on the SCSI bus.

7 Repeat Steps 3, 5, and 6 for the last disk system on the SCSI bus if three disk systems are being installed. Keep in mind the maximum bus length (12 meters).

8 Install a terminator if the device is the last one on the bus. For supported terminators, see Table 10 on page 33 or Table 12 on page 34.

Note The disk system, when connected at the end of a SCSI bus, requires a terminator. Refer to the documentation that came with your wide SCSI device to determine if it needs a terminator or not. Narrow SCSI devices at the end of a daisy-chain always require a terminator.

19385
E

Connecting the Power Cable

CAUTION **Ensure that the connection of multiple units to the supply circuit does not overload the supply overcurrent protection or the supply wiring. Refer to the storage electrical ratings when determining the correct branch circuit rating for your installation. See Table 3 on page 12**

FIGURE 9 AC Power Connector Location

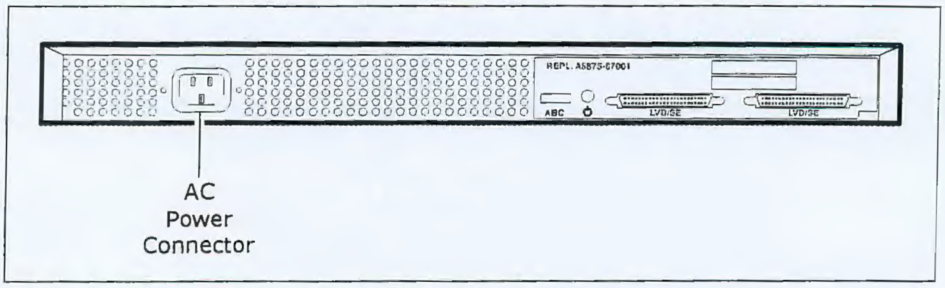

Powering On and Off

Power On the Disk System

1 Press the power button and release it. The power button will stay in a depressed position, indicating that the power is on

See Figure 10 on page 18 for the location of the power button.

2 Confirm that the disk system is running properly by checking the system power LED and the disk module LEDs at the front of your disk system. A system reboot may be necessary to assure that the host system recognizes the disk drives within the disk system.

Refer to "Front Panel" on page 3 for explanations of the LED functions and their meanings.

Power Off the Disk System

1 Back up all data if the disk system is still operational.

2 Ensure that no data on disk modules within the disk system is being accessed.

3 Unmount any file systems associated with the disk system that is going to be powered off. Refer to the system administration manual for your host system's operating system for the correct procedure for unmounting a file system or stopping access to the disk modules within the disk system.

> **Note** If your host system's boot or root disk drive resides in the disk system, the host system must be brought down before the power to the disk system can be turned off. See your host system's system administrator's manual for the correct procedures.

4 Turn the disk system off by pressing the power button. See Figure 10 (below) for the power button location.

5 Confirm that the storage disk system System Power LED is off. See Figure 2 on page 9

6 Disconnect the power cable from the power connector at the rear of the disk system.

FIGURE 10 Power Button Location

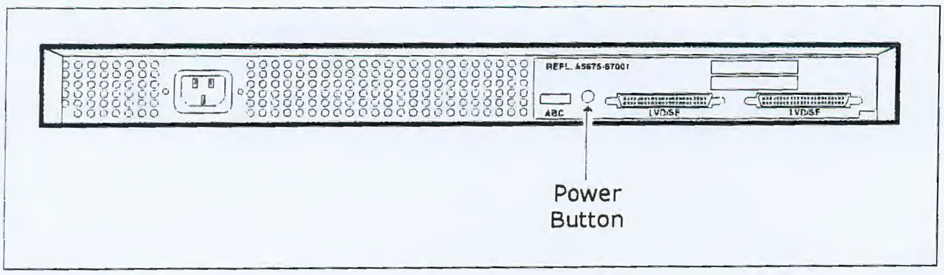

Adding Disk Modules

Disk modules can be added, removed, and replaced while the disk system is running. Because the disk modules can be handled in this way, they are called hot-pluggable.

The SCSI addresses for the disks are set using the addressing switches on the back of the disk system. You can determine the assigned SCSI addresses by looking at the SCSI address switch settings at the rear of the disk system.

If your storage system contains less than 4 disk modules, the remaining empty slots require filler panels. These filler panels ensure that the proper cooling is maintained within the storage system.

The system administration procedure for adding a disk modules is operating system specific. You must decide where the disk module is to be installed and install it. Then the operating system has to be configured to recognize the new disk drive.

Note The disk system supports the 18 and 36 GB disk capacity points. It supports these capacities in 10,000 RPM disk modules only at this time.

Add a Disk Module

1 Select an available slot for the new disk module. Note the slot chosen for application configuration.

2 Remove the disk module filler panel. Store the filler panel for future use.

3 Hold the locking handle open on the disk drive. Push the locking lever to the left to release the latch.

4 Slide the disk into the appropriate slot.

5 Gently push the drive until the locking mechanism engages. When the disk modules is completely installed, an audible click can be heard.

6 Close the locking handle completely, using gentle downward pressure.

FIGURE 11 Installing a disk module

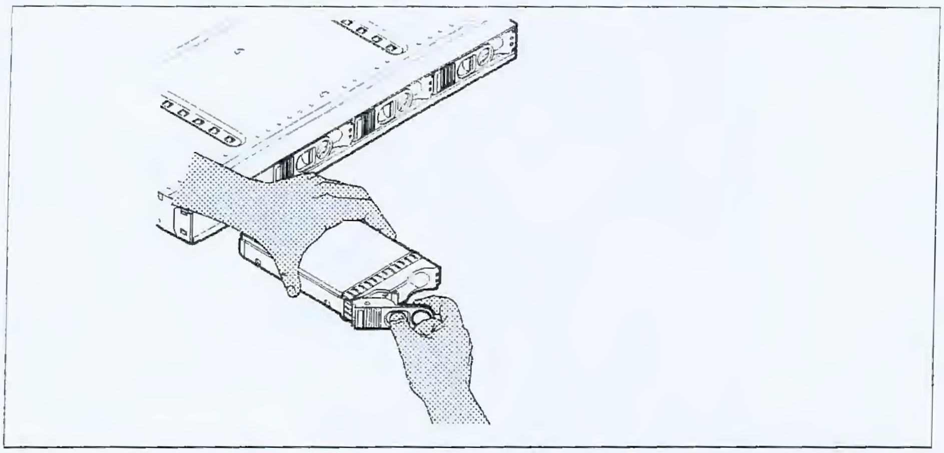

Configure the new Disk Module

Configure a New Disk Module within Your Application

CAUTION **These procedures should be performed only by qualified system administrator. Performing hot-plug operations on an active disk drive can result in data loss or corruption.**

Remove and Replace A Disk Module

Removing a Disk Module

CAUTION **If the disk system is running and a disk module slot does not contain a disk drive or filler panel, the disk system will not cool properly and may overheat. If you are not replacing the disk module immediatelty, a filler panel must be installed to maintain proper cooling.**
Ensure that you have a replacement disk or filler panel before removing the disk module from the disk system unless you are returning a defective disk system.

1 Ensure there is no activity on the drives.

If the Disk Drive LED for that disk module is flashing green, the disk drive is active. When the Disk Drive LED off, the disk module is either idle or completely off (see Table 1 on page 9). In this state, it is safe to remove the disk module.

2 Prepare the software environment to remove the drive.

See your operating system documentation for instructions and procedures required to remove a disk module. It may be necessary to unmount file systems associated with the disk drives installed in the disk system.

3 Unlatch the disk module handle and remove the disk module.

Squeeze the latch to unlock the disk module. Put a finger behind the latch and pull the disk module out until it disengages with the backplane. Leave it in its slot for two to three minutes to allow the disk drive to spin down before removing the disk module from the enclosure

If you are going to replace the disk module immediately, refer to "Replace a Disk Module" below.

If you are not going to replace the disk module immediately insert a filler panel into the open slot to assure the required cooling if you are not returning the disk system.

Replace a Disk Module

Note The disk system supports the 18 and 36 GB disk capacity. It supports these capacities in 10,000 RPM disk modules only at this time.

1 Refer to your operating system documentation for requirements and procedures for replacing disk modules.

2 Align the replacement disk module with the disk module slot.

3 Slide the disk module into the slot until it engages with the backplane.

4 Press on the locking handle until it locks. You will hear an audible click when the locking mechanism is fully engaged.

1938 O
E

Remove the Disk System

CAUTION **Some data paths may be slot-dependent. Be sure to note the slot from which the disk module is removed so it can be installed in the same slot in the replacement disk system.**

Reverse the installation instructions enclosed with the rackmount kit for your cabinet:

- A5679A for Hewlett-Packard Rack Systems/E for Enterprise Systems
- A5680A for all other Hewlett-Packard rack systems for Enterprise Systems purchased before November 1998.
- A6532A for Hewlett-Packard Rack Systems/E for Commercial Systems
- A6533A for all other Hewlett-Packard rack systems for Commerial Systems purchased before November 1998.

If your disk system is a factory-integrated unit, see the figure below for removal instructions.

FIGURE 12 Removing the Disk System

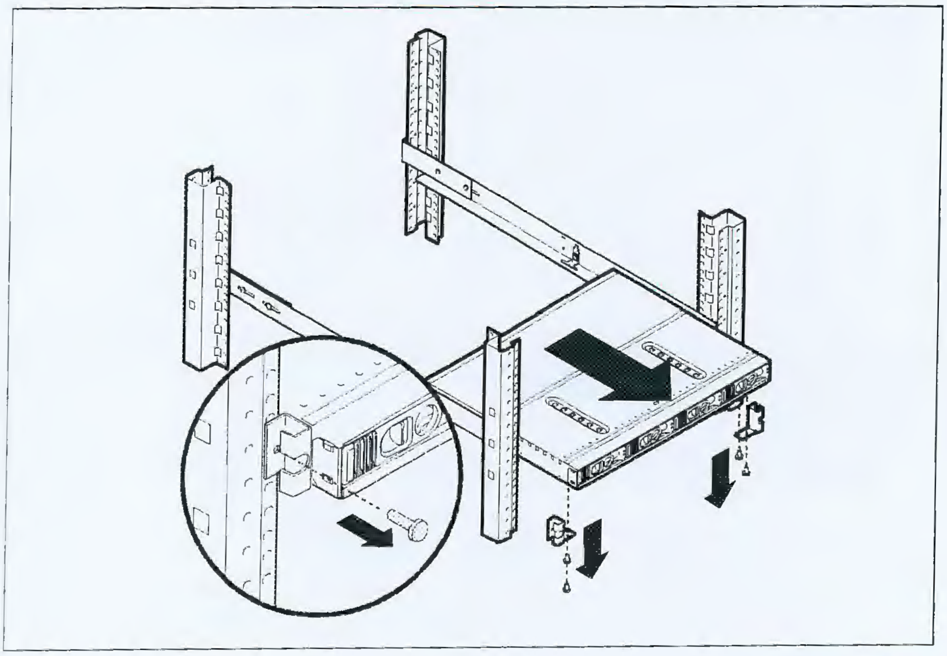

If you are returning a desktop disk system, be sure to remove and store the plastic Desktop Disk System Cover for use on the replacement disk system.

Setting Up the Hardware Event Monitor (HP-UX Only)

Hardware event monitors run on HP-UX hosts, versions 10.20 and 11.0. The Disk Monitor (disk_em) monitors all disks bound to sdisk drivers. Consequently, if the Disk Monitor is active on your host, it is already set up to monitor the disks of a new disk system. If you need to install or activate the Disk Monitor, refer to the *EMS Hardware Monitors User's Guide* in the latest IPR Support Media or on the Web (*http://www.docs.hp.com/hpux/systems/#ems*)..

> **Note** This Disk Monitor should not be confused with the EMS disk monitor that is used to monitor LVM resources.

The way you configure the monitor determines, among other things, where event messages will be sent and what level of severity will be reported.

Configuration Overview

The following steps will help you identify and resolve disk system failures:

1. Gather information from all sources:
   - Hardware event notifications (page 26)
   - Disk system LED status (page 26)
   - Online information tools (page 27)
2. Isolate the cause of the problem (Table 7 on page 31).
3. Correct the problem. (See page 20 for disk modulecremoval and replacement.)
4. Verify operational status with IOSCAN or other host utilities.

19378
E

Event Notification (HP-UX Only)

The Disk Monitor, an EMS hardware event monitor, reports changes in disk status. Depending on how the monitor is set up, it sends messages to the console, an e-mail address, a log file, or a third-party application. These messages give early notice of a disk problem. Events include media errors, failed read and write attempts, invalid commands, changed operating parameters, failed diagnostics, and many others.

Event severity ranges from critical to informational:

Critical An event that causes data loss, host system downtime, or other loss of service. Host system operation will be affected if the disk system continues to be used without correction. Immediate action is required. For example, read data could not be recovered.

Serious An event that may cause data loss, host system downtime, or other loss of service if left uncorrected. Host system and hardware operation may be adversely affected. The problem needs repair as soon as possible. For example, the request queue is full.

Warning An event that could escalate to a serious condition if not corrected. Host system operation should not be affected and normal use of the disk system can continue. Repair is needed but at a convenient time. For example, the bus failed to reset.

Information An event that is expected as part of the normal operation of the hardware. No action is required. For example, write protection was switched on or off.

Event messages (see Figure 13) contain the following types of information:

- Message Data – Date and time the message was sent, the source and destination of the message, and the severity level
- Event Data – Date and time of the event, the host, event ID, name of the monitor, event number, event class, severity level, hardware path, associated OS error log entry ID
- Error Description – Narrative information indicating the component that experienced the event and the nature of the event
- Probable Cause/Recommended Action – The cause of the event and suggested steps toward a solution. This information should be the first step in troubleshooting.

FIGURE 13 Sample Hardware Event Notification

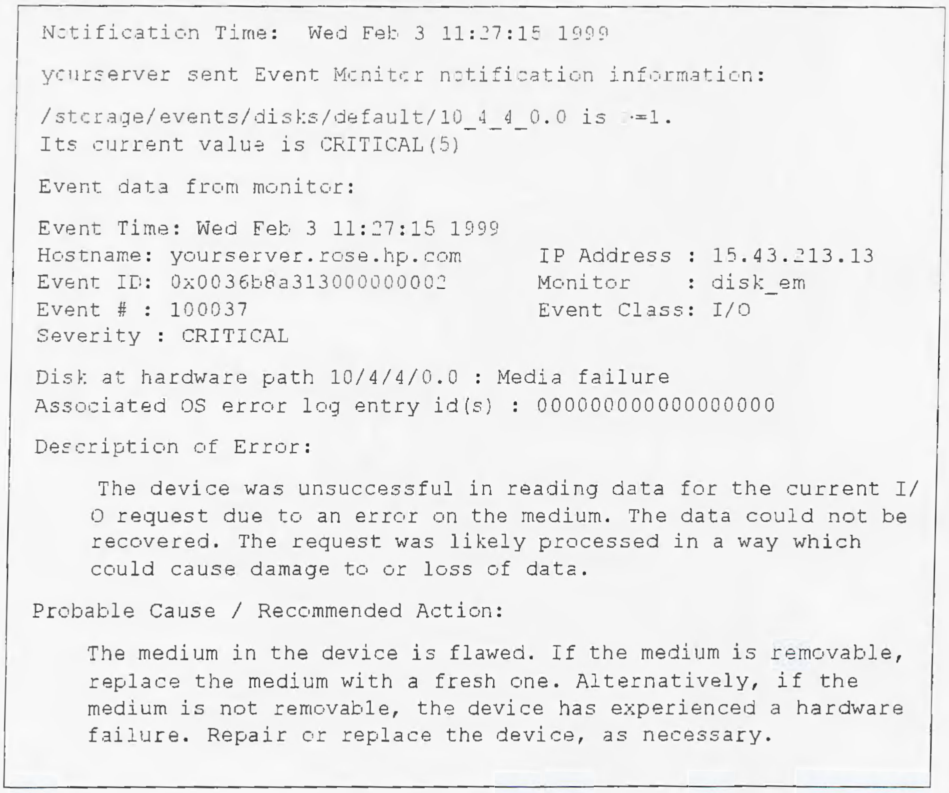

```
Notification Time:  Wed Feb 3 11:27:15 1999

yourserver sent Event Monitor notification information:

/storage/events/disks/default/10_4_4_0.0 is >=1.
Its current value is CRITICAL(5)

Event data from monitor:

Event Time: Wed Feb 3 11:27:15 1999
Hostname: yourserver.rose.hp.com      IP Address : 15.43.213.13
Event ID: 0x0036b8a3130000000002      Monitor    : disk_em
Event # : 100037                      Event Class: I/O
Severity : CRITICAL

Disk at hardware path 10/4/4/0.0 : Media failure
Associated OS error log entry id(s) : 000000000000000000

Description of Error:

     The device was unsuccessful in reading data for the current I/
     O request due to an error on the medium. The data could not be
     recovered. The request was likely processed in a way which
     could cause damage to or loss of data.

Probable Cause / Recommended Action:

     The medium in the device is flawed. If the medium is removable,
     replace the medium with a fresh one. Alternatively, if the
     medium is not removable, the device has experienced a hardware
     failure. Repair or replace the device, as necessary.
```

19376
ε

Status LEDs

LEDs indicate the status of the disk system itself and each of the disk modules. A green system LED is visible on the front of the disk system. It shows that a fault has occurred or that power is on. Disk activity LEDs are on the front of the disk modules (see Figure 14 and Table 6 below).

FIGURE 14 LED Status Indicators

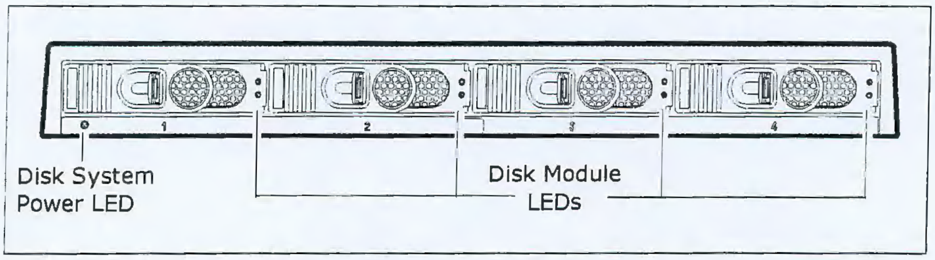

TABLE 6 LED Status Indicators

LED	State	Indication
System Power	Green	Power is on
	Blinking	Malfunction - either a fan is not operating properly or internal voltage is too low.
	OFF	Power is off
Disk Activity	ON	Installed and spinning up. If the LED is still on 3 minutes after power is engaged, the disk may be faulty.
	Flashing	I/O activity on the disk
	OFF	Not installed, not operating, or no I/O activity

View Disk Status

HP-UX and MPE/iX utilities provide descriptive and diagnostic information about disks, including disk type, firmware revision, and errors. On HP-UX and MPE/iX 6.5, the disk utility is Support Tools Manager (STM). For all other operating systems, consult the appropriate system administration manual for disk module status checking procedures.

STM Disk Information: HP-UX

STM displays the last-generated Information Log for a selected disk. Start STM and run the Information tool as follows.

1. Log on the system.
2. At the system prompt, type **xstm&**. STM starts and displays a graphic of the devices on the system.
3. Select the desired disk.
4. Select **Information** from the Tools menu.
5. To generate a current log, select **Run**. The log will be displayed as soon as it is generated
6. To view a log without updating the contents, select **Information Log**.
7. Select **Done** when you have finished viewing the information.
8. To quit STM, type **exit**.
9. Figure 15 shows a sample Information Log.

19374
ε

FIGURE 15 Sample STM Information Log (HP-UX)

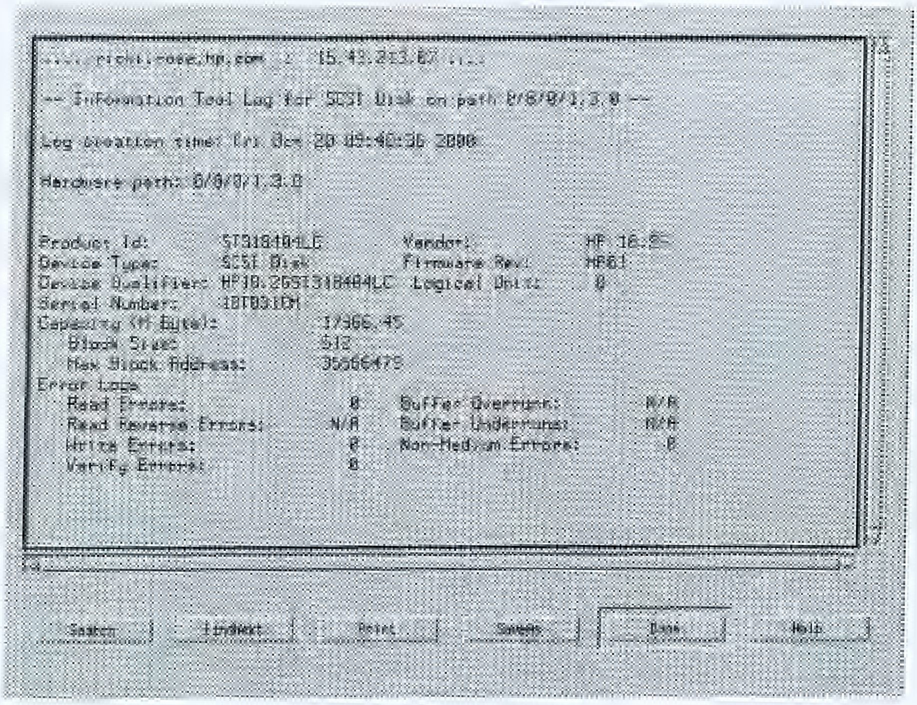

```
.....richt.rose.hp.com  :  15.43.213.87 .....

-- Information Tool Log for SCSI Disk on path 0/0/0/1.3.0 --

Log creation time: Fri Oct 20 09:40:36 2000

Hardware path: 0/0/0/1.3.0


Product Id:        ST318404LC        Vendor:            HP 18.2G
Device Type:       SCSI Disk         Firmware Rev:      HP61
Device Qualifier: HP18.2GST318404LC  Logical Unit:       0
Serial Number:     3BT031EM
Capacity (M Byte):             17366.45
   Block Size:                 512
   Max Block Address:          35566479
Error Logs
   Read Errors:                    0   Buffer Overruns:          N/A
   Read Reverse Errors:          N/A   Buffer Underruns:         N/A
   Write Errors:                   0   Non-Medium Errors:          0
   Verify Errors:                  0
```

1093-73
E

STM Disk Information: MPE/iX 6.5 or 7.0

STM displays the last-generated Information Log for a selected disk. Start STM and run the Information tool as follows.

1 Log on the system.

2 At the system prompt (:), type **vsclose** <physical volume number>. This removes the disk from use.

3 At the system prompt (:), type **cstm**. STM starts.

4 At the cstm prompt, type **map**. STM displays a list of all the disks installed on the system.

5 Select the desired disk by typing **select device** <number>; for example, **select device 15**.

6 Type **information**. STM updates the system map.

7 To display the information log, type **infolog**.

FIGURE 16 Sample STM Expert Tool Disk Error Log (MPE/iX 6.5 or 7.0)

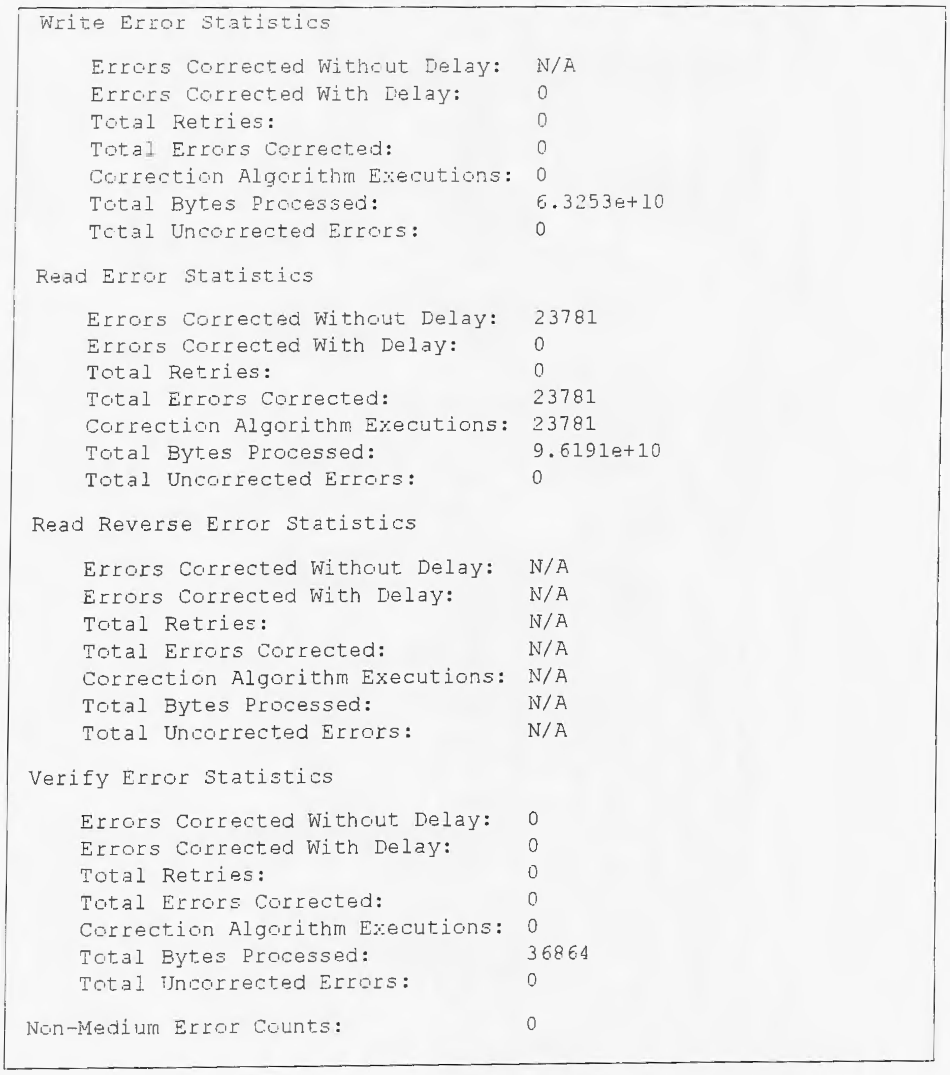

```
Write Error Statistics

     Errors Corrected Without Delay:   N/A
     Errors Corrected With Delay:      0
     Total Retries:                    0
     Total Errors Corrected:           0
     Correction Algorithm Executions:  0
     Total Bytes Processed:            6.3253e+10
     Total Uncorrected Errors:         0

Read Error Statistics

     Errors Corrected Without Delay:   23781
     Errors Corrected With Delay:      0
     Total Retries:                    0
     Total Errors Corrected:           23781
     Correction Algorithm Executions:  23781
     Total Bytes Processed:            9.6191e+10
     Total Uncorrected Errors:         0

Read Reverse Error Statistics

     Errors Corrected Without Delay:   N/A
     Errors Corrected With Delay:      N/A
     Total Retries:                    N/A
     Total Errors Corrected:           N/A
     Correction Algorithm Executions:  N/A
     Total Bytes Processed:            N/A
     Total Uncorrected Errors:         N/A

Verify Error Statistics

     Errors Corrected Without Delay:   0
     Errors Corrected With Delay:      0
     Total Retries:                    0
     Total Errors Corrected:           0
     Correction Algorithm Executions:  0
     Total Bytes Processed:            36864
     Total Uncorrected Errors:         0

Non-Medium Error Counts:               0
```

1123-11
E

Isolating Faults

Table 7 lists the probable causes and solutions for problems you may detect on the disk system. When more than one problem describes your situation, investigate the first solution that applies. The table lists the most basic problems first and excludes them from subsequent problem descriptions.

Erratic LED behavior on the disk system could be observed in the following situations:

- When a server that is connected to a Disk System 2100 is powered down or loses power (and the disk system remains powered on)
- When the disk system is powered up when connected to a downed server
- When an unconnected disk system is powered up

An example of this LED behavior is the LEDs staying lit solidly. This condition is caused by the disk system being deprived of term power when the server loses power or is powered down. The disk system does not provide its own term power. It relies on the host bus adapter to which it is connected for term power.

TABLE 7 Troubleshooting Table

Problem Description	**LED State**	**Probable Cause/Solution**
Disk system fails to power on when installed	System power LED is off	Power cord is not plugged in. The power button is not pressed. AC breaker is tripped or AC power source has failed. The PDU/PDRU is defective. Replace. Enclosure chassis is faulty. Replace.
	System power LED is blinking	Power supply is defective - replace the enclosure Cooling fans are not spinning at the correct speed - replace the enclosure.
Operating system reports errors on a device	Disk module LED is on solid or off	Use diagnostic utilities to determine disk status. Depending on the results, monitor or replace disk module.
IOSCAN (HP-UX) lists disk as NO_HW, or Mapper or DSTAT ALL (MPE/iX) lists no device type	Disk module LED is on solid or off	Disk module is faulty - Replace. Enclosure is faulty - Replace. If the all disks on the bus have this problem, the cable is faulty. Replace the cable.
Unable to configure device for use by operating system	Disk module LED is on solid or off	Disk module possibly not correctly seated. If at a new installation, remove and reinstall the disk module. If troubleshooting an existing installation, run diagnostics to determine drive status. Replace the disk module.

193-70
E

Reference Information

Product Numbers and Options

Three models of the disk system are available:

- A5675A field-racked disk system (Enterprise Customer Installable)
- A5675AD desktop disk system (Enterprise Systems)
- A5675AZ factory racked disk system (Enterprise Systems)
- A5676A field-racked disk system (Commerial Customer Installable)
- A5676AD desktop disk system (Commerial Systems)

TABLE 8 Products and Options

Disk System Product Number	Disk Module Product (with Option)	Description
A5675A		Field installed disk system (Enterprise)
	A6537A (0D1)	Drive integrated into the field installed disk system
	A6538A (0D1)	Drive integrated in the same field installed disk system
A5675AD		Desktop disk system (Enterprise)
	A6537A (0D1)	Drive integrated into the factory installed disk system
	A6538A (0D1)	Drive integrated in the same factory installed disk system
A5675AZ		Factory installed disk system (Enterprise)
	A6537A (0D1)	Drive integrated into the desktop disk system
	A6538A (0D1)	Drive integrated in the same desktop disk system
A5676A		Field installed disk system (Commercial)
	A6544A (0D1)	Drive integrated into the field installed disk system
	A6545A (0D1)	Drive integrated in the same field installed disk system
A5676AD		Desktop disk system (Commercial)
	A6544A (0D1)	Drive integrated into the factory installed disk system
	A6545A (0D1)	Drive integrated in the same factory installed disk system

19349
E

Upgrade Disk Products Available

TABLE 9 Available Upgrade Disk Modules

Product Number	Description
A6537A	18-Gbyte 10,000 rpm LVD disk module (for Enterprise Systems)
A6538A	36-Gbyte 10,000 rpm LVD disk module (for Enterprise Systems)
A6544A	18-Gbyte 10,000 rpm LVD disk module (for Commerial Systems)
A6545A	36-Gbyte 10,000 rpm LVD disk module (for Commerial Systems)

TABLE 10 Upgrade Products for Hewlett-Packard Systems

Product Number	Description
A4999A	Single Port Ultra 2 SCSI Host Bus Adapter for Hewlett-Packard Workstations
A5149A	Single Port Ultra 2 SCSI Host Bus Adapter for HP-UX Systems
A5150A	Dual Port Ultra 2 SCSI Host Bus Adapter for HP-UX Systems
A5856A	Internal Disk Array Controller for HP-UX Systems
A5675A	HP Surestore Disk System 2100
A5679A	Rail kit for the Rittal Style Rack Systems and HP Rack Systems/E
A5680A	Rail kit for HP Rack Systems
A6198A	Disk Filler Panel
C2364A	LVD/SE SCSI Terminator (HDTS68)
C7430A	Ultra 2 SCSI Host Bus Adapter for HP Netservers
D2140A	Netraid 1si Host Bus Adapter for HP Netservers
D5025A	Ultrawide Host Bus Adapter for HP Netserver
D5955A	Netraid 3si Host Bus Adapter for HP Netservers
D9161A	Netraid 4M/64 Host Bus Adapter for HP Netservers
D9351A	Netraid 4M/128 Host Bus Adapter for HP Netservers
P3413A	Single Port Ultra 3 SCSI Host Bus Adapter for HP Netservers

19308
E

TABLE 11 Replaceable Parts

Product Number	Description	Replaceable Part Numbers
A5675A	Base Disk System Assembly	A5675-67001
A6198A	Disk Filler Panel	A6198-60001
A6537A	18-Gbyte 10,000 rpm LVD disk module	A6537-69001
A6538A	36-Gbyte 10,000 rpm LVD disk module	A6538-69001
A6544A	18-Gbyte 10,000 rpm LVD disk module	A6544-69001
A6545A	36-Gbyte 10,000 rpm LVD disk module	A6545-69001
	Desktop Disk System Cover	5065-5217
	Desktop Disk System Feet	0403-0285

TABLE 12 Supported HP Cables and Terminators (for Enterprise Systems)

Order Number	Description	Part Number
C2978B	0.5-meter HDT S68 SCSI multimode cable	5183-2670
C2911C	1.0-meter HDT S68 SCSI multimode cable	5183-2671
C2979B	1.5-meter HDT S68 SCSI multimode cable	5183-2672
C2924C	2.5-meter HDT S68 SCSI multimode cable	5183-2673
C7521A	5.0-meter HDT S68 SCSI multimode cable	5183-2678
C2361B	1.0-meter VHDT S68/HDT S68 SCSI multimode cable	5183-2674
C2362B	2.5-meter VHDT S68/HDT S68 SCSI multimode cable	5183-2675
C2365B	5.0-meter VHDT S68/HDT S68 SCSI multimode cable	5183-2676
C2364A	SCSI Terminator LVD/SE HDTS68	5183-2657

TABLE 13 Supported HP Cables and Terminators (for Commercial Systems)

Order Number	Description	Part Number
A6523A	0.5-meter HDT S68 SCSI multimode cable	5183-2670
A6524A	1.0-meter HDT S68 SCSI multimode cable	5183-2671
A6525A	1.5-meter HDT S68 SCSI multimode cable	5183-2672
A6526A	2.5-meter HDT S68 SCSI multimode cable	5183-2673
A6527A	5.0-meter HDT S68 SCSI multimode cable	5183-2678
A6528A	1.0-meter VHDT S68/HDT S68 SCSI multimode cable	5183-2674
A6529A	2.5-meter VHDT S68/HDT S68 SCSI multimode cable	5183-2675
A6530A	5.0-meter VHDT S68/HDT S68 SCSI multimode cable	5183-2676
A6531A	SCSI Terminator LVD/SE HDTS68	5183-2657

Specifications

TABLE 14 Physical Dimensions

Measure	Metric	English
Width	45.08 cm	17.75 in
Depth	38.10 cm	15.0 in
Height	4.32 cm	1.7 in
Weight without disk modules	4.94 kg	10.90 lbs
Weight fully loaded	8.11 kg	17.89 lbs

Electrical Specifications

TABLE 15 AC Power Requirements

Electrical Element	Requirement
AC Operating Voltage Range	88-269 VACRMs
Line Frequency Range	47-63 Hz
Published Product Operating Range	100 - 240 VACRMs
Published Product Frequency Range	50 - 60 HZ

19366
ε

TABLE 16 Environmental Specifications

Climate Control		
Temperature (dry bulb)	Operating	0°C - +40°C
	Storage	-40°C - +70°C
Relative Humidity (non-condensing)	Operating	5% - 95% Relative Humidity at 40°C
	Storage	90% Relative Humidity at 65°C
Altitude (based on disks)	Operating	-1000 ft to +10,000 ft (3048 M)
	Storage	40,000 ft (12,092 M)
Heat Dissipation (maximum)	Operating	< 100 watts

Operating Temperatures

If the storage system is installed in a multi-unit rack assembly, the operating ambient temperature of the rack environment may exceed room ambient temperature. The rack environment ambient temperature cannot exceed 40° Celsius (104° Fahrenheit).

If your storage system contains less than 4 disk modules, the remaining empty slots require filler panels. These filler panels (part number A6198-60001) ensure that the proper cooling is maintained within the storage system.

Regulatory Statements

Safety Certifications

UL listed, UL 1950:1995 – 3rd Edition

CSA certified, C22.2 No. 950:1989

TUV certified with GS mark, EN 60950:1992 + A1:1993, A2:1993, A3:1995, A4:1997

CE mark (see on page 39)

EMC Compliance

Australia: AS/NZS 3548, Class B

Canada: ICES-003, Class B

China: GB9254-88

European Union: EN55022 Class B, EN55024

19365
E

Japan: VCCI Class B

Taiwan: CNS 13438, Class A

US: 47 CFR Parts 2 & 15, Class B

A. FCC Notice for United States

The Federal Communications Commission (in 47 CFR 15.105) has specified that the following notice be brought to the attention of the users of this product.

This equipment has been tested and found to comply with the limits for a Class B digital device, pursuant to Part 15 of the FCC Rules. These limits are designed to provide reasonable protection against harmful interference when the equipment is operated in a residential installation. This equipment generates, uses, and can radiate radio frequency energy and, if not installed and used in accordance with the instruction manual, may cause harmful interference to radio communications. However, there is no guarantee that interference will not occur in a particular installation. If this equipment does cause harmful interference to radio or television reception, which can be determined by turning the equipment off and on, the user is encouraged to try to correct the interference by one or more of the following measures:

- Reorient or relocate the receiving antennna.
- Increase the separation between the equipment and the receiver.
- Connect the equipment into an outlet on a circuit different from that to which the receiver is connected.
- Consult the dealer or an experienced radio/television technician for help.

Hewlett-Packard's cerification tests were conducted with a Hewlett-Packard supported computer system and Hewlett-Packard shielded cables, such as those you received with your storage product. Changes of modificaitons not expressly approved by Hewlett_packard could void the user's authority to operate the equipment. cables used with this device must be properly sheilded to comply with the requirements of the FCC.

B. FCC Notice for Canada

This Class B digital apparatus complies with Canadian ICES-003.

Cet appareil numérique de la classe B est conforme à la norme NMB-003 du Canada.

19364
E

C. Notice for Japan

この装置は、情報処理装置等電波障害自主規制協議会（ＶＣＣＩ）の基準に基づくクラスＢ情報技術装置です。この装置は、家庭環境で使用することを目的としていますが、この装置がラジオやテレビジョン受信機に近接して使用されると受信障害を引き起こすことがあります。
取り扱い説明書に従って正しい取り扱いをして下さい。

Harmonics Conformance (Japan)

高調波ガイドライン適合品

D. Notice for Taiwan

警告使用者：這是甲類的資訊產品，在居住的環境中使用時，可能會造成射頻干擾，在這種情況下，使用者會被要求採取某些適當的對策。

E. Notice for Korea

사용자 안내문 (B급 기기)

이 기기는 비업무용으로 전자파장해검정을 받은기기로서, 주거지역에서는 물론 모든 지역에서 사용 할 수 있읍니다.

F. Notice for Germany

Schalldruckpegel Lp = 55.0 dB(A)

Am Arbeitsplatz (operator position)

Normaler Betrieb (normal operation)

Nach ISO 7779:1999 (Typprüfung)

19363

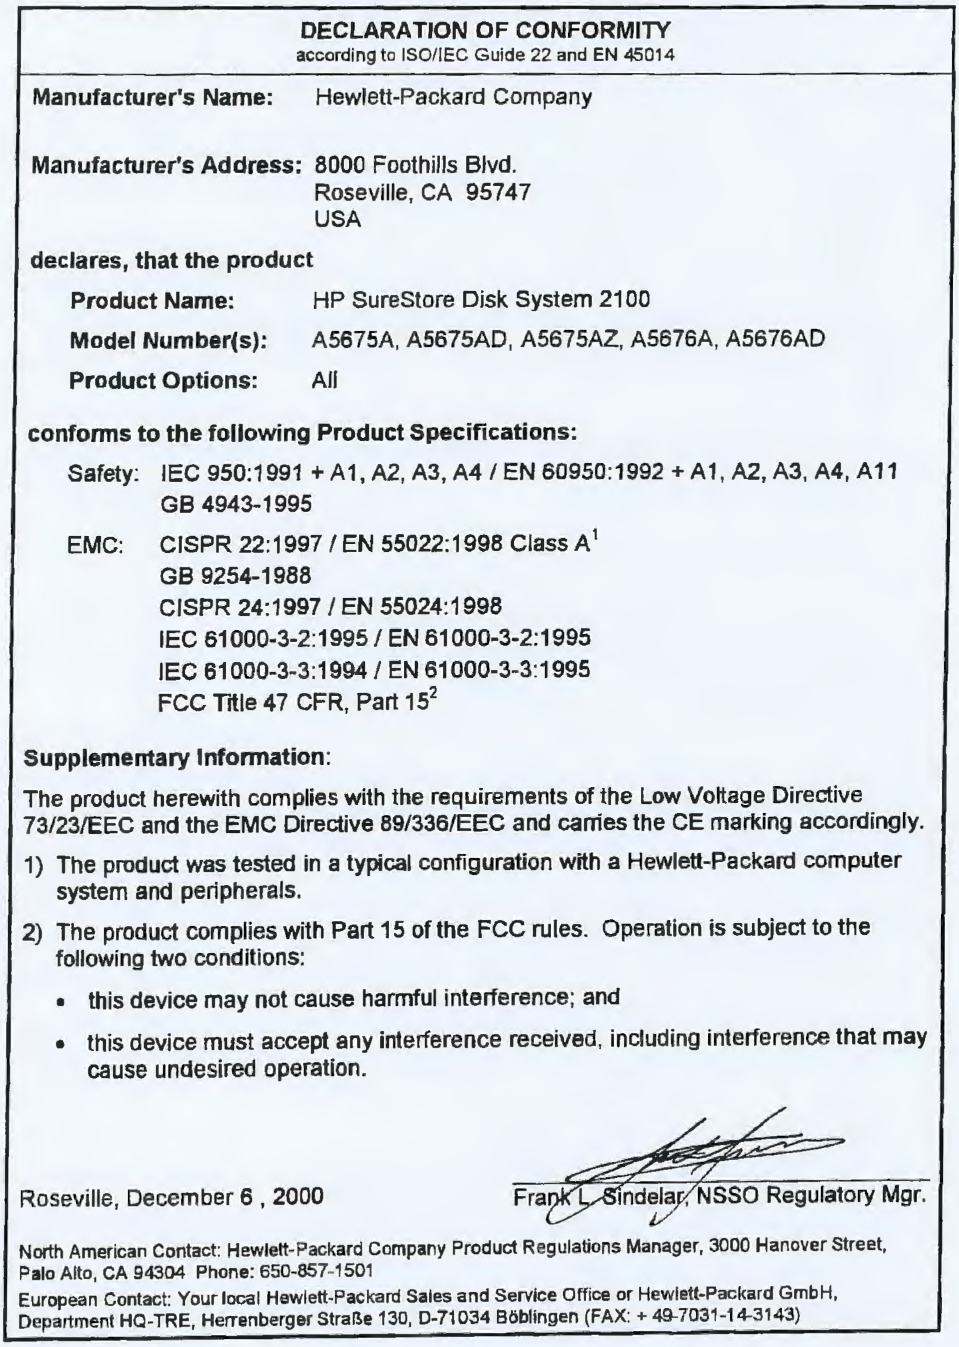

DECLARATION OF CONFORMITY

according to ISO/IEC Guide 22 and EN 45014

Manufacturer's Name: Hewlett-Packard Company

Manufacturer's Address: 8000 Foothills Blvd.
Roseville, CA 95747
USA

declares, that the product

Product Name: HP SureStore Disk System 2100

Model Number(s): A5675A, A5675AD, A5675AZ, A5676A, A5676AD

Product Options: All

conforms to the following Product Specifications:

Safety: IEC 950:1991 + A1, A2, A3, A4 / EN 60950:1992 + A1, A2, A3, A4, A11
GB 4943-1995

EMC: CISPR 22:1997 / EN 55022:1998 Class A[1]
GB 9254-1988
CISPR 24:1997 / EN 55024:1998
IEC 61000-3-2:1995 / EN 61000-3-2:1995
IEC 61000-3-3:1994 / EN 61000-3-3:1995
FCC Title 47 CFR, Part 15[2]

Supplementary Information:

The product herewith complies with the requirements of the Low Voltage Directive 73/23/EEC and the EMC Directive 89/336/EEC and carries the CE marking accordingly.

1) The product was tested in a typical configuration with a Hewlett-Packard computer system and peripherals.

2) The product complies with Part 15 of the FCC rules. Operation is subject to the following two conditions:

- this device may not cause harmful interference; and
- this device must accept any interference received, including interference that may cause undesired operation.

Roseville, December 6 , 2000

Frank L. Sindelar, NSSO Regulatory Mgr.

North American Contact: Hewlett-Packard Company Product Regulations Manager, 3000 Hanover Street, Palo Alto, CA 94304 Phone: 650-857-1501

European Contact: Your local Hewlett-Packard Sales and Service Office or Hewlett-Packard GmbH, Department HQ-TRE, Herrenberger Straße 130, D-71034 Böblingen (FAX: + 49-7031-14-3143)

19362
E

Product Web Site

This guide is available in Adobe® Acrobat® format on the HP Customer Care web site for enterprise storage (*http://www.hp.com/essd/*). The complete URL is *www.hp.com/support/ds2100*

Related Documents

The following manuals explain how to use the system software interfaces to the HP Surestore Disk System 2100:

- *Online Diagnostics (for HP 9000): Support Tools Manager Overview*, available at *http://www.docs.hp.com/hpux/systems/*
- *HP-UX System Administration Tasks Manual*, HP Order No. B2355-90079
- Window's NT installation guide for HP Netservers: (NT 4.0) *http://netserver.hp.com/netserver/docs/download.asp?file=it_nt_5_5_00.pdf*
- Window's NT installation guide for HP Netservers: (NT 2000) *http://netserver.hp.com/netserver/docs/download.asp?file=it_win2k_5_1_00.pdf*
- Window's NT installation guide for Legacy HP Netservers: (NT 4.0) *http://netserver.hp.com/netserver/docs/download.asp?file=ig_mwnt40.pdf*

1136.1
E

Reader Comment Sheet

Hewlett-Packard SureStore Disk System 2100 User Guide

We welcome your evaluation of this manual. Your comments and suggestions will help us improve our publications. Remove this page and mail or FAX it to 916-785-2299. Use and attach additional pages if necessary.

	Agree				Disagree	N/A
The manual is well organized.	○	○	○	○	○	○
The information is technically accurate.	○	○	○	○	○	○
Information is easy to find.	○	○	○	○	○	○
Step-by-step procedures are easy to perform.	○	○	○	○	○	○
There are enough examples and pictures.	○	○	○	○	○	○
The examples and pictures are useful.	○	○	○	○	○	○

Comments ____

____

____

____

Name: ____________________ Phone: ____________________

Title: ____________________ FAX: ____________________

Company: ____________________ E-mail: ____________________

Address: ____________________ ZIP: ____________________

City & State: ____________________ Country: ____________________

___ Check here if you would like a reply.

Hewlett-Packard has the right to use submitted suggestions without obligation, with all such ideas becoming the property of Hewlett-Packard.

HPSO/MSO Roseville Information Engineering A5675-90901 E0301

NO POSTAGE
NECESSARY
IF MAILED
IN THE
UNITED STATES

BUSINESS REPLY MAIL

FIRST CLASS MAIL PERMIT NO. 256 ROSEVILLE, CA

POSTAGE WILL BE PAID BY ADDRESSEE

Attention: Information Engineering (MS5668)

Hewlett-Packard Company
HP Storage Organization/ MSO
8000 Foothills Boulevard
Roseville, CA 95747-9987

b9574799875b

Fold Here

Please Do Not Staple

Tape

19359
ε

ANEXO UNIDADE DE BACKUP ROBOTIZADO
PARTE U

hp-ux 11i
oes

hp-ux 11i operating environments
enterprise release

10358
3

hp-ux 11i
oes

hp-ux 11i operating environments benefits

greatly simplified software deployment

- Only one reboot needed to install the Operating Environment (OE) of your choice
- No codewords are necessary to access any of the functionality/application products resident on the OE media
- Comprehensive offering of Network, Mass Storage, and I/O Drivers available during install process
- Online Diagnostics loaded during cold install

simple to purchase license

- Each OE license product contains licensing for the base HP-UX O/S and all of the included HP applications

attractive pricing

- Pricing of the OE licenses reflects a built-in advantage over purchasing individual OE components separately

published testing results

- Testing results of application products in the OEs will be published on docs.hp.com for worldwide access both inside and outside HP

simple to purchase software support

- Simplification in Software Support ordering and contract administration has been achieved in parallel with the introduction of HP-UX 11i Operating Environments
- For more information, please visit: http://nternet.fc.hp.com/catscore/communic.htm

10357

hp-ux 11i
oes

getting to know hp-ux operating environments

- The following series of slides will show how HP-UX 11i Operating Environments are constructed from pieces of the total 11i Software Solution

19356
3

hp-ux 11i
oes

getting to know hp-ux operating environments

design overview

- Each of the Commercial Server Operating Environment license and media products are designed to be supersets of one another
- The TCOE represents a singular solution for Technical Servers and Workstations
- Base HP-UX and Application content common across all four OEs is synchronized with the same revision level

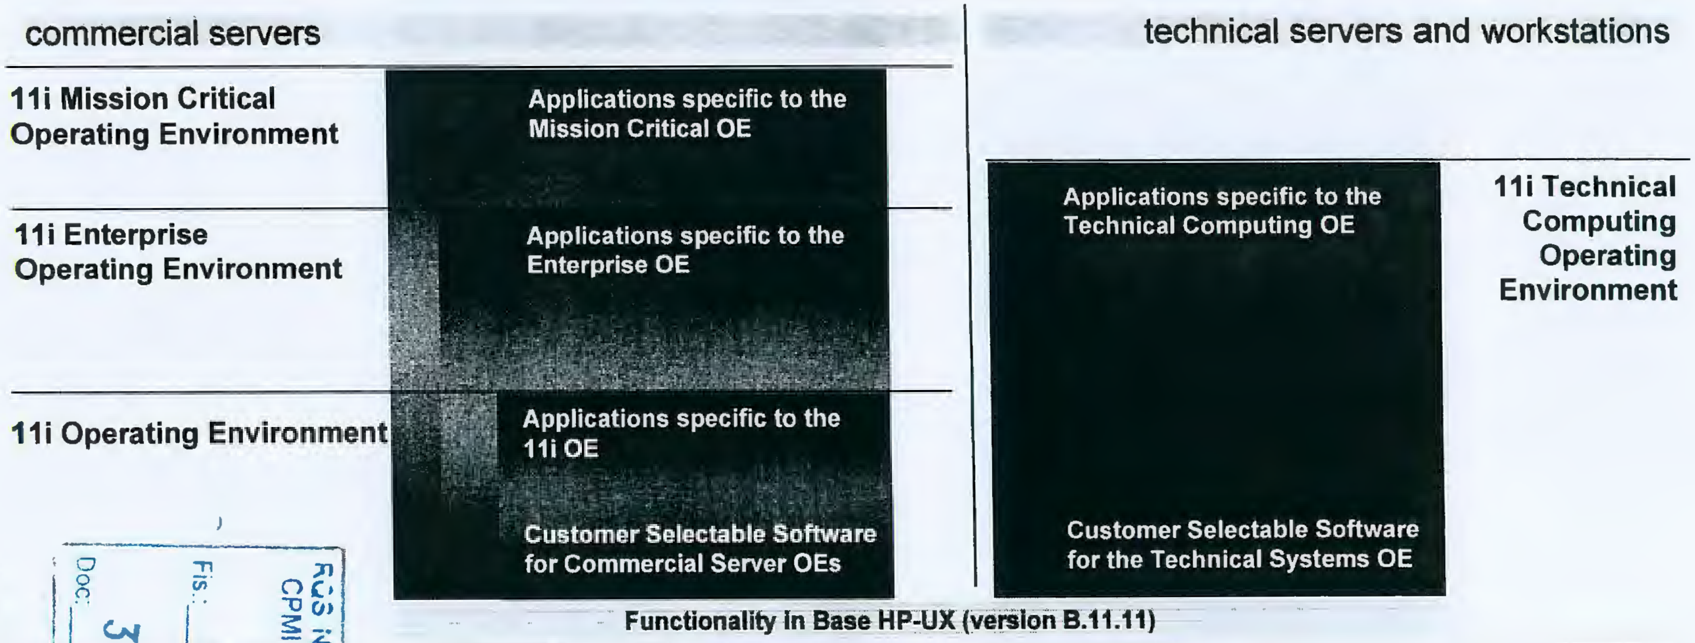
commercial servers
technical servers and workstations
11i Mission Critical Operating Environment
Applications specific to the Mission Critical OE
11i Enterprise Operating Environment
Applications specific to the Enterprise OE
11i Operating Environment
Applications specific to the 11i OE
Customer Selectable Software for Commercial Server OEs
Applications specific to the Technical Computing OE
11i Technical Computing Operating Environment
Customer Selectable Software for the Technical Systems OE
Functionality In Base HP-UX (version B.11.11)

19355

hp-ux 11i
oes

getting to know hp-ux operating environments

- HP-UX version B.11.11 is at the heart of each HP-UX 11i Operating Environment and provides the sound foundation onto which each OE Solution is built
- Two global base HP-UX bundles are delivered at B.11.11 and are differentiated by bitness: HPUXBase64 for 64-bit capable hardware, and HPUXBase32 for 32-bit capable hardware

HP-UX 11i Core Functionality
HPUXBase64 (64-bit)
HPUXBase32 (32-bit)

19354
3

11i (B.11.11) operating system features

major new features of hp-ux 11i (version B.11.11 and future)

- Built-in Unlimited Simultaneous User Level License for HP-UX
- Includes Linux APIs
- Introduces online replacement and addition of I/O cards for N-Class and L-Class servers
- Fully supports the low-end A500 and A400 servers, designed for ISP operation, where performance per rack and serviceability are keys
- Supports L2000 and L1000, an economical N-Class server with an upgrade path to IA-64
- Supports the V2500/V2600 platform
- Offers VERITAS (JFS 3.3) file system support
- Has improved and expanded file system support from CacheFS
- Supports NFS over TCP/IP
- Adds systems management improvements (PRM enhancements)
- Offers secure defaults and intrusion detection
- Has extensive performance tuning for one-way to 32-way configurations

1935

more information on the www

- HP-UX 11i Quick Reference Card
  http://esp.cup.hp.com:2000/nav24/ppos/358/hPUX/11iQRC.pdf

- 11i Technical Overview paper on ESP
  http://esp.cup.hp.com:2000/nav24/ppos/358/hPUX/technOverv/techOV.doc

LTO
LINEAR TAPE OPEN

19351
E

ULT

About The Technology: FAQs

Overview
Technology Primer
FAQs
White Papers
Public Documents

About LTO Technology

Q1: What is LTO Technology?
A1: LTO Technology, or Linear Tape-Open Technology, combines the advanta of linear multi-channel bi-directional tape formats in common usage today wi enhancements in the areas of timing based servo, hardware data compressio optimized track layouts and high efficiency error correction code to maximize capacity and performance.

Q2: What did Certance (formerly Seagate Removable Storage Solutio Hewlett-Packard, and IBM contribute to LTO Technology
A2: Certance, HP and IBM each provided expert knowledge of customer need and complementary technology that provided for delivery of a best-of-breed technology and promotes a strong foundation for data interchange. It is their intent that other companies can participate in this tape industry opportunity through the open licensing process.

Q3: When will products based on LTO Technology be available, and hc will they be priced?
A3: Ultrium Format Generation 1 and Generation 2 licenses are now available Licenses became available on April 7, 1998 for Generation 1 implementation (the LTO Technology. Tape storage products based on the Ultrium format Generation 1 began to appear in the market in late 2000. Ultrium format Generation 2 licenses became available on April 9, 2002. Typically, products based on new storage technology specifications begin to emerge 12 to 18 months after technology licenses become available. For specific pricing questions, please contact the individual licensing companies, and be sure to v the LTO Technology Website at www.ultrium.com for updated information on licensing companies from whom specific product information can be expected

About The Ultrium Format

Q4: What is the Ultrium format?
A4: The Ultrium tape format specification is the implementation of LTO (Linea Tape-Open) Technology optimized for high capacity and performance with outstanding reliability, in either a stand-alone or an automated environment. Using a single-reel tape cartridge to maximize capacity, the Ultrium tape forr is ideally suited for backup, restore and archive applications. It is our intent tl the Ultrium tape format will meet the needs of the enterprise through single server user on a roadmap, or migration path, which extends well into the nex decade. The Ultrium tape format establishes a new benchmark for large volur backup and archive.

Q5: Are products based on the Ultrium format compatible with any ot storage products that exist today?
A5: The new Ultrium tape format specification has been designed to incorpor the best-of-breed strengths of several existing technologies. Compatibility wit existing formats was not a primary consideration but is not prohibited by the

19350

Ultrium tape format specification.

Q6: How can users be sure of data interchangeability among Ultrium format products?

A6: The Ultrium format trademark identifies products that are subject to annu 3rd party compliance verification sample testing to verify that these products comply with the Ultrium tape format specification. This compliance will maxim the likelihood for data interchange of Ultrium tape cartridges between differer manufacturers' Ultrium tape mechanisms. Use of this trademark will be grant only after an Ultrium tape format product has passed the compliance verificat testing process.

Q7-A: When did products based on the Ultrium format Generation 1 become available, and how are they priced?

A7-A: Ultrium format Generation 1 licenses became available April 7, 1998. T storage products based on the Ultrium format Generation 1 began to appear i the market in late 2000. Product specific questions should be directed to individual licensing companies. Please revisit the LTO Technology Website at www.ultrium.com for updated information on licensing companies from whom specific product announcements can be expected.

Q7-B: When will products based on the Ultrium format Generation 2 b available, and how will they be priced?

A7-B: Ultrium format Generation 2 licenses became available April 9, 2002. Typically, products based on new storage technology specifications begin to emerge 12 to 18 months after technology licenses become available. Plans fo Ultrium format Generation 2 tape storage products by individual licensing companies are scheduled for introduction in 2003. Please revisit the LTO Technology Website a www.ultrium.com for updated information on licensing companies from whom specific product announcements can be expected.

Q8: What are the key feature differences between Ultrium format Generation 1 and Generation 2?

A8: As stated in the LTO Technology four-generation roadmap, speed and capacity is expected to double with each generation of the Ultrium format. Th Ultrium format Generation 1 allows for up to 100GB capacity (uncompressed) a single cartridge, while Ultrium format Generation 2 doubles the capacity up 200GB (uncompressed) using Generation 2 data cartridges. The data transfer peak performance speed doubles from up to 20MB/second native in the Ultriu format Generation 1 up to 40MB/second native in Generation 2.

Users will benefit from a greater than 30 percent increase in the number of tracks and average tape speed, as well as an improved recording method. Th Ultrium format Generation 2 provides users with outstanding data and investment protection because of its capability to read-and-write Generation : cartridges in the Ultrium Generation 1 format.

Q9: What are the backward compatibility characteristics of the Ultriur format Generation 2?

A9: The LTO Ultrium compatibility is defined with two concepts demonstrating investment protection:

1) An Ultrium drive is expected to read data from a cartridge in its own generation and at least the two prior generations.

2) An Ultrium drive is expected to write data to a cartridge in its own generation and to a cartridge from the immediate prior generation in the prior generation format.

For example:

- An Ultrium format Generation 2 drive will read and write data on an Ultrium

format Generation 1 cartridge as specified by the Generation 1 format.
- An Ultrium format Generation 2 drive will read and write data on an Ultrium format Generation 2 cartridge as specified by the Generation 2 format.

Q10: Is LTO on track with its roadmap with the introduction of Generation 2?
A10: Yes. In fact, Generation 2 licenses are now available and on schedule wi the roadmap. LTO has followed its four-generation roadmap set forth in 1997 and has not changed it. IT managers and storage administrators can trust tha LTO will hold true to its promises as we have so far, and by adhering closely t our roadmap, ensure that storage solution investments are future-proofed.

Licensing

Q11: How can manufacturers who wish to develop products based on the Ultrium tape format specification obtain a license?
A11: Specifications and license packages are available for both Generation 1 a 2 licenses. Full details can be obtained by contacting Ladas & Parry at (323) 9 2300, or by e-mail at LTO_info@ladasparry.com.

Q12: How can I best determine if I should buy a license and which on to buy?
A12: Deciding whether to buy a license and which one is an individual matter the Licensing Information available from Ladas & Parry -- Phone: (323) 934-2300 or FAX: (323) 934-0202 -- or this Web site (http://www.lto-technology.com/) does not provide you with enough information about the format specifications, you may want to consider purchasing a Basic Package, Ultrium Specification Document (AP-6). Each of these licenses provides the ta mechanism and the tape cartridge specifications for the applicable format. Th Basic Packages are available for a nominal fee. If you are interested in either these packages, please contact Ladas & Parry for licensing at LTO_info@ladasparry.com.

Q13-A: I have purchased the Ultrium format Generation 1 Basic Package, Ultrium Specification Document (AP-6), and have decided to purchase one of the Enhanced Packages, Ultrium Tape Mechanism License (AP-4) or Ultrium Tape Cartridge License (AP-5). Do I get cre for the license fee paid for the Basic Package against the fee for the Enhanced Package license?
A13-A: No.

Q13-B: I have purchased the Ultrium format Generation 2 Basic Package, Ultrium Specification Document (AP-6), and have decided to purchase one of the Enhanced Packages, Ultrium Tape Mechanism License (AP-4) or Ultrium Tape Cartridge License (AP-5). Do I get cre for the license fee paid for the Basic Package against the fee for the Enhanced Package license?
A13-B: No

Q14: Does the LTO Program still support the Accelis format? If not, w
A14: No. Accelis was developed in 1997 to meet customer demands for fast access to data by using a two-reel cartridge that loads in the middle of the ta to minimize access time. However, because of technological developments the have occurred since it was first developed, coupled with the fact that custome have demonstrated strong satisfaction with the performance and growth of th LTO Program's Ultrium open tape format, there is no longer a demand for Accelis. Thus, no license will be offered.

About Open Tape Format Specifications

8
1934

Q15: Why did Certance, Hewlett-Packard and IBM develop new open tape format specifications?
A15: Customers, OEMs and industry analysts believe that the fragmentation i the tape industry and proliferation of formats and technologies has overly complicated customer-buying decisions. With these issues in mind, Certance, and IBM have taken steps to define a technology that has led to open tape format specifications that can better serve multiple market needs and be supported by multiple suppliers.

Q16: What is an "open tape format specification?"
A16: An open tape format specification is one that is readily available to all potential manufacturers and offers licenses for the IP (Intellectual Property) a reasonable price. We feel that the Ultrium tape format specification meets the criteria.

Q17: Why is tape so popular for data protection today? Considering n technologies like DVD-RAM, OAW (optically-assisted Winchester) and near-field recording, how long will tape be a viable and desirable solution?
A17: Tape remains unrivaled in terms of cost and capacity for data storage, a should play an increasing crucial role in corporate data protection strategies. I other technology offers the same combined low cost and high capacity advantage of tape. While other technologies may offer strengths in one or mc areas, overall, they do not meet the entire set of customer needs that tape addresses.

Q18: How does an open tape format specification ultimately benefit customers?
A18: The end-user customer and OEMs benefit from this open tape format specification approach through the availability of multiple sources of product, multiple sources of tape cartridges and common tape format specifications foi interchange. The pro-competitive environment fostered by multiple sources o product leads to faster paced innovation and enhancements, generally means lower prices and provides for data compatibility. For OEMs, open tape format specifications shorten their qualification cycles and reduce complexities in system configuration.

1934 7
E

ANEXO CONCORD SYSTEMEDGE (PARTE 1)

COMPROVAÇÃO DAS ESPECIFICAÇÕES EXIGIDAS NO EDITAL

2.1. ASPECTOS GERAIS

REQUISITO	ATRIBUTOS OFERTADOS	ATRIBUTOS OFERTADOS ADICIONALMENTE	CONFIRMA ATENDIMENTO (SIM / NÃO)	PÁGINA DA DOCUMENTAÇÃO TÉCNICA
2.1 – Comprovação de Performance para os Equipamentos da plataforma RISC				
2.2 – Comprovação de Performance para os Equipamentos da plataforma INTEL				
2.3 – Requisitos Gerais				
2.4 – Garantia				
2.5 – Alimentação Elétrica				
2.6 – Assistência técnica				
2.7 – Recursos Mínimos de Particionamento para os Servidores RISC				19346 E
2.8 – Recursos Mínimos de Gerenciamento do Hardware requeridos para os Servidores INTEL				
2.9 – Recursos Mínimos de hardware e software para os Servidores INTEL	**107 licenças de Enhanced SystemEdge**		**SIM**	**PROPOSTA TÉCNICA ITEM 11.16 página 20**
2.10 – Recursos Mínimos de hardware e software para os servidores RISC adicionais	**88 licenças de Enhanced SystemEdge**		**SIM**	**PROPOSTA TÉCNICA ITEM 11.16 página 20**
2.11 – Comprovação dos Requisitos Técnicos				
2.12 – Organização da				

Os itens não preenchidos estão comprovados nos demais anexos.

Documentação Técnica				
2.13 – Suporte Remoto				
2.14 – Configuração das Ferramentas de Gerenciamento	**1) Instalação e configuração dos agentes Enhanced SystemEdge de acordo com os padrões da plataforma de Gerência já existente na ECT;** **2) Licenças e configuração do Concord Service Availability para gerenciar DNS, HTTP, HTTPS, SMTP, POP3, FTP e TCP;** **3) Serão configurados as seguintes extensões do Concord SystemEdge: Exchange, Oracle, SQL, Apache e IIS;** *4) A ser respondido pela Integradora* **5) Serão realizados os serviços de integração dos eventos dos elementos desta licitação com o HP OpenView OVO e NNM; Criação e integração de regras no Live Health da**	**2) Possibilidade de gerenciamento de PING e NNTP;**		2) Manual **eHealth Service Response User Guide (Arquivo svcrsp.pdf)**, página 11. Obs: Service Availability é o novo nome comercial para o Service Response; 3) Manuais: **ADVANTEDGE FOR MICROSOFT® EXCHANGE User Guide (Arquivo ExchangeUser.pdf)**, página 1-1; **eHealth AIM for Oracle User Guide (Arquivo OracleUser.pdf)**, página 11; **ADVANTEDGE FOR MICROSOFT SQL SERVER User Guide (Arquivo SQLUser.pdf)**, página 1-1; **eHealth AIM for Apache User Guide (Arquivo ApacheUser.pdf)**, página 12; **ADVANTEDGETM FOR MICROSOFT® IIS User Guide (Arquivo IISUser.pdf)**, página 1-1.

19345

Os itens não preenchidos estão comprovados nos demais anexos.

	Concord, enviando alarmes para o ambiente HP OpenView; Criação de relatórios e agendamento de geração de relatórios do Concord eHealth.			

1a344
E

Os itens não preenchidos estão comprovados nos demais anexos.

19343
E

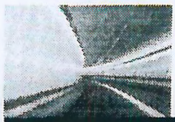

eHealth AIM for Apache User Guide

09-16050-003
December 2001

19342
ε

Important Notice

Concord Communications, Inc., *e*Health, *e*Health Suite, the Concord Logo, eROI, AdvantEDGE, SystemEDGE, Live Health, Network Health, Live Status, System Health, Application Health, Automating Technology Management, Constellation, Enterprise, Enterprise Monitor, Firstsense, FirstSense and design, FirstSense Enterprise, Pulse, Pulsecheck, Token/Net, Token/Scope, We See It Happening, Fault Manager, Empire, Empire Technologies and/or other Concord marks or products referenced herein are either registered trademarks or trademarks of Concord Communications, Inc. Other trademarks are the property of their respective owners.

Adobe, Acrobat, and the Acrobat logo are trademarks, which may be registered in certain jurisdictions, and PostScript is a registered trademark of Adobe Systems Incorporated.
Linux is a registered trademark of Linus Torvalds.
Microsoft, Windows, Windows NT, Windows 2000, and Windows XP are trademarks of Microsoft Corporation.
Intel is a registered trademark of Intel Corporation.
Red Hat is a registered trademark of Red Hat, Inc.
Sun, Solaris, and SunOS are trademarks of Sun Microsystems, Inc.
UNIX is a trademark in the United States and other countries licensed exclusively through X/Open Company Ltd.

This product includes software developed by the Apache Group for use in the Apache HTTP server project (http://www.apache.org/).

© Copyright Bigelow and Holmes 1986, 1985. Lucida is a registered trademark of Bigelow & Holmes. Sun Microsystems Inc., AT&T, and Bigelow & Holmes make no representations about the suitability of the source code for any purpose. It is provided "as is" without express or implied warranty of any kind.

General Notice: Some of the product names used herein have been used for identification purposes only and may be trademarks of their respective companies.

Proprietary Notice

The information and descriptions contained herein are the property of Concord Communications, Inc. Such information and descriptions may not be copied, disseminated, or distributed without the express written consent of Concord Communications, Inc. Concord Communications, Inc., assumes no responsibility for any inaccuracies that may appear in this document. Concord Communications, Inc., reserves the right to improve its products and change specifications at any time without notice.

U. S. Government Restricted Rights

Use, reproduction, and disclosure by the U.S. Government are subject to the restrictions set forth in FAR §52.227-19 (c) (1) and (2) and DFARS §252.227-7013 (c) (1) (ii).

Patent Information

U. S. Patent 5,615,323
Patents Pending

19341
E

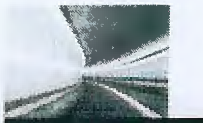

Table of Contents

Preface **7**

Audience. 7
About This Guide 7
    Revision Information. 8
    Documentation Conventions 8
Technical Support 9
Professional Services 9

Chapter 1 **Introduction** **11**

Introducing *e*Health AIM for Apache 11
Features 12
    Using the mod_info and mod_status Facilities 13
    Caching Disk Space Information 13
Using *e*Health AIM for Apache 13
    Using *e*Health AIM for Apache with AdvantEDGE View. . . . 14
    Using *e*Health AIM for Apache with *e*Health 16
    Using *e*Health AIM for Apache with Live Health 16

Chapter 2 **Installing *e*Health AIM for Apache** **17**

Installation Requirements. 17
Installing the Software. 18
    Installing the Software for UNIX Systems 18
    Installing the Software for Windows NT Systems 19

10340
ε

*e*Health AIM for Apache Files. 20
Configuring *e*Health AIM for Apache 21
Editing the sysedge.cf File. 21
Editing the httpd.conf File 22
Editing the apachemod.cf File 24
Licensing *e*Health AIM for Apache 25
Obtaining a License. 26
Generating the License through the Web-Based License Form. 27
Generating a License through AdvantEDGE View Event Processing 29
Generating a License through AdvantEDGE View Host Administration 31
Sample License File. 33

Chapter 3 **Using the *e*Health AIM for Apache MIB** **35**

Configuration Section 36
Server Configuration 36
Performance Section 38
Apache Footprint 38
Server Performance. 41

Chapter 4 **Using *e*Health AIM for Apache** **45**

Overview. 45
*e*Health AIM for Apache MIB Branch 46
Assigning Entry Rows for the SystemEDGE Self-Monitoring Tables 46
Setting Local Policy. 46
Reserving Blocks of Rows. 46
Using the SystemEDGE Self-Monitoring Features 47
Using SystemEDGE Process Monitoring 48
Using SystemEDGE Threshold Monitoring. 49
Using SystemEDGE History Collection 50

19338

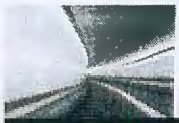

Preface

This guide describes how to install and use the *e*Health application insight module (AIM) for Apache for Solaris SPARC, Linux x86, HP-UX, and Windows NT 4.0, Windows 2000, and Windows XP systems. This guide supports the following:

- *e*Health AIM for Apache Release 1.0 Patchlevel 2 and later
- *e*Health SystemEDGE Release 4.0 Patchlevel 3 and later

This product supports the Apache Web server Version 1.3.2 and later.

Audience

This guide is intended for the person responsible for installing and configuring *e*Health AIM for Apache. To use *e*Health AIM for Apache, you should have a basic understanding of the Apache Web server, *e*Health SystemEDGE, and your host's operating system environment. For more information, refer to Apache documentation (http://www.apache.org) and the *eHealth SystemEDGE User Guide.*

About This Guide

This section describes the changes and enhancements that have been made since the last release of this guide. It also includes the documentation conventions used in this guide.

19337
E

Revision Information

This guide now includes installation and configuration instructions for *e*Health AIM for Apache on Windows NT, Windows 2000, Windows XP, and HP-UX 10.x and 11.x systems.

Documentation Conventions

Table 1 lists the conventions used in this document.

Table 1. Documentation Conventions (Page 1 of 2)

Convention	Description
File or Directory Name	File or directory names.
`code`	System, code, or operating system command line examples.
emphasis	Emphasis and guide titles.
`enter`	Text that you must type exactly as shown.
Name	Text that defines menus, fields in dialog boxes, or keyboard keys.
New Term	A new term, that is, one that is being introduced.
`Variable`	Variable values that you substitute.
→	A sequence of menus or menu options. For example, **File** → **Exit** means "Choose **Exit** from the **File** menu."

19336 E

Table 1. Documentation Conventions (Page 2 of 2)

Convention	Description
NOTE	Important information, tips, or other noteworthy details.
CAUTION	Information that helps you avoid data corruption or system failures.
WARNING	Information that helps you avoid physical danger.

Technical Support

If you need any assistance with this product or the SystemEDGE agent, contact Technical Support at the following:

Phone:	(888) 832-4340
	(508) 303-4300
Fax:	(508) 303-4343
E-mail:	support@concord.com
Web site:	http://www.concord.com

Professional Services

If you need any assistance with customizing this product, contact Professional Services at the following:

Phone:	(800) 851-8725 (Choose option 7)
Fax:	(508) 486-4555
E-mail:	proserv@concord.com
Web site:	http://www.concord.com

19335
E

19334

Introduction

This chapter provides an overview of *e*Health AIM for Apache.

Introducing *e*Health AIM for Apache

*e*Health AIM for Apache is a plug-in for the SystemEDGE agent that enables information technology (IT) operators to monitor the Apache Web server. The power and flexibility of the Apache Web server make it the server of choice for many corporations.

*e*Health AIM for Apache enables you to monitor the Apache server's performance and availability on these operating systems:

- Sun Solaris Release 2.5 and later
- RedHat Linux Release 6.0 and later
- HP-UX Release 10.x and 11.x (not including 10.01)
- Microsoft Windows NT 4.0, Windows 2000, and Windows XP

To use *e*Health AIM for Apache, you must install it on every Apache server that you want to monitor. For more information, refer to Chapter 2, "Installing eHealth AIM for Apache."

19333
E.

NOTE

This guide is not intended to describe how to install, administer, or use the Apache Web server. For help with Apache, refer to your Apache Web server documentation.

Features

*e*Health AIM for Apache monitors the following:

- Apache server processes
- Apache server log files
- Performance metrics that are specific to Apache
- Web service response and availability

*e*Health AIM for Apache monitors httpd process attributes. For example, it monitors whether each process is running (alive). It can also restart processes, if necessary. In addition, it monitors memory usage, memory size, and page faults.

Because the Apache Web server records Web accesses and errors in log files, *e*Health AIM for Apache can use the SystemEDGE agent log-file monitoring capability to scan those logs and forward certain events as Simple Network Management Protocol (SNMP) traps when appropriate.

*e*Health AIM for Apache is designed to monitor one or more Apache servers running on a single system. To support multiple servers, the data presented in the *e*Health AIM for Apache management information base (MIB) is organized into tables that are indexed by server port number. For instance, if you have a server running on port 80 and another on port 8080, entries appear in each table for index 80 and index 8080. For more information, refer to Chapter 3, "Using the eHealth AIM for Apache MIB."

1193 32
E

Using the mod_info and mod_status Facilities

*e*Health AIM for Apache obtains server status and behavior information from the Apache server through the Apache **mod_info** and **mod_status** facilities. To access this type of information, you must enable these facilities in your Apache configuration file (httpd.conf). If these facilities are secured through the use of an authentication mechanism, you must also provide a valid user name and password for accessing these resources in the apachemod.cf file. For more information, refer to "Editing the httpd.conf File" on page 23 and "Editing the apachemod.cf File" on page 25.

Caching Disk Space Information

As part of the application footprint calculations, *e*Health AIM for Apache can calculate the total amount of disk space used by your Apache server. On some systems with large amounts (greater than 1 GB) of Web data, this calculation can take a long time. For this reason, *e*Health AIM for Apache is designed to cache the disk space information to avoid frequent recalculations. You can control the frequency of the disk space calculations in the *e*Health AIM for Apache configuration file (apachemod.cf). For more information, refer to "Editing the apachemod.cf File" on page 25.

Using *e*Health AIM for Apache

*e*Health AIM for Apache provides you with the tools and information that are necessary for monitoring the health and availability of the Apache Web server. It makes important information about Apache available to management software through the SystemEDGE agent and SNMP.

19331
E

*e*Health AIM for Apache works with the SystemEDGE agent to closely manage the Apache Web server, providing real-time fault detection and automatically correcting problems, if necessary. You can use *e*Health AIM for Apache with any SNMP-compliant management software, including Concord's *e*Health suite of products, AdvantEDGE View, HP OpenView, and others.

*e*Health AIM for Apache and the SystemEDGE agent can provide you with the following types of information:

- Number of "hits" your Web server is receiving, which can help you keep up with daily volume and set monitor points to watch for unusual traffic loss or denial of service attacks
- Amount of space your Web log and Web server files are consuming
- How effectively the Apache processes monitor idle services, warn you when the number of idle services is too low, and monitor the number of active processes
- How much of your system resources (Central Processing Unit [CPU] and memory) Apache is using on your server
- Whether bottlenecks on your Web servers are caused by problems with the CPU, memory, disk, or network

Using *e*Health AIM for Apache with AdvantEDGE View

You can use AdvantEDGE View with *e*Health AIM for Apache to monitor the performance, configuration, availability, and health of the Apache Web server.

To run an AdvantEDGE View query for Apache:

1. Select the target system or group from the **System** or **Group** list in the AdvantEDGE View interface.
2. Select **Apache** from the **Applications** list.
3. Click the **Applications** icon.

19330
ε

Figure 1 shows a sample image map that AdvantEDGE View displays when you run a query on the target Apache workstation. Click the area for which you want to display information.

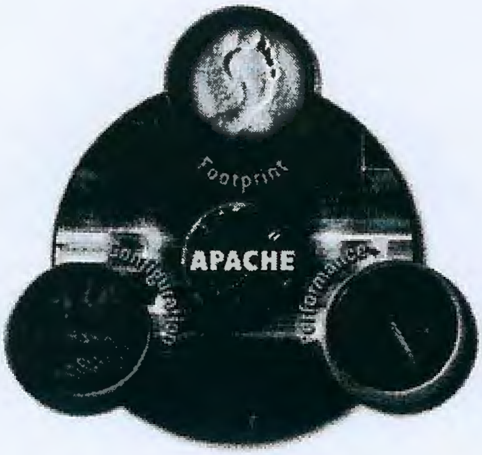

Configuration Footprint Performance

Figure 1. AdvantEDGE View Image Map for Apache Queries

For example, if you click the **Footprint** area, AdvantEDGE View displays information about Apache's CPU, memory, and disk resource consumption. Figure 2 shows the Disk Usage section of the AdvantEDGE View Footprint query for an Apache Web server.

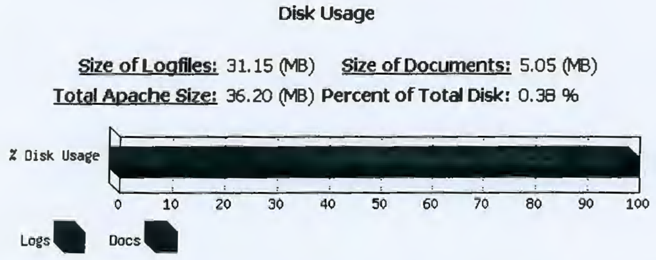

Disk Usage

Size of Logfiles: 31.15 (MB) Size of Documents: 5.05 (MB)
Total Apache Size: 36.20 (MB) Percent of Total Disk: 0.38 %

Figure 2. Sample Section of an AdvantEDGE View Footprint Query for Apache

Using *e*Health AIM for Apache with *e*Health

In previous releases, eHealth – Application Insight was called eHealth – Application Assessment.

You can use *e*Health AIM for Apache and the SystemEDGE agent with the *e*Health product suite to provide historical data for long-term trending analysis and capacity planning. With *e*Health – Application Insight, you can run At-a-Glance, Trend, Top N, and MyHealth reports for the following types of information:

- Amount of Central Processing Unit (CPU), total memory, and disk space the Apache Web server is using
- Number of users who are connected to the Apache server
- Number and type of processes that are running
- Amount of Web traffic
- Total size of the Apache service logs
- End-to-end performance summaries for your Apache application, host system, and network

For more information about the variables you can monitor and reports that you can run when you integrate *e*Health AIM for Apache with *e*Health, refer to the *e*Health Web Help.

Using *e*Health AIM for Apache with Live Health

You can also use *e*Health AIM for Apache and the SystemEDGE agent with Live Health for real-time detection of potential problems. Live Health applies intelligent algorithms to the data, resulting in precise assessments of application health and performance. For more information about how Live Health can detect "brownouts" and service delays across applications, systems, and networks, refer to the Live Health Web Help.

19328

2

Installing eHealth AIM for Apache

This chapter explains how to install, configure, and license *e*Health AIM for Apache.

NOTE

> For the most current information about installing this module, refer to the relnotes.txt file on the *e*Health AIM for Apache installation CD-ROM.

Installation Requirements

Before you install *e*Health AIM for Apache, you must first install, license, and configure the SystemEDGE agent Release 4.0, Patchlevel 3 or later. For more information, refer to the *eHealth SystemEDGE User Guide*. Also, your system must be running the Apache Web server Release 1.3.2 or later on one of these operating systems:

- Sun Solaris (SPARC) Release 2.5 or later
- Red Hat Linux (x86) Release 6.0 or later
- HP-UX Release 10.x and 11.x (not including 10.01)
- Microsoft Windows NT 4.0, Windows 2000, or Windows XP

Installing the Software

This section describes how to install *e*Health AIM for Apache for Solaris, Linux, HP-UX, and Windows operating systems.

Installing the Software for UNIX Systems

For Solaris, Linux, and HP-UX systems, *e*Health AIM for Apache is distributed as a tar file.

To install *e*Health AIM for Apache:

1. Locate the Apache server that you want to monitor.
2. Log in to the UNIX system where that server is located as **root**.
3. Copy the apachemod.tar file from the CD-ROM to the /tmp directory.
4. Change directory to the SystemEDGE agent directory on your system by entering the following:

   ```
   cd /opt/EMPsysedge
   ```
5. Create the plugins directory, if it is not already present, by entering the following:

   ```
   mkdir plugins
   ```
6. Change directory to the plugins directory by entering the following:

   ```
   cd plugins
   ```
7. Enter the following to start the installation.

   For Solaris, enter the following command:

   ```
   tar xvf /tmp/apachemod_1.0p2_sol.tar
   ```

   For Linux, enter the following command:

   ```
   tar xvf /tmp/apachemod_1.0p2_linux.tar
   ```

   For HP-UX, enter the following command:

   ```
   tar xvf /tmp/apachemod_1.0p2_hpux.tar
   ```

   The installation creates files in the plugins/apachemod directory. *e*Health AIM for Apache is now installed.

1737
Paula

to the *<system_name.company.com>*, but it cannot because it does not know this system as *<system_name.company.com>*, but only as *<system_name>*. The system name cannot be expanded to the full name because the DNS is improperly configured. This situation can also be the other way around, where DNS is configured on the Cell Manager and not on the Application Client.

Action Set up the TCP/IP protocol and configure DNS properly. Refer to Appendix B in the *HP OpenView Storage Data Protector Installation and Licensing Guide* for information.

Problems with non-ASCII Characters in Filenames

In mixed platform environments, there are some limitations regarding filenames containing non-ASCII characters.

Action See Appendix B in the *HP OpenView Storage Data Protector Software Release Notes* for types of problems that occur, for situations in which they occur, and for workarounds.

19325
G

*e*Health AIM for Apache Files

Table 2 describes the files created by the *e*Health AIM for Apache installation procedure.

Table 2. Files Installed by *e*Health AIM for Apache

File Name	Description
apachemod.asn1	*e*Health AIM for Apache MIB specification
apachemod.cf	*e*Health AIM for Apache configuration file
apachemod.pdf	*eHealth AIM for Apache User Guide*
apachemod.dll (Windows only)	*e*Health AIM for Apache dynamic link library (DLL) module for Windows NT, Windows 2000, and Windows XP
apachemod.so (Solaris and Linux only)	*e*Health AIM for Apache shared library for 32-bit Solaris and Linux operating systems
apachemod-sparcv9.so (Solaris only)	*e*Health AIM for Apache shared library for 64-bit Solaris operating systems
apachemod-hpux.so (HP-UX only)	*e*Health AIM for Apache shared library for 32-bit HP-UX (10.x and 11.x) operating systems
apachemod-hpux11-64.so (HP-UX only)	*e*Health AIM for Apache shared library for 64-bit HP-UX (11.x) operating systems
relnotes.txt	Release notes for *e*Health AIM for Apache

19324
E

Configuring *e*Health AIM for Apache

After you install *e*Health AIM for Apache, you must configure it by editing the following files:

- sysedge.cf
- httpd.conf
- apachemod.cf

Editing the sysedge.cf File

By default, the SystemEDGE agent does not load any plug-ins at initialization time, but you can edit the sysedge.cf file to configure the agent to load any *e*Health AIMs that you have installed. To enable the SystemEDGE agent to load *e*Health AIM for Apache at startup, you must edit the sysedge.cf configuration file. This file is located in your system directory by default; for example, it is located in the /etc/sysedge.cf directory on UNIX systems and in C:\winnt\system32 for Windows NT and 2000 systems. Use the **sysedge_plugin** keyword as described in the following sections to configure SystemEDGE to load *e*Health AIM for Apache at startup.

NOTE

To configure the SystemEDGE agent to start *e*Health AIM for Apache, you must provide the *complete pathname* to the shared library file for your system.

Enabling *e*Health AIM for Apache for UNIX Systems

Add one of the following lines to the sysedge.cf file.

For Solaris or Linux systems in 32-bit mode, add the following line:

```
sysedge_plugin /opt/EMPsysedge/plugins/apachemod/apachemod.so
```

fl322
ε

For Solaris 2.7 and later systems in 64-bit mode, add the following line:

```
sysedge_plugin /opt/EMPsysedge/plugins/apachemod/apachemod-sparcv9.so
```

For HP-UX 10.x systems in 32-bit mode, add the following line:

```
sysedge_plugin /opt/EMPsysedge/plugins/apachemod/apachemod-hpux.so
```

For HP-UX 10.x and 11.x systems in 64-bit mode, add the following line:

```
sysedge_plugin /opt/EMPsysedge/plugins/apachemod/apachemod-hpux11-64.so
```

Enabling *e*Health AIM for Apache for Windows Systems

To enable *e*Health AIM for Apache for Windows NT, Windows 2000, and Windows XP systems, you must provide the complete path name to apachemod.dll. The actual path depends on the location you selected when you installed *e*Health AIM for Apache. For example, enter the following command if you installed the files in the C:\sysedge\plugins\apachemod directory:

```
sysedge_plugin C:\sysedge\plugins\apachemod\apachemod.dll
```

For more information about the sysedge.cf file, refer to the *eHealth SystemEDGE User Guide*.

Editing the httpd.conf File

Edit the httpd.conf file to enable the Apache mod_info and mod_status facilities. These facilities control access to HTML pages that provide information about the server's status and behavior. When these modules are enabled, *e*Health AIM for Apache can use them to obtain information about the Apache server's status and behavior. For more information about the mod_info and mod_status modules, refer to Apache documentation (at http://www.apache.org).

NOTE

You must ensure that your server has server-status and server-info built into Apache by enabling the mod_status and mod_info facilities, as described in the following section. If you do not, you may experience process-termination errors.

To enable the mod_info and mod_status features in your Apache server:

1. Remove the pound sign (#) in front of the following lines in httpd.conf:

   ```
   LoadModule status_module modules/mod_status.so
   LoadModule info_module modules/mod_info.so
   ```

2. Add the following lines to httpd.conf:

```
# Turn on Extended Status Information
ExtendedStatus On
# Enable server-status access from the local host
<Location /server-status>
   SetHandler server-status
   Order deny,allow
   Deny from all
  Allow from 127.0.0.1
</Location>
  # Enable server-info access from the local host
  <Location /server-info>
     SetHandler server-info
      Order deny,allow
      Deny from all
      Allow from 127.0.0.1
  </Location>
```

NOTE

You must restart the Apache server after you make these changes to ensure that they take effect.

19320
E

Editing the apachemod.cf File

You can edit the apachemod.cf file to do the following:

- Indicate which Transmission Control Protocol (TCP) port your Apache server is using.
- Assign user names and passwords for the mod_info and mod_status facilities. You must include user names and passwords to the apachemod.cf file, even if they are not specifically required by the Apache Web server. The Apache Web server requires these fields *only* if you have enabled the mod_info and mod_status facilities as described in the previous section, "Editing the httpd.conf File," and you are using authentication (through user names and passwords) to restrict access to those modules. The apachemod.cf file, however, *always* requires these fields.

NOTE

Even if you do not password protect your server-info and server-status pages, you must configure artificial user name/password combinations in apachemod.cf. The Apache Web server and *e*Health AIM for Apache then ignore those user names and passwords.

- Indicate which Apache server(s) *e*Health AIM for Apache should monitor if you are running more than one Apache server on your system through port numbers. Port 80 is the default Web server port.
- Set the interval for calculating the total amount of disk space being used by the Apache server, or disable this checking.

19319
ε

Sample apachemod.cf File

The following is a sample apachemod.cf file:

```
# apachemod.cf
# Configuration file for eHealth AIM for Apache
#
# For each apache server running on your system, specify the following:
# apache port username password filestat-interval
#  port - port number on which the server is running
#  username - the username for accessing server-status and server-info pages
#  password - the password for accessing server-status and server-info pages
#  filestat-interval - interval in seconds between checks of the file sizes
#                      specify "0" to disable file size checking
# Primary server - example
apache 80 status statpass1 3600
# Application server - example
# apache 8080 status statpass1 3600
```

NOTE

After you make any changes to the apachemod.cf file, you must restart the SystemEDGE agent to ensure that the changes take effect.

Licensing *e*Health AIM for Apache

Like the SystemEDGE agent, *e*Health AIM for Apache utilizes a host-based license method. Copies of *e*Health AIM for Apache can run only on systems that possess a valid license key. This license is separate from the one used for the SystemEDGE agent.

1318

The first time that you attempt to start the SystemEDGE agent after installing *e*Health AIM for Apache, the agent displays a message stating that it could not find a valid license for *e*Health AIM for Apache. It then provides you with a **public key** that is used to generate a permanent license key for your host machine.

A license key is composed of four space-separated, 8-character sequences, totaling 32 characters. The sysedge.lic file contains the *e*Health AIM for Apache license, as well as the SystemEDGE agent license and other *e*Health AIM licenses. For an example, refer to the sample license file in "Sample License File" on page 34.

Obtaining a License

To obtain a license, you can do any of the following:

- Complete the online license form through the Internet, as described in the next section, "Generating the License through the Web-Based License Form."
- Use AdvantEDGE View to receive an SNMP license trap or to query and license the plug-in without a trap. For more information, refer to "Generating a License through AdvantEDGE View Event Processing" on page 30 or "Generating a License through AdvantEDGE View Host Administration" on page 32.
- Send an e-mail request to licenses@concord.com, and place the returned license key in the sysedge.lic file. Always include your user name in license requests that you send through e-mail.
- Run the Concord-supplied licenseutil.pl script.
- Run the licenseme.exe license utility.

For more information about licensing, refer to the *eHealth SystemEDGE User Guide* and the *Automating the Licensing of the AdvantEDGE Point Plug-in Modules* white paper.

1a3\7
E

Generating the License through the Web-Based License Form

This section describes how to generate the *e*Health AIM for Apache license through the Web-based license form.

NOTE

If you are using an evaluation copy of *e*Health AIM for Apache, you must request a temporary license to enable it to operate during the evaluation period.

To generate a license:

1. Start the SystemEDGE agent as follows:
   a. Log in as **root**.
   b. Change directory (cd) to /opt/EMPsysedge.
   c. Enter the following:

      ```
      bin/sysedge
      ```

      The SystemEDGE agent displays a message indicating that you need a license for the *e*Health AIM for Apache module on this host system. It then displays a message similar to the following:

```
SystemEDGE Version 4.1 Patchlevel 1
Copyright 2001 by Concord Communications, Inc.
Please contact Concord Communications, Inc. to obtain a license
http://www.concord.com/support, Email: licenses@concord.com
Provide this: apachemod neptune sol2 5.9 346561363366b19c 1.0 Patchlevel 2
```

2. Using a Web browser, go to the licensing Web site at http://license.concord.com, and select the **Create License** option that matches your use of the agent:
   - **Create SystemEDGE/AdvantEDGE Eval License** (if you are evaluating the AIM or are a Concord partner or reseller)
   - **Create SystemEDGE Outsource License** (if you are outsourcing the AIM)
   - **Create SystemEDGE/AdvantEDGE License** (if you have purchased the AIM)

1a316
ε

Note

You must supply a user name and password to access the license form.

3. Complete the license form, entering the information that was printed by the SystemEDGE agent. You must supply the following information:
    - Name
    - E-mail address
    - Software version number (1.0 in the example)
    - Patchlevel (2 in the example)
    - System name (neptune in the example)
    - Operating system name (sol2 in the example)
    - Operating system version (5.9 in the example)

    Note

    When you request a license, select the option for *e*Health AIM for Apache in the product field of the licensing form.

    After you submit the license request form, the Concord Web server generates a license and displays it on your Web browser. It also e-mails the license to the contact person in your organization.
4. Copy the license into the sysedge.lic file. This file is located in the /etc directory for UNIX operating systems and in the \winnt\system32 directory for Windows NT, Windows 2000, and Windows XP operating systems.

    The license key is case sensitive. Copy it exactly as it appears. If possible, use your system's copy-and-paste facility instead of typing it by hand. If you are entering the license key manually, be careful not to confuse characters such as the letters l and I and the number 1, or the letter O and the number 0.
5. Save the sysedge.lic file.

11315
ε

6. Restart the SystemEDGE agent.

   For UNIX systems, restart the SystemEDGE agent by entering the following when you are logged in as root:

   ```
   bin/sysedge -b
   ```

   For Windows systems, stop and restart the Windows NT Master agent by entering these commands at the command prompt:

   ```
   net stop snmp
   net start snmp
   ```

Generating a License through AdvantEDGE View Event Processing

In order to use AdvantEDGE View event processing to license *e*Health AIM for Apache, your system must meet the following requirements:

- You must be using SystemEDGE Release 4.0 Patchlevel 3 or later with AdvantEDGE View.
- You must configure the SystemEDGE agent to send SNMP traps to AdvantEDGE View. For more information, refer to the section on configuring the SystemEDGE agent in the *eHealth SystemEDGE User Guide.*
- You must configure the SystemEDGE agent with a read-write community so that AdvantEDGE View can issue an SNMP Set to transmit the license key to it. For more information, refer to the section on configuring the SystemEDGE agent in the *eHealth SystemEDGE User Guide.*
- Your AdvantEDGE View system must have access to the Internet, either directly or through a Web proxy.
- The AdvantEDGE View User who is generating the license must have either write or admin permissions.

19314
ε

To generate a license through AdvantEDGE View:

1. Start the SystemEDGE agent with *e*Health AIM for Apache in unlicensed mode. SystemEDGE sends a license trap to AdvantEDGE View for that module.
2. Start AdvantEDGE View, and click the **Events** icon to display the Event Processing screen.

   AdvantEDGE View displays a license trap for the system that requires a license.
3. Click the index number for that system to view the Trap Details form for **License Software** to display the AdvantEDGE View Software Licensing form.
4. Complete the licensing form, and click **Get License**.

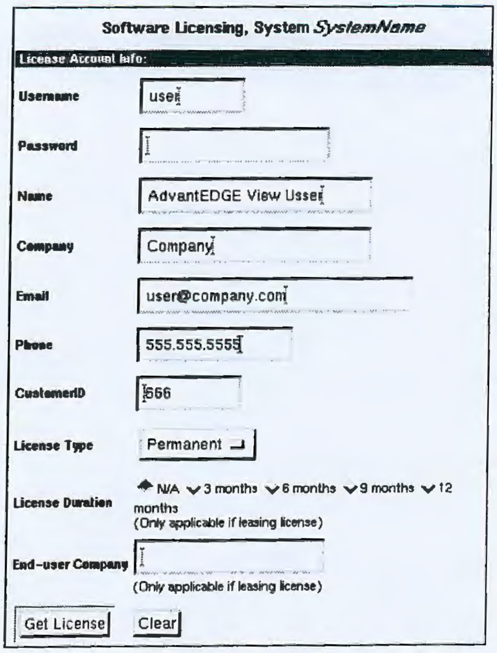

Software Licensing, System *SystemName*

License Account Info:

Username	user
Password	
Name	AdvantEDGE View Usser
Company	Company
Email	user@company.com
Phone	555.555.5555
CustomerID	666
License Type	Permanent
License Duration	◆ N/A ✓ 3 months ✓ 6 months ✓ 9 months ✓ 12 months (Only applicable if leasing license)
End-user Company	 (Only applicable if leasing license)

19313
E

NOTE

If you have configured AdvantEDGE View preferences, AdvantEDGE View fills in all of the information (except password) on this form. You must enter the password each time you use the form for security purposes.

AdvantEDGE View contacts the Web-based license server, obtains a license for *e*Health AIM for Apache and issues an SNMP Set to the target SystemEDGE agent to inform it of the new software license key.

Generating a License through AdvantEDGE View Host Administration

You can also license systems through AdvantEDGE View Host Administration.

To access Host Administration:

1. Start AdvantEDGE View, and click the **Administration** icon. AdvantEDGE View displays the Administration page.
2. Click the **Host Administration** icon. AdvantEDGE View displays the host list.

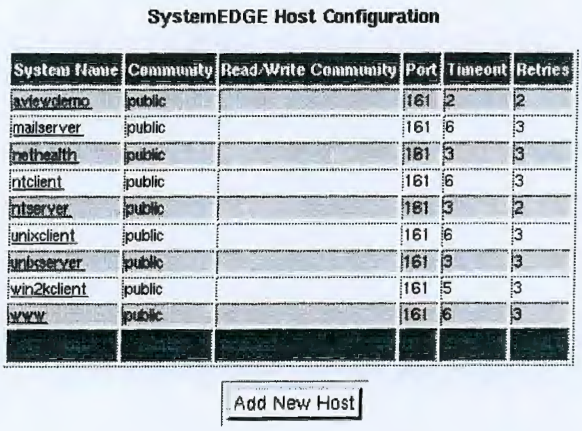

SystemEDGE Host Configuration

System Name	Community	Read/Write Community	Port	Timeout	Retries
aviewdemo	public		161	2	2
mailserver	public		161	6	3
nethealth	public		161	3	3
ntclient	public		161	6	3
ntserver	public		161	3	2
unixclient	public		161	6	3
unixserver	public		161	3	3
win2kclient	public		161	5	3
www	public		161	6	3

19312
ε

3. Click the name of the system that you want to license from the **System Name** column. AdvantEDGE View displays the Modify Host form.

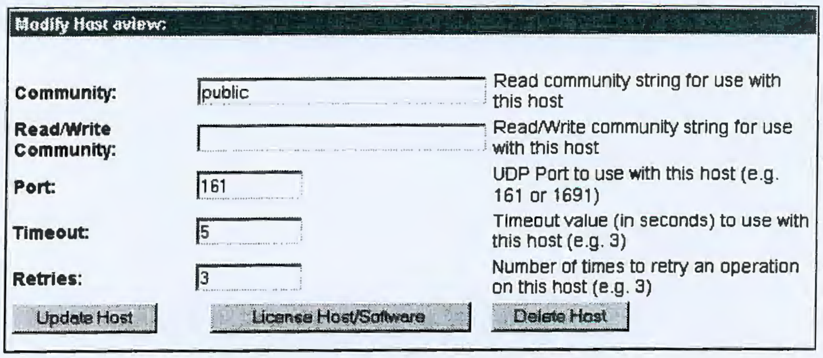

Modify Host aview:

Community:	public	Read community string for use with this host
Read/Write Community:		Read/Write community string for use with this host
Port:	161	UDP Port to use with this host (e.g. 161 or 1691)
Timeout:	5	Timeout value (in seconds) to use with this host (e.g. 3)
Retries:	3	Number of times to retry an operation on this host (e.g. 3)

4. Click **License Host/Software** to display the licensing form.

5. Select the product you want to license from the **Product** list, and then click **License Software.**

   AdvantEDGE View contacts the Web-based license server, obtains a license for the software, and issues an SNMP Set to the target SystemEDGE agent, informing it of the new software license key.

19311
ε

Sample License File

The following is a sample SystemEDGE agent license file. A pound character (#) in column 1 indicates that the entire line is a comment.

```
# license file for SystemEDGE Agent
# Concord Communications, Inc.
# http://www.concord.com
#
# file /etc/sysedge.lic or %SystemRoot%\system32\sysedge.lic
# A valid license key has four parts of 8 characters per part
# parts are separated by space(s) with one license key per line

# sysedge neptune sol2 5.8 807cb1da007cb1da 4.1 Patchlevel 1
e13311d3 0F2a7cb1 abC512dc fF8C923a
# apachemod neptune sol2 5.8 807cb1da007cb1da 1.0 Patchlevel 2
a7943fde 098a87ij a4kiuf39 afafEkj4
```

19310
E

3

Using the eHealth AIM for Apache MIB

This chapter explains the organization and content of the Concord Communications MIB for the Apache Web server. The MIB specification (apachemod.asn1) defines a collection of objects for monitoring and managing Apache. You must configure the SystemEDGE agent to monitor the MIB objects that are relevant for your configuration. For more information, refer to Chapter 4, "Using eHealth AIM for Apache." Figure 3 shows the organization of the Apache MIB.

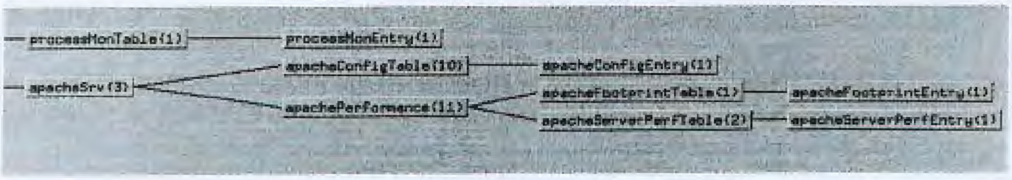

Figure 3. *e*Health AIM for Apache MIB

The figures in the following sections represent AdvantEDGE View queries on information that is available through the AdvanatEDGE for Apache MIB.

The MIB is organized into sections for server configuration and performance. Within the performance section, a footprint section defines MIB objects that convey how much of the underlying system's resources are consumed by Apache. This chapter defines all sections of the Apache MIB and highlights important MIB objects from each section. For a complete list of MIB objects, refer to the *e*Health AIM for Apache MIB specification (apachemod.asn1).

19309
6

Configuration Section

The Configuration section of the *e*Health AIM for Apache MIB contains configuration parameters and settings that are important for streamlining the health and performance of the Apache Web server. It also includes information about server configuration.

Server Configuration

The Server Configuration group contains configuration parameters, process IDs, and version numbers. Table 3 defines important Server Configuration parameters.

Table 3. Selected MIB Objects – Apache Server Configuration Group (Page 1 of 2)

MIB Object	Description
apacheConfigTable	Table of configuration settings for each Apache installation.
apacheConfigEntry	Entry in the Apache configuration table.
apacheConfigPort	TCP port currently used by this Apache configuration. **NOTE** This value serves as the index for this table.
apacheConfigVersion	Apache Web server version.
apacheConfigPID	Process ID (PID) of the master Apache process; zero if the server is not running.
apacheConfigRunMode	Current operating mode (for example, standalone).
apacheConfigUser	User who is currently running the server processes.
apacheConfigGroup	Group that is currently running the server processes.
apacheConfigHostname	Host name used by this Apache configuration.
apacheConfigStartProcs	Number of server processes started by the Apache server at startup.
apacheConfigMinIdleProcs	Minimum number of idle server processes maintained by the Apache server.

19308
ç.

Table 3. Selected MIB Objects – Apache Server Configuration Group (Page 2 of 2)

MIB Object	Description
apacheConfigMaxIdleProcs	Maximum number of idle server processes allowed by the Apache server.
apacheConfigMaxProcs	Maximum number of server processes allowed by the Apache server.
apacheConfigRequestsMaxPerChild	Number of requests handled by a server process before it is recycled by the Apache server.
apacheConfigRequestsKeepAlive	Status of whether keep-alive mode is enabled for persistent connections: disabled = 0; enabled = 1.
apacheConfigRequestsMaxPerConn	Number of requests handled in a single connection if keep-alive mode is enabled.
apacheConfigThreadsPerChild	Maximum number of threads per child process.
apacheConfigConnectionTimeout	Timeout value for closing inactive connections.
apacheConfigKeepAliveTimeout	Keep-alive timeout value for open connections.
apacheConfigServerRoot	Root directory for this Apache installation.
apacheConfigConfigFile	Current configuration file for this Apache installation.
apacheConfigPIDFile	Current PID file for this Apache installation.
apacheConfigScoreboardFile	Current scoreboard file for this Apache installation.
apacheConfigDocumentRoot	Current document root directory for this Apache installation.
apacheConfigAccessLogFile	Current access or transaction log file for this Apache installation.
apacheConfigErrorLogFile	Current error log file for this Apache installation.
apacheConfigScriptLogFile	Current script log file for this Apache installation.

19307e

Figure 4 shows a sample AdvantEDGE View Apache Server Configuration query.

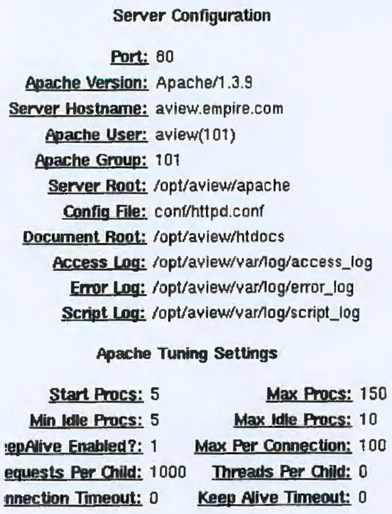

Server Configuration

Port:	80
Apache Version:	Apache/1.3.9
Server Hostname:	aview.empire.com
Apache User:	aview(101)
Apache Group:	101
Server Root:	/opt/aview/apache
Config File:	conf/httpd.conf
Document Root:	/opt/aview/htdocs
Access Log:	/opt/aview/var/log/access_log
Error Log:	/opt/aview/var/log/error_log
Script Log:	/opt/aview/var/log/script_log

Apache Tuning Settings

Start Procs:	5	Max Procs:	150
Min Idle Procs:	5	Max Idle Procs:	10
:epAlive Enabled?:	1	Max Per Connection:	100
equests Per Child:	1000	Threads Per Child:	0
nnection Timeout:	0	Keep Alive Timeout:	0

Figure 4. Sample AdvantEDGE View Server Configuration Query for Apache

Performance Section

The Performance section of the *e*Health AIM for Apache MIB contains performance data that is necessary for capacity planning and trend analysis, as well as real-time performance and availability monitoring. The Performance group is divided into subgroups for footprint data and server performance.

Apache Footprint

The Footprint group provides information about the Apache CPU, memory, data flow, and disk-resource consumption, which is more commonly called the *footprint.* Long-term trending analysis of footprint information is useful for anticipating and avoiding problems due to resource exhaustion.

You can also monitor footprint information in real time to detect and correct temporary resource exhaustion due to viruses, security incidents, and hardware failures. Table 4 defines important Apache Footprint metrics.

Table 4. Selected MIB Objects – Apache Footprint Group (Page 1 of 2)

MIB Object	**Description**
apacheFootprintTable	Table that reports the performance footprint for each Apache service.
apacheFootprintEntry	Entry in the Apache Performance Footprint table.
apacheFootprintPort	TCP port that is currently used by this Apache configuration. **NOTE** This value serves as the index for this table.
apacheFootprintCPUTime	CPU time, in seconds, accumulated by the Apache server, including all Apache processes.
apacheFootprintPercentCPU	Percentage of CPU utilization by the Apache server over the last sample interval; the value reported is percentage multiplied by 100.
apacheFootprintTotalMEMSize	Combined size of Apache's text, data, and stack segments in KB; summation of the process sizes for all Apache server processes.
apacheFootprintTotalRSS	Real memory (resident set) size (RSS) of the Apache server in KB; summation of all process RSS for all Apache server processes.
apacheFootprintPercentMEM	Percentage (0 to 100) of real memory used by the Apache server, which includes all Apache server processes.
apacheFootprintNumThreads	Number of threads executing within all Apache processes of which the operating system is aware.
apacheFootprintInBlks	Number of blocks of data input by the processes.
apacheFootprintOutBlks	Number of blocks of data output by the processes.
apacheFootprintMsgsSent	Number of messages sent by the processes.
apacheFootprintMsgsRecv	Number of messages received by the processes.
apacheFootprintSysCalls	Number of system calls invoked by the processes.

19305
E

Table 4. Selected MIB Objects – Apache Footprint Group (Page 2 of 2)

MIB Object	Description
apacheFootprintMinorPgFlts	Number of minor page faults (which do not require input/output [I/O] to retrieve the page) incurred by the processes.
apacheFootprintMajorPgFlts	Number of major page faults (which require I/O to retrieve the page) incurred by the processes.
apacheFootprintNumSwaps	Number of times the processes have been swapped.
apacheFootprintVolCtx	Number of voluntary context switches incurred by the processes.
apacheFootprintInvolCtx	Number of involuntary context switches incurred by the processes.
apacheFootprintTotalLogSize	Size in KB of the Apache service logs; sum of the access and error log file sizes.
apacheFootprintDocSize	Size in KB of the Apache document root directory and all files beneath it.
apacheFootpringTotalDiskSize	Size in KB of all the Apache disk storage areas; summation of log sizes and service directories.

The following figures show sample AdvantEDGE View Footprint queries for an Apache system that is serving a medium-sized company. Figure 5 shows a sample Disk Usage query.

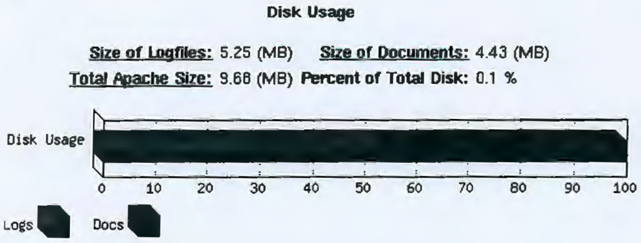

Disk Usage

Size of Logfiles: 5.25 (MB) Size of Documents: 4.43 (MB)
Total Apache Size: 9.68 (MB) Percent of Total Disk: 0.1 %

Figure 5. AdvantEDGE View Disk Usage Query for Apache

19304
ε

Figure 6 shows a sample CPU and Memory Footprint query.

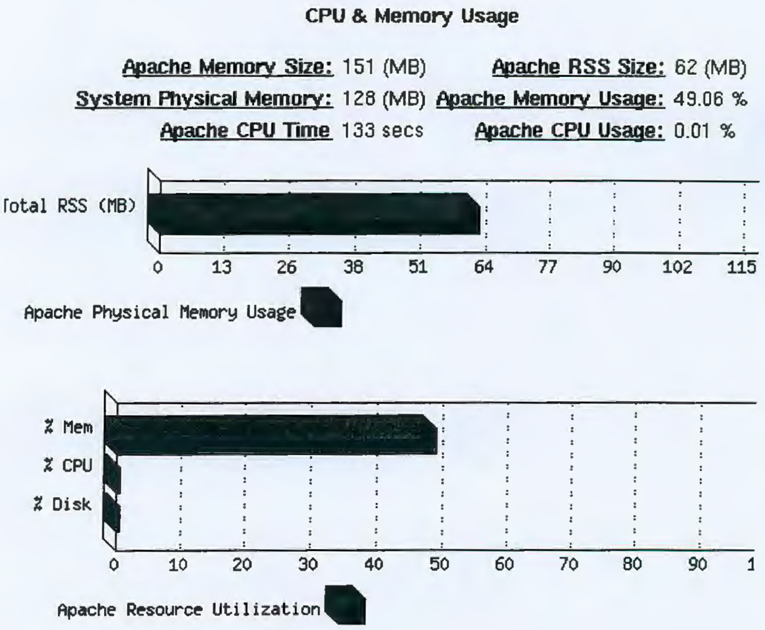

CPU & Memory Usage

Apache Memory Size: 151 (MB) Apache RSS Size: 62 (MB)
System Physical Memory: 128 (MB) Apache Memory Usage: 49.06 %
Apache CPU Time 133 secs Apache CPU Usage: 0.01 %

Figure 6. AdvantEDGE View CPU and Memory Footprint Query for Apache

Server Performance

The Server Performance group provides performance metrics and counters for the Apache Web server, including user statistics and transfer statistics. These metrics include those useful for real-time management and longer-term capacity planning and trend analysis. Table 5 defines important Server Performance metrics.

19303

Table 5. Selected MIB Objects – Apache Server Performance Group (Page 1 of 2)

MIB Object	Description
apacheServerPerfTable	Table that reports the performance of each Apache service.
apacheServerPerfEntry	Entry in the Apache Server Performance table.
apacheServerPerfPort	TCP port that is currently used by this Apache configuration; index for this table.
apacheServerPerfUptime	Number of seconds that the Apache server has been running.
apacheServerPerfTotalAccesses	Number of accesses (hits) to this server since it was last started.
apacheServerPerfTotalTraffic	Number of KB transferred by this server since it was last started.
apacheServerPerfCurrentUsers	Number of current active users (connections) maintained by the Apache server.
apacheServerPerfCurrentIdleProcs	Number of current idle processes available on the Apache server.
apacheServerPerfCurrentStartupProcs	Number of processes that are currently in startup mode on the Apache server.
apacheServerPerfCurrentReadProcs	Number of processes that are currently reading requests on the Apache server.
apacheServerPerfCurrentReplyProcs	Number of processes that are currently replying to requests on the Apache server.
apacheServerPerfCurrentKeepAliveProcs	Number of processes that are currently in keep-alive mode on the Apache server.
apacheServerPerfCurrentDNSProcs	Number of processes that are currently doing a Domain Name System (DNS) lookup on the Apache server.

19302 ε

Table 5. Selected MIB Objects – Apache Server Performance Group (Page 2 of 2)

MIB Object	Description
apacheServerPerfCurrentLoggingProcs	Number of processes that are currently logging transactions on the Apache server.
apacheServerPerfCurrentFinishingProcs	Number of processes that are currently finishing transactions on the Apache server.
apacheServerPerfCurrentTotalProcs	Total number of Apache processes that are currently running on the Apache server.

Figure 7 shows an Apache server Performance summary.

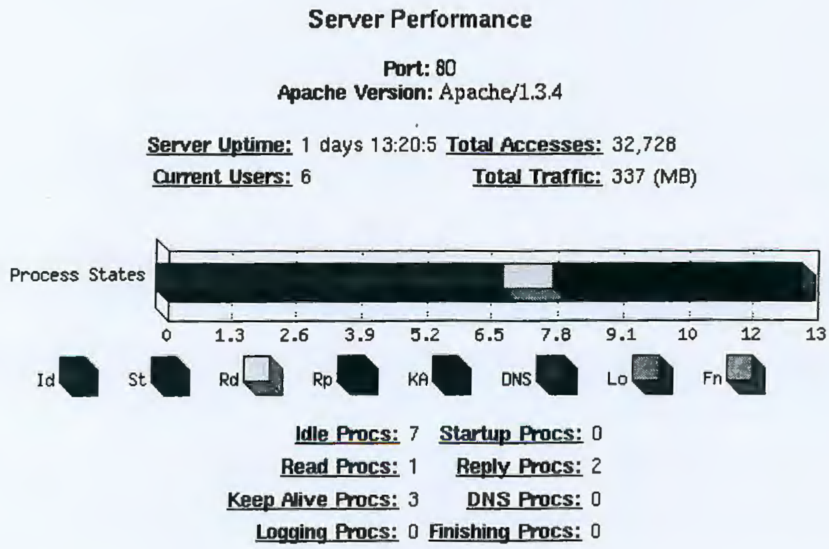

Server Performance

Port: 80
Apache Version: Apache/1.3.4

Server Uptime: 1 days 13:20:5 **Total Accesses:** 32,728
Current Users: 6 **Total Traffic:** 337 (MB)

Idle Procs: 7 **Startup Procs:** 0
Read Procs: 1 **Reply Procs:** 2
Keep Alive Procs: 3 **DNS Procs:** 0
Logging Procs: 0 **Finishing Procs:** 0

Figure 7. AdvantEDGE View Server Performance Summary for Apache

19301
3

4

Using eHealth AIM for Apache

This chapter describes how to use *e*Health AIM for Apache. This module is implemented as a SystemEDGE agent plug-in. After you enable this plug-in in the **sysedge.cf** file and obtain a license for it, it will load automatically at SystemEDGE start time. For more information, refer to "Editing the sysedge.cf File" on page 21 and "Licensing eHealth AIM for Apache" on page 25.

Overview

The *e*Health AIM for Apache plug-in implements additional MIB objects that provide advanced information about the health and availability of the Apache Web server. It can operate with any SNMP-compliant management software such as Concord's *e*Health suite of products, AdvantEDGE View, HP OpenView, and others. If you are using *e*Health AIM for Apache with *e*Health, refer to the *e*Health Web Help for more information about the reports you can generate.

The default configuration settings of *e*Health AIM for Apache enable you to use the advanced self-monitoring capabilities of the SystemEDGE agent in conjunction with *e*Health AIM for Apache.

19300
E

*e*Health AIM for Apache MIB Branch

You can use AdvantEDGE View or another SNMP tool to edit the SystemEDGE configuration file to utilize the MIB objects found in *e*Health AIM for Apache with the process-monitoring, threshold-monitoring, and history-collection features of the SystemEDGE agent. All MIB objects related to *e*Health AIM for Apache exist at object identifier (OID) branch *1.3.6.1.4.1.546.16.3* in the Concord Systems Management MIB. The MIB is defined in the apachemod.asn1 file, which is available in the *e*Health AIM for Apache product installation.

Assigning Entry Rows for the SystemEDGE Self-Monitoring Tables

All SystemEDGE self-monitoring tables require the use of unique row numbers. Each table contains an Index column which acts as a key field to distinguish rows in the table. This section describes the benefits of reserving a block of rows (in the range of 11 to the maximum number of rows in your table) for use by the system or application administrator.

Setting Local Policy

You may choose, as a matter of local policy, to reserve a block of rows for system administration. This policy allows you to define row entries within a reserved block of rows without being concerned that the row might already be taken by another user's entry. In compliance with the local policy, all other users should use row indices that are outside the reserved range when defining user-configured entries.

Reserving Blocks of Rows

By reserving a block of rows, you can define a consistent set of conditions (row entries) to be monitored across all machines such that the same condition is defined in the same row number on each machine. For example, you can use row 3000 in each table to define entries monitoring the number of accesses to the server (apacheServerPerfTotalAccesses). You can

19299
ε

then distribute this configuration to every host so that every system that is running Apache uses row 3000 for monitoring the number of accesses to the server in any of the SystemEDGE agent monitoring tables.

To reserve a block of rows for monitoring Apache:

1. Decide which block of rows you want to reserve for use with monitoring Apache.
2. Use that block of rows to define a set of row entries for each SystemEDGE self-monitoring table. For more information, refer to the chapter on self monitoring in the *eHealth SystemEDGE User Guide.*
3. Distribute configuration file entries to all hosts that are running the Apache Web server and *e*Health AIM for Apache. For more information, refer to the *Automating the Deployment of SystemEDGE and the AdvantEDGE Point Plug-in Modules* white paper.

   NOTE

   As an alternative, you can use this row-number assignment policy with AdvantEDGE View for group-configuration operations.

4. Require end users to avoid your block of rows when defining their own self-monitoring table entries.

Using the SystemEDGE Self-Monitoring Features

This section provides examples of how to use SystemEDGE process, threshold, and history monitoring to monitor the Apache Web server. Add these commands to the sysedge.cf file to enable monitoring of the MIB objects that they specify. Modify these examples as necessary to monitor the MIB objects and thresholds that are relevant for your configuration.

The examples in the following sections present row numbers in the 5000 range; select a row number for your configuration that conforms to local policies. For more information on row assignment, refer to “Assigning Entry Rows for the SystemEDGE Self-Monitoring Tables” on page 46.

The following command, for example, instructs the SystemEDGE agent to monitor whether the Apache process is alive at 30-second intervals and to store the data in row 5000 of the Process Monitoring table:

```
watch process procAlive ’httpd’ 5000 0x0100 30 ’Apache Web Server’ ’’
```

For more information about the syntax for the commands in this section, refer to the *eHealth SystemEDGE User Guide.*

NOTE

Enter the commands throughout this chapter on one line. Do *not* use a carriage return to match the formatting shown here.

Using SystemEDGE Process Monitoring

This section provides an example of how to use the SystemEDGE agent to monitor the availability of a critical Apache process. For more information, refer to the section on process and service monitoring in the *eHealth SystemEDGE User Guide.*

To ensure that the Apache Web server is running, enter the following command in the sysedge.cf file:

```
watch process procAlive ’httpd’ 5000 0x0100 30 ’Apache Web Server’ ’’
```

You must include the 0x0100 flag to force the SystemEDGE agent to monitor the parent process of the Apache process group.

19297

Using SystemEDGE Threshold Monitoring

This section provides examples of how to use the SystemEDGE agent to monitor thresholds for important Apache metrics. Add the commands in the following sections to the sysedge.cf file to monitor thresholds for these MIB objects. For more information, refer to the section on threshold monitoring in the *eHealth SystemEDGE User Guide.*

NOTE

> The thresholds in this section may not be appropriate for your Apache Web server. Select thresholds that are appropriate for your environment.

Monitoring Idle Apache Processes

To monitor the number of idle Apache processes on the server that is running on port 80, enter the following:

```
monitor oid apacheServerPerfCurrentIdleProcs.80 5000 0x0 60
absolute < 10 'Idle Server Processes' ''
```

Monitoring Total Resident Memory Size of an Apache Service

To monitor the total resident memory size of the Apache service that is running on port 80, enter the following:

```
monitor oid apacheFootprintTotalRSS.80 5001 0x0 60 absolute > 50000
'Total Resident Memory' ''
```

Monitoring Total Size of Apache Service Log Files

To monitor the total size of the log files for the Apache service that is running on port 80, enter the following:

```
monitor oid apacheFootprintTotalLogSize.80 5002 0x0 60
absolute > 100000 'Total Log Size' ''
```

19296
E

Monitoring Total Size of Apache Service Document Files

To monitor the total size of the document files for the Apache service that is running on port 80, enter the following:

```
monitor oid apacheFootprintTotalDocSize.80 5003 0x0 60
absolute > 500000 'Total Document Size' ''
```

Using SystemEDGE History Collection

This section provides examples of how to use SystemEDGE history collection to track the value of important Apache metrics over time. Add the commands in the following sections to the sysedge.cf file to collect history for these MIB objects. For more information, refer to the section on history collection in the *eHealth SystemEDGE User Guide.*

NOTE

The number of samples and the interval between samples used in this section may not be appropriate for your Apache system; select values that are suitable for your environment.

Collecting History on Number of Hits to the Server

To collect history on the number of hits to the server on port 8080, enter the following:

```
emphistory 5000 60 apacheServerPerfTotalAccesses.8080 300 'Total Hits
To Port 8080 Server'
```

Collecting History on the Number of Current Active Users

To collect history on the number of current active users that are being maintained by the Apache server on port 8080, enter the following:

```
emphistory 5001 60 apacheServerPerfCurrentUsers.8080 300 'Total Users
On Port 8080 Server'
```

19245
E

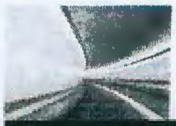

Glossary

Abstract Notation One (ASN.1) A language that describes data types independent of computer structures and representations. For more information refer to ISO International Standard 8824.

Active Server Page (ASP) An HTML file that includes one or more scripts written using VBScript, JavaScript, or ActiveX Data Objects program statements. ASP files (named with the ".asp" suffix) receive a user request and create a customized Web page for the user (usually based on database information).

AdvantEDGE View A Web-based management interface for use with the SystemEDGE agent that enables an administrator to use a Web browser to manage systems and applications.

agent In network management, a program that provides information from a management information base (MIB) for SNMP agents. *e*Health or a network management system (NMS) use the information about managed devices and take corrective action when appropriate.

American Standard Code for Information Interchange (ASCII) The most common format for character representation in computers and the Internet. Characters fit into a single byte. It was developed by the American National Standards Institute (ANSI).

19294
E.

Apache A freely distributed Web server that runs on most UNIX, Linux, and Windows NT operating systems. For details about the Apache server, refer to the Apache Web site, http://www.apache.org.

application A program that performs a specific function for one or more users or for another application program. Types of applications include communication programs, management programs, word processors, databases, and drawing programs.

ASCII *See* American Standard Code for Information Interchange (ASCII).

ASN.1 *See* Abstract Notation One (ASN.1).

ASP *See* Active Server Page (ASP).

buffer A temporary storage area for data. Often implemented as holding areas between the backplane and an interface; data remains in the buffer until it can be transmitted on the interface or processed by the central processing unit (CPU).

central processing unit (CPU) The component within a device that performs the instruction execution for the applications and programs that run on the device. Also referred to as a processor or microprocessor.

CGI *See* Common Gateway Interface (CGI).

client A computer system, usually a desktop computer or laptop, that presents data directly to a user and accepts input. They drive the computing process, supporting local processing and accessing remote servers as needed for data access and analysis.

Also refers to the application software residing on a machine that is used by an end user.

19293
4

Common Gateway Interface (CGI) A server-side interface for Web-based applications that defines how a Web server can exchange data with an application. The Active Server Pages (ASP) and Internet Server Application Program Interface (ISAPI) technologies are alternatives to CGI.

congestion A condition in which the network traffic is greater than the amount that the network can carry. Often causes performance problems and delays on a network.

CPU *See* central processing unit (CPU).

DHCP *See* Dynamic Host Configuration Protocol

DNS *See* domain name system (DNS).

domain name system (DNS) The system that locates and translates Internet domain names such as concord.com into Internet Protocol (IP) addresses. A DNS server is typically a device that translates domain names to IP addresses within your network.

Dynamic Host Configuration Protocol A protocol that enables dynamic allocation of IP addresses so that they can be reused.

eHealth AIM *See* eHealth application insight module.

eHealth application insight module A plug-in (supplementary program) that extends the functionality of the SystemEDGE agent. AIMs add the capability to manage application-specific events, processes, thresholds, and health.

event An occurrence on a system that typically results in a message, such as an SNMP trap, being sent to a configured management system. Common events include system failures, system reboots, exceeded thresholds, or any user-configurable situation that the user wants to identify.

file cache A block of memory that holds frequently or recently used data. A system can read those blocks at memory speed rather than the slower disk access speed.

19292

File Transfer Protocol (FTP) A means for uploading and downloading files on the Internet (the oldest Internet protocol for retrieving files). You can use an FTP client application to request files from or transfer files to an FTP server.

FTP *See* File Transfer Protocol (FTP).

Host Resources MIB A MIB (management information base) that defines a set of objects that are useful for the management of host computers. For example, it defines host storage areas, devices, and file systems. This MIB is defined in RFC 1514.

hostname The name for an individual IP (Internet Protocol) address on a computer. While many computers have only one hostname, some machines, such as network servers have multiple hostnames.

HTML *See* Hypertext Markup Language (HTML).

HTTP *See* Hypertext Transfer Protocol (HTTP).

Hypertext Markup Language (HTML) A programmatic language used for controlling the way that text and images appear when a file is displayed on the World Wide Web.

Hypertext Transfer Protocol (HTTP) An application protocol that defines the set of rules for exchanging files (text, graphics, multimedia, and other files) on the World Wide Web.

I/O *See* input/output (I/O).

Information Technology (IT) A widely-used term to describe all of the technologies used for creating, exchanging, managing, and using information in various forms.

input/output (I/O) Any operation, program, or device that transfers data to or from a computer.

Internet Control Message Protocol (ICMP) A protocol between a server and a gateway to the Internet.

19291
3

Internet Protocol (IP) The method (or protocol) by which packets of information are sent across the Internet. IP defines addressing, error handling, routing, and option codes for data transmission. IP requires no continuing connection between the endpoints that are communicating.

IP *See* Internet Protocol (IP).

IT *See* Information Technology (IT).

management information base (MIB) A formal description of a set of network objects that can be managed using Simple Network Management Protocol (SNMP).

MIB *See* management information base (MIB).

network A collection of computers, printers, routers, switches, and other devices that are able to communicate using a common transmission media such as TCP/IP.

network management system (NMS) An application program usually residing on a computer that manages at least part of a network, including systems and applications. The NMS communicates with agents to monitor network statistics and resources, control network device configuration, and analyze network problems. *See also* agent.

NMS *See* network management system (NMS).

object identifier (OID) a unique identifier of a managed object in a MIB hierarchy. *See also* management information base (MIB).

OID *See* object identifier (OID).

operating system (OS) The program that manages all other programs (applications or application programs) on a computer. Provides the following services: determining the order in which each application runs and the time allotted for that application, managing the sharing of internal memory among multiple applications and handling input to and output from attached hardware devices.

19290

OS *See* operating system (OS).

packet A logical unit of data routed between an origin and a destination on the Internet or any other packet-switched network. On the Internet, the Transmission Control Protocol (TCP) layer of TCP/IP divides a file into packets of manageable size for routing.

packet-switched network A communications network in which data is transferred in small units called packets. Individual packets for a file may travel different routes. When all packets for a file reach their destination, the file is reassembled.

packet-switching A networking technology in which network nodes share bandwidth by sending packets. The same data path can be shared by many users in the network. Packet-switching is widely used throughout the Internet.

page In computers that utilize virtual memory, a unit of data storage. Systems transfer pages of data from disk storage to memory and back again.

On the World Wide Web, a file written using Hypertext Markup Language (HTML) that specifies how text, images, and other multimedia will be presented to the user. A Web site delivers information to the user one page at a time.

paging The process by which a computer moves portions of programs between random access memory and auxiliary storage (on disk).

partition A logical division of a hard disk on a PC that is created so that each partition can have a different operating system or can be used for different purposes (for example, file management or multiple users).

performance threshold The upper limit of acceptable response time.

19289
3

ping An Internet echo message used to confirm the reachability of a network device. An abbreviation for Packet Internet or Inter-Network Groper.

process Typically, an instance of a program or application that is running on a server. Applications can have one or more associated processes.

protocol The set of rules by which the endpoints in a telecommunication connection communicate. The protocol defines the packet format of the transmitted information. On the Internet, common protocols are TCP, IP, HTTP and FTP.

queue In a system, a set of jobs awaiting resources. In a network device such as a router, a collection of packets waiting to be processed or forwarded. Insufficient central processing unit (CPU) speed, memory, or interface speeds can contribute to long queues, and therefore, to delay on the network.

real-time A level of computer responsiveness that an end user would deem as immediate or fast enough to show incremental changes of an external process (for example, to present visualizations of the weather as it constantly changes).

Request For Comments (RFC) The name of the document series regarding Internet design. Most RFCs define protocol specifications such as Telnet and FTP. RFCs are widely available online.

RFC *See* Request For Comments (RFC).

server A program that provides services to other programs in the same and other computers.

Also a computer that performs file storage and application hosting as well as provides computing services to other devices and users on the network. Typically has one or more central processing units (CPUs), disks, interfaces, and storage partitions.

11928 E

Simple Network Management Protocol (SNMP) The network management protocol used almost exclusively in data networks. A method for monitoring and controlling network devices, as well as managing configurations, statistics collection, performance, and security.

SNMP *See* Simple Network Management Protocol (SNMP).

SNMP agent A program such as the SystemEDGE agent that conforms to a management information base (MIB) specification to collect information about managed devices and to take corrective action (using SNMP traps) when appropriate.

SystemEDGE agent Concord's SNMP agent that autonomously monitors system configuration, status, performance, users, applications, file systems, and other critical resources.

Systems Management MIB A set of MIB (management information base) objects that extends the capabilities of the Host Resources MIB. It provides greater visibility into systems and specific information about Windows NT and UNIX systems.

TCP/IP *See* Transmission Control Protocol (TCP) and "Internet Protocol (IP).

threshold *See* performance threshold.

Transmission Control Protocol (TCP) A connection-based protocol used along with the Internet Protocol (IP) to send data in the form of message units between computers over the Internet. While IP is responsible for the actual delivery of the data, TCP is responsible for dividing data into packets at the sending system and constructing the data message from individual packets at the receiving system.

trap A message sent by an SNMP agent to a console or network management system (NMS) to indicate that a threshold has been reached or another user-defined condition has occurred. The SystemEDGE agent defines a number of traps for system and application management.

1986

UDP *See* User Datagram Protocol (UDP).

User Datagram Protocol (UDP) A communications protocol that uses Internet Protocol (IP) to send and receive data and is similar to Transmission Control Protocol (TCP), but provides fewer packet management services.

variable A performance metric for an element. A characteristic or behavior upon which *e*Health gathers data and evaluates the performance of the element. SystemEDGE agents can also monitor local variables to reduce network polls and increase scalability.

Web *See* World Wide Web (WWW, Web).

workstation A powerful computer that is equipped with a fast processor, a large amount of random access memory, and other features such as high-speed graphical rendering that make it suitable for business users such as engineers, graphic designers, and architects.

World Wide Web (WWW, Web) All of the resources on the Internet that use Hypertext Transfer Protocol (HTTP). Users of the Web access information through browser software.

19285

Index

A

AdvantEDGE View 51
  application query for Apache 14
apachemod.asn1 file 20
apachemod.cf file 20
  editing 24
  example 25
apachemod.dll file 20, 22
apachemod.exe file 19
apachemod.pdf file 20
apachemod.so file 20
apachemod-hpux.so file 20
apachemod-hpux-64.so file 20
apachemod-sparcv9.so file 20
application query, Apache 14
assigning entry rows for self-monitoring tables 46

C

caching disk space calculations 24
collecting history
  current active users 50
  hits to the server 50
configuring *e*Health AIM for Apache 21
controlling frequency of disk space calculations 24

D

disk space calculations 24

E

editing
  apachemod.cf file 24
  httpd.conf file 22
  sysedge.cf file 21
*e*Health
  application integration modules 53
  using 16
*e*Health AIM for Apache
  components 12
  files 20
  installation requirements 17
  overview 11
  using 45
enabling
  *e*Health AIM for Apache 21
  mod_info and mod_status 23

19283
3

evaluation copy, licensing 27

F
files installed 20
footprint 39

G
generating a license 27

H
history collection 50
httpd.conf file 22

I
installation
    directory 19
    files 20
    requirements 17
installing *e*Health AIM for Apache
    UNIX systems 18
    Windows systems 19

L
license
    generating 27
    key 29
    sample 33
licensing
    *e*Health AIM for Apache 25
    evaluation copy 27
    host administration 31
Live Health 16
local policy for monitoring tables 46

M
Management Information Base. *See* MIB
MIB
    configuration section 36
    performance section 38
    specification 35
    using 45
mod_info module 22
mod_status module 22
monitoring
    assignment policy for tables 46
    idle Apache processes 49
    multiple Apache servers 24
    service document files 50
    service log files 49
    thresholds 49
    total resident memory size 49

O
obtaining a license 26

P
policy for monitoring tables 46

R
relnotes.txt file 20
requirements for installation 17
reserving rows in monitoring tables 46

S
server
    configuration 36
    performance 41
setting policy for monitoring tables 46

SNMP 12
sysedge.cf file 22
sysedge.lic file 26, 29
SystemEDGE agent 11, 13
  configuring 21
  history capability 50
  monitoring processes 46
  monitoring thresholds 49

T

temporary license 27
threshold monitoring 49

U

using *e*Health AIM for Apache
  configuring self monitoring 45
  overview 13
  with AdvantEDGE View 14
  with *e*Health 16
  with Live Health 16

19281
E

CONTACT CONCORD COMMUNICATIONS AT:

CONCORD COMMUNICATIONS, INC.
NORTH AMERICA
600 NICKERSON ROAD
MARLBORO, MASSACHUSETTS 01752
800-851-8725
P 508-460-4646
F 508-481-9772

CONCORD COMMUNICATIONS EUROPE
DELFTECHPARK 26
2628 XH DELFT
THE NETHERLANDS
P +31 (0) 15 2600 930
F +31 (0) 15 2600 929

CONCORD COMMUNICATIONS ASIA PACIFIC
LEVEL 7, 53 WALKER STREET
NORTH SYDNEY NSW 2060
AUSTRALIA
P 61-2-9965-0600
F 61-2-9929-0411

FRANCE: +33 (0) 1 4692 2420
GERMANY +49 (0) 89 944 90 105
UK: 00800 3283 2888

JAPAN 813-5778-7629
SINGAPORE: 65-4309533

CONCORD.COM

19280
E

ADVANTEDGE™

FOR MICROSOFT® EXCHANGE

User Guide

Release 1.1 or Later

09-16030-002
April 2001

19279
3

Important Notice

Application Health, System Health, Live Health, SystemEDGE, AdvantEDGE, *e*Health, Infrasense, MIB Manager, State-of-the-World, and the FirstSense logo are trademarks of Concord Communications, Inc.

The Total View, Concord, the Concord logo, Empire Technologies, Network Health, TRAKKER, FirstSense, and We See It Happening are registered trademarks of Concord Communications, Inc.

Microsoft, Active Directory, Windows, Windows 2000, and Windows NT are trademarks or registered trademarks of Microsoft Corporation.

© Copyright Bigelow and Holmes 1986, 1985. Lucida is a registered trademark of Bigelow & Holmes. Sun Microsystems Inc., AT&T, and Bigelow & Holmes make no representations about the suitability of the source code for any purpose. It is provided "as is" without express or implied warranty of any kind.

General Notice: Some of the product names used herein have been used for identification purposes only and may be trademarks of their respective companies.

Proprietary Notice

The information and descriptions contained herein are the property of Concord Communications, Inc. Such information and descriptions may not be copied, disseminated, or distributed without the express written consent of Concord Communications, Inc. Concord Communications, Inc., assumes no responsibility for any inaccuracies that may appear in this document. Concord Communications, Inc., reserves the right to improve its products and change specifications at any time without notice.

U. S. Government Restricted Rights

Use, reproduction, and disclosure by the U.S. Government are subject to the restrictions set forth in FAR §52.227-19 (c) (1) and (2) and DFARS §252.227-7013 (c) (1) (ii).

Patent Information

U. S. Patent 5,615,323
Patents Pending

19278
ε

Table of Contents

About This Guide

How This Guide Is Organized v
Conventions Used in This Guide vi
Contact Information vii

Chapter 1
Introduction

Features of AdvantEDGE for Microsoft Exchange 1-1
Microsoft Exchange Architecture. 1-2
Using AdvantEDGE for Microsoft Exchange. 1-3
    Using AdvantEDGE for Microsoft Exchange with eHealth 1-4
    Using AdvantEDGE for Microsoft Exchange with
    Live Health. 1-4

Chapter 2
Installing AdvantEDGE for Microsoft Exchange

Installation Requirements 2-1
Installing the Software. 2-1
AdvantEDGE for Microsoft Exchange Files 2-2
Configuring AdvantEDGE for Microsoft Exchange 2-3
Licensing AdvantEDGE for Microsoft Exchange 2-3
    Obtaining a License. 2-4
    Generating the License 2-4
    Sample License File. 2-6

1277
E

Chapter 3
Using the AdvantEDGE for Microsoft Exchange MIB

Configuration Section 3-2
Server Configuration. 3-2
Connector Configuration 3-4
Performance. 3-7
Exchange Footprint 3-8
Server Performance 3-10
Connector Performance 3-11

Chapter 4
Using AdvantEDGE for Microsoft Exchange

Editing the SystemEDGE Configuration File 4-1
Assigning Entry Rows in the SystemEDGE Self-Monitoring Tables . . . 4-2
Setting Local Policy 4-2
Reserving Blocks of Rows 4-2
Using the SystemEDGE Self-Monitoring Features 4-3
Using SystemEDGE Process Monitoring 4-4
Using SystemEDGE Threshold Monitoring 4-5
Using SystemEDGE History Collection 4-7
Using SystemEDGE Windows NT Event Monitoring 4-9

Index

19276
ε

About This Guide

The *AdvantEDGE for Microsoft Exchange User Guide* provides instructions for installing and using AdvantEDGE for Microsoft Exchange for Windows NT x86 and Windows 2000 systems.

NOTE

AdvantEDGE for Microsoft Exchange supports Exchange 5.5 and Exchange 2000.

This guide is intended for the person responsible for installing and configuring AdvantEDGE for Microsoft Exchange. This version supports AdvantEDGE for Microsoft Exchange Release 1.1 or later, and the SystemEDGE Agent Release 4.0, Patchlevel 3 and later.

To use AdvantEDGE for Microsoft Exchange, you should be familiar with the Microsoft Exchange application and the Concord SystemEDGE agent. Refer to Microsoft documentation and the *SystemEDGE Agent User Guide* for more information.

NOTE

Unless otherwise specified, the content of this guide applies to both Exchange 5.5 and Exchange 2000. In areas where there are differences, this guide specifies to which version the content applies.

How This Guide Is Organized

This guide is organized as follows:

- Chapter 1, "Introduction," provides an overview of AdvantEDGE for Microsoft Exchange and its features.
- Chapter 2, "Installing AdvantEDGE for Microsoft Exchange," explains how to install, configure, and license the AdvantEDGE for Microsoft Exchange software on a host system.

19275
3

- Chapter 3, "Using the AdvantEDGE for Microsoft Exchange MIB," describes the information that is available through the AdvantEDGE for Microsoft Exchange MIB.
- Chapter 4, "Using AdvantEDGE for Microsoft Exchange," explains how to configure and use AdvantEDGE for Microsoft Exchange in your host environment.

Conventions Used in This Guide

This section describes conventions used in this guide.

System Root Text Convention

This guide uses the system root text convention that is used by Microsoft to denote the Windows NT root directory:

```
%SystemRoot%\system32\
```

where `%SystemRoot%` is `C:\winnt` for Windows NT 4.0 and Windows 2000.

Conventions for Commands

The following paragraph shows a sample command. Due to space limitations in this guide, some commands wrap from one line to the next. Disregard these line breaks, and **enter each command as one line**. Otherwise, your command syntax will be incorrect. For example, when you see a command such as the following:

```
watch process procAlive 'emsmta|EMSMTA' 5000 0x0 30
   'Exchange MTA' ''
```

You **must** enter the command on one line, as shown here:

```
watch process procAlive 'emsmta|EMSMTA' 5000 0x0 30 'Exchange MTA' ''
```

19274
ε

Contact Information

If you need any assistance with the SystemEDGE agent or the AdvantEDGE for Microsoft Exchange Point module, contact Customer Support, using one of the following methods:

Phone:	(888) 832-4340 (for calls from the USA and Canada) (508) 303-4300 (for calls from other countries)
Fax:	(508) 303-4343
E-mail:	support@concord.com
Web site:	http://www.concord.com
Licensing:	http://license.concord.com

10273
E

K1275
E

1

Introduction

This chapter provides an overview of the AdvantEDGE Point module for Microsoft Exchange and of the Microsoft Exchange architecture.

NOTE

Unless otherwise stated, the term *Exchange,* as used throughout this guide and the Management Information Base (MIB) specification, refers to the Microsoft Exchange application in its entirety, which encompasses all the core services, connectors, and optional components.

Features of AdvantEDGE for Microsoft Exchange

AdvantEDGE for Microsoft Exchange is a plug-in for the SystemEDGE agent. It enables information technology (IT) operators to monitor the performance and availability of Microsoft Exchange. Microsoft Exchange is a groupware application that enables communication and collaborative work. At its core is an e-mail routing, distribution, and storage facility. Exchange serves as the e-mail backbone for many corporations. Therefore, monitoring its health and availability is crucial to ensuring the smooth functioning of today's corporate information infrastructure.

The AdvantEDGE for Microsoft Exchange Point module makes important information about Microsoft Exchange available to management software through the SystemEDGE agent and Simple Network Management Protocol (SNMP). The SystemEDGE agent's self-monitoring capabilities enable the monitoring of important Exchange metrics, processes, and services, as well as the sending of SNMP traps when exceptions or exception conditions occur.

This release of AdvantEDGE for Microsoft Exchange supports Exchange 5.5 and Exchange 2000.

NOTE

This document does not explain how to install, administer, or use Microsoft Exchange. For help with Microsoft Exchange, refer to your Microsoft documentation.

Microsoft Exchange Architecture

Microsoft Exchange is a complex piece of software containing many components. At the core of Microsoft Exchange are the following services:

- Mail transfer agent (MTA), which is responsible for routing e-mail messages to and from users.
- Directory service (DS), which maintains information about recipients, mailboxes, public and private folders, mailing lists, and other distribution lists (Exchange 5.5 only; Exchange 2000 uses Active Directory).
- Information store (IS), which serves as the repository of all messages on the Exchange server, and is composed of private and public areas.
- System attendant, which maintains the Exchange application's database and directory integrity, and on which all other Exchange services depend.
- Connectors, which transfer messages between sites, organizations, and non-native Exchange e-mail formats (for example, Internet, Lotus Notes, and Microsoft Mail). Several connectors are available from Microsoft and from third parties.
- Internet Information Services (IIS), which provides SMTP, IMAP, and POP services for Exchange 2000.

Figure 1-1 shows the relationships between the components of Microsoft Exchange 5.5. The architecture of Exchange 2000 is similar.

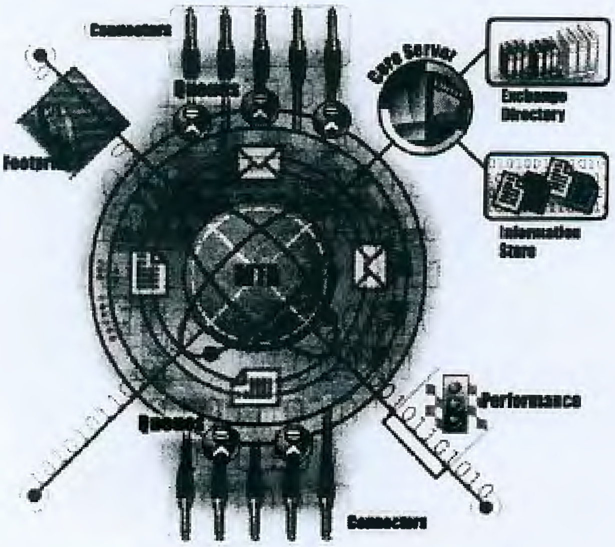

Figure 1-1: Microsoft Exchange 5.5 Architecture

Because Microsoft Exchange is a complex application, monitoring Exchange is more complex than ensuring that a single process or Windows NT service is up and running. The many components that make up Exchange can function properly only through the availability of a variety of configuration parameters, settings, processes, Windows NT services, queues, and system resources.

Using AdvantEDGE for Microsoft Exchange

AdvantEDGE for Microsoft Exchange can provide you with the tools and information that you need to monitor the health and availability of your Exchange server. It works with the SystemEDGE agent to closely manage the Microsoft Exchange application, providing real-time fault detection and automatically correcting problems, if necessary. You can use AdvantEDGE for Microsoft Exchange with any SNMP-compliant management software, including Concord's *e*Health suite of products, AdvantEDGE View, HP OpenView, and others.

You can use this plug-in with the SystemEDGE agent to perform the following tasks:

- Detect failed Exchange services (such as the Information Store or Message Transfer Agent) and restart them automatically.
- Alert an administrator when message queues become dangerously large, indicating potential security violations through SystemEDGE intelligent self-monitoring.
- Watch for Information Stores and disk partitions that are nearing capacity.
- Automatically delete temporary files when a threshold is reached to free up disk space and ensure continuous availability and performance.
- Detect types of Windows NT Events and forward them as SNMP traps to your network management system (NMS).

19272
E

Using AdvantEDGE for Microsoft Exchange with *e*Health

You can use AdvantEDGE for Microsoft Exchange and the SystemEDGE agent with the *e*Health product suite to provide the historical data for long-term trending analysis and capacity planning. With *e*Health – Application Assessment, you can run At-a-Glance, Trend, Top N, or MyHealth reports for the following types of variables:

- Amount of Central Processing Unit (CPU), total memory, and disk space that the Exchange application is using
- Size of the MTA and DS logs (Exchange 5.5 only)
- Amount of data, and number and type (public or private) of messages being processed by the MTA service
- Number of messages waiting to be processed by the MTA service
- Number of users
- SMTP traffic

For more information about the variables that you can monitor and reports you can run when you integrate AdvantEDGE for Microsoft Exchange with *e*Health, refer to the *e*Health Web Help.

Using AdvantEDGE for Microsoft Exchange with Live Health

You can also use AdvantEDGE for Microsoft Exchange and the SystemEDGE agent with Live Health for real-time detection of potential problems. Live Health applies intelligent algorithms to the data, resulting in precise assessments of application health and performance. For more information about how Live Health can detect "brownouts" and service delays across applications, systems, and networks, refer to the Live Health Web Help.

19271
ε

2

Installing AdvantEDGE for Microsoft Exchange

This chapter explains how to install, configure, and license AdvantEDGE for Microsoft Exchange.

Installation Requirements

Before you install AdvantEDGE for Microsoft Exchange, you must first install, license, and configure the SystemEDGE agent Release 4.0, Patchlevel 3 or later. Refer to the *SystemEDGE Agent User Guide* for more information. Also, you must be running Windows NT 4.0 or Windows 2000. For more information, refer to your Microsoft documentation.

Installing the Software

AdvantEDGE for Microsoft Exchange is distributed as a self-extracting executable named **xchgmod.exe**.

Follow these steps to install AdvantEDGE for Microsoft Exchange:

1. Log on to the Windows NT system as administrator.
2. Click **Start**.
3. Select **Programs → Command Prompt**.
4. Insert the CD containing the Concord software distributions into the CD-ROM drive.

   Windows NT or Windows 2000 automatically mounts the drive using the CD-ROM drive's corresponding drive letter. The particular drive letter is specific to your system and depends on the number and types of disks attached to your system.

5. Determine which directory you want to use as the installation directory for AdvantEDGE for Microsoft Exchange. If the SystemEDGE agent is installed in C:\sysedge, the recommended installation directory is C:\sysedge\plugins.
6. Run the self-extracting executable by entering the following at the command prompt, where *D:* is the CD-ROM drive for your system, and *C:sysedge\plugins* is the installation directory:

   ```
   D:\xchgmod\ntx86\xchgmod.exe -dir C:\sysedge\plugins
   ```

   The -dir option instructs the self-extracting executable to create the intended subdirectory hierarchy that is described throughout this guide. It then places the distribution in an xchgmod subdirectory within the specified target directory (for example, C:\sysedge\plugins).

NOTE

You cannot run xchgmod.exe directly from the CD-ROM.

AdvantEDGE for Microsoft Exchange is now installed.

AdvantEDGE for Microsoft Exchange Files

Table 2-1 describes the files that are installed during the AdvantEDGE for Microsoft Exchange installation.

Table 2-1: Files Installed by AdvantEDGE for Microsoft Exchange

File Name	Description
xchgmod.dll	AdvantEDGE for Microsoft Exchange dynamic link library (DLL) module for Windows NT and Windows 2000
xchgmod.pdf	*AdvantEDGE for Microsoft Exchange User Guide*
xchgmod.asn1	AdvantEDGE for Microsoft Exchange MIB specification
examples	AdvantEDGE for Microsoft Exchange monitoring examples
relnotes.txt	Release notes for AdvantEDGE for Microsoft Exchange

19269
e

Configuring AdvantEDGE for Microsoft Exchange

The SystemEDGE agent reads the configuration file sysedge.cf and uses the *sysedge_plugin* keyword to specify which AdvantEDGE Point modules to load at system initialization. By default, the SystemEDGE agent does not load any plug-ins at initialization time, but you can configure the agent to load any AdvantEDGE Point modules that you have installed by editing the sysedge.cf file as follows.

To configure the SystemEDGE agent to start AdvantEDGE for Microsoft Exchange, provide the complete path name to xchgmod.dll, the AdvantEDGE for Microsoft Exchange DLL. The actual path depends on the location you selected when installing AdvantEDGE for Microsoft Exchange files. For example, enter this command if you installed the files in the C:\sysedge\plugins\xchgmod directory:

```
sysedge_plugin C:\sysedge\plugins\xchgmod\xchgmod.dll
```

For more information about the sysedge.cf file, refer to the *SystemEDGE Agent User Guide.*

Licensing AdvantEDGE for Microsoft Exchange

Like the SystemEDGE agent, AdvantEDGE for Microsoft Exchange utilizes a *host-based* license method. Copies of AdvantEDGE for Microsoft Exchange can run only on systems that possess a valid license key. This license is separate from the one used for the SystemEDGE agent.

The first time that you attempt to start the SystemEDGE agent after installing AdvantEDGE for Microsoft Exchange, the agent displays a message that says that a valid license was not found for AdvantEDGE for Microsoft Exchange. It then provides you with a *public key* that is used to generate a permanent license key for your host machine.

A license key is made up of four space-separated, 8-character sequences, totaling 32 characters. The AdvantEDGE for Microsoft Exchange license is stored in the sysedge.lic file, the same file that is used for SystemEDGE agent licenses. Refer to the sample license file on page 2-6.

19268
ε

Obtaining a License

To obtain a license, you can do any of the following:

- Run the Concord-supplied licenseutil.pl script.
- Run the licenseme.exe license utility.
- Use the AdvantEDGE View licensing procedure, which is based on SNMP traps. For more information, refer to the AdvantEDGE View Web Help.
- Send an e-mail request to license@concord.com and place the returned license key in the appropriate license file.

NOTE

Always include the Customer ID and user name in license requests that you send through e-mail.

- Complete the online license form through the Internet, as described in the next section, "Generating the License".

For more information about licensing, refer to the *SystemEDGE Agent User Guide* and the *Automating the Licensing of SystemEDGE and AdvantEDGE Point Plug-in Modules* white paper.

NOTE

If you are using an evaluation copy of AdvantEDGE for Microsoft Exchange, you must request a temporary license that will enable it to operate during the evaluation period.

Generating the License

This section describes how to generate the license using the Web-based license form. For Windows NT and Windows 2000, the setup program generates the licensing information for your system.

1. Run the SystemEDGE agent setup command to request licensing information by entering the following at the command prompt:

   ```
   sysedge\setup -l
   ```

   The setup program displays a message similar to the following:

```
SystemEDGE Version 4.0 Patchlevel 3
Copyright 2001 by Concord Communications, Inc.
Please contact Concord Communications, Inc. to obtain a license
http://www.concord.com/support, Email: license@concord.com
Provide this: sysedge neptune NTx86 4.0 346561363366b19c 4.0 Patchlevel 3
```

19267

2. Fill out the online Web-based license form that is available from Concord's licensing Web server at the following URL:

   ```
   http://license.concord.com
   ```

   NOTE

   You must supply a user name and password to access the license form.

   The license form asks you to supply the following information:

   - Customer ID
   - Name
   - E-mail address
   - Software version number (4.0 in the example above)
   - Patchlevel (3 in the example above)
   - System name (neptune in the example above)
   - Operating system name (NTx86 in the example above)
   - Version (4.0 in the example above)
   - System identifier (346561363366b19c in the example above)

   NOTE

   When you are licensing AdvantEDGE for Microsoft Exchange, select **xchgmod** as the product on the licensing form.

   After you submit the license request, the Concord Web server generates a license and displays it to your Web browser. It also e-mails the license to the contact person in your organization.

3. Copy the generated license key into the sysedge.lic file in the system32 subdirectory (for example, C:\winnt\system32), and then save the file.

   The license key is case sensitive. Copy it exactly as it appears. If possible, use your system's cut-and-paste facility instead of typing it by hand. If you are entering the license key by hand, be careful not to confuse characters such as the letters l and I and the number 1, or the letter O and the number 0.

4. Stop and restart the Windows NT Master agent by entering these commands at the command prompt:

   ```
   net stop snmp
   net start snmp
   ```

19266
E

The AdvantEDGE for Microsoft Exchange Point module is now licensed and ready to use.

Sample License File

The following is a sample SystemEDGE agent license file. A pound character (#) in column 1 indicates that the entire line is a comment.

```
# license file for SystemEDGE Agent
# Empire Technologies, Inc.
# A Concord Communications Company
# http://www.concord.com
#
# file /etc/sysedge.lic or %SystemRoot%\system32\sysedge.lic
# A valid license key has four parts of 8 characters per part
# parts are separated by space(s) with one license key per line
# sysedge jupiter NTx86 4.0 807cb1da007cb1da 4.0
e13311d3 0F2a7cb1 abC512dc fF8C923a

# xchgmod jupiter NTx86 4.0 807cb1da007cb1da 4.0
a7943fde 098a87ij a4kiuf39 afafEkj4
```

19265
E

3

Using the AdvantEDGE for Microsoft Exchange MIB

This chapter outlines the organization and content of the Concord Communications MIB for Microsoft Exchange. The MIB specification (xchgmod.asn1) defines a collection of objects for monitoring and managing Microsoft Exchange. You must configure the SystemEDGE agent to monitor the MIB objects that are relevant for your configuration. For more information, refer to Chapter 4, "Using AdvantEDGE for Microsoft Exchange."

Figure 3-1 shows the organization of the AdvantEDGE for Microsoft Exchange MIB.

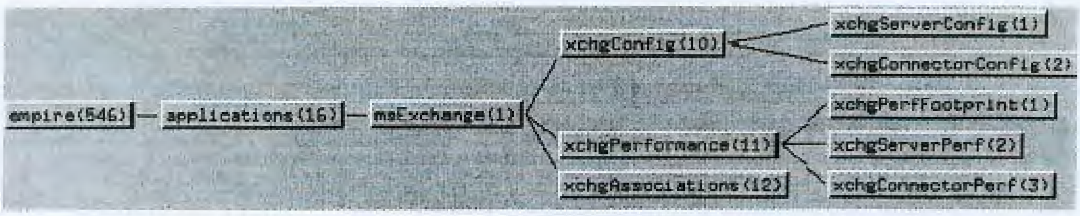

Figure 3-1: AdvantEDGE for Microsoft Exchange MIB

The MIB is organized into broad sections for configuration and performance. Within those broad sections are subsections for connector configuration and performance, and for core server configuration and performance. Within the performance section, a footprint section defines MIB objects that convey how much of the underlying system's resources are consumed by the Microsoft Exchange application.

The following sections define important MIB objects from the Exchange MIB. This chapter defines all sections of the AdvantEDGE for Microsoft Exchange MIB, but it does **not** define all of the MIB objects. For a complete list of MIB objects, refer to the AdvantEDGE for Microsoft Exchange MIB Specification (xchgmod.asn1).

NOTE

Unless otherwise noted, these MIB objects are supported for both Exchange 5.5 and Exchange 2000.

Configuration Section

The Configuration section of the AdvantEDGE for Microsoft Exchange MIB contains configuration parameters and settings that are important for streamlining the health and performance of your Exchange server. It also includes configuration information about core servers and connectors.

Server Configuration

The server configuration MIB group contains configuration parameters, process identifiers (IDs), and version and build numbers, as well as log and database locations. Table 3-1 defines important Server Configuration parameters.

Table 3-1: Selected MIB Objects – Exchange Server Configuration Group (Page 1 of 2)

MIB Object	Description
xchgVersion	Exchange version.
xchgBuildNumber	Exchange build number.
xchgInstallLocation	Location where Exchange is installed.
xchgStoreBuffers	Number of Exchange storage buffers configured. (Exchange 5.5 only)
xchgMinStoreThreads	Minimum number of information store (IS) threads. (Exchange 5.5 only)
xchgMaxStoreThreads	Maximum number of IS threads. (Exchange 5.5 only)
xchgPubStoreFile	Filename of the public IS.
xchgPrivStoreFile	Filename of the private IS.

19263
E

Table 3-1: Selected MIB Objects – Exchange Server Configuration Group (Page 2 of 2)

MIB Object	Description
xchgISWorkDir	IS working directory.
xchgMTARunDir	Mail Transfer Agent (MTA) run directory where temporary and working files are stored.
xchgMTADBPath	Directory containing the MTA database file(s).
xchgDSDBFile	Directory database filename. (Exchange 5.5 only)
xchgDSWorkDir	Exchange working directory where temporary and working files are stored. (Exchange 5.5 only)
xchgSApid	Process ID of the Exchange system attendant.
xchgISpid	Process ID of the Exchange information store.
xchgMTApid	Process ID of the Exchange MTA.
xchgDSpid	Process ID of the Exchange Directory.
xchgCoreIISPID	Process ID of the core IIS service, which provides SMTP, IMAP4, POP3, NNTP, and the core message routing engine. (Exchange 2000 only)

Figure 3-2 shows an example of an AdvantEDGE View core server status for Microsoft Exchange 5.5.

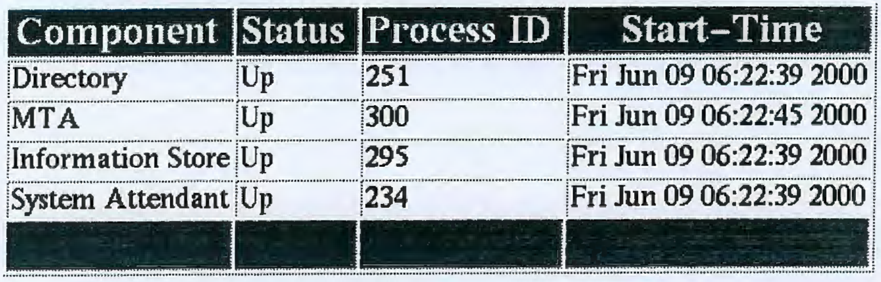

Component	Status	Process ID	Start-Time
Directory	Up	251	Fri Jun 09 06:22:39 2000
MTA	Up	300	Fri Jun 09 06:22:45 2000
Information Store	Up	295	Fri Jun 09 06:22:39 2000
System Attendant	Up	234	Fri Jun 09 06:22:39 2000

Figure 3-2: Core Server Listing for Exchange 5.5

19262
E.

Figure 3-3 shows an example of an AdvantEDGE View core server status for Microsoft Exchange 2000.

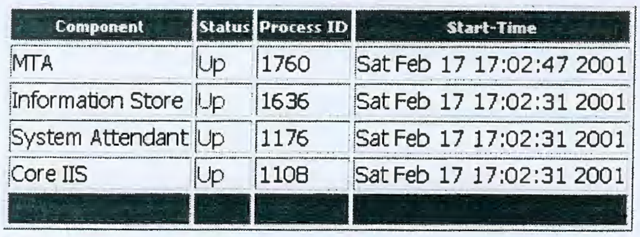

Component	Status	Process ID	Start-Time
MTA	Up	1760	Sat Feb 17 17:02:47 2001
Information Store	Up	1636	Sat Feb 17 17:02:31 2001
System Attendant	Up	1176	Sat Feb 17 17:02:31 2001
Core IIS	Up	1108	Sat Feb 17 17:02:31 2001

Figure 3-3: Core Server Listing for Exchange 2000

Connector Configuration

The Connector Configuration MIB group contains the configuration parameters, process IDs, and installation status of the various Exchange connectors. Table 3-2 defines important Connector Configuration parameters.

Table 3-2: Selected MIB Objects – Exchange Connector Configuration Group (Page 1 of 2)

MIB Object	Description
xchgMMCInstalled	Indicates whether MS Mail connector is installed.
xchgMMCpid	Provides the process ID of the MS Mail connector.
xchgCCMCInstalled	Indicates whether Lotus cc:Mail connector is installed.
xchgCCMCpid	Provides the process ID of the cc:Mail connector.
xchgIMSInstalled	Indicates whether the Internet Mail connector is installed. On Exchange 2000, this object indicates whether the IIS/SMTP service is installed.
xchgIMSpid	Provides the process ID of the Internet Mail connector. On Exchange 2000, this object reports the PID of the IIS/SMTP service.
xchgKMSInstalled	Indicates whether Key Management Service connector is installed.
xchgKMSpid	Provides the process ID of the Key Management Service connector.

1926 1
E

Table 3-2: Selected MIB Objects – Exchange Connector Configuration Group (Page 2 of 2)

MIB Object	Description
xchgNEWSInstalled	Indicates whether the Microsoft Exchange USENET/News connector is installed. On Exchange 2000, this object reports on whether the IIS/NNTP service is installed.
xchgNEWSpid	Provides the process ID of the USENET/News connector. On Exchange 2000, this object reports the PID of the IIS/NNTP service.
xchgIMAP4Installed	Indicates whether the Exchange 2000 IIS/IMAP 4 service is installed. (Exchange 2000 only)
xchgPOP3Installed	Indicates whether the Exchange 2000 IIS/POP3 service is installed. (Exchange 2000 only)
xchgRouteInstalled	Indicates whether the Exchange 2000 IIS/RoutingEngine service is installed. (Exchange 2000 only)
xchgSRSInstalled	Indicates whether the Exchange 2000 Site Replication Service is installed. SRS enables Exchange 2000 to emulate Exchange 5.5 directory services. (Exchange 2000 only)
xchgSRSPID	Provides the process ID of the Exchange 2000 Site Replication Service.

19260
E

Figure 3-4 shows an example of an AdvantEDGE View connector configuration status for Microsoft Exchange 5.5.

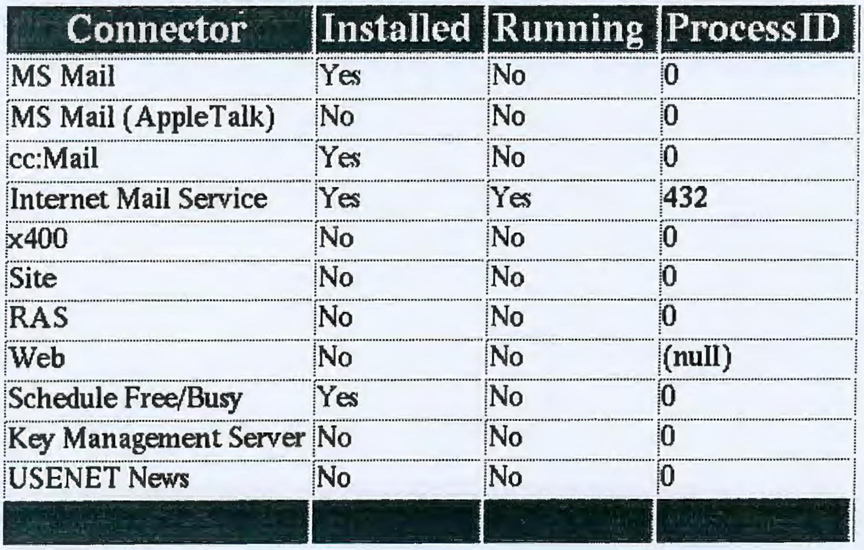

Connector	Installed	Running	ProcessID
MS Mail	Yes	No	0
MS Mail (AppleTalk)	No	No	0
cc:Mail	Yes	No	0
Internet Mail Service	Yes	Yes	432
x400	No	No	0
Site	No	No	0
RAS	No	No	0
Web	No	No	(null)
Schedule Free/Busy	Yes	No	0
Key Management Server	No	No	0
USENET News	No	No	0

Figure 3-4: Exchange 5.5 Connector Configuration Listing

19259
ε

Figure 3-5 shows an example of an AdvantEDGE View connector configuration status for Microsoft Exchange 2000.

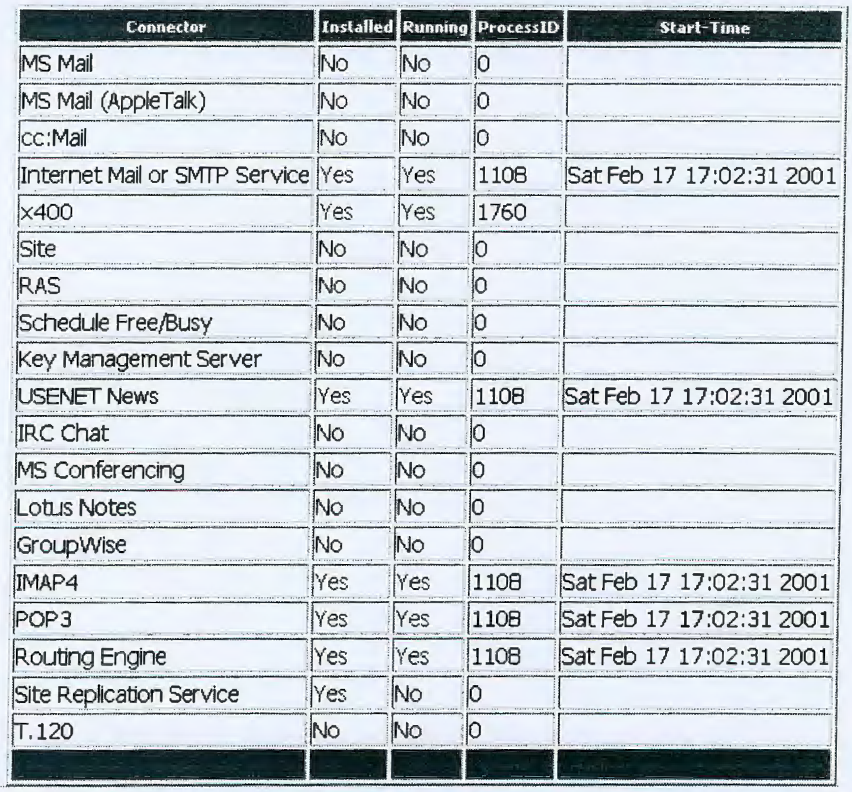

Connector	Installed	Running	ProcessID	Start-Time
MS Mail	No	No	0	
MS Mail (AppleTalk)	No	No	0	
cc:Mail	No	No	0	
Internet Mail or SMTP Service	Yes	Yes	1108	Sat Feb 17 17:02:31 2001
x400	Yes	Yes	1760	
Site	No	No	0	
RAS	No	No	0	
Schedule Free/Busy	No	No	0	
Key Management Server	No	No	0	
USENET News	Yes	Yes	1108	Sat Feb 17 17:02:31 2001
IRC Chat	No	No	0	
MS Conferencing	No	No	0	
Lotus Notes	No	No	0	
GroupWise	No	No	0	
IMAP4	Yes	Yes	1108	Sat Feb 17 17:02:31 2001
POP3	Yes	Yes	1108	Sat Feb 17 17:02:31 2001
Routing Engine	Yes	Yes	1108	Sat Feb 17 17:02:31 2001
Site Replication Service	Yes	No	0	
T.120	No	No	0	

Figure 3-5: Exchange 2000 Connector Configuration Listing

Performance

The Performance section of the Exchange MIB contains performance data that is necessary for capacity planning and trend analysis, as well as real-time performance and availability monitoring. The Performance group is divided into several subgroups for footprint data (page 3-8), server performance (page 3-10), and connector performance (page 3-11).

14258
E

Exchange Footprint

The Exchange Footprint group provides information about the Exchange CPU, memory, and disk resource consumption, more commonly called its *footprint*. Long-term trending analysis of footprint information is useful for anticipating and avoiding email problems due to resource exhaustion. Footprint information can also be monitored in real time to detect and correct temporary resource exhaustion due to viruses, security incidents, and hardware failures. Table 3-3 defines important Footprint metrics.

Table 3-3: Selected MIB Objects – Exchange Footprint Group

MIB Object	Description
xchgCPUTime	Total accumulated central processing unit (CPU) time for Exchange.
xchgPercentCPU	Percentage of CPU, over the last interval, used by Exchange.
xchgTotalRSS	Total real memory currently in use by Exchange.
xchgPercentMEM	Percentage of real memory currently in use by Exchange.
xchgDirSize	Current size of the Exchange directory. (Exchange 5.5 only)
xchgPrivStoreSize	Current size of the private IS.
xchgPubStoreSize	Current size of the public IS.
xchgTotalDiskSize	Estimate of the current total disk space used by Exchange.
xchgTotalThreads	Total number of system threads used by Exchange.

The following figures show sample footprints for a live Exchange application that is serving a medium-sized company. They represent real data collected from live Exchange servers and displayed in AdvantEDGE View reports.

11925 7
3

Figure 3-6 shows a sample AdvantEDGE View footprint for Exchange disk usage.

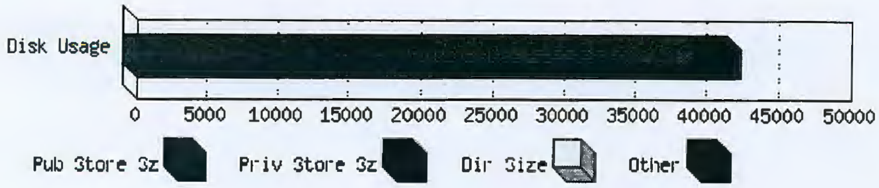

Figure 3-6: Exchange Disk Usage Footprint

Figure 3-7 shows a sample AdvantEDGE View footprint for Exchange memory usage (resident set size [RSS]).

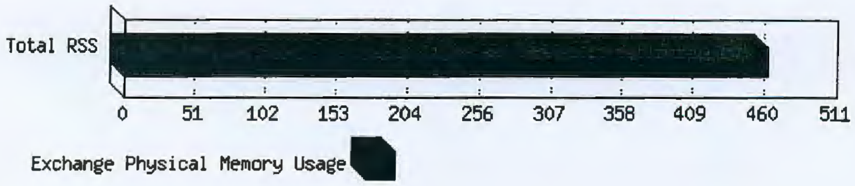

Figure 3-7: Exchange Memory Usage (RSS) Footprint

Figure 3-8 shows a sample AdvantEDGE View footprint summary for Exchange.

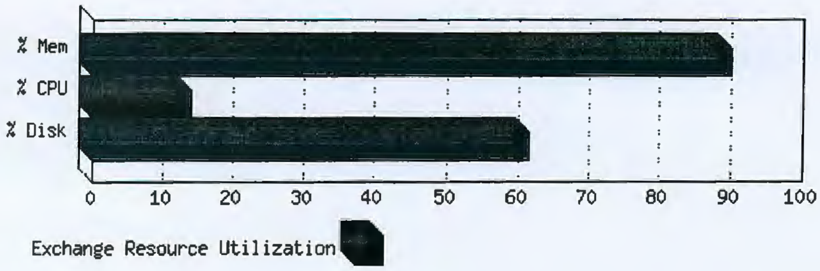

Figure 3-8: Exchange Footprint Summary

11256
ε

Server Performance

The Server Performance group provides performance metrics and counters for the core Exchange server including the information store, directory, MTA, and system attendant. These metrics include those useful for real-time management and longer-term capacity planning and trend analysis. Table 3-4 defines important Server Performance metrics.

Table 3-4: Selected MIB Objects – Exchange Server Performance Group (Page 1 of 2)

MIB Object	Description
xchgStoreUserCount	Current number of information store users/connections.
xchgMTAWorkQueueLen	Current number of messages waiting to be processed by the MTA.
xchgMTAAssoc	Current number of MTA-to-MTA associations.
xchgMTAMessages	Total number of messages sent and received by the MTA.
xchgISPubInMessages	Total number of public messages submitted to clients.
xchgISPubOutMessages	Total number of public messages delivered to recipients.
xchgISPubSendQueLen	Current length of the public message send queue.
xchgISPrivSendQueLen	Current length of the private message send queue.
xchgISPrivInMessages	Total number of private messages submitted to clients.
xchgISPrivOutMessages	Total number of private messages delivered to recipients.
xchgDirABbrowse	Number of address book browses processed by the Microsoft Exchange directory service. (Exchange 5.5 only; Exchange 2000 uses the Active Directory service.)
xchgDirABreads	Number of address book browses reads by the Microsoft Exchange directory service. (Exchange 5.5 only; Exchange 2000 uses the Active Directory service.)

1925
3

Table 3-4: Selected MIB Objects – Exchange Server Performance Group (Page 2 of 2)

MIB Object	Description
xchgDirEXDSReads	Number of extended directory service client reads processed by this Exchange service. (Exchange 5.5 only; Exchange 2000 uses the Active Directory service.)
xchgDirReplUpdates	Number of replication updates processed by this Exchange server. (Exchange 5.5 only; Exchange 2000 uses the Active Directory service.)
xchgDirThreads	Number of directory threads currently allocated. (Exchange 5.5 only; Exchange 2000 uses the Active Directory service.)

Connector Performance

The Connector Performance group provides performance metrics and counters for Exchange connectors including the Internet Mail connector, Lotus Notes cc:Mail, and others. These metrics include those useful for real-time management and longer-term capacity planning and trend analysis. Table 3-5 defines important Connector Performance metrics.

Table 3-5: Selected MIB Objects – Exchange Connector Performance Group

MIB Object	Description
xchgMMCMTAMsgs	Total number of messages moved through the Microsoft Mail Connector.
xchgIMSInQueLen	Number of Internet messages awaiting delivery in the Exchange server. (Exchange 5.5 only; Exchange 2000 uses the IIS/SMTP service.)
xchgIMSOutQueLen	Number of messages awaiting conversion to Internet mail format. (Exchange 5.5 only; Exchange 2000 uses the IIS/SMTP service.)
xchgIMSTotQueLen	Total number of messages waiting in Internet Mail Service (IMS) queues. On Exchange 2000, this number represents the sum of the local and remote SMTP server queue lengths, plus the SMTP server local and remote retry queue lengths.

19254
3

Table 3-5: Selected MIB Objects – Exchange Connector Performance Group

MIB Object	Description
xchgIMSTotalConn	Total number of successful Internet mail connections. On Exchange 2000, this object represents the sum of the SMTP server total incoming and outgoing connections.
xchgIMSQueOut	Number of messages waiting for delivery to the Internet. On Exchange 2000, this object represents the SMTP server local queue length.
xchgIMSTotalInMsgs	Total number of Internet messages delivered to Exchange. On Exchange 2000, this object represents the total number of messages received by the SMTP server.
xchgIMSTotalOutMsgs	Total number of outbound messages delivered to Exchange server. On Exchange 2000, this object represents the total number of messages sent by the SMTP server.
xchgCCMCQueIn	Number of messages in the cc:Mail connector queue awaiting delivery to Exchange.
xchgCCMCQueOut	Number of messages in Exchange awaiting delivery to cc:Mail connector.

19253
ε

Figure 3-9 shows a sample AdvantEDGE View Exchange 5.5 queue.

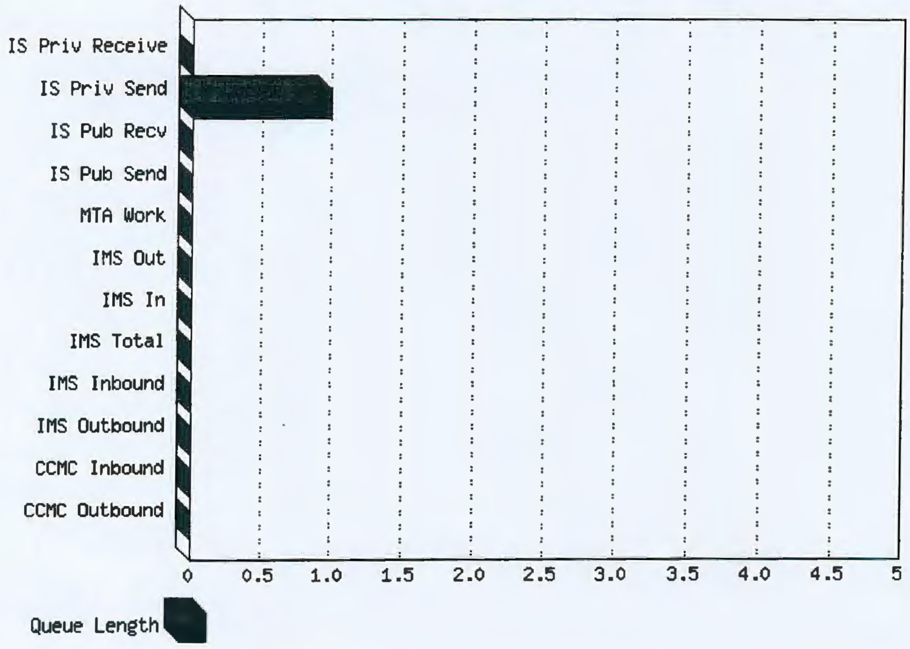

Figure 3-9: Exchange Queues for Exchange 5.5 Server

19252
ε

1925-3

4

Using AdvantEDGE for Microsoft Exchange

This chapter describes how to configure and use AdvantEDGE for Microsoft Exchange. This Point module is implemented as a SystemEDGE plug-in. After you enable this Point module in the sysedge.cf file and license it, it will load automatically at SystemEDGE start time. For more information, refer to "Configuring AdvantEDGE for Microsoft Exchange" and "Licensing AdvantEDGE for Microsoft Exchange" on page 2-3.

The AdvantEDGE for Microsoft Exchange plug-in implements additional MIB objects that provide advanced information about the health and availability of the Microsoft Exchange groupware application. It can operate with any SNMP-compliant management software, such as Concord's *e*Health suite of products, AdvantEDGE View, HP OpenView, and others. If you are using AdvantEDGE for Microsoft Exchange with *e*Health, refer to the *e*Health Web Help for more information about the reports that are available.

The default configuration settings of the AdvantEDGE for Microsoft Exchange plug-in enable you to use the advanced self-monitoring capabilities of SystemEDGE in conjunction with AdvantEDGE for Microsoft Exchange.

Editing the SystemEDGE Configuration File

You can use AdvantEDGE View or another SNMP tool to edit the SystemEDGE configuration file to utilize the MIB objects found in AdvantEDGE for Microsoft Exchange with the process-monitoring, threshold-monitoring, Windows NT event-monitoring, and history-collection features of the SystemEDGE agent. All MIB objects that are related to AdvantEDGE for Microsoft Exchange exist at object identifier (OID) branch 1.3.6.1.4.1.546.16.1 in the Concord Systems Management MIB. The MIB is defined in the xchgmod.asn1 file, which is available in the AdvantEDGE for Microsoft Exchange product installation.

19250
E

Assigning Entry Rows in the SystemEDGE Self-Monitoring Tables

All SystemEDGE self-monitoring tables (for example, log monitoring, Windows NT event monitoring, process/service monitoring, threshold monitoring, and history collection) require the use of unique row numbers. Each table contains an *Index* column which acts as a *key field* to distinguish rows in the table. This section describes the benefits of reserving a block of rows (somewhere in the range of 11 to the maximum number of rows in your table) for use by the system or application administrator.

Setting Local Policy

You may choose, as a matter of local policy, to reserve a block of rows for system administration. This policy allows you to define row entries within a reserved block of rows without worrying about the row already being taken by another user's entry. In compliance with the local policy, all other users should use row indices that are outside of the reserved range when they define user-configured entries.

By reserving a block of rows, you can define a consistent set of conditions (row entries) to be monitored across all machines such that the same condition is defined in the same row number on each of the machines. For example, you might use row 3000 in each table to define entries monitoring the Exchange MTA work queue length (xchgMTAWorkQueueLen). You can then distribute this configuration to every host so that every machine that is running Microsoft Exchange uses row 3000 for monitoring MTA work queue length, whether it is the threshold monitoring table or the history table. Further, every machine can also use row 3000 for monitoring the MTA service in the process/service monitoring table.

Reserving Blocks of Rows

To reserve a block of rows for monitoring Microsoft Exchange:

1. Decide on a block of rows that you want to reserve for your use with monitoring Microsoft Exchange.
2. Use that block of rows to define a set of row entries for each of the respective SystemEDGE self-monitoring tables. For more information, refer to the chapter on self-monitoring in the *SystemEDGE Agent User Guide.*

19249
E

3. Distribute configuration file entries out to all hosts that are running Microsoft Exchange and AdvantEDGE for Microsoft Exchange. For more information, refer to the *Automating the Deployment of SystemEDGE and AdvantEDGE Point Plug-in Modules* white paper.

NOTE

As an alternative, you can use this row-number assignment policy with AdvantEDGE View for group configuration operations.

4. Require end-users to avoid your block of rows when defining their own self-monitoring table entries.

Using the SystemEDGE Self-Monitoring Features

The examples in this section show SystemEDGE configuration-file commands for monitoring Microsoft Exchange. Add these commands to the sysedge.cf file to enable monitoring of the MIB objects they specify. Modify these examples as necessary to monitor the MIB objects that are relevant for your configuration.

The examples in the following sections present row numbers in the 5000 range; select a row number for your configuration that conforms to local policies. For more information on row assignment, refer to "Assigning Entry Rows in the SystemEDGE Self-Monitoring Tables" on page 4-2.

The following command, for example, instructs the SystemEDGE agent to monitor whether the Exchange MTA process is alive every 30 seconds and to store the data in row 5000 of the Process Monitoring table:

```
watch process procAlive 'emsmta|EMSMTA' 5000 0x0 30
 'Exchange Dir' ''
```

For more information about the syntax for the commands in this section, refer to the *SystemEDGE Agent User Guide.*

NOTE

Enter the commands throughout this chapter on one line. Do not use a carriage return to match the formatting shown here.

10248
E

Using SystemEDGE Process Monitoring

This section provides examples for monitoring the availability of critical Microsoft Exchange processes and services through SystemEDGE process and service monitoring. Enter the following commands in the sysedge.cf file to monitor these processes. For more information, refer to the chapter on process and service monitoring in the *SystemEDGE Agent User Guide.*

Monitoring the Exchange 5.5 Directory Service

To make sure the Exchange 5.5 Directory Service is running, enter the following command:

```
watch process procAlive 'dsamain|DSAMAIN' 5000 0x0 30
  'Exchange Dir' ''
```

Monitoring the Exchange MTA

To make sure the Exchange MTA is running, enter the following command:

```
watch process procAlive 'emsmta|EMSMTA' 5001 0x0 30 'Exchange
  MTA' ''
```

Monitoring the Exchange Information Store

To make sure the Exchange Information Store is running, enter the following command:

```
watch process procAlive 'store|STORE' 5002 0x0 30 'Exchange
  Info Store' ''
```

Monitoring the Exchange Attendant

To make sure the Exchange Attendant is running, enter the following command:

```
watch process procAlive 'mad|MAD' 5003 0x0 30 'Exchange
  Attendant' ''
```

Monitoring the Exchange Event Service

To make sure the Exchange Event Service is running, enter the following command:

```
watch process procAlive 'events|EVENTS' 5004 0x0 30 'Exchange
  Event Service' ''
```

19247
ε

Monitoring the Exchange 2000 SMTP Service

To make sure the Exchange 2000 SMTP service is running, enter the following command:

```
watch process procAlive 'smtp|SMTP' 5005 0x0 30 'Exchange 2000
  SMTP Service' ''
```

Monitoring the Core IIS Service

To make sure the Core IIS Service is running, enter the following command:

```
watch process procAlive 'iis|IIS' 5006 0x0 30 'Core IIS
  Service' ''
```

Using SystemEDGE Threshold Monitoring

This section provides examples for monitoring important Exchange metrics through SystemEDGE threshold monitoring. Add the commands that are provided in the following sections to the sysedge.cf file to monitor thresholds for these MIB objects. For more information, refer to the chapter on threshold monitoring in the *SystemEDGE Agent User Guide.*

NOTE

The thresholds used in these examples may not be appropriate for your Microsoft Exchange server; select thresholds that are appropriate for your environment.

Monitoring the MTA Work Queue Length

To monitor MTA work queue length, enter the following command:

```
monitor oid xchgMTAWorkQueueLen.0 5002 0x0 60 absolute > 15
  'MTA Queue Len exceeds threshold' ''
```

Monitoring Messages Received by the MTA

To monitor the number of messages received by the MTA, enter the following command:

```
monitor oid xchgMTAInMessages.0 5003 0x0 60 delta > 35 'MTA
  In Messages exceeds threshold' ''
```

19246
ε

Monitoring Messages Sent by the MTA

To monitor the number of messages sent by the MTA, enter the following command:

```
monitor oid xchgMTAOutMessages.0 5004 0x0 60 delta > 35 'MTA
  Out Messages Exceeds threshold' ''
```

Monitoring Information Store Users

To monitor the number of Information Store users, enter the following command:

```
monitor oid xchgStoreUserCount.0 5005 0x0 60 absolute > 750
  'Store User Cnt exceeds threshold' ''
```

Monitoring Private Store Messages Submitted by Clients

To monitor the number of Private Store messages submitted by clients, enter the following command:

```
monitor oid xchgISPrivInMessages.0 5006 0x0 60 delta > 35
  'Priv Store In Msg crosses threshold' ''
```

Monitoring Private Store Messages Delivered to Recipients

To monitor the number of Private Store messages delivered to recipients, enter the following command:

```
monitor oid xchgISPrivOutMessages.0 5007 0x0 60 delta > 35
  'Priv Store Out Msg crosses threshold' ''
```

Monitoring SMTP Queue Length

To monitor the SMTP Queue Length, enter the following command:

```
monitor oid xchgIMSTotQueLen.0 5008 0x0 60 delta > 35 'SMTP
  Queue Length crosses threshold' ''
```

19245
ε

Using SystemEDGE History Collection

This section provides examples for tracking the value of important Microsoft Exchange metrics over time through SystemEDGE history collection. Add the commands in the following sections to the sysedge.cf file to collect history for these MIB objects. For more information, refer to the chapter on history collection in the *SystemEDGE Agent User Guide.*

NOTE

The number of samples and the interval between samples used in these examples may not be appropriate for your Microsoft Exchange server; choose values that are appropriate for your environment.

Collecting History for MTA Work Queue Length

To collect history for MTA work queue length, enter the following command:

```
emphistory 5002 60 xchgMTAWorkQueueLen.0 480 'MTA Queue Len
  History'
```

Figure 4-1 shows a sample AdvantEDGE View Exchange Work Queue Length History.

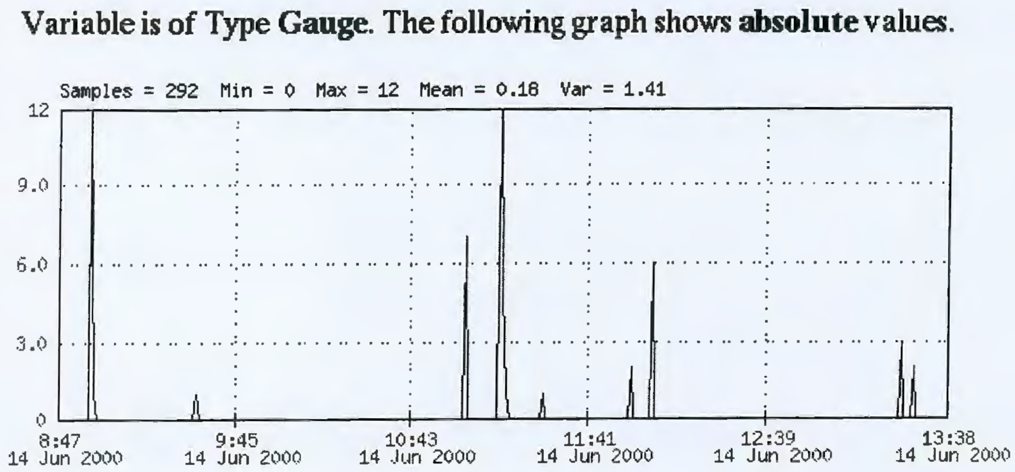

Variable is of Type Gauge. The following graph shows **absolute** values.

Figure 4-1: Exchange MTA Work Queue Length History

19244
ε

Collecting History for MTA Message Reception

To collect history for MTA message reception, enter the following command:

```
emphistory 5003 60 xchgMTAInMessages.0 180 'MTA In Messages
 History'
```

Collecting History for MTA Message Delivery

To collect history for MTA message delivery, enter the following command:

```
emphistory 5004 60 xchgMTAOutMessages.0 180 'MTA Out Messages
 History'
```

Collecting History for Information Store User Count

To collect history for the Information Store user count, enter the following command:

```
emphistory 5005 60 xchgStoreUserCount.0 120 'Store User Cnt
 History'
```

Figure 4-2 shows an AdvantEDGE View sample history for user count.

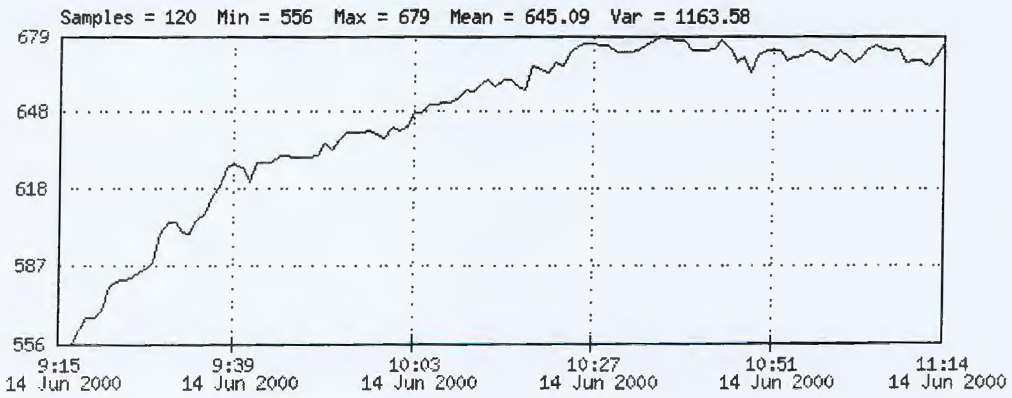

Figure 4-2: Sample History for Exchange User Count

Collecting History for Private Store Message Reception

To collect history for Private Store message reception, enter the following command:

```
emphistory 5006 60 xchgISPrivInMessages.0 120 'Priv Store In
 Msg History'
```

Collecting History for Private Store Message Delivery

To collect history for Private Store message delivery, enter the following command:

```
emphistory 5007 60 xchgISPrivOutMessages.0 120 'Priv Store
  Out Msg History'
```

Using SystemEDGE Windows NT Event Monitoring

This section provides examples for using the SystemEDGE Windows NT event-monitoring capabilities to capture important Microsoft Exchange-related Windows NT events and forward them to the appropriate configuration-management software as SNMP traps. Add the commands in the following sections to the sysedge.cf file to monitor these Windows NT events. For more information, refer to the chapter on Windows NT event monitoring in the *SystemEDGE Agent User Guide.*

Monitoring Exchange Events in the System Event Log

To watch for Exchange events in the system event log, enter the following command:

```
watch ntevent 5000 0x00 System All 'MSExchange' '.*'
  'Monitor Exchange System Events' ''
```

Monitoring Exchange Events in the Security Event Log

To watch for Exchange events in the security event log, enter the following command:

```
watch ntevent 5000 0x00 System All 'MSExchange' '.*'
  'Monitor Exchange Security Events' ''
```

Monitoring Exchange Events in the Application Event Log

To watch for Exchange events in the application event log, enter the following command:

```
watch ntevent 5000 0x00 System All 'MSExchange' '.*'
  'Monitor Exchange Application Events' ''
```

19224
3

Monitoring Exchange Database Error Events in the System Event Log

To watch for Exchange database error events in the system event log, enter the following command:

```
watch ntevent 5003 0x00 System Error 'EDB' ' '.*' 'Monitor
  Exchange Database Events' ''
```

Monitoring Exchange Database Error Events in the Security Event Log

To watch for Exchange database error events in the security event log, enter the following command:

```
watch ntevent 5004 0x00 Security Error 'EDB' '.*' 'Monitor
  Exchange Database Events' ''
```

Monitoring Exchange Database Error Events in the Application Event Log

To watch for Exchange database error events in the application event log, enter the following command:

```
watch ntevent 5005 0x00 Application Error 'EDB' '.*' 'Monitor
  Exchange Database Events' ''
```

192211 E

Index

A

architecture, Exchange 1-2
assigning rows for self-monitoring tables 4-2

C

capacity planning 3-10
collecting history
  Information Store user count 4-8
  MTA message delivery 4-8
  MTA message reception 4-8
  MTA Work Queue length 4-7
  Private Store message delivery 4-9
  Private Store message reception 4-8
configuration file. *See* sysedge.cf file
connector performance 3-11
core services 1-2

D

database error events 4-10

E

eHealth 1-4
Exchange
  architecture 1-2
  core services 1-2
  MIB 3-1

G

generating the AdvantEDGE for Microsoft Exchange license 2-4

I

IIS
  core services 1-2
  IMAP4 service 3-5
  NNTP service 3-5
  POP3 service 3-5
  RoutingEngine service 3-5
  SMTP service 3-4, 3-11
installation
  directory 2-2
  files 2-2
  requirements 2-1
installing
  AdvantEDGE for Microsoft Exchange 2-1
  SystemEDGE agent 2-1
Internet Information Services. *See* IIS

L

license
  key 2-5
  sample file 2-6
licensing 2-3
Live Health 1-4

A9240
E

INDEX

M

- Management Information Base. *See* MIB
- MIB 3-1
  - footprint group 3-8
  - performance group 3-7
  - server configuration group 3-2
- monitoring
  - application event log
    - database error events 4-10
    - Exchange events 4-9
  - core IIS service 4-5
  - database events 4-10
  - Exchange 2000 SMTP Service 4-5
  - Exchange Attendant 4-4
  - Exchange directory service 4-4
  - Exchange Event Service 4-4
  - Exchange Information Store 4-4
  - Exchange MTA 4-4
  - Information Store users 4-6
  - MTA
    - messages received 4-5
    - messages sent 4-6
  - MTA work queue length 4-5
  - Private Store
    - messages delivered 4-6
    - messages submitted 4-6
  - processes 4-4
  - security event log
    - database error events 4-10
    - Exchange events 4-9
  - servers in real time 3-10
  - SMTP queue length 4-6
  - system event log
    - database error events 4-10
    - Exchange events 4-9

P

- process monitoring 4-4

R

- real-time server monitoring 3-10
- requirements 2-1
- reserving blocks of rows for self-monitoring tables 4-2

S

- server
  - configuration MIB group 3-2
  - performance MIB group 3-10
- setting local policy 4-2
- sysedge.cf file 4-3
  - loading plug-ins 2-3
  - monitoring Windows NT events 4-9
- sysedge.lic file 2-3, 2-5
- SystemEDGE agent 1-1

T

- trend analysis 3-10

U

- using AdvantEDGE for Microsoft Exchange 4-1
  - with *e*Health 1-4
  - with Live Health 1-4

X

- xchgmod.asn1 file 2-2, 3-2, 4-1
- xchgmod.dll file 2-2, 2-3
- xchgmod.exe file 2-1
- xchgmod.pdf file 2-2

19239
E

ADVANTEDGE™

FOR MICROSOFT® IIS

User Guide

Release 1.0 Patchlevel 2
or Later

09-16040-003
April 2001

19238
a

Important Notice

Application Health, System Health, Live Health, SystemEDGE, AdvantEDGE, eHealth, Infrasense, MIB Manager, State-of-the-World, and the FirstSense logo are trademarks of Concord Communications, Inc.

The Total View, Concord, the Concord logo, Empire Technologies, Network Health, TRAKKER, FirstSense, and We See It Happening are registered trademarks of Concord Communications, Inc.

Adobe, Acrobat, and the Acrobat logo are trademarks, which may be registered in certain jurisdictions, and PostScript is a registered trademark of Adobe Systems Incorporated.
Intel is a registered trademark of Intel Corporation.
Windows, Windows NT, Windows 2000, Internet Information Server, and Internet Information Services are trademarks of Microsoft Corporation.

© Copyright Bigelow and Holmes 1986, 1985. Lucida is a registered trademark of Bigelow & Holmes. Sun Microsystems Inc., AT&T, and Bigelow & Holmes make no representations about the suitability of the source code for any purpose. It is provided "as is" without express or implied warranty of any kind.

General Notice: Some of the product names used herein have been used for identification purposes only and may be trademarks of their respective companies.

Proprietary Notice

The information and descriptions contained herein are the property of Concord Communications, Inc. Such information and descriptions may not be copied, disseminated, or distributed without the express written consent of Concord Communications, Inc. Concord Communications, Inc., assumes no responsibility for any inaccuracies that may appear in this document. Concord Communications, Inc., reserves the right to improve its products and change specifications at any time without notice.

U. S. Government Restricted Rights

Use, reproduction, and disclosure by the U.S. Government are subject to the restrictions set forth in FAR §52.227-19 (c) (1) and (2) and DFARS §252.227-7013 (c) (1) (ii).

Patent Information

U. S. Patent 5,615,323
Patents Pending

1237
4

Table of Contents

About This Guide

How This Guide Is Organized v
Conventions in This Guide vi
Contact Information. vi

Chapter 1
Introduction

Introducing AdvantEDGE for Microsoft IIS 1-1
Microsoft IIS Architecture. 1-2
Using AdvantEDGE for Microsoft IIS. 1-3
    Using AdvantEDGE for Microsoft IIS with eHealth 1-3
    Using AdvantEDGE for Microsoft IIS with Live Health. 1-4

Chapter 2
Installing AdvantEDGE for Microsoft IIS

Installation Requirements. 2-1
Installing the Software. 2-2
AdvantEDGE for Microsoft IIS Files 2-3
Configuring AdvantEDGE for Microsoft IIS 2-3
Licensing AdvantEDGE for Microsoft IIS 2-3
    Obtaining a License. 2-4
    Generating the License 2-4
    Sample License File. 2-6

19236
ε

Table of Contents

Chapter 3
Using the AdvantEDGE for Microsoft IIS MIB

Configuration Section 3-2
Server Configuration. 3-2
Performance Section. 3-3
IIS Footprint 3-3
Server Performance 3-5
Services Group 3-6
WWW Group 3-6
FTP Group 3-8
SMTP Group. 3-10
NNTP Group 3-11
Extensions Group 3-13
ASP Group 3-13
CGI Group 3-14
ISAPI Group 3-15

Chapter 4
Using AdvantEDGE for Microsoft IIS

Editing the SystemEDGE Configuration File 4-1
Assigning Entry Rows for the SystemEDGE Self-Monitoring Tables. . . 4-2
Setting Local Policy 4-2
Reserving Blocks of Rows 4-2
Using the SystemEDGE Self-Monitoring Features 4-3
Using SystemEDGE Process Monitoring 4-4
Using SystemEDGE Threshold Monitoring 4-5
Using SystemEDGE History Collection 4-7
Using SystemEDGE Windows NT Event Monitoring 4-9
Using the Windows NT RegPerf Extensions 4-12

Index

About This Guide

The *AdvantEDGE for Microsoft IIS User Guide* provides instructions for installing and using AdvantEDGE for Microsoft IIS for Windows NT x86 and Windows 2000 systems. This guide is intended for the person responsible for installing and configuring AdvantEDGE for Microsoft IIS.

This version supports AdvantEDGE for Microsoft IIS Release 1.0 Patchlevel 2 or later, and the SystemEDGE agent Release 4.0, Patchlevel 3 and later.

NOTE

> The acronym IIS stands for Internet Information Server in version 4.0 for Windows NT, and for Internet Information Services in version 5.0 for Windows 2000.

To use AdvantEDGE for Microsoft IIS, you should have a basic understanding of the Microsoft IIS application, the Concord SystemEDGE agent, and your host's operating systems environment. Refer to Microsoft documentation and the *SystemEDGE Agent User Guide* for more information.

How This Guide Is Organized

This guide is organized as follows:

- Chapter 1, "Introduction," provides an overview of AdvantEDGE for Microsoft IIS and its capabilities for monitoring Microsoft IIS.
- Chapter 2, "Installing AdvantEDGE for Microsoft IIS," explains how to install, configure, and license the Concord AdvantEDGE for Microsoft IIS software on a host system.
- Chapter 3, "Using the AdvantEDGE for Microsoft IIS MIB," describes the information that is available through the Concord AdvantEDGE for Microsoft IIS MIB.
- Chapter 4, "Using AdvantEDGE for Microsoft IIS," explains how to configure and use Concord's AdvantEDGE for Microsoft IIS in your host environment.

Conventions in This Guide

The following paragraph shows a sample command. Due to space limitations in this guide, some commands wrap from one line to the next. Disregard these line breaks, and **enter each command as one line**. Otherwise, your command syntax will be incorrect.

For example, when you see a command such as the following:

```
watch process procAlive 'inetinfo|INETINFO' 5000 0x0
 30 'IIS' ''
```

You **must** enter the command on one line, as shown here:

```
watch process procAlive 'inetinfo|INETINFO' 5000 0x0 30 'IIS' ''
```

Contact Information

If you need any assistance with the SystemEDGE agent or the AdvantEDGE for Microsoft IIS Point module, contact Customer Support, using one of the following methods:

Phone:	(888) 832-4340 (for calls from the USA and Canada) (508) 303-4300 (for calls from other countries)
Fax:	(508) 303-4343
E-mail:	support@concord.com
Web site:	http://www.concord.com
Licensing:	http://license.concord.com

1

Introduction

This chapter provides an overview of how you can use the AdvantEDGE for Microsoft IIS Point module to monitor the Microsoft IIS application. You must install this Point module on every IIS workstation that you want to monitor.

NOTE

Unless otherwise stated, the term *IIS*, as used throughout this guide and the Management Information Base (MIB) specification, refers to the Microsoft IIS application in its entirety, which encompasses all of the services and optional components.

Introducing AdvantEDGE for Microsoft IIS

AdvantEDGE for Microsoft IIS is a plug-in module for the SystemEDGE agent. This plug-in enables information technology (IT) operators to monitor the performance and availability of Microsoft IIS on Microsoft Windows NT 4.0 and Windows 2000. Microsoft IIS is a super server that consists of several services, which are described in the next section, "Microsoft IIS Architecture".

AdvantEDGE for Microsoft IIS makes important information about IIS available to management software through the SystemEDGE agent and Simple Network Management Protocol (SNMP). The SystemEDGE agent enables the monitoring of important IIS metrics, processes, and services, as well as the sending of SNMP traps when exceptions or exception conditions occur.

NOTE

This document is not intended as a manual on how to install, administer, or use Microsoft IIS. For help with IIS, refer to your Microsoft documentation.

19237

Microsoft IIS Architecture

Microsoft IIS is a complex piece of software with many components. At the core of IIS are the following services:

- World Wide Web (WWW) service, which services the Web content to and from users
- File Transfer Protocol (FTP) service, which services file transfer requests to and from users
- Simple Mail Transfer Protocol (SMTP) service, which provides a mail transport mechanism for IIS
- Network News Transfer Protocol (NNTP) service, which provides news groups for IIS

Figure 1-1 shows the core components of AdvantEDGE for Microsoft IIS.

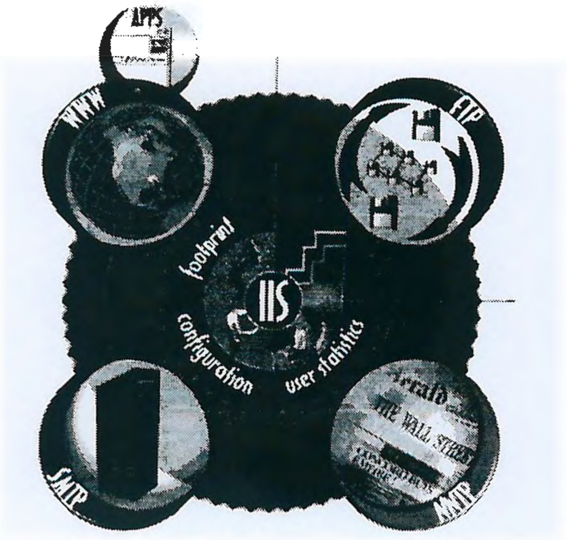

Figure 1-1: AdvantEDGE for Microsoft IIS Components

IIS controls worker threads, and it can use each worker thread to provide a specific service that the user has requested. This architecture requires the monitoring of the IIS process and other processes in order to obtain sufficient information on the performance of IIS. Because IIS includes so many components, its proper functioning requires the availability of a variety of processes; configuration parameters and settings; and Windows NT services, queues, and system resources.

19231
3

Using AdvantEDGE for Microsoft IIS

AdvantEDGE for Microsoft IIS provides you with the tools and information that you need to monitor and respond to the IIS application and its use of your system resources. With AdvantEDGE for Microsoft IIS, you can fix potential problems before users—and your business—are affected.

You can use AdvantEDGE for Microsoft IIS with any SNMP-compliant management software, including Concord's *e*Health suite of products, AdvantEDGE View, HP OpenView, and others. With AdvantEDGE for Microsoft IIS and the SystemEDGE agent, you can perform the following types of tasks:

- Monitor the availability of IIS and its various services.
- Automatically restart any service that fails.
- Alert IT staff when IIS starts to consume significant levels of system resources, including CPU, disk space, and memory.
- Monitor logs for security, system, and application events across the Web, FTP, SMTP, and NNTP services.
- Detect error statistics across the Active Server Pages (ASP), Common Gateway Interface (CGI), and Internet Server Application Program Interface (ISAPI) application extension pages, including Web 404 (page not found) errors and ASP script errors.

Using AdvantEDGE for Microsoft IIS with *e*Health

You can use AdvantEDGE for Microsoft IIS and the SystemEDGE agent with the *e*Health product suite to provide the historical data for long-term trending analysis and capacity planning. With *e*Health – Application Assessment, you can run At-a-Glance, Trend, Top N, and MyHealth reports for the following types of information:

- Amount of Central Processing Unit (CPU), total memory, and disk space the IIS application is using
- Size of the IIS service logs
- Number of times the IIS application is using the database cache to redisplay information
- Number of bytes and number of files processed by the WWW and FTP services
- Number of users for the WWW and FTP services

19230
ε

- Number of WWW requests the IIS application is processing per second
- Number of page not found errors the WWW service is encountering

For more information about the variables that you can monitor and reports that you can run when you integrate AdvantEDGE for Microsoft IIS with *e*Health, refer to the *e*Health Web Help.

Using AdvantEDGE for Microsoft IIS with Live Health

You can also use AdvantEDGE for Microsoft IIS and the SystemEDGE agent with Live Health for real-time detection of potential problems. Live Health applies intelligent algorithms to the data, resulting in precise assessments of application health and performance. For more information about how Live Health can detect "brownouts" and service delays across applications, systems, and networks, refer to the Live Health Web Help.

2

Installing AdvantEDGE for Microsoft IIS

This chapter explains how to install, configure, and license AdvantEDGE for Microsoft IIS.

NOTE

> For the most current information about installing the AdvantEDGE for Microsoft IIS module, refer to the release notes (relnotes.txt) that ship on the installation CD.

Installation Requirements

Before you install AdvantEDGE for Microsoft IIS, you must first install, license, and configure the SystemEDGE agent Release 4.0, Patchlevel 3 or later. For more information, refer to the *SystemEDGE Agent User Guide.*

Your system must also be running the Microsoft IIS application. Use Table 2-1 to determine which version of IIS and which service packs are required for your operating system.

Table 2-1: Required IIS Version and Service Pack by Operating System

Operating System	Required Software	Required Service Pack
Windows NT 4.0 x86	Microsoft Internet Information Server Version 4.0	6a
Windows 2000	Microsoft Internet Information Services Version 5.0	1

For more information, refer to Microsoft documentation.

Installing the Software

AdvantEDGE for Microsoft IIS is distributed as a self-extracting executable named **iismod.exe** for Windows NT and Windows 2000.

Follow these steps to install AdvantEDGE for Microsoft IIS:

1. Locate the IIS workstation that needs to be monitored.
2. Log on to the system as administrator.
3. Click **Start**.
4. Select **Programs → Command Prompt**.
5. Insert the CD containing the Concord software distributions into the CD-ROM drive.

   The operating system automatically mounts the drive using the CD-ROM drive's corresponding drive letter. The particular drive letter is specific to your system and depends on the number and types of disks attached to your system.
6. Determine which directory you want to use as the installation directory for AdvantEDGE for Microsoft IIS. If the SystemEDGE agent is installed at C:\sysedge, the recommended installation directory is C:\sysedge\plugins.
7. Run the self-extracting executable by entering the following at the command prompt, where `D:` is the CD-ROM drive for your system, and `C:\sysedge\plugins` is the installation directory:

   ```
   D:\iismod\ntx86\iismod.exe -dir C:\sysedge\plugins
   ```

   The -dir option is important because it instructs the self-extracting executable to create the intended subdirectory hierarchy that is used throughout this guide. The distribution is then placed in an iismod subdirectory within the specified target directory (for example, C:\sysedge\plugins\iismod).

NOTE
You cannot execute iismod.exe directly from the CD-ROM.

AdvantEDGE for Microsoft IIS is now installed.

19227
E

AdvantEDGE for Microsoft IIS Files

Table 2-2 describes the files created by the installation process.

Table 2-2: Files Installed by AdvantEDGE for Microsoft IIS

File Name	Description
iismod.dll	AdvantEDGE for Microsoft IIS dynamic link library (DLL) module for Windows NT and Windows 2000
iismod.pdf	*AdvantEDGE for Microsoft IIS User Guide*
iismod.asn1	AdvantEDGE for Microsoft IIS MIB specification
examples	AdvantEDGE for Microsoft IIS monitoring examples
relnotes.txt	Release notes for AdvantEDGE for Microsoft IIS

Configuring AdvantEDGE for Microsoft IIS

The SystemEDGE agent uses the configuration file sysedge.cf and the *sysedge_plugin* keyword to specify which AdvantEDGE modules to load at system initialization. By default, the SystemEDGE agent does not load any plug-ins at initialization time, but you can edit the sysedge.cf file to configure the agent to load any AdvantEDGE Point modules that you have installed.

The sysedge.cf file is located, by default, in your system directory; for example, C:\winnt\system32. To configure the SystemEDGE agent to start AdvantEDGE for Microsoft IIS, you must provide the complete path name to iismod.dll, the AdvantEDGE for Microsoft IIS dynamic link library. The actual path depends on the location you selected when you installed the AdvantEDGE for Microsoft IIS files. For example, enter this command if you installed the files in the C:\sysedge\plugins\iismod directory:

```
sysedge plugin C:\sysedge\plugins\iismod\iismod.dll
```

For more information about the sysedge.cf file, refer to the *SystemEDGE Agent User Guide.*

Licensing AdvantEDGE for Microsoft IIS

Like the SystemEDGE agent, AdvantEDGE for Microsoft IIS utilizes a host-based license method. Copies of AdvantEDGE for Microsoft IIS can run only on systems that possess a valid license key. This license is separate from the one used for the SystemEDGE agent.

19226
3

The first time that you attempt to start the SystemEDGE agent after installing AdvantEDGE for Microsoft IIS, the agent will display a message saying that a valid license was not found for AdvantEDGE for Microsoft IIS. It then provides you with a *public key* that is used to generate a permanent license key for your host machine.

A license key is made up of four space-separated, 8-character sequences, totaling 32 characters. The AdvantEDGE for Microsoft IIS license is stored in the sysedge.lic file, the same file that contains the SystemEDGE agent licenses. Refer to the sample license file on page 2-6.

Obtaining a License

To obtain a license, you can do any of the following:

- Run the Concord-supplied licenseutil.pl script.
- Run the licenseme.exe license utility.
- Use the AdvantEDGE View licensing procedure, which uses SNMP traps. For more information, refer to the AdvantEDGE View Web Help.
- Send an e-mail request to license@concord.com and place the returned license key in the appropriate license file.

NOTE

Always include the Customer ID and user name in license requests that you send through e-mail.

- Complete the online license form through the Internet, as described in the next section, "Generating the License"."

For more information about licensing, refer to the *SystemEDGE Agent User Guide* and the *Automating the Licensing of SystemEDGE and AdvantEDGE Point Plug-in Modules* white paper.

Generating the License

This section describes how to generate the license using the Web-based license form. The SystemEDGE setup program can generate the licensing information for your system.

NOTE

If you are using an evaluation copy of AdvantEDGE for Microsoft IIS, you must request a temporary license that will enable it to operate during the evaluation period.

19225 e

To generate a license, follow these steps:

1. Run the SystemEDGE agent setup command by entering the following at the command prompt:

   `sysedge\setup -l`

   The setup program displays a message similar to the following:

```
SystemEDGE Version 4.0 Patchlevel 3
Copyright 2001 by Concord Communications, Inc.
Please contact Concord Communications, Inc. to obtain a license
http://www.concord.com/support, Email: license@concord.com
Provide: sysedge neptune NTx86 4.0 346561363366b19c 4.0 Patchlevel 3
```

2. To obtain a license for AdvantEDGE for Microsoft IIS, fill out the online Web-based license form available from Concord's licensing Web server at the following URL:

   `http://license.concord.com`

   NOTE

   You must supply a user name and password to access the license form.

   The license form asks you to supply the following information:

   - Customer ID
   - Name
   - E-mail address
   - Software version number (4.0 in the example above)
   - Patchlevel
   - System name (neptune in the example above)
   - Operating system name (NTx86 in the example above)
   - Operating system version (4.0 in the example above)
   - System identifier (346561363366b19c in the example above)

   NOTE

   When you are licensing AdvantEDGE for Microsoft IIS, select **iismod** as the product on the licensing form.

   After you submit the license request, the Concord Web server generates a license and displays it to your Web browser. It also e-mails the license to the contact person in your organization.

19224

3. Copy the generated license key into the sysedge.lic file in the system32 subdirectory (for example, C:\winnt\system32), and then save the file.

   The license key is case sensitive. Copy it exactly as it appears. If possible, use your system's cut-and-paste facility instead of entering it by hand. If you do enter the license key by hand, be careful not to confuse characters such as the letters l and I and the number 1,or the letter O and the number **0**.

4. Stop and then restart the Windows NT Master agent by entering the following commands at the command prompt:

   ```
   net stop snmp
   net start snmp
   ```

The AdvantEDGE for Microsoft IIS module is now licensed and ready to use.

Sample License File

The following is a sample SystemEDGE agent license file. A pound character (#) in column 1 indicates that the entire line is a comment.

```
# license file for SystemEDGE Agent
# Empire Technologies, Inc.
# A Concord Communications Company
# http://www.concord.com
#
# file /etc/sysedge.lic or %SystemRoot%\system32\sysedge.lic
# A valid license key has four parts of 8 characters per part
# parts are separated by space(s) with one license key per line
# sysedge neptune NT/x86 4.0 807cb1da007cb1da 4.0
e13311d3 0F2a7cb1 abC512dc fF8C923a

# iismod neptune NT/x86 4.0 807cb1da007cb1da 1.0
a7943fde 098a87ij a4kiuf39 afafEkj4
```

19223
E

3

Using the AdvantEDGE for Microsoft IIS MIB

This chapter explains the organization and content of the Concord Communications MIB for Microsoft IIS. This MIB specification (iismod.asn1) defines a collection of objects for monitoring and managing IIS. You must configure the SystemEDGE agent to monitor the AdvantEDGE for Microsoft IIS MIB objects that are relevant for your configuration. For more information about configuring the SystemEDGE agent to monitor the IIS application, refer to Chapter 4, "Using AdvantEDGE for Microsoft IIS."

Figure 3-1 shows the organization of the AdvantEDGE for Microsoft IIS MIB.

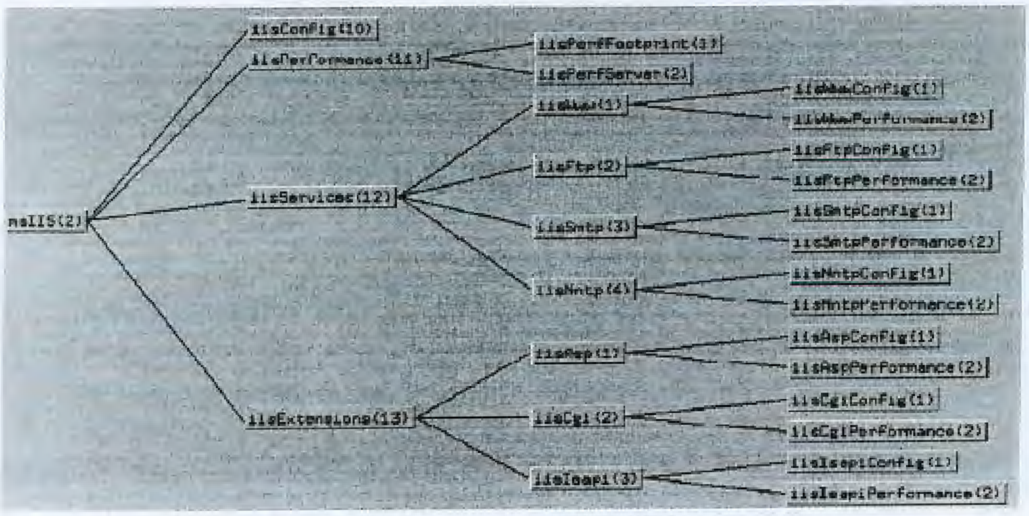

Figure 3-1: AdvantEDGE for Microsoft IIS MIB

19222
ε

The MIB is organized into broad sections for server, services, and extension information. Within those broad sections are subsections for configuration and performance. Within the performance section, a footprint section defines MIB objects that convey how much of the underlying system's resources are consumed by IIS.

The following sections highlight important MIB objects from the IIS MIB. This chapter defines all sections of the IIS MIB, but it does not define all of the MIB objects. For a complete list of MIB objects, refer to the AdvantEDGE for Microsoft IIS MIB specification (iismod.asn1).

Configuration Section

The Configuration section of the AdvantEDGE for Microsoft IIS MIB contains configuration parameters and settings that are important for streamlining the health and performance of IIS. It also includes information about server configuration.

Server Configuration

The Server Configuration group contains configuration parameters, process IDs, and version numbers. Table 3-1 defines important Server Configuration parameters.

Table 3-1: Selected MIB Objects – IIS Server Configuration Group

MIB Object	Description
iisVersion	IIS version.
iisPid	IIS process ID.
iisObjectCacheTTL	How often the cache scavenger runs.
iisMAXPoolThreads	How many threads IIS will use to perform tasks.
iisListenBackLog	Maximum number of connection requests in the queue for each service.
iisOpenFilesInCache	Number of files IIS will keep open in the cache.

1922

Figure 3-2 shows a sample AdvantEDGE View process-status window for IIS.

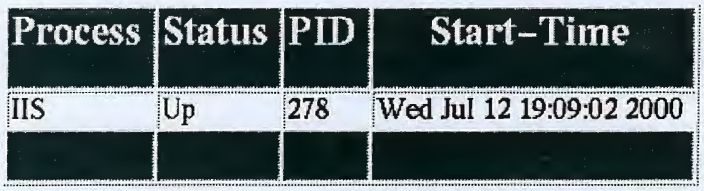

Process	Status	PID	Start-Time
IIS	Up	278	Wed Jul 12 19:09:02 2000

Figure 3-2: Sample IIS Process ID and Status

Performance Section

The Performance section of the AdvantEDGE for Microsoft IIS MIB contains performance data that is necessary for capacity planning and trend analysis, as well as for real-time performance and availability monitoring. The Performance group is divided into subgroups for footprint data and server performance.

IIS Footprint

This group provides information about IIS CPU, memory, and disk resource consumption, more commonly called its *footprint*. Long-term trending analysis of footprint information is useful for anticipating and avoiding problems due to resource exhaustion. You can also monitor footprint information in real time to detect and correct temporary resource exhaustion due to viruses, security incidents, and hardware failures. Table 3-2 defines important IIS Footprint metrics.

Table 3-2: Selected MIB Objects – IIS Footprint Group

MIB Object	Description
iisCPUTime	IIS total accumulated CPU time.
iisPercentCPU	Percentage of CPU, over the last interval, used by IIS.
iisTotalRSS	Total real memory currently in use by IIS.
iisPercentMEM	Percentage of real memory currently in use by IIS.
iisTotalLogSize	Estimate of the current disk space used by all IIS services.
iisTotalDiskSize	Estimate of the current total disk space used by IIS.
iisTotalThreads	Total number of system threads used by IIS.

The following figures show sample AdvantEDGE View footprints for an IIS machine serving a medium-sized company. They represent real data collected from live IIS servers and displayed in AdvantEDGE View reports.

Figure 3-3 shows an example of disk usage by service:

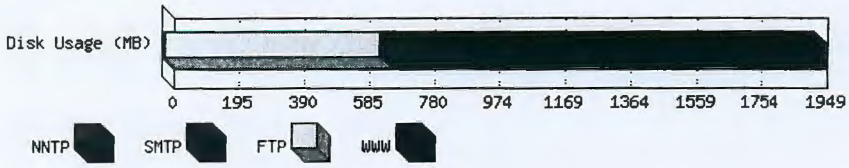

Figure 3-3: IIS Disk Usage by Service

Figure 3-4 shows an example of memory usage:

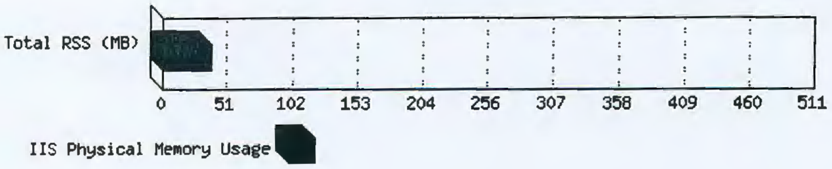

Figure 3-4: IIS Memory Usage

Figure 3-5 shows a sample IIS footprint summary:

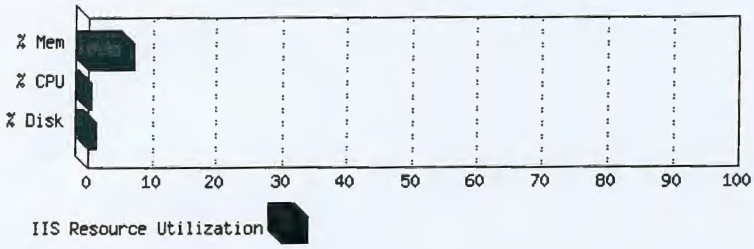

Figure 3-5: IIS Footprint Summary

19219
E

Server Performance

The Server Performance group provides performance metrics and counters for IIS, including user statistics and transfer statistics. These metrics include those useful for real-time management and longer-term capacity planning and trend analysis. Table 3-3 defines important Server Performance metrics.

Table 3-3: Selected MIB Objects – IIS Server Performance Group (Page 1 of 2)

MIB Object	Description
iisCacheHits	Total number of times an item was found by IIS in the object cache.
iisCacheMisses	Total number of times an item was not found by IIS in the object cache.
iisCacheFlushes	Number of times an item was deleted from the IIS object cache.
iisTtlCurAnonUsr	Total number of current anonymous users maintained by all IIS services.
iisTtlCurNAnonUsr	Total number of current non-anonymous users maintained by all IIS services.
iisTtlCurUsr	Total number of users maintained by all IIS services.
iisTtlMaxAnonUsr	Total maximum number of anonymous users maintained by all IIS services.
iisTtlMaxNAnonUsr	Total maximum number of non-anonymous users maintained by all IIS services.
iisTtlMaxUsr	Total maximum number of users maintained by all IIS services.
iisTtlAnonUsr	Running count of anonymous users maintained by all IIS services.
iisTtlNAnonUsr	Running count of non-anonymous users maintained by all IIS services.
iisTtlUsr	Running count of all users maintained by all IIS services.
iisTtlBytesSent	Count of KB sent by all IIS services.
iisTtlBytesRecv	Count of KB received by all IIS services.
iisTtlBytes	Count of KB transferred by all IIS services.
iisTtlFilesSent	Count of files sent by all IIS services.

19218

Table 3-3: Selected MIB Objects – IIS Server Performance Group (Page 2 of 2)

MIB Object	Description
iisTtlFilesRecv	Count of files received by all IIS services.
iisTtlFiles	Count of files transferred by all IIS services.
iisWebAppReqTtl	Total number of Web requests made by IIS.

Figure 3-6 shows a sample user statistics summary:

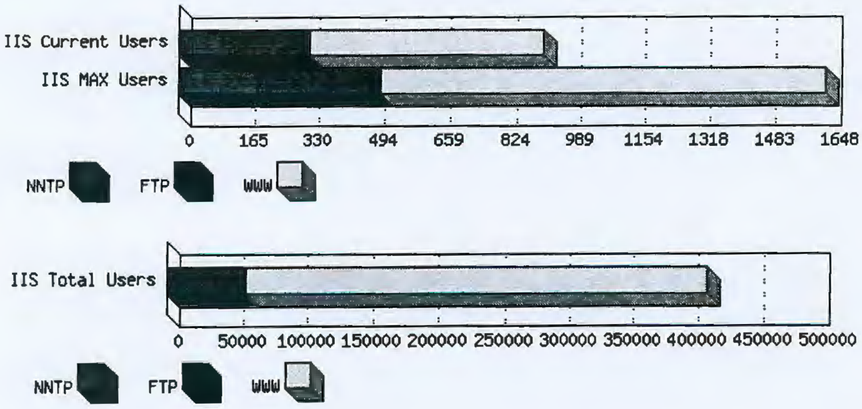

Figure 3-6: IIS User Statistics Summary

Services Group

The Services group provides configuration information, performance metrics, and counters for IIS services, including WWW, FTP, SMTP, and NNTP. These metrics include those useful for real-time management and longer-term capacity planning and trend analysis.

WWW Group

The World Wide Web (WWW) group provides configuration information, performance metrics, and counters for the IIS WWW Service. It looks at the WWW Service as a whole and provides totals of all WWW service activities in IIS.

19217

WWW Configuration Group

The WWW Configuration group provides version information for the IIS WWW service.

WWW Performance Group

The WWW Performance group provides performance metrics and counters for the IIS WWW service, including user statistics and transfer statistics. These metrics include those useful for real-time management and longer-term capacity planning and trend analysis. Table 3-4 defines important WWW Performance metrics.

Table 3-4: Selected MIB Objects – IIS WWW Performance Group (Page 1 of 2)

MIB Object	Description
iisWwwLogSize	Estimate of the current disk space used by the WWW service logs.
iisWwwTtlDiskSize	Total size in KB of log and service directories.
iisWwwRezSysSize	Size in KB of the WWW service resident system code.
iisWwwTtlNotFoundErr	Total count of the page-not-found errors for the WWW service.
iisWwwTtlConnAtempt	Total number of connections to the well-known port of the WWW service.
iisWwwTtlLogonAtempt	Total number of logins that have been attempted to the WWW service.
iisWwwTtlCurAnonUsr	Total number of current anonymous users maintained by the WWW service.
iisWwwTtlCurNAnonUsr	Total number of current non-anonymous users maintained by the WWW service.
iisWwwTtlCurUsr	Total number of users maintained by the WWW service.
iisWwwTtlMaxAnonUsr	Maximum number of anonymous users maintained by the WWW service.
iisWwwTtlMaxNAnonUsr	Maximum number of non-anonymous users maintained by the WWW service.
iisWwwTtlMaxUsr	Maximum number of users maintained by the WWW service.

Table 3-4: Selected MIB Objects – IIS WWW Performance Group (Page 2 of 2)

MIB Object	Description
iisWwwTtlAnonUsr	Running count of anonymous users maintained by the WWW service.
iisWwwTtlNAnonUsr	Running count of non-anonymous users maintained by the WWW service.
iisWwwTtlUsr	Running count of all users maintained by the WWW service.
iisWwwTtlBytesSent	Count of KB sent by the WWW service.
iisWwwTtlBytesRecv	Count of KB received by the WWW service.
iisWwwTtlBytes	Count of KB transferred by the WWW service.
iisWwwTtlFilesSent	Count of files sent by the WWW service.
iisWwwTtlFilesRecv	Count of files received by the WWW service.
iisWwwTtlFiles	Count of files transferred by the WWW service.
iisWwwTtlGetReq	Total GET methods on the WWW service.
iisWwwTtlHeadReq	Total HEAD methods on the WWW service.
iisWwwTtlPostReq	Total POST methods on the WWW service.
iisWwwTtlPutReq	Total PUT methods on the WWW service.
iisWwwTtlTraceReq	Total TRACE methods on the WWW service.
iisWwwTtlDeleteReq	Total DELETE methods on the WWW service.
iisWwwTtlOtherReq	Total methods that are not using GET, POST, PUT, DELETE, TRACE, or HEAD.
iisWwwTtlMethodReq	Total methods that are using GET, POST, PUT, DELETE, TRACE, and HEAD.

FTP Group

The FTP group provides performance metrics and counters for the IIS FTP service. It includes information such as user statistics and transfer statistics. This group looks at the FTP Service as a whole and provides totals of all FTP service activities in IIS.

19215
ε

FTP Configuration Group

The FTP Configuration group provides version information for the IIS FTP service.

FTP Performance Group

The FTP Performance group provides performance metrics and counters for the IIS FTP service, including the total current users, total users, total files transferred, and total KB transferred. These metrics include those useful for real-time management and longer-term capacity planning and trend analysis. Table 3-5 defines important FTP Performance metrics.

Table 3-5: Selected MIB Objects – IIS FTP Performance Group (Page 1 of 2)

MIB Object	Description
iisFtpLogSize	Estimate of the current disk space used by the FTP service logs.
iisFtpTtlDiskSize	Total size in KB of log and service directories.
iisFtpTtlConnAtempt	Total number of connections to the well-known port of the FTP service.
iisFtpTtlLogonAtempt	Total number of logins that have been attempted to the FTP service.
iisFtpTtlCurAnonUsr	Total number of current anonymous users maintained by the FTP service.
iisFtpTtlCurNAnonUsr	Total number of current non-anonymous users maintained by the FTP service.
iisFtpTtlCurUsr	Total number of users maintained by the FTP service.
iisFtpTtlMaxAnonUsr	Maximum number of anonymous users maintained by the FTP service.
iisFtpTtlMaxNAnonUsr	Maximum number of non-anonymous users maintained by the FTP service.
iisFtpTtlMaxUsr	Maximum number of users maintained by the FTP service.
iisFtpTtlAnonUsr	Running count of anonymous users maintained by the FTP service.
iisFtpTtlNAnonUsr	Running count of non-anonymous users maintained by the FTP service.

1192124
4

Table 3-5: Selected MIB Objects – IIS FTP Performance Group (Page 2 of 2)

MIB Object	Description
iisFtpTtlUsr	Running count of all users maintained by the FTP service.
iisFtpTtlBytesSent	Count of KB sent by the FTP service.
iisFtpTtlBytesRecv	Count of KB received by the FTP service.
iisFtpTtlBytes	Count of KB transferred by the FTP service.
iisFtpTtlFilesSent	Count of files sent by the FTP service.
iisFtpTtlFilesRecv	Count of files received by the FTP service.
iisFtpTtlFiles	Count of files transferred by the FTP service.

SMTP Group

The SMTP group provides configuration information, performance metrics, and counters for the IIS SMTP Service. This group contains information such as transfer statistics and queue lengths. It looks at the SMTP service as a whole and provides totals of all SMTP service activities in IIS.

SMTP Configuration Group

The SMTP Configuration group provides version information for the IIS SMTP service.

SMTP Performance Group

The SMTP Performance group provides performance metrics and counters for the IIS SMTP service, including total messages transferred, total KB transferred, and various queue lengths. These metrics include those useful for real-time management and longer-term capacity planning and trend analysis. Table 3-6 defines important SMTP Performance metrics.

Table 3-6: Selected MIB Objects – IIS SMTP Performance Group

MIB Object	Description
iisSmtpLogSize	Estimate of the current disk space used by the SMTP service logs.
iisSmtpTtlDiskSize	Total size in KB of log and service directories.
iisSmtpTtlConErr	Total connection errors for the SMTP service.

19213
3

Table 3-6: Selected MIB Objects – IIS SMTP Performance Group

MIB Object	Description
iisSmtpTtlBytesSent	Count of KB sent by the SMTP service.
iisSmtpTtlBytesRecv	Count of KB received by the SMTP service.
iisSmtpTtlBytes	Count of KB transferred by the SMTP service.
iisSmtpTtlMsgSent	Count of messages sent by the SMTP service.
iisSmtpTtlMsgRecv	Count of messages received by the SMTP service.
iisSmtpTtlMsg	Count of messages transferred by the SMTP service.
iisSmtpCurInConn	Number of connections that are currently inbound.
iisSmtpCurOutConn	Number of connections that are currently outbound.
iisSmtpTtlInConn	Total number of inbound connections received.
iisSmtpTtlOutConn	Total number of outbound connections attempted.
iisSmtpLocQueLen	Number of messages in the local queue.
iisSmtpRtryQueLen	Number of messages in the local retry queue.
iisSmtpRemtQueLen	Number of messages in the remote queue.
iisSmtpRemtRtryQueLen	Number of messages in the remote retry queue for remote delivery.
iisSmtpPicupQueLen	Number of messages in the directory pickup queue.

NNTP Group

The NNTP group provides configuration information, performance metrics, and counters for the IIS NNTP service. This group contains information such as transfer statistics and article statistics. It looks at the NNTP service as a whole and provides totals of all NNTP service activities in IIS.

NNTP Configuration Group

The NNTP Configuration group provides version information for the IIS NNTP service.

1921
2

NNTP Performance Group

The NNTP Performance group provides performance metrics and counters for the IIS NNTP service, including total current users, total users, total KB transferred, and total articles transferred. Table 3-7 defines important NNTP Performance metrics.

Table 3-7: Selected MIB Objects – IIS NNTP Performance Group (Page 1 of 2)

MIB Object	Description
iisNntpLogSize	Estimate of the current disk space used by the NNTP service logs.
iisNntpTtlDiskSize	Total size in KB of log and service directories.
iisNntpTtlLogonAtempt	Total number of logins that have been attempted to the NNTP service.
iisNntpTtlCurAnonUsr	Total number of current anonymous users maintained by the NNTP service.
iisNntpTtlCurNAnonUsr	Total number of current non-anonymous users maintained by the NNTP service.
iisNntpTtlCurUsr	Total number of users maintained by the NNTP service.
iisNntpTtlMaxAnonUsr	Maximum number of anonymous users maintained by the NNTP service.
iisNntpTtlMaxNAnonUsr	Maximum number of non-anonymous users maintained by the NNTP service.
iisNntpTtlMaxUsr	Maximum number of users maintained by the NNTP service.
iisNntpTtlAnonUsr	Running count of anonymous users maintained by the NNTP service.
iisNntpTtlNAnonUsr	Running count of non-anonymous users maintained by the NNTP service.
iisNntpTtlUsr	Running count of all users maintained by the NNTP service.
iisNntpTtlBytesSent	Count of KB sent by the NNTP service.
iisNntpTtlBytesRecv	Count of KB received by the NNTP service.
iisNntpTtlBytes	Count of KB transferred by the NNTP service.

19211
3

Table 3-7: Selected MIB Objects – IIS NNTP Performance Group (Page 2 of 2)

MIB Object	Description
iisNntpTtlArtclSent	Total articles sent for the NNTP service.
iisNntpTtlArtclRecv	Total articles received for the NNTP service.
iisNntpTtlArtcl	Total articles for the NNTP service.

Extensions Group

The Extensions group provides performance metrics and counters that extend IIS services. This group contains metrics for ASP, CGI, and ISAPI.

ASP Group

The ASP group contains ASP extension configuration and performance data. ASP is a server-side scripting environment used in Web-based applications. This group looks at ASP as a whole and provides totals of all ASP activities in IIS.

ASP Performance Group

The ASP Performance group provides performance metrics and counters for Active Server Pages, including total requests, total requests that succeeded, and total number of script errors. These metrics include those useful for real-time management and longer-term capacity planning and trend analysis. Table 3-8 defines important ASP Performance metrics.

Table 3-8: Selected MIB Objects – IIS ASP Performance Group (Page 1 of 2)

MIB Object	Description
iisAspMemAlcated	Total memory in KB that is currently allocated to ASP.
iisAspReqExecuting	Current number of ASP requests that are executing.
iisAspReqQued	Current number of ASP requests that are waiting for service from the queue.
iisAspReqDiscnt	Number of requests that were disconnected due to communications failure.
iisAspReqNAuth	Number of requests that failed due to insufficient access rights.

Table 3-8: Selected MIB Objects – IIS ASP Performance Group (Page 2 of 2)

MIB Object	Description
iisAspReqNFnd	Number of requests for files that were not found.
iisAspReqRej	Number of requests that were not executed because of insufficient processing resources.
iisAspReqWaitTime	Number of milliseconds that the most recent request waited in the queue.
iisAspReqExecTime	Total number of milliseconds to execute the most recent request.
iisAspReqSucceeded	Number of requests that executed successfully.
iisAspReqTtl	Total number of requests since the service was started.
iisAspSessCur	Number of sessions being serviced.
iisAspSessDur	Number of milliseconds that the most recent session persisted.
iisAspSessTmdOut	Number of sessions that timed out.
iisAspSessTtl	Number of sessions since the service was started.
iisAspTrnsPending	Number of transactions in progress.
iisAspTrnsCommitted	Number of transactions committed.
iisAspTrnsAbort	Number of transactions aborted.
iisAspTrnsTtl	Total number of transactions since the service started.
iisAspErrScrptRun	Number of requests that failed due to runtime errors.
iisAspErrScrptCmp	Number of requests that failed due to script-compilation errors.
iisAspErrScrptTtl	Total number of script errors, both runtime and compilation.

CGI Group

The CGI group contains CGI extension configuration and performance data. CGI is a server-side gateway interface used in Web-based applications. This group looks at CGI as a whole and provides totals of all CGI activities in IIS.

19209

CGI Performance Group

The CGI Performance group provides performance metrics and counters for the Common Gateway Interface, including number of current requests, maximum number of requests, and total number of requests. These metrics include those useful for real-time management and longer-term capacity planning and trend analysis. Table 3-9 defines important CGI Performance metrics.

Table 3-9: Selected MIB Objects – IIS CGI Performance Group

MIB Object	Description
iisCgiReqCur	Number of current CGI requests.
iisCgiReqMax	Maximum number of simultaneous CGI requests.
iisCgiReqTtl	Total number of CGI requests since the service started.

ISAPI Group

The ISAPI group contains ISAPI extension configuration and performance data. ISAPI is a server-side program interface used in Web-based applications. This group looks at ISAPI as a whole and provides totals of all ISAPI activities in IIS.

ISAPI Performance Group

The ISAPI group provides performance metrics and counters for the ISAPI, including number of current requests, maximum number of requests, and total number of requests. These metrics include those useful for real-time management and longer-term capacity planning and trend analysis. Table 3-10 defines important ISAPI Performance metrics.

Table 3-10: Selected MIB Objects – IIS ISAPI Performance Group

MIB Object	Description
iisIsapiReqCur	Number of current ISAPI requests.
iisIsapiReqMax	Maximum number of simultaneous ISAPI requests.
iisIsapiReqTtl	Total number of ISAPI requests since service startup.

19208
ε

4

Using AdvantEDGE for Microsoft IIS

This chapter describes how to configure and use AdvantEDGE for Microsoft IIS. This Point module is implemented as a SystemEDGE agent plug-in. After you enable this plug-in in the sysedge.cf file and obtain a license for it, it will load automatically at SystemEDGE start time. For more information, refer to "Configuring AdvantEDGE for Microsoft IIS" on page 2-3 and "Licensing AdvantEDGE for Microsoft IIS" on page 2-3.

The AdvantEDGE for Microsoft IIS plug-in implements additional MIB objects that provide advanced information about the health and availability of IIS. AdvantEDGE for Microsoft IIS can operate with any SNMP-compliant management software, such as Concord's *e*Health suite of products, AdvantEDGE View, HP OpenView, and others. If you are using AdvantEDGE for Apache with *e*Health, refer to the *e*Health Web Help for more information about the reports that you can generate.

The default configuration settings of the AdvantEDGE for Microsoft IIS plug-in enable you to use the advanced self-monitoring capabilities of the SystemEDGE agent in conjunction with AdvantEDGE for Microsoft IIS.

Editing the SystemEDGE Configuration File

You can use AdvantEDGE View or another SNMP management tool to edit the SystemEDGE configuration file to utilize the MIB objects that exist in AdvantEDGE for Microsoft IIS with the process-monitoring, threshold-monitoring, Windows NT event-monitoring, and history-collection features of the SystemEDGE agent. All MIB objects related to AdvantEDGE for Microsoft IIS exist at object identifier (OID) branch 1.3.6.1.4.1.546.16.2 in the Concord Systems Management MIB. The MIB is defined in the iismod.asn1 file, which is available in the AdvantEDGE for Microsoft IIS product installation.

19207 ~~ε~~

Assigning Entry Rows for the SystemEDGE Self-Monitoring Tables

All SystemEDGE self-monitoring tables (for example, log monitoring, Windows NT event monitoring, process/service monitoring, threshold monitoring, and history collection) require the use of unique row numbers. Each table contains an *Index* column which acts as a *key field* to distinguish rows in the table. This section describes the benefits of reserving a block of rows (somewhere in the range of 11 to the maximum number of rows in your table) for use by the system or application administrator.

Setting Local Policy

You may choose, as a matter of local policy, to reserve a block of rows for system administration. This policy allows you to define row entries within a reserved block of rows without worrying about the row already being taken by another user's entry. In compliance with the local policy, all other users should use row indices that are outside of the reserved range when they define user-configured entries.

By reserving a block of rows, you can define a consistent set of conditions (row entries) to be monitored across all machines such that the same condition is defined in the same row number on each of the machines. For example, you can use row 3000 in each table to define entries monitoring the page-not-found errors (iisWwwTtlNotFoundErr). You can then distribute this configuration out to every host so that every machine running IIS uses row 3000 for monitoring page-not-found errors, whether it is the threshold monitoring table or the history table. Further, every machine can use row 3000 for monitoring the WWW service in the process/service monitoring table.

Reserving Blocks of Rows

To reserve a block of rows for monitoring Microsoft IIS:

1. Decide on a block of rows that you want to reserve for use with monitoring Microsoft IIS.
2. Using that block of rows, define a set of row entries for each of the respective SystemEDGE self-monitoring tables. For more information, refer to the chapter on self-monitoring in the *SystemEDGE Agent User Guide*.

19206
E

3. Distribute configuration file entries out to all hosts that are running IIS and AdvantEDGE for Microsoft IIS. For more information, refer to the *Automating the Deployment of SystemEDGE and AdvantEDGE Point Plug-in Modules* white paper.

NOTE

As an alternative, you can use this row-number assignment policy with AdvantEDGE View for group configuration operations.

4. Require end users to avoid your block of rows when defining their own self-monitoring table entries.

Using the SystemEDGE Self-Monitoring Features

The examples in this section show SystemEDGE configuration-file commands for monitoring Microsoft IIS. Add these commands to the sysedge.cf file to enable monitoring of the MIB objects that they specify. Modify these examples as necessary to monitor the MIB objects that are relevant for your configuration.

The examples in the following sections present row numbers in the 5000 range; select a row number for your configuration that conforms to local policies. For more information on row assignment, refer to “Assigning Entry Rows for the SystemEDGE Self-Monitoring Tables” on page 4-2.

The following command, for example, instructs the SystemEDGE agent to monitor whether the IIS process is alive every 30 seconds and to store the data in row 5000 of the Process Monitoring table:

```
watch process procAlive 'inetinfo|INETINFO' 5000 0x0
  30 'IIS' ''
```

For more information about the syntax for the commands in this section, refer to the *SystemEDGE Agent User Guide.*

NOTE

Enter the commands throughout this chapter as one line. Do not use a carriage return to match the formatting shown here.

19205
Ce

Using SystemEDGE Process Monitoring

This section provides examples of how to use the SystemEDGE process- and service-monitoring capabilities to monitor the availability of crucial IIS processes and services. Add the commands in the following sections to the sysedge.cf file to monitor these processes. For more information on the SystemEDGE process- and service-monitoring capabilities, refer to the chapter on process and service monitoring in the *SystemEDGE Agent User Guide.*

Monitoring IIS

To make sure IIS is running, enter the following command:

```
watch process procAlive 'inetinfo|INETINFO' 5000 0x0
  30 'IIS' ''
```

Monitoring the IIS WWW Service

To make sure the IIS WWW Service is running, enter the following command:

```
watch ntservice 'World Wide Web Publishing Service' 5002 0x0
  30 'IIS WWW Service' ''
```

Monitoring the IIS FTP Service

To make sure the IIS FTP Service is running, enter the following command:

```
watch ntservice 'FTP Publishing Service' 5003 0x0 30 'IIS FTP
  Service' ''
```

Monitoring the IIS SMTP Service

To make sure the IIS SMTP Service is running, enter the following command:

```
watch ntservice 'Microsoft SMTP Service' 5004 0x0 30 'IIS SMTP
  Service' ''
```

Monitoring the IIS NNTP Service

To make sure the IIS NNTP Service is running, enter the following command:

```
watch ntservice 'Microsoft NNTP Service' 5005 0x0 30 'IIS NNTP
  Service' ''
```

19204
ε

Using SystemEDGE Threshold Monitoring

This section outlines the use of SystemEDGE threshold-monitoring capabilities to monitor important IIS metrics. Add the commands in the following sections to the sysedge.cf file to monitor thresholds for these MIB objects. For more information on SystemEDGE threshold monitoring, refer to the chapter on threshold monitoring in the *SystemEDGE Agent User Guide*.

NOTE

The choice of thresholds used in these examples may not be appropriate for your configuration; select thresholds that are appropriate for your environment.

Monitoring the Total Current Users

To monitor the total number of current users, enter the following command:

```
monitor oid iisTtlCurUsr.0 5002 0x0 60 delta > 100 'Total
  Current Users exceeds threshold'
```

Monitoring the KB Transferred

To monitor the number of KB transferred, enter the following command:

```
monitor oid iisTtlBytes.0 5003 0x0 60 delta > 524288 'Total
  KBytes Transferred exceeds threshold'
```

Monitoring the Number of WWW 404 Errors

To monitor the number of WWW 404 (page-not-found) errors, enter the following command:

```
monitor oid iisWwwTtlNotFoundErr.0 5004 0x0 60 delta > 100
  'WWW Total 404 Errors exceeds threshold'
```

Monitoring the Number of Web Application Requests

To monitor the number of Web application requests, enter the following command:

```
monitor oid iisWebAppReqTtl.0 5005 0x0 60 delta > 100 'Web
  Application Requests exceeds threshold'
```

Monitoring the Number of ASP Script Errors

To monitor the number of ASP script errors, enter the following command:

```
monitor oid iisAspErrScript.0 5006 0x0 60 absolute > 40 'ASP
  Total Script Errors exceeds threshold'
```

Monitoring WWW 404 Errors on an Individual Web Site

To monitor the number of WWW 404 (page not found) errors on an individual Web site, enter the following command:

```
monitor oid ntRegPerf.50.0 5007 0x0 60 absolute > 200 'WWW
  Total 404 Errors on WebA exceeds threshold'
```

NOTE

If you are monitoring WWW 404 errors on an individual Web site, you must use the corresponding ntRegPerf item. Refer to "Using the Windows NT RegPerf Extensions" on page 4-12.

Monitoring the Number of Web Logon Failures Over Time

To monitor for 20 Web Logon Failure events in the system log within one minute, enter the following command:

```
monitor oid ntEventMonMatches.5012 50010 0x0 60 delta > 20
  'More than 20 Logon Failures in Web Server in 1 Min' ''
```

NOTE

To use this Web Logon Failure example, you need a corresponding Windows NT Event Monitor 5012 to monitor Web Logon Failures. For the corresponding entry, refer to "Monitoring IIS Web Logon Failure Events for IIS WWW Service in the System Event Log" on page 4-11.

Monitoring the Number of FTP Logon Failures Over Time

To monitor for 20 FTP Logon Failure events in the system log within one minute, enter the following command:

```
monitor oid ntEventMonMatches.5013 50011 0x0 60 delta > 20
  'More than 20 Logon Failures on FTP Server in 1 Min' ''
```

NOTE

To use this FTP Logon Failure example, you need a corresponding Windows NT Event Monitor 5013 looking at FTP Logon Failures. For the corresponding entry, refer to "Monitoring IIS FTP Logon Failure Events for IIS FTP Service in the System Event Log" on page 4-11.

19202
ε

Using SystemEDGE History Collection

This section outlines the use of SystemEDGE history capabilities to track the value of important IIS metrics over time. Add the commands in the following sections to the sysedge.cf file to collect history for these MIB objects. For more information, refer to the chapter on history collection in the *SystemEDGE Agent User Guide.*

NOTE

The number of samples and the interval between samples used in these examples may not be appropriate for your IIS system; select values that are appropriate for your environment.

Collecting History for Current Number of Users for the IIS Service

To collect history for the current number of users for the WWW service, enter the following command:

```
emphistory 5002 60 iisWwwCurUsrs.0 400 'WWW Num Current Users
  History'
```

Collecting History for Current Number of Users for the IIS FTP Service

To collect history for the current number of users for the FTP service, enter the following command:

```
emphistory 5003 60 iisFtpCurUsrs.0 200 'FTP Num Current Users
  History'
```

19201
ε

Figure 4-1 shows a sample history for total current users.

Object Identifier: iisTtlCurUsr

Variable is of Type **Gauge**. The following graph shows **absolute** values.

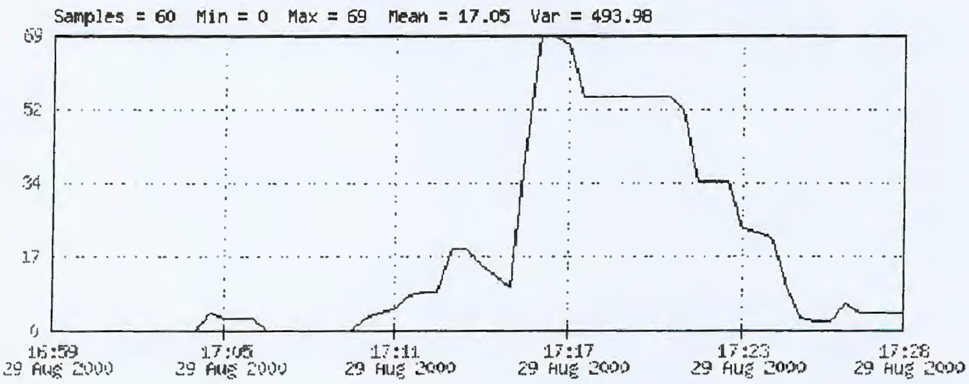

Figure 4-1: IIS Total Current Users Count History

Collecting History for Local Queue Length for the SMTP Service

To collect history for the local queue length for the SMTP service, enter the following command:

```
emphistory 5004 60 iisSmtpLocQueLen.0 480 'SMTP Local Queue
  Len History'
```

Collecting History for the Local Retry Queue Length for the SMTP Service

To collect history for the local retry queue length for the SMTP service, enter the following command:

```
emphistory 5005 60 iisSmtpRtryQueLen.0 480 'SMTP Local Retry
  Queue Len History'
```

Using SystemEDGE Windows NT Event Monitoring

This section outlines the use of SystemEDGE Windows NT event monitoring to capture important IIS-related Windows NT events and forward them to the appropriate management software as SNMP traps. The examples in this section show SystemEDGE configuration-file commands that instruct the SystemEDGE agent to monitor for certain types of Windows NT events that are related to IIS. Add the commands in the following sections to the sysedge.cf file to monitor for these Windows NT events. For more information, refer to the chapter on Windows NT event monitoring in the *SystemEDGE Agent User Guide.*

Monitoring IIS WWW Service Events in the System Event Log

To monitor for WWW service events in the system event log, enter the following command:

```
watch ntevent 5000 0x00 System All 'W3SVC' '.*' 'Monitor IIS
  WWW Service System Events' ''
```

Monitoring IIS WWW Service Events in the Security Event Log

To monitor for WWW service events in the security event log, enter the following command:

```
watch ntevent 5001 0x00 Security All 'W3SVC' '.*' 'Monitor IIS
  WWW Service Security Events' ''
```

Monitoring IIS WWW Service Events in the Application Event Log

To monitor for WWW service events in the application event log, enter the following command:

```
watch ntevent 5002 0x00 Application All 'W3SVC' '.*' 'Monitor
  IIS WWW Service Application Events' ''
```

Monitoring IIS FTP Service Events in the System Event Log

To monitor for FTP service events in the system event log, enter the following command:

```
watch ntevent 5003 0x00 System All 'MSFTPSVC' '.*' 'Monitor
  IIS FTP Service System Events' ''
```

19199
ε

Monitoring IIS FTP Service Events in the Security Event Log

To monitor for FTP service events in the security event log, enter the following command:

```
watch ntevent 5004 0x00 Security All 'MSFTPSVC' '.*' 'Monitor
  IIS FTP Service Security Events' ''
```

Monitoring IIS FTP Service Events in the Application Event Log

To monitor for FTP service events in the application event log, enter the following command:

```
watch ntevent 5005 0x00 Application All 'MSFTPSVC' '.*' 'Monitor
IIS FTP Service Application Events' ''
```

Monitoring IIS SMTP Service Events in the System Event Log

To monitor for SMTP service events in the system event log, enter the following command:

```
watch ntevent 5006 0x00 System All 'SMTPSVC' '.*' 'Monitor IIS
  SMTP Service System Events' ''
```

Monitoring IIS SMTP Service Events in the Security Event Log

To monitor for SMTP service events in the security event log, enter the following command:

```
watch ntevent 5007 0x00 Security All 'SMTPSVC' '.*' 'Monitor
  IIS SMTP Service Security Events' ''
```

Monitoring IIS SMTP Service Events in the Application Event Log

To monitor for SMTP service events in the application event log, enter the following command:

```
watch ntevent 5008 0x00 Application All 'SMTPSVC' '.*'
  'Monitor IIS SMTP Service Application Events' ''
```

19198

Monitoring IIS NNTP Service Events in the System Event Log

To monitor for NNTP service events in the system event log, enter the following command:

```
watch ntevent 5009 0x00 System All ’NNTPSVC’ ’.*’ ’Monitor IIS
  NNTP Service System Events’ ’’
```

Monitoring IIS NNTP Service Events in the Security Event Log

To monitor for NNTP service events in the security event log, enter the following command:

```
watch ntevent 5010 0x00 Security All ’NNTPSVC’ ’.*’ ’Monitor
  IIS NNTP Service Security Events’ ’’
```

Monitoring IIS NNTP Service Events in the Application Event Log

To monitor for NNTP service events in the application event log, enter the following command:

```
watch ntevent 5011 0x00 Application All ’NNTPSVC’ ’.*’
  ’Monitor IIS NNTP Service Application Events’ ’’
```

Monitoring IIS Web Logon Failure Events for IIS WWW Service in the System Event Log

To monitor for Web Logon Failure events in the system event log, enter the following command:

```
watch ntevent 5012 0x00 System All ’W3SVC’
  ’.*Logon.*failure.*’ ’Web Logon Failure’ ’’
```

Monitoring IIS FTP Logon Failure Events for IIS FTP Service in the System Event Log

To monitor for FTP Logon Failure events in the system event log, enter the following command:

```
watch ntevent 5013 0x00 System All ’MSFTPSVC’
  ’.*Logon.*failure.*’ ’FTP Logon Failure’ ’’
```

19197
ɛ

Monitoring IIS FTP Timeouts for IIS FTP Service in the System Event Log

To monitor for FTP Timeout events in the system event log, enter the following command:

```
watch ntevent 5014 0x00 System All 'MSFTPSVC'
  '.*User.*timed-out.*' 'FTP User Time-out' ''
```

Using the Windows NT RegPerf Extensions

This section outlines the use of SystemEDGE Windows NT RegPerf extension-monitoring capabilities to capture extended IIS-related information. The examples in this section show the SystemEDGE configuration-file commands that instruct the SystemEDGE agent to provide information on an individual Web site. You can use this variable with threshold monitoring to provide traps on specific user sites. Place these commands in the SystemEDGE configuration file, sysedge.cf.

Add the commands in the following sections to the sysedge.cf file to monitor these Windows NT RegPerf extensions. The following examples use values of 50 and 51; select values for your system that conform to local policies.

For more information, refer to the chapter on Windows NT Registry MIB objects in the *SystemEDGE Agent User Guide*.

Monitoring IIS WWW 404 Errors

To monitor for WWW 404 errors on Web site A (WebA), enter the following command:

```
ntregperf 50 Integer Performance 'Web Service' 'Total Not
  Found Errors' 'WebA'
```

Monitoring Unauthorized ASP Requests

To monitor for ASP requests that are not authorized on Web site B (Web B), enter the following command:

```
ntregperf 51 Integer Performance 'Active Server Pages'
  'Requests Not Authorized' 'WebB'
```

19196

Index

A

AdvantEDGE for Microsoft IIS
- components 1-2
- files 2-3
- installation requirements 2-1
- overview 1-1

ASP group 3-13

assignment policy for monitoring tables 4-2

C

CGI group 3-14

collecting history
- FTP current users 4-7
- SMTP
  - local queue length 4-8
  - local retry queue length 4-8
- WWW current users 4-7

collecting history information 4-7

configuring AdvantEDGE for Microsoft IIS
- startup 2-3

E

eHealth 1-3

evaluation copy, licensing 2-4

example file 2-3

extensions group 3-13

F

files installed 2-3

footprint 3-3

FTP
- event monitoring 4-9
- MIB objects 3-9
- monitoring 4-4
  - logon failures 4-6, 4-11
  - user timeouts 4-12

G

generating a license 2-4

H

history collection 4-7

I

IIS
- architecture 1-2
- footprint 1
- monitoring 4-4
- service packs 2-1
- services 1-1
- version 2-1

iismod.asn1 file 2-3

iismod.dll 2-3

iismod.pdf file 2-3

installation
- files 2-3
- requirements 2-1

installing
- AdvantEDGE IIS 2-2
- SystemEDGE agent 2-1

19195
6

- Internet Information Server. *See* IIS
- Internet Information Services. *See* IIS
- ISAPI group 3-15

L

- license
  - generating 2-4
  - key 2-6
  - sample 2-6
- licensing AdvantEDGE for Microsoft IIS 2-3

M

- Management Information Base. *See* MIB
- MIB
  - configuration section 3-2
  - extension 3-13
  - performance section 3-3
  - services 3-6
  - specification 3-1
  - using 4-1
- monitoring
  - application event log
    - FTP service events 4-10
    - NNTP service events 4-11
    - SMTP service events 4-10
    - WWW service events 4-9
  - ASP script errors 4-6
  - assignment policy for tables 4-2
  - FTP service 4-4
  - IIS process 4-4
  - KB transferred 4-5
  - logon failures
    - FTP 4-6
    - WWW 4-6
  - NNTP service 4-4
  - page not found errors 4-5, 4-6, 4-12
  - processes 4-4
  - security event log
    - FTP service events 4-10
    - NNTP service events 4-11
    - SMTP service events 4-10
    - WWW service events 4-9
  - services 4-4
  - SMTP service 4-4
  - system event log
    - FTP Logon failure 4-11
    - FTP service events 4-9
    - NNTP service events 4-11
    - SMTP service events 4-10
    - WWW logon failure 4-11
    - WWW service events 4-9
  - thresholds 4-5
  - total current users 4-5
  - unauthorized ASP requests 4-12
  - Web application requests 4-5
  - Windows NT RegPerf extensions
    - unauthorized ASP requests 4-12
    - WWW 404 errors 4-12
  - WWW
    - 404 errors 4-5, 4-6, 4-12
    - service 4-4

N

- NNTP
  - event monitoring 4-11
  - MIB objects 3-12
  - monitoring 4-4

O

- obtaining a license 2-4

P

- process monitoring 4-4

19194
ε

R

- RegPerf extension 4-12
- relnotes.txt file 2-3
- requirements for installation 2-1
- reserving rows in monitoring tables 4-2

S

- server
  - configuration 3-2
  - performance 3-5
- service monitoring 4-4
- setting policy for monitoring tables 4-2
- SMTP
  - event monitoring 4-10
  - MIB objects 3-10
  - monitoring 4-4
- sysedge.cf file 2-3
- sysedge.lic file 2-4, 2-6
- SystemEDGE agent 1-1
  - configuring 2-3
  - history capability 4-7
  - monitoring processes 4-1
  - monitoring Windows NT events 4-9
  - monitoring Windows NT RegPerf extension 4-12
  - threshold monitoring 4-5

T

- temporary license 2-4
- threshold monitoring 4-5

U

- using AdvantEDGE for Microsoft IIS 4-1
  - with eHealth 1-3

W

- Web logon failures 4-6, 4-11
- Windows NT event monitoring 4-9
- WWW
  - event monitoring 4-9
  - monitoring 4-4

19193
E

ANEXO CONCORD SYSTEMEDGE (PARTE 2)

COMPROVAÇÃO DAS ESPECIFICAÇÕES EXIGIDAS NO EDITAL

2.1. ASPECTOS GERAIS

REQUISITO	ATRIBUTOS OFERTADOS	ATRIBUTOS OFERTADOS ADICIONALMENTE	CONFIRMA ATENDIMENTO (SIM / NÃO)	PÁGINA DA DOCUMENTAÇÃO TÉCNICA
2.1 – Co mprovação de Performance para os Equipamentos da plataforma RISC				
2.2 – Co mprovação de Performance para os Equipamentos da plataforma INTEL				
2.3 – Requisitos Gerais				
2.4 – Garantia				
2.5 – Alimentação Elétrica				
2.6 – Assistência técnica				
2.7 – Recursos Mínimos de Particionamento para os Servidores RISC				
2.8 – Recursos Mínimos de Gerenciamento do Hardware requeridos para os Servidores INTEL				
2.9 – Recursos Mínimos de hardware e software para os Servidores INTEL	**107 licenças de Enhanced SystemEdge**		**SIM**	
2.10 – Recursos Mínimos de hardware e software para os servidores RISC adicionais	**88 licenças de Enhanced SystemEdge**		**SIM**	
2.11 – Co mprovação dos Requisitos Técnicos				
2.12 – Organização da Documentação Técnica				
2.13 – Suporte Remoto				
2.14 – Configuração das Ferramentas	**1) Instalação e configuração**	**2) Possibilidade**		2) Manual **eHealth Service**

Os itens não preenchidos estão comprovados nos demais anexos.

de Gerenciamento	dos agentes Enhanced SystemEdge de acordo com os padrões da plataforma de Gerência já existente na ECT; 2) Licenças e configuração do Concord Service Availability para gerenciar DNS, HTTP, HTTPS, SMTP, POP3, FTP e TCP; 3) Serão configurados as seguintes extensões do Concord SystemEdge: Exchange, Oracle, SQL, Apache e IIS; *4) A ser respondido pela Integradora* 5) Serão realizados os serviços de integração dos eventos dos elementos desta licitação com o HP OpenView OVO e NNM; Criação e integração de regras no Live Health da Concord, enviando alarmes para o ambiente HP OpenView; Criação de relatórios e agendamento de geração de relatórios do Concord eHealth.	de gerenciamento de PING e NNTP;		**Response User Guide (Arquivo svcrsp.pdf)**, página 11. Obs: Service Availability é o novo nome comercial para o Service Response; 3) Manuais: **ADVANTEDGE FOR MICROSOFT® EXCHANGE User Guide (Arquivo ExchangeUser.pdf)**, página 1-1; **eHealth AIM for Oracle User Guide (Arquivo OracleUser.pdf)**, página 11; **ADVANTEDGE FOR MICROSOFT SQL SERVER User Guide (Arquivo SQLUser.pdf)**, página 1-1; **eHealth AIM for Apache User Guide (Arquivo ApacheUser.pdf)**, página 12; **ADVANTEDGETM FOR MICROSOFT® IIS User Guide (Arquivo IISUser.pdf)**, página 1-1.

Os itens não preenchidos estão comprovados nos demais anexos.

19190
ε

ÍNDICE
(Parte I e II)

eHealth AIM for Apache - *User Guide*

AdvantEDGE for Microsoft Exchange – *User Guide*

AdvantEDGE for Microsoft Internet Information Services – *User Guide*

eHealth AIM for Oracle – *User Guide*

AdvantEDGE for Microsoft SQL Server – *User Guide*

eHealth Service Response – *User Guide*

19189

eHealth AIM for Oracle User Guide

09-16060-003
December 2001

19188
e

Important Notice

Concord Communications, Inc., *e*Health, *e*Health Suite, the Concord Logo, eROI, AdvantEDGE, SystemEDGE, Live Health, Network Health, Live Status, System Health, Application Health, Automating Technology Management, Constellation, Enterprise, Enterprise Monitor, Firstsense, FirstSense and design, FirstSense Enterprise, Pulse, Pulsecheck, Token/Net, Token/Scope, We See It Happening, Fault Manager, Empire, Empire Technologies and/or other Concord marks or products referenced herein are either registered trademarks or trademarks of Concord Communications, Inc. Other trademarks are the property of their respective owners.

Adobe, Acrobat, and the Acrobat logo are trademarks, which may be registered in certain jurisdictions, and PostScript is a registered trademark of Adobe Systems Incorporated.
CA-OpenIngres and Ingres II are registered trademarks of Computer Associates International, Inc.
Hewlett-Packard, HP OpenView, HP, and HP-UX are trademarks or registered trademarks of Hewlett-Packard Company in the United states and other countries.
Intel is a registered trademark of Intel Corporation.
Motif is a trademark of Open Systems Foundation, Inc.
Oracle is a registered trademark of Oracle Corporation.
Sun, Sun-4, SunOS, SunNet, and OpenWindows are trademarks, and Sun Microsystems, Sun Workstation, and NFS are registered trademarks of Sun Microsystems, Inc.
SPARCstation and SPARCserver are trademarks, and SPARC is a registered trademark of SPARC International, Inc.
UNIX is a registered trademark in the United States and other countries licensed exclusively through X/Open Company Ltd.
Windows and Windows NT are trademarks of Microsoft Corporation.

© Copyright Bigelow and Holmes 1986, 1985. Lucida is a registered trademark of Bigelow & Holmes. Sun Microsystems Inc., AT&T, and Bigelow & Holmes make no representations about the suitability of the source code for any purpose. It is provided "as is" without express or implied warranty of any kind.

General Notice: Some of the product names used herein have been used for identification purposes only and may be trademarks of their respective companies.

Proprietary Notice

The information and descriptions contained herein are the property of Concord Communications, Inc. Such information and descriptions may not be copied, disseminated, or distributed without the express written consent of Concord Communications, Inc. Concord Communications, Inc., assumes no responsibility for any inaccuracies that may appear in this document. Concord Communications, Inc., reserves the right to improve its products and change specifications at any time without notice.

U. S. Government Restricted Rights

Use, reproduction, and disclosure by the U.S. Government are subject to the restrictions set forth in FAR §52.227-19 (c) (1) and (2) and DFARS §252.227-7013 (c) (1) (ii).

Patent Information

U. S. Patent 5,615,323
Patents Pending

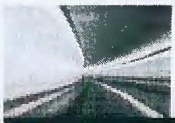

Table of Contents

About This Guide 7

Audience. 7
About This Guide 8
  Revision Information 8
  Documentation Conventions 8
Technical Support 9
Professional Services 10

Chapter 1 **Introduction** 11

Introducing *e*Health AIM for Oracle. 11
Features 12
Using *e*Health AIM for Oracle 13
  Using *e*Health AIM for Oracle with AdvantEDGE View 13
  Using *e*Health AIM for Oracle with *e*Health 15
  Using *e*Health AIM for Oracle with Live Health 16

Chapter 2 **Installing *e*Health AIM for Oracle** 17

Installation Requirements. 17
Installing the Software. 18
*e*Health AIM for Oracle Files 20

19186
ε

Configuring *e*Health AIM for Oracle 20
- Editing the sysedge.cf File. 21
- Editing the oramod.cf File 22
- Editing the SystemEDGE Startup File. 24

Licensing *e*Health AIM for Oracle 27
- Obtaining a License. 27
- Generating the License through the Web-based License Form. 28
- Generating a License through AdvantEDGE View Event Processing 30
- Generating a License through AdvantEDGE View Host Administration 32
- Sample License File. 33

Chapter 3 **Using the *e*Health AIM for Oracle MIB** **35**

Configuration Section 36
- Database Configuration Group 36
- Database File Configuration Group 39
- Redo Log File Configuration Group 40
- System Global Area Configuration Group 41

Performance Section 41
- Oracle Footprint Group 41
- Metrics Group 45
- SGA Group 48
- Redo Logs Group 51
- Rollback Segment Group 51
- Waits Group 53
- Expensive SQL Group. 54
- Tablespace Group 55
- Lock Group 55
- Backup Group 57
- Archive Group. 58

Chapter 4 **Using *e*Health AIM for Oracle** 59

*e*Health AIM for Oracle MIB Branch 60
Assigning Entry Rows in the SystemEDGE Self-Monitoring Tables 60
Setting Local Policy 60
Reserving Blocks of Rows 60
Using the SystemEDGE Self-Monitoring Features. 61
Using SystemEDGE Process Monitoring 62
Using SystemEDGE Threshold Monitoring. 63
Using SystemEDGE History Collection 64

Glossary 67

Index 77

19184

About This Guide

This guide describes how to install and use the *e*Health application insight module (AIM) for Oracle. It is intended for the person responsible for installing and configuring *e*Health AIM for Oracle. This guide supports *e*Health AIM for Oracle Release 1.3 Patchlevel 1 or later, and *e*Health SystemEDGE Release 4.0, Patchlevel 3 or later. *e*Health AIM for Oracle supports the following operating systems:

- Sun Solaris (SPARC) Release 2.6 or later in *32-bit mode* and one of the following versions of Oracle in *32-bit mode*:
  - Oracle 8.0.4, 8.0.5, 8.0.6
  - Oracle 8i (8.1.5, 8.1.6, or 8.1.7)
- AIX Releases 4.2 and 4.3 on Oracle 8i
- HP-UX Releases 11.0 and 11.1 on Oracle 8i

Audience

This guide is intended for the person who is installing and using *e*Health AIM for Oracle. To use *e*Health AIM for Oracle, you should have a basic understanding of the Oracle database, *e*Health SystemEDGE, and your host's operating system environment. For more information, refer to Oracle documentation (http://www.oracle.com) and the *eHealth SystemEDGE User Guide*.

About This Guide

This section describes the changes and enhancements that have been made since the last release of this guide. It also includes the documentation conventions used in this guide.

Revision Information

This guide describes information that is specific to *e*Health AIM for Oracle 1.0 Patchlevel 2. The following is new in this version of the guide:

- Installation instructions for the AIX and HP-UX versions of *e*Health AIM for Oracle
- Licensing through AdvantEDGE View Event Processing and Host Administration
- Glossary

Documentation Conventions

Table 1 lists the conventions used in this document.

Table 1. Documentation Conventions (Page 1 of 2)

Convention	Description
File or Directory Name	File or directory names.
`code`	System, code, or operating system command line examples.
emphasis	Emphasis and guide titles.
`enter`	Text that you must type exactly as shown.
Name	Text that defines menus, fields in dialog boxes, or keyboard keys.
New Term	A new term, that is, one that is being introduced.
`Variable`	Variable values that you substitute.

19182
4

Table 1. Documentation Conventions (Page 2 of 2)

Convention	Description
→	A sequence of menus or menu options. For example, **File** → **Exit** means "Choose **Exit** from the **File** menu."
NOTE	Important information, tips, or other noteworthy details.
CAUTION	Information that helps you avoid data corruption or system failures.
WARNING	Information that helps you avoid physical danger.

Technical Support

If you need any assistance with this product, contact Technical Support at the following:

Phone:	(888) 832-4340
	(508) 303-4300
Fax:	(508) 303-4343
E-mail:	support@concord.com
Web site:	http://www.concord.com

Professional Services

If you need any assistance with customizing this product, contact Professional Services at the following:

Phone:	(800) 851-8725 (Choose option 7)
Fax:	(508) 486-4555
E-mail:	proserv@concord.com
Web site:	http://www.concord.com

19180

1

Introduction

This chapter provides an overview of *e*Health AIM for Oracle.

Introducing *e*Health AIM for Oracle

*e*Health AIM for Oracle is a plug-in for the SystemEDGE agent that enables information technology (IT) operators to monitor the health and availability of Oracle databases and servers. You can configure this plug-in to monitor the Oracle processes and features that are relevant to your organization. You can also configure *e*Health AIM for Oracle to alert you to any potential issues with the application or the system on which it is running before those issues become problems.

To use *e*Health AIM for Oracle, you must install it on every Oracle system that you want to monitor. For more information, refer to Chapter 2, “Installing eHealth AIM for Oracle.”

NOTE

This guide is not intended to describe how to install, administer, or use Oracle databases. For help with Oracle, refer to your Oracle documentation.

fl 1479

Features

*e*Health AIM for Oracle monitors the following:

- Oracle database configuration
- Oracle log files (including alert, redo, database, and configuration logs)
- Performance metrics
- Database response and availability

*e*Health AIM for Oracle monitors database process attributes. For example, it monitors whether each process is alive; and it can restart processes, if necessary. In addition, it monitors memory use, log sizes, number of transactions, and efficiency of the library and the data dictionary.

Because the Oracle application records error information in alert log files, *e*Health AIM for Oracle can use the log-file monitoring capability of the SystemEDGE agent to scan the logs and forward certain events as SNMP traps when appropriate.

For more information about the management information base (MIB) objects that you can monitor with *e*Health AIM for Oracle, refer to Chapter 3, "Using the eHealth AIM for Oracle MIB." For more information about configuring the SystemEDGE agent to monitor these MIB objects, refer to Chapter 4, "Using eHealth AIM for Oracle."

*e*Health AIM for Oracle is designed to monitor one or more Oracle databases of the same Oracle release that are running on a single system. To support multiple databases, the data in the *e*Health AIM for Oracle MIB is organized into tables that are indexed by database session identification (SID) number. Each table includes an entry for the SID. For more information, refer to Chapter 3, "Using the eHealth AIM for Oracle MIB."

19178
ε

Using *e*Health AIM for Oracle

*e*Health AIM for Oracle provides important information about Oracle to management software through the SystemEDGE agent and Simple Network Management Protocol (SNMP). This AIM works with the SystemEDGE agent to closely manage the Oracle application, providing real-time fault detection and automatically correcting problems, if necessary.

You can use *e*Health AIM for Oracle with any SNMP-compliant management software, including Concord's *e*Health suite of products, AdvantEDGE View, HP OpenView, and others. *e*Health AIM for Oracle and the SystemEDGE agent can provide you with the following types of information:

- Number and type of users connected to the database
- Amount of work that each transaction performs
- Database workload per transaction
- Number of client and database requests per transaction
- Rate at which application systems reference the database
- Effectiveness of the database buffer cache and library cache
- Number of changes to the database
- Amount of memory allocated to sessions
- Sort statistics and efficiency percentages
- Server statistics

Using *e*Health AIM for Oracle with AdvantEDGE View

You can use *e*Health AIM for Oracle with AdvantEDGE View to run queries for monitoring the performance, configuration, availability, and health of the Oracle application.

To run an AdvantEDGE View Application query for Oracle:

1. Select the target system or group from the **System** or **Group** list.
2. Select **Oracle** from the **Applications** list.

19177
3

3. Click the **Applications** icon.

AdvantEDGE View runs the query for the specified application on the system or group you selected.

NOTE

> If you run a query for a group of systems, AdvantEDGE View may request additional information before running the query. For more information, refer to the AdvantEDGE View Web Help.

Figure 1 shows the image map that AdvantEDGE View displays when you run an Application query for Oracle. Click the area for which you want to display information.

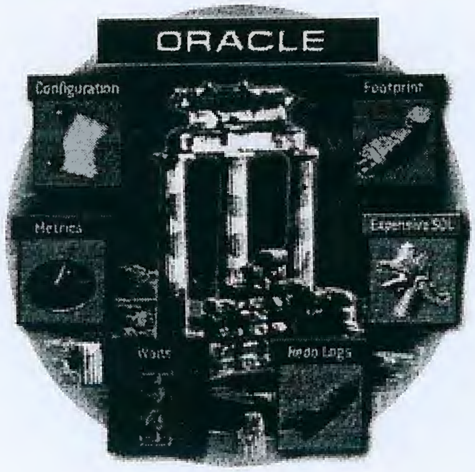

Figure 1. AdvantEDGE View Image Map for Oracle Queries

For example, if you click the **Metrics** area, AdvantEDGE View displays the Oracle metrics. Figure 2 shows a sample Metrics query for the Oracle application.

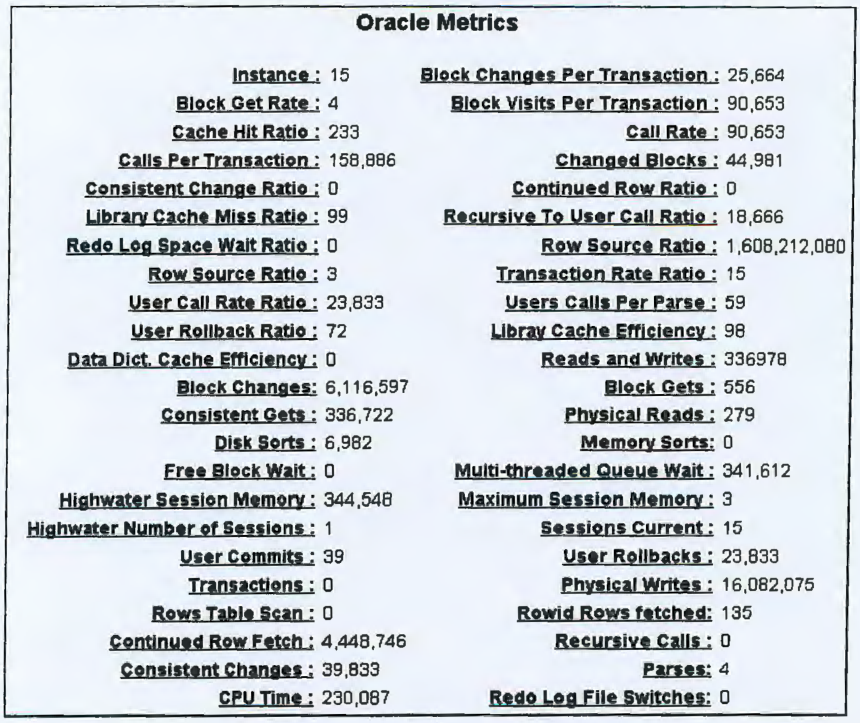

Oracle Metrics

Instance :	15	Block Changes Per Transaction :	25,664
Block Get Rate :	4	Block Visits Per Transaction :	90,653
Cache Hit Ratio :	233	Call Rate :	90,653
Calls Per Transaction :	158,886	Changed Blocks :	44,981
Consistent Change Ratio :	0	Continued Row Ratio :	0
Library Cache Miss Ratio :	99	Recursive To User Call Ratio :	18,666
Redo Log Space Wait Ratio :	0	Row Source Ratio :	1,608,212,080
Row Source Ratio :	3	Transaction Rate Ratio :	15
User Call Rate Ratio :	23,833	Users Calls Per Parse :	59
User Rollback Ratio :	72	Libray Cache Efficiency :	98
Data Dict. Cache Efficiency :	0	Reads and Writes :	336978
Block Changes:	6,116,597	Block Gets :	556
Consistent Gets :	336,722	Physical Reads :	279
Disk Sorts :	6,982	Memory Sorts:	0
Free Block Wait :	0	Multi-threaded Queue Wait :	341,612
Highwater Session Memory :	344,548	Maximum Session Memory :	3
Highwater Number of Sessions :	1	Sessions Current :	15
User Commits :	39	User Rollbacks :	23,833
Transactions :	0	Physical Writes :	16,082,075
Rows Table Scan :	0	Rowid Rows fetched:	135
Continued Row Fetch :	4,448,746	Recursive Calls :	0
Consistent Changes :	39,833	Parses:	4
CPU Time :	230,087	Redo Log File Switches:	0

Figure 2. Sample AdvantEDGE View Oracle Metrics Query

Using *e*Health AIM for Oracle with *e*Health

In previous releases, eHealth – Application Insight was called eHealth – Application Assessment.

You can use *e*Health AIM for Oracle and the SystemEDGE agent with the *e*Health product suite to provide the historical data for long-term trending analysis and capacity planning. With *e*Health – Application Insight, you can run At-a-Glance, Trend, Top N, and MyHealth reports for the following types of variables:

- Amount of Central Processing Unit (CPU), memory, and disk space that the Oracle application is using
- Size of the Oracle configuration and database logs

- Number of transactions that the Oracle application is performing each second
- Amount of CPU used, disk input/output, work performed, and user calls for each transaction
- Ratio of transaction rollbacks and failures in writing to the redo logs
- Frequency with which the Oracle application uses the database cache to redisplay information

For more information about the variables that you can monitor and reports that you can run when you integrate *e*Health AIM for Oracle with *e*Health, refer to the *e*Health Web Help.

Using *e*Health AIM for Oracle with Live Health

You can also use *e*Health AIM for Oracle and the SystemEDGE agent with Live Health for real-time detection of potential problems. Live Health applies intelligent algorithms to the data, resulting in precise assessments of application health and performance. For more information about how Live Health can detect "brownouts" and service delays across applications, systems, and networks, refer to the Live Health Web Help.

19174
ε

2

Installing eHealth AIM for Oracle

This chapter explains how to install, configure, and license *e*Health AIM for Oracle.

NOTE

> For the most current information about installing this AIM, refer to the relnotes.txt file on the *e*Health AIM for Oracle installation CD-ROM.

Installation Requirements

Before you install *e*Health AIM for Oracle, you must first install, configure, and license the SystemEDGE agent Release 4.0, Patchlevel 3 or later. For more information, refer to the *eHealth SystemEDGE User Guide.*

To use *e*Health AIM for Oracle, your system must be running one of the following:

- Sun Solaris (SPARC) Release 2.6 or later in *32-bit mode* and one of the following versions of Oracle in *32-bit mode*:
  - Oracle 8.0.4, 8.0.5, 8.0.6
  - Oracle 8i (8.1.5, 8.1.6, or 8.1.7)
- AIX Releases 4.2 and 4.3 and Oracle 8i
- HP-UX Releases 11.0 and 11.1 and Oracle 8i

19173
E

NOTE

If you are using the Solaris operating system, you *must* run both Solaris and Oracle in *32-bit mode only*.

Installing the Software

*e*Health AIM for Oracle is distributed as a tape archive (tar) file. The name of the tar file varies for each operating system, as follows:

- oramod_1.3p1_sol.tar for Solaris systems
- oramod_1.3p1_aix.tar for AIX systems
- oramod_1.3p1_hpux.tar for HP-UX systems

To install *e*Health AIM for Oracle:

1. Locate the Oracle database that you need to monitor.
2. Log on to that system as the root user.
3. Copy the tar file for your system from the CD-ROM to the /tmp directory.
4. Change directory to the SystemEDGE agent directory on your system by entering one of the following commands.

   For Solaris and HP-UX systems, enter the following:

   ```
   cd /opt/EMPsysedge
   ```

   For AIX systems, enter the following:

   ```
   cd /usr/lpp/EMPsysedge
   ```

5. Create the plugins directory, if it is not already present, by entering the following:

   ```
   mkdir plugins
   ```

19172
E

6. Change directory to the plugins directory by entering the following:

   ```
   cd plugins
   ```

7. Enter one of the following commands to launch the installation, depending on your operating system.

   For Solaris systems, enter the following:

   ```
   tar xvf /tmp/oramod_1.3pl_sol.tar
   ```

   For AIX systems, enter the following:

   ```
   tar xvf /tmp/oramod_1.3pl_aix.tar
   ```

   For HP-UX systems, enter the following:

   ```
   tar xvf /tmp/oramod_1.3pl_hpux.tar
   ```

   The installation creates files in the plugins/oramod directory.

8. Edit the sysedge.cf file as described in "Editing the sysedge.cf File" on page 21. After you install the files, you must edit the sysedge.cf file to specify the correct shared library for your version of Oracle.

9. Edit the oramod.cf file to indicate which databases, SIDs, and Oracle home directories to use, as described in "Editing the oramod.cf File" on page 22.

10. Edit the SystemEDGE startup file to include Oracle environment variables so that the agent and AIM can find the Oracle home and lib directories, as described in "Editing the SystemEDGE Startup File" on page 24.

11. Start the SystemEDGE agent, using one of the following commands.

    For Solaris systems, enter the following:

    ```
    /etc/rc2.d/S99sysedge start
    ```

    For AIX systems, enter the following:

    ```
    bin/sysedge -b
    ```

19171 ε

For HP-UX systems, enter the following:

```
/sbin/rc2.d/S990sysedge start
```

12. License the AIM, as described in “Licensing eHealth AIM for Oracle” on page 27.
13. Restart the SystemEDGE agent.

*e*Health AIM for Oracle Files

Table 2 describes the files created by the *e*Health AIM for Oracle installation procedure.

Table 2. Files Installed by *e*Health AIM for Oracle

File Name	Description
oramod.asn1	*e*Health AIM for Oracle MIB specification
oramod.cf	*e*Health AIM for Oracle configuration file
oramod.pdf	*eHealth AIM for Oracle User Guide*
oramod8.so	eHealth AIM for Oracle shared library for Oracle 8.0.4, 8.0.5, and 8.0.6
oramod8i.so	*e*Health AIM for Oracle shared library for Oracle 8i (verions 8.1.5, 8.1.6, and 8.1.7)
relnotes.txt	Release notes for the *e*Health AIM for Oracle module

Configuring *e*Health AIM for Oracle

After you install *e*Health AIM for Oracle, you must configure it by editing the following files:

- sysedge.cf (page 21)
- oramod.cf (page 22)
- SystemEDGE startup file (page 24)

Editing the sysedge.cf File

You must edit the sysedge.cf file to use the correct shared library file for your system and to enable the SystemEDGE agent to load *e*Health AIM for Oracle. You can use the **sysedge_plugin** keyword in the sysedge.cf configuration file to specify which *e*Health AIMs the SystemEDGE agent will load at system initialization. By default, the SystemEDGE agent does not load any AIMs at initialization, but you can edit the sysedge.cf file to configure the agent to load any *e*Health AIMs that you have installed.

The sysedge.cf file is located in your system directory by default; for example, it is located in the /etc/sysedge.cf directory on UNIX systems. For more information about the sysedge.cf file, refer to the *eHealth SystemEDGE User Guide*.

NOTE

To configure the SystemEDGE agent to start *e*Health AIM for Oracle, you must provide the complete pathname to the shared library file for your version of Oracle.

Enabling *e*Health AIM for Oracle for Oracle 8.0.4, 8.0.5, or 8.0.6 (Solaris)

To enable the *e*Health AIM for Oracle module in the SystemEDGE agent for Oracle 8.0.4, 8.0.5, or 8.0.6 on Solaris systems, add the following line to /etc/sysedge.cf:

```
sysedge_plugin /opt/EMPsysedge/plugins/oramod/oramod8.so
```

Enabling *e*Health AIM for Oracle for Oracle 8i (Solaris and HP-UX)

To enable the *e*Health AIM for Oracle module in the SystemEDGE agent for Oracle 8i (versions 8.1.5, 8.1.6, or 8.1.7) on Solaris and HP-UX systems, add the following line to /etc/sysedge.cf:

```
sysedge_plugin /opt/EMPsysedge/plugins/oramod/oramod8i.so
```

1a169
E

Enabling *e*Health AIM for Oracle for Oracle 8i (AIX)

To enable the *e*Health AIM for Oracle module in the SystemEDGE agent for Oracle 8i for AIX systems, add the following line to /etc/sysedge.cf:

```
sysedge_plugin /usr/lpp/EMPsysedge/plugins/oramod/oramod8i.so
```

Editing the oramod.cf File

The oramod.cf file describes the *e*Health AIM for Oracle configuration. You can edit the oramod.cf file to indicate the following:

- User name and password.
- Name (SID) of the database that you are monitoring. This value is contained in the oramodDbCfgSID MIB object.

NOTE

> If you are using Transparent Network Substrate (TNS) names, use the TNS name that matches your database name in the oramod.cf file.

- SID index. This value is also the value of the oramodDbCfgSIDINDX MIB object. This value is user-configurable. You can set it to any integer greater than zero. Each instance must have a unique index for the platform.
- SID-home. Installation directory for Oracle, or ORACLE_HOME. This value is also the value of the oramodDbCfgHOME MIB object.
- SID-base. Base directory for Oracle, or ORACLE_BASE. This value is also the value of the oramodDbCfgBASE MIB object.
- File Check Interval. This value is the interval for checking files, in seconds.
- Tablespace Check Interval. This value is the interval for checking tablespaces.

19168
4

To use *e*Health AIM for Oracle, you must have a user ID that has either DBA privileges or SELECT_CATALOG_ROLE permissions for Oracle Version 8i and later. For earlier versions of Oracle, the SELECT_ANY_TABLE permissions should suffice.

*e*Health AIM for Oracle can monitor different versions of Oracle databases on the same platform *only* if you are running multiple versions of the SystemEDGE agent on different ports. Otherwise, *e*Health AIM for Oracle can simultaneously monitor *only* instances of Oracle that are the *same Oracle version.*

Sample oramod.cf File

The following is a sample oramod.cf file.

```
# oramod.cf
# Configuration file for the eHealth AIM for Oracle Module
# Concord Communications, Copyright 2001
#
# For each Oracle database running on your system, specify the following:
#
# oramod username passwd sid sid-index sid-home sid-base
#  file-check-interval table-space-check-interval
#
oramod fred elTsdim VIS 15 /u01/app/oracle/8.0.5 /u01/app/oracle 60 60
```

NOTE

After you make any changes to the oramod.cf file, you must restart the SystemEDGE agent to ensure that the changes take effect.

Editing the SystemEDGE Startup File

The SystemEDGE startup file is the shell script for starting and stopping the SystemEDGE agent. This file was installed with the SystemEDGE agent. You must edit the file for your system to match your Oracle installation environment.

For Solaris, edit the following file:

```
/etc/rc2.d/S99sysedge
```

For HP-UX, edit the following file:

```
/sbin/rc2.d/S990sysedge
```

For AIX, edit the following file:

```
/etc/rc.tcpip
```

You must set the following environment variables in the SystemEDGE startup file to export the environment variables that Oracle requires:

- ORACLE_BASE=/oracle; export ORACLE_BASE
- ORACLE_SID=NHTD; export ORACLE_SID
- ORACLE_HOME=/oracle/product/8.1.5; export ORACLE_HOME
- LD_LIBRARY_PATH=$ORACLE_HOME/lib; export LD_LIBRARY_PATH

For more information, refer to the following sections.

NOTE

The *e*Health AIM for Oracle installation includes a sample SystemEDGE startup file. Do *not* install this sample file in place of your existing file. Use the sample file only as an example to help you edit the existing startup file to export the Oracle environment variables.

19166

Setting Environment Variables for Solaris Systems

Before you start SystemEDGE, add the following to your site's /etc/rc2.d/S99sysedge to customize directory locations for your Oracle installation for Solaris systems:

```
ORACLE_BASE=/export/ahab3/oracle/oracle.8.1.7
export ORACLE_BASE
ORACLE_SID=EH55
export ORACLE_SID
ORACLE_HOME=$ORACLE_BASE
export ORACLE_HOME
LD_LIBRARY_PATH=$ORACLE_HOME/lib:$ORACLE_HOME/jdbc/lib:/usr/dt/lib
export LD_LIBRARY_PATH
```

Setting Environment Variables for AIX Systems

Before you start SystemEDGE, add the following to your site's /etc/rc.tcpip file to customize directory locations for your Oracle installation for AIX systems:

```
ORACLE_BASE=/oracle/u01/app/oracle
export ORACLE_BASE
ORACLE_SID=POS1
export ORACLE_SID
ORACLE_HOME=$ORACLE_BASE/8i
export ORACLE_HOME
LD_LIBRARY_PATH=$ORACLE_HOME/lib
export LD_LIBRARY_PATH
```

Setting Environment Variables for HP-UX Systems

Before you start SystemEDGE, add the following to your site's /etc/rc2.d/S990sysedge file to customize directory locations for your Oracle installation:

```
ORACLE_BASE=/export/ahab3/oracle/oracle.8.1.7
```

19165

```
export ORACLE_BASE
ORACLE_SID=EH55
export ORACLE_SID
ORACLE_HOME=$ORACLE_BASE
export ORACLE_HOME
LD_LIBRARY_PATH=$ORACLE_HOME/lib:$ORACLE_HOME/jdbc/lib:/usr/dt/lib
export LD_LIBRARY_PATH
SHLIB_PATH=$ORACLE_HOME/lib:$ORACLE_HOME/jdbc/lib:/usr/dt/lib
export SHLIB_PATH
LD_PRELOAD=/usr/lib/libpthread.sl:/usr/lib/libcl.sl:/export/ahab3/
oracle/oracle.8.1.7/JRE/lib/PA_RISC/native_threads/libjava.sl
export LD_PRELOAD
```

Verifying that the Startup File is Set Up Correctly

If your SystemEDGE startup file is not set up correctly, you will receive an error message similar to the following:

```
brahma# ../../bin/sysedge
SystemEDGE Version 4.0 Patchlevel 3
Copyright 2001 by Concord Communications, Inc.
load_plugin: dlopen of /opt/EMPsysedge/plugins/oramod/oramod8.so failed, 9,
ld.so.1: ../../bin/sysedge: fatal: libclntsh.so.1.0: open failed: No such file
or directory
sysedge: load plugin /opt/EMPsysedge/plugins/oramod/oramod8.so failed
sysedge: using port 161, config file /etc/sysedge.cf
```

If you receive an error message, verify that you have set up the SystemEDGE startup file correctly and that you have loaded the correct shared library file for your version of Oracle.

Licensing *e*Health AIM for Oracle

Like the SystemEDGE agent, *e*Health AIM for Oracle utilizes a host-based license method. Copies of *e*Health AIM for Oracle can run only on systems that possess a valid license key. This license key is separate from the one used for the SystemEDGE agent.

19164
a.

The first time that you attempt to start the SystemEDGE agent after installing *e*Health AIM for Oracle, the agent displays a message stating that it could not find a valid license for *e*Health AIM for Oracle. It then provides you with a **public key** that is used to generate a permanent license key for your host machine.

A license key is composed of four space-separated, 8-character sequences, totaling 32 characters. The sysedge.lic file contains the *e*Health AIM for Oracle license, as well as the SystemEDGE agent license and other *e*Health AIM licenses. For an example, refer to the sample license file in "Sample License File" on page 33.

Obtaining a License

To obtain a license, you can do any of the following:

- Run the Concord-supplied licenseutil.pl script.
- Run the licenseme.exe license utility.
- Use AdvantEDGE View to receive an SNMP license trap or to query and license the plug-in without a trap. For more information, refer to "Generating a License through AdvantEDGE View Event Processing" on page 30 or "Generating a License through AdvantEDGE View Host Administration" on page 32.
- Send an e-mail request to licenses@concord.com and place the returned license key in the appropriate license file.
- Complete the online license form through the Internet, as described in the next section, "Generating the License through the Web-based License Form."

For more information, refer to the *eHealth SystemEDGE User Guide* and the *Automating the Licensing of SystemEDGE and AdvantEDGE Point Plug-in Modules* white paper.

Generating the License through the Web-based License Form

This section describes how to generate the license using the Web-based license form.

NOTE

If you are using an evaluation copy of *e*Health AIM for Oracle, you must request a temporary license that will enable it to operate during the evaluation period.

To generate a license for *e*Health AIM for Oracle:

1. Start the SystemEDGE agent as follows:
   a. Log in as **`root`**.
   b. Change directory (cd) to /opt/EMPsysedge.
   c. Enter the following:

   ```
   ./bin/sysedge -b
   ```

   The SystemEDGE agent displays a message indicating that you need a license for the *e*Health AIM for Oracle module on this host machine. It displays a message similar to the following:

```
SystemEDGE Version 4.1 Patchlevel 1
Copyright 2001 by Concord Communications, Inc.
Please contact Concord Communications, Inc. to obtain a license
http://www.concord.com/support, Email: licenses@concord.com
Provide this: sysedge neptune sol2 5.8 346561363366b19c 1.3 Patchlevel 1
```

2. Using a Web browser, go to the licensing Web site at http://license.concord.com, and select the **Create License** option that matches your use of the agent:
   - **Create SystemEDGE/AdvantEDGE Eval License** (if you are evaluating the AIM or are a Concord partner or reseller)
   - **Create SystemEDGE Outsource License** (if you are outsourcing the AIM)

- **Create SystemEDGE/AdvantEDGE License** (if you have purchased the AIM)

NOTE

> You must supply a user name and password to access the license form.

If you do not have Web access, fill out the license request form, /config/license.txt (available as part of the *e*Health AIM for Oracle installation), with the complete string generated by the SystemEDGE agent, and e-mail the completed form to licenses@concord.com.

3. Fill out the license form, entering the information that was printed by the SystemEDGE agent. You must supply the following information:

   - Name
   - E-mail address
   - Software version number (1.3 in the example on page 28)
   - Patchlevel (1 in the example on page 28)
   - System name (neptune in the example on page 28)
   - Operating system name (sol2 in the example on page 28)
   - Operating system version (5.8 in the example on page 28)
   - System identifier (346561363366b19c in the example on page 28)

   NOTE

   > Select the option for *e*Health AIM for Oracle from the product list on the licensing form.

   After you submit the license request form, the Concord Web server generates a license, displays it on your Web browser, and e-mails it to the contact person in your organization.

4. Copy the license into /etc/sysedge.lic, and save that file.
5. Restart the SystemEDGE agent as root by entering the following:

```
./bin/sysedge -b
```

Generating a License through AdvantEDGE View Event Processing

In order to use AdvantEDGE View event processing to license *e*Health AIM for Oracle, your system must meet the following requirements:

- You must be using SystemEDGE Release 4.0 Patchlevel 3 or later with AdvantEDGE View.
- You must configure the SystemEDGE agent to send SNMP traps to AdvantEDGE View. For more information, refer to the section on configuring the SystemEDGE agent in the *eHealth SystemEDGE User Guide.*
- You must configure the SystemEDGE agent with a read-write community so that AdvantEDGE View can issue an SNMP Set to transmit the license key to it. For more information, refer to the section on configuring the SystemEDGE agent in the *eHealth SystemEDGE User Guide.*
- Your AdvantEDGE View system must have access to the Internet, either directly or through a Web proxy.
- The AdvantEDGE View user who is generating the license must have either write or admin permissions.

To generate a license through AdvantEDGE View:

1. Start the SystemEDGE agent with *e*Health AIM for Oracle in unlicensed mode. SystemEDGE sends a license trap to AdvantEDGE View for that module.
2. Start AdvantEDGE View, and click the **Events** icon to display the Event Processing screen.

   AdvantEDGE View displays a license trap for the system that requires a license.

19160

3. Click the index number for that system to view the Trap Details form for **License Software** to display the AdvantEDGE View Software Licensing form.
4. Complete the licensing form, and click **Get License**.

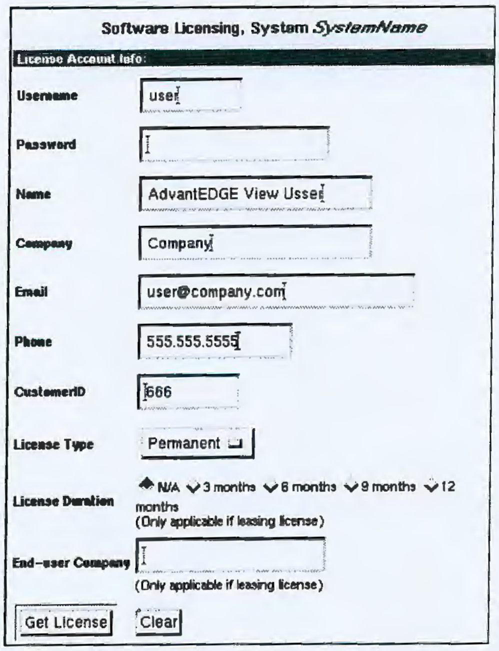

Software Licensing, System *SystemName*

License Account Info:

Username	user
Password	
Name	AdvantEDGE View Usser
Company	Company
Email	user@company.com
Phone	555.555.5555
CustomerID	666
License Type	Permanent
License Duration	● N/A ○ 3 months ○ 6 months ○ 9 months ○ 12 months (Only applicable if leasing license)
End-user Company	 (Only applicable if leasing license)

NOTE

If you have configured AdvantEDGE View preferences, AdvantEDGE View fills in all of the information (except password) on this form.

AdvantEDGE View contacts the Web-based license server, obtains a license for *e*Health AIM for Oracle, and issues an SNMP Set to the target SystemEDGE agent to inform it of the new software license key.

19159
E

Generating a License through AdvantEDGE View Host Administration

You can also license systems through AdvantEDGE View Host Administration.

To access Host Administration:

1. Start AdvantEDGE View, and click the **Administration** icon. AdvantEDGE View displays the Administration page.
2. Click the **Host Administration** icon. AdvantEDGE View displays the host list.

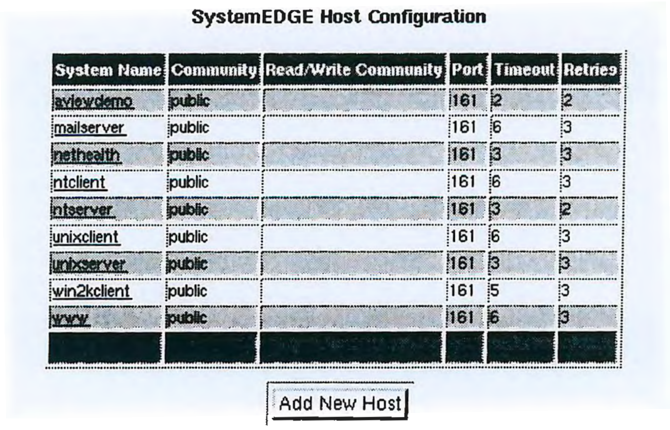

SystemEDGE Host Configuration

System Name	Community	Read/Write Community	Port	Timeout	Retries
aviewdemo	public		161	2	2
mailserver	public		161	6	3
nethealth	public		161	3	3
ntclient	public		161	6	3
ntserver	public		161	3	2
unixclient	public		161	6	3
unixserver	public		161	3	3
win2kclient	public		161	5	3
www	public		161	6	3

19158
ε

3. Click the name of the system that you want to license from the **System Name** column. AdvantEDGE View displays the Modify Host form.

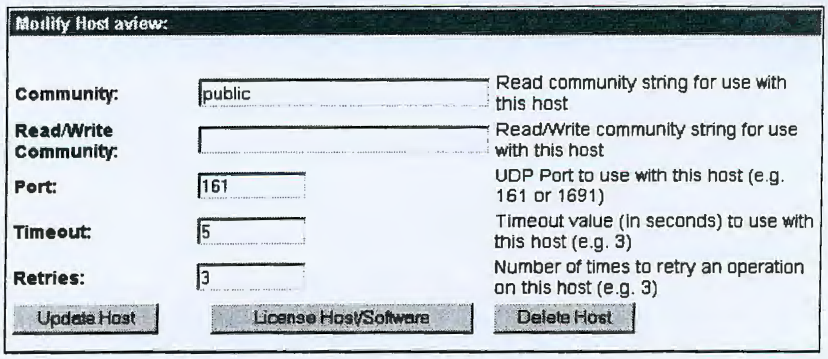

Modify Host aview:

Community:	public	Read community string for use with this host
Read/Write Community:		Read/Write community string for use with this host
Port:	161	UDP Port to use with this host (e.g. 161 or 1691)
Timeout:	5	Timeout value (in seconds) to use with this host (e.g. 3)
Retries:	3	Number of times to retry an operation on this host (e.g. 3)

4. Click **License Host/Software** to display the licensing form.

5. Select the option for *e*Health AIM for Oracle from the **Product** list, and then click **License Software**.

   AdvantEDGE View contacts the Web-based license server, obtains a license for the software, and issues an SNMP Set to the target SystemEDGE agent, informing it of the new software license key.

Sample License File

The following is a sample SystemEDGE agent license file. A pound character (#) in column 1 indicates that the entire line is a comment.

```
# license file for SystemEDGE Agent
# Concord Communications, Inc.
# http://www.concord.com
# file /etc/sysedge.lic or %SystemRoot%\system32\sysedge.lic
# A valid license key has four parts of 8 characters per part
# parts are separated by space(s) with one license key per line

# sysedge jupiter sol2 5.8 807cb1da007cb1da 4.1 PL 1
e13311d3 0F2a7cb1 abC512dc fF8C923a

# oramod jupiter sol2 5.8 807cb1da007cb1da 1.3 PL 1
a7943fde 098a87ij a4kiuf39 afafEkj4
```

3

Using the eHealth AIM for Oracle MIB

This chapter explains the organization and content of the Concord Communications MIB for the Oracle application. The MIB specification (oramod.asn1) defines a collection of objects for monitoring and managing Oracle. You must configure the SystemEDGE agent to monitor the *e*Health AIM for Oracle MIB objects that are relevant for your configuration. For more information, refer to Chapter 4, "Using eHealth AIM for Oracle." Figure 3 shows part of the *e*Health AIM for Oracle MIB.

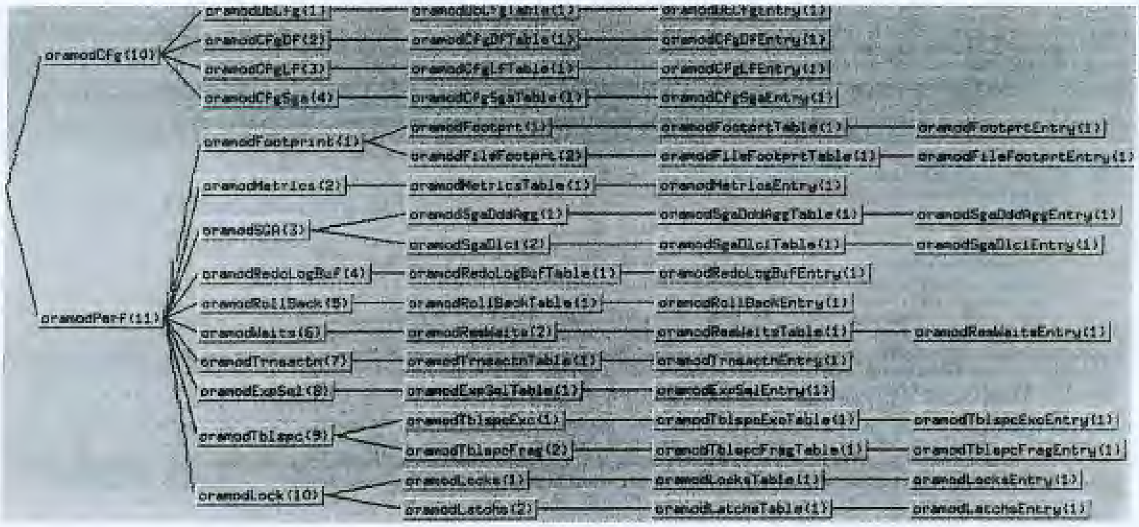

Figure 3. *e*Health AIM for Oracle MIB

19156
E

The *e*Health AIM for Oracle MIB is organized into sections for database configuration and performance. Within the performance section, a footprint section defines MIB objects that show how much of the underlying system's resources Oracle is consuming. The following sections describe the *e*Health AIM for Oracle MIB. For the most current list of MIB objects, refer to the *e*Health AIM for Oracle MIB specification (oramod.asn1).

Configuration Section

The Configuration section of the *e*Health AIM for Oracle MIB contains configuration parameters and settings that are important for monitoring the health and performance of the Oracle database. The Configuration section includes groups for database, database files, redo logs, and the System Global Area (SGA).

Database Configuration Group

The Database Configuration group contains instance names, index numbers, and initial startup information. Table 3 defines the Database Configuration parameters.

Table 3. MIB Objects – Oracle Database Configuration Group (Page 1 of 3)

MIB Object	Description
oramodDbCfgSIDINDX	Specifies the SID index number.
oramodDbCfgSID	Specifies the database name.
oramodDbCfgVERSION	Provides the description and version of this Oracle installation.
oramodDbCfgHOME	Specifies the installation directory for Oracle.
oramodDbCfgBASE	Specifies the base directory for Oracle, or ORACLEBASE.

19155
E

Table 3. MIB Objects – Oracle Database Configuration Group (Page 2 of 3)

MIB Object	Description
oramodDbCfgID	Specifies the database ID that was calculated when the database was created; this value is stored in all file headers.
oramodDbCfgCRTDT	Specifies the original creation date and time of the database.
oramodDbCfgLOGMODE	Specifies the archive log mode: NOARCHIVELOG or ARCHIVELOG.
oramodDbCfgCTRLFILETYPE	Specifies the control file type: CURRENT, STANDBY, CLONE, BACKUP, or CREATED.
oramodDbCfgOPENMODE	Indicates whether the database is set to READ WRITE or READ ONLY mode.
oramodDbCfgMAXPROCESS	Specifies the maximum number of user processes that can simultaneously connect to an Oracle server for a multi-process operation.
oramodDbCfgMAXSESSION	Specifies the maximum number of user and system sessions.
ormodDbCfgTIMEDSTATISTICS	Specifies whether timing is on. Set this variable to FALSE for optimum performance.
oramodDbCfgCPUCNT	Specifies the maximum number of CPUs for this instance.
oramodDbCfgSHAREDPOOLSIZE	Specifies the shared pool size in KB. This variable controls the size of the memory area that is used for parsing and executing SQL statements.
oramodDbCfgSHAREDPOOLRSSSIZE	Specifies the shared pool reserved size in KB.
oramodDbCfgLARGEPOOLSIZE	Specifies the large pool reserved size in KB.

19154

Table 3. MIB Objects – Oracle Database Configuration Group (Page 3 of 3)

MIB Object	Description
oramodDbCfgCPUJAVAPOOLSIZE	Specifies the Java pool size in KB. The default value is 20,000 KB.
oramodDbCfgCPUCNTRLFILES	Specifies the name and path of the control files.
oramodDbCfgCPUDBBLKBUFF	Specifies the number of database block buffers in blocks.
oramodDbCfgBLKSIZE	Specifies the size of the database blocks.
oramodDbCfgCKPTINTRVL	Specifies the redo log checkpoint interval.
oramodDBCfgDBFILES	Specifies the maximum number of database files.
oramodDbCfgSORTAREASIZE	Specifies the maximum sort area size. This object controls the amount of memory allocated to each process in the Process Global Area for any sorting activity.
oramodDBCfgOPENCURSORS	Specifies the maximum number of simultaneous open cursors that a single-user process can have.
oramodDBCfgTRNSACTNS	Specifies the maximum number of simultaneous, concurrent transactions.
oramodDBCfgTRNSACTNSPERSEG	Specifies the maximum number of concurrent transactions per segment.
oramodDBCfgMAXROLLSEG	Specifies the maximum number of rollback segments.

Database File Configuration Group

The Database File Configuration group describes the Oracle database files. Table 4 defines the Database File Configuration parameters.

Table 4. MIB Objects – Oracle Database File Configuration Group (Page 1 of 2)

MIB Object	Description
oramodCfgDfFILENUM	Specifies the file identification number.
oramodCfgDfSTATUS	Specifies the type of file (system or user) and its status. The value can be OFFLINE, ONLINE, SYSTEM, RECOVER, or SYSOFF.
oramodCfgDfENABLED	Describes how accessible a file is from SQL. The value can be DISABLED, READ ONLY, READ WRITE, or UNKNOWN.
oramodCfgDfUNRCVRBLECHG	Specifies the last unrecoverable change number that was made to this data file.
oramodCfgDfUNRCVRBLETIME	Specifies the last unrecoverable change time that was made to this data file.
oramodCfgDfKBYTES	Specifies the current size of the file in kilobytes (KB).
oramodCfgDfCRTKBYTES	Specifies the size of the data file when it was created.
oramodCfgDfFNAME	Specifies the file name.
oramodCfgDfCRTTIME	Specifies the time at which the data file was created.
oramodCfgDfTBLSPACENUM	Specifies the tablespace number.
oramodCfgDfTBLESPACERFILENUM	Specifies the tablespace relative data file number.
oramodCfgDfBLOCKS	Specifies the current size of the data file in blocks.

19152
ε

Table 4. MIB Objects – Oracle Database File Configuration Group (Page 2 of 2)

MIB Object	Description
oramodCfgDfBLOCKSIZE	Specifies the block size of the data file.
oramodCfgDfERROR	Indicates whether the datafile header read an error. If so, a restore is required. A value of NULL indicates success.
oramodCfgDfRECOVER	Indicates whether the file needs media recovery. The value can be YES or NO.
oramodCfgDfRSTLOGSCHGNUM	Specifies the reset log change number.
oramodCfgDfRSTLOGSTIME	Specifies the reset log timestamp.

Redo Log File Configuration Group

The Redo Log File Configuration group describes the redo log file locations and status for each Oracle installation. Table 5 defines the Redo Log File Configuration parameters.

Table 5. MIB Objects – Oracle Redo Log File Configuration Group

MIB Object	Description
oramodCfgLfMEMBERINDX	Specifies the log file member index.
oramodCfgLfGROUPNUM	Specifies the log file group number.
oramodCfgLfSTATUS	Specifies the log file status.
oramodCfgLfMEMBER	Specifies the log file name and path.

1915
3

System Global Area Configuration Group

The System Global Area (SGA) Configuration group summarizes the Oracle system global area. Table 6 defines the SGA Configuration parameters.

Table 6. MIB Objects – Oracle SGA Configuration Group

MIB Object	Description
oramodCfgSgaTOTALMEMALLOC	Specifies the total real memory allocated.
oramodCfgSgaFIXEDSGA	Specifies the fixed memory allocated in KB.
oramodCfgSgaVARIABLE	Specifies the variable memory that is allocated in the SGA.
oramodCfgSgaDBBUFF	Specifies the database buffers allocated in KB. The typical range is from 4 KB to 2 GB.
oramodCfgSgaREDOBUFF	Specifies the redo buffers allocated in KB.

Performance Section

The Performance section of the *e*Health AIM for Oracle MIB contains performance data that is necessary for capacity planning and trend analysis, as well as real-time performance and availability monitoring. The Performance group is divided into subgroups for footprint data, metrics, the SGA, redo logs, rollback segments, waits, expensive SQL, tablespaces, and locks.

Oracle Footprint Group

The Footprint group provides information about the Oracle CPU, memory, data flow, and disk-resource consumption, more commonly called the **footprint**. Long-term trending analysis of footprint information is useful for anticipating and avoiding problems due to resource exhaustion. You can also monitor footprint information in real time to detect and correct temporary resource exhaustion due to viruses, security incidents, and hardware failures. The Footprint group includes subgroups for footprint and file footprint.

10150

Footprint Group

The Footprint group describes the performance footprint for each Oracle service. Table 7 defines the Oracle Footprint parameters.

Table 7. MIB Objects – Oracle Footprint Group (Page 1 of 2)

MIB Object	Description
oramodFootprtCPUTIME	Specifies the CPU time in seconds accumulated by Oracle. This time includes all core components and the database instance.
oramodFootprtPERCENTCPU	Specifies the percentage of CPU utilization by Oracle over the last sample interval. This value is reported as the percentage multiplied by 100. It is the sum of the CPU utilization of all Oracle core services and components, and provides a good indication of the CPU resource consumption by the entire Oracle application. Do **not** use this variable as a counter.
oramodFootprtMEMSIZE	Specifies the combined size of all Oracle text, data, and stack segments in KB. This value is the sum of all process sizes for all core Oracle services and optional connectors.
oramodFootprtRSS	Specifies the total real memory (resident size set [RSS]) of Oracle in KB. This value is the sum of process RSS for all core Oracle services and optional connectors.
oramodFootprtPERCENTMEM	Specifies the percentage (0-100) of real memory used by Oracle, including core services and optional connectors.
oramodFootprtTHREADS	Specifies the total number of threads operating within the respective Oracle components as seen by the operating system.
oramodFootprtINBLKS	Specifies the number of blocks of data input by the processes.
oramodFootprtOUTBLKS	Specifies the number of blocks of data output by the processes.
oramodFootprtMSGSSENT	Specifies the number of messages sent by the processes.

19149

Table 7. MIB Objects – Oracle Footprint Group (Page 2 of 2)

MIB Object	Description
oramodFootprtMSGSRCVD	Specifies the number of messages received by the processes.
oramodFootprtSYSCALLS	Specifies the number of system calls invoked by the processes.
oramodFootprtMINORPGFLTS	Specifies the number of minor page faults incurred by the process.
oramodFootprtMAJORPGFLTS	Specifies the number of major page faults incurred by the process.
oramodFootprtNUMSWAPS	Specifies the number of times the processes have been swapped.
oramodFootprtVOLCNTX	Specifies the number of voluntary context switches incurred by the processes.
oramodFootprtINVOLCNTX	Specifies the number of involuntary context switches incurred by the processes.
oramodFootprtHOMESIZE	Specifies the size in KB of the Oracle root directory and all files beneath it.
oramodFootprtDBDISKSIZE	Specifies the size in KB of all Oracle database disk storage areas. This value is a sum of all database files.
oramodFootprtSTARTUPTIME	Specifies the date and time that the Oracle instance was started.
oramodFootprtSTATUS	Specifies the status of the instance: STARTED, MOUNTED, or OPEN.
oramodFootprtDATABASESTATUS	Specifies the database status of the instance.
oramodFootprtSHUTDOWNPENDING	Specifies whether the database is about to be shut down: YES or NO.

19148

File Footprint Group

The File Footprint group describes the Oracle file footprint. Table 7 defines the Oracle File Footprint parameters.

Table 8. MIB Objects – Oracle File Footprint Group

MIB Object	Description
oramodFFootprtFILEINDX	Specifies the file index number.
oramodFFootprtFILETYPE	Specifies the file type.
oramodFFootprtFILENAME	Specifies the file name.
oramodFFootprtCRTTS	Specifies the timestamp of the last update.
oramodFFootprtCRTKBYTES	Specifies the file creation size in KB.
oramodFFootprtKBYTES	Specifies the file size in KB.
oramodFFootprtBLOCKS	Specifies the file blocks.
oramodFFootprtSTATUS	Specifies the file status.
oramodFFootprtRECOVER	Indicates whether the file needs recovering.
oramodFFootprtAVGIOTIM	Specifies the average input/output time for the file.

19147
E

Metrics Group

The Metrics Group describes the Oracle database metrics. Table 9 defines the Oracle Metrics parameters.

Table 9. MIB Objects – Oracle Metrics Group (Page 1 of 4)

MIB Object	Description
oramodMetricsBCPT	Block Changes Per Transaction (scaled by 100). Measures the amount of data manipulation language (DML) work that each transaction performs.
oramodMetricsBGR	Block Get Rate per second (scaled by 100). Measures the rate at which application systems reference the database.
oramodMetricsBVPT	Block Visits Per Transaction (scaled by 100). Measures the database load imposed for each transaction.
oramodMetricsCHR	Cache Hit Ratio (scaled by 100). Measures the effectiveness of the hits against the buffer cache.
oramodMetricsCR	Call Rate (user and recursive) over time. Measures the work demand rate that all work sources are placing on the database instance.
oramodMetricsCPT	Calls Per Transaction (scaled by 100). Indicates the number of database requests per committed transaction.
oramodMetricsCBR	Changed Blocks (scaled by 100). Measures the balance between queries and DML and can indicate changes in application use or indexes.
oramodMetricsCCR	Consistent Change Ratio (scaled by 100). Measures consistent database changes and gets, or the extent to which applications are exercising the read-consistency mechanism.
oramodMetricsCRR	Continued Row Ratio (scaled by 100). Measures the space usage of data blocks.
oramodMetricsLCM	Library Cache Miss Ratio (scaled by 100). Measures the missed requests and reloads to access library-cached objects since instance startup. Monitor the library cache to help determine whether you should resize the shared pool.

Table 9. MIB Objects – Oracle Metrics Group (Page 2 of 4)

MIB Object	Description
oramodMetricsRTUC	Recursive To User Call Ratio (scaled by 100). Measures the ratio of recursive and user to application calls.
oramodMetricsRLSW	Redo Log Space Wait Ratio (scaled by 100). Shows the rate at which failures are occurring in writing to the redo log space. It can measure memory allocation.
oramodMetricsRSR	Row Source Ratio (scaled by 100). Specifies the percentage of total rows retrieved from full table scans.
oramodMetricsSOR	Sort Overflow Ratio (scaled by 100). Indicates the number of sorts that are using temporary segments.
oramodMetricsTRR	Transaction Rate Ratio (scaled by 100). Indicates the rate at which users or applications are committing transactions.
oramodMetricsUCR	User Call Rate Ratio (scaled by 100). Indicates the rate at which client users or applications are requesting SQL statements.
oramodMetricsUCPP	User Calls Per Parse Ratio (scaled by 100). Specifies the ratio of user calls to parses in the shared SQL area. This value indicates how well the application is managing its context area.
oramodMetricsURR	User Rollback Ratio (scaled by 100). Indicates the rate at which application transactions are failing.
oramodMetricsSGALCE	Library Cache Efficiency Percentage (scaled by 100). Specifies the percentage of times that an SQL statement did not need to be reloaded because it was already in the library cache.
oramodMetricsSGADDCE	Data Dictionary Cache Efficiency Percentage (scaled by 100). Shows the ratio of gets to misses within the shared pool.
oramodMetricsDBTOTALRW	Specifies the total number of physical reads and writes since the instance started.
oramodMetricsDBBLKCHG	Specifies the total number of database block changes.

Table 9. MIB Objects – Oracle Metrics Group (Page 3 of 4)

MIB Object	Description
oramodMetricsDBBLKGET	Specifies the total number of database block gets since the instance started.
oramodMetricsDBCNSTGET	Specifies the total number of consistent gets since the instance started.
oramodMetricsDBPHYSREAD	Specifies the total number of physical reads since the instance started.
oramodMetricsDBSORTDISK	Specifies the total number of database disk sorts since the instance started.
oramodMetricsDBSORTMEM	Specifies the total number of database memory sorts since the instance started.
oramodMetricsBLKFREEWAIT	Specifies the percentage of times that a request resulted in a wait for a free block.
oramodMetricsMTHRDQUEUEWAIT	Specifies the number of seconds that a request waited in the queue for multi-threaded servers.
oramodMetricsSESSHIWTRMEM	Specifies the maximum high-water amount of session memory that has been used.
oramodMetricsSESSCURRMEM	Specifies the current maximum amount of session memory that has been used.
oramodMetricsSESSHIWTR	Specifies the high-water mark for the number of sessions.
oramodMetricsSESSCURRENT	Specifies the current number of sessions.
oramodMetricsUSERCOMMITS	Specifies the number of successfully completed and aborted database transactions.
oramodMetricsUSERROLLBACK	Specifies the number of transactions that have been rolled back since the instance started.
oramodMetricsUSERCALLS	Specifies the number of transactions that have been started since the instance started.
oramodMetricsDBPHYSWRTS	Specifies the total number of physical writes since the instance started.

19144
ε

Table 9. MIB Objects – Oracle Metrics Group (Page 4 of 4)

MIB Object	Description
oramodMetricsTBLSCANROWS	Specifies the number of rows that were retrieved through table scans since the instance started.
oramodMetricsTBLFTCHROWID	Specifies the number of rows that were retrieved through table fetches by row ID since the instance started.
oramodMetricsTBLFTCHCROW	Specifies the number of rows that were retrieved through table fetches by row since the instance started.
oramodMetricsRECRSVCALLS	Specifies the number of recursive calls since the instance started.
oramodMetricsCNSTCHGS	Specifies the number of consistence changes since the instance started.
oramodMetricsPARSECNT	Specifies the number of parses since the instance started.
oramodMetricsCPUTM	Specifies the amount of time that was taken to work requests.
oramodMetricsLOGFILESWTCH	Specifies the number of redo log file switches.

SGA Group

The SGA area of the Oracle MIB stores aggregate information about the performance of the dictionary cache. Each user and data request that connects to the database references this data dictionary. Retaining as much of the data dictionary information as possible in memory contributes to database efficiency. The SGA Group includes subgroups for Data Dictionary Aggregate (DDA) and Detailed Library Cache Information (DLCI).

1343
E.

SGA DDA Group

The SGA DDA group describes the data dictionary aggregates. Table 10 defines the Oracle SGA DDA parameters.

Table 10. MIB Objects – Oracle SGA DDA Group

MIB Object	Description
oramodSGAAGGCNT	Specifies the total number of entries in the data dictionary cache since the instance started.
oramodSGAAGGUSGE	Specifies the total number of cache entries with valid data since the instance started.
oramodSGAAGGFIX	Specifies the total number of fixed entries in the data dictionary cache since the instance started.
oramodSGAAGGGET	Specifies the total number of information requests since the instance started.
oramodSGAAGGGETMISS	Specifies the total number of information requests that resulted in cache misses since the instance started.
oramodSGAAGGSCAN	Specifies the total number of scan requests since the instance started.
oramodSGAAGGSCANMISS	Specifies the total number of scan requests that resulted in misses since the instance started.
oramodSGAAGGSCANCPLT	Specifies the total number of times that the list was scanned completely since the instance started.
oramodSGAAGGMODS	Specifies the total number of insertions, updates, and deletions since the instance started.
oramodSGAAGGFLUSH	Specifies the total number of times the data was flushed to disk since the instance started.

SGA DLCI Group

The SGA DLCI group describes the detailed library cache information. Table 11 defines the Oracle SGA DLCI parameters.

Table 11. MIB Objects – Oracle SGA DLCI Group

MIB Object	Description
oramodSGADLCINAME	Specifies the Library Cache Namespace: SQL AREA, TABLE/PROCEDURE, BODY, TRIGGER, INDEX, CLUSTER, OBJECT, or PIPE.
oramodSGADLCIGET	Specifies the total number of request handles to the library namespace since instance startup.
oramodSGADLCIGETHIT	Specifies the total number of times that handles were already allocated in the cache.
oramodSGADLCIGETHITRT	Specifies the Get/Hit ratio that indicates the percentage of parse calls that find a cursor to share.
oramodSGADLCIPIN	Specifies the number of pin requests for objects in the cache since the instance started.
oramodSGADLCIPINHIT	Specifies the number of pin hits or executions of objects that are already allocated and initiated in the cache.
oramodSGADLCIPINHITRT	Specifies the ratio of pin hits to pins.
oramodSGADLCIRELOAD	Specifies the number of times a statement had to be reparsed in order to be executed.
oramodSGADLCIINVALID	Specifies the number of times that non-persistent library objects have been invalidated.

Redo Logs Group

Oracle stores data that is being changed by transactions in redo logs. The application can later use these redo logs to perform database recoveries, if necessary. Table 12 defines the Oracle Redo Log parameters.

Table 12. MIB Objects – Oracle Redo Log Group

MIB Object	Description
oramodREDOBLKWRT	Specifies the total number of redo blocks written since instance startup.
oramodREDOENTRIES	Specifies the total number of redo entries in the redo log since instance startup.
oramodREDOSIZE	Specifies the size of the redo log.
oramodREDOSPACERQST	Specifies the number of redo log space requests. If this value is not close to zero, the buffer is too small.
oramodREDOSPACEWAIT	Specifies the total number of redo log space wait times since the instance started.
oramodREDOSYNCHWRT	Specifies the total number of redo synch writes since the instance started.
oramodREDOSSYNCHTM	Specifies the total amount of redo synch time.
oramodREDOWASTAGE	Specifies the total amount of redo log waste.
oramodREDORETRIES	Specifies the total number of redo retries.

Rollback Segment Group

The Rollback Segment section of the Oracle MIB records cumulative statistics about the use of all rollback segments since the database was last started. These segments contain records of data that have not yet been modified in a transaction. The database can use these segments to remove changes that were performed for a transaction that was canceled. That is, they can return the record to the state it was in before any changes were

19140

made. For best performance, try to size the rollback segment so that each of its extents is large enough to support a typical transaction. Table 13 defines the Oracle Rollback Segment parameters.

Table 13. MIB Objects – Oracle Rollback Segment Group (Page 1 of 2)

MIB Object	**Description**
oramodRollBackSEGNAME	Specifies the rollback segment name.
oramodRollBackEXTENTS	Specifies the total number of rollback segment extends since the instance started.
oramodRollBackRSSIZE	Specifies the size of the rollback segment.
oramodRollBackWRITES	Specifies the total number of segment Writes since the instance started.
oramodRollBackGETS	Specifies the total number of rollback segment Gets since the instance started.
oramodRollBackWAITS	Specifies the total number of rollback segment Waits since the instance started.
oramodRollBackOPTSIZE	Specifies the optimal size of the rollback segment. Base this value on application use to minimize space management issues.
oramodRollBackSHRINKS	Specifies the total number of rollback segment shrinks (return to normal size after an extend) since the instance started.
oramodRollBackWRAPS	Specifies the total number of rollback segment wraps since instance startup. If this value is not equal to zero, recreate the rollback segments with larger extent sizes.
oramodRollBackEXTENDS	Specifies the number of rollback segments extends (extensions for handling larger transactions that what the segment was designed to handle).

19139

Table 13. MIB Objects – Oracle Rollback Segment Group (Page 2 of 2)

MIB Object	Description
oramodRollBackAVESHRINK	Specifies the average shrink size of this rollback segment since the instance started.
oramodRollBackAVEACTIVE	Specifies the average activity for this rollback segment since the instance started.
oramodRollBackSTATUS	Specifies the rollback segment status.

Waits Group

Wait time contributes most significantly to total response time. The Waits table identifies the events that are contributing to the longest wait times. Table 14 defines the Oracle Waits parameters.

Table 14. MIB Objects – Oracle Waits Group

MIB Object	Description
oramodWaitsEVENT	Specifies the name of the resource that is experiencing waits.
oramodWaitsTOTALWAITS	Specifies the total number of waits for the resource.
oramodWaitsTOTALTIMOUT	Specifies the total number of timeouts for the resource.
oramodWaitsTIMEWAITED	Specifies the total time the resource waited.
oramodWaitsAVGWAIT	Specifies the average time the resource waited.

19138
e

Expensive SQL Group

SQL queries that are not optimized can significantly reduce the overall performance of a system. The Expensive SQL table identifies statements that require too many disk reads or parses. Table 15 defines the Oracle Expensive SQL parameters.

Table 15. MIB Objects – Oracle Expensive SQL Group

MIB Object	Description
oramodExpSqlEXECUTION	Specifies the number of executions of expensive SQL.
oramodExpSqlDISKREAD	Specifies the number of disk reads.
oramodExpSqlPARSECNT	Specifies the number of parses. A large number of hard parses requires more memory for the shared pool.
oramodExpSqlBUFFGETS	Specifies the number of Buffer Gets.
oramodExpSqlSORTS	Specifies the number of Sorts.
oramodExpSqlSQL	Specifies the SQL statement.

19137
E

Tablespace Group

The Tablespace group identifies the fragmentation within each database tablespace. For best performance, you must have enough free block space to fill database requests. The Tablespace table describes the amount of space left in a tablespace after the next largest extent of any table in the tablespace is allocated. Table 16 defines the Oracle Tablespace parameters.

Table 16. MIB Objects – Oracle Tablespace Group

MIB Object	Description
oramodTblspTBLSPINDX	Specifies the tablespace index number.
oramodTblspTBLSPNAME	Specifies the tablespace name.
oramodTblspKBYTESLEFT	Specifies the number of KB that are left after the next largest extent is allocated.
oramodTblspFILNAME	Specifies the tablespace file name.
oramodTblspEXTENTS	Specifies the number of extents.
oramodTblspLRGEXTENT	Specifies the size in bytes of the largest extent.
oramodTblspSMEXTENT	Specifies the size in bytes of the smallest extent.
oramodTblspINCREMENTBY	Specifies the size in bytes of the next extent increment.
oramodTblspBYTESCOALSD	Specifies the number of free bytes that were coalesced.
oramodTblspBYTESFREE	Specifies the number of free bytes that are available.

Lock Group

Locks are access restrictions in a multi-user environment. They can maintain security and protect the integrity of the data by remaining in place until a commit (which saves all changes to a database since the start of a transaction) or rollback (which removes updates performed by a partially completed transaction) takes place. The Lock Group includes subgroups for locks and latches.

19136
t

Locks Group

The Locks table identifies lock holders and requestors. Table 17 defines the Oracle Locks parameters.

Table 17. MIB Objects – Oracle Locks Group

MIB Object	Description
oramodLockUSERINDX	Specifies the user's index number.
oramodLockUSERNAME	Specifies the name of the user who is holding the lock.
oramodLockOBJECT	Specifies the name of the locked object.
oramodLockTYPE	Specifies the object type of the locked object.
oramodLockMODE	Specifies the type of lock mode that is being held: Null, Row share, Row exclusive, Share, Share row exclusive, or Exclusive.
oramodLockCTIME	Specifies the amount of time since the current mode was granted.
oramodLockBLOCK	Indicates whether the lock is blocking another lock.

Latches Group

Latches are subsets of locks that can prevent access to Oracle internal memory structures while a process is accessing them. If the database spends too much time waiting for latches, it can affect performance. The Latch table identifies latches held per instance. Table 18 defines the Oracle Latch parameters.

Table 18. MIB Objects – Oracle Latch Group (Page 1 of 2)

MIB Object	Description
oramodLatchLATCHNUM	Specifies the latch number.
oramodLatchLATCHNAME	Specifies the latch name.
oramodLatchGETS	Specifies the number of times that Oracle satisfied requests for latches without waiting.

Table 18. MIB Objects – Oracle Latch Group (Page 2 of 2)

MIB Object	Description
oramodLatchMISSES	Specifies the number of times that Oracle did not satisfy latch requests.
oramodLatchSLEEPS	Specifies the number of latch sleeps.
oramodLatchIMDGETS	Specifies the number of latch immediate gets.
oramodLatchIMDMISSES	Specifies the number of latch immediate misses.
oramodLatchWAITSHOLDING	Specifies the number of latch waits that are holding.
oramodLatchSPINGETS	Specifies the number of latch spin gets.

Backup Group

The Backup group provides physical backup information. Table 19 defines the Backup parameters.

Table 19. MIB Objects – Oracle Backup Group

MIB Object	Description
oramodBackupFILENUM	Specifies the backup file index number.
oramodBackupSTATUS	Specifies the status of the backup.
oramodBackupCHANGENUM	Specifies the change number of the backup file.
oramodBackupDATE	Specifies the date of the backup.

19134
E

Archive Group

The Archive group provides physical backup information. Table 20 defines the Backup parameters.

Table 20. MIB Objects – Oracle Archive Group

MIB Object	Description
oramodArcDestARCMODE	Specifies the mode of the archive log.
oramodArcDestSTATUS	Specifies the status of the archive log file.
oramodArcDestDEST	Specifies the archive log file destination.

19133

4

Using eHealth AIM for Oracle

This chapter describes how to configure and use *e*Health AIM for Oracle. This module is implemented as a SystemEDGE agent plug-in. After you enable this plug-in in the sysedge.cf file and obtain a license for it, SystemEDGE will load it automatically at startup. For more information, refer to “Editing the sysedge.cf File” on page 21 and “Licensing eHealth AIM for Oracle” on page 26.

*e*Health AIM for Oracle implements additional MIB objects that provide advanced information about the health and availability of Oracle databases. It can operate with any SNMP-compliant management software such as Concord’s *e*Health suite of products, AdvantEDGE View, HP OpenView, and others. If you are using *e*Health AIM for Oracle with *e*Health, refer to the *e*Health Web Help for more information about the reports that you can generate.

The default configuration settings of the *e*Health AIM for Oracle plug-in enable you to use the advanced self-monitoring capabilities of the SystemEDGE agent in conjunction with *e*Health AIM for Oracle.

19132
ε

*e*Health AIM for Oracle MIB Branch

You can use AdvantEDGE View or another SNMP tool to edit the SystemEDGE configuration file to utilize the *e*Health AIM for Oracle MIB objects with the process-monitoring, threshold-monitoring, and history-collection features of the SystemEDGE agent. All MIB objects that are related to *e*Health AIM for Oracle exist at object identifier (OID) branch *1.3.6.1.4.1.546.16.4* in the Concord Systems Management MIB. The MIB is defined in the oramod.asn1 file, which is available in the *e*Health AIM for Oracle product installation.

Assigning Entry Rows in the SystemEDGE Self-Monitoring Tables

All SystemEDGE self-monitoring tables (including log monitoring, process/service monitoring, threshold monitoring, and history collection) require the use of unique row numbers. Each table contains an **Index** column which acts as a **key field** to distinguish rows in the table. This section describes the benefits of reserving a block of rows (in the range of 11 to the maximum number of rows in your table) for use by the system or application administrator.

Setting Local Policy

You may choose, as a matter of local policy, to reserve a block of rows for system administration. This policy allows you to define entries within a reserved block of rows without being concerned that the row might already be taken by another user's entry. In compliance with the local policy, all other users should use row indices that are outside the reserved range when they define user-configured entries.

Reserving Blocks of Rows

By reserving a block of rows, you can define a consistent set of conditions (row entries) to be monitored across all machines such that the same condition is defined in the same row number on each machine. For example, you can use row 3000

19131
ε

in each table to define entries monitoring the MIB object for total number of SQL disk reads (oramodExpSqlDISKREAD). You can then distribute this configuration to every host so that every system that is running Oracle uses row 3000 for monitoring the number of SQL disk reads, for each SystemEDGE self-monitoring table.

To reserve a block of rows for monitoring Oracle:

1. Decide which block of rows you want to reserve for use with monitoring Oracle.
2. Use that block of rows to define a set of row entries for each SystemEDGE self-monitoring table. For more information, refer to the chapter on self-monitoring in the *eHealth SystemEDGE User Guide*.
3. Distribute configuration file entries to all hosts that are running Oracle and *e*Health AIM for Oracle. For more information, refer to the *Automating the Deployment of SystemEDGE and the AdvantEDGE Point Plug-in Modules* white paper.

NOTE

As an alternative, you can use this row-number assignment policy with AdvantEDGE View for group-configuration operations.

4. Require end users to avoid your block of rows when defining their own self-monitoring table entries.

Using the SystemEDGE Self-Monitoring Features

This section provides examples of how to use SystemEDGE process, threshold, and history monitoring to monitor the Oracle application. Add these commands to the sysedge.cf file to enable monitoring of the MIB objects that they specify. Modify these examples as necessary to monitor the MIB objects that are relevant for your configuration.

fl.130
E

The examples in this section use row numbers in the 5000 range; use a row number for your configuration that conforms to local policies. For more information on row assignment, refer to "Assigning Entry Rows in the SystemEDGE Self-Monitoring Tables" on page 60.

The following command, for example, instructs the SystemEDGE agent to monitor the RSS of the Oracle database every 60 seconds for values that are greater than 50,000 and to store the data in row 5001 of the Threshold Monitoring table:

```
monitor oid oramodFootprtRSS 5001 0x0 60 absolute > 50000 'Total
Resident Memory' ''
```

For more information about the syntax for the commands in this section, refer to the *eHealth SystemEDGE User Guide*. For more information about *e*Health AIM for Oracle MIB objects, refer to Chapter 3, "Using the eHealth AIM for Oracle MIB," or to the MIB specification (oramod.asn1).

NOTE

Enter the commands throughout this chapter on one line. Do *not* use a carriage return to match the formatting shown here.

Using SystemEDGE Process Monitoring

This section provides an example of how to use the SystemEDGE agent to monitor the availability of a critical Oracle process. For more information, refer to the chapter on process monitoring in the *eHealth SystemEDGE User Guide*.

Monitoring the Oracle database

To ensure that the Oracle database is running, enter the following command in the sysedge.cf file:

```
watch process procAlive 'oracle' 5000 30 'Oracle Database' ''
```

19129
E

Using SystemEDGE Threshold Monitoring

This section provides examples of how to use the SystemEDGE agent to monitor important Oracle metrics. Add the commands that are provided in the following sections to the sysedge.cf file to monitor thresholds for these MIB objects. For more information, refer to the chapter on threshold monitoring in the *eHealth SystemEDGE User Guide.*

NOTE

The thresholds used in this section may not be appropriate for your Oracle application; select thresholds that are appropriate for your environment.

Monitoring Total Resident Memory Size of an Oracle Service

To monitor the total resident memory size of an Oracle service, enter the following:

```
monitor oid oramodFootprtRSS 5001 0x0 60 absolute > 50000 'Total
Resident Memory' ''
```

Monitoring Combined Size of Oracle Text, Data, and Stack Segments

To monitor the total size of the Oracle text, data, and stack segments, enter the following:

```
monitor oid oramodFootprtMEMSIZE 5002 0x0 60 absolute > 1000000 'Total
Memory Size' ''
```

Monitoring Total Size of Oracle Redo Log File

To monitor the total size of the Oracle redo log file, enter the following:

```
monitor oid oramodREDOSIZE 5003 0x0 60 absolute > 500000 'Total Redo
Log Size' ''
```

Fl128 ε

Monitoring Total Number of Database Disk Sorts

To monitor the total number of database disk sorts, enter the following:

```
monitor oid oramodMetricsDBSORTDISK 5004 0x0 60 absolute > 50000 'Total
Number Database Disk Sorts' ''
```

Using SystemEDGE History Collection

This section provides examples of how to use SystemEDGE history capabilities to track the value of important Oracle metrics over time. Add the commands that are provided in the following sections to the sysedge.cf file to collect history for these MIB objects. For more information about history collection, refer to the chapter on history collection in the *eHealth SystemEDGE User Guide.*

NOTE

> The number of samples and the interval between samples used in this section may not be appropriate for your Oracle system; select values that are suitable for your environment.

Collecting History on Number of Hits to the Database

To collect history on the number hits to the database, enter the following:

```
emphistory 5000 60 oramodMetricsCHR 300 'Total Hits To Oracle Database'
```

Collecting History on Block Changes Per Transaction

To collect history on the number of block changes per transaction for the Oracle database, enter the following:

```
emphistory 5001 60 oramodMetricsBCPT 300 'Oracle Database Block Changes
Per Transaction'
```

19127
ε

Collecting History on the Number of Transactions Started

To collect history on the number of transactions started since the last instance startup, enter the following:

```
emphistory 5002 60 oramodMetricsUSERCALLS 300 'Total Transactions
Started Since Instance Startup'
```

19126

Glossary

Abstract Notation One (ASN.1) A language that describes data types independent of computer structures and representations. For more information refer to ISO International Standard 8824.

AdvantEDGE View A Web-based management interface for use with the SystemEDGE agent that enables an administrator to use a Web browser to manage systems and applications.

agent In network management, a program that provides information from a management information base (MIB) for SNMP agents. *e*Health or a network management system (NMS) use the information about managed devices and take corrective action when appropriate.

American Standard Code for Information Interchange (ASCII) The most common format for character representation in computers and the Internet. Characters fit into a single byte. It was developed by the American National Standards Institute (ANSI).

application A program that performs a specific function for one or more users or for another application program. Types of applications include communication programs, management programs, word processors, databases, and drawing programs.

ASCII *See* American Standard Code for Information Interchange (ASCII).

ASN.1 *See* Abstract Notation One (ASN.1).

baseline A level of performance that is considered normal, average, or typical over a period of time such as a day, week, or month. Compare current performance metrics against baseline data to identify trends in performance levels and service delivery.

buffer A temporary storage area for data. Often implemented as holding areas between the backplane and an interface; data remains in the buffer until it can be transmitted on the interface or processed by the central processing unit (CPU).

capacity A measurement of the volume that an element can support. For interfaces, this is the bandwidth that can be carried. For hard disks, this is the disk size or the amount of information that can be stored on the disks. *See also* traffic.

central processing unit (CPU) The component within a device that performs the instruction execution for the applications and programs that run on the device. Also referred to as a processor or microprocessor.

client A computer system, usually a desktop computer or laptop, that presents data directly to a user and accepts input. They drive the computing process, supporting local processing and accessing remote servers as needed for data access and analysis.

Also refers to the application software residing on a machine that is used by an end user.

congestion A condition in which the network traffic is greater than the amount that the network can carry. Often causes performance problems and delays on a network.

CPU *See* central processing unit (CPU).

19124
E

Database Management System (DBMS) A program such as Oracle, Microsoft SQL Server, or Sybase for creating and providing access to one or more databases.

delay The time required for a packet or frame to travel from the sending station (source) to the receiving station (destination).

disk thrashing A condition that results when a server performs high disk input /output (I/O) operations—reads and writes to the disk—without producing actual work. Often occurs when a server performs excessive paging and swapping due to physical memory limitations.

DNS *See* domain name system (DNS).

domain name system (DNS) The system that locates and translates Internet domain names such as concord.com into Internet Protocol (IP) addresses. A DNS server is typically a device that translates domain names to IP addresses within your network.

eHealth AIM *See* eHealth application insight module.

eHealth application insight module A plug-in (supplementary program) that extends the functionality of the SystemEDGE agent. AIMs add the capability to manage application-specific events, processes, thresholds, and health.

event An occurrence on a system that typically results in a message, such as an SNMP trap, being sent to a configured management system. Common events include system failures, system reboots, exceeded thresholds, or any user-configurable situation that the user wants to identify.

File Transfer Protocol (FTP) A means for uploading and downloading files on the Internet (the oldest Internet protocol for retrieving files). You can use an FTP client application to request files from or transfer files to an FTP server.

FTP *See* File Transfer Protocol (FTP).

Host Resources MIB A MIB (management information base) that defines a set of objects that are useful for the management of host computers. For example, it defines host storage areas, devices, and file systems. This MIB is defined in RFC 1514.

hostname The name for an individual IP (Internet Protocol) address on a computer. While many computers have only one hostname, some machines, such as network servers have multiple hostnames.

HTML *See* Hypertext Markup Language (HTML).

HTTP *See* Hypertext Transfer Protocol (HTTP).

Hypertext Markup Language (HTML) A programmatic language used for controlling the way that text and images appear when a file is displayed on the World Wide Web.

Hypertext Transfer Protocol (HTTP) An application protocol that defines the set of rules for exchanging files (text, graphics, multimedia, and other files) on the World Wide Web.

Hertz (Hz) A unit of frequency of one cycle per second that measures the change in the state of an alternating current, sound wave, or other cyclical wave form.

I/O *See* input/output (I/O).

Information Technology (IT) A widely-used term to describe all of the technologies used for creating, exchanging, managing, and using information in various forms.

input/output (I/O) Any operation, program, or device that transfers data to or from a computer.

internet infrastructure The applications, systems, and networks that a company uses to run its business, for both internal use and for interfaces to the outside world.

19122
E

Internet Protocol (IP) The method (or protocol) by which packets of information are sent across the Internet. IP defines addressing, error handling, routing, and option codes for data transmission. IP requires no continuing connection between the endpoints that are communicating.

IP *See* Internet Protocol (IP).

IT *See* Information Technology (IT).

KB Kilobytes.

latency A measure of delay, often network delay. Depending on the type of element, *e*Health reports can show two types of latency: round-trip latency, which is the length of time in milliseconds for a ping packet to travel from the *e*Health system to a polled element and back. Alternate latency, which is the length of time in milliseconds for a ping packet to travel from a network resource (the alternate latency source) such as a router to other critical network resources such as routers and servers (the alternate latency partner).

management information base (MIB) A formal description of a set of network objects that can be managed using Simple Network Management Protocol (SNMP).

MB Megabytes.

MIB *See* management information base (MIB).

network A collection of computers, printers, routers, switches, and other devices that are able to communicate using a common transmission media such as TCP/IP.

network management system (NMS) An application program usually residing on a computer that manages at least part of a network, including systems and applications. The NMS communicates with agents to monitor network statistics and resources, control network device configuration, and analyze network problems. *See also* agent.

NMS *See* network management system (NMS).

object identifier (OID) a unique identifier of a managed object in a MIB hierarchy. *See also* management information base (MIB).

OID *See* object identifier (OID).

operating system (OS) The program that manages all other programs (applications or application programs) on a computer. Provides the following services: determining the order in which each application runs and the time allotted for that application, managing the sharing of internal memory among multiple applications and handling input to and output from attached hardware devices.

OS *See* operating system (OS).

packet A logical unit of data routed between an origin and a destination on the Internet or any other packet-switched network. On the Internet, the Transmission Control Protocol (TCP) layer of TCP/IP divides a file into packets of manageable size for routing.

page In computers that utilize virtual memory, a unit of data storage. Systems transfer pages of data from disk storage to memory and back again.

On the World Wide Web, a file written using Hypertext Markup Language (HTML) that specifies how text, images, and other multimedia will be presented to the user. A Web site delivers information to the user one page at a time.

paging The process by which a computer moves portions of programs between random access memory and auxiliary storage (on disk).

partition A logical division of a hard disk on a PC that is created so that each partition can have a different operating system or can be used for different purposes (for example, file management or multiple users).

19120
E

path In networking, a path is a route from one location to another in a network. In an Asynchronous Transfer Mode (ATM) network, a path is a virtual pipe that can carry a number of channels.

PC *See* personal computer (PC).

performance threshold The upper limit of acceptable response time.

personal computer (PC) A computer designed for individual use. Prior to the PC, computers were designed to be used by many individuals and system resources were shared by all. A PC often refers to a computer with an Intel microprocessor architecture and an operating system such as Microsoft DOS or Windows.

ping An Internet echo message used to confirm the reachability of a network device. An abbreviation for Packet Internet or Inter-Network Groper.

port The physical (hardware) connection on a device that connects the device to a network.

process Typically, an instance of a program or application that is running on a server. Applications can have one or more associated processes. *See also* Database Management System (DBMS).

process set A collection of one or more processes that relate to a specific application. Using *e*Health – System At-a-Glance reports, you can obtain information about the impact and performance of process sets running on systems that have Concord SystemEDGE agents.

protocol The set of rules by which the endpoints in a telecommunication connection communicate. The protocol defines the packet format of the transmitted information. On the Internet, common protocols are TCP, IP, HTTP and FTP.

19119

queue In a system, a set of jobs awaiting resources. In a network device such as a router, a collection of packets waiting to be processed or forwarded. Insufficient central processing unit (CPU) speed, memory, or interface speeds can contribute to long queues, and therefore, to delay on the network.

real-time A level of computer responsiveness that an end user would deem as immediate or fast enough to show incremental changes of an external process (for example, to present visualizations of the weather as it constantly changes).

Request For Comments (RFC) The name of the document series regarding Internet design. Most RFCs define protocol specifications such as Telnet and FTP. RFCs are widely available online.

RFC *See* Request For Comments (RFC).

server A program that provides services to other programs in the same and other computers.

Also a computer that performs file storage and application hosting as well as provides computing services to other devices and users on the network. Typically has one or more central processing units (CPUs), disks, interfaces, and storage partitions.

server process A server-side part of a distributed application.

Simple Network Management Protocol (SNMP) The network management protocol used almost exclusively in data networks. A method for monitoring and controlling network devices, as well as managing configurations, statistics collection, performance, and security.

SNMP *See* Simple Network Management Protocol (SNMP).

SNMP agent A program such as the SystemEDGE agent that conforms to a management information base (MIB) specification to collect information about managed devices and to take corrective action (using SNMP traps) when appropriate.

speed The capacity (bandwidth) of an interface in bits per second (bps).

swapping The process in which a computer moves entire programs in and out of random access memory to and from auxiliary storage (swap partition or pagefile).

SystemEDGE agent Concord's SNMP agent that autonomously monitors system configuration, status, performance, users, applications, file systems, and other critical resources.

Systems Management MIB A set of MIB (management information base) objects that extends the capabilities of the Host Resources MIB. It provides greater visibility into systems and specific information about Windows NT and UNIX systems.

TCP/IP *See* Transmission Control Protocol (TCP) and "Internet Protocol (IP).

threshold *See* performance threshold.

throughput The rate of data transfer on an interface over time. At each poll, *e*Health calculates throughput by dividing the total number of bits for an interface by the elapsed time in seconds since the previous poll.

traffic The data that travels over a network.

Transmission Control Protocol (TCP) A connection-based protocol used along with the Internet Protocol (IP) to send data in the form of message units between computers over the Internet. While IP is responsible for the actual delivery of the data, TCP is responsible for dividing data into packets at the sending system and constructing the data message from individual packets at the receiving system.

19117

trap A message sent by an SNMP agent to a console or network management system (NMS) to indicate that a threshold has been reached or another user-defined condition has occurred. The SystemEDGE agent defines a number of traps for system and application management.

UDP *See* User Datagram Protocol (UDP).

User Datagram Protocol (UDP) A communications protocol that uses Internet Protocol (IP) to send and receive data and is similar to Transmission Control Protocol (TCP), but provides fewer packet management services.

variable A performance metric for an element. A characteristic or behavior upon which *e*Health gathers data and evaluates the performance of the element. SystemEDGE agents can also monitor local variables to reduce network polls and increase scalability.

Web *See* World Wide Web (WWW, Web).

workstation A powerful computer that is equipped with a fast processor, a large amount of random access memory, and other features such as high-speed graphical rendering that make it suitable for business users such as engineers, graphic designers, and architects.

World Wide Web (WWW, Web) All of the resources on the Internet that use Hypertext Transfer Protocol (HTTP). Users of the Web access information through browser software.

19116

Index

A

Abstract Notation One 67
AdvantEDGE View
  application query 13
  definition 67
  licensing 30, 32
agents 67
American National Standards Institute 67
application
  definition 67
  query for Oracle 13
archive MIB group 58
ASN.1. *See* Abstract Notation One
assigning entry rows for self-monitoring tables 60

B

backup MIB group 57

C

caching disk space calculations 22
collecting history information 64
configuration section of MIB
  database file MIB group 39
  database MIB group 36
  redo logs MIB group 40
  SGA MIB group 41
configuring *e*Health AIM for Oracle 21
controlling frequency of disk space calculations 22

D

Data Dictionary Aggregage. *See* DDA
database
  configuration MIB group 36
  file configuration MIB group 39
DDA 49
Detailed Library Cache Information. *See* DLCI
disk space calculations 22
DLCI 50

E

editing
  oramod.cf file 22
  sysedge.cf file 21, 22

19115
ε

*e*Health 15
*e*Health AIM for Oracle
    components 12
    files 20
    installation requirements 17
    overview 11
    using 59
enabling *e*Health AIM for Oracle 21, 22
evaluation copy, licensing 28
expensive SQL MIB group 54

F

files installed 20
footprint
    file footprint MIB group 44
    MIB group 42

G

generating a license 28

H

history collection 64

I

installation
    files 20
    requirements 17
installing
    *e*Health AIM for Oracle 18
    SystemEDGE agent 17

L

latches MIB group 56
license
    generating 28
    sample 33
licensing
    *e*Health AIM for Oracle 27
    evaluation copy 28
    host administration 32
Live Health 16
local policy for monitoring tables 60
locks MIB group 56

M

Management Information Base. *See* MIB
metrics MIB group 45
MIB
    configuration section 36
    performance section 41
    specification 20, 35
    using 59
monitoring
    assignment policy for tables 60
    multiple Oracle applications 22
    processes 62
    thresholds 63

O

obtaining a license 27
oramod.asn1 file 20
oramod.cf file 20
    editing 22
    example 23
oramod.pdf file 20
oramod8.so file 20
oramod8i.so file 20

1a114
E

P

performance section of MIB
archive MIB group 58
backup MIB group 57
expensive SQL MIB group 54
file footprint MIB group 44
footprint MIB group 42
latches MIB group 56
locks MIB group 56
metrics MIB group 45
redo logs MIB group 51
rollback segments MIB group 52
SGA DDA MIB group 49
SGA DLCI MIB group 50
tablespace MIB group 55
waits MIB group 53
policy for monitoring tables 60

R

redo log
configuration MIB group 40
performance MIB group 51
relnotes.txt file 20
requirements for installation 17
reserving rows in monitoring tables 60
rollback segment MIB group 52

S

setting policy for monitoring tables 60
SGA
configuration MIB group 41
DDA MIB group 49
DLCI MIB group 50
performance MIB group 48
SNMP 13
sysedge.cf file 21
sysedge.lic file 27
System Global Area. *See* SGA
SystemEDGE agent
configuring 21
definition 75
history collection 64
process monitoring 62
startup file 24
threshold monitoring 63
using with *e*Health AIM for Oracle 13

T

tablespace MIB group 55
temporary license 28
threshold monitoring 63

U

using *e*Health AIM for Oracle 13, 59
with AdvantEDGE View 13
with *e*Health 15
with Live Health 16

W

waits MIB group 53

19113

ADVANTEDGE

FOR MICROSOFT SQL SERVER

User Guide

Release 1 and Later

09-16070-002
May 2001

19112
3

Important Notice

Application Health, Server Health, Live Health, SystemEDGE, AdvantEDGE, eHealth, FirstSense Enterprise, FSagent, FServer, FStransaction, FSterminal, FSconsole, FScontroller, FSpro, Infrasense, MIB Manager, and the FirstSense logo are trademarks of Concord Communications, Inc.

The Total View, Concord, the Concord logo, Network Health, TRAKKER, FirstSense, and We See It Happening are registered trademarks of Concord Communications, Inc.

Adobe, Acrobat, and the Acrobat logo are trademarks, which may be registered in certain jurisdictions, and PostScript is a registered trademark of Adobe Systems Incorporated.
Intel is a registered trademark of Intel Corporation.
Microsoft, Microsoft SQL Server, Windows, and Windows 2000 are either registered trademarks or trademarks of Microsoft Corporation in the United States and/or other countries.

SMIC. Copyright 1992 SynOptics Communications, Inc. All Rights Reserved. SynOptics makes no representations about the suitability of this software for any particular purpose. The software is supplied "as is", and SynOptics makes no warranty, either express or implied, as to the use, operation, condition, or performance of the software. SynOptics retains all title and ownership in the software.

Network Health incorporates compression code by the Info-ZIP group. There are no extra charges or costs due to the use of this code, and the original compression sources are freely available from ftp://ftp.cdrom.com/pub/infozip/ on the Internet and from the Concord Communications Web site: http://www.concord.com.

© Copyright Bigelow and Holmes 1986, 1985. Lucida is a registered trademark of Bigelow & Holmes. Sun Microsystems Inc., AT&T, and Bigelow & Holmes make no representations about the suitability of the source code for any purpose. It is provided "as is" without express or implied warranty of any kind.

General Notice: Some of the product names used herein have been used for identification purposes only and may be trademarks of their respective companies.

Proprietary Notice

The information and descriptions contained herein are the property of Concord Communications, Inc. Such information and descriptions may not be copied, disseminated, or distributed without the express written consent of Concord Communications, Inc. Concord Communications, Inc., assumes no responsibility for any inaccuracies that may appear in this document. Concord Communications, Inc., reserves the right to improve its products and change specifications at any time without notice.

U. S. Government Restricted Rights

Use, reproduction, and disclosure by the U.S. Government are subject to the restrictions set forth in FAR §52.227-19 (c) (1) and (2) and DFARS §252.227-7013 (c) (1) (ii).

Patent Information

U. S. Patent 5,615,323
Patents Pending

Table of Contents

About This Guide

How This Guide Is Organized v
Conventions in This Guide vi
Contact Information vi

Chapter 1
Introduction

Introducing AdvantEDGE for Microsoft SQL Server 1-1
Purpose of AdvantEDGE for Microsoft SQL Server 1-2
Functionality of AdvantEDGE for Microsoft SQL Server 1-2
    Using AdvantEDGE for Microsoft SQL Server 1-2
    Using AdvantEDGE for Microsoft SQL Server with *e*Health 1-3
    Using AdvantEDGE for Microsoft SQL Server with Live Health . . 1-3

Chapter 2
Installing AdvantEDGE for Microsoft SQL Server

Installation Requirements 2-1
Installing the Software 2-1
    AdvantEDGE for Microsoft SQL Server Files 2-3
Configuring AdvantEDGE for Microsoft SQL Server 2-3
    Editing the sysedge.cf File 2-3
    Editing the sqlmod.cf File 2-4
    Creating an ODBC Connection to Microsoft SQL Server 2-5
Licensing AdvantEDGE for Microsoft SQL Server 2-6
    Obtaining a License 2-6
    Generating the License 2-7
    Sample License File 2-9

1411
E

Chapter 3
Using the AdvantEDGE for Microsoft SQL Server MIB

MIB Overview 3-1
Configuration Section 3-2
Database Description Group 3-2
SQL Server General Information Group 3-3
Advanced Configuration Group 3-4
Databases Information Group 3-4
Database Options Group 3-5
Database Devices Group 3-6
Footprint Section 3-6
General Footprint Group 3-7
Physical Disk Group 3-7
SQL Server Process Footprint Group 3-8
Performance Section 3-8
Locks Group 3-9
Access Methods Group 3-9
Log Management Group 3-10
Memory Management Group 3-11
SQL Statistics Group 3-12
Buffer Management Group 3-13
General Statistics Group 3-13
Cache Management Group 3-14
Flat Information Group 3-15
Flat Transaction Performance Group 3-16

Chapter 4
Using AdvantEDGE for Microsoft SQL Server

Assigning Entry Rows for the SystemEDGE Self-Monitoring Tables . . . 4-2
Setting Local Policy 4-2
Reserving Blocks of Rows 4-2
Using the SystemEDGE Self-Monitoring Features 4-3
Using SystemEDGE Process Monitoring 4-3
Using SystemEDGE Threshold Monitoring 4-4
Using SystemEDGE History Collection 4-5

About This Guide

The *AdvantEDGE for Microsoft SQL Server User Guide* provides instructions for installing and using AdvantEDGE for Microsoft SQL Server. This guide is intended for the person responsible for installing and configuring AdvantEDGE for Microsoft SQL Server. This guide supports AdvantEDGE for Microsoft SQL Server Release 1.0 and later and the SystemEDGE agent Release 4.0, Patchlevel 3 and later.

To use AdvantEDGE for Microsoft SQL Server, you should have a basic understanding of the Microsoft SQL Server application, the Concord SystemEDGE agent, Simple Network Management Protocol (SNMP), the Host Resources management information base (MIB), and your host's operating systems environment. For more information, refer to Microsoft documentation and the *SystemEDGE Agent User Guide*.

How This Guide Is Organized

This guide is organized as follows:

- Chapter 1, "Introduction," provides an overview of AdvantEDGE for Microsoft SQL Server and its capabilities for monitoring Microsoft SQL Server.
- Chapter 2, "Installing AdvantEDGE for Microsoft SQL Server," explains how to install and configure the Concord AdvantEDGE for Microsoft SQL Server software on a host system.
- Chapter 3, "Using the AdvantEDGE for Microsoft SQL Server MIB," describes the information that is available through the Concord AdvantEDGE for Microsoft SQL Server MIB.
- Chapter 4, "Using AdvantEDGE for Microsoft SQL Server," explains how to configure and use Concord's AdvantEDGE for Microsoft SQL Server in your host environment.

09/08

Conventions in This Guide

The following paragraph shows a sample command. Throughout this guide, commands are formatted this way to distinguish them from other information. Due to space limitations in this guide, some commands wrap from one line to the next. Disregard these line breaks and **enter each command on one line.** Otherwise, your command syntax will be incorrect.

```
watch process procAlive 'sqlservr' 5000 0x0
30 'Microsoft SQL Server' ''
```

You must enter the command on one line, as shown here:

```
watch process procAlive 'sqlservr' 5000 0x0 30 'Microsoft SQL Server' ''
```

Contact Information

If you need any assistance with the SystemEDGE agent or the AdvantEDGE for Microsoft SQL Server Point module, contact Customer Support, using one of the following methods:

Phone:	(888) 832-4340 (for calls from the USA and Canada) (508) 303-4300 (for calls from other countries)
Fax:	(508) 303-4343
Email:	support@concord.com
Web site:	http://www.concord.com

19/07

1

Introduction

This chapter provides an overview of the purpose and features of AdvantEDGE for Microsoft SQL Server.

Introducing AdvantEDGE for Microsoft SQL Server

AdvantEDGE for Microsoft SQL Server is a plug-in for the SystemEDGE agent that brings the intelligent self-management capabilities of the SystemEDGE agent to the application level. You can configure this plug-in to monitor the Microsoft SQL Server processes and features that are relevant to your organization. You can also configure it to alert you to any potential issues with the application or the system on which it is running before those issues become problems.

AdvantEDGE for Microsoft SQL Server makes important information about Microsoft SQL Server available to network management software through the SystemEDGE agent and Simple Network Management Protocol (SNMP). It can provide information at the system, application, and database levels.

AdvantEDGE for Microsoft SQL Server enables you to monitor the health and availability of Microsoft SQL Server running on the Windows NT or Windows 2000 operating system. It must be installed on every SQL Server system that needs to be monitored. For more information, refer to Chapter 2, "Installing AdvantEDGE for Microsoft SQL Server."

NOTE

This document is not intended as a manual on how to install, administer, or use Microsoft SQL Server. For assistance, refer to Microsoft documentation.

19106
ε

Purpose of AdvantEDGE for Microsoft SQL Server

At its best, an instance of Microsoft SQL Server can process queries from thousands of concurrent users who connect over a network. It can support data warehouses or data marts that process complex queries required to discover trends and analyze critical factors in enterprise activity.

Despite advances in the design of database management systems in the last ten years, issues remain that frequently impact the health and availability of Microsoft SQL Server.

These issues include the following:

- Performance tuning, which involves trade-offs between resources for queries and resources for indexing columns
- Table locking, which can prevent users from being able to update records
- Tuning the maximum size of transaction logs to ensure that records can be updated

The purpose of AdvantEDGE for Microsoft SQL Server is to provide you with the tools and information necessary for monitoring the health and availability of Microsoft SQL Server.

Functionality of AdvantEDGE for Microsoft SQL Server

AdvantEDGE for Microsoft SQL Server accomplishes its purpose depending on the way it is deployed. It can operate with any SNMP-compliant management software such as Concord's *e*Health suite of products, AdvantEDGE View, HP OpenView, and others.

Using AdvantEDGE for Microsoft SQL Server

As a plug-in for the SystemEDGE agent, AdvantEDGE for Microsoft SQL Server works to closely manage the Microsoft SQL Server application. Right out-of-the-box, it provides real-time fault detection and is capable of correcting problems without human intervention.

Using AdvantEDGE for Microsoft SQL Server with *e*Health

When teamed with the *e*Health product suite, AdvantEDGE for Microsoft SQL Server and the SystemEDGE agent provide the historical data for long-term trending analysis and capacity planning.

With *e*Health – Application Assessment, you can run At-a-Glance, Trend, Top N, and MyHealth reports for the following types of variables:

- Amount of central processing unit (CPU), memory, and disk space the Microsoft SQL Server application is using
- Size of the Microsoft SQL Server configuration and database logs
- Transaction log size and percentage of transaction log space used
- Frequency of hits within Microsoft SQL Server caches

For more information about the variables that you can monitor and the reports that you can run when you integrate AdvantEDGE for Microsoft SQL Server with *e*Health, refer to the *e*Health Web Help.

Using AdvantEDGE for Microsoft SQL Server with Live Health

In a deployment with Live Health, AdvantEDGE for Microsoft SQL Server and the SystemEDGE agent provide the data for real-time detection of faults, potential outages, and delays associated with Microsoft SQL Server database activity. Unlike other real-time monitoring solutions and network management systems, Live Health applies intelligent algorithms to the data, resulting in precise assessments of application health and database performance. For more information about how Live Health can detect brownouts and service delays across applications, systems, and networks, refer to the Live Health Web Help.

19104

2

Installing AdvantEDGE for Microsoft SQL Server

This chapter explains how to install, configure, and license the AdvantEDGE for Microsoft SQL Server Point module.

NOTE

> For the most current information about installing AdvantEDGE for Microsoft SQL Server, refer to the release notes (relnotes.txt) on the installation CD-ROM.

Installation Requirements

Before you install AdvantEDGE for Microsoft SQL Server, you must first install, license, and configure the SystemEDGE agent Release 4.0, Patchlevel 3 or later. For more information, refer to the *SystemEDGE Agent User Guide*. Also, your system must be running Microsoft SQL Server 7.0 or Microsoft SQL Server 2000 on Windows NT 4.0 or later. For more information, refer to the Microsoft documentation.

Installing the Software

AdvantEDGE for Microsoft SQL Server is distributed as a self-extracting executable named sqlmod.exe.

To install it:

1. Locate the SQL Server system that you need to monitor.
2. Log on to the system as the administrator.
3. Click **Start**.
4. Select **Programs** → **Command Prompt**.

19/03
3

5. Insert the CD-ROM containing the Concord software distributions into the CD-ROM drive.

   Windows NT automatically mounts the drive using the CD-ROM drive's corresponding drive letter.

   NOTE
   The particular drive letter is specific to your system and depends on the number and types of disks attached to your system.

6. Determine the directory that you want to use as the installation directory for AdvantEDGE for Microsoft SQL Server. If the SystemEDGE agent is installed at C:\sysedge, you should use C:\sysedge\plugins.

7. Run the self-extracting executable by entering the following at the command prompt:

   ```
   D:\sqlmod\ntx86\sqlmod.exe -dir C:\sysedge\plugins
   ```

   where *D* is the CD-ROM drive for your system, and `C:\sysedge\plugins` is the installation directory.

   NOTE
   The -dir option is important because it instructs the self-extracting executable to recreate the intended subdirectory hierarchy used throughout this manual.

   The executable then places the distribution in an sqlmod subdirectory within the specified target directory (for example, C:\sysedge\plugins\sqlmod).

NOTE
You cannot execute sqlmod.exe directly from the CD-ROM.

AdvantEDGE for Microsoft SQL Server is now installed.

19102
3

AdvantEDGE for Microsoft SQL Server Files

Table 2-1 describes the files created by the installation process.

Table 2-1: Files Installed by AdvantEDGE for Microsoft SQL Server

Filename	Description
sqlmod.asn1	AdvantEDGE for Microsoft SQL Server MIB specification
sqlmod.cf	AdvantEDGE for Microsoft SQL Server configuration file
sqlmod.dll	AdvantEDGE for Microsoft SQL Server dynamic link library (DLL) for Windows NT and Windows 2000
sqlmod.pdf	AdvantEDGE for Microsoft SQL Server User Guide
examples	AdvantEDGE for Microsoft SQL Server monitoring examples
relnotes.txt	Release notes for AdvantEDGE for Microsoft SQL Server

Configuring AdvantEDGE for Microsoft SQL Server

After you install AdvantEDGE for Microsoft SQL Server, you must configure the Point module by editing the following files:

- sysedge.cf
- sqlmod.cf

Editing the sysedge.cf File

The SystemEDGE agent uses the configuration file sysedge.cf and the *sysedge_plugin* keyword to specify which AdvantEDGE modules to load at system start time. By default, the SystemEDGE agent does not load any plug-ins at initialization time, but you can edit the sysedge.cf file to configure the agent to load any AdvantEDGE Point modules that you have installed. The sysedge.cf file is located in your system directory (for example, C:\winnt\system32). For more information about the sysedge.cf file, refer to the *SystemEDGE Agent User Guide.*

To configure the SystemEDGE agent to start AdvantEDGE for Microsoft SQL Server at system start time, open the sysedge.cf file for editing; then add the following lines:

```
# SQL Server
sysedge_plugin C:\sysedge\plugins\sqlmod\sqlmod.dll
```

where C:\sysedge\plugins\sqlmod is the directory in which you installed the AdvantEDGE for Microsoft SQL Server files. This command line provides the complete pathname to sqlmod.dll, the AdvantEDGE for Microsoft SQL Server dynamic link library.

Editing the sqlmod.cf File

The sqlmod.cf file describes the AdvantEDGE for Microsoft SQL Server Point module configuration. Make sure that this file is located in the same directory as sqlmod.dll (for example, C:\sysedge\plugins\sqlmod). If you are running more than one instance of Microsoft SQL Server on your system, edit the configuration file to associate the Open Database Connectivity (ODBC) connections with the instances that you want to monitor.

The following is a sample sqlmod.cf file.

```
sqlmod 1 "{SQL Server}" SLOWPOKE MSSQLODBC sa password
sqlmod 2 "{SQL Server}" SLOWPOKE\INSTANCE1 ODBCInstance1 sa password
sqlmod 3 "{SQL Server}" SLOWPOKE\INSTANCE2 ODBCInstance2 sa password
```

- Position 1: sqlmod.
- Position 2: The index number uniquely assigned to the instance.
- Position 3: The name of the driver to connect with Microsoft SQL Server; it will usually be "{SQL Server}". The driver name must be enclosed in brackets.
- Position 4: The name of the Microsoft SQL Server to which to connect.
- Position 5: The name of the ODBC connection.
- Position 6: A user name; for example, "sa" would be an appropriate choice for a user with a system administrator role.
- Position 7: An appropriate password.

NOTE

The user that you specify in this file does **not** need to be a system administrator and does **not** need db_owner level access.

19100
e

Editing Guidelines

When editing the sqlmod.cf file, follow these guidelines:

- Enclose an entry in quotation marks if it contains a space character.
- After you modify the sqlmod.cf file, save it and restart the SystemEDGE agent.
- After restarting the SystemEDGE agent, license the module as described in “Licensing AdvantEDGE for Microsoft SQL Server” on page 2-6.
- The user that you specify in sqlmod.cf does not need to have a server role, but **must** have at least db_datareader access to the master database. You can set up and verify this access from the Enterprise Manager under Security -> Logins.

Creating an ODBC Connection to Microsoft SQL Server

The system administrator should set up an ODBC connection to use Microsoft SQL Server authentication with a login ID and password. This connection stays open after it is established.

NOTE

> It is important that the username assigned for the ODBC connection have the proper Microsoft SQL Server access permissions. Failure to do so may cause SQL Server operation to lock up when attempting the connection.

To create a new ODBC connection using Microsoft SQL Server authentication:

1. Log on to the system as the administrator.
2. Click **Start**.
3. Select **Settings → Control Panel**. The Control Panel window appears.
4. Double-click the **ODBC Data Sources** icon. (In Windows 2000, this is in the Administrative Tools folder of the Control Panel window.) The ODBC Data Sources Administrator window appears.
5. Select the **User DSN** tab (if it is not already selected.)
6. Select **Add**.
7. Select **SQL Server** from the list of ODBC drivers and click **Finish**. The Create a New Data Source to SQL Server wizard appears.
8. Specify a connection name in the **Name** field.

19099
ε

9. Optionally, specify a description of the data source in the **Description** field.
10. Specify the server for the ODBC connection in the **Server** field.
11. Select **Next**.
12. Select **With SQL Server authentication using a login ID and password entered by the user**.
13. Select **Next**.
14. Select **Next** to accept the default values. The ODBC Microsoft SQL Server Setup window appears.
15. Select **Test Data Source**.

Licensing AdvantEDGE for Microsoft SQL Server

Like the SystemEDGE agent, AdvantEDGE for Microsoft SQL Server utilizes a host-based license method. Copies of AdvantEDGE for Microsoft SQL Server can run only on systems that possess a valid license key. This license is separate from the one used for the SystemEDGE agent.

The first time that you attempt to start the SystemEDGE agent after installing AdvantEDGE for Microsoft SQL Server, the agent displays a message stating that it did not find a valid license for AdvantEDGE for Microsoft SQL Server. It then provides you with a *public key* that you can use to generate a permanent license key for your host machine.

A license key is made up of four space-separated, 8-character sequences, totaling 32 characters. The sysedge.lic file contains the AdvantEDGE for Microsoft SQL Server license, as well as the SystemEDGE agent license and other AdvantEDGE Point module licenses. For an example of a license file, refer to "Sample License File" on page 2-9.

NOTE

If you are using an evaluation copy of AdvantEDGE for Microsoft SQL Server, you must request a temporary license that will enable AdvantEDGE for Microsoft SQL Server to operate during the evaluation period.

19098
ε

Obtaining a License

To obtain a license, choose one of the following methods:

- Run the Concord licenseutil.pl script.
- Run the licenseme.exe license utility.
- Use the AdvantEDGE View licensing procedure; refer to the AdvantEDGE View Web Help.
- Send an e-mail request to license@empire.com.
- Complete the online license form through the Internet, as described in "Generating the License" on page 2-7.

You can obtain a license at any time from the Concord licensing Web server (http://license.concord.com). If you use AdvantEDGE View licensing, AdvantEDGE View can automatically retrieve and set up licenses for the SystemEDGE agent and the plug-ins. For more information, refer to the *Automating the Licensing of SystemEDGE and AdvantEDGE Point Plugin Modules* white paper and the *SystemEDGE Agent User Guide.*

Generating the License

The SystemEDGE setup program generates the licensing information for your system.

To generate a license:

1. Run the SystemEDGE agent setup command by entering the following at the command prompt:

   ```
   sysedge\setup -l
   ```

   The setup program displays a message similar to the following:

```
SystemEDGE Version 4.0 Patchlevel 3
Copyright 2001 by Concord Communications, Inc.
Please contact Concord Communications, Inc. to obtain a license
http://www.concord.com/support, Email: license@concord.com
Provide: sysedge neptune NTx86 5.0 346561363366b19c 4.0 Patchlevel 3
```

2. Complete the Web-based license form for AdvantEDGE for Microsoft SQL Server available from Concord's licensing Web server at the following URL:

   http://license.concord.com

14097
E

NOTE

You must supply a user name and password to access the license form.

Supply the following information on the form:

- Customer ID
- Name
- E-mail address
- Software version number (4.0 in the example above)
- Patchlevel
- System name (neptune in the example above)
- Operating system name (NTx86 in the example above)
- Version (5.0 in the example above)
- System identifier (346561363366b19c in the example above)

NOTE

When you are licensing AdvantEDGE for Microsoft SQL Server, select **sqlmod** as the product on the licensing form.

After you submit the license request, the Concord Web server generates a license and displays it to your Web browser. It also e-mails the license to the contact person in your organization.

3. Copy the generated license key into the sysedge.lic file in the system32 subdirectory (C:\winnt\system32) and save this file.

NOTE

The license key is case-sensitive. Copy it exactly as it appears. If possible, use your system's cut-and-paste feature instead of entering it manually. If you do enter the license key manually, ensure that you do not confuse characters such as the letters l and I and the number 1, or the letter O and the number 0.

4. Stop and restart the Windows NT Master agent by entering these commands at the command prompt:

```
net stop snmp
net start snmp
```

AdvantEDGE for Microsoft SQL Server is now licensed and ready to use.

19096
ε

Sample License File

The following is a sample SystemEDGE agent license file. A pound character (#) in column 1 indicates that the entire line is a comment.

```
# license file for SystemEDGE Agent
# Empire Technologies, Inc.
# A Concord Communications Company
# http://www.empire.com
#
# file /etc/sysedge.lic or %SystemRoot%\system32\sysedge.lic
# A valid license key has four parts of 8 characters per part
# parts are separated by space(s) with one license key per
  line
# sysedge neptune NTx86 5.0 807cb1da007cb1da 4.0
e13311d3 0F2a7cb1 abC512dc fF8C923a
# sqlmod neptune NTx86 5.0 807cb1da007cb1da 1.0
a7943fde 098a87ij a4kiuf39 afafEkj4
```

19095
E

3

Using the AdvantEDGE for Microsoft SQL Server MIB

This chapter outlines the information available from the Concord Communications Management Information Base (MIB) for the Microsoft SQL Server. The MIB specification (sqlmod.asn1) defines a collection of objects for monitoring and managing Microsoft SQL Server. All MIB objects related to AdvantEDGE for Microsoft SQL Server exist at object identifier (OID) branch 1.3.6.1.4.1.546.16.8 in the Concord Systems Management MIB.

This chapter explains the organization and content of the AdvantEDGE for Microsoft SQL Server MIB. You must configure the SystemEDGE agent to monitor the MIB objects that are relevant to your configuration. For that information, refer to Chapter 4, "Using AdvantEDGE for Microsoft SQL Server."

MIB Overview

The MIB is organized into broad sections for server configuration, footprint, and performance. The server configuration section describes MIB objects that capture data such as database names, database creation dates, database status, and database device names, as well as overall server-level activity indicators. The footprint section defines MIB objects that convey how much of the underlying system's resources are consumed by Microsoft SQL Server. The performance section contains MIB objects that capture data regarding lock requests, access methods, transaction log activity, memory management, SQL statistics, buffer management, and cache management.

SQL Server 2000 supports multiple instances of the sqlserver process. The Concord Communications MIB for Microsoft SQL Server was specifically designed to handle data from multiple instances. The information set contained in this MIB is unique.

19094
E

The following sections briefly describe the MIB objects. Refer to the AdvantEDGE for Microsoft SQL Server MIB Specification (sqlmod.asn1) for a list of MIB objects and their syntax sequences and statuses.

Configuration Section

The Configuration section of the AdvantEDGE for Microsoft SQL Server MIB contains configuration parameters and settings that are important for monitoring the Microsoft SQL Server process and configuring databases.

Database Description Group

The Database Description group contains high-level information about the Microsoft SQL Server application including the status of the database. Table 3-1 describes the Database Description MIB objects.

Table 3-1: SQL Server Database Description MIB Objects

MIB Object	Description
sqlmodDBDescInstanceIdx	SQL Server instance InstanceIdx
sqlmodDBDescVendor	Database server vendor
sqlmodDBDescProdName	Product name for the database server
sqlmodDBDescVers	SQL Server version
sqlmodDBDescContact	Contact individual or organization
sqlmodDBDescStatus	Number indicating SQL Server status where 1=initialized, 2=ready, 3=running, 4=standby, 5=terminated, 6=waiting, 7=transitioning, 8=unknown
sqlmodDBDescUptime	SQL Server uptime in msec
sqlmodDBDescIsClustered	Number indicating SQL Server clustering status where 1=not clustered, 2=clustered, 3=not determined

19093

SQL Server General Information Group

The SQL Server General Information group contains MIB objects that capture basic information about an instance of Microsoft SQL Server and the status of this process. Table 3-2 describes the SQL Server General Information MIB objects.

Table 3-2: SQL Server General Information MIB Objects

MIB Object	Description
sqlmodDBGenProcInstanceIdx	SQL Server instance InstanceIdx
sqlmodDBGenProcServName	Name of SQL Server instance
sqlmodDBGenProcStartTime	Date and time that the SQL Server instance started
sqlmodDBGenProcCompPageRead	Number of physical page reads completed for this instance since it started
sqlmodDBGenProcCompPageWrite	Number of physical page writes completed for this instance since it started
sqlmodDBGenProcTDSPackRead	Number of TDS packets read from network for this instance since it started
sqlmodDBGenProcTDSPackWrite	Number of TDS packets written to network for this instance since it started
sqlmodDBGenProcBlkngLocks	Total number of blocking locks for this instance
sqlmodDBGenProcUsersBlkd	Number of users blocked by other users for this instance

19092

Advanced Configuration Group

The Advanced Configuration group contains MIB objects that capture data about advanced configuration parameters such as locks, query wait, and maximum number of worker threads. Table 3-3 describes the Advanced Configuration MIB objects.

Table 3-3: Advanced Configuration MIB Objects

MIB Object	Description
sqlmodDBAdvCfgInstanceIdx	SQL Server instance InstanceIdx
sqlmodDBAdvCfgParams	SQL Server configuration values (For a list of values, refer to the sqlmod.asn1 file.)
sqlmodDBAdvCfgParamName	Description of parameter
sqlmodDBAdvCfgMin	Minimum value of parameter
sqlmodDBAdvCfgMax	Maximum value of parameter
sqlmodDBAdvCfgCfg	Configured value for parameter
sqlmodDBAdvCfgRun	Runtime value for parameter

Databases Information Group

The Databases Information group includes MIB objects that report on total database size, transaction log size, and other measures of database management activity. Table 3-4 describes the Database Information MIB objects.

Table 3-4: Database Information MIB Objects (Page 1 of 2)

MIB Object	Description
sqlmodDBDbasesInfoInstanceIdx	SQL Server instance InstanceIdx
sqlmodDBDbasesInfoDBID	SQL Server DBID
sqlmodDBDbasesInfoName	SQL Server database name
sqlmodDBDbasesInfoState	SQL Server database state. (For a list of values, refer to the sqlmod.asn1 file.)
sqlmodDBDbasesInfoCreat	Database creation date and time
sqlmodDBDbasesInfoOwn	Database owner
sqlmodDBDbasesInfoSz	Database size in KB

19091

Table 3-4: Database Information MIB Objects (Page 2 of 2)

MIB Object	Description
sqlmodDBDbasesInfoUnalloc	Unallocated space in KB
sqlmodDBDbasesInfoReserv	Reserved space in KB
sqlmodDBDbasesInfoUseddata	Used data space in KB
sqlmodDBDbasesInfoUsedIdx	Used index space in KB
sqlmodDBDbasesInfoUnused	Unused space in KB
sqlmodDBDbasesInfoLasttranslog	Date and time of last transaction log dump
sqlmodDBDbasesInfoTranslogsz	Transaction log size in KB
sqlmodDBDbasesInfoTranslogsp	Percent of transaction log space used
sqlmodDBDbasesInfoLastfull	Date and time of last transaction log full backup
sqlmodDBDbasesInfoLastdifferential	Date and time of last transaction log differential backup
sqlmodDBDbasesInfoLastfilegroup	Date and time of last transaction log file/filegroup backup

Database Options Group

The Database Options group relates the database ID and database name to a database option which has been set. Table 3-5 describes the Database Options MIB objects.

Table 3-5: Database Options MIB Objects

MIB Object	Description
sqlmodDBDbasesOptInstanceIdx	SQL Server instance InstanceIdx
sqlmodDBDbasesOptDBID	SQL Server database ID
sqlmodDBDbasesOptOptID	Option index
sqlmodDBDbasesOptName	Name of SQL Server database to which option applies
sqlmodDBDbasesOptOption	Name of the database option which is set

19090
E

Database Devices Group

The Database Devices group includes MIB objects for device-specific information such as device description, device status, and device size. Table 3-6 describes the Database Devices MIB objects.

Table 3-6: Database Devices MIB Objects

MIB Object	Description
sqlmodDBDbasesDevInstanceIdx	SQL Server instance InstanceIdx
sqlmodDBDbasesDevDevice	SQL Server device index
sqlmodDBDbasesDevName	SQL Server device name
sqlmodDBDbasesDevPhysName	SQL Server device physical name
sqlmodDBDbasesDevDescrip	SQL Server device description
sqlmodDBDbasesDevStatus	SQL Server device status
sqlmodDBDbasesDevCtrlType	SQL Server device controller type
sqlmodDBDbasesDevSize	SQL Server device size

Footprint Section

The Footprint section of the AdvantEDGE for Microsoft SQL Server MIB contains footprint statistics that are important for monitoring the SQL Server process and its impact on the system. Long-term trend analysis of footprint information is useful for anticipating and avoiding problems due to resource exhaustion. You can also monitor footprint information in real time to detect and correct temporary resource exhaustion due to viruses, security incidents, and hardware failures.

General Footprint Group

The General Footprint group contains MIB objects that pertain to the system, not just to an instance of SQL Server. Table 3-7 describes the General Footprint MIB objects.

Table 3-7: General Footprint MIB Objects

MIB Object	Description
sqlmodFootprintTotCPUTime	Total CPU time for all SQL Server instances
sqlmodFootprintTotPercentCPU	Total percent of CPU usage for all SQL Server instances
sqlmodFootprintTotMEMSize	Total memory usage for all SQL Server instances
sqlmodFootprintTotRSS	Total resident set size for all SQL Server instances
sqlmodFootprintTotPercentMEM	Total percent of memory usage for all SQL Server instances
sqlmodFootprintTotThreads	Total number of threads for all SQL Server instances
sqlmodFootprintTotFaults	Total number of page faults for all SQL Server instances

Physical Disk Group

The Physical Disk group contains MIB objects that describe the physical disks in which database files are stored. Table 3-8 describes the Physical Disk MIB objects.

Table 3-8: Physical Disk MIB Objects (Page 1 of 2)

MIB Object	Description
sqlmodFootprintPhysicalDiskIdx	Number indicating the disk drive being monitored where 1 = total, 2 = drive C, 3 = drive D, 4 = drive E, 5 = drive F, etc.
sqlmodFootprintPhysicalDiskWhichdr	Name of the disk drive (for example, "C", "D") or "Total"

19088
ε

Table 3-8: Physical Disk MIB Objects (Page 2 of 2)

MIB Object	Description
sqlmodFootprintPhysicalDiskDisktrans	Total disk transfers
sqlmodFootprintPhysicalDiskCurrdskquelen	Current disk queue length
sqlmodFootprintPhysicalDiskDskbytes	Total number of bytes moved

SQL Server Process Footprint Group

The SQL Server Process Footprint group contains MIB objects that report the impact of the SQL Server process on the CPU and virtual memory. Table 3-9 describes the SQL Server Process Footprint MIB objects.

Table 3-9: SQL Server Process Footprint MIB Objects

MIB Object	Description
sqlmodFootprintProcessInstanceIdx	SQL Server InstanceIdx
sqlmodFootprintProcessMEM	Percent of real memory used by the process
sqlmodFootprintProcessSize	Combined size of the text, data, and stack segments in bytes
sqlmodFootprintProcessRSS	Size of resident set used by the process in kilobytes
sqlmodFootprintProcessTime	Accumulated CPU time used by the process in seconds
sqlmodFootprintProcessThreads	Number of threads used by the process
sqlmodFootprintProcessFaults	Number of page faults related to the process
sqlmodFootprintProcessPercentCPU	Percent of CPU usage by the process

Performance Section

The Performance section of the AdvantEDGE for Microsoft SQL Server MIB contains MIB objects that are indicators of the availability and performance of the application. This information is necessary for capacity planning and trend analysis.

19087
ε

Locks Group

The Locks group contains the MIB objects that indicate the frequency of users being prevented from updating records in the database. Table 3-10 describes the Locks MIB objects.

Table 3-10: Locks MIB Objects

MIB Object	Description
sqlmodLocksInstanceIdx	SQL Server InstanceIdx
sqlmodLocksWhichlock	Number indicating lock type where 1 = extent, 2 = key, 3 = page, 4 = table, 5 = rid, 6 = database, and 7 = total
sqlmodLocksAvgwaittime	Average wait time (in msec) for lock request
sqlmodLocksLockreq	Number of lock requests
sqlmodLocksLocktimeouts	Number of lock timeouts
sqlmodLocksLockwaittime	Total lock wait time (in msec)
sqlmodLocksLockwaits	Number of lock requests that require a wait
sqlmodLocksDeadlocks	Number of deadlocked lock requests

Access Methods Group

The Access Methods group contains MIB objects that reflect the types of searching that the SQL Server process performs during execution. Table 3-11 describes the Access Methods MIB objects.

Table 3-11: Access Methods MIB Objects (Page 1 of 2)

MIB Object	Description
sqlmodAccessMethodsInstanceIdx	SQL Server InstanceIdx
sqlmodAccessMethodsExtalloc	Number of extents allocated to database objects
sqlmodAccessMethodsFrwdrec	Number of records fetched through forwarded pointers
sqlmodAccessMethodsFullscans	Number of unrestricted full scans
sqlmodAccessMethodsIdxsearch	Number of index searches

Table 3-11: Access Methods MIB Objects (Page 2 of 2)

MIB Object	Description
sqlmodAccessMethodsPagesplits	Number of page splits (overflowing index pages)
sqlmodAccessMethodsPagesalloc	Number of pages allocated to database objects
sqlmodAccessMethodsProbescans	Number of probe scans
sqlmodAccessMethodsRangescans	Number of range scans
sqlmodAccessMethodsSkipghostrec	Number of ghosted records skipped
sqlmodAccessMethodsTablelockescl	Number of times locks on a table were escalated
sqlmodAccessMethodsWorktblcreat	Number of worktables created

Log Management Group

The Log Management group includes MIB objects that relate instances of the database to key indicators of transaction log activity. Table 3-12 describes the Log Management MIB objects.

Table 3-12: Log Management MIB Objects (Page 1 of 2)

MIB Object	Description
sqlmodDatabasePerfInstanceIdx	SQL Server InstanceIdx
sqlmodDatabasePerfDBidx	SQL Server DBidx
sqlmodDatabasePerfName	Name of SQL Server database or "Total"
sqlmodDatabasePerfActtrans	Number of active transactions
sqlmodDatabasePerfBulkcopyrows	Number of rows bulk copied
sqlmodDatabasePerfBulkcopythroughput	Size of bulk data copied (in KB)
sqlmodDatabasePerfDatafilesizes	Cumulative size of data files (in KB)
sqlmodDatabasePerfLogcachehitratio	Percent of reads satisfied from the log cache
sqlmodDatabasePerfLogfilesize	Cumulative size of the transaction log files

19085
E

Table 3-12: Log Management MIB Objects (Page 2 of 2)

MIB Object	Description
sqlmodDatabasePerfLogflushes	Number of flushes of the transaction log
sqlmodDatabasePerfLoggrowths	Number of times that the transaction log has expanded
sqlmodDatabasePerfLogshrinks	Number of times that the transaction log has shrunk
sqlmodDatabasePerfLogtruncs	Number of times that the transaction log has been truncated
sqlmodDatabasePerfLogutilized	Percent of log space used
sqlmodDatabasePerfShrinkdata	Amount of data being moved by autoshrink operations (in bytes)
sqlmodDatabasePerfTrans	Number of transactions

Memory Management Group

The Memory Management group contains MIB objects that characterize management of dynamic memory. Table 3-13 describes the Memory Management MIB objects.

Table 3-13: Memory Management MIB Objects (Page 1 of 2)

MIB Object	Description
sqlmodMemoryManagerInstanceIdx	SQL Server InstanceIdx
sqlmodMemoryManagerConnmem	Total dynamic memory used for connections
sqlmodMemoryManagerLockmem	Total dynamic memory used for locks
sqlmodMemoryManagerMaxworkspacemem	Maximum memory for executing processes
sqlmodMemoryManagerMemgrantsout	Number of processes with a workspace memory grant
sqlmodMemoryManagerMemgrantspend	Number of processes awaiting a workspace memory grant
sqlmodMemoryManagerOptmem	Total dynamic memory for query optimization

Table 3-13: Memory Management MIB Objects (Page 2 of 2)

MIB Object	Description
sqlmodMemoryManagerSQLcachemem	Total dynamic memory for the SQL cache
sqlmodMemoryManagerTargservmem	Total dynamic memory that the server may consume
sqlmodMemoryManagerTotservmem	Total dynamic memory that the server is using

SQL Statistics Group

The SQL Statistics group contains the MIB objects that characterize the autoparameterizations (comparisons of SQL text to a template as part of query optimization) that take place in response to batch Transact-SQL (TSQL) requests. Table 3-14 describes the SQL Statistics MIB objects.

Table 3-14: SQL Statistics MIB Objects

MIB Object	Description
sqlmodSQLStatisticsInstanceIdx	SQL Server InstanceIdx
sqlmodSQLStatisticsAutoparamattmpt	Total failed, safe, and unsafe autoparameterizations
sqlmodSQLStatisticsBatchreq	Number of TSQL batch requests
sqlmodSQLStatisticsFailedautoparam	Number of failed autoparameterizations
sqlmodSQLStatisticsSafeautoparam	Number of safe autoparameterizations
sqlmodSQLStatisticsSQLComp	Number of SQL compilations
sqlmodSQLStatisticsUnsafeautoparam	Number of unsafe autoparameterizations

19083

Buffer Management Group

The Buffer Management group contains the MIB objects that track statistics that are relevant to the efficiency of buffer management. Table 3-15 describes the Buffer Management MIB objects.

Table 3-15: Buffer Management MIB Objects

MIB Object	Description
sqlmodBufferManagerInstanceIdx	SQL Server InstanceIdx
sqlmodBufferManagerPagereads	Number of SQL Server (not total system) page reads
sqlmodBufferManagerChkptpages	Number of checkpoint pages
sqlmodBufferManagerPagewrites	Number of SQL Server page writes (inclusive)
sqlmodBufferManagerBuffercachehitratio	Percent of data being retrieved from cache
sqlmodBufferManagerLazWriteFlush	Number of pages flushed by LazyWriter for this instance since it started
sqlmodBufferManagerPhysRead	Number of physical reads by Read Ahead Manager for this instance since it started

General Statistics Group

The General Statistics group contains MIB objects that track the level of user login activity. Table 3-16 describes the General Statistics MIB objects.

Table 3-16: General Statistics MIB Objects

MIB Object	Description
sqlmodGeneralStatisticsInstanceIdx	SQL Server InstanceIdx
sqlmodGeneralStatisticsLogins	Number of logins
sqlmodGeneralStatisticsLogouts	Number of logouts
sqlmodGeneralStatisticsUserConns	Number of user connections

Cache Management Group

The Cache Management group contains MIB objects that are factors in the efficiency and effectiveness of cache management. Table 3-17 describes the Cache Management MIB objects.

Table 3-17: Cache Management MIB Objects

MIB Object	Description
sqlmodCacheManagerInstanceIdx	SQL Server InstanceIdx
sqlmodCacheManagerAdhocSQLCachehitratio	Percent of data found in cache for adhoc SQL queries
sqlmodCacheManagerAdhocSQLNumobj	Number of objects for adhoc SQL queries
sqlmodCacheManagerAdhocSQLNumpages	Number of pages for adhoc SQL queries
sqlmodCacheManagerProcplansCachehitratio	Percent of data found in cache for procedure plans
sqlmodCacheManagerProcplansNumobj	Number of objects for procedure plans
sqlmodCacheManagerProcplansNumpages	Number of pages for procedure plans
sqlmodCacheManagerTrigplansCachehitratio	Percent of data found in cache for trigger plans
sqlmodCacheManagerTrigplansNumobj	Number of objects for trigger plans
sqlmodCacheManagerTrigplansNumpages	Number of pages for trigger plans
SqlmodCacheManagerPrepSQLplansCachehitratio	Percent of data found in cache for prepared SQL plans
sqlmodCacheManagerPrepSQLplansNumobj	Number of objects for prepared SQL plans
sqlmodCacheManagerPrepSQLplansNumpages	Number of pages for prepared SQL plans

1908

Flat Information Group

The Flat Information group contains MIB objects that are totals across the set of databases for each instance of Microsoft SQL Server. Table 3-18 describes the Flat Information MIB objects.

Table 3-18: Flat Information MIB Objects

MIB Object	Description
sqlmodDBDbasesFlatInfoInstanceIdx	SQL Server instance
sqlmodDBDbasesFlatInfoANDedState	SQL Server database state where 64=pre-recovery, 128=recovering, 256=not recovered, 32768=emergency mode
sqlmodDBDbasesFlatInfoSzTotal	Database size in kilobytes
sqlmodDBDbasesFlatInfoUnallocTotal	Unallocated space in kilobytes
sqlmodDBDbasesFlatInfoTranslogszTotal	Transaction log size in kilobytes
sqlmodDBDbasesFlatInfoTranslogspTotal	Percent of transaction log space used
sqlmodDBDbasesFlatInfoLasttranslog	Date and time of last transaction log backup
sqlmodDBDbasesFlatInfoLastfull	Date and time of last transaction log full backup
sqlmodDBDbasesFlatInfoLastdifferential	Date and time of last transaction log differential backup
sqlmodDBDbasesFlatInfoLastfilegroup	Date and time of last transaction log file/filegroup backup

19080

Flat Transaction Performance Group

The Flat Transaction Performance group contains MIB objects that are totals across the set of databases for each instance of Microsoft SQL Server. Table 3-19 describes the Flat Transaction Performance MIB objects.

Table 3-19: Flat Information MIB Objects

MIB Object	Description
sqlmodDatabaseFlatPerfInstanceIdx	SQL Server instance
sqlmodDatabaseFlatPerfTransTotal	Number of total transactions
sqlmodDatabaseFlatPerfLogcachehitsTotal	Total number of log cache hits
sqlmodDatabaseFlatPerfLogcachemissesTotal	Total number of log cache misses

4

Using AdvantEDGE for Microsoft SQL Server

This chapter describes how to configure and use AdvantEDGE for Microsoft SQL Server. After you license this plug-in and enable it in the sysedge.cf file (refer to "Configuring AdvantEDGE for Microsoft SQL Server" on page 2-3), it will load automatically at SystemEDGE start time.

This plug-in implements additional MIB objects that provide advanced information about the health and availability of Microsoft SQL Server. AdvantEDGE for Microsoft SQL Server can operate with any SNMP-compliant management software such as Concord's *e*Health suite of products, AdvantEDGE View, HP OpenView, and others. If you are using AdvantEDGE for Microsoft SQL Server with *e*Health, refer to the *e*Health Web Help for more information about the reports that you can generate.

The purpose of this chapter is to illustrate how you can edit the SystemEDGE configuration file to utilize the new MIB objects with the process-monitoring, threshold-monitoring, and history-collection features of the SystemEDGE agent. For more examples of configuration file commands, refer to the examples file, which is available in the AdvantEDGE for Microsoft SQL Server product installation.

19078

Assigning Entry Rows for the SystemEDGE Self-Monitoring Tables

All of the SystemEDGE self-monitoring tables (for example, Log Monitor table, Windows NT Event Monitor table, Process/Service Monitor table, Threshold Monitor table, and History table) require the use of unique row numbers. Each table contains an *Index* column which acts as a *key field* to distinguish rows in the table. The following sections describe the benefits of reserving a block of rows (somewhere in the range of 11 to the maximum number of rows in your table) for use by the system or application administrator.

Setting Local Policy

System administrators may choose, as a matter of local policy, to reserve a block of rows for system administration. In compliance with this policy, all other users should use row indices that are outside of the reserved range when defining user-configured entries. This policy prevents users from using rows reserved for system administration.

Reserving Blocks of Rows

This policy also allows system administrators to define a consistent set of conditions (row entries) to be monitored across all machines such that the same condition is defined in the same row number on each machine. For example, you can use row 3000 in each table to define entries monitoring the total number of page faults (sqlmodFootprintGeneralPageFaults) MIB object. You can then distribute this configuration to every host so that every machine running Microsoft SQL Server uses row 3000 for monitoring the total number of page faults, whether it is the Threshold Monitoring table or the History table.

To reserve a block of rows for monitoring Microsoft SQL Server:

1. Identify a block of rows that you want to reserve for use with monitoring Microsoft SQL Server.
2. Use that block of rows to define a set of row entries for each SystemEDGE self-monitoring tables. For more information, refer to the chapter on self-monitoring in the *SystemEDGE Agent User Guide*.

3. Distribute the configuration file entries to all hosts that are running Microsoft SQL Server and AdvantEDGE for Microsoft SQL Server. For more information, refer to the *Automating the Deployment of SystemEDGE and AdvantEDGE Point Plugin Modules* white paper.

NOTE

Alternatively, you can use this row-number assignment policy with AdvantEDGE View for group configuration operations.

4. Require end users to avoid your block of rows when defining their own self-monitoring table entries.

Using the SystemEDGE Self-Monitoring Features

This section provides examples of how to use SystemEDGE process-monitoring, threshold-monitoring, and history-collection features to monitor the Microsoft SQL Server application. Add these examples to the sysedge.cf file to enable monitoring of the MIB objects they specify. Modify these examples as necessary to monitor the MIB objects that are relevant for your configuration.

The examples in this section use row numbers in the 5000 range; use a row number for your configuration that conforms to local policies. For more information on row assignment, refer to "Assigning Entry Rows for the SystemEDGE Self-Monitoring Tables" on page 4-2. For more information on SystemEDGE process monitoring and service monitoring, refer to the *SystemEDGE Agent User Guide.*

NOTE

Enter the commands in this section and throughout this chapter as one line. Do **not** use a carriage return to match the formatting shown here.

Using SystemEDGE Process Monitoring

This section provides an example of how to use the SystemEDGE agent to monitor the availability of a critical Microsoft SQL Server process. For more information, refer to the chapter on process monitoring in the *SystemEDGE Agent User Guide.*

19076

Monitoring the Microsoft SQL Server Process

To ensure that Microsoft SQL Server is running, enter the following command:

```
watch process procAlive 'sqlservr' 5000 0x0
 30 'SQL Server' ''
```

This entry instructs the SystemEDGE agent to monitor the run-status (or liveness) of the Microsoft SQL Server process every 30 seconds and to store the data in row 5000 of the Process Monitor table.

Using SystemEDGE Threshold Monitoring

This section provides examples of how to use SystemEDGE threshold-monitoring capabilities to monitor important Microsoft SQL Server metrics. Add the following commands to the sysedge.cf file to monitor thresholds for these MIB objects. For more information on SystemEDGE threshold-monitoring, refer to the chapter on threshold monitoring in the *SystemEDGE Agent User Guide.*

NOTE

The thresholds used in these examples may not be appropriate for your system; select thresholds that are appropriate for your environment.

Monitoring the Status of the SQL Server Process

To monitor if the SQL Server process has paused, enter the following command:

```
monitor oid sqlmodDBDescStatus 5002 0x0 60 absolute = 2 'SQL
 Server has paused'
```

This entry instructs the SystemEDGE agent to track the object, sqlmodDBDescStatus, and save the data to row 5002 of the agent's monitor table. The agent will sample the status of the SQL Server process every 60 seconds. The sample-type is `'absolute'` since the object is an integer. The operator type `'='` instructs the agent to send a trap whenever the status of the SQL Server process is equal to `'2'` which is specified in Concord's MIB for Microsoft SQL Server as indicating a paused process.

19075
3

Monitoring the Number of Blocked Users

To monitor the number of blocked users, enter the following command:

```
monitor oid sqlmodDBGenProcUsersBlkd 5004 0x0 60 absolute >
  20 'Number of Blocked Users exceeds threshold'
```

Monitoring the Percentage of Log Space Used

To monitor the percentage of the transaction log that is used, enter the following command:

```
monitor oid sqlmodDatabasePerfLogutilitzed 5005 0x0 60
  absolute > 85 'Percentage of Transaction Log Utilized
  exceeds threshold'
```

Monitoring the Number of Processes Awaiting a Workspace Memory Grant

To monitor the number of processes awaiting a workspace memory grant, enter the following command:

```
monitor oid sqlmodMemoryManagerMemgrantspend 5006 0x0 60
  absolute > 20 'Number of Processes Awaiting a Workspace
  Memory Grant exceeds threshold'
```

Using SystemEDGE History Collection

This section provides examples of how to use SystemEDGE history collection to track the value of important Microsoft SQL Server metrics over time. Add the following commands to the sysedge.cf file to collect history for these MIB objects. For more information on SystemEDGE history capabilities, refer to the chapter on history collection in the *SystemEDGE Agent User Guide*.

NOTE

The number of samples and the interval between samples used in these examples may not be appropriate for your system; select values that are appropriate for your environment.

19074

Collecting Number of Blocked Users History

To collect history for the number of blocked users, enter the following command:

```
emphistory 5010 60 sqlmodDBGenProcUsersBlkd 400 'Num Blocked
  Users History'
```

This entry instructs the SystemEDGE agent to track the value of the object, sqlmodDBGenProcUsersBlkd, and save the data in row 5010 of the empireHistoryCtrlTable. The agent will sample the value every 60 seconds and store the last 400 samples.

Collecting Lock Requests Requiring a Wait History

To collect history for the number of lock requests requiring a wait, enter the following command:

```
emphistory 5011 60 sqlmodLocksLockwaits 480 'Number of Lock
  Requests Requiring a Wait History'
```

Collecting Average Length of Disk Queue History

To collect history for the average length of the disk queue, enter the following command:

```
emphistory 5012 60 sqlmodPhysicalDiskAvgdskquelen 480
  'Average Length of Disk Queue History'
```

Collecting Number of Database Transactions History

To collect history for the number of database transactions, enter the following command:

```
emphistory 5013 60 sqlDatabasePerfTrans 480 'Number of
  Database Transactions History'
```

Collecting Total Amount of Dynamic Memory Used History

To collect history for the total amount of dynamic memory used, enter the following command:

```
emphistory 5014 60 sqlMemoryManagerTotservmem 480 'Total
  Amount of Dynamic Memory Used History'
```

19073
ε

Index

A

access methods group 3-9
advanced configuration group 3-4
AdvantEDGE for Microsoft SQL Server
    compatibility with network management systems 4-1
    files 2-3
    installation requirements 2-1
    overview 1-1
    using 4-1
    using with *e*Health 1-3
    using with Live Health 1-3
assigning entry rows for self-monitoring tables 4-2

B

buffer management group 3-13

C

cache management group 3-14
collecting history information 4-5
configuration section of MIB
    advanced configuration group 3-4
    database description group 3-2
    database devices group 3-6
    database options group 3-5
    databases information group 3-4
    SQL Server general information group 3-3

D

database description group 3-2
database devices group 3-6
database options group 3-5
databases information group 3-4

E

editing
    sqlmod.cf file 2-4
    sysedge.cf file 2-3
*e*Health 1-3
evaluation copy, licensing 2-6
examples file 2-3, 4-1

F

files installed 2-3
flat information group 3-15
flat transaction performance group 3-16
footprint section of MIB
    general footprint group 3-7
    physical disk group 3-7
    SQL Server process footprint group 3-8

G

general footprint group 3-7
general statistics group 3-13
generating a license 2-7

19072
ε

H

- history collection 4-5
- Host Resources MIB v
- HP OpenView 1-2

I

- installation
  - files 2-3
  - requirements 2-1
- installing
  - AdvantEDGE for Microsoft SQL Server 2-1
  - SystemEDGE agent 2-1

L

- license
  - generating 2-7
  - key 2-8
  - sample 2-9
- licensing
  - AdvantEDGE for Microsoft SQL Server 2-6
  - evaluation copy 2-6
- Live Health 1-3
- locks group 3-9
- log management group 3-10

M

- Management Information Base. *See* MIB
- memory management group 3-11
- MIB
  - configuration section 3-2
  - footprint section 3-6
  - performance section 3-8
  - specification 3-1
  - using 4-1
- monitoring
  - assignment policy for tables 4-2
  - processes 4-3
  - thresholds 4-4

O

- object identifier (OID) branch 3-1
- obtaining a license 2-6, 2-7

P

- performance section of MIB 3-8
  - access methods group 3-9
  - buffer management group 3-13
  - cache management group 3-14
  - flat information group 3-15
  - flat transaction performance group 3-16
  - general statistics group 3-13
  - locks group 3-9
  - log management group 3-10
  - memory management group 3-11
  - SQL statistics group 3-12
- physical disk group 3-7
- process monitoring 4-3

R

- relnotes.txt file 2-3
- requirements for installation 2-1
- reserving rows in monitoring tables 4-2

S

- setting policy for monitoring tables 4-2
- SNMP v, 1-1, 4-1
- SQL Server
  - general information group 3-3
  - process footprint group 3-8
- SQL statistics group 3-12
- sqlmod.asn1 file 2-3
- sqlmod.cf file 2-4
- sqlmod.dll file 2-3
- sqlmod.pdf file 2-3
- supported releases v
- sysedge.cf file
  - adding commands 4-3
  - enabling plug-in 2-3

19071
3

sysedge.lic file 2-6, 2-8
SystemEDGE agent
  configuring 2-3
  described 1-1
  history collection 4-5
  threshold monitoring 4-4

T
temporary license 2-6
threshold monitoring 4-4

U
using AdvantEDGE for Microsoft SQL Server 4-1

eHealth Service Response User Guide

09-16020-005
December 2001

1069
ε

Important Notice

Network Health is a registered trademark of Concord Communications, Inc. Concord, the Concord logo, *e*Health, *e*Health Suite, Application Health, System Health, Live Health, and the 24x7 logo are trademarks of Concord Communications, Inc. Other trademarks are the property of their respective owners.

Adobe, Acrobat, and the Acrobat logo are trademarks, which may be registered in certain jurisdictions, and PostScript is a registered trademark of Adobe Systems Incorporated.
CA-OpenIngres and Ingres II are registered trademarks of Computer Associates International, Inc.
Hewlett-Packard, HP OpenView, HP, and HP-UX are trademarks or registered trademarks of Hewlett-Packard Company in the United states and other countries.
Intel is a registered trademark of Intel Corporation.
Motif is a trademark of Open Systems Foundation, Inc.
Netscape, Netscape Certificater Server, Netscape FastTrack Server, Netscape Navigator, Netscape ONE, SuiteSpot, and the Netscape N and Ship's Wheel logos are registered trademarks of Netscape Communications Corporation in the United States and other countries. Netscape® Enterprise Server is also a trademark of Netscape Communications Corporation, which may be registered in other countries.
Sun, Sun-4, SunOS, SunNet, and OpenWindows are trademarks, and Sun Microsystems, Sun Workstation, and NFS are registered trademarks of Sun Microsystems, Inc.
SPARCstation and SPARCserver are trademarks, and SPARC is a registered trademark of SPARC International, Inc.
UNIX is a registered trademark in the United States and other countries licensed exclusively through X/Open Company Ltd.
Windows and Windows NT are trademarks of Microsoft Corporation.

SMIC. Copyright 1992 SynOptics Communications, Inc. All Rights Reserved. SynOptics makes no representations about the suitability of this software for any particular purpose. The software is supplied "as is", and SynOptics makes no warranty, either express or implied, as to the use, operation, condition, or performance of the software. SynOptics retains all title and ownership in the software.

*e*Health incorporates compression code by the Info-ZIP group. There are no extra charges or costs due to the use of this code, and the original compression sources are freely available from ftp://ftp.cdrom.com/pub/infozip/ on the Internet and from the Concord Communications Web site: http://www.concord.com.

© Copyright Bigelow and Holmes 1986, 1985. Lucida is a registered trademark of Bigelow & Holmes. Sun Microsystems Inc., AT&T, and Bigelow & Holmes make no representations about the suitability of the source code for any purpose. It is provided "as is" without express or implied warranty of any kind.

General Notice: Some of the product names used herein have been used for identification purposes only and may be trademarks of their respective companies.

Proprietary Notice

The information and descriptions contained herein are the property of Concord Communications, Inc. Such information and descriptions may not be copied, disseminated, or distributed without the express written consent of Concord Communications, Inc. Concord Communications, Inc., assumes no responsibility for any inaccuracies that may appear in this document. Concord Communications, Inc., reserves the right to improve its products and change specifications at any time without notice.

U. S. Government Restricted Rights

Use, reproduction, and disclosure by the U.S. Government are subject to the restrictions set forth in FAR §52.227-19 (c) (1) and (2) and DFARS §252.227-7013 (c) (1) (ii).

Patent Information

U. S. Patent 5,615,323
Patents Pending

19068

Table of Contents

Preface 7

Audience 7
About This Guide 8
Revision Information 8
Documentation Conventions 8
Customer Support 9
Professional Services 9

Chapter 1 **Introduction** 11

Introducing *e*Health Service Response 11
Using *e*Health Service Response 12
Performance Criteria Measured by *e*Health Service Response . . . 13
Configuration Options 14
Using *e*Health Service Response with AdvantEDGE View 15

Chapter 2 **Installing *e*Health Service Response** 17

Installing *e*Health Service Response on UNIX Systems 17
Installing the Software 17
Installing *e*Health Service Response on Windows NT Systems . . . 18
Installing the Software 18

19067
ε

*e*Health Service Response Files 20
Files Installed for UNIX Systems 20
Files Installed for Windows NT Systems. 20
Files Created for UNIX and Windows NT Systems 21
Configuring *e*Health Service Response Startup. 21
Licensing *e*Health Service Response 23
Obtaining a License. 23
Generating the License through the Web-based License Form 24
Generating a License through AdvantEDGE View Event Processing 27
Generating a License through AdvantEDGE View Host Administration 29
Sample License File 30

Chapter 3 Configuring *e*Health Service Response 31

The Service Response Table 31
Sample Entry in the Service Response Table 37
Assigning Entry Rows for the Service Response Table 37
Configuring the Service Response Table 38
Dynamic Configuration During Operation 38
Initial Configuration During Start-Up 39
Sample Entries for the svcrsp.cf Configuration File 41
Measuring Web Server Response (HTTP) 41
Measuring Web Server Response by Proxy (HTTP) 42
Measuring Secure Web Server Response (HTTPS) 43
Measuring Domain Name Service Response (DNS) 44
Measuring Sendmail Response (SMTP) 45
Measuring Post Office Protocol Service Response (POP). . . . 46
Measuring Network News Service Response (NNTP) 47
Measuring File Transfer Service Response (FTP) 48
Measuring Network Reachability (PING) 49
Measuring TCP Service Connections 50
Measuring Custom Services 51
Using the svcwatch Utility 52
svcwatch Command Arguments 53
svcwatch Example 54

Removing Service Response Entries 54
Manually Removing the Entry 54
Dynamically Removing the Entry 55
Using SystemEDGE Self-Monitoring to Monitor Service Response Metrics 56
Using SystemEDGE Threshold Monitoring. 56
Using SystemEDGE History Collection 58

Glossary **59**

Index **73**

A1065
E

14064
c

Preface

This guide explains how to install, license, and configure *e*Health Service Response Release 1.2 Patchlevel 2. This release of *e*Health Service Response supports the monitoring of common Internet applications on the following operating systems:

- Solaris 2.x and later
- HP-UX 10.x and 11.x
- IRIX 6.x
- AIX 4.2 and later
- Microsoft Windows 4.0, Windows 2000, and Windows XP
- Red Hat Linux 6.0 and later

This release of *e*Health Service Response supports *e*Health SystemEDGE Release 4.0 Patchlevel 3 and later.

Audience

This guide is intended for the person who is installing and configuring *e*Health Service Response to monitor the response time and availability of critical Internet applications. To use this guide, you must have a basic familiarity with the SystemEDGE agent, the Internet applications you are monitoring, and your host's operating system environment.

About This Guide

This section describes the changes and enhancements that have been made since the last release of this guide. It also includes the documentation conventions used in this guide.

Revision Information

Since Release 1.2 Patchlevel 1, this guide has been updated to use a new documentation template and to include a glossary.

Documentation Conventions

Table 1 lists the conventions used in this document.

Table 1. Documentation Conventions (Page 1 of 2)

Convention	Description
File or Directory Name	File or directory names.
`code`	System, code, or operating system command line examples.
emphasis	Emphasis and guide titles.
`enter`	Text that you must type exactly as shown.
Name	Text that defines menus, fields in dialog boxes, or keyboard keys.
New Term	A new term, that is, one that is being introduced.
`Variable`	Variable values that you substitute.
→	A sequence of menus or menu options. For example, **File → Exit** means "Choose **Exit** from the **File** menu."

Table 1. Documentation Conventions (Page 2 of 2)

Convention	Description
NOTE	Important information, tips, or other noteworthy details.
CAUTION	Information that helps you avoid data corruption or system failures.
WARNING	Information that helps you avoid physical danger.

Customer Support

If you need any assistance with *e*Health Service Response, contact Customer Support at the following:

Phone:	(888) 832-4340
	(508) 303-4300
Fax:	(508) 303-4343
E-mail:	support@concord.com
Web site:	http://www.concord.com

Professional Services

If you need any assistance with customizing *e*Health Service Response, contact Professional Services at the following:

Phone:	(800) 851-8725 (Choose option 7)
Fax:	(508) 486-4555
E-mail:	proserv@concord.com
Web site:	http://www.concord.com

19061

1

Introduction

Introducing *e*Health Service Response

*e*Health Service Response is a plug-in to the SystemEDGE agent that monitors the response and availability of critical network services. Through the Service Response plug-in, the SystemEDGE agent performs real, active test transactions to measure response time and track availability for the following Internet applications from a user's perspective:

- Domain Name System (DNS)
- File Transfer Protocol (FTP)
- Hypertext Transfer Protocol (HTTP)
- Secure HTTP (HTTPS)
- Packet Inter-Network Groper (PING)
- Network News Transfer Protocol (NNTP)
- Post Office Protocol version 3 (POP3)
- Simple Mail Transfer Protocol (SMTP)
- Transmission Control Protocol connections (TCPConnect)

*e*Health Service Response also enables you to define your own service tests through custom scripts and programs. You can perform these measurements from any system within the enterprise network.

14060
E

Using *e*Health Service Response

*e*Health Service Response provides the flexibility you need to monitor service delivery. You can use *e*Health Service Response to monitor all critical applications on a single system or to monitor a particular service or application across a group of systems. You can also modify Service Response tests in real time so that you always get the information you need.

The SystemEDGE agent provides configuration and reporting for the Service Response module through Simple Network Management Protocol (SNMP). Its self-monitoring and autonomous management capabilities work with the data that *e*Health Service Response gathers.

You can configure *e*Health Service Response and the SystemEDGE agent to do the following:

- Monitor the response times of various Web servers and send a warning when the servers become unavailable.
- Warn you of response slowdowns or unavailable applications.
- Test site access and issue an alarm if it detects a service disruption.

19059
E

Performance Criteria Measured by *e*Health Service Response

Table 2 shows the performance criteria that *e*Health Service Response measures and reports for services.

Table 2. Performance Criteria for Each Service

Criterion	**Description**
Availability	Percentage of successful service requests.
Name Lookup Time	Time required to resolve the server name for the service to a network address. *e*Health Service response provides values for Last Sample, Mean, Minimum, Maximum, and Sample Variance.
Connection Time	Time required to connect to the server providing this service. *e*Health Service response provides values for Last Sample, Mean, Minimum, Maximum, and Sample Variance.
Transaction Time	Time required to perform the requested transaction, after the connection is established. For example, this value could be the amount of time required to download the Web page or check the mail status. *e*Health Service response provides values for Last Sample, Mean, Minimum, Maximum, and Sample Variance.
Total Response Time	Total time required for the given service to correctly respond to the request. This value is simply the sum of the other three measurements: Name Lookup Time, Connection Time and Transaction Time. *e*Health Service response provides values for Last Sample, Mean, Minimum, Maximum, and Sample Variance.

19058

Configuration Options

*e*Health Service Response allows you to configure service measurement mechanisms to best meet the needs of your network environment. Table 3 shows the options that you can specify for each measurement entry.

Table 3. Configuration Options for Each Measurement Entry

Option	**Description**
Target Host	Host that provides the service to be tested.
Timeout	Time (in seconds) to wait before canceling the service request.
Sample Interval	Interval at which to perform the sample operation. For example, set this value to 60 to instruct Service Response to query the DNS server every 60 seconds.
Samples Per Interval	Number of samples to take at each interval. For example, set this option to 3 to instruct Service Response to PING the server three times in succession at each interval.
Sample Window	Time (in seconds) over which to make the statistics calculations. Service Response calculates the statistics over a sliding window of the most samples. For example, set this value to 600 to instruct Service Response to use all samples it took during the past 10 minutes to calculate the results.

For several of the services, *e*Health Service Response requires additional configuration information to complete the sample transactions. Table 4 shows the additional information that is required for each service.

Table 4. Information Required for Each Service (Page 1 of 2)

Service	**Additional Information Required**
HTTP	Target URL [proxy host][username:user password:pass]
HTTPS	Target URL [proxy host][username:user password:pass]
SMTP	None
POP3	User name and password for valid POP user

19057
ε

Table 4. Information Required for Each Service (Page 2 of 2)

Service	Additional Information Required
DNS	Internet name to lookup
NNTP	None
FTP	User name and password for valid FTP user
PING	Packet size
TCPCONNECT	Port number
CUSTOM	Name of script to run for each sample

Using *e*Health Service Response with AdvantEDGE View

You can use *e*Health Service Response with AdvantEDGE View to run queries for monitoring the response and availability of Internet applications.

To run an AdvantEDGE View Application query for Service Response:

1. Select the target system or group from the **System** or **Group** list.
2. Select **Service Response** from the **Applications** list.
3. Click the **Applications** icon.

AdvantEDGE View runs the query for the specified application on the system or group you selected.

NOTE

If you run a query for a group of systems, AdvantEDGE View may request additional information before running the query. For more information, refer to the AdvantEDGE View Web Help.

19056
ε

AdvantEDGE View displays the response time, service availability, and configuration details for all of the services that are running on the system or group that you selected. Figure 1 shows the Mean Service Response Times portion of the AdvantEDGE View Service Response query.

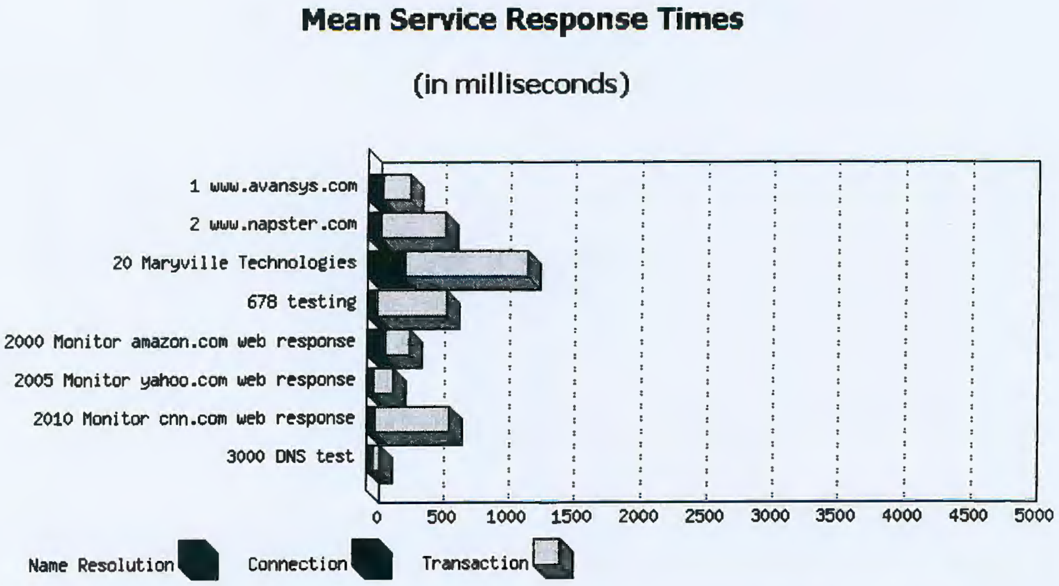

Figure 1. Sample Portion of AdvantEDGE View Service Response Query

19055
ε

Installing eHealth Service Response

Throughout this guide, the term Windows NT encompasses Windows NT 4.0, Windows 2000, and Windows XP.

This chapter describes how to install and license *e*Health Service Response for UNIX and Windows NT systems. If you are installing the software on a UNIX system, refer to “Installing eHealth Service Response on UNIX Systems.” If you are installing the software on a Windows NT system, refer to “Installing eHealth Service Response on Windows NT Systems” on page 2-18.

Before installing *e*Health Service Response, you must install, license, and configure the SystemEDGE agent Release 4.0 Patchlevel 3 or later. For more information, refer to the *eHealth SystemEDGE User Guide.*

Installing *e*Health Service Response on UNIX Systems

This section describes how to install *e*Health Service Response on UNIX systems.

Installing the Software

*e*Health Service Response for UNIX systems is distributed as a tar file named svcrsp.tar.

To install the software:

1. Log in as root by entering **su** and the root password at the command prompt.

19054
ε

2. Verify that the plugins directory exists in the SystemEDGE agent distribution area. For most systems, the recommended directory is /opt/EMPsysedge/plugins. If that directory does not exist, verify that you have SystemEDGE agent Release 4.0 Patchlevel 3 or later installed, and then create the directory manually.
3. Insert the CD containing the software distributions into the CD-ROM drive and mount it on the partition /cdrom. For mounting instructions, refer to your system documentation. For example, enter the following command for Solaris systems:

   ```
   mount -r -t hsfs /dev/sr0 /cdrom
   ```

4. Change directory to the *e*Health Service Response home directory, and load the files from the CD-ROM. For example, enter the following commands for Solaris systems:

   ```
   cd /opt/EMPsysedge/plugins
   tar xvof /cdrom/svcrsp/sol2/svcrsp.tar
   ```

   *e*Health Service Response is now installed.
5. Review the installed files. For more information, refer to "eHealth Service Response Files" on page 2-20.
6. License *e*Health Service Response. For more information, refer to "Licensing eHealth Service Response" on page 23.

Installing *e*Health Service Response on Windows NT Systems

This section describes how to install *e*Health Service Response on Windows NT systems.

Installing the Software

*e*Health Service Response for Windows NT is distributed as a self-extracting executable named svcrsp.exe.

19053

To install *e*Health Service Response:

1. Log on to the Windows NT system as **`administrator`**.
2. Select **Start → Programs → Command Prompt**.
3. Insert the CD-ROM that contains the Concord software distributions into the CD-ROM drive.

   Windows NT automatically mounts the drive using the CD-ROM drive's corresponding drive letter. The drive letter is specific to your system and depends on the number and types of disks attached to your system. Step 5 in this procedure uses *`D:`* as the CD-ROM drive. Modify that step if necessary to use the drive letter for your system's CD-ROM drive.
4. Determine which directory you want to use as the installation directory for *e*Health Service Response. If the SystemEDGE agent is installed in C:\sysedge, the recommended installation directory is C:\sysedge\plugins.
5. Run the self-extracting executable by typing the following at the command prompt:

```
D:\svcrsp\ntx86\svcrsp.exe -dir C:\sysedge\plugins
```

   where *`D:`* is the CD-ROM drive for your system, and *`C:\sysedge\plugins`* is the installation directory.

   NOTE

   The -dir option instructs the self-extracting executable to recreate the intended sub-directory hierarchy described throughout this guide. This command places the distribution in a svcrsp subdirectory within the specified target directory (for example, C:\sysedge\plugins\svcrsp).

   *e*Health Service Response is now installed.
6. Review the installed files. For more information, refer to "eHealth Service Response Files" on page 2-20.
7. License *e*Health Service Response. For more information, refer to "Licensing eHealth Service Response" on page 23.

19052
4

*e*Health Service Response Files

This section describes the files installed with *e*Health Service Response.

Files Installed for UNIX Systems

Table 5 shows the files that the *e*Health Service Response installation program installs on UNIX systems.

Table 5. Files Installed on UNIX Systems

File	Description
collector.exe	*e*Health Service Response collector program.
svcrsp.asn1	*e*Health Service Response MIB specification.
svcrsp.cf.example	Sample configuration file for *e*Health Service Response.
svcrsp.pdf	*eHealth Service Response User Guide.*
svcrsp.so	Shared library module for *e*Health Service Response.
svcwatch	Configuration utility program for *e*Health Service Response.
svcwatch.1	Manual page that explains how to use the svcwatch utility.

Files Installed for Windows NT Systems

Table 6 shows the files that the *e*Health Service Response installation program installs on Windows NT systems.

Table 6. Files Installed on Windows NT Systems (Page 1 of 2)

File	Description
collector.exe	*e*Health Service Response collector program.
svcrsp.asn1	*e*Health Service Response MIB specification.
svcrsp.dll	Dynamic link library (DLL) module for *e*Health Service Response.
svcrsp.cf.example	Sample configuration file for *e*Health Service Response.

1905-1
3

Table 6. Files Installed on Windows NT Systems (Page 2 of 2)

File	Description
svcrsp.pdf	*eHealth Service Response User Guide.*
svcwatch.exe	Configuration utility program for *e*Health Service Response.
svcwatch.txt	Text file that explains how to use the svcwatch utility.

Files Created for UNIX and Windows NT Systems

*e*Health Service Response also creates two files while it is running. It creates these files in the directory that contains the executables and configuration files. Table 7 shows the files that *e*Health Service Response creates.

Table 7. Files Created for UNIX and Windows NT Systems

File Name	Description
svcrsp.dat	Shared data file used by the collector.
svcrsp.lck	Lock file for access control.

Configuring *e*Health Service Response Startup

You must edit the sysedge.cf file to use the correct shared library file for your system and to enable the SystemEDGE agent to load *e*Health Service Response. You can use the **sysedge_plugin** keyword in the sysedge.cf configuration file to specify which plug-in modules the SystemEDGE agent will load at system initialization. By default, the SystemEDGE agent does not load any plug-ins at initialization, but you can edit the sysedge.cf file to configure the agent to load any plug-ins that you have installed.

The sysedge.cf file is located in your system directory by default; for example, it is located in the /etc/sysedge.cf directory on UNIX systems and in the C:\winnt\system32 directory on Windows NT systems. For more information about the sysedge.cf file, refer to the *eHealth SystemEDGE User Guide.*

NOTE

To configure the SystemEDGE agent to start *e*Health Service Response, you must provide the complete pathname to the shared library file for your version of *e*Health Service Response. The path depends on the location you selected when you installed *e*Health Service Response.

Table 8 shows the recommended path names for the shared library file for each operating system. Add the command shown in the right column to your sysedge.cf file to instruct the SystemEDGE agent to load *e*Health Service Response at system initialization.

Table 8. sysedge.cf Entries

Platform	Recommended sysedge.cf Entry
Solaris SPARC (32-bit)	sysedge_plugin /opt/EMPsysedge/plugins/svcrsp/svcrsp.so
Solaris SPARC (64-bit)	sysedge_plugin /opt/EMPsysedge/plugins/svcrsp/svcrsp-sparcv9.so
Windows NT, Windows 2000, or Windows XP (x86)	sysedge_plugin \sysedge\plugins\svcrsp\svcrsp.dll
HPUX 10.x and 11.x	sysedge_plugin /opt/EMPsysedge/plugins/svcrsp/svcrsp.so
Linux	sysedge_plugin /opt/EMPsysedge/plugins/svcrsp/svcrsp.so
AIX	sysedge_plugin /usr/lpp/EMPsysedge/plugins/svcrsp/svcrsp.so
IRIX	sysedge_plugin /opt/EMPsysedge/plugins/svcrsp/svcrsp.so

19049
E

Licensing *e*Health Service Response

Like the SystemEDGE agent, *e*Health Service Response utilizes a host-based license method. Copies of *e*Health Service Response can run only on systems that possess a valid license key. This license key is separate from the one used for the SystemEDGE agent.

The first time that you attempt to start the SystemEDGE agent after installing *e*Health Service Response, the agent displays a message stating that it could not find a valid license for *e*Health Service Response. It then provides you with a **public key** that is used to generate a permanent license key for your host machine.

A license key is composed of four space-separated, 8-character sequences, totaling 32 characters. The sysedge.lic file contains the license for *e*Health Service Response, as well as the SystemEDGE agent license and licenses for any *e*Health application insight modules (AIMs) that you have installed. For an example, refer to the sample license file in "Sample License File" on page 30.

Obtaining a License

To obtain a license, you can do any of the following:

- Run the Concord-supplied licenseutil.pl script.
- Run the licenseme.exe license utility.
- Use AdvantEDGE View to receive an SNMP license trap or to query and license the plug-in without a trap. For more information, refer to "Generating a License through AdvantEDGE View Event Processing" on page 27 or "Generating a License through AdvantEDGE View Host Administration" on page 29.
- Send an e-mail request to licenses@concord.com and place the returned license key in the appropriate license file.

19048
ε.

NOTE

Always include the user name in license requests that you send through e-mail.

- Complete the online license form through the Internet, as described in the next section, "Generating the License through the Web-based License Form."

For more information about licensing, refer to the *eHealth SystemEDGE User Guide*.

Generating the License through the Web-based License Form

This section describes how to generate the license using the Web-based license form.

NOTE

If you are using an evaluation copy of *e*Health Service Response, you must request a temporary license that will enable it to operate during the evaluation period.

To generate a license for *e*Health Service Response:

1. Start the SystemEDGE agent.

   Do the following for UNIX systems:

   a. Log in as **`root`**.

   b. Change directory (cd) to /opt/EMPsysedge.

   c. Enter the following:

   ```
   ./bin/sysedge -b
   ```

   Do the following for Windows NT systems:

   a. Log in as **`administrator`**.

   b. Open a command prompt window, and enter the following:

   ```
   C:\sysedge\setup -l
   ```

The SystemEDGE agent displays a message indicating that you need a license for the *e*Health Service Response module on this system. It displays a message similar to the following:

```
SystemEDGE Version 4.0 Patchlevel 3
Copyright 2001 by Concord Communications, Inc.
Please contact Concord Communications, Inc. to obtain a license
http://www.concord.com/support, Email: license@concord.com
Provide this: svcrsp pluto SunOS 5.8 8035b1f8f643ab43 1.2 Patchlevel 2
```

2. Using a Web browser, go to the licensing Web site at http://license.concord.com, and select the **Create License** option that matches your use of the module:
   - **Create SystemEDGE/AdvantEDGE Eval License** (if you are evaluating the module or are a Concord partner or reseller)
   - **Create SystemEDGE Outsource License** (if you are outsourcing the module)
   - **Create SystemEDGE/AdvantEDGE License** (if you have purchased the module)

   NOTE

   You must specify a user name and password to access the license form.

   If you do not have Web access, fill out the license request form, /config/license.txt (available as part of the *e*Health Service Response installation), with the complete string generated by the SystemEDGE agent, and e-mail the completed form to licenses@concord.com.

3. Fill out the license form, entering the information that was printed by the SystemEDGE agent. You must supply the following information:
   - Name
   - E-mail address

19046
4

- Software version number (1.2 in the example on page 25)
- Patchlevel (2 in the example on page 25)
- System name (pluto in the example on page 25)
- Operating system name (SunOS in the example on page 25)
- Operating system version (5.8 in the example on page 25)
- System identifier (`8035b1f8f643ab43` in the example on page 25)

NOTE

Select the option for *e*Health Service Response from the product list on the licensing form.

After you submit the license request form, the Concord Web server generates a license and displays it on your Web browser. It also e-mails the license to the contact person in your organization.

4. Copy the license into the sysedge.lic file (located in /etc or C:\winnt\system32), and save that file.
5. Restart the SystemEDGE agent.

   For UNIX systems, enter the following:

   ```
   ./bin/sysedge -b
   ```

   For Windows NT systems, stop and start the Windows NT Master agent by entering the following:

   ```
   C:\net stop snmp
   C:\net start snmp
   ```

*e*Health Service Response is now licensed and ready to use.

1904S
E.

Generating a License through AdvantEDGE View Event Processing

In order to use AdvantEDGE View event processing to license *e*Health Service Response, your system must meet the following requirements:

- You must be using SystemEDGE Release 4.0 Patchlevel 3 or later with AdvantEDGE View.
- You must configure the SystemEDGE agent to send SNMP traps to AdvantEDGE View. For more information, refer to the section on configuring the SystemEDGE agent in the *eHealth SystemEDGE User Guide*.
- You must configure the SystemEDGE agent with a read-write community so that AdvantEDGE View can issue an SNMP Set to transmit the license key to it. For more information, refer to the section on configuring the SystemEDGE agent in the *eHealth SystemEDGE User Guide*.
- Your AdvantEDGE View system must have access to the Internet, either directly or through a Web proxy.
- The AdvantEDGE View user who is generating the license must have either write or admin permissions.

To generate a license through AdvantEDGE View:

1. Start the SystemEDGE agent with *e*Health Service Response in unlicensed mode. SystemEDGE sends a license trap to AdvantEDGE View for that module.
2. Start AdvantEDGE View, and click the **Events** icon to display the Event Processing screen.

   AdvantEDGE View displays a license trap for the system that requires a license.
3. Click the index number for that system to view the Trap Details form for **License Software** to display the AdvantEDGE View Software Licensing form.

19044
4

4. Complete the licensing form, and click **Get License**.

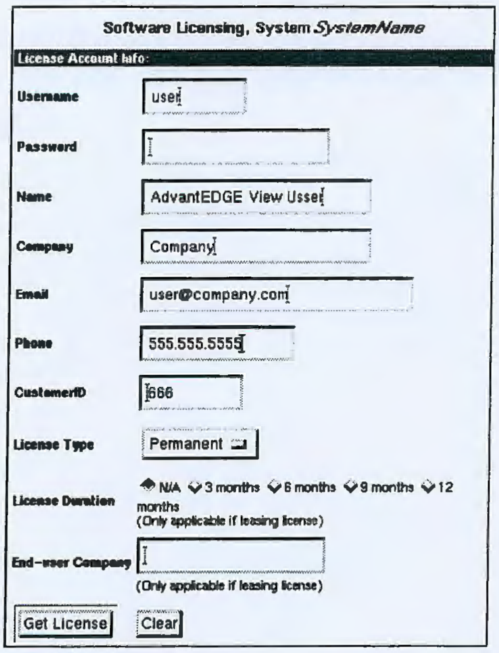

Software Licensing, System *SystemName*

License Account Info:

Username	user
Password	
Name	AdvantEDGE View Usser
Company	Company
Email	user@company.com
Phone	555.555.5555
CustomerID	666
License Type	Permanent
License Duration	◆ N/A ◇ 3 months ◇ 6 months ◇ 9 months ◇ 12 months (Only applicable if leasing license)
End-user Company	 (Only applicable if leasing license)

NOTE

If you have configured AdvantEDGE View preferences, AdvantEDGE View fills in all of the information (except password) on this form.

AdvantEDGE View contacts the Web-based license server, obtains a license for *e*Health Service Response, and issues an SNMP Set to the target SystemEDGE agent to inform it of the new software license key.

Generating a License through AdvantEDGE View Host Administration

You can also license systems through AdvantEDGE View Host Administration.

To access Host Administration:

1. Start AdvantEDGE View, and click the **Administration** icon. AdvantEDGE View displays the Administration page.
2. Click the **Host Administration** icon. AdvantEDGE View displays the host list.

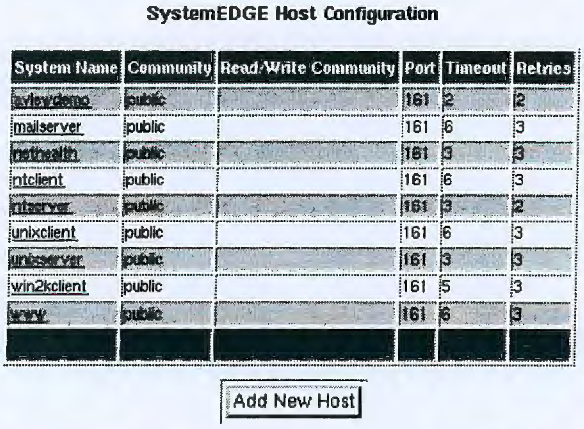

SystemEDGE Host Configuration

System Name	Community	Read/Write Community	Port	Timeout	Retries
aviewdemo	public		161	2	2
mailserver	public		161	6	3
nethealth	public		161	3	3
ntclient	public		161	6	3
ntserver	public		161	3	2
unixclient	public		161	6	3
unixserver	public		161	3	3
win2kclient	public		161	5	3
www	public		161	6	3

3. Click the name of the system that you want to license from the **System Name** column. AdvantEDGE View displays the Modify Host form.

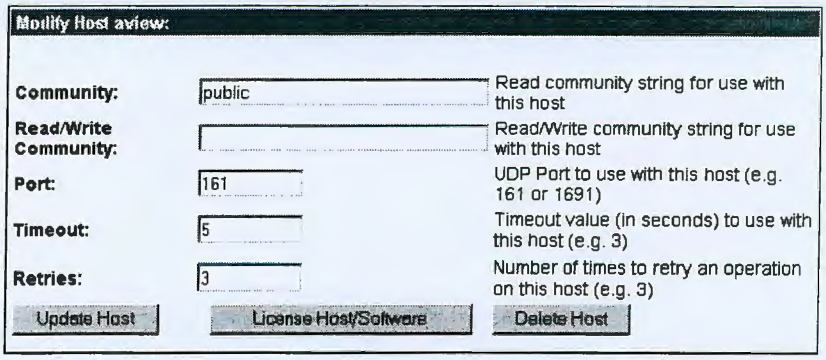

Modify Host aview:

Community:	public	Read community string for use with this host
Read/Write Community:		Read/Write community string for use with this host
Port:	161	UDP Port to use with this host (e.g. 161 or 1691)
Timeout:	5	Timeout value (in seconds) to use with this host (e.g. 3)
Retries:	3	Number of times to retry an operation on this host (e.g. 3)

4. Click **License Host/Software** to display the licensing form.
5. Select the product you want to license from the **Product** list, and then click **License Software**.

   AdvantEDGE View contacts the Web-based license server, obtains a license for the software, and issues an SNMP Set to the target SystemEDGE agent, informing it of the new software license key.

Sample License File

The following is a sample SystemEDGE agent license file. A pound character (#) in column 1 indicates that the entire line is a comment.

```
# license file for SystemEDGE Agent
# Concord Communications, Inc.
# http://www.concord.com
#
# file /etc/sysedge.lic or %SystemRoot%\system32\sysedge.lic
# A valid license key has four parts of 8 characters per part
# parts are separated by space(s) with one license key per line
# sysedge jupiter sol2 5.8 807cb1da007cb1da 4.1 PL 1
e13311d3 0F2a7cb1 abC512dc fF8C923a
# svcrsp pluto SunOS 5.8 807cb1da007cb1da 1.2 PL 2
a7943fde 098a87ij a4kiuf39 afafEkj4
```

1904

3

Configuring eHealth Service Response

This chapter explains how to configure and use *e*Health Service Response.

The Service Response Table

*e*Health Service Response is implemented as an SNMP table in the Systems Management MIB. The Service Response table provides information about each of the services that the SystemEDGE agent is currently monitoring. Each row of the table represents a single monitored service. You can specify as many service monitor entries (rows) as necessary for your Service Response implementation. For each entry, the table provides information such as the service being monitored, specific arguments (for example, the URL to query), how often the agent checks the service, and the measurement results.

In addition to the Service Response table, *e*Health Service Response provides two additional MIB objects:

- svcRspVersion, which reports the *e*Health Service Response version information.
- svcRspPID, which reports the process identifier of the *e*Health Service Response collector process.

19040
E

All MIB objects that are related to *e*Health Service Response exist at object identifier (OID) branch *1.3.6.1.4.1.546.16.6* in the Systems Management MIB. The MIB is defined in the svcrsp.asn1 file, which is installed as part of the *e*Health Service Response installation. Table 9 shows the columns of the Service Response table.

Table 9. Service Response Table (Page 1 of 5)

MIB Object	Permissions	Description
svcRspTableIndex	Read-Only	An integer (1 through MAXINT) that indicates the row index for this entry.
svcRspTableDescr	Read-Only	A textual description of the row entry. This field is entirely for the user and is not interpreted by the software.
svcRspTableSvc	Read-Write	An integer that indicates the type of service to be sampled. The following are possible values: • NNTP(1) • DNS(2) • POP3(3) • HTTP(4) • FTP(5) • SMTP(6) • PING(7) • TCPCONNECT(8) • CUSTOM(9) • HTTPS(10) Additional values will be defined in the future as new services are implemented.

19039

Table 9. Service Response Table (Page 2 of 5)

MIB Object	Permissions	Description
svcRspTableArgs	Read-Write	A quoted string (0 through 128 characters) that specifies the service-specific arguments that the module uses for measuring purposes. The following are examples of service arguments: • DNS - dns-server hostname • HTTP - Target URL [proxy host][username:user password:pass] • HTTPS - Target URL [proxy host][username:user password:pass] • FTP - ftp-server username passwd • POP3 - pop3-server username passwd • NNTP - nntp-server • SMTP - smtp-server • PING - target-host [size] • TCPCONNECT - target-host port-number • CUSTOM - script-name
svcRspTableInterval	Read-Write	An integer value (30 through MAXINT) that indicates how often (in seconds) the agent should measure the service response. For example, the value 30 instructs the agent to sample theservice every 30 seconds. This value *must* be a multiple of 30 seconds
svcRspTableSamplesPerInterval	Read-Write	An integer value (1 through MAXINT) that indicates the number of times that the agent should perform the sample query at each interval. For example, you can specify 3 to perform a PING measurement three times each interval.
svcRspTableTimeout	Read-Write	An integer value (1 through MAXINT) that indicates the time (in seconds) that this measurement should wait for a response. A sample that does not return within the timeout value is recorded as "unavailable" for the purposes of the availability measurement.

19038
E

Table 9. Service Response Table (Page 3 of 5)

MIB Object	Permissions	Description
svcRspTableStatsWindow	Read-Write	An integer value (1 through MAXINT) that indicates the time in seconds to be used in the statistical calculations. For instance, a value of 1800 specifies that the agent will calculate all statistical results (for example, mean and availability) over the last 30 minutes.
svcRspTableStatus	Read-Write	Row status; one of the following values: • active • notInService • notReady • createAndGo • createAndWait These values are identical in meaning to the SNMPv2 SMI RowStatus textual convention. Normally, a row is either active or notInService.
svcRspTableLastUpdate	Read-Only	Time (based on sysUpTime) at which the agent last sampled this service. A value of 0 indicates that this service has not yet been sampled.
svcRspTableNumSamples	Read-Only	Total number of samples that the agent has taken for this response time entry since the row was initialized.
svcRspTableTotalLastSample	Read-Only	Last recorded total response time (in milliseconds) for this service. A value of 0 indicates that the last sample failed to respond within the given timeout.
svcRspTableTotalMin	Read-Only	Smallest successful total response time (in milliseconds) for this service during the current measurement window.
svcRspTableTotalMax	Read-Only	Largest successful total response time (in milliseconds) for this service during the current measurement window.

19037

Table 9. Service Response Table (Page 4 of 5)

MIB Object	Permissions	Description
svcRspTableTotalMean	Read-Only	Sample mean of the successful total response times (in milliseconds) for this service during the current measurement window.
svcRspTableTotalVariance	Read-Only	Sample variance of the successful total response times (in milliseconds) for this service during the current measurement window.
svcRspTableTotalAvailability	Read-Only	Percentage of the total response measurement attempts that were successful during the current measurement window.
svcRspTableNameLastSample	Read-Only	Last recorded name lookup time (in milliseconds) for this service. A value of 0 indicates that the last sample failed to respond within the given timeout.
svcRspTableNameMin	Read-Only	Smallest successful name lookup time (in milliseconds) for this service during the current measurement window.
svcRspTableNameMax	Read-Only	Largest successful name lookup time (in milliseconds) for this service during the current measurement window.
svcRspTableNameMean	Read-Only	Sample mean of the successful name lookup times (in milliseconds) for this service during the current measurement window.
svcRspTableNameVariance	Read-Only	Sample variance of the successful name lookup times (in milliseconds) for this service during the current measurement window.
svcRspTableConnLastSample	Read-Only	Last recorded connection time (in milliseconds) for this service. A value of 0 indicates that the last sample failed to respond within the given timeout.
svcRspTableConnMin	Read-Only	Smallest successful connection time (in milliseconds) for this service during the current measurement window.

11903 6 e

Table 9. Service Response Table (Page 5 of 5)

MIB Object	Permissions	Description
svcRspTableConnMax	Read-Only	Largest successful connection time (in milliseconds) for this service during the current measurement window.
svcRspTableConnMean	Read-Only	Sample mean of the successful connection times (in milliseconds) for this service during the current measurement window.
svcRspTableConnVariance	Read-Only	Sample variance of the successful connection times (in milliseconds) for this service during the current measurement window.
svcRspTableTranLastSample	Read-Only	Last recorded transaction time (in milliseconds) for this service. A value of 0 indicates that the last sample failed to respond within the given timeout.
svcRspTableTranMin	Read-Only	Smallest successful transaction time (in milliseconds) for this service during the current measurement window.
svcRspTableTranMax	Read-Only	Largest successful transaction time (in milliseconds) for this service during the current measurement window.
svcRspTableTranMean	Read-Only	Sample mean of the successful transaction times (in milliseconds) for this service during the current measurement window.
svcRspTableTranVariance	Read-Only	Sample variance of the successful transaction times (in milliseconds) for this service during the current measurement window.

19035
E.

Sample Entry in the Service Response Table

The following shows a sample entry in the Service Response table for monitoring the HTTP service.

Index	Service	Arguments	Interval	SamplesPerInterval	Timeout	StatsWindow	Status
10	HTTP(4)	"http://www.empire.com"	60	1	10	3600	ACTIVE(1)

The entry is the 10th row in the table, and its purpose is to monitor the HTTP service by retrieving the Web page that is located at http://www.empire.com. The agent performs this query once every 60 seconds with a timeout of 10 seconds. The value of 3600 in the statistics window column indicates that the agent uses only the last 3600 seconds (or 1 hour) of samples to calculate statistical results. The current status of this row is active.

Assigning Entry Rows for the Service Response Table

The svcRspTableIndex column is the row index of the Service Response table, and it acts as a key field to distinguish rows in the table. You may choose, as a matter of local policy, to reserve a block of rows to be used solely for system administration. By reserving a block of rows, you can define a consistent set of conditions (row entries) to be monitored across all systems such that the same condition is defined in the same row number on each of the systems. For example, you might use row 11 (svcRspTableIndex = 11) to define an entry for monitoring the DNS service throughout the enterprise. You can then distribute this configuration to every system so that they all use row 11 to measure the DNS service.

To reserve a block of rows:

1. Decide on a block of rows that you want to reserve for your use in the Service Response table.

119034
E

2. Define a set of row entries (services to be measured) in the svcrsp.cf configuration file (in the block of rows you decided to reserve). For more information, refer to the next section, "Configuring the Service Response Table."
3. Distribute the svcrsp.cf configuration file to all systems on which *e*Health Service Response is installed.
4. Require end users to avoid your block of rows when defining their own Service Response table entries.

Configuring the Service Response Table

You can control the services that *e*Health Service Response monitors by adding, deleting, or modifying entries in the Service Response table.

You can configure the Service Response table in one of the following ways:

- **Dynamically**. Use SNMP commands from a management station to modify the table. For more information, refer to the next section, "Dynamic Configuration During Operation." For information about how to use the SystemEDGE agent svcwatch utility to dynamically configure the agent to monitor MIB variables, refer to "Using the svcwatch Utility" on page 52.
- **At start-up initialization**. Specify the process attributes to monitor through the svcrsp.cf configuration file. For more information, refer to "Initial Configuration During Start-Up" on page 39.
- **Through AdvantEDGE View configuration**. For more information, refer to the AdvantEDGE View Web Help.

Dynamic Configuration During Operation

You can use your network management system (NMS) to issue SNMP SetRequest messages to the agent to modify the entries in the Service Response table. The agent uses the SNMPv2 SMI Row-Status textual convention for creating, deleting, and modifying rows in the table.

19033

Each time the Service Response table is successfully modified, the agent updates the svcrsp.cf file to record the changes so that when the agent is restarted, it starts up with the same Service Response table configuration as it had when it was stopped.

NOTE

Service Response table entries are saved to the svcrsp.cf configuration file so that any changes made during the operation of the agent are preserved across agent and system restarts. You can edit the svcrsp.cf file *only* when the SystemEDGE agent is not running.

Initial Configuration During Start-Up

On start-up, *e*Health Service Response reads the svcrsp.cf file. You can use this file to specify the services that you want the agent to measure. If you are configuring several systems to measure services throughout an enterprise, you can create a single svcrsp.cf file and distribute that file to all of your systems.

The Service Response configuration file consists of a series of entries that are delimited by brackets ({ }). Within each entry, fields exist on separate lines. The format for an entry is as follows:

```
{
    Index
    Description
    Service
    Arguments
    Interval
    SamplesPerInterval
    Timeout
    Window Size
    SNMP Row Status
}
```

1903 2
ε

Table 10 describes each field of an entry in the svcrsp.cf file.

Table 10. Configuration File Entries

Entry	Description
Index	Row (index) of the Service Response table for this entry. Each row in the Service Response table is uniquely identified by an index number.
Description	Quoted string of up to 128 characters that describes the entry.
Service	Name of the service to be tested. Currently one of the following: • NNTP • DNS • POP3 • HTTP • HTTPS • FTP • SMTP • PING • TCPCONNECT • CUSTOM
Arguments	Quoted string of up to 128 characters that contains the service-specific parameters.
Interval	Measurement interval in seconds. This value must be a multiple of 30.
SamplesPerInterval	Number of samples to be taken at each interval.
Timeout	Sample timeout in seconds.
Window Size	Time window in seconds to use for calculating statistical results.
Row Status	SNMP row status for this row. Values of active or notInService are recommended.

Sample Entries for the svcrsp.cf Configuration File

This section contains several examples for using *e*Health Service Response to monitor services through entries in the svcrsp.cf file.

Measuring Web Server Response (HTTP)

The following entry instructs *e*Health Service Response to monitor the amount of time required to access the main Web page at www.cnn.com.

```
{
    6
    "Test CNN Web Server"
    HTTP
    "http://www.cnn.com/"
    60
    1
    20
    300
    active
    }
```

The entry is created as row 6 in the Service Response table. The agent tests the service once every 60 seconds and waits up to 20 seconds for a successful response. The agent calculates statistics over the last 300 seconds (5 minutes). This entry is active.

Measuring Web Server Response by Proxy (HTTP)

The following entry instructs *e*Health Service Response to monitor the amount of time required to access the main Web page at www.weather.com. In this case, the testing system does not access the site directly; instead, it uses the Web proxy host myproxy that is running on port 8080.

```
{
    6
    "Test Weather Channel Server Via Proxy"
    HTTP
    "http://www.weather.com/ myproxy:8080"
    60
    1
    20
    300
    active
}
```

The entry is created as row 6 in the Service Response table. The agents tests the service once every 60 seconds and waits up to 20 seconds for a successful response. The agent calculates statistics over the last 300 seconds (5 minutes). This entry is active.

19028
4

Measuring Secure Web Server Response (HTTPS)

The following entry instructs *e*Health Service Response to monitor the amount of time required to access the main Web page at charge.mycredit.

```
  {
       12
       "Test Secure Web Server"
  "https://charge.mycredit/commit.exe username:empire password:tech"
  60
  1
  20
  300
  active
}
```

The entry is created as row 12 in the Service Response table. The agent tests the service once every 60 seconds and waits up to 20 seconds for a successful response. The agent calculates statistics over the last 300 seconds (5 minutes). This entry is active.

Measuring Domain Name Service Response (DNS)

The following entry instructs *e*Health Service Response to monitor the amount of time required to resolve the IP address for www.cnn.com using the name server at 194.13.12.92. In this case, the agent is testing the response time for lookups at host 194.13.12.92. In your environment, use the IP address of your local DNS server. The actual name that is being resolved is not the most critical parameter.

```
{
    11
    "Test DNS Lookup"
    DNS
    "194.13.12.92 www.cnn.com"
    300
    1
    10
    86400
    active
}
```

The entry is created as row 11 in the Service Response table. The agent tests the service once every 300 seconds (5 minutes) and waits up to 10 seconds for a successful response. The agent calculates statistics over the last 86,400 seconds (1 day). This entry is active.

Measuring Sendmail Response (SMTP)

The following entry instructs *e*Health Service Response to monitor the amount of time required to connect to the SMTP service on host mailserver.yourdomain and to perform a null transaction. This test is a good measure of the baseline time that is required to send a mail message.

```
{
    7
    "Test Sendmail Response"
    SMTP
    "mailserver.yourdomain"
    60
    1
    10
    300
    active
}
```

The entry is created as row 7 in the Service Response table. The agent tests the service once every 60 seconds and waits up to 10 seconds for a successful response. The agent calculates statistics over the last 300 seconds (5 minutes). This configuration works well if you are using a polling station to sample the mean and availability values for this entry. This entry is active.

1A025
E

Measuring Post Office Protocol Service Response (POP)

The following entry instructs eHealth Service Response to monitor the amount of time required to log in to and test the status of the POP mail service at host popserver.yourdomain. The username poptest and password pop123 must be a valid username-password combination for a POP user on this server. Any valid account works, and the sample query does not affect the contents of the mailbox.

```
{
    5
    "Test POPmail Response"
    POP3
    "popserver.yourdomain poptest pop123"
    300
    1
    10
    21600
    active
}
```

The entry is created as row 5 in the Service Response table. The agent tests the service once every 300 seconds (5 minutes) and waits up to 10 seconds for a successful response. The agent calculates statistics over the last 21,600 seconds (6 hours). This entry is active.

19024
E

Measuring Network News Service Response (NNTP)

The following entry instructs *e*Health Service Response to monitor the amount of time required to connect to the NNTP service at host news.yourdomain and to perform a simple transaction.

```
{
    1
    "Test Net News Response"
    NNTP
    "news.yourdomain"
    3600
    1
    10
    86400
    active
}
```

The entry is created as row 1 in the Service Response table. The agent tests the service once every 3600 seconds (1 hour) and waits up to 10 seconds for a successful response. The agent calculates statistics over the last 86,400 seconds (1 day). This entry is active.

19223

Measuring File Transfer Service Response (FTP)

The following entry instructs *e*Health Service Response to monitor the amount of time required to log in to and test the status of the FTP service at host ftpserver.yourdomain. The username ftptest and password ftp123 must be a valid username-password combination for an FTP user on this server. Any valid account works.

```
{
    2
    "Test FTP Service Response"
    FTP
    "ftpserver.yourdomain ftptest ftp123"
    3600
    1
    10
    604800
    active
}
```

The entry is created as row 2 in the Service Response table. The agent tests the service once every 3600 seconds (1 hour) and waits up to 10 seconds for a successful response. The agent calculates statistics over the last 604,800 seconds (1 week). This entry is active.

Measuring Network Reachability (PING)

The following entry instructs *e*Health Service Response to monitor the amount of time required to perform a network-level ping of host server.yourdomain. This test is an excellent way to determine whether the system is up and network connectivity exists.

```
{
    3
    "Test PING Response"
    PING
    "server.yourdomain"
    60
    3
    5
    86400
    active
}
```

The entry is created as row 3 in the Service Response table. The agent tests the service 3 times every 60 seconds (1 minute) and waits up to 5 seconds for a successful response. The agent calculates statistics over the last 86,400 seconds (1 day). This entry is active.

17021
Cu

Measuring TCP Service Connections

The following entry instructs *e*Health Service Response to monitor the amount of time required to connect to port 2049 on host server.yourdomain. This tes is an excellent way to determine whether the service is up and network connectivity exists.

```
{
    9
    "Test TCP Connection"
    TCPCONNECT
    "nfsserver.yourdomain 2049"
    60
    1
    5
    600
    active
}
```

The entry is created as row 9 in the Service Response table. The agent tests the service once every 60 seconds (1 minute) and waits up to 5 seconds for a successful response. The agent calculates statistics over the last 600 seconds (10 minutes). This entry is active.

Measuring Custom Services

The following entry causes *e*Health Service Response to execute the /local/bin/custom-response script and to use the output as the response time values for this custom service.

```
{
    99
       "Test Custom Service"
CUSTOM
    "/local/bin/custom-response"
    120
    1
    20
    3600
    active
}
```

The entry is created as row 99 in the Service Response table. The agent test the service once every 120 seconds (2 minutes) and waits up to 20 seconds for a successful response. The agent calculates statistics over the last 3600 seconds (1 hour). This entry is active.

You can create a custom script to perform any desired test or operation. You can write the script as a binary executable or in a scripting language such as UNIX shell or Perl. Custom response modules work very much like SystemEDGE agent extension objects. For more information, refer to the section on extension objects in the *eHealth SystemEDGE User Guide.*

*e*Health Service Response expects the custom script to provide a single line of output with three values followed by a line feed. It interprets the values as the name lookup time, connection time, and transaction time. The script must report all times in milliseconds.

19009
6.

Using the svcwatch Utility

svcwatch is a command-line utility that automatically configures *e*Health Service Response to monitor the service that you specify. You identify the service, arguments, measurement interval, timeout, and statistics window, and the svcwatch utility issues an SNMP SetRequest to create the appropriate entry in the target Service Response table.

Use svcwatch as follows:

```
svcwatch hostname[:port][,timeout] community command
```

Table 11 describes the svcwatch arguments.

Table 11. svcwatch Arguments

Argument	Description
hostname[:port][,timeout]	Specifies the hostname or IP address (in dotted quad notation) of the system that contains the agent and MIB object to be monitored. If the agent is running on an alternative UDP port (for example, 1691), specify that port number along with the hostname/address with a colon-separator. In addition, you can specify an optional SNMP timeout value (in seconds) using a command-separator.
community	Specifies the community string that svcwatch uses in its SNMP requests to the agent. Because svcwatch uses SNMP SetRequests, the community string must provide read-write access to the target agent.
command	Specifies the command and associated arguments. Supported commands include the following: oid – for monitoring an object filesystem – for monitoring a file system list – for listing the current entries setstatus – for setting the status of an entry delete – for deleting an entry For more information about these commands, refer to the next section, "svcwatch Command Arguments."

svcwatch Command Arguments

These are the svcwatch commands and associated arguments:

- add *index "descr" service "arguments" interval samples timeout window*
- setstatus *index status*
- delete *index*
- list
- version

Table 12 describes the svcwatch arguments that are associated with the svcwatch commands.

Table 12. svcwatch Arguments Associated with Commands (Page 1 of 2)

Argument	Description
index	Specifies the row (index) of the Service Response table for this entry.
"descr"	Describes the rows in a quoted string of up to 128 characters.
service	Specifies the service to monitor. One of the following: • HTTP • HTTPS • FTP • NNTP • DNS • SMTP • POP3 • PING • TCPCONNECT • CUSTOM
"arguments"	Specifies the service-specific arguments in a quoted string of up to 128 characters.
interval	Specifies an integer value (30 to MAXINT) that indicates how often (in seconds) the service should be performed tested. This value *must* be a multiple of 30 seconds.

19007

Table 12. svcwatch Arguments Associated with Commands (Page 2 of 2)

Argument	Description
samples	Specifies an integer value (1 to MAXINT) that indicates how many times the agent should monitor the service at each interval.
status	Row status; one of the following: • active – Activates a table row. • notInService – Deactivates but preserves a row. • destroy – Deletes a row.
timeout	Specifies the time in seconds to wait for the service (in an integer value).
window	Specifies the time window of samples to include in statistical calculations.

svcwatch Example

Enter the following to create an entry, at index 11 in the Service Response table, that tests the network reachability to the pingtarget system:

```
svcwatch 143.45.0.12 private add 11 "Test PING" PING
"pingtarget" 120 1 10 3600
```

Removing Service Response Entries

To stop the sampling of a particular service, you must remove the appropriate entry from the Service Response table. There are two options for removing these table entries:

- Manually removing the entry from the svcrsp.cf file
- Dynamically removing the entry with the svcwatch utility

Manually Removing an Entry

You can remove an entry from the Service Response table by removing the entry from the svcrsp.cf configuration file.

19006

NOTE

Before you edit the svcrsp.cf file, you must stop the SystemEDGE agent.

To remove an entry from the svcrsp.cf file:

1. Stop the SystemEDGE agent. For more information, refer to the *eHealth SystemEDGE User Guide.*
2. Edit the file svcrsp.cf. Locate and remove the entry you want to delete. Remove the entire entry, including the bracket characters ({ and }).
3. Save the svcrsp.cf file.
4. Restart the SystemEDGE agent.

Dynamically Removing an Entry

To dynamically remove an entry from the Service Response table, use the svcwatch utility. The following example deletes row 14 from the Service Response table on the 143.45.0.12 system. Enter the following to remove that row from memory and from the svcrsp.cf file.

```
svcwatch 143.45.0.12 private delete 14
```

In some cases, it may not be possible to use the svcwatch utility to delete entries. For example, if you have configured the SystemEDGE agent to prevent SNMP SET operations, the svcwatch utility does not work. In this situation, you need to remove the Service Response entry from the table manually. For more information, refer to “Manually Removing an Entry” on page 54.

19005
ε

Using SystemEDGE Self-Monitoring to Monitor Service Response Metrics

This section describes how to use SystemEDGE threshold monitoring and history collection to monitor service-response metrics. In addition to using the svcrsp.cf file to add monitoring entries to the Service Response table, you can add entries directly to the SystemEDGE agent configuration file, sysedge.cf.

Each of the following examples presents a row number in the 5000 range; select a row number for your configuration that conforms to local policies. The metrics used in these examples were chosen for illustrative purposes only; you may choose to measure other metrics. You may also choose to use other thresholds, numbers of samples, and intervals between samples. Use values that make sense for your environment.

NOTE

Enter the commands throughout this section as one line. Do *not* use a carriage return to match the formatting shown here.

Using SystemEDGE Threshold Monitoring

This section provides examples for using SystemEDGE threshold monitoring to monitor Service Response metrics. Enter the following examples into the sysedge.cf file to instruct the agent to perform the monitoring they describe. For more information about SystemEDGE threshold monitoring, refer to the *eHealth SystemEDGE User Guide.*

There are nearly 30 useful values recorded for each Service Response entry. The most common values are the Mean Response Time (svcRspTableTotalMean) and Mean Availability (svcRspTableTotalAvailability). Other values also provide interesting real-time monitoring solutions. For example, you can monitor the variance (svcRspTableTotalVariance) to watch for periods of large variation in response.

19004
4

Sending a Trap when a Service Fails to Respond

To configure the agent to send a trap when a service fails to respond, you must monitor the svcRspTableTotalLastSample MIB variable. This value records the last sampled response time (in milliseconds) for this service entry. If the last test failed, the value is zero.

If, for example, you have created a Web server response-monitoring entry at row index 100 of the Service Response table, and this entry tests the server every 60 seconds, set up a SystemEDGE self-monitoring entry to watch the samples for that row and send a trap if the value is zero. To do so, enter the following in sysedge.cf:

```
monitor oid svcRspTableTotalAvailability.100 5001 0x0 60
absolute = 0 ’Web Server Down’ ’’
```

Sending a Trap when a Response Sample is Greater than 7000

To send a trap if any response sample is greater than 7000, enter the following in sysedge.cf:

```
monitor oid svcRspTableTotalLastSample.100 5002 0x0 60 absolute
> 7000 ’Web Server Too Slow’ ’’
```

Sending a Trap when the Mean Response Time is Greater than 5000

To send a trap when the *mean* response time is greater than 5000, enter the following in sysedge.cf:

```
monitor oid svcRspTableTotalMean.100 10 0x0 60 absolute > 5000
’Web Server Too Slow On Average’ ’’
```

19003
&

Sending a Trap when a New Maximum Value Appears

To send a trap whenever a new maximum value appears, regardless of the value, enter the following in sysedge.cf:

```
monitor oid svcRspTableTotalMax.100 10 0x0 60 delta > 0 'New
Maximum Web Server Response' ''
```

Using SystemEDGE History Collection

This section outlines the use of SystemEDGE history collection to track the value of important Service Response metrics over time. For more information about SystemEDGE history collection, refer to the *eHealth SystemEDGE User Guide*.

Collecting History on Mean Response Time

To collect historical data on the mean service response time, you must gather history on the svcRspTableTotalMean MIB variable. This value records the mean over the sample window for this service entry.

For example, if you have created a Web server response-monitoring entry at row index 100 of the Service Response table, and this entry tests the server every 60 seconds, you can set up a SystemEDGE History table entry to record the samples for that row. To do so, enter the following command in sysedge.cf:

```
emphistory 5002 60 svcRspTableTotalMean.100 400 'Web Response
History'
```

Collecting History on Connect Time

Enter this command to collect history on the connect time:

```
emphistory 5002 60 svcRspTableConnMean.100 400 'Web Connection
History'
```

19002
ε

Glossary

Abstract Notation One (ASN.1) A language that describes data types independent of computer structures and representations. For more information refer to ISO International Standard 8824.

access list A list of devices or IP addresses that can use a router, device, or application for particular services.

AdvantEDGE View A Web-based management interface for use with the SystemEDGE agent that enables an administrator to use a Web browser to manage systems and applications.

agent In network management, a program that provides information from a management information base (MIB) for SNMP agents. *e*Health or a network management system (NMS) use the information about managed devices and take corrective action when appropriate.

American Standard Code for Information Interchange (ASCII) The most common format for character representation in computers and the Internet. Characters fit into a single byte. It was developed by the American National Standards Institute (ANSI).

19001

application A program that performs a specific function for one or more users or for another application program. Types of applications include communication programs, management programs, word processors, databases, and drawing programs.

ASCII *See* American Standard Code for Information Interchange (ASCII).

ASN.1 *See* Abstract Notation One (ASN.1).

availability The percentage of time that an element is operational during the report period.

bandwidth The throughput of a communications line usually measured in megabits per second (Mbps). Also refers to the difference between the highest and lowest frequencies in a communications channel, expressed in units of hertz (Hz), or cycles per second.

baseline A level of performance that is considered normal, average, or typical over a period of time such as a day, week, or month. Compare current performance metrics against baseline data to identify trends in performance levels and service delivery.

buffer A temporary storage area for data. Often implemented as holding areas between the backplane and an interface; data remains in the buffer until it can be transmitted on the interface or processed by the central processing unit (CPU).

capacity A measurement of the volume that an element can support. For interfaces, this is the bandwidth that can be carried. For hard disks, this is the disk size or the amount of information that can be stored on the disks.

central processing unit (CPU) The component within a device that performs the instruction execution for the applications and programs that run on the device. Also referred to as a processor or microprocessor.

19000
ε

client A computer system, usually a desktop computer or laptop, that presents data directly to a user and accepts input. They drive the computing process, supporting local processing and accessing remote servers as needed for data access and analysis.

Also refers to the application software residing on a machine that is used by an end user.

client process The client-side part of a distributed application.

connect time The total time that a user is connected to a network.

CPU *See* central processing unit (CPU).

delay The time required for a packet or frame to travel from the sending station (source) to the receiving station (destination).

DHCP *See* Dynamic Host Configuration Protocol

disk thrashing A condition that results when a server performs high disk input /output (I/O) operations—reads and writes to the disk—without producing actual work. Often occurs when a server performs excessive paging and swapping due to physical memory limitations.

DNS *See* domain name system (DNS).

domain name system (DNS) The system that locates and translates Internet domain names such as concord.com into Internet Protocol (IP) addresses. A DNS server is typically a device that translates domain names to IP addresses within your network.

Dynamic Host Configuration Protocol A protocol that enables dynamic allocation of IP addresses so that they can be reused.

***e*Health AIM** *See* eHealth application insight module.

18999
E

eHealth application insight module A plug-in (supplementary program) that extends the functionality of the SystemEDGE agent. AIMs add the capability to manage application-specific events, processes, thresholds, and health.

event An occurrence on a system that typically results in a message, such as an SNMP trap, being sent to a configured management system. Common events include system failures, system reboots, exceeded thresholds, or any user-configurable situation that the user wants to identify.

fault tolerance A mechanism that protects networks and devices against downtime due to system failure. Fault tolerant solutions typically rely on redundancy in hardware and mirroring of applications and data.

file cache A block of memory that holds frequently or recently used data. A system can read those blocks at memory speed rather than the slower disk access speed.

File Transfer Protocol (FTP) A means for uploading and downloading files on the Internet (the oldest Internet protocol for retrieving files). You can use an FTP client application to request files from or transfer files to an FTP server.

filter A set of selection criteria used to focus a report on the desired data.

FTP *See* File Transfer Protocol (FTP).

Gbps An acronym representing gigabits per second, a common measurement of data transfer rates. One Gbps is equivalent to 10^9 bits per second.

group A collection of monitored elements. Typically, groups are used to organize elements by geographic location, department, market segment, vendor, or customer. Users can enter localized text for group names.

group list A set of one or more groups. Users can enter localized text for group list names.

18998

Host Resources MIB A MIB (management information base) that defines a set of objects that are useful for the management of host computers. For example, it defines host storage areas, devices, and file systems. This MIB is defined in RFC 1514.

hostname The name for an individual IP (Internet Protocol) address on a computer. While many computers have only one hostname, some machines, such as network servers have multiple hostnames.

HTML *See* Hypertext Markup Language (HTML).

HTTP *See* Hypertext Transfer Protocol (HTTP).

Hypertext Markup Language (HTML) A programmatic language used for controlling the way that text and images appear when a file is displayed on the World Wide Web.

Hypertext Transfer Protocol (HTTP) An application protocol that defines the set of rules for exchanging files (text, graphics, multimedia, and other files) on the World Wide Web.

Hertz (Hz) A unit of frequency of one cycle per second that measures the change in the state of an alternating current, sound wave, or other cyclical wave form.

I/O *See* input/output (I/O).

ICMP *See* Internet Control Message Protocol (ICMP).

Information Technology (IT) A widely-used term to describe all of the technologies used for creating, exchanging, managing, and using information in various forms.

input/output (I/O) Any operation, program, or device that transfers data to or from a computer.

Integrated Services Digital Network (ISDN) A high-speed carrier service offered by telecommunications companies.

Internet Control Message Protocol (ICMP) A protocol between a server and a gateway to the Internet.

18997
ε.

internet infrastructure The applications, systems, and networks that a company uses to run its business, for both internal use and for interfaces to the outside world.

Internet Protocol (IP) The method (or protocol) by which packets of information are sent across the Internet. IP defines addressing, error handling, routing, and option codes for data transmission. IP requires no continuing connection between the endpoints that are communicating.

Internet Service Provider (ISP) A company that provides individuals and companies with access to the Internet. ISPs also provide related services such as Web site building and virtual hosting.

IP *See* Internet Protocol (IP).

ISDN *See* Integrated Services Digital Network (ISDN).

ISP *See* Internet Service Provider.

IT *See* Information Technology (IT).

LAN *See* local area network (LAN).

latency A measure of delay, often network delay. Depending on the type of element, *e*Health reports can show two types of latency: round-trip latency, which is the length of time in milliseconds for a ping packet to travel from the *e*Health system to a polled element and back. Alternate latency, which is the length of time in milliseconds for a ping packet to travel from a network resource (the alternate latency source) such as a router to other critical network resources such as routers and servers (the alternate latency partner).

local area network (LAN) A shared communication medium that connects computers and devices over a limited area. The area limitations of a LAN usually result from the electrical signal limits of the medium.

18996
ε

management information base (MIB) A formal description of a set of network objects that can be managed using Simple Network Management Protocol (SNMP).

MB Megabytes.

Mbits Megabits.

MBps An acronym representing megabytes per second.

Mbps An acronym representing megabits per second, a common measurement of data transfer rates.

MIB *See* management information base (MIB).

MIB Translation File (MTF) A file that normalizes the data collected from standard and proprietary SNMP (Simple Network Management Protocol) agents. *e*Health uses an MTF to translate MIB variables into its own variables. Each MTF consists of the associated MIB and its filename, an agent for the element type, and a set of statements that map MIB variables to the appropriate *e*Health database column.

mirroring A process by which data is duplicated on separate disk systems to provide faster access and fault tolerance in the event of a disk failure.

MTF *See* MIB Translation File (MTF).

network A collection of computers, printers, routers, switches, and other devices that are able to communicate using a common transmission media such as TCP/IP.

network management system (NMS) An application program usually residing on a computer that manages at least part of a network, including systems and applications. The NMS communicates with agents to monitor network statistics and resources, control network device configuration, and analyze network problems. *See also* agent.

18995

Network News Transfer Protocol (NNTP) The predominant protocol used by computers for managing messages posted on Usenet newsgroups.

network operations center (NOC) The place where network administrators manage a telecommunications network or networks. It usually contains visualizations of the networks and workstations that are used to distribute software, troubleshoot problems, and monitor performance.

network time The time spent establishing network connections to complete a transaction.

NMS *See* network management system (NMS).

NNTP *See* Network News Transfer Protocol (NNTP).

object identifier (OID) a unique identifier of a managed object in a MIB hierarchy. *See also* management information base (MIB).

OID *See* object identifier (OID).

operating system (OS) The program that manages all other programs (applications or application programs) on a computer. Provides the following services: determining the order in which each application runs and the time allotted for that application, managing the sharing of internal memory among multiple applications and handling input to and output from attached hardware devices.

operational support system (OSS) A network management system (NMS) with a specific focus such as provisioning services or alarm surveillance.

OS *See* operating system (OS).

OSS *See* operational support system (OSS).

18994
ε.

packet A logical unit of data routed between an origin and a destination on the Internet or any other packet-switched network. On the Internet, the Transmission Control Protocol (TCP) layer of TCP/IP divides a file into packets of manageable size for routing.

page In computers that utilize virtual memory, a unit of data storage. Systems transfer pages of data from disk storage to memory and back again.

On the World Wide Web, a file written using Hypertext Markup Language (HTML) that specifies how text, images, and other multimedia will be presented to the user. A Web site delivers information to the user one page at a time.

partition A logical division of a hard disk on a PC that is created so that each partition can have a different operating system or can be used for different purposes (for example, file management or multiple users).

path In networking, a path is a route from one location to another in a network.

PC *See* personal computer (PC).

personal computer (PC) A computer designed for individual use. Prior to the PC, computers were designed to be used by many individuals and system resources were shared by all. A PC often refers to a computer with an Intel microprocessor architecture and an operating system such as Microsoft DOS or Windows.

ping An Internet echo message used to confirm the reachability of a network device. An abbreviation for Packet Internet or Inter-Network Groper.

port The physical (hardware) connection on a device that connects the device to a network.

process Typically, an instance of a program or application that is running on a server. Applications can have one or more associated processes.

18993
2

process set A collection of one or more processes that relate to a specific application. Using *e*Health – System At-a-Glance reports, you can obtain information about the impact and performance of process sets running on systems that have Concord SystemEDGE agents.

protocol The set of rules by which the endpoints in a telecommunication connection communicate. The protocol defines the packet format of the transmitted information. On the Internet, common protocols are TCP, IP, HTTP and FTP.

queue In a system, a set of jobs awaiting resources. In a network device such as a router, a collection of packets waiting to be processed or forwarded. Insufficient central processing unit (CPU) speed, memory, or interface speeds can contribute to long queues, and therefore, to delay on the network.

RAID *See* Redundant Array of Inexpensive Disks (RAID).

RAS *See* remote access server (RAS).

real time A level of computer responsiveness that an end user would deem as immediate or fast enough to show incremental changes of an external process (for example, to present visualizations of the weather as it constantly changes).

Redundant Array of Inexpensive Disks (RAID) A technology that merges several inexpensive disks into a single large disk to increase speed, capacity, and reliability. The RAID controller manipulates disks to share the work on file reads and writes for large files or to perform multiple simultaneous reads or writes for small files.

remote access server (RAS) A device that provides remote users with dial-up access to a network. RAS devices usually contain modem or Integrated Services Digital Network (ISDN) cards that provide the connection services.

18992

remote network monitoring (RMON) A type of device that collects nine kinds of network management information, including packets sent, bytes sent, packets dropped, statistics by host, by conversations between two sets of addresses, and certain kinds of events that have occurred. A probe is an example of an RMON device.

Request For Comments (RFC) The name of the document series regarding Internet design. Most RFCs define protocol specifications such as Telnet and FTP. RFCs are widely available online.

RFC *See* Request For Comments (RFC).

RMON *See* remote network monitoring (RMON).

RMON2 A type of device that collects network management information as specified in the latest version of the MIB (management information base) specification, RMON, version 2. For more information, refer to RFC 2021, a document widely available on the Internet.

router A device that connects networks. Routers learn the addresses of the network points that send data by reading the address information in the data frames. Hardware vendors often use the terms router and switch interchangeably.

routing The process of finding paths through a network to a destination.

server A program that provides services to other programs in the same and other computers. Also, a computer that performs file storage and application hosting as well as provides computing services to other devices and users on the network. Typically has one or more central processing units (CPUs), disks, interfaces, and storage partitions.

server process A server-side part of a distributed application.

18991
ε

server time The amount of time that a server requires to process a transaction. It is calculated by determining network time and subtracting it from remote time.

server type The kind of server process associated with a particular server request protocol.

Simple Network Management Protocol (SNMP) The network management protocol used almost exclusively in data networks. A method for monitoring and controlling network devices, as well as managing configurations, statistics collection, performance, and security.

SNMP *See* Simple Network Management Protocol (SNMP).

SNMP agent A program such as the SystemEDGE agent that conforms to a management information base (MIB) specification to collect information about managed devices and to take corrective action (using SNMP traps) when appropriate.

speed The capacity (bandwidth) of an interface in bits per second (bps).

swapping The process in which a computer moves entire programs in and out of random access memory to and from auxiliary storage (swap partition or pagefile).

SystemEDGE agent Concord's SNMP agent that autonomously monitors system configuration, status, performance, users, applications, file systems, and other critical resources.

Systems Management MIB A set of MIB (management information base) objects that extends the capabilities of the Host Resources MIB. It provides greater visibility into systems and specific information about Windows NT and UNIX systems.

TCP/IP *See* Transmission Control Protocol (TCP) and "Internet Protocol (IP).

throughput The rate of data transfer on an interface over time.

18990
ε

Transmission Control Protocol (TCP) A connection-based protocol used along with the Internet Protocol (IP) to send data in the form of message units between computers over the Internet. While IP is responsible for the actual delivery of the data, TCP is responsible for dividing data into packets at the sending system and constructing the data message from individual packets at the receiving system.

trap A message sent by an SNMP agent to a console or network management system (NMS) to indicate that a threshold has been reached or another user-defined condition has occurred. The SystemEDGE agent defines a number of traps for system and application management.

Trivial File Transfer Protocol (TFTP) An Internet utility that uses User Datagram Protocol (UDP) instead of Transmission Control Protocol (TCP) to transfer files. TFTP is simpler than FTP, but does not support user authentication and directory visibility.

UDP *See* User Datagram Protocol (UDP).

User Datagram Protocol (UDP) A communications protocol that uses Internet Protocol (IP) to send and receive data and is similar to Transmission Control Protocol (TCP), but provides fewer packet management services.

variable A performance metric for an element. A characteristic or behavior upon which *e*Health gathers data and evaluates the performance of the element. The SystemEDGE agents can also monitor local variables to reduce network polls and increase scalability.

variance A statistical term that indicates how closely most of the data points differ from the average of the data points.

WAN *See* wide area network (WAN).

Web *See* World Wide Web (WWW, Web).

wide area network (WAN) A network that interconnects multiple systems or networks over unlimited distances.

workstation A powerful computer that is equipped with a fast processor, a large amount of random access memory, and other features such as high-speed graphical rendering that make it suitable for business users such as engineers, graphic designers, and architects.

World Wide Web (WWW, Web) All of the resources on the Internet that use Hypertext Transfer Protocol (HTTP). Users of the Web access information through browser software.

18988

Index

A

Abstract Notation One 59
access lists 59
AdvantEDGE View 59
  licensing 27, 29
agents 59
American National Standards Institute 59
American Standard Code for Information Interchange 60
application query for Service Response 15
ASCII, *See* American Standard Code for Information Interchange.
ASN.1, *See* Abstract Notation One.
availability 60

B

bandwidth 60
baselines 60
buffers 60

C

capacity 60
central processing unit 60
client processes 61
clients 61
configuration examples
  measuring DNS response 44
  measuring FTP response 48
  measuring HTTP response 41
  measuring HTTPS response 43
  measuring network reachability 49
  measuring NNTP response 47
  measuring ping response 49
  measuring POP response 46
  measuring response for custom services 51
  measuring secure Web server response 43
  measuring Sendmail response 45
  measuring SMTP response 45
  measuring TCP connection 50
  measuring Web server response 41
configuring *e*Health Service Response
  enabling the plug-in 21
  Service Response table 31
  using the svcrsp.cf file 39
  using the sysedge.cf file 56
connect time 61
CPU, *See* central processing unit.
custom services, measuring 51

18987

D

delay 61
DHCP. *See* Dynamic Host Configuration Protocol.
disk thrashing 61
DNS
  definition 61
  measuring response 44
  servers, definition 61
domain name system 61
Dynamic Host Configuration Protocol 61

E

*e*Health application integration modules 62
*e*Health Service Response
  AdvantEDGE View 15
  configuring 31, 38
  configuring in sysedge.cf 21
  initial configuration 39
  installation
    UNIX 17
    Windows NT 18
  installed files 20
  licensing 23
  overview 11
evaluation copy, licensing 24
events 62

F

fault tolerance 62
file cache 62
File Transfer Protocol 62
filters 62
FTP
  definition 62
  measuring response 48

G

Gbps 62
generating a license 24

H

Host Resources MIB 63
hostnames 63
HTTP
  definition 63
  measuring response 41
HTTPS, measuring response 43

I

I/O 63
Information Technology 63
input/output 63
installation directory 19
installing *e*Health Service Response 17
  UNIX 17
  Windows 18
internet infrastructure 64
Internet Protocol 64
IP, *See* Internet Protocol.
IT, *See* Information Technology.

L

LAN, *See* local area network.
latency 64

18986

license
  generating 24
  sample 30
licensing
  AdvantEDGE View event processing 27
  AdvantEDGE View host
    administration 29
  *e*Health Service Response 23
  evaluation copy 24
local area networks 64

M

management information base 65
MIB, *See* management information base.
mirroring 65

N

network
  definition 65
  management system 65
  operations center 66
  reachability, measuring 49
  time 66
Network News Transfer Protocol 66
NMS, *See* network management system.
NNTP
  definition 66
  measuring response 47
NOC, *See* network operations center.

O

object identifier 66
obtaining a license 23
OID, *See* object identifier.

P

packet internet groper 67
ping
  definition 67
  measuring 49
POP, measuring response 46
ports 67
protocols 68

Q

queues 68

R

real time 68
removing entries from Service Response
    table 54
Request For Comments 69
RFC, *See* Request For Comments.
routers 69
routing 69

S

secure Web server response, example 43
Sendmail, measuring response 45
Service Response table
  adding entries 39
  overview 32
  removing entries 54
Simple Network Management Protocol 70
SMTP response, example 45
SNMP
  agents 70
  definition 70
  traps 71
speeds 70
svcrsp.exe 19

18985
ε

svcwatch utility 52
    example 54
swapping 70
sysedge.cf file 22
    SystemEDGE self-monitoring 56
sysedge.lic file 23
SystemEDGE agent
    configuring 22
    definition 70
    monitoring service response 56
    sysedge.cf file 56
Systems Management MIB 70

T

TCP
    definition 68
    measuring response 50
TCP/IP 64, 71
temporary license 24
throughput 70
Transmission Control Protocol 71
traps 71

U

UDP, *See* User Datagram Protocol.
User Datagram Protocol 71
using
    AdvantEDGE View 15
    svcwatch utility 52

V

variables 71
variance 71

W

WAN, *See* wide area network.
Web server response, example 41
Web, *See* World Wide Web.
wide area networks 71
World Wide Web 72
WWW, *See* World Wide Web.

COBRA TECNOLOGIA

Integradora de Soluções em TI e Serviços

ANEXO SERVIDOR PARA DETECÇÃO DE INTRUSÃO
PARTE 1

18784

COBRA Tecnologia S.A.
Estrada dos Bandeirantes 7966
CEP 22783-110 Rio de Janeiro RJ
Tel 21 2442-8800
www cobra com.br

Sistema para Detecção de Intruso	IDS

ATRIBUTO	REQUISITOS DO EDITAL	ATRIBUTOS OFERTADOS	ATRIBUTOS OFERTADOS ADICIONALMENTE	CONFIRMA ATENDIMENTO (SIM / NÃO)	PÁGINA DA DOCUMENTAÇÃO TÉCNICA
1 - Descrição	O Servidor de detecção deverá operar como IDS, sendo formado por dois módulos que deverão integrar-se ao Switch tipo 1, tomando o tráfego diretamento no painel traseiro desse Switch. Um módulo em BSB e outro em SPO	Não estao sendo ofertados modulos internos ao chassis	N/A	SIM	(Ver Proposta Técnica e Proposta Comercial)
	Será permitida a utilização de 02 equipamentos externos. Um em BSB e outro SPO. Desde que a tomada de tráfego seja realizada através de uma das portas SPAN do Switch tipo 1	Todos os exigidos	N/A	SIM	Ver Proposta Técnica e Comercial
	O ambiente do IDS deve ser composto por 02 elementos complementares, quais sejam: **Sensor** (responsável por monitorar a rede a que está conectado) de análise de cabeçalho e da área de dados de cada pacote que trafega pela rede citada	Conforme Edital	N/A	SIM	Anexo 25 - pág. 1 Anexo 26A (todas as páginas)
	O sensor deve ser capaz de monitorar o tráfego de rede TCP/IP observando no mínimo uma interface de rede com 500 Mbps em tráfego Full	Conforme Edital	N/A	SIM	Anexo 25 - pág. 2 Anexo 26F (todas as páginas)

18183

O sensor deve oferecer 01 interface de controle 10/100TX (em caso de equipamento externo ao tipo 1)	Todos os exigidos	A interface é do tipo 10/100/1000T	SIM	Anexo 25 - pág. 3
O sensor deve funcionar de forma transparente, ou seja, não deve causar impacto no desempenho do Switch, monitorando o tráfego no painel traseiro de comutação do Switch (em caso da utilização de módulos para o Switch tipo 1)	Todos os exigidos	N/A	SIM	Anexo 26A - pág. 3
O sensor deve possuir a capacidade de monitorar diversas VLANs simultâneamente utilizando recurso de captura de VLANs ACL ou SPAN	Todos os exigidos	N/A	SIM	Anexo 26A - pág. 4 Anexo 26G (todas as paginas)

18982

Software de Controle: Baseado em sistema operacional HP-UX, Solaris, AIX, Windows 2000 ou LINUX, capaz de monitorar de forma centralizada a atividade de múltiplos sensores, estejam estes localizados em segmentos de redes locais ou remotos. Deve ser possível configurar remotamente os sensores utilizando-se o software de controle. O software de controle deve registrar em uma base de dados as configurações de cada um dos sensores que lhes são subordinados. A comunicação entre os sensores e o software de controle deve ser autenticada. O software de controle deve permitir a integração com a base de dados ORACLE 8i ou superior, ou SQL Server versão 2000 ou superior.	Todos os exigidos	N/A	SIM	Anexo 16B (todas as páginas) Anexo 26I (todas as paginas) Anexo 26J (todas as paginas) Anexo 26K (todas as paginas) Anexo 26L (todas as paginas)

08981
3

O sistema deve analisar cada um dos pacotes que trafegam pela rede a que está conectado e também as relações de tais pacotes pacotes, como os adjacentes a ele no fluxo de dados da rede. Imediatamente após a identificação de uma eventual violação da política de segurança, o sensor deve enviar um alarme para o software de controle.	Todos os exigidos	N/A	SIM	Anexo 26G (todas as paginas) Anexo 26M (todas as paginas)
O bloqueio de uma tentativa de invasão não deve afetar os demais usuários.	Todos os exigidos	N/A	SIM	Anexo 26 (todas as páginas) Anexo 26G (todas as paginas) Anexo 26M (todas as paginas)
Capacidade de detecção de PHF e SMURF	Conforme Edital	N/A	SIM	Anexo 26C (todas as páginas) Anexo 26D (todas as páginas) Anexo 26H (todas as paginas) Anexo 26N - pag. 1
Capacidade de detecção de ataques genéricos de pacotes IP fragmentados e TEARDROP	Conforme Edital	N/A	SIM	Anexo 25 - pág. 4
Capcidade de detecção de ataques que utilizam sistemas Whisker anti-IDS	Conforme Edital	N/A	SIM	Ver Carta do Fabricante
Capacidade de detecção de ataques com assinaturas complexas tais como SIMPLEX MODE TCP hijacking e e-mail SPAN	Conforme Edital	N/A	SIM	Anexo 25 - pág. 4 Anexo 26O - pag. 1

2 - Características e Funcionalidades	Capacidade de detecção de ataques a servidores de WEB, servidores de e-mail (Denial of Services), FTP exploits, acesso não autorizado a servidores de banco de dados, servidores de aplicação e serviço (NT, Solaris, HP-UX, AIX e LINUX), servidores DNS e PROBIN ATTACKS	Conforme Edital	N/A	SIM	Anexo 25 - pág. 3
	Deve permitir a criação de regras personalizadas de identificação de invasões para que possa ser adaptado à estrutura particular disponível na CONTRATANTE. À medida que novos ataques forem sendo descobertos deve ser possível criar "assinaturas" associadas ao mesmo, de modo a previnir a reincidência	Conforme Edital	N/A	SIM	Anexo 26B (todas as páginas) Anexo 26K (todas as paginas)
	O Software de controle deve ser capaz de enviar alarmes via e-mail para notificar violação de uma dada regra de segurança.	Conforme Edital	N/A	SIM	Anexo 25 -pág. 6 Anexo 26Q (todas as paginas)
	O sistema deve registar informações tais como origem e destino, horário e tipo dos ataques ocorridos	Conforme Edital	N/A	SIM	Anexo 25 - pág. 6 Anexo 26R (todas as paginas)
	O sistema deve possibiltar atualização automática das assinaturas através de download seguro via WEB.	Conforme Edital	N/A	SIM	Anexo 25 - pág. 4 Anexo 26S (todas as paginas)
	O sitema deve possuir suporte a SSH para comunicação e comunicação segura	Conforme Edital	N/A	SIM	Anexo 25 - pág. 4 Anexo 26T (todas as paginas)

18973

O sistema deve realizar a verificação do TCP (Three Way Handshake)	Conforme Edital	N/A	SIM	Anexo 26 (todas as páginas) Anexo 26T (todas as paginas)
Deve oferecer respostas em tempo real para ataques via rede, com possibilidade de término da sessão e reconfiguração de regras de acesso em Firewall, que seja possível configurar	Conforme Edital	N/A	SIM	Anexo 25 - pág. 5 Anexo 26M (todas as paginas)
Deve suportar captura de LOG de sessão no formato padrão TCP DUMP	Conforme Edital	N/A	SIM	Anexo 26A - pág. 6 Anexo 26G (todas as paginas)
Deve bloquer tentativas de invasão que seja possível configurar	Conforme Edital	N/A	SIM	Anexo 25 - pág. 4 Anexo 26K (todas as paginas)
Deve permitir verificação de codificação de protocolos, scripts CGI, DNS e acesso remoto via BIND e DAEMONS	Conforme Edital	N/A	SIM	Anexo 25 - pág.3, 4 Anexo 26U (todas as paginas)
Deve permitir customização de rsposta a intrusões, mascaramento de tráfego, criação de conexões e modificação de ação de resposta	Conforme Edital	N/A	SIM	Anexo 25 - pág. 3, 4, 5 Anexo 26K (todas as assinaturas)
Deve possuir ferramentas de configuração com interface gráfica, controlando múltiplos sistemas de detecção	Conforme Edital	N/A	SIM	Anexo 25 - pág. 6 Anexo 26T (todas as paginas)
Deve registrar sessões de ataques com a possibilidade de PLAYBACK	Conforme Edital	N/A	SIM	Anexo 26A - pág. 6 Anexo 25 (todas as paginas)

3898

Deve utilizar base de dados em tempo real para ajuda a resposta a incidentes ocorridos	Conforme Edital	N/A	SIM	Anexo 26E - pág 1 Anexo 26V (todas as paginas)
Deve gerar sumários de relatórios das atividades registradas	Conforme Edital	N/A	SIM	Anexo 25 - pág. 5 Anexo 26X (todas as paginas)
Deve utilizar recursos de e-mail e traps SNMP para envio de informações à consoles com plataforma de gerenciamento	Conforme Edital	N/A	SIM	Anexo 25
Deve detectar eventos em ambientes computacionais com os sistemas operacionais Windows NT, 4.0 (server e Workstation) e UNIX (AIX, Solaris, LINUX, HP-UX)	Conforme Edital	N/A	SIM	Anexo 25 Anexo 26Z (todas as paginas)
Deve fornecer suporte ao gerenciamento de riscos através de relatórios técnicos e gerenciais pré-definidos, com detalhamento das informações coletadas	Conforme Edital	N/A	SIM	Anexo 26E - pág. 1, 2 Anexo 26X (todas as paginas)
Deve permitir a criação de relatórios técnicos e gerenciais personalizados, em formas textuais e gráficas	Conforme Edital	N/A	SIM	Anexo 26E - pág. 1, 2 Anexo 26X (todas as paginas)
Deve possuir documentação detalhada para possibilitar a configuração	Conforme Edital	N/A	SIM	Anexo 26E - pág. 1, 2 Anexo 26AA (todas as paginas)
Deve ser robusto para suportar qualquer tipo de ataque contra si próprio	Conforme Edital	N/A	SIM	Anexo 26A - pág. 5 Anexo 26O (todas as paginas)

	Deve permitir instalação sem necessidade de alteração na infra-estrutura de rede e sem causar degradação do seu desempenho	Conforme Edital	N/A	SIM	Anexo 26A - pág. 5 Anexo 26G (todas as paginas)
	Deve permitir configuração remota à partir da console de gerenciamento	Conforme Edital	N/A	SIM	Anexo 25 - pág.6 Anexo 26T (todas as paginas)
3 - Características Adicionais	O ambiente de gerenciamento dos servidores IDS deverá ser disponibilizado pela CONTRATADA em Hardware exclusivo, à parte, podendo estar integrado em um mesmo framework de gerenciamento dos servidores de segurança lógica, Switch tipo 5 e roteador tipo 1	Conforme Edital	N/A	SIM	A CONTRATADA disponibilizara o hardware exclusivo, à parte. Anexo 26L (todas as paginas)
	Todo o ambiente deverá ser acompanhado de documentação de instalação e configuração	Conforme Edital	N/A	SIM	A CONTRATADA Fornecerá a documentação pertinente e acompanhará toda a instalação. Anexo 26L (todas as paginas)
	O produto deverá integrar-se de forma transparente com o ambiente da gerência dos aservidores de segurança lógica, sendo capaz de reconfigurar automáticamente as regras destes equipamentos	Conforme Edital	N/A	SIM	Anexo 25 (todas as páginas) Anexo 26M (todas as paginas)

18976

Anexo 25

18975

CISCO SYSTEMS

Data Sheet

Cisco **IDS 4200** Series Sensors

Cisco integrated network security solutions enable organizations to protect productivity gains and reduce operating costs.

The Cisco IDS 4200 Series sensors are used in the Cisco Intrusion Protection System. These intrusion detection system sensors work in concert with the other components to efficiently protect your data and information infrastructure. With the increased complexity of security threats, achieving efficient network intrusion security solutions is critical to maintaining a high level of protection. Vigilant protection ensures business continuity and minimizes the effect of costly intrusions.

Additionally, Cisco's flexible deployment options allow businesses to minimize the total cost of ownership of their IDS deployments by delivering:

- unprecedented price/performance ratios
- the ability to simultaneously protect multiple network subnets through the support for multiple sniffing interfaces, thereby delivering up to five sensors in one
- a wide array of performance options
- investment protection by delivering modular, upgradable components
- support for multi-VLAN traffic
- embedded web-based management solutions packaged with the IDS sensors

Please refer to Table 1 for information on the characteristics of the Cisco IDS 4200 Series Sensors.

For details on the complete Cisco Intrusion Protection System, go to http://www.cisco.com/go/ids.

Deploying the Cisco IDS 4200 Series Sensors

The Cisco IDS 4200 Series includes four products: the Cisco IDS 4215, IDS 4235, IDS 4250 and IDS 4250-XL sensors. The Cisco IDS product line delivers a broad range of solutions that allow easy integration into many different environments, including enterprise and service provider environments. Each sensor addresses the bandwidth requirements at one of several speeds, from 80 Mbps to gigabits per second.

The Cisco IDS 4215 can monitor up to 80 Mbps of traffic and is suitable for T1/E1 and T3 environments. Additionally, multiple sniffing interfaces are supported on the IDS-4215 which allow the ability to simultaneously protect multiple subnets, thereby delivering five sensors in a single unit.

At 250 Mbps, the Cisco IDS 4235 can be deployed to provide protection in switched environments, on multiple T3 subnets, and with the support of 10/100/1000 interfaces it can also be deployed on partially utilized gigabit links.

The Cisco IDS 4250 supports a 500 Mbps speed and can be used to protect gigabit subnets and traffic traversing switches that are being used to aggregate traffic from numerous subnets. In addition, the Cisco IDS 4250 provides the flexibility to accommodate a simple hardware upgrade to scale to full line-rate gigabit performance.

At 1 Gbps, the Cisco IDS 4250-XL provides unprecedented performance by providing customized hardware acceleration to protect fully-saturated gigabit links as well as multiple partially-utilized gigabit subnets.

As shown in Figure 1, sensors can be placed on almost any network segment of the enterprise-wide network where security visibility is required.

Please refer to Table 2 for ordering information for the Cisco IDS 4200 Series Sensors.

Figure 1
Deployment Scenarios for the 4200 Series Appliance Sensors

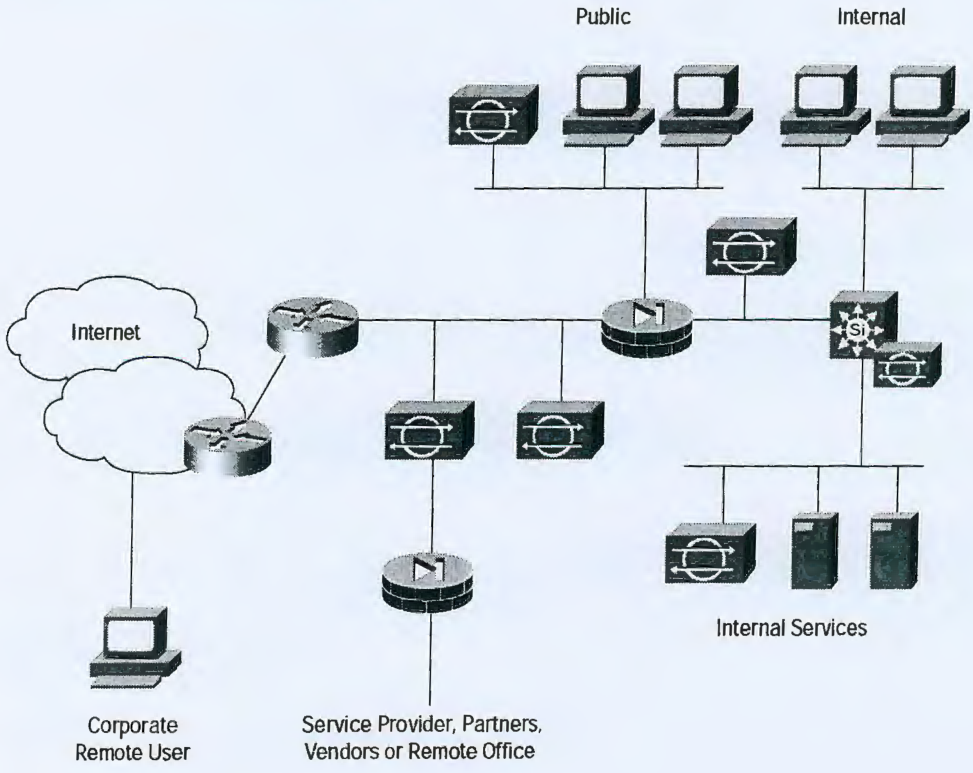

18773

Product Specifications

Table 1 Characteristics of Cisco IDS 4215, 4235, 4250, and 4250-XL Sensors

	Cisco IDS 4215	Cisco IDS 4235	Cisco IDS 4250	Cisco IDS 4250-XL
Performance	80 Mbps	250 Mbps	500 Mbps	1000 Mbps
Standard monitoring interface	10/100BASE-Tx	10/100/1000BASE-TX	10/100/1000BASE-TX	Dual 1000BASE-SX interface with MTRJ
Standard command and control interface	10/100BASE-Tx	10/100/1000BASE-TX	10/100/1000BASE-TX	10/100/1000BASE-TX
Optional interface	Four 10/100BaseTx (4FE) sniffing interfaces (allowing a total of 5 sniffing interfaces).	Four 10/100BaseTx (4FE) sniffing interfaces (allowing a total of 5 sniffing interfaces).	-1000BASE-SX (fiber) -Four 10/100BaseTx (4FE) sniffing interfaces (allowing a total of 5 sniffing interfaces).	1000BASE-SX (fiber)
Performance upgradable	No	No	Yes	No
Form factor	One rack unit	One rack unit	One rack unit	One rack unit
Advanced protection algorithms				
Stateful pattern recognition	Yes	Yes	Yes	Yes
Protocol parsing	Yes	Yes	Yes	Yes
Heuristic detection	Yes	Yes	Yes	Yes
Anomaly detection	Yes	Yes	Yes	Yes
Attack protection				
Sweeps or floods	Yes	Yes	Yes	Yes
Denial-of-service (DoS) mitigation	Yes	Yes	Yes	Yes
Worms or viruses	Yes	Yes	Yes	Yes
Common gateway interface (CGI) or WWW attacks	Yes	Yes	Yes	Yes
Buffer overflow protection	Yes	Yes	Yes	Yes

18972
ε

Table 1 Characteristics of Cisco IDS 4215, 4235, 4250, and 4250-XL Sensors

	Cisco IDS 4215	Cisco IDS 4235	Cisco IDS 4250	Cisco IDS 4250-XL
Remote-procedure call (RPC) attack detection	Yes	Yes	Yes	Yes
IP fragmentation attacks	Yes	Yes	Yes	Yes
Internet Control Message Protocol (ICMP) attacks	Yes	Yes	Yes	Yes
Simple Message Transfer Protocol (SMTP), Sendmail, Internet Message Access Protocol (IMAP), or Post Office Protocol (POP) attacks	Yes	Yes	Yes	Yes
File Transfer Protocol (FTP), Secure Shell Protocol (SSH), Telnet, and rlogin attacks	Yes	Yes	Yes	Yes
Domain Name System (DNS) attacks	Yes	Yes	Yes	Yes
TCP hijacks	Yes	Yes	Yes	Yes
Windows or NetBios attacks	Yes	Yes	Yes	Yes
TCP application protection	Yes	Yes	Yes	Yes
BackOrifice attacks	Yes	Yes	Yes	Yes
Network Timing Protocol (NTP) attacks	Yes	Yes	Yes	Yes
Customizable signatures using Signature Micro-Engine technology	Yes	Yes	Yes	Yes
Automated signature updates	Yes	Yes	Yes	Yes

18971
3

Table 1 Characteristics of Cisco IDS 4215, 4235, 4250, and 4250-XL Sensors

	Cisco IDS 4215	Cisco IDS 4235	Cisco IDS 4250	Cisco IDS 4250-XL
Alarm summarization	Yes	Yes	Yes	Yes
Support for 802.1q traffic	Yes	Yes	Yes	Yes
P2P / file sharing detection techniques	Yes	Yes	Yes	Yes
Secure communication				
IP Security (IPSec) or Secure Sockets Layer (SSL) between sensor and management console	Yes	Yes	Yes	Yes
Encrypted signature packages	Yes	Yes	Yes	Yes
SSH for remote administration	Yes	Yes	Yes	Yes
Serial Control Protocol (SCP) support for secure file transfer	Yes	Yes	Yes	Yes
IDS evasion protection				
IP fragmentation re-assembly	Yes	Yes	Yes	Yes
TCP stream re-assembly	Yes	Yes	Yes	Yes
Unicode deobfuscation	Yes	Yes	Yes	Yes
Active response actions				
Router access-control-list (ACL) modifications	Yes	Yes	Yes	Yes
Firewall policy modifications	Yes	Yes	Yes	Yes
Switch ACL modifications	Yes	Yes	Yes	Yes
Session termination via TCP resets	Yes	Yes	Yes	Yes
IP session logging or session replay	Yes	Yes	Yes	Yes

18970
E

Table 1 Characteristics of Cisco IDS 4215, 4235, 4250, and 4250-XL Sensors

	Cisco IDS 4215	Cisco IDS 4235	Cisco IDS 4250	Cisco IDS 4250-XL
Active notification actions				
Alarm display	Yes	Yes	Yes	Yes
E-mail alerts	Yes	Yes	Yes	Yes
E-page alerts	Yes	Yes	Yes	Yes
Customizable script execution	Yes	Yes	Yes	Yes
Multiple alarm destinations	Yes	Yes	Yes	Yes
Third-party tool integration	Yes	Yes	Yes	Yes
IDS active update bulletins	Yes	Yes	Yes	Yes
Administration				
Web user interface (Secure Hypertext Transfer Protocol [HTTPS])	Yes	Yes	Yes	Yes
Command-line interface (CLI) (console)	Yes	Yes	Yes	Yes
CLI (Telnet or SSH)	Yes	Yes	Yes	Yes
CiscoWorks VPN Security Management Solution (VMS) support	Yes	Yes	Yes	Yes
High availability				
Redundant power supply	No	Yes	Yes	Yes
Failure detection				
Monitoring link failure detection	Yes	Yes	Yes	Yes
Communications failure detection	Yes	Yes	Yes	Yes
Services failure detection	Yes	Yes	Yes	Yes
Device failure detection	Yes	Yes	Yes	Yes
Dimensions				

18969
e

Table 1 Characteristics of Cisco IDS 4215, 4235, 4250, and 4250-XL Sensors

	Cisco IDS 4215	Cisco IDS 4235	Cisco IDS 4250	Cisco IDS 4250-XL
Height	1.7 in. (4.37 cm)	1.67 in. (4.24 cm)	1.67 in. (4.24 cm)	1.67 in. (4.24 cm)
Width	16.8 in. (42.72 cm)	17.6 in. (44.70 cm)	17.6 in. (44.70 cm)	17.6 in. (44.70 cm)
Depth	11.8 in. (29.97 cm)	27.0 in. (68.58 cm)	27.0 in. (68.58 cm)	27.0 in. (68.58 cm)
Weight	11.5 lb (4.11 kg)	35 lb (15.88 kg)	35 lb (15.88 kg)	35 lb (15.88 kg)
Rack-mountable	Yes	Yes	Yes	Yes
Power				
Autoswitching	100V to 240V AC	110–220 VAC	110–220 VAC	110–220 VAC
Frequency	50 to 60 Hz	50–60 Hz	50–60 Hz	50–60 Hz
Operating current	1.5A	2.7A at 115V 1.3A at 220V	2.7A at 115V 1.3A at 220V	2.7A at 115V 1.3A at 220V
Operating environment				
Operating temperature	+5°C to +40°C (+41°F to +104°F)	10 to 35°C (50 to 95°F)	10 to 35°C (50 to 95°F)	10 to 35°C (50 to 95°F)
Nonoperating temperature	-25°C to +701/4°C (–13F to +1581/4°F)	–40 to 65°C (–40 to 149°F)	–40 to 65°C (–40 to 149°F)	–40 to 65°C (–40 to 149°F)
Operating relative humidity	5 to 95% (noncondensing)	8 to 80% (noncondensing)	8 to 80% (noncondensing)	8 to 80% (noncondensing)
Nonoperating relative humidity	5 to 95% (noncondensing)	5 to 95% (noncondensing)	5 to 95% (noncondensing)	5 to 95% (noncondensing)
Heat dissipation (most severe case with full power usage)	410 Btu/hr (full power usage (65W))	983 Btu/hr (maximum)	983 Btu/hr (maximum)	983 Btu/hr (maximum)

Note:

- This 80-Mbps performance for the Cisco IDS 4215 is based on the following conditions:
  - 800 new TCP connections per second
  - 800 HTTP transactions per second
  - Average packet size of 445 bytes,
  - Running Cisco IDS 4.0 Sensor Software
- This 250-Mbps performance for the Cisco IDS 4235 is based on the following conditions:
  - 3000 new TCP connections per second
  - 3000 HTTP transactions per second
  - Average packet size of 445 bytes

18968
E.

- Running Cisco IDS 4.0 Sensor Software

- This 500-Mbps performance for the Cisco IDS 4250 is based on the following conditions:
  - 5000 new TCP connections per second
  - 5000 HTTP transactions per second
  - Average packet size of 445 bytes
  - Running Cisco IDS 4.1 Sensor Software
- This 1000-Mbps performance for the Cisco IDS 4250-XL is based on the following conditions:
  - 5000 new TCP connections per second
  - 5000 HTTP transactions per second
  - Average packet size of 595 bytes
  - Running Cisco IDS 4.0 Sensor Software

Regulatory Compliance

- EMC—FCC (CFR 47 Part 15) Class A, CISPR 22 Class A, EN 55022 Class A, EN 55024, EN61000-3-2, EN61000-3-3, VCCI Class A, AS/NZS 3548 Class A, CE marking
- Safety UL 60950, CSA 22.2 No.60950, IEC 60950, EN 60950, AS/NZS 3260, CE marking.

Table 2 Ordering Information for the Cisco IDS 4200 Series Sensor

Product number	Product description
IDS-4215-K9	Cisco IDS 4215 Sensor (chassis, software, SSH, 2 onboard 10/100BASE-Tx interfaces with RJ-45 connector), 80-Mbps
IDS-4215-4FE-K9	Cisco IDS 4215 Sensor (chassis, software, SSH, 2 onboard 10/100BASE-Tx interfaces with RJ-45 connector plus 4FE interface card), 80-Mbps
IDS-4235-K9	Cisco IDS 4235 Sensor (chassis, software, SSH, 10/100/1000BASE-T with RJ-45 connector)
IDS-4250-TX-K9	Cisco IDS 4250 Sensor (chassis, software, SSH, 10/100/1000BASE-T with RJ-45 connector)
IDS-4250-SX-K9	Cisco IDS 4250 Sensor (chassis, software, SSH, 1000BASE-SX with SC connector)
IDS-4250-XL-K9	Cisco IDS 4250-XL Sensor (chassis, software, SSH, hardware accelerator, with dual 1000BASE-SX and MTRJ connectors)
IDS-XL-INT=	Cisco IDS Accelerator Card with dual 1000BASE-SX interfaces and MTRJ connectors
IDS-4250-SX-INT=	1000BASE-SX monitoring interface with SC connector
IDS-4FE-INT=	Spare 4FE (10/100 BaseTx) sniffing interfaces for 4215, 4235, & 4250
IDS-PWR=	Spare power supply for the Cisco IDS 4235 and 4250 sensors
IDS-SCSI=	Spare Small Computer Systems Interface (SCSI) hard disk drive for Cisco IDS 4250 Sensor
IDS-RAIL-2=	Two post rail kits for the Cisco IDS 4235 and 4250 sensor platforms
IDS-RAIL-4=	Four post rail kits for the Cisco IDS 4235 and 4250 sensor platforms

18967

Table 2 Ordering Information for the Cisco IDS 4200 Series Sensor

Product number	Product description
CON-SNT-IDS4215XK	Cisco SMARTnet support 8 x 5 x NBD (Cisco IDS 4215-K9)
CON-SNTE-IDS4215XK	Cisco SMARTnet support 8 x 5 x 4 (Cisco IDS 4215-K9)
CON-SNTP-IDS4215XK	Cisco SMARTnet support 24 x 7 x 4 (Cisco IDS 4215-K9)
CON-OS-IDS4215XK	Cisco SMARTnet Onsite support 8 x 5 x NBD (Cisco IDS 4215-K9)
CON-OSE-IDS4215XK	Cisco SMARTnet Onsite support 8 x 5 x 4 (Cisco IDS 4215-K9)
CON-OSP-IDS4215XK	Cisco SMARTnet Onsite support 24 x 7 x 4 (Cisco IDS 4215-K9)
CON-SNT-IDS4215-4FEXK	Cisco SMARTnet support 8 x 5 x NBD (Cisco IDS 4215-4FE-K9)
CON-SNTE-IDS4215-4FEXK	Cisco SMARTnet support 8 x 5 x 4 (Cisco IDS 4215-4FE-K9)
CON-SNTP-IDS4215-4FEXK	Cisco SMARTnet support 24 x 7 x 4 (Cisco IDS 4215-4FE-K9)
CON-OS-IDS4215-4FEXK	Cisco SMARTnet Onsite support 8 x 5 x NBD (Cisco IDS 4215-4FE-K9)
CON-OSE-IDS4215-4FEXK	Cisco SMARTnet Onsite support 8 x 5 x 4 (Cisco IDS 4215-4FE-K9)
CON-OSP-IDS4215-4FEXK	Cisco SMARTnet Onsite support 24 x 7 x 4 (Cisco IDS 4215-4FE-K9)
CON-SNT-IDS4235K9	Cisco SMARTnet support 8 x 5 x NBD (Cisco IDS 4235)
CON-SNTE-IDS4235K9	Cisco SMARTnet support 8 x 5 x 4 (Cisco IDS 4235)
CON-SNTP-IDS4235K9	Cisco SMARTnet support 24 x 7 x 4 (Cisco IDS 4235)
CON-OS-IDS4235K9	Cisco SMARTnet onsite support 8 x 5 x NBD (Cisco IDS 4235)
CON-OSE-IDS4235K9	Cisco SMARTnet onsite support 8 x 5 x 4 (Cisco IDS 4235)
CON-OSP-IDS4235K9	Cisco SMARTnet onsite support 24 x 7 x 4 (Cisco IDS 4235)
CON-SNT-IDS4250TK	Cisco SMARTnet support 8 x 5 x NBD (Cisco IDS 4250-TX)
CON-SNTE-IDS4250TK	Cisco SMARTnet support 8 x 5 x 4 (Cisco IDS 4250-TX)
CON-SNTP-IDS4250T	Cisco SMARTnet support 24 x 7 x 4 (Cisco IDS 4250-TX)
CON-OS-IDS4250TK	Cisco SMARTnet onsite support 8 x 5 x NBD (Cisco IDS 4250-TX)
CON-OSE-IDS4250TK	Cisco SMARTnet onsite support 8 x 5 x 4 Cisco (IDS 4250-TX)
CON-OSP-IDS4250TK	Cisco SMARTnet onsite support 24 x 7 x 4 (Cisco IDS 4250-TX)
CON-SNT-IDS4250SK	Cisco SMARTnet support 8 x 5 x NBD Cisco (IDS 4250-SX)
CON-SNTE-IDS4250SK	Cisco SMARTnet support 8 x 5 x 4 (Cisco IDS 4250-SX)
CON-SNTP-IDS4250SK	Cisco SMARTnet support 24 x 7 x 4 (Cisco IDS 4250-SX)
CON-OS-IDS4250SK	Cisco SMARTnet onsite support 8 x 5 x NBD (Cisco IDS 4250-SX)
CON-OSE-IDS4250SK	Cisco SMARTnet onsite support 8 x 5 x 4 (Cisco IDS 4250-SX)
CON-OSP-IDS4250SK	Cisco SMARTnet onsite support 24 x 7 x 4 (Cisco IDS 4250-SX)
CON-SNT-IDS4250XK	Cisco SMARTnet support 8 x 5 x NBD Cisco (IDS 4250-XL)
CON-SNTE-IDS4250XK	Cisco SMARTnet support 8 x 5 x 4 (Cisco IDS 4250-XL)
CON-SNTP-IDS4250XK	Cisco SMARTnet support 24 x 7 x 4 (Cisco IDS 4250-XL)
CON-OS-IDS4250XK	Cisco SMARTnet onsite support 8 x 5 x NBD (Cisco IDS 4250-XL)
CON-OSE-IDS4250XK	Cisco SMARTnet onsite support 8 x 5 x 4 (Cisco IDS 4250-XL)
CON-OSP-IDS4250XK	Cisco SMARTnet onsite support 24 x 7 x 4 (Cisco IDS 4250-XL)

18966
E

Table 2 Ordering Information for the Cisco IDS 4200 Series Sensor

Product number	Product description
CON-SNT-IDS4FE	Cisco SMARTnet support 8 x 5 x NBD (IDS-4FE-INT=)
CON-SNTE-IDS4FE	Cisco SMARTnet support 8 x 5 x 4 (IDS-4FE-INT=)
CON-SNTP-IDS4FE	Cisco SMARTnet support 24 x 7 x 4 (IDS-4FE-INT=)
CON-OS-IDS4FE	Cisco SMARTnet onsite support 8 x 5 x NBD (IDS-4FE-INT=)
CON-OSE-IDS4FE	SMARTnet onsite support 8 x 5 x 4 (IDS-4FE-INT=)
CON-OSP-IDS4FE	SMARTnet onsite support 24 x 7 x 4 (IDS-4FE-INT=)
CON-SNT-IDSXL	Cisco SMARTnet support 8 x 5 x NBD (IDS-XL-INT=)
CON-SNTE-IDSXL	Cisco SMARTnet support 8 x 5 x 4 (IDS-XL-INT=)
CON-SNTP-IDSXL	Cisco SMARTnet support 24 x 7 x 4 (IDS-XL-INT=)
CON-OS-IDSXL	Cisco SMARTnet onsite support 8 x 5 x NBD (IDS-XL-INT=)
CON-OSE-IDSXL	SMARTnet onsite support 8 x 5 x 4 (IDS-XL-INT=)
CON-OSP-IDSXL	SMARTnet onsite support 24 x 7 x 4 (IDS-XL-INT=)

Export Considerations

The Cisco IDS 4200 Series sensors are subject to export controls. Refer to the export compliance Web site for guidance at: http://www.cisco.com/wwl/export/crypto/.

For specific export questions, contact export@cisco.com.

Additional Information

For more information about the Cisco Intrusion Protection System, go to: http://www.cisco.com/go/ids

For more information about the CiscoWorks VMS Solutions (IDS management), go to: http://www.cisco.com/go/vms

Solutions | Products | Ordering | Support | Partners | Training | Corporate

Data Sheets

How to Buy Cisco Products | PDF | Get Acrobat Reader

English ▾ Select

18965

Data Sheet

Cisco Intrusion Protection

Efficient Intrusion Protection

Cisco® integrated network security solutions enable organizations to protect productivity gains and reduce operating costs.

The Cisco Intrusion Protection is designed to efficiently protect your data and information infrastructure. With the increased complexity of security threats, achieving efficient network intrusion security solutions is critical to maintaining a high level of protection. Vigilant protection ensures business continuity and minimizes the effect of costly intrusions.

Cisco's advancements deliver an efficient intrusion protection system through four critical elements:

1. *Accurate threat detection*—Cisco Intrusion Detection System Version 4.0 (Cisco IDS 4.0) delivers the first step in providing a secure environment by comprehensively detecting all potential threats.

2. *Intelligent threat investigation*—Cisco Threat Response technology virtually eliminates false alarms, and automatically determines which threats need immediate attention to avoid costly intrusions.

3. *Ease of management*—Browser-based tools simplify the user interaction, while providing powerful analytical tools that allow for a rapid and efficient response to threats.

4. *Flexible deployment options*—A range of high-availability devices provide the flexible backbone for creating the secure and efficient intrusion protection system.

All four elements combine to achieve a secure, efficient, and comprehensive intrusion protection solution.

Accurate Threat Detection

Providing unprecedented security, the Cisco IDS 4.0 is the core of the Cisco Intrusion Protection. It is designed to accurately identify and classify known and unknown threats targeting your network, including worms, denial-of-service (DoS), and application attacks. The first step in delivering an efficient and secure intrusion protection system is accurately detecting all possible threats. To achieve this goal, multiple detection methods are employed, thus ensuring comprehensive coverage. The methods include stateful pattern recognition, protocol analysis, traffic anomaly detection, and protocol anomaly detection. In addition, Cisco IDS 4.0 enhances the capability to prevent detected attacks from reaching their targets. And, several ease-of-use features are integrated to maximize efficiency.

Comprehensive Threat Protection

18964
6

- *Multiple Detection Methods*—Cisco IDS uses an array of detection methods to accurately detect nearly all potential threats. Building on seven years of IDS experience, Cisco delivers a hybrid system using detection methods most appropriate for the threat including stateful pattern recognition, protocol analysis, traffic anomaly detection, and protocol anomaly detection. Cisco IDS 4.0 delivers enhancements to these detection methods, most notably in the area of protocol anomaly detection. Additionally, Cisco IDS 4.0 delivers a Layer 2 signature engine to provide protection from ARP spoofing techniques in switched environments. These advanced detection techniques coupled with anti-IDS evasion techniques such as IP defragmentation, TCP streams reassembly, and deobfuscation, provide comprehensive protection against an array to threats allowing users to quickly identify and mitigate potential damage to data or networked assets.
- *Extensive protocol monitoring*—Cisco IDS 4.0 can monitor all the major TCP/IP protocols including, but not limited to, IP, Internet Control Message Protocol (ICMP), TCP, and User Datagram Protocol (UDP). It can also statefully decode application-layer protocols such as File Transfer Protocol (FTP), Simple Mail Transfer Protocol (SMTP), HTTP, Domain Name System (DNS), remote procedure call (RPC), NetBIOS, NNTP, and Telnet.
- *Comprehensive attack detection*—The Cisco IDS 4.0 has the most extensive and comprehensive capability to detect attacks in all the following categories:
  - *Exploits activity*—Indicative of someone attempting to gain access or compromise systems on your network, such as Back Orifice, failed login attempts, and TCP hijacking
  - *DoS activity*—Indicative of someone attempting to consume bandwidth or computing resources to disrupt normal operations, such as Trinoo, TFN, and SYN floods
  - *Reconnaissance activity*—Indicative of someone probing or mapping your network to identify "targets of opportunity" such as ping sweeps and port sweeps; usually a precursor to an actual exploit attempt
  - *Misuse activity*—Indicative of someone attempting to violate corporate policy; this can be detected by configuring the sensor to look for custom text strings in the network traffic; for example, XYZ Corporation could easily configure the Cisco IDS to send an alarm on and eliminate any connection that transmits the phrase "XYZ Confidential" in e-mail or FTP
  - Protection against file-sharing threats with support for advanced P2P attack mitigation techniques

Damage Prevention

Cisco IDS uses multilayer protection options to prevent an attack from successfully reaching the target. After the attack is accurately identified and classified, the system can stop the attack before damage occurs. Whether dropping the packet, terminating the session, reconfiguring access control lists (ACLs) on routers and switches, or dynamically modifying the firewall policy to "shun" the intruder, Cisco IDS offers an array of immediate response actions to stop attacks that can cost you time and money. Cisco IDS 4.0 enhances these techniques by extending its capability to include shunning by source and destination port number in addition to source and destination IP address. This provides added levels of granularity to the way in which response actions can be configured.

Easy to Use

- *Flexible policy language*—Because the security objectives for each IDS deployment are unique, Cisco IDS allows users to create and modify policies to specifically suit the environment in which they are deployed. Using the innovative Cisco Threat Analysis Micro Engine (T.A.M.E.) policy language, users have the flexibility to create new policies or modify existing policies to meet their unique security objectives. Because T.A.M.E. policies are decoupled from the sensing application, changes do not affect the sensor performance or reliability. Unlike other security languages that rely on simple pattern matching, Cisco T.A.M.E language allows users to take advantage of the underlying protocol analysis capabilities. Cisco IDS 4.0 simplifies the policy management with improved navigation, allowing global changes to be implemented across categories. Additionally, Cisco IDS 4.0 now provides detailed information about the alarm trigger, providing the user with forensics data and advanced analysis data to simplify the support process.

- *Automated updates to streamline management*—Cisco IDS Active Update technology automates the process of updating deployed sensors, thus reducing the operating costs. This process provides a facility to automatically distribute new signature files and application upgrades to sensors without operator involvement. Using a secure staging technique, new signature files are placed on a central server and passed to the sensor at scheduled intervals. After verifying the integrity of the package, the sensor automatically installs the update. Alternatively, the same degree of automation is delivered through the IDS Management solution (VMS-VPN & Security Management Solution) where updates can be automatically pushed to sensors using secure file transfer capabilities. These capabilities significantly streamline the process of regularly updating remote sensors, thereby lowering the recurring operational costs associated with this task. Additionally, users can subscribe to Cisco IDS Active Update notification services to stay informed about breaking vulnerability news and posted countermeasures. These policy updates are developed and maintained by the Cisco Countermeasures Research Team (Cisco-CRT). This elite team of security professionals is dedicated to rapid response to new and evolving threats.

Intelligent Threat Investigation

Cisco's Threat Response technology works with Cisco Network IDS sensors to provide an efficient intrusion protection solution. Cisco Threat Response virtually eliminates false alarms, escalates real attacks, and aids in the remediation of costly intrusions.

Unlike other intrusion-management solutions, only Cisco Threat Response technology provides an automated, just-in-time analysis of each targeted host to determine whether a compromise has actually occurred. Only by investigating the host under attack can you efficiently uncover the real intrusions and address them quickly. The automated, real-time capabilities of this technology help protect your network environment around the clock.

The result? False alarms are eliminated and real intrusions are quickly identified and addressed, saving you time, resources, and the high costs associated with recovering from a successful attack.

Increased Efficiency, Reduced Costs

- *Elimination of false alarms and escalation of real attacks*—With its innovative intrusion-investigation process focused on the targeted host, Cisco Threat Response accurately determines whether or not an IDS alarm needs your attention. Designed to handle intrusion response in the same way an experienced security officer would, Cisco Threat Response carefully examines the targeted host to determine if the attack worked. This technology uses a three-phased approach:
    - *Basic investigation of target vulnerability*—The first step involves a noninvasive real-time check of the OS of the targeted system, its patch levels, and its Web services as applicable to determine if the attack could have succeeded. For example, a Linux attack against a Windows system would be downgraded, and indicated as a failed attack, whereas a Windows attack against a Windows system would be indicated as a potentially successful attack.
    - *Advanced investigation of target*—The second step is a detailed system-level

18962

investigation that includes the capture and analysis of Web logs, system logs, and other relevant data. Based on this detailed level of investigation, the Cisco Threat Response technology can determine if an attack succeeded or failed. Failed attacks are downgraded so staff can focus on the critical events. (*This capability is available only as part of a CiscoWorks VPN Security Management Solution [VMS] technology bundle*).

  - *Forensic data capture*—It is not enough to simply tell you that a confirmed problem exists. Cisco Threat Response goes further to actively collect relevant forensic evidence and provide you with intelligent information so you can make informed decisions. This technology immediately copies and safely stores audit trails, log files, and intrusion traces from the targeted system. In this way, the intruder cannot avoid detection by tampering with these files. Cisco Threat Response real-world advice and recovery procedures guide you in dealing with the incident effectively. (*This capability is available only as part of a CiscoWorks VMS technology bundle.*)

- *Fast, consistent, and automated process*—Twenty-four hours a day, seven days a week, Cisco Threat Response robustly, consistently, and automatically investigates attacks that threaten your network. Responding in seconds to a detected network attack, Cisco Threat Response leaves an intruder wondering what happened.
- *Easy deployment*—Cisco Threat Response technology allows for host investigation without the need to deploy software agents on each system within the enterprise. This means rapid deployment and ease of maintenance.

Ease of Management

Cisco provides effective security monitoring and configuration regardless of deployment size using a range of management options. All management tools are designed with an intuitive user interface, and easy navigation that enables rapid installation, configuration, and management of security events and devices. In addition, Cisco IDS 4.0 delivers a full featured IOS-like command-line interface (CLI) over a secure SSH connection. Intuitive Event Display.

- *Secure, browser-based graphical user interface (GUI)*—Alarms can be easily viewed from practically any desktop, no matter which operating system is being used on the desktop. The result is rapid access to data from systems throughout the enterprise. The familiar browser interface enhances ease of use. And with Secure Sockets Layer (SSL), security of data is maintained.
- *Unified, scalable view of all security events*—With the CiscoWorks VMS Solution, events from all types of security devices, including firewall, virtual private network (VPN), and IDS can be viewed from a single console. Multiple data sources can be supported and managed. This enhances the ability to view security across the enterprise.

Easy Alarm Processing

- *Forensic data*—With the Cisco Threat Response technology, the GUI provides a view into the steps taken to investigate and confirm intrusion events. To aid in remediation of intrusions, forensic data collected by this technology is accessible. Examples include Web logs, system logs, and other relevant data. (*This capability is available only as part of a CiscoWorks VMS technology bundle.*)
- *Correlation of events*—CiscoWorks VMS provides event correlation to enable improved confidence in alarm data by corroborating data from multiple security devices.
- *Network security database (NSDB)*—The NSDB provides instant access to specific information about the attacks, hyperlinks, potential countermeasures, and related vulnerabilities. Because the NSDB is an HTML database, it can be personalized for each user to include operation-specific information such as response and escalation procedures for specific attacks.

Flexible Reporting and Notification

18961
E.

- *Default reports*—Default reports about network activities monitored by sensors on your include summary reports based on alarms, sources, or destinations. Because these reports are HTML based, they can be sent in e-mail to key administration personnel.
- *Custom reporting*—Custom reports can be created to meet the specific needs of your environment.

Simple Configuration

- *Wizard-based configuration*—Wizards guide the user through the configuration process, enabling quick and easy configuration of sensors. This wizard-based capability is extended to deliver easy-to-use and intuitive signature customization.
- *Automatic updates*—Automatic update capabilities maintain the effectiveness of the intrusion protection system, and simplify the regular maintenance.
- *Remote management*—Because you are not always at the same computer, or where the IDS system is located, remote access capabilities via a secure Web browser connection allow for easy remote connectivity.

Scalable Enterprise Management

- *Multitiered architecture*—CiscoWorks VMS promotes a three-tiered architecture that meets the enhanced scalability needs of enterprise security deployments.
- *Flexible device grouping*—Easily manage large IDS deployments by grouping devices by function, by location, or by configuration to perform mass configuration changes.
- *Role-based access control*—Control administrative access to ensure proper device authorization.
- *Tiered approval model (optional)*—Separate configuration definition and deployment authorities to provide proper audit and control.

Flexible Deployment Options

Cisco offers the widest range of network IDS deployment options, providing customers with the ability to choose the intrusion solution that is most cost-effective for their environments. All solutions are designed for high availability and backed by outstanding customer support from Cisco. Network IDS solutions are available in a range of performance levels from 45 Mbps up to 1 Gbps. Network intrusion protection is available as dedicated appliances using the Cisco 4200 IDS Sensor Series, or as integrated solutions using the Cisco Catalyst® 6500 IDS modules and the network module for the Cisco 2600, 3600, and 3700 series routers. Additionally, a subset of network IDS functions is available as an integrated solution in routers and firewall systems.

Easy Installation

The Cisco Intrusion Protection has been designed for easy, rapid deployment. Appliance sensors were designed for technician-level installation, requiring the sensors to be plugged into the network, turned on, and configured with a few initialization parameters. Installation of the Catalyst 6500 Series 600 Mbps IDS Service Module is as easy as sliding the module into an open chassis slot, configuring the module with the initialization parameters, and configuring the switch to recognize the card and send traffic to it. The Cisco IDS network modules fit into a single network module slot on the Cisco 2600XM Series, Cisco 3660, and Cisco 3700 Series platforms, while taking advantage of the power options of the router, including DC power and redundant power. After the sensors are initialized and running, configurations can be modified and pushed to them from any of the management consoles.

18760

Low Total Cost of Ownership

Cisco's flexible deployment options allow businesses to minimize the total cost of ownership of their IDS deployments by delivering:

- Unprecedented price-performance ratios
- The ability to simultaneously protect multiple network subnets through the support for multiple sniffing interfaces, thereby delivering up to five sensors in one
- A wide array of performance options
- Investment protection by delivering modular, upgradable components
- Support for VLANs that enables the optimization of traffic inspection based on bandwidth utilizations
- Embedded Web-based management solutions packaged with the IDS sensors

Worldwide, World Class Support

Cisco provides leading-edge services to extend and enhance the operations of your Cisco Systems® products. The Cisco IDS 4200 Series sensors and the Catalyst 6500 IDS Module are covered by a full suite of SMARTnet maintenance options. This includes hardware support as well as sensor software upgrades and access to the latest signatures posted on Cisco.com. The management console software is covered by Software Application Support (SAS) and Software Application Support Plus Upgrades (SASU).

Network IDS Deployment Options

- *Cisco IDS 4200 Series sensors*—Dedicated IDS solutions enable deployment of IDS sensors wherever they are needed in the network architecture. Four performance levels are available:
  - Cisco IDS 4215—80 Mbps
  - Cisco IDS 4235—250 Mbps
  - Cisco IDS 4250—500 Mbps
  - Cisco IDS 4250 XL—1000 Mbps
- *Cisco IDSM-2 Module for the Cisco Catalyst 6500 chassis*—This product efficiently integrates full IDS capabilities into the Cisco Catalyst Switch via a dedicated module, providing integrated protection at 600 Mbps.
- *Cisco IDS Network Module for the Cisco 2600, 3600, and 3700 series router*—Each sensor addresses the bandwidth requirements of different routers up to 10 Mbps in the Cisco 2600XM, and up to 45 Mbps in the Cisco 3700 Series. By integrating IDS and branch office routing, Cisco reduces the complexity of securing WAN links, while reducing operational costs. Additionally, by delivering full-featured intrusion protection to remote offices and branch offices, network administrators can now mitigate threats at these remote locations and effectively isolate them from the corporate network.
- *Router sensor*—The router sensor provides a focused set of IDS capabilities via a software solution integrated into the router OS.
- *Firewall sensor*—The firewall sensor provides a focused set of IDS capabilities via a software solution integrated into the firewall OS.

Ordering Information

For ordering details of the Cisco Intrusion Protection elements, visit:

http://www.cisco.com/go/ids

18959 ε

For More Information

For more information, visit:

http://www.cisco.com/go/ids

| Home | What's New | How to Buy | Login | Register | Feedback | Search | Map/Help |

Posted: Wed Jun 11 22:02:47 PDT 2003

Anexo (2.53)

18958
E

CISCO SYSTEMS

DATA SHEET

Cisco **IDS 3.0** Sensor Software for Cisco 4200 Series Appliances

The industry-leading Cisco IDS Intrusion Detection System solution provides customers with unmatched intrusion protection technology through the Cisco Active Defense System. The integrated hardware and software solution delivers best-of-breed protection for both perimeter and internal resources.

Version 3.0 is the latest release of the Cisco IDS application software for the Cisco 4200 Series appliances. It delivers the latest in innovative IDS features, including Active Update signature distribution mechanisms, customizable signature language, extensions to the Active Response capabilities, and secure administration.

Key Features and Benefits

Signature Definition and Distribution Enhancements

Active Update Mechanism

Cisco IDS now provides a facility to automatically distribute new signature files and application upgrades to sensors without operator involvement. Utilizing a secure staging technique, new signature files are placed on a central server and passed to the sensor at scheduled intervals. After verifying the integrity of the package, the sensor automatically installs the update. This new capability significantly streamlines the process of regularly updating remote sensors, thereby lowering the recurring operational costs associated with this task. Additionally, users can subscribe to Cisco Active Update notification services to stay informed about breaking vulnerability news and posted countermeasures.

Signature Instruction Language for Verifying Exploits and Reconnaissance (SILVER)

Cisco IDS continues to innovate with the release of the Cisco SILVER signature language. Users can now develop custom signatures spanning Layer 2 through Layer 7 attributes. Developed by the Cisco Countermeasures Research Team, SILVER is the mechanism that this team uses to develop all new signatures for the IDS product. Now Cisco is exposing this capability to its customers, providing them unparalleled flexibility to develop new signatures.

This new language also decreases the time between signature releases because the signatures can be released as update files rather than compiled in the core sensor binary. Similar to antivirus products, the signature engines and the data files can be released separately.

Cisco Systems, Inc.

18957

Tunable Signatures

One of the challenges of event monitoring for IDS is to reduce the occurrence of false positives, thereby allowing the user to focus only on alarms that are deemed security relevant. However, because the security objectives for each IDS deployment are unique, Cisco IDS 3.0 adds granularity to the way in which sensors may be tuned to specifically suit the environment in which they are deployed. Users may now tune signature threshold parameters to meet their unique security objectives. For example, for the ping sweep signature, the user can now define the number of hosts pinged before the signature fires.

Alarm Summarization

Attackers may use tools such as Stick to cause sensors to send numerous alarms to management consoles as a way of flooding the management console. The Alarm Throttle feature in Cisco IDS 3.0 allows users to consolidate alarms to a single event, thereby preventing the management console from being overwhelmed with numerous events. For example, if the rate of alarms being sent by the sensor exceeds a certain threshold value, the signature may be set to automatically throttle to a summarization mode to produce a single alarm that comprises numerous alarm firings. A time interval may also be specified that determines the rate at which subsequent consolidated alarms are sent to the management console.

Detection of Network DoS Attacks

Cisco IDS 3.0 adds functionality to detect network denial-of-service (DoS) attacks that are targeted at consuming network resources. An inspector monitors the data capture by maintaining a counter that tracks the number of packets being captured per second. If the rate of packets being captured exceeds a specified rate and such peaks occur at a certain frequency, an alarm is generated to alert the user of a possible DoS attempt.

Active Response Extensions

Shunning with the Cisco PIX Firewall and Cisco Catalyst Switches

Cisco IDS incorporates proactive response functionality into the sensor appliances by allowing users to configure the IDS to dynamically respond to an attacking host by preventing new connections and disallowing packets from any existing connection. Cisco IDS 3.0 extends this Active Response capability beyond Cisco routers into a wide range of Cisco's high-performance networking devices, including the Cisco PIX® Firewall (running PIX 6.0+), Cisco Catalyst® 6000 Switch, and Cisco Catalyst 5000 Switch. This shunning gives security operators unparalleled control to reach across the network to quickly stop misuse and end intruders' access to the network.

Interoperability with Existing ACLs

Shunning is accomplished by the dynamic modification of access control lists (ACLs) on managed devices. Cisco IDS 3.0 allows the user to configure the sensor to shun while maintaining any user-defined access control entries on the specified interface and direction.

Secure Administration

In addition to existing functionality that provides secure communications between the sensor and the management console using IP Security (IPSec), Cisco IDS 3.0 also supports the Secure Shell Protocol (SSH) to allow users to remotely access sensors over a secure connection.

IP Session Replay

IP session logging provides extensive logging that is important for system troubleshooting as well as for reconstructing system events before and after attacks. Cisco IDS 3.0 augments this existing capability by converting these logs to a standard TCP dump format that allows them to be viewed and replayed using public domain utilities, such as Ethereal and TCPReplay.

Enhanced Filtering

Users have the ability to specify which source or destination IP addresses that specified signatures must alarm on. Because certain signatures may be classified as security relevant for certain source or destination IP addresses and yet do not need to be analyzed for other IP addresses, this level of enhanced filtering minimizes the occurrence of false positives.

Alarm Generation When Sniffing Interface is Disconnected

Sensors are equipped with a monitoring interface for data packet capture and a command and control interface for transmitting alarms to the management console and receiving configuration information from the

18956
E

management console. With Cisco IDS 3.0, when the sniffing interface is disconnected, an alarm is triggered. This setup provides the user with an alert mechanism when the interface is tampered with and hence assures persistent operation.

Technical Specifications

Ordering Information

SMARTnet™ customers may download Cisco IDS 3.0 from Software Center on CCO at http://www.cisco.com

Part number for IDS Software Upgrade for Non-Support Customers: IDS-SW-U

Export Considerations

IDS sensors are subject to export controls. Please refer to the export compliance Web site at

http://www.cisco.com/wwl/export/crypto/ for guidance.

For specific export questions, please contact export@cisco.com.

Additional Information

Cisco Secure Intrusion Detection System:

http://www.cisco.com/go/ids/

Cisco Systems

Corporate Headquarters
Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA
www.cisco.com
Tel: 408 526-4000
800 553-NETS (6387)
Fax: 408 526-4100

European Headquarters
Cisco Systems Europe
11, Rue Camille Desmoulins
92782 Issy-les-Moulineaux
Cedex 9
France
www-europe.cisco.com
Tel: 33 1 58 04 60 00
Fax: 33 1 58 04 61 00

Americas Headquarters
Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA
www.cisco.com
Tel: 408 526-7660
Fax: 408 527-0883

Asia Pacific Headquarters
Cisco Systems Australia, Pty., Ltd
Level 9, 80 Pacific Highway
P.O. Box 469
North Sydney
NSW 2060 Australia
www.cisco.com
Tel: +61 2 8448 7100
Fax: +61 2 9957 4350

Cisco Systems has more than 200 offices in the following countries and regions. Addresses, phone numbers, and fax numbers are listed on the

Cisco Web site at www.cisco.com/go/offices

Argentina • Australia • Austria • Belgium • Brazil • Bulgaria • Canada • Chile • China PRC • Colombia • Costa Rica • Croatia Czech Republic • Denmark • Dubai, UAE • Finland • France • Germany • Greece • Hong Kong SAR • Hungary • India • Indonesia Ireland • Israel • Italy • Japan • Korea • Luxembourg • Malaysia • Mexico • The Netherlands • New Zealand • Norway • Peru Philippines • Poland • Portugal • Puerto Rico • Romania • Russia • Saudi Arabia • Scotland • Singapore • Slovakia • Slovenia • South Africa Spain • Sweden • Switzerland • Taiwan • Thailand • Turkey • Ukraine • United Kingdom • United States • Venezuela • Vietnam • Zimbabwe

Printed in the USA

Anexo 26A

18955
E

Cisco Intrusion **Detection** System

Q. What is a "network-based" IDS?

A. Two basic types of IDSs are on the market today: host-based and network-based systems. The fundamental difference between them is the source of the activity that they monitor and analyze to detect intrusions. Host-based IDSs monitor activity on a host or end system, while network-based IDSs monitor network traffic. Host-based IDSs are used to protect critical network servers or other individual systems containing sensitive information.

Network-based IDSs are used to monitor activity on a specific network segment. Whereas a host-based IDS resides on a workstation and shares CPU with other user applications, a network-based solution is a dedicated platform. Network-based IDSs perform a rule-based or expert system analysis of traffic using parameters set up by the security manager, and the signatures, which flag suspicious or attack activity. The systems analyze network packet headers to make security decisions based on source, destination, and packet type. They also analyze packet data to make decisions based on the actual data being transmitted. These systems scale well for network protection because the number of actual workstations, servers, or user systems on the network is not critical, the amount of traffic is what matters. In addition, sensors placed around the globe can be configured to report back to a central site, enabling a small team of security experts to support a large enterprise. The Cisco® network-based Intrusion Detection System provides network administrators with enhanced security technology and capabilities to secure their networks.

Q. If I already have a firewall, do I really need an IDS?

A. Absolutely. Although an IDS will not replace your firewalls or other security devices for that matter, it serves a very complementary role and addresses certain risks that firewalls cannot. The primary function of the firewall is to control access to services and hosts based on your site security policy. If a service or connection to a specific host is permitted, firewalls typically permit all such traffic, and they do not inspect the content of the permitted traffic. An example is permitting public access to a Web server on a DMZ. All connection requests to the Hypertext Transfer Protocol (HTTP) port on that Web server will be permitted by the firewall, including malicious traffic directed at the HTTP server to exploit a buffer overflow vulnerability. Although most firewalls will not protect against data/content-driven attacks (for example, buffer overflow), IDSs will. Furthermore, firewalls typically will not protect you against attacks originating from inside your network or entering your environment from other ingress points not protected by firewalls (for example, remote access servers). IDSs can be strategically deployed to monitor activity from internal sources and other network ingress points without impacting your network. Deploying an IDS to complement your firewall(s) will significantly enhance your security posture.

18954
ε

Q. Is there a mechanism by which users may contact the IDS Product Team at Cisco?

A. Yes. Users may pose questions, requests, and comments to the following e-mail address:

ids-news@cisco.com

In addition, users have the ability to share experiences with other users and also pose questions to the Cisco IDS Engineering & Product Marketing teams at the IDS Networking Professionals Forum at:

http://forums.cisco.com/eforum/servlet/
NetProf?page=netprof&CommCmd=MB%3FcmdDdisplay_messages26mode3Dnew26location%3D.ee6e1fc

Q. Does anyone offer a managed IDS service using the Cisco IDS?

A. Yes, numerous managed service providers offer a managed IDS service using the Cisco IDS. These managed service providers include AT&T, Counterpane, IBM Emergency Response Services, NetSolve, Riptech, RedSiren, and Ubizen.

Sensors

Q. What are the new features of the Cisco IDS 4.0 Sensor software?

A. The Cisco IDS 4.0 Sensor software delivers a number of new features and enhancements to the network-based IDS portfolio. These features include:

- Re-architecture of communications protocol to enhance the efficiency of message transactions
- Common code base to allow feature parity between the appliance sensor and the switch sensor
- Delivery of a Layer 2 signature engine to mitigate issues such as man in the middle attacks and ARP spoofing in switched environments
- Introduction of an SMB engine to efficiently address attacks related to SMB
- Ability to capture the trigger packet that caused an alarm
- Enhanced shunning capabilities to allow shunning by port address
- Major enhancements to our existing protocol anomaly techniques
- Provision of Analysis Statistics Engine to deliver information of metrics such as bad checksums, bytes processed, data rates Mbps, TCP nodes per sec, and other analysis metrics
- Introduction of a full featured Cisco IOS-like CLI (command-line interface) for unprecedented sensor management over a secure SSH connection
- Capability of capture and display of the VLAN ID of the malicious traffic that was detected
- Enhancements to IP Fragmentation Reassembly
- Higher levels of granularity for the alarm information that is transmitted to the management console
- Support for ntp
- NAT support
- Logical signature groupings to allow global changes across the groupings
- Ability to implement exceptions to filter events to be displayed
- Tunability of IP session logging parameters

18953
E.

Q. What performance numbers (Mbps) are supported by the Cisco IDS Sensors?

A. The Cisco IDS 4215 supports 80 Mbps of performance and can be used to protect T1/E1/T3 environments.

At 250 Mbps, the Cisco IDS 4235 can be deployed to provide protection in switched environments, on multiple T3 subnets, and with the support of 10/100/1000 interfaces, it can also be deployed on partially utilized gigabit links.

The Cisco IDS 4250 supports superior performance at 500 Mbps and can be used to protect gigabit subnets and traffic traversing switches that are being used to aggregate traffic from numerous subnets.

Intrusion protection for fully saturated gigabit links is delivered by the Cisco IDS 4250-XL. Using customized hardware acceleration, the IDS-4250-XL can be used to protect gigabit subnets and multiple partially utilized gigabit links.

The Cisco Catalyst® 6500 Series Intrusion Detection System (IDSM-2) Services Module supports 600 Mbps. This module operates within the Catalyst 6500 Series and provides protection for traffic traversing the switch, which could be traffic from a single subnet or from numerous subnets that are being aggregated through the switch.

The Cisco IDS Network Module provides full-featured Intrusion Protection that is integrated into the Cisco 2600, 3600, and 3700 series routers. Each sensor addresses the bandwidth requirements of different routers up to 10 Mbps in the Cisco 2600XM, and up to 45 Mbps in the Cisco 3700 Series. By integrating IDS and branch office routing, Cisco reduces the complexity of securing WAN links and at the same time reduces operational costs. Additionally, by delivering full-featured intrusion protection to remote offices and branch offices, network administrators can now mitigate threats at these remote locations and effectively isolate them from the corporate network. The Network Module has the capability of inspecting GRE/IPsec encrypted packets that are traversing the router into which it integrates.

Q. How does the IDS sensor work?

A. Sensors monitor the network traffic by directly "tapping" the line (for example, via a shared-media hub) or by receiving copies of the traffic (for example, Switched Port Analyzer [SPAN] port on a switch) using a passive, promiscuous interface (the "monitoring interface"). The sensor analyzes the captured packets and compares them against a rule set of typical intrusion activity (that is, "signatures"). If the captured packets match a defined intrusion pattern in the rule set, the sensor sends an alarm to the management console and automatically responds (if configured to do so). The alarms are sent out a separate management interface so as not to impede continual packet capture by the monitoring interface.

Q. What kind of a performance impact does the sensor impose on the monitored network?

A. None. Sensors operate by "tapping" the network (for example, via a shared-media hub) or off copies of the packets (for example, via a switch SPAN port). The monitoring interfaces on the sensors are passive and do not source packets onto the network (the one exception is TCP reset packets for automatic response).

Q. How do you deploy sensors in a switched environment?

A. With most IDS products on the market today, sensors must be placed on the switch SPAN port to monitor network traffic. Although the SPAN port can provide access to network traffic, it does have certain limitations (for example, limited number of SPAN sessions). The Catalyst 6000 IDS Module was designed specifically to address switched environments by integrating the IDS functionality directly into the switch and taking traffic right off the switch backplane.

18952
ε

Q. What is the Cisco IDS Network Module for the Cisco 2600, 3600, and 3700 series routers?

A. The Cisco IDS Network Module is a network module that is installed in a Cisco 2600, 3600, or 3700 series chassis to provide full-featured intrusion protection services within the router. The Cisco 2600, 3600, and 3700 Series IDS Network Module provides the ability to inspect all traffic traversing the router, to identify unauthorized or malicious activity such as hacker attacks, worms, or denial-of-service attacks, and to terminate this traffic to suppress or contain threats.

Q. How does the Cisco IDS Network Module work?

A. The Cisco 2600, 3600, and 3700 Series IDS Network Module receives copies of packets directly from the router's backplane in a passive or promiscuous mode. The packets are passed through the internal monitoring interface for classification and processing. The Cisco 2600, 3600, and 3700 Series IDS Network Module analyzes the captured packets and compares them against a rule set of typical intrusion activity. If the captured packets match a defined intrusion pattern in the rule set, the IDS Network Module can take one of two actions. It can send a command to the router to either shut down the interface or it can send a TCP reset packet to the sender to stop the TCP session causing the attack.

Q. What is the rated performance of the Cisco 2600, 3600, and 3700 Series IDS Network Module?

A. The Cisco 2600, 3600, and 3700 Series IDS Network Module provides up to 10 Mbps for the 2600XM Series and up to 45 Mbps for the 3700 Series, depending on the platform in which the network module is inserted.

Q. Can the Cisco Network IDS Sensors monitor trunked traffic?

A. Both IDSM and the appliance sensors can monitior 802.1q traffic and, hence, are VLAN aware.

Q. What type of interfaces are supported on the appliance sensors?

A. Copper interfaces are supported on the IDS 4215 and IDS 4235. Both Copper and Fiber interfaces are supported on the IDS 4250 Sensor. The 4250-XL supports dual fiber interfaces with MTRJ connectors.

Q. Does Cisco IDS provide multi-interface support?

A. Yes. Dual sniffing interfaces are supported on the IDS 4250-XL. Up to 5 interfaces are supported on the IDS 4215, 4235, and 4250 Sensor appliances.

A configurable four Fast Ethernet interface card is provided for other models of the Cisco 4200 Series sensors to deliver a total of five sniffing interfaces for each sensor—one onboard sniffing interface plus four Fast Ethernet configurable interfaces.

The Cisco IDSM-2 can be used to monitor traffic from multiple interfaces. The network module for the Cisco access routers can monitor traffic from any of the router interfaces.

Q. Is the user notified when the sniffing interface of a sensor is disconnected?

A. Sensors are equipped with a monitoring interface for data packet capture and a command and control interface for transmitting alarms to the management console and receiving configuration information from the management console. When the sniffing interface is disconnected, an alarm is triggered. This setup provides the user with an alert mechanism when the interface is tampered with and, hence, assures persistent operation.

18951
ε

Q. Does the Cisco IDS Sensor provide an indication of when it is oversubscribed?

A. Yes. The sensor, IDSM, and network module issue an alarm when their respective performance ratings are exceeded.

Q. Where can I find more details on the IDS signature algorithms?

A. For more information on the signature algorithms, please refer to a white paper that may be downloaded at:

http://www.cisco.com/warp/public/cc/pd/sqsw/sqidsz/prodlit/idssa_wp.htm

Q. Does Cisco IDS protect against common IDS evasion techniques?

A. Cisco IDS also includes protection from a number of advanced, anti-IDS evasion techniques including:

- IP fragmentation reassembly
- TCP streams reassembly
- Unicode Web deobfuscation

Q. Can the sensor itself be attacked and compromised?

A. A properly configured and installed sensor cannot be compromised. The monitoring interface (connected to the production network) cannot be detected, and packets cannot be directed at it. The interface is in promiscuous mode, and has neither a protocol stack nor an IP address bound to it. It is not susceptible to "antisniff" detection techniques. The separate management interface does have an IP address, but Cisco recommends that a separate, isolated management subnet be used to provide connectivity from the management interface on the sensor to the IDS management console. In addition, only a very limited number of services are available from the management interface, and access controls can be configured to allow only designated management systems to connect to the sensor.

Q. What is Cisco Countermeasures Research Team (C-CRT)?

A. The core of the Cisco IDS solution—the advanced protection capabilities—is developed and maintained by C-CRT. This team of elite security professionals is dedicated to:

- Advancing countermeasures research
- Identifying and responding to new threats
- Distributing proactive signature files and signature micro-engines
- Maintaining our network security database (NSDB)
- Contributing research to the Cisco Security Encyclopedia (CSEC)
- Improving the state of threat mitigation science

18950
3

The C-CRT is comprised of elite "white hat" personnel. C-CRT's esteemed credentials include:

- *Heritage and tenure*—most joined Cisco through the WheelGroup acquisition.
- *Government clearance*—greater than 65 percent have held Secret and Top Secret Department of Defense security clearances.
- *Military backgrounds*—experience logged from USAF Information Warfare Center, Department of Defense (DOD), Department of Energy (DOE), National Security Agency (NSA), Central Intelligence Agency (CIA), or other notable government organizations.
- *Security experience*—average member of the C-CRT has over six years of computer security experience, allowing Cisco to deliver the most mature, accurate, and industry-proven intrusion protection solutions.

Q. Can I create my own signatures?

A. Because the security objectives for each IDS deployment are unique, Cisco IDS adds granularity to the way in which sensors may be tuned to specifically suit the environment in which they are deployed. Using our innovative TAME (Threat Analysis Micro Engine) policy language, users can create new policies or modify existing policies to meet their unique security objectives. Since the TAME policies are decoupled from the sensing application, changes can be made without affecting sensor performance or reliability.

Q. Is it possible to record and replay the IP session of the source IP address that triggered an IDS alarm?

A. IP session logging provides extensive logging that is important for system troubleshooting as well as for reconstructing system events before and after attacks. Cisco IDS augments this capability by converting these logs to a standard TCP dump format that allows them to be viewed and replayed using public domain utilities, such as Ethereal and TCPReplay.

Q. Can the sensor detect attacks if the traffic is encrypted, for example IPsecurity (IPsec) or Secure Sockets Layer (SSL)?

A. The Cisco IDS Sensor analyzes both packet header information (context data) and packet data information (content data) to determine if suspicious activity is occurring. Encryption algorithms encrypt the data portion of the packet for confidentiality. Because it can process only what it can "see," the Sensor cannot detect attacks that require inspection of the payload or data fields within a packet. It will, however, still alarm and respond to attacks, which are detected from the unencrypted packet header information. All network-based IDSs suffer this problem. Therefore, in networks carrying encrypted traffic, Sensor placement is critical. To take advantage of their full intrusion-detection capability, the Cisco IDS Sensors should be installed where the traffic has already been decrypted. Otherwise, the Sensor can be placed on an encrypted segment and will detect all but the packet data or payload-based attacks.

Q. What techniques does Cisco use for mitigating threats?

A. Several techniques provide comprehensive protection against the latest cyber threats, including simple pattern matching, stateful pattern matching, protocol anomaly detection, heuristic-based detection, and anomaly detection.

Q. Does Cisco IDS deliver Peer to Peer signatures?

A. Yes. Cisco IDS delivers protection against file-sharing threats with support for advanced P2P attack mitigation techniques.

18944
ε

Q. How are the Cisco IDS signatures updated?

A. Cisco posts signature updates on Cisco Connection Online (CCO) approximately every 14 days. Cisco IDS provides a facility to automatically distribute new signature files and application upgrades to sensors without operator involvement. Utilizing a secure staging technique, new signature files are placed on a central server and passed to the sensor at scheduled intervals. After verifying the integrity of the package, the sensor automatically installs the update. This new capability significantly streamlines the process of regularly updating remote sensors, thereby lowering the recurring operational costs associated with this task. Additionally, users can subscribe to Cisco Active Update notification services to stay informed about breaking vulnerability news and posted countermeasures at:

http://www.cisco.com/warp/public/779/largeent/issues/security/idsnws/archive.html

Users may refer to the following site for a chronological listing of the Cisco IDS Active Update Notification Bulletins:

http://www.cisco.com/warp/public/779/largeent/issues/security/idsnws/archive.html

Q. How do the Cisco IDS Sensors and management consoles communicate with each other?

A. Communication between the Cisco IDS 4.0 Sensors and management consoles is provided by a secure (SSL) XML based messaging format. All alarm transmissions from the sensor to the management console are acknowledged.

If connectivity from the sensor to the management console is disrupted, the sensor will continue to monitor the network, and will queue alarms and retransmit until successful.

Q. How much additional network traffic does the Cisco IDS generate?

A. Because each alarm and acknowledgment is contained in a single UDP packet, there is negligible impact on network traffic.

Q. Is there a site that lists all the supported Cisco IDS signatures?

A. Yes. Users may access the latest Cisco IDS signatures at the Cisco Secure Encyclopedia site at:

http://www.cisco.com/pcgi-bin/front.x/csec/idsHome.pl

Q. Does Cisco support a centralized site that contains a compiled listing of the latest vulnerabilities?

A. Yes. Cisco's Security Encyclopedia is a one-of-a-kind clearinghouse of security and vulnerability information. Unlike other security databases that simply consolidate vulnerability information published on a number of public-source Web sites, the CSEC contains statistics on the vulnerabilities by industry or by sector. These statistics are compiled from over 400 actual Security Posture Assessments (SPA) performed by the Cisco Security Consulting team. CSEC is developed and maintained by the elite C-CRT. You may visit the CSEC site at:

http://www.cisco.com/go/csec

Q. Where can I download the latest IDS software?

A. Both current and archived IDS Sensor software can be downloaded at the Software Center on CCO (CCO login required):

http://www.cisco.com/public/sw-center/ciscosecure/ids/crypto/

Q. Where can I access documentation on the Cisco IDS Sensor Software?

A. Documentation for sensor software updates are available at:

http://www.cisco.com/univercd/cc/td/doc/product/iaabu/csids/index.htm

18948

Threat Response Technology

Q. Is Threat Response an event correlator?

A. Threat Response is not an event correlator. Event correlation involves analyzing data from NIDS sensor, firewalls, routers, and other sources. Instead of correlating this type of data, Threat Response investigates the actual target of an attack. This is the same process an expert network security specialist would use, and is the best way to determine if a system has been compromised.

Q. Does Threat Response conduct network vulnerability scans?

A. No, Threat Response does not conduct enterprise wide scans of your environment, nor does Threat Response inventory your network. The designers of Threat Response have been network administrators for large mission-critical enterprises such as the U.S. Pentagon and know that downtime is unacceptable. Because of this, Threat Response conducts a low-impact investigation of targeted systems only when needed. Threat Response is able to work in dynamic network environments (including DHCP and wireless) without the need to run regular vulnerability scans that can disrupt your enterprise.

Q. Does Threat Response require deployment of software across the enterprise?

A. No, Threat Response does not require the deployment of software across the enterprise. Threat Response accesses these systems in the same way a security network administrator would—with read access privileges.

Q. How does Threat Response stay up to date with the latest attacks?

A. Cisco releases updates to keep the Threat Response IDS signature database up to date with the IDS vendors, as well as corresponding forensic signature updates to investigate IDS events. When an update is available, the administrator will be notified via the Threat Response GUI and can use the integrated auto-update feature to keep the product current.

Q. What does Threat Response do once real attacks are identified?

A. Threat Response will provide the user with detailed information on how the event was investigated, as well as any forensic data gathered showing details on the actual attack. This information can then be used by an administrator to quickly remediate an intrusion.

Q. What type of systems can Threat Response investigate?

A. Threat Response can conduct a full active investigation of systems running Windows-based operating systems. For systems running Linux, Solaris, and other forms of UNIX, Threat Response will perform passive checks as a first line of investigation. Based on this initial analysis, Threat Response can eliminate many alarms that are not targeted to those specific platform types.

Q. How do I get Threat Response technology?

A. Threat Response technology is currently available as a full featured, 90-day free trial software solution. The trial version ships with every Cisco IDS sensor. Once the trial has expired, customers have the choice of:

- switching to a reduced capabilities free version (only conducts basic level investigation of the targeted system)
- purchasing the full featured version, which will be offered as part of a VMS bundle

18947
3

Management

Q. What management console options are available for the Cisco IDS?

A. See Table 1 and Table 2.

Table 1 Event Management

	IEV	IEV with CTR Technology	VMS SecMon	VMS SecMon with CTR Technology
Deployment method	Dedicated system required	Dedicated system required	Dedicated system required	Dedicated system required
GUI type	Java desktop application	Browser-based GUI	Browser-based GUI	Browser-based GUI
# sensors	5	5	Unlimited	Unlimited
Event types	IDS	IDS	IDS, Firewall, Router	IDS, Firewall, Router

Table 2 Device Management

	IDM	CLI	VMS Management Center
Deployment method	Integrated on sensor	Integrated on sensor	Dedicated system required
GUI type	Browser-based GUI	Browser-based GUI	Browser-based GUI
# sensors	Unlimited, by sensor	Unlimited, by sensor	Unlimited, by sensor groups
Event types	IDS	IDS	IDS, Firewall, Router

IEV = Cisco IDS Event Viewer, included free of charge with IDS sensor
VMS SecMon = CiscoWorks Monitoring Center for IDS, part of the CiscoWorks VMS bundle
IDM = IDS Device Manager, included free of charge with the IDS sensor
CLI = Command-line interface, included free of charge with the IDS sensor
VMS Mgmt Center = CiscoWorks Management Center for IDS, part of the CiscoWorks VMS bundle

For More Information

More information on Cisco's VMS solutions can be found at:

http://www.cisco.com/go/vms

In addition, augmentation to the alarm viewing, analysis, and reporting capabilities of the Cisco IDS Management solution are provided through third-party applications that are available from Cisco Security Associates partners.

For more details see:

http://www.cisco.com/warp/public/778/security/sap/management.html

Q. Are Cisco IDS communications encrypted?

A. IPsec functionality is included on the appliance Sensors to allow customers to encrypt traffic to management consoles with IPsec capabilities.

Q. If I lose connectivity to a remote sensor, can I tell from the management console?

A. The Cisco IDS management consoles monitor the health of a Sensor via a continuous heartbeat. If communications with the Sensor are lost for more than one minute (by default), a visual indicator is displayed on the management console, indicating a communications failure with the Sensor. If it is determined that a sensor has failed, it can be quickly replaced with another sensor and the configuration, stored on the management console, and can be quickly pushed to the new sensor.

Q. How many sensors can one Cisco IDS Management console manage?

A. Although the technical limit is very large (greater than 1000), Cisco typically recommends a ratio of 20 to 25 sensors per management console for practical reasons. With ratios greater than this, operators can be easily overwhelmed with the volume of information that they may be required to analyze, thereby diminishing the overall effectiveness of the IDS. For deployments larger than 25 sensors, multiple management consoles can be installed to scale the number of sensors.

Q. Can I have multiple Cisco IDS management consoles?

A. The Cisco IDS architecture supports the deployment of multiple management platforms. Sensors can send alarms to multiple management consoles simultaneously, and management consoles can forward alarms to other management consoles, allowing customers to build large, hierarchical management infrastructures.

Cisco Systems

Corporate Headquarters
Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA
www.cisco.com
Tel: 408 526-4000
800 553-NETS (6387)
Fax: 408 526-4100

European Headquarters
Cisco Systems International BV
Haarlerbergpark
Haarlerbergweg 13-19
1101 CH Amsterdam
The Netherlands
www-europe.cisco.com
Tel: 31 0 20 357 1000
Fax: 31 0 20 357 1100

Americas Headquarters
Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA
www.cisco.com
Tel: 408 526-7660
Fax: 408 527-0883

Asia Pacific Headquarters
Cisco Systems, Inc.
Capital Tower
168 Robinson Road
#22-01 to #29-01
Singapore 068912
www.cisco.com
Tel: +65 6317 7777
Fax: +65 6317 7799

Cisco Systems has more than 200 offices in the following countries and regions. Addresses, phone numbers, and fax numbers are listed on the **Cisco Web site at www.cisco.com/go/offices**

Argentina • Australia • Austria • Belgium • Brazil • Bulgaria • Canada • Chile • China PRC • Colombia • Costa Rica • Croatia
Czech Republic • Denmark • Dubai, UAE • Finland • France • Germany • Greece • Hong Kong SAR • Hungary • India • Indonesia • Ireland
Israel • Italy • Japan • Korea • Luxembourg • Malaysia • Mexico • The Netherlands • New Zealand • Norway • Peru • Philippines • Poland
Portugal • Puerto Rico • Romania • Russia • Saudi Arabia • Scotland • Singapore • Slovakia • Slovenia • South Africa • Spain • Sweden
Switzerland • Taiwan • Thailand • Turkey • Ukraine • United Kingdom • United States • Venezuela • Vietnam • Zimbabwe

Anexo (26C)

18945

Preface

This document describes how to install IDS appliances and IDS modules and provides basic configuration procedures using the CLI.

This preface contains the following topics:

- Audience, page xiii
- Conventions, page xiv
- Obtaining Documentation, page xv
- Obtaining Technical Assistance, page xvi
- Obtaining Additional Publications and Information, page xviii

Audience

This guide is intended for audiences who need to do the following:

- Install IDS appliances and IDS modules.
- Secure their network with IDS sensors.
- Detect intrusion on their networks and monitor subsequent alarms.

18944
E

Conventions

This guide uses the following conventions:

- Filenames, directory names, and arguments for which you supply values are in italics.
- The symbol ^ represents the key labeled Ctrl (control). To enter a control key; for example, ^z, hold down the Ctrl key while you press the z key.
- Command names, buttons, and keywords in text are shown in boldface.
- Command statements are shown in boldface italic screen font.
- Variables in command syntax descriptions are shown in italics. Command options in square brackets [ ] can be optionally entered, and parameters separated by a vertical bar (|) require you to enter one parameter, but not the other(s).
- Examples depict screen displays and the command line in screen font.
- Information you need to enter in examples is shown in boldface screen font.
- Variables for which you must supply a value are shown in italic screen font.
- Selecting a menu item (or screen) is indicated by the following convention: Select **screen1 > screen 2 > screen3**.

Note Means reader take note. Notes contain helpful suggestions or references to material not covered in the manual.

Caution Means reader be careful. In this situation, you might do something that could result in equipment damage or loss of data.

Warning **This warning symbol means danger. You are in a situation that could cause bodily injury. work on any equipment, be aware of the hazards involved with electrical circuitry and with standard practices for preventing accidents. To see translations of the warnings that in this publication, refer to the Regulatory Compliance and Safety Information document accompanied this device.**

18943
ε

Obtaining Documentation

Cisco provides several ways to obtain documentation, technical assistance, and other technical resources. These sections explain how to obtain technical information from Cisco Systems.

Cisco.com

You can access the most current Cisco documentation on the World Wide Web at this URL:

http://www.cisco.com/univercd/home/home.htm

You can access the Cisco website at this URL:

http://www.cisco.com

International Cisco websites can be accessed from this URL:

http://www.cisco.com/public/countries_languages.shtml

Documentation CD-ROM

Cisco documentation and additional literature are available in a Cisco Documentation CD-ROM package, which may have shipped with your product. The Documentation CD-ROM is updated regularly and may be more current than printed documentation. The CD-ROM package is available as a single unit or through an annual or quarterly subscription.

Registered Cisco.com users can order a single Documentation CD-ROM (product number DOC-CONDOCCD=) through the Cisco Ordering tool:

http://www.cisco.com/en/US/partner/ordering/ordering_place_order_ordering_tool_launch.html

All users can order annual or quarterly subscriptions through the online Subscription Store:

http://www.cisco.com/go/subscription

Ordering Documentation

You can find instructions for ordering documentation at this URL:

http://www.cisco.com/univercd/cc/td/doc/es_inpck/pdi.htm

You can order Cisco documentation in these ways:

- Registered Cisco.com users (Cisco direct customers) can order Cisco product documentation from the Networking Products MarketPlace:

  http://www.cisco.com/en/US/partner/ordering/index.shtml

- Nonregistered Cisco.com users can order documentation through a local account representative by calling Cisco Systems Corporate Headquarters (California, USA.) at 408 526-7208 or, elsewhere in North America, by calling 800 553-NETS (6387).

Documentation Feedback

You can submit comments electronically on Cisco.com. On the Cisco Documentation home page, click **Feedback** at the top of the page.

You can send your comments in e-mail to bug-doc@cisco.com.

You can submit comments by using the response card (if present) behind the front cover of your document or by writing to the following address:

Cisco Systems
Attn: Customer Document Ordering
170 West Tasman Drive
San Jose, CA 95134-9883

We appreciate your comments.

Obtaining Technical Assistance

For all customers, partners, resellers, and distributors who hold valid Cisco service contracts, the Cisco Technical Assistance Center (TAC) provides 24-hour, award-winning technical support services, online and over the phone. Cisco.com features the Cisco TAC website as an online starting point for technical assistance.

18941
ε

Cisco TAC Website

The Cisco TAC website (http://www.cisco.com/tac) provides online documents and tools for troubleshooting and resolving technical issues with Cisco products and technologies. The Cisco TAC website is available 24 hours a day, 365 days a year.

Accessing all the tools on the Cisco TAC website requires a Cisco.com user ID and password. If you have a valid service contract but do not have a login ID or password, register at this URL:

http://tools.cisco.com/RPF/register/register.do

Opening a TAC Case

The online TAC Case Open Tool (http://www.cisco.com/tac/caseopen) is the fastest way to open P3 and P4 cases. (Your network is minimally impaired or you require product information). After you describe your situation, the TAC Case Open Tool automatically recommends resources for an immediate solution. If your issue is not resolved using these recommendations, your case will be assigned to a Cisco TAC engineer.

For P1 or P2 cases (your production network is down or severely degraded) or if you do not have Internet access, contact Cisco TAC by telephone. Cisco TAC engineers are assigned immediately to P1 and P2 cases to help keep your business operations running smoothly.

To open a case by telephone, use one of the following numbers:

Asia-Pacific: +61 2 8446 7411 (Australia: 1 800 805 227)
EMEA: +32 2 704 55 55
USA: 1 800 553-2447

For a complete listing of Cisco TAC contacts, go to this URL:

http://www.cisco.com/warp/public/687/Directory/DirTAC.shtml

TAC Case Priority Definitions

To ensure that all cases are reported in a standard format, Cisco has established case priority definitions.

18940
6

Priority 1 (P1)—Your network is "down" or there is a critical impact to your business operations. You and Cisco will commit all necessary resources around the clock to resolve the situation.

Priority 2 (P2)—Operation of an existing network is severely degraded, or significant aspects of your business operation are negatively affected by inadequate performance of Cisco products. You and Cisco will commit full-time resources during normal business hours to resolve the situation.

Priority 3 (P3)—Operational performance of your network is impaired, but most business operations remain functional. You and Cisco will commit resources during normal business hours to restore service to satisfactory levels.

Priority 4 (P4)—You require information or assistance with Cisco product capabilities, installation, or configuration. There is little or no effect on your business operations.

Obtaining Additional Publications and Information

Information about Cisco products, technologies, and network solutions is available from various online and printed sources.

- The *Cisco Product Catalog* describes the networking products offered by Cisco Systems, as well as ordering and customer support services. Access the *Cisco Product Catalog* at this URL:

  http://www.cisco.com/en/US/products/products_catalog_links_launch.html

- Cisco Press publishes a wide range of networking publications. Cisco suggests these titles for new and experienced users: Internetworking Terms and Acronyms Dictionary, Internetworking Technology Handbook, Internetworking Troubleshooting Guide, and the Internetworking Design Guide. For current Cisco Press titles and other information, go to Cisco Press online at this URL:

  http://www.ciscopress.com

- Packet magazine is the Cisco quarterly publication that provides the latest networking trends, technology breakthroughs, and Cisco products and solutions to help industry professionals get the most from their networking investment. Included are networking deployment and troubleshooting tips,

18939
ε

configuration examples, customer case studies, tutorials and training, certification information, and links to numerous in-depth online resources. You can access Packet magazine at this URL:

http://www.cisco.com/go/packet

- iQ Magazine is the Cisco bimonthly publication that delivers the latest information about Internet business strategies for executives. You can access iQ Magazine at this URL:

  http://www.cisco.com/go/iqmagazine

- Internet Protocol Journal is a quarterly journal published by Cisco Systems for engineering professionals involved in designing, developing, and operating public and private internets and intranets. You can access the Internet Protocol Journal at this URL:

  http://www.cisco.com/en/US/about/ac123/ac147/about_cisco_the_internet_protocol_journal.html

- Training—Cisco offers world-class networking training. Current offerings in network training are listed at this URL:

  http://www.cisco.com/en/US/learning/index.html

18938
ε

Anexo 26D

18937
ε

CHAPTER 1

Preparing for Installation

This chapter provides information you should know before you install Intrusion Detection System (IDS) sensors. In this guide, the term "sensor" refers to all models unless specifically noted otherwise. See Supported Sensors, page 1-2, for a complete list of supported sensors and their model numbers.

- Installation Preparation, page 1-1
- Supported Sensors, page 1-2
- Setting the Time on Sensors, page 1-4
- IDS Appliances, page 1-6
- IDS Modules, page 1-16

Installation Preparation

To prepare for installing IDS sensors, follow these steps:

Step 1 Review the safety precautions outlined in the *Regulatory Compliance and Safety Information for the Cisco Intrusion Detection System 4200 Series Appliance Sensor* that shipped with your IDS sensor.

Step 2 To familiarize yourself with the location of IDS documentation on Cisco.com, read the *Cisco Intrusion Detection System (IDS) Hardware and Software Version 4.1 Documentation Guide* that shipped with your IDS sensor.

Step 3 Obtain the *Release Notes for the Cisco Intrusion Detection System Version 4.1* from Cisco.com and completely read them before proceeding with the installation.

Step 4 Unpack the IDS sensor.

Step 5 Place the IDS sensor on a stable work surface.

Supported Sensors

Table 1-1 lists the sensors (IDS appliances and IDS modules) that are supported in this document and that are supported by the most recent Cisco IDS software.

Note For instructions on how to obtain the most recent Cisco IDS software, see Obtaining Cisco IDS Software, page 8-1.

Caution Installing the most recent Cisco IDS software (version 4.1) on unsupported sensors may yield unpredictable results. We do not support software installed on unsupported platforms.

Table 1-1 Supported Sensors

Model Name	Part Number	Optional Interfaces
IDS Appliances		
IDS-4210	IDS-4210	—
	IDS-4210-K9	—
	IDS-4210-NFR	—
IDS-4215	IDS-4215-K9	IDS-4FE-INT=
	IDS-4215-4FE-K9	—
IDS-4220	IDS-4220-E	—
IDS-4230	IDS-4230-FE	—

18935
ε

Table 1-1 Supported Sensors (continued)

Model Name	Part Number	Optional Interfaces
IDS Appliances		
IDS-4235	IDS-4235-K9	IDS-4FE-INT=
IDS-4250	IDS-4250-TX-K9	IDS-4FE-INT=, IDS-4250-SX-INT=, IDS-XL-INT=
	IDS-4250-SX-K9	IDS-XL-INT=
	IDS-4250-XL-K9	—
IDS Network Modules		
NM-CIDS	NM-CIDS-K9	—
IDS Services Modules		
IDSM-2	WS-SVC-IDSM2-K9	—

Note The IDS-4215-4FE-K9 is the IDS-4215-K9 with the optional 4FE card (IDS-4FE-INT=) installed at the factory.

Note The following IDS appliance models are legacy models and are not supported in this document: NRS-2E, NRS-2E-DM, NRS-2FE, NRS-2FE-DM, NRS-TR, NRS-TR-DM, NRS-SFDDI, NRS-SFDDI-DM, NRS-DFDDI, NRS-DFDDI-DM, IDS-4220-TR, IDS-4230-SFDDI, and IDS-4230-DFDDI.

Note The WS-X6381, the IDSM, is a legacy model and is not supported in this document.

18934
ε

Setting the Time on Sensors

The IDS requires a reliable time source. All events (alerts) must have the correct GMT and local time stamp, otherwise, you cannot correctly analyze the logs after an attack. When you initialize your sensor, you set up the time zones and summer time settings. See Initializing the Sensor, page 9-2, for more information.

Here is a summary of ways to set the time on IDS sensors:

- For IDS appliances
  - Use the **clock set** command to set the time. This is the default.

    Refer to the *Cisco Intrusion Detection System Command Reference Version 4.1* for information on the **clock set** command.
  - Use Network Timing Protocol (NTP).

    You can configure your IDS appliance to get its time from an NTP time synchronization source. You can set up NTP during initialization. You will need the NTP server IP address, the NTP key ID, and the NTP key value.

    Note We recommend that you use an NTP time synchronization source.
- For IDSM-2
  - The IDSM-2 can automatically synchronize its clock with the switch time. This is the default.

    Note The GMT time is synchronized between the switch and the IDSM-2. The time zone and summer time settings are not synchronized between the switch and the IDSM-2.

Caution Be sure to set the time zone and summertime settings on both the switch and the IDSM-2 to ensure that the GMT time settings are correct.

18933

- Use NTP.

  You can configure your IDSM-2 to get its time from an NTP time synchronization source. You can configure the IDSM-2 to use NTP during initialization. You will need the NTP server IP address, the NTP key ID, and the NTP key value.

  Note We recommend that you use an NTP time synchronization source.

- For NM-CIDS
  - The NM-CIDS can automatically synchronize its clock with the router time. This is the default.

    Note The GMT time is synchronized between the router and the NM-CIDS. The time zone and summer time settings are not synchronized between the router and the NM-CIDS.

Caution Be sure to set the time zone and summertime settings on both the router and the NM-CIDS to ensure that the GMT time settings are correct.

  - Use NTP.

    You can configure your NM-CIDS to get its time from an NTP time synchronization source. You can configure the NM-CIDS to use NTP during initialization. You will need the NTP server IP address, the NTP key ID, and the NTP key value.

    Note We recommend that you use an NTP time synchronization source.

18932
ε

IDS Appliances

This section describes the IDS appliance, and contains these topics:

- Introducing the IDS Appliance, page 1-6
- How the IDS Appliance Functions, page 1-7
- Your Network Topology, page 1-8
- Placing an IDS Appliance on Your Network, page 1-10
- Deployment Considerations, page 1-12
- Recommended Keyboards and Monitors, page 1-13
- IDS Appliance Restrictions, page 1-14
- Setting Up a Terminal Server, page 1-14

Introducing the IDS Appliance

The IDS appliance is a high-performance, plug-and-play appliance. The IDS appliance is a component of the Intrusion Detection System (IDS), a network-based, real-time intrusion detection system. See Supported Sensors, page 1-2, for a list of the supported IDS appliances.

You can use the Command Line Interface (CLI), IDS Device Manager, or Management Center for IDS Sensors to configure the IDS appliance. For instructions on configuring IDS accessing IDS documentation on Cisco.com, refer to *Cisco Intrusion Detection System (IDS) Hardware and Software Version 4.1 Documentation Guide* that shipped with your IDS appliance.

You can configure the IDS appliance to respond to recognized signatures as it captures and analyzes network traffic. These responses include logging the event, forwarding the event to the IDS manager, performing a TCP reset, generating an IP log, and/or reconfiguring a router.

After being installed at key points in the network, the IDS appliance monitors and performs real-time analysis of network traffic by looking for anomalies and misuse based on an extensive, embedded signature library. When the system detects unauthorized activity, IDS appliances can terminate the specific connection, permanently block the attacking host, log the incident, and send an alarm to the IDS manager. Other legitimate connections continue to operate independently without interruption.

IDS appliances can also monitor and analyze syslog messages from Cisco routers to detect and report network security policy violations.

IDS appliances are optimized for specific data rates and are packaged in Ethernet, Fast Ethernet, and Gigabit Ethernet configurations. In switched environments, IDS appliances must be connected to the switch's Switched Port Analyzer (SPAN) port or VLAN Access Control list (VACL) capture port.

How the IDS Appliance Functions

The next step in protecting your network is understanding how the IDS appliance captures network traffic.

Each IDS appliance comes with two interfaces. In a typical installation, one interface monitors (sniffs) the desired network segment, and the other interface (command and control) communicates with the IDS manager and other network devices. The monitoring interface is in promiscuous mode, meaning it has no IP address and is not visible on the monitored segment.

Note With the addition of the 4-port Fast Ethernet NIC card, the IDS-4235, IDS-4250, and the IDS-4215 have six interfaces. With the addition of the 2-port XL card, the IDS-4250 has four interfaces. With the addition of the SX card, the IDS-4250 has three interfaces.

The command and control interface is always Ethernet. This interface has an assigned IP address, which allows it to communicate with the IDS manager workstation or network devices (typically a Cisco router). Because this interface is visible on the network, you should use encryption to maintain data privacy. Secure Shell (SSH) is used to protect the Command Line Interface (CLI) and the Transaction Layer Security/Secure Sockets Layer (TLS/SSL) is used to protect the IDS manager workstation. Both SSH and TLS/SSL are enabled by default on the IDS manager workstations.

When responding to attacks, the IDS appliance can do the following:

- Insert TCP resets via the monitoring interface.

  Note The TCP reset action is only appropriate as an action selection on those signatures that are associated with a TCP-based service. If selected as an action on non-TCP-based services, no action is taken. Additionally, TCP resets are not guaranteed to tear down an offending session because of limitations in the TCP protocol. On the IDS-4250-XL, TCP resets are sent through the TCP Reset interface.

- Make access control list (ACL) changes on routers that the IDS appliance manages.

  Note ACLs may block only future traffic, not current traffic.

- Generate IP session logs

  IP session logs are used to gather information about unauthorized use. IP log files are written when a certain event or events occur that you configure on the sensor.

The last step in understanding how a IDS appliance functions is the data speed or load on the monitored network. Because the IDS appliance is not in the data path, it has a negligible impact on network performance. However, there are limitations on the data speeds it can monitor.

Your Network Topology

Before you deploy and configure your IDS appliances, you should understand the following about your network:

- The size and complexity of your network.
- Connections between your network and other networks (and the Internet).
- The amount and type of network traffic on your network.

This knowledge will help you determine how many IDS appliances are required, the hardware configuration for each IDS appliance (for example, the size and type of network interface cards), and how many IDS managers are needed.

18929
3

The IDS appliance monitors all traffic across a given network segment. With that in mind, you should consider all the connections to the network you want to protect. These connections fall into four categories, or locations, as illustrated in Figure 1-1.

Figure 1-1 Major Types of Network Connections

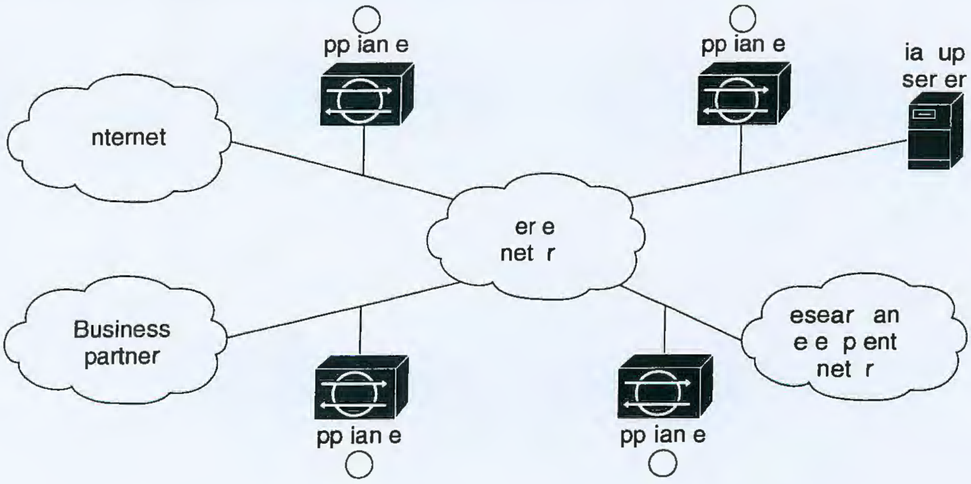

In location one, the IDS appliance is placed to monitor traffic between the protected network and the Internet. This is referred to as perimeter protection and is the most common deployment for an IDS appliance. This location can be shared with firewall protection and is discussed in Placing an IDS Appliance on Your Network. page 1-10.

In location two, the IDS appliance is monitoring an extranet connection with a business partner. Although most companies have defined policies on the use and security of this type of connection, there is no guarantee that the network of a partner is adequately protected. Consequently, an outsider may enter your network through this type of connection. These extranet connections may have firewalls as well.

In location three, the IDS appliance is monitoring the network side of a remote access server. Although this connection may be only for employee use, it could be vulnerable to external attack.

18928
ε

In location four, the IDS appliance is monitoring an intranet connection. For example, the protected network of one department may contain an e-commerce site where all the access types described so far are required. The network of another department may contain company-specific research and development or other engineering information and should be given additional protection.

Determine the type of location you have to determine which segments of the network you want to monitor. Remember, each IDS appliance maintains a security policy configured for the segment it is monitoring. The security policies can be standard across the organization or unique for each IDS appliance. You may consider changing your network topology to force traffic across a given monitored network segment. There are always operational trade-offs when going through this process. The end result should be a rough idea of the number of IDS appliances required to protect the desired network.

Placing an IDS Appliance on Your Network

You can place an IDS appliance in front of or behind a firewall. Each position has its benefits and drawbacks.

Placing an IDS appliance in front of a firewall allows the IDS appliance to monitor all incoming and outgoing network traffic. However, when deployed in this manner, the IDS appliance does not normally detect traffic that is internal to the network. An internal attacker taking advantage of vulnerabilities in network services would remain undetected by the external IDS appliance (see Figure 1-2).

18927
ε

Figure 1-2 IDS Appliance in Front of a Firewall

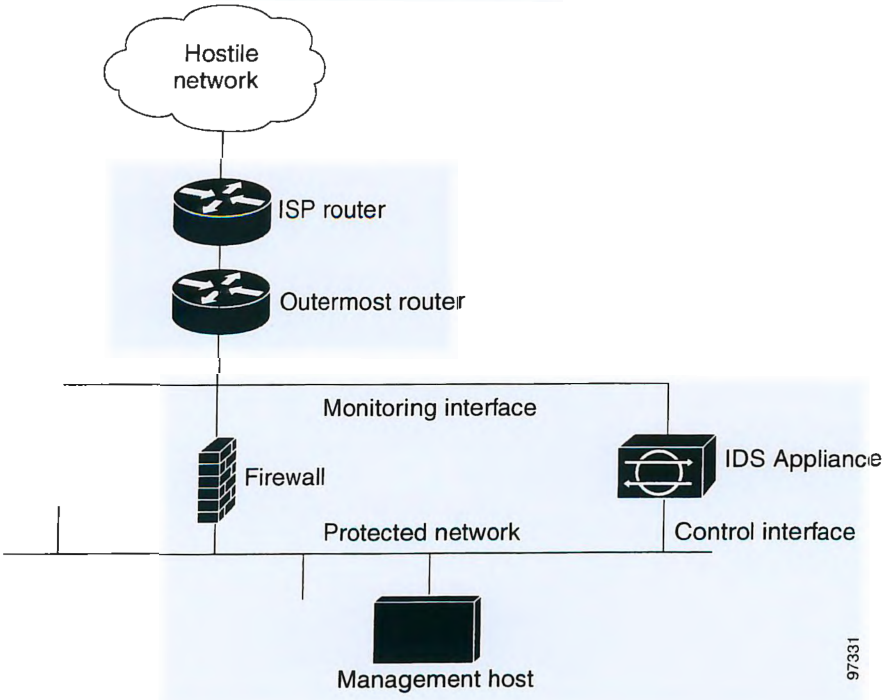

P:acing an IDS appliance (a monitoring or sniffing interface) behind a firewall shields the IDS appliance from any policy violations that the firewall rejec:s (see Figure 1-3).

Figure 1-3 IDS Appliance Behind a Firewall

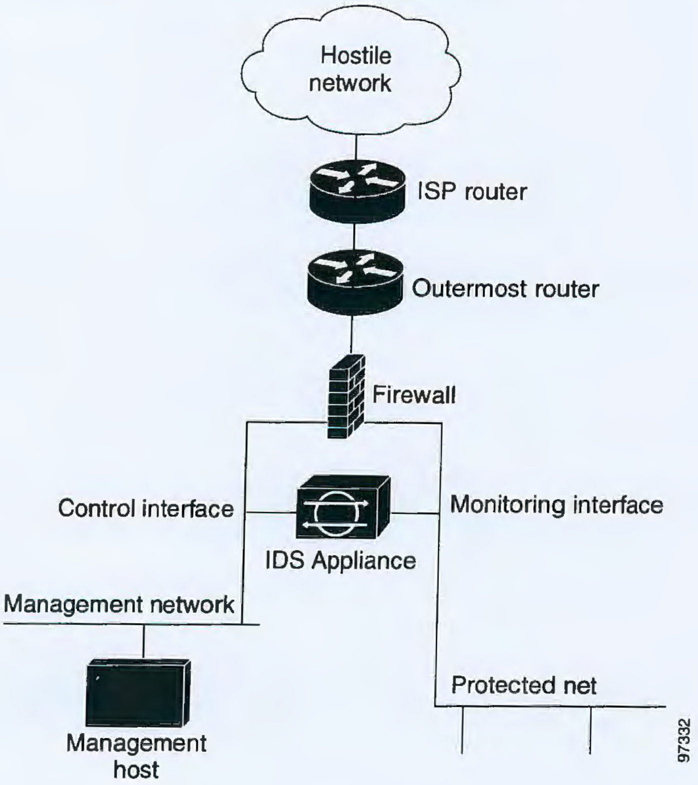

Deployment Considerations

For the IDS appliance to effectively defend a network with a router and firewall configuration, you must do the following:

- Enable SSH services on the router if available, otherwise, enable Telnet.
- Add the router to the device management list of the IDS appliance (via the IDS manager).

18925
E.

- Configure the firewall to permit the following traffic:
  - SSH or Telnet traffic from the control interface of the IDS appliance to the router.
  - Syslog (UDP port 514) traffic from the router to the IDS appliance.

    Note To capture policy violations on the router, the IDS appliance must also be configured to accept syslog messages.

  - Communications (TCP ports 443 for TLS/SSL and 22 for SSH) between the IDS appliance and any IDS manager workstation, if the firewall comes between them.

  Essentially, the firewall implements policy filtering. The IDS appliance captures packets between the Cisco router and the firewall, and can dynamically update the ACLs of the Cisco router to deny unauthorized activity.

  Note You can also configure the sensor to manage a PIX Firewall instead of the Cisco router.

Recommended Keyboards and Monitors

Some keyboards and monitors are not compatible with the IDS appliance. This incompatibility could cause the IDS appliance to boot improperly. We recommend that you use the following keyboards and monitors with the IDS appliance:

Note You can also use a serial cable to connect to the sensor's console port. You cannot use a monitor and keyboard with the IDS-4215.

- Keyboards
  - KeyTronic E03601QUS201-C
  - KeyTronic LT DESIGNER

18924

- Monitors
  - MaxTech XT-7800
  - Dell D1025HT

Caution The IDS appliance does not function properly with some HP keyboards and with IBM model G50 monitors.

IDS Appliance Restrictions

The following restrictions apply to the use and operation of the IDS appliance:

- The IDS appliance is not a general purpose workstation.
- Cisco Systems prohibits using the IDS appliance for anything other than operating Cisco IDS.
- Cisco Systems prohibits modifying or installing any hardware or software in the IDS appliance that is not part of the normal operation of the Cisco IDS.

Setting Up a Terminal Server

A terminal server is a router with multiple, low speed, asynchronous ports that are connected to other serial devices. Terminal servers can be used to remotely manage network equipment, including sensors.

To set up a Cisco terminal server with RJ-45 or hydra cable assembly connections, follow these steps:

Step 1 Connect to a terminal server using one of the following methods:

- For the IDS-4215:
  - For RJ-45 connections, connect a 180/rollover cable from the console port on the sensor to a port on the terminal server.
  - For hydra cable assemblies, connect a straight-through patch cable from the console port on the sensor to a port on the terminal server.

18923 ε

- For all other sensors, connect the M.A.S.H. adapter (part number 29-4077-01) to COM1 on the sensor and:
  - For RJ-45 connections, connect a 180/rollover cable from the M.A.S.H. adapter to a port on the terminal server.
  - For hydra cable assemblies, connect a straight-through patch cable from the M.A.S.H. adapter to a port on the terminal server.

Step 2 Configure the line/port on the terminal server as follows:

a. In enable mode, enter the following configuration, where # is the line number of the port to be configured:

```
config t
line #
login
transport input all
stopbits 1
flowcontrol hardware
exit
exit
wr mem
```

b. If you are configuring a terminal server for an IDS-4215, skip to Step 3.

Otherwise, for all other supported IDS appliances, to direct all output to the terminal server, log in to the IDS CLI and type the following commands:

```
sensor# configure terminal
sensor(config)# display-serial
```

Output is directed to the serial port. Use the **no display-serial** command to re-direct output to the keyboard/monitor.

Note You can set up a terminal server and use the IDS CLI **display-serial** command to direct all output from the sensor to the serial port. This option allows you to view system messages on a console connected to the serial port, even during the boot process. When you use this option, all output is directed to the serial port and any local keyboard/monitor connection is disabled. However, BIOS and POST messages are still displayed on the local keyboard/monitor.

18922
ε

Note There is only one console port on an IDS-4215; therefore, the **display-serial** and **no display-serial** commands do not apply to that platform.

Step 3 Be sure to properly close a terminal session to avoid unauthorized access to the sensor.

If a terminal session is not stopped properly, that is, if it does not receive an exit(0) signal from the application that initiated the session, the terminal session can remain open. When terminal sessions are not stopped properly, authentication is not performed on the next session that is opened on the serial port.

Tip Always exit your session and return to a login prompt before terminating the application used to establish the connection.

Note If a connection is dropped or terminated by accident, you should re-establish the connection and exit normally to prevent unauthorized access to the sensor.

IDS Modules

This section describes the IDS modules, and contains these topics:

- Introducing the Cisco Intrusion Detection System Network Module, page 1-17
- Introducing the Cisco Catalyst 6500 Series Intrusion Detection System Services Module, page 1-19

18921
3

Introducing the Cisco Intrusion Detection System Network Module

The Cisco Intrusion Detection System Network Module (NM-CIDS) integrates the Cisco IDS functionality into the branch office router. With the NM-CIDS, you can implement full-featured IDS at your remote branch offices. You can install the NM-CIDS in any one of the network module slots on the Cisco 2600, 3600, and 3700 series routers. The NM-CIDS provides up to 45 Mbps. See Software and Hardware Requirements, page 6-2, for a list of supported routers. Only one NM-CIDS is supported per router and is not hot-swappable. Figure 1-4 shows the IDS router in a branch office environment.

Figure 1-4 NM-CIDS in the Branch Office Router

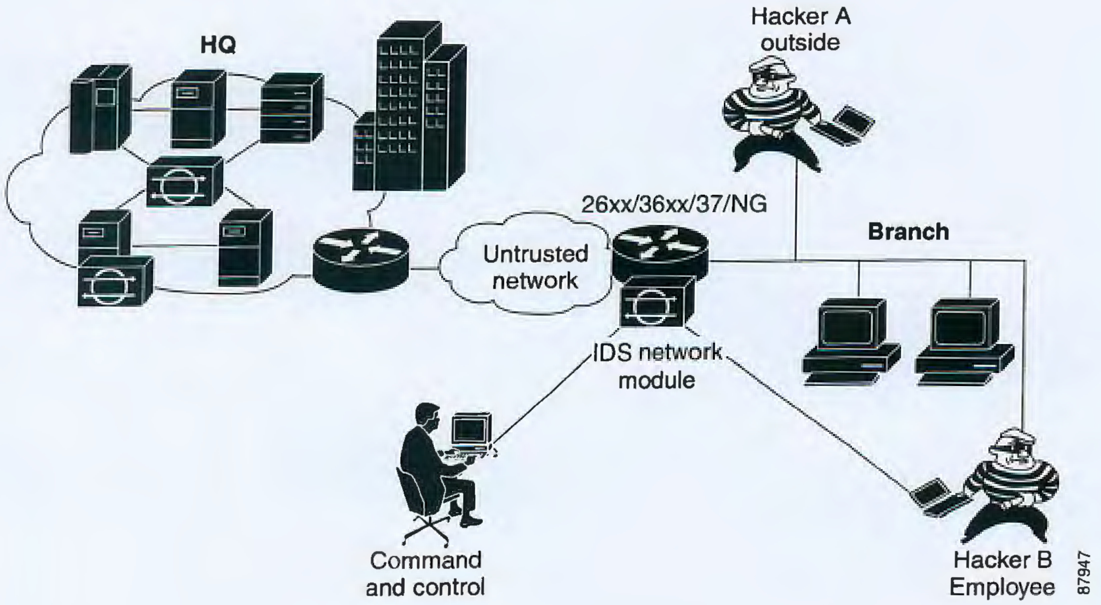

The NM-CIDS has one internal 10/100 Ethernet port that connects to the router's backplane. There is also one external 10/100-based Ethernet port that is used for device management (management of other routers and/or PIX Firewalls to perform shunning) and command and control of the NM-CIDS by the IDS managers.

The NM-CIDS communicates with the router to exchange control and state information for bringing up and shutting down the NM-CIDS and for exchanging version and status information with the router. The NM-CIDS processes packets that are forwarded from the selected interfaces on the router to the IDS interface on the NM-CIDS. The NM-CIDS analyzes the captured packets and compares them against a rule set of typical intrusion activity called signatures. If the captured packets match a defined intrusion pattern in the signatures, the NM-CIDS can take one of two actions: it can make ACL changes on the router to block the attack, or it can send a TCP reset packet to the sender to stop the TCP session that is causing the attack.

In addition to the packet capture that analyzes the traffic to identify malicious activity, the NM-CIDS can also perform IP session logging that can be configured as a response action on a per-signature basis. When the signature fires, session logs are created over a pre-specified time period in a TCPDump format. You can view these logs using Ethereal or replay the IP session using tools such as TCP Replay.

Note The NM-CIDS does not support sending syslog messages to a syslog server if there is an intrusion event, nor does it support Simple Network Management Protocol (SNMP) traps.

You can manage and retrieve events from the NM-CIDS through the CLI or through one of these IDS managers—IDS Device Manager or Management Center for IDS Sensors. For instructions on accessing IDS documentation on Cisco.com, refer to *Cisco Intrusion Detection System (IDS) Hardware and Software Version 4.1 Documentation Guide* that shipped with your NM-CIDS.

The IDS requires a reliable time source. All the events (alarms) must have the correct time stamp, otherwise, you cannot correctly analyze the logs after an attack. You cannot set the time on the NM-CIDS. The NM-CIDS gets its time from the Cisco router. Routers do not have a battery so they cannot preserve a time setting when they are powered off. You must set the router's clock each time you power up or reset the router, or you can configure the IDS router to use NTP time synchronization. We recommend NTP time synchronization. See your Cisco IOS documentation for more information on setting the time on your Cisco router.

18919
ε

Introducing the Cisco Catalyst 6500 Series Intrusion Detection System Services Module

The Cisco Catalyst 6500 Series Intrusion Detection System Services Module (IDSM-2) is a switching module that is easy to install and maintain in the Catalyst 6500 series switch. It is part of the IDS. You can use the CLI, IDS Device Manager, or Management Center for IDS Sensors to configure the IDSM-2. For instructions on accessing the IDS documentation on Cisco.com, refer to the *Cisco Intrusion Detection System (IDS) Hardware and Software Version 4.1 Documentation Guide* that shipped with your IDSM-2.

The IDSM-2 performs network sensing—real-time monitoring of network packets through packet capture and analysis. The IDSM-2 captures network packets and then reassembles and compares the packet data against attack signatures indicating typical intrusion activity. Network traffic is either copied to the IDSM-2 based on security VLAN access control lists (VACLs) in the switch or is copied to the IDSM-2 through the switch's Switched Port Analyzer (SPAN) port feature. Both methods permit user-specified traffic based on switch ports, VLANs, or traffic type to be inspected. (See Figure 1-5.)

Figure 1-5 IDSM-2 Block Diagram

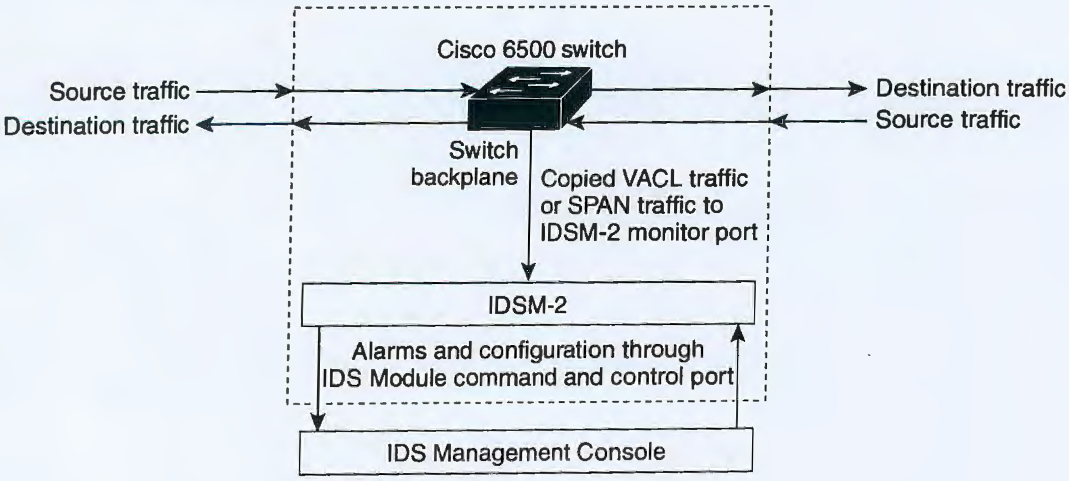

18918
ε

The IDSM-2 searches for patterns of misuse by examining either the data portion and/or the header portion of network packets. Content-based attacks derive from the data portion, and context-based attacks derive from the header portion.

After the IDSM-2 detects an attack, it generates an alarm. Additionally, you can configure the IDSM-2 to transmit TCP resets on the source VLAN, generate an IP log, and/or initiate blocking countermeasures on a firewall or other managed device. Alarms are generated by the IDSM-2 through the Catalyst 6500 series switch backplane to the IDS manager, where they are logged or displayed on a graphical user interface.

ANEXO (262)

CISCO SYSTEMS

Home | Log In | Register | Contacts & Feedb

Select a

Products & Services GO

18917

PRODUCTS & SERVICES
CISCO SECURITY AND VPN SOFTWARE
CISCO INTRUSION DETECTION SYSTEM
PRODUCT LITERATURE
BUSINESS CASES
Ease of Management

Search:

Search A

Toolkit:

CISCO INTRUSION DETECTION SYSTEM

Ease of Management

Ease of Management

Cisco provides effective security monitoring and configuration regardless of deployment size using a range of management options. All management tools are designed with an intuitive user interface, and easy navigation enabling rapid installation, configuration, and management of security events and devices.

Intuitive Event Display

- **Secure, browser-based GUI**--Alarms can be easily viewed from practically any desktop, no matter which operating system is being used on the desktop. The result is rapid access to data from systems throughout the enterprise. The familiar browser interface enhances usability. And with Secure Sockets Layer (SSL), security of data is maintained.
- **Unified, scalable view of all security events**--With the CiscoWorks VPN/Security Management Solution (CiscoWorks VMS), events from all types of security devices, including firewall, VPN, and IDS can be viewed from a single console. Multiple data sources can be supported and managed. This enhances the ability to view security across the enterprise.

Easy Alarm Processing

- **Forensic data**--With Cisco Threat Response technology, the GUI provides a view into the steps taken to investigate and confirm intrusion events. To aid in remediation of intrusions, forensic data collected by the Threat Response technology is accessible. Examples include Web logs, system logs, and other relevant data. (This capability is available only as part of a CiscoWorks VMS technology bundle.)
- **Correlation of events**--CiscoWorks VMS provides event correlation to enable improved confidence in alarm data by corroborating data from multiple security devices.
- **Network security database**--The network security database (NSDB) provides instant access to specific information about the attacks, hot links, potential countermeasures, and related vulnerabilities Because the NSDB is an HTML database, it can be personalized for each user to include operation-specific information such as response and escalation procedures for specific attacks.

Flexible Reporting and Notification

- **Default reports**--Reporting allows the creation and viewing of reports about network activities monitored by sensors on your network. The reports include summary reports based on alarms, sources, or destinations. Because these reports are HTML based, they can be sent in e-mail messages to key administration personnel.
- **Custom reporting**--Custom reports can be created to meet the specific needs of your environment.

Simple Configuration

- **Wizard-based configuration**--Wizards guide the user through the configuration process, allowing sensors to be configured quickly and easily.

ANEXO 26F

18915
E

CISCO SYSTEMS

White Paper

Capacity Verification for High-Speed Network Intrusion Detection Systems

Abstract. Commercially available network intrusion detection systems (NIDS) came onto the market over six years ago. These systems have gained acceptance as a viable means of monitoring the security of consumer networks, yet no commercial standards exist to help consumers understand the capacity characteristics of these devices. Existing NIDS tests are flawed. These tests resemble the same tests used with other networking equipment, such as switches and routers. However, switches and routers do not conduct the same level of deep packet inspection, nor do they require the higher-level protocol awareness that a NIDS demands. Therefore, the current testing does not allow consumers to infer any expected performance in their environments. Designing a new set of tests that specific to the weak areas, or bottlenecks, of a NIDS is the key to discovering metrics meaningful to the consumers. Consumers of NIDS technology can then examine the metrics used in the tests and profile their network traffic based on these same metrics. Consumers can use standard test results to accurately predict performance on their networks. This paper proposes a test methodology for standardized capacity benchmarking of NIDS. The test methodology starts with examination of the bottlenecks in a NIDS, then maps these bottlenecks to metrics that can be tested, and finally explores some results from tests conducted.

Introduction and Scope

Currently, no industry standards exist for testing any aspect of network intrusion detection systems (NIDS). The NIDS industry is maturing along the same lines as the router, switch, and firewall industries that came before it, and has now reached the point where standardization of testing and benchmarking is possible. Attempting to define a testing standard is beyond the scope of this paper. Instead, the metrics and methodology used to properly verify the capacity of high-speed NIDS are explored. Performance of NIDS is usually defined by false positive and false negative ratios, and speed or capacity. This paper addresses the issue of benchmarking the capacity of a NIDS. This paper uses capacity to refer to the ability of a NIDS *to capture, process, and perform at the same level of accuracy under a given network load as it does on a quiescent network.*

Gauging the capacity of a NIDS is difficult. Several variables in the characteristics of the network traffic affect the performance of a NIDS. The last year has seen claims of NIDS performing at or near gigabit speeds. In every case, however, further investigation by reasonably sophisticated NIDS practitioners revealed critical flaws in the testing methodology.

The variety of technology used to perform network-based intrusion detection further complicates finding the proper metrics. The following technologies are used for NIDS:

- Stateless inspection of the packets or packet headers
- Protocol decode and analysis
- Regular expression matching of packet data
- Anomaly detection

- **Automatic updates**--Automatic update capabilities maintain the effectiveness of the intrusion protection system, and simplify the regular maintenance.
- **Remote management**--Because you are not always at the same computer, or where the IDS system is located, remote access capabilities using a secure Web browser connection allow for easy remote connectivity.

18916
ε

Scalable Enterprise Management

- **Multi-tiered architecture**--CiscoWorks VMS promotes a three-tiered architecture providing enhanced scalability needs to enterprise security deployments.
- **Flexible device grouping**--Easily manage large IDS deployments by grouping devices by function, location, or by configuration to perform mass configuration changes.
- **Roles-based access control**--Control administrative access to ensure proper device authorization.
- **Tiered approval model (optional)**--Separate configuration definition and deployment authorities to provide proper audit and control.

For more details on the CiscoWorks VMS solution go to: www.cisco.com/go/vms.

BUSINESS INDUSTRIES & SOLUTIONS | NETWORKING SOLUTIONS & PROVISIONED SERVICES | PI
TECHNOLOGIES | ORDERING | TECHNICAL SUPPORT | LEARNING & EVENTS | PARTNERS & RESEL
Home | Log In | Register | Contacts & Feedback | Help | Site Map

18914
3

Most NIDS employ a mix of all these methods. Some of the metrics discussed in this paper do not apply to all the technologies. Choosing metrics and test methods valid for all NIDS in existence is impossible. Choosing a broad set of metrics that is generally applicable to most NIDS is possible. This paper focuses on two questions: What are the proper metrics for performance testing? What testing methodology best evaluates these metrics? The testing metrics and methodology described are intended for use on a NIDS located at the edge of an internal network functioning near the firewall or border router. The focus is further refined by looking at how these metrics apply to a NIDS using a combination of the technologies listed previously. However, many of the metrics and methods included also apply to the performance of a NIDS inside the core of an enterprise network and to a NIDS employing other methods of detecting intrusions such as pure anomaly-based systems.

History

Most NIDS capacity benchmarks to date have been run by independent third parties either for publication in a trade magazine or at the request of the vendor for inclusion in marketing material. The test methodologies were developed based on experiences in the router- and switch-testing arenas. These tests are generally not adequate for the purposes of developing a NIDS performance profile because the benchmark tests for switch and router capacity often forward packets of various sizes without regard for any protocol above IP or even the validity of the packets used. Although routers and switches are typically not concerned with Layer 4 and above, NIDS may discard packets that are not interesting. A NIDS also needs to look much deeper in a packet than a switch or a router to follow Layer 4 and above. For example, a NIDS may discard TCP streams that are not opened using a valid three-way handshake. If a switch or router test is used, most of the traffic might be ignored. The NIDS then performs very little deep packet inspection.

Because the results of a NIDS performance test based on these types of test methodologies are often skewed in the favor of the vendor, a consumer may believe these results are valid for a deployment and encounter strikingly different performance characteristics after the NIDS is fielded on the network.

For example, NIDS tests to date from Mier Communications are flawed [1]. Mier Labs concluded that two different NIDS could perform at gigabit lines rates. Although the lab report is technically accurate, there is no mention anywhere in the report that a test using TCP port 0 packets was not representative of the performance most consumers would experience. Using this type of testing methodology for NIDS is, therefore, flawed. Marcus Ranum also mentions a few other flawed testing methodologies in "Experiences Benchmarking Intrusion Detection Systems" [2]. Ranum does an excellent job explaining why benchmarking NIDS is difficult.

Defining the Metrics

Defining metrics for any type of testing is difficult. For example, the materials used in bridge construction need to be tested to ensure the integrity of the bridge. A common approach to defining metrics for these types of tests involves asking the bridge engineer to identify the weak spots in the bridge design. Where is the most stress concentrated? What has the highest potential for failure when the load exceeds design specifications? The same premise can be used for defining the metrics used for testing the performance of NIDS. The stress points for most protocol decode and pattern matches in a NIDS are the same.

Fixed Resource Limitations

All computing devices have a similar list of fixed resources. Only so many cycles are available on the CPU, and only so many bits are available to store the program code, state information, and runtime conditions. In addition, only so many transmission cycles are available on the various buses of the computing architecture. The upper limit for system performance is approached as one or more of these resources approaches its upper limit. Therefore, the test methods should include metrics that apply to these resources. The following resources are the most important to the performance of NIDS:

- Memory size
- Memory bus bandwidth
- Memory latency
- Bus bandwidth for the network interfaces (Ethernet card)
- Persistent storage bandwidth (hard drive, Flash memory, etc.)
- CPU speed and bandwidth

Packet Capture Architecture

NIDS products monitor network traffic, and NIDS packet capture architectures impose physical limits on the type of traffic that can be observed. For example, a NIDS built with a standard Gigabit Ethernet card cannot observe all minimum-sized Ethernet frames sent at line rate. This equates to approximately 1.4 million packets per second, and no currently shipping standard gigabit network adapter can handle many more than 700 to 800 thousand packets per second. However, if the NIDS uses dedicated hardware or some network processing units (NPUs), then it can handle more than 1.4 million packets. Also, the host software platform for the NIDS can have significant impact on the ability of the NIDS to capture packets. Many NIDS running on a host operating system do not bind the network interface to the IP stack of an operating system, and the architectures include custom network interface drivers.

Many of the more recently published NIDS performance tests actually tested only the interface bandwidth of the NIDS. This type of testing has limited use because it shows only the upper limit of how the NIDS performs if no other fixed resources are used. Typically this type of test lets the consumer understand how quickly a NIDS can ignore traffic that is uninteresting. Knowing the performance of only the packet capture architecture is useful, but it does not provide the information needed to quantify the performance of the entire system.

The metrics that affect the performance of the packet capture architecture include packets per second, bytes (or bits) per second, and physical network interface.

Packet Flow Architecture

Packet flow architecture is the overall architecture for data flow within the NIDS and includes the packet capture architecture. The metrics used in the packet capture architecture section are also valid for the packet flow architecture, assuming the packets used are of interest to the NIDS, that they cause deep packet inspection, and that they make proper use of protocols of interest to the NIDS. Using Hypertext Transfer Protocol (HTTP) traffic to test the packet flow architecture is generally a good choice. For a NIDS that employs some method of protocol decode and state aware reassembly, HTTP traffic flows through a major portion of the packet flow architecture.

In addition, not all packets take the same amount of time to process. Buffering in the packet flow architecture allows a NIDS to recover from the packets that take a long time to inspect. Packet buffering is an important feature for reducing the number of dropped packets. Therefore, when testing a NIDS with buffering in the packet flow architecture, it is important to test with sustained rates for a length of time to ensure that the buffer is not inflating performance.

State Tracking

Any NIDS that performs TCP state tracking, IP fragment reassembly, and detection of sweeps and floods must keep track of the state of the traffic on the network. Many of the signatures used in this type of NIDS are based on a certain threshold of events that occur within a specified period of time. The only way to assess event thresholds is to keep a count until the time has expired or the threshold has been exceeded. A NIDS must track information about both the source and destination of the packets, the state of TCP connections, and the type and number of packets going to or from the hosts. All this information must be stored somewhere within the NIDS. The storage medium for this information is the *state database.*

Database benchmarking is very mature. The database industry understands the weak points in database-like applications. The largest metrics include the time needed to insert, delete, and search through data and how those transaction times scale with the size of the data set and frequency of transactions. How do we correlate those database metrics to NIDS metrics?

State information must be inserted into the state database as new network hosts and unique connections are observed. The state information is typically removed from the database after either an alarm event has occurred or some time has elapsed. State database searches are conducted any time the incoming packet may need to refer to prior state information. Depending on the types of signatures used, searching the database may need to be done for each packet.

The size of the state database derives from the unique hosts and connections that the NIDS considers interesting and maintains prior information about. The following metrics directly affect the performance of the state database:

- Number of unique hosts
- Number of new connections over time (that is, TCP connections per second)
- Number of concurrent connections at any given time
- Efficiency of expiring data in the database

The connection duration in number of packets or time can be used as an indirect metric for testing the performance of the state database because the duration of a session is related to the number of connections over time as well as the number of concurrent connections. Therefore, to accurately measure the capacity of a NIDS, one must vary the number of new connections per second, the number of simultaneous open connections, and the total number of unique hosts that the NIDS must track. The ability of the NIDS to handle network loads varies as these variables are adjusted.

18912
e

Packet Analysis

Memory bandwidth and memory latency are large factors in the performance of a NIDS. Much of the memory bandwidth use and latency are caused by access to memory while inspecting the packets. Different NIDS architectures exhibit different use patterns for memory. A NIDS that relies solely on regular expression matching consumes the most bandwidth and induces the most latency in the system. Inspecting each character in the packet payload and advancing a regular expression state are expensive operations.

Protocol analysis helps reduce the number of bytes that must be inspected. Every NIDS does some type of protocol decode, even if it is limited to just the IP header. Many of the commercial NIDS decode most Layer 7 protocols and perform regular expression inspection on only a subset of the entire packet.

The size of the packets, therefore, plays an important role in determining the capacity of a NIDS. Testing performance with the smallest possible average packet size reduces the amount of time available per packet for inspections. Increasing the average packet size allows more time for inspection and increases the use of memory.

The average packet size of typical traffic on the Internet is about 450 to 550 bytes per packet [3,4]. However, some networks contain averages much larger and much smaller. The average packet size is an important metric in capacity testing.

Event or Alarm Reporting

The generation of the alarm event expends CPU cycles that would otherwise be available for analysis. Additionally, the event needs to be stored in nonvolatile storage. This usually means that it must be written to disk, a relatively slow operation, or sent over a network connection. Under normal circumstances this does not affect the operation of a NIDS. However, as the rate of alarm production increases and/or the load on the network increases, alarm event production and log maintenance can have a significant effect on NIDS performance. The event generation component of a NIDS must be able to handle the events generated by the high rates of traffic. The ability of the NIDS to notify the user varies as the alarm event rate is adjusted.

The metric used to test this component of a NIDS is simply the number of alarms per second. Tools such as stick and nessus easily set off alarm events in NIDS products. In addition, packet generators can be used to generate single packets that cause an alarm event. Testing the alarm channel does not require the traffic causing the alarms to originate from real hosts.

The Metrics

With the major stress points of a NIDS identified, it is now possible to focus on defining the metrics that can be used to quantify the capacity of a NIDS. Table 1 defines the test metrics and how they are related to the use of the fixed resources described in the section "Fixed Resource Limitations."

Table 1

NIDS Test Metrics and Corresponding Resources Used

Test metrics	Resources used
Packets per second	CPU cycles, network interface bandwidth, and memory bus bandwidth
Bytes per second (average packet size)	CPU cycles, network interface bandwidth, and memory bus bandwidth
Protocol mix	CPU cycles and memory bus bandwidth

18911
ε

NIDS Test Metrics and Corresponding Resources Used

Test metrics	Resources used
Number of unique hosts	Memory size, CPU cycles, and memory bus bandwidth
Number of new connections per second	CPU cycles and memory bus bandwidth
Number of concurrent connections	Memory size, CPU cycles, and memory bus bandwidth
Alarms per second	Memory size, CPU cycles, and memory bus bandwidth

In practice, all these metrics are related. For example, a network does not typically raise only the number of hosts without also increasing the packets and bytes per second. However, to the extent possible, the tests should attempt to stress only one component or metric of the NIDS at a time. Prospective consumers of a NIDS can then use the test results for each metric and a profile of their networks to determine if the NIDS is even capable of sustaining inspection of the traffic.

Developing the Tests

Developing the tests to quantify the metrics for the potential weak points of a NIDS is a significant task. This section explores traffic mix selection and qualification, potential problems with a test network, and a set of tests and the intended stress points under test.

Traffic Mix

Many of the metrics defined in the previous section are directly and indirectly derived from the network traffic. What is the correct mix of traffic to measure? The correct mix for each user is one that best matches the traffic where that user plans to deploy the NIDS. Obviously a test cannot be designed that contains all the traffic mixes for all potential consumers of NIDS technology. However, after the NIDS industry agrees on a standard methodology for testing the stress points, consumers could profile their traffic mix and get a reasonable idea about how well the various products perform.

How do we define the tests to be used? Studies have been performed that describe the mix of traffic seen on the major network trunks. "The Nature of the Beast: Recent Traffic Measurements from an Internet Backbone" [3] and "Trends in Wide Area IP Traffic" [4] are two good resources for defining a general test. The information in these papers is from 1998 and 2000. For more recent data, Table 2 includes results from profiling three data sets from 2002. A major university, a major U.S. government site, and a large online retailer provided the data sets. Although this data is not necessarily representative of a traffic mix found on a corporate network, it is representative of the mix that would be seen at the edge of most large networks. The metrics in this traffic provide a good starting point for the mix of the loading traffic for the test.

Table 2

Traffic Metrics from Three Customer Sites (Site 1 is a major university, Site 2 is a major U.S. government site, and Site 3 is an online retailer. In the Layer 4 OTHER field, no individual protocol grouped into this field consisted of more than 3 percent of the total traffic.

	Site 1	Site 2	Site 3
Average packet size	543	501	557
Average bandwidth	25.0 Mbps	36.2 Mbps	31.4 Mbps

18910
ε

Traffic Metrics from Three Customer Sites (Site 1 is a major university, Site 2 is a major U.S. government site, and Site 3 is an online retailer. In the Layer 4 OTHER field, no individual protocol grouped into this field consisted of more than 3 percent of the total traffic.

	Site 1	Site 2	Site 3
Layer 3			
Percent TCP	94.8	97.6	98.7
Percent User Datagram Protocol (UDP)	5.0	1.2	1.3
Percent Internet Control Message Protocol (ICMP)	0.2	0.1	0
Percent OTHER	0.1	1.1	0
Layer 4			
TCP connections per second	201	277	118
Percent HTTP	49.4	64.6	61.6
Percent Simple Mail Transfer Protocol (SMTP)	5.1	9.4	9.3
Percent Microsoft Network News Transport Protocol (NNTP)	5.5	0	0
Percent OTHER	40.0	26.0	29.1

Table 2 shows the same general traffic characteristics found in the CAIDA data [3,4]. Currently there are no data sets for networks running at or near gigabit-per-second speeds. Further research is needed in this area. Because of the limited scope of this paper, it is assumed that the traffic mix will scale evenly with the increased bandwidth.

With a general understanding of the type of traffic found at the edge of protected networks, it is now possible to explore crafting tests that quantify the metrics at a level useful for NIDS consumers.

HTTP Traffic for Testing

Typically the HTTP protocol is not blocked outbound from firewalls and is a dominant portion of the traffic on the Internet. The servers and clients that implement HTTP have garnered the attention of many crackers and security professionals. HTTP-based signatures make up the majority of signatures in NIDS. Fortunately this situation allows for simplification of the testing in the general case. When HTTP traffic is used to test the capacity of a NIDS, it obviously stresses a large portion of the packet flow architecture.

The use of HTTP traffic has a few other advantages as well. HTTP traffic is relatively easy and inexpensive (in time and money) to produce. Web server testing tools, such as Standees WebBench, can be used to generate traffic to reproduce tests inexpensively. In addition, several vendors are selling network test equipment that utilizes real TCP/IP stack implementations instead of using "canned" traffic. We are most familiar with the products from Caw Networks and Antera. For the test conducted in the section "Example Test Results," Caw Networks *WebReflector* and *WebAvalanche* products were used.

Generating Real Traffic vs. Replaying Traffic

Most NIDS shipping today perform some level of protocol decode and state tracking. Therefore, it is very important that any load traffic exhibit the same characteristics found on a consumer's network. Most of the products that allow for high-speed traffic generation have some critical flaws that make them unsuitable for testing NIDS at high speeds. Some of these issues include:

- Inability to create valid checksums for all layers at high speeds
- Inability to vary the IP addresses in a more random manner than a straight increment
- Inability to maintain state of TCP connections and issue resets if packets are dropped
- Inability to play a large mix of traffic due to the limitations in buffer size for the transmitters

These issues do not plague the replay devices at slower speeds. At high speeds, the buffer size for the replay devices prohibits large traffic samples. Using replay requires the tester to use more replay interfaces. Adding more source interfaces when testing a high aggregate rate presents problems for the test network, as described in the next section. Therefore, using test devices that use real TCP/IP implementation to generate the traffic is preferred.

18.909

Inter-Packet Arrival Gap on High-Speed Test Networks

The typical network setup for testing a gigabit NIDS includes several traffic generators, an attack network, and a victim network. All these devices are typically connected to a switch, and the traffic is then port mirrored, spanned, or copy captured to the NIDS. For high-speed tests, the interface for the NIDS is Gigabit Ethernet. The inter-packet arrival gap on Gigabit Ethernet is 96 ns. Inter-packet arrival gap becomes important as more interfaces are added to the switch. Each interface used to generate traffic, regardless of interface speed, increases the chances that traffic destined for the NIDS will be dropped at the switch.

Imagine that ten Fast Ethernet ports, each generating approximately 80 Mbps of traffic, are used during testing. Eventually several of these ports begin to transmit packets all within a few nanoseconds of each other. Because each of these transmitted packets must be copied to the NIDS, the switch forwards each of the packets to the port where the NIDS is connected. Unfortunately the NIDS is using Ethernet, which demands the 96-ns delay between packets. Because several packets are arriving at the port at very nearly the same time, the port buffer fills up and the switch port drops packets. This problem does not manifest itself if a choke point, such as a firewall or router, is used in the test network. But, if the industry tests require a router or firewall capable of the same high traffic speeds to reproduce the tests, it raises the cost of testing significantly. *It is, therefore, better to use fewer ports for generating traffic on the switch when testing.*

Potential Test Suite

No single test provides all the information needed to quantify the capacity of a NIDS. A suite of tests is used to quantify each portion of a NIDS. Only then looking at the output from all these tests can a consumer infer performance on his network.

Establishing the Peak—Testing the network interface bandwidth establishes the peak for packet capture for the NIDS. The NIDS is never able to perform above this absolute peak on any further tests. Testing the network interface bandwidth is simple. Choose a packet of no interest to the NIDS and resend it at a high rate until the NIDS cannot count all the packets. Repeat the tests for minimum-sized packets and for maximum-sized packets. This reveals the maximum packets per second and the maximum bytes per second, respectively. A good example packet for this test is a UDP packet with ports set to 0 (assuming the NIDS does not send alarms on such a packet).

The Alarm Channel—Testing the alarm channel capacity of a NIDS can be accomplished with a similar test. Choose a packet that causes an alarm. The "Ping of Death" is a good packet for this test. Send this packet at different rates of speed and check packet and alarm counts. Some NIDS buffer alarms when under a heavy load, so a quiescent period after the packets are sent may be necessary before collecting counts.

Stressing the State Database—Inserting, searching, and deleting the state information from the state database are all potential bottlenecks for the NIDS. Varying the IP addresses of traffic requiring state tracking adds load to the database. Opening a large number of TCP connections causes the state database to contain many records. The search performance for a database is affected by the size of the data set. For this type of test, open a large number of concurrent TCP connections and then run one of the more general tests described. The open connections will stress both the database and the overall system architecture.

General Tests with Configurable Metrics—Establishing a baseline of traffic and then varying one of the metrics can expose a NIDS weakness in certain environments. Because each specific component of the NIDS is being tested, it is not necessary to ensure that the traffic looks exactly like the traffic of the end user. If users can extract the same metrics from their traffic, then performance on their networks can be inferred from test results using simpler data. In the example tests found in the section "Example Test Results," the traffic mix consists of only HTTP. HTTP rides on TCP, which requires state tracking. Depending on the level of protocol decoding used, HTTP may also require state tracking. In addition, most of the signatures found in a NIDS are HTTP signatures. Therefore, using an HTTP-only traffic mix still stresses the NIDS in many areas. The example tests in the section "Example Test Results" could have also included additional protocols such as Domain Name System (DNS), SMTP, and NNTP. However, due to time and space constraints these protocols were omitted for this test.

Table 3 shows the characteristics of the traffic mix when using Caw Networks WebReflector and WebAvalanche products. This test equipment allows for high-speed testing using only two ports on the switch for generating traffic. The Caw Networks equipment also has the ability to randomly drop packets. The dropped packets cause their systems to retransmit and the traffic looks more like real-world traffic. By simply varying the HTTP transaction size, many characteristics of the traffic can be manipulated. Using HTTP transaction size is just one example. The maximum segment size (MSS[1]) for the server or client can also be varied to affect the characteristics in other ways.

[1] Used in TCP to specify the maximum amount of TCP data in a single IP datagram that the local system can accept. The MSS is typically the maximum transmission unit of the outgoing interface minus 40 bytes for the IP and TCP headers.

18.908

Table 3

Traffic Mix Characteristics when Using Caw Networks WebReflector and WebAvalanche to Generate HTTP Traffic for General Stress Tests

Transaction size	Packets in stream	Bytes in stream	Average packet size
1,000	8	1,448	181
5,000	13	5,640	434
10,000	20	10,832	542
20,000	33	21,280	645
40,000	60	42,112	702
50,000	73	52,560	720
100,000	140	104,672	748
240,000	326	250,624	769
360,000	486	375,744	773
[illegible]0,000	540	417,472	773

Example Test Results

Using the lessons from the previous section, this section explores two simple tests that evaluate some aspects of the capacity of a NIDS. Other tests are then suggested as a way to further refine the knowledge gained from the example tests.

Test Network

Figure 1 shows the test network used for the example tests. All links to the switch use a single 1-Gbps full-duplex connection.

Figure 1

A Block Diagram Showing the Network Layout for the Example Tests (The WebReflector acts as all the Web servers. The WebAvalanche acts as all the Web clients. The Cisco Catalyst(r) Switch copy captures the traffic to the NIDS.

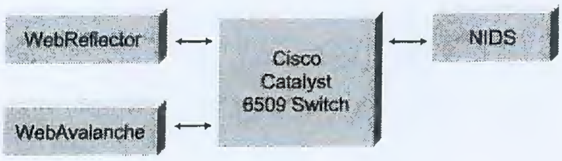

Baseline Test

The first test baselines the capture efficiency of a NIDS in a pure HTTP environment. A full analysis load is assumed by turning on all default signatures. However, the traffic generated does not trigger alarms or events. The number of client hosts on the network is fixed at 5080, and the number of servers is fixed at 2. In this first test, TCP sessions are allowed to run to completion as quickly as possible; therefore, the number of simultaneous open sessions is fixed at fewer than 30 for all cases. The average packet size is varied through manipulation of the HTTP transaction size. This also results in variations in the average number of packets per TCP connection, the packets per second (kpps), and the overall bandwidth used. The final variable manipulated is the number of new TCP connections per second. Each test is run for three minutes before capacity measurements are made. The results of the baseline test are given in Table 4.

18.907

Table 4

Results for Baseline Test (Traffic was HTTP only with 5080 client IP addresses and 2 server IP addresses. The test runs for 3 minutes, with no server delay.

Average packet size	Bandwidth @ 1000 cps	Capture efficiency	Bandwidth @ 2500 cps	Capture efficiency	Bandwidth @ 5000 cps	Capture efficiency
434	50 Mbps	100%	125 Mbps	100%	230 Mbps	99.95%
482	67 Mbps	100%	167 Mbps	100%	335 Mbps	97.10%
542	93 Mbps	100%	232 Mbps	100%		
645	180 Mbps	100%	446 Mbps	100%		
688	275 Mbps	100%	680 Mbps	99.30%		
702	380 Mbps	100%				
720	444 Mbps	100%				

ding Simultaneous Open TCP sessions

The second test introduces simultaneous open TCP sessions. A 4-second delay is introduced on the server response to cause sessions to remain open. All other test variables remain constant. The bandwidth consumed, the packets per second, and the average packet size at each data point are somewhat affected by the open connections. In the absence of a correlation study, it is unclear if these factors are statistically significant. For the purposes of this paper, they are assumed to be insignificant. The results from the open connection test are given in Table 5.

Table 5

Results for Open Connection Test (Traffic was HTTP only with 5080 client IP addresses and 2 server IP addresses. The test runs for 3 minutes, with a 4-second forced server delay.)

Average packet size	Bandwidth @ 1000 cps and 4000 open streams	Capture efficiency	Bandwidth @ 2500 cps and 10000 open streams	Capture efficiency	Bandwidth @ 5000 cps and 20000 open streams	Capture efficiency
434	50 Mbps	100%	125 Mbps	100%	245 Mbps	68.75%
482	69 Mbps	100%	170 Mbps	100%	350 Mbps	72.26%
542	93 Mbps	100%	241 Mbps	100%		
645	181 Mbps	100%	446 Mbps	99.95%		
688	268 Mbps	100%	655 Mbps	95.01%		
702	355 Mbps	100%				
720	440 Mbps	100%				

The Results

The most significant variations in capture efficiency are in the 5000-connections-per-second tests. Capture efficiency appears affected in the 2500-connections-per-second tests as well, but this does not appear to be statistically meaningful.

The two tests differ only in the number of concurrent open TCP sessions. This implies that the state database is the component under stress. With only these results it is not possible to precisely identify what operation within the database is causing the drop in capacity. The traffic may have exceeded the insertion rate of the database, its time to search capacity, or the ability of the system to perform maintenance on the database and delete aging entries.

Regardless, it is apparent that a consumer whose network has an average number of open TCP sessions near 10,000, an average new TCP connections-per-second rate of no more than 2500 per second, and a bandwidth consumption less than approximately 400 Mbps could field the tested NIDS with confidence that the capture efficiency would be at or near 100 percent.

18.906

Further Tests

The example tests do not provide enough information on which to base a full-confidence decision. Nevertheless, by using the same methodology for developing further tests, the NIDS industry or independent labs could establish a suite of tests that could provide quantifiable results for each of the different stress points.

For example, the total number of database insertions per second for the database can be quantified by establishing a test that runs at a very low rate for all other stress points of the NIDS. The traffic must be crafted such that the database needs to start maintaining state on many different key values. One possible way of shaping the traffic is to use a packet generator and generate a valid TCP session with a full three-way handshake. The rate at which the connections are introduced must be ramped upward until the NIDS starts dropping traffic. Because these simple TCP connections consist of small packets, the bandwidth should remain low. Results need to be cross checked with the raw packet capture architecture evaluation results as described in the section "Potential Test Suite" to ensure that it is the database inserts and not the packets-per-second limit that has been reached.

Conclusion

As the NIDS industry matures, standardized testing will become prevalent. Developing these tests can be done using the same concepts of standardized testing found in other industries. Hopefully, the information found in this paper will stimulate the development of standardized tests, providing the NIDS consumer the information that is currently missing. The same techniques used for capacity testing can be extended to other performance areas such as positive ratios.

References

1. Mier Communications: Test report for ManHunt from Recourse Inc. and test report for Intrusion.com's NIDS. At: http://www.mier.com/reports/vendor.html

2. Ranum, M.: Experiences Benchmarking Intrusion Detection Systems. At: http://www.nfr.com/forum/white-papers/Benchmarking-IDS-NFR.pdf

3. Claffy, K., Miller, G., and Thompson, K.: The Nature of the Beast: Recent Traffic Measurements from an Internet Backbone. At: http://www.caida.org/outreach/-papers/1998/Inet98/ (1998)

4. McCreary, S., and Claffy, K.: Trends in Wide Area IP Traffic Patterns: A View from Ames Internet Exchange. At: http://www.caida.org/outreach/papers/2000/-AIX0005/ (2000)

18.905

CISCO SYSTEMS

Corporate Headquarters
Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA
www.cisco.com
Tel: 408 526-4000
800 553-NETS (6387)
Fax: 408 526-4100

European Headquarters
Cisco Systems Europe
11, Rue Camille Desmoulins
92782 Issy-les-Moulineaux
Cedex 9
France
www-europe.cisco.com
Tel: 33 1 58 04 60 00
Fax: 33 1 58 04 61 00

Americas Headquarters
Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA
www.cisco.com
Tel: 408 526-7660
Fax: 408 527-0883

Asia Pacific Headquarters
Cisco Systems, Inc.
Capital Tower
168 Robinson Road
#22-01 to #29-01
Singapore 068912
www.cisco.com
Tel: 65 317 7777
Fax: 65 317 7799

Cisco Systems has more than 200 offices in the following countries and regions. Addresses, phone numbers, and fax numbers are listed on the Cisco Web site at www.cisco.com/go/offices

Argentina • Australia • Austria • Belgium • Brazil • Bulgaria • Canada • Chile • China PRC • Colombia • Costa Rica • Croatia Czech Republic • Denmark • Dubai, UAE • Finland • France • Germany • Greece • Hong Kong SAR • Hungary • India • Indonesia • Ireland Israel • Italy • Japan • Korea • Luxembourg • Malaysia • Mexico • The Netherlands • New Zealand • Norway • Peru • Philippines • Poland Portugal • Puerto Rico • Romania • Russia • Saudi Arabia • Scotland • Singapore • Slovakia • Slovenia • South Africa • Spain • Sweden Switzerland • Taiwan • Thailand • Turkey • Ukraine • United Kingdom • United States • Venezuela • Vietnam • Zimbabwe

TS/JSI/12.02

Anexo (266)

fls. 904

CISCO SYSTEMS

Q & A

Cisco Intrusion **Detection** System

Q. What is a "network-based" IDS?

A. Two basic types of IDSs are on the market today: host-based and network-based systems. The fundamental difference between them is the source of the activity that they monitor and analyze to detect intrusions. Host-based IDSs monitor activity on a host or end system, while network-based IDSs monitor network traffic. Host-based IDSs are used to protect critical network servers or other individual systems containing sensitive information.

Network-based IDSs are used to monitor activity on a specific network segment. Whereas a host-based IDS resides on a workstation and shares CPU with other user applications, a network-based solution is a dedicated platform. Network-based IDSs perform a rule-based or expert system analysis of traffic using parameters set up by the security manager, and the signatures, which flag suspicious or attack activity. The systems analyze network packet headers to make security decisions based on source, destination, and packet type. They also analyze packet data to make decisions based on the actual data being transmitted. These systems scale well for network protection because the number of actual workstations, servers, or user systems on the network is not critical, the amount of traffic is what matters. In addition, sensors placed around the globe can be configured to report back to a central site, enabling a small team of security experts to support a large enterprise. The Cisco® network-based Intrusion Detection System provides network administrators with enhanced security technology and capabilities to secure their networks.

Q. If I already have a firewall, do I really need an IDS?

A. Absolutely. Although an IDS will not replace your firewalls or other security devices for that matter, it serves a very complementary role and addresses certain risks that firewalls cannot. The primary function of the firewall is to control access to services and hosts based on your site security policy. If a service or connection to a specific host is permitted, firewalls typically permit all such traffic, and they do not inspect the content of the permitted traffic. An example is permitting public access to a Web server on a DMZ. All connection requests to the Hypertext Transfer Protocol (HTTP) port on that Web server will be permitted by the firewall, including malicious traffic directed at the HTTP server to exploit a buffer overflow vulnerability. Although most firewalls will not protect against data/content-driven attacks (for example, buffer overflow), IDSs will. Furthermore, firewalls typically will not protect you against attacks originating from inside your network or entering your environment from other ingress points not protected by firewalls (for example, remote access servers). IDSs can be strategically deployed to monitor activity from internal sources and other network ingress points without impacting your network. Deploying an IDS to complement your firewall(s) will significantly enhance your security posture.

18.903

Q. Is there a mechanism by which users may contact the IDS Product Team at Cisco?

A. Yes. Users may pose questions, requests, and comments to the following e-mail address:

ids-news@cisco.com

In addition, users have the ability to share experiences with other users and also pose questions to the Cisco IDS Engineering & Product Marketing teams at the IDS Networking Professionals Forum at:

http://forums.cisco.com/eforum/servlet/
NetProf?page=netprof&CommCmd=MB%3FcmdDdisplay_messages26mode3Dnew26location%3D.ee6e1fc

Q. Does anyone offer a managed IDS service using the Cisco IDS?

A. Yes, numerous managed service providers offer a managed IDS service using the Cisco IDS. These managed service providers include AT&T, Counterpane, IBM Emergency Response Services, NetSolve, Riptech, RedSiren, and Ubizen.

Sensors

Q. What are the new features of the Cisco IDS 4.0 Sensor software?

A. The Cisco IDS 4.0 Sensor software delivers a number of new features and enhancements to the network-based IDS portfolio. These features include:

- Re-architecture of communications protocol to enhance the efficiency of message transactions
- Common code base to allow feature parity between the appliance sensor and the switch sensor
- Delivery of a Layer 2 signature engine to mitigate issues such as man in the middle attacks and ARP spoofing in switched environments
- Introduction of an SMB engine to efficiently address attacks related to SMB
- Ability to capture the trigger packet that caused an alarm
- Enhanced shunning capabilities to allow shunning by port address
- Major enhancements to our existing protocol anomaly techniques
- Provision of Analysis Statistics Engine to deliver information of metrics such as bad checksums, bytes processed, data rates Mbps, TCP nodes per sec, and other analysis metrics
- Introduction of a full featured Cisco IOS-like CLI (command-line interface) for unprecedented sensor management over a secure SSH connection
- Capability of capture and display of the VLAN ID of the malicious traffic that was detected
- Enhancements to IP Fragmentation Reassembly
- Higher levels of granularity for the alarm information that is transmitted to the management console
- Support for ntp
- NAT support
- Logical signature groupings to allow global changes across the groupings
- Ability to implement exceptions to filter events to be displayed
- Tunability of IP session logging parameters

18.902

Q. What performance numbers (Mbps) are supported by the Cisco IDS Sensors?

A. The Cisco IDS 4215 supports 80 Mbps of performance and can be used to protect T1/E1/T3 environments.

At 250 Mbps, the Cisco IDS 4235 can be deployed to provide protection in switched environments, on multiple T3 subnets, and with the support of 10/100/1000 interfaces, it can also be deployed on partially utilized gigabit links.

The Cisco IDS 4250 supports superior performance at 500 Mbps and can be used to protect gigabit subnets and traffic traversing switches that are being used to aggregate traffic from numerous subnets.

Intrusion protection for fully saturated gigabit links is delivered by the Cisco IDS 4250-XL. Using customized hardware acceleration, the IDS-4250-XL can be used to protect gigabit subnets and multiple partially utilized gigabit links.

The Cisco Catalyst® 6500 Series Intrusion Detection System (IDSM-2) Services Module supports 600 Mbps. This module operates within the Catalyst 6500 Series and provides protection for traffic traversing the switch, which could be traffic from a single subnet or from numerous subnets that are being aggregated through the switch.

The Cisco IDS Network Module provides full-featured Intrusion Protection that is integrated into the Cisco 2600, 3600, and 3700 series routers. Each sensor addresses the bandwidth requirements of different routers up to 10 Mbps in the Cisco 2600XM, and up to 45 Mbps in the Cisco 3700 Series. By integrating IDS and branch office routing, Cisco reduces the complexity of securing WAN links and at the same time reduces operational costs. Additionally, by delivering full-featured intrusion protection to remote offices and branch offices, network administrators can now mitigate threats at these remote locations and effectively isolate them from the corporate network. The Network Module has the capability of inspecting GRE/IPsec encrypted packets that are traversing the router into which it integrates.

Q. How does the IDS sensor work?

A. Sensors monitor the network traffic by directly "tapping" the line (for example, via a shared-media hub) or by receiving copies of the traffic (for example, Switched Port Analyzer [SPAN] port on a switch) using a passive, promiscuous interface (the "monitoring interface"). The sensor analyzes the captured packets and compares them against a rule set of typical intrusion activity (that is, "signatures"). If the captured packets match a defined intrusion pattern in the rule set, the sensor sends an alarm to the management console and automatically responds (if configured to do so). The alarms are sent out a separate management interface so as not to impede continual packet capture by the monitoring interface.

Q. What kind of a performance impact does the sensor impose on the monitored network?

A. None. Sensors operate by "tapping" the network (for example, via a shared-media hub) or off copies of the packets (for example, via a switch SPAN port). The monitoring interfaces on the sensors are passive and do not source packets onto the network (the one exception is TCP reset packets for automatic response).

Q. How do you deploy sensors in a switched environment?

A. With most IDS products on the market today, sensors must be placed on the switch SPAN port to monitor network traffic. Although the SPAN port can provide access to network traffic, it does have certain limitations (for example, limited number of SPAN sessions). The Catalyst 6000 IDS Module was designed specifically to address switched environments by integrating the IDS functionality directly into the switch and taking traffic right off the switch backplane.

18.901

Q. What is the Cisco IDS Network Module for the Cisco 2600, 3600, and 3700 series routers?

A. The Cisco IDS Network Module is a network module that is installed in a Cisco 2600, 3600, or 3700 series chassis to provide full-featured intrusion protection services within the router. The Cisco 2600, 3600, and 3700 Series IDS Network Module provides the ability to inspect all traffic traversing the router, to identify unauthorized or malicious activity such as hacker attacks, worms, or denial-of-service attacks, and to terminate this traffic to suppress or contain threats.

Q. How does the Cisco IDS Network Module work?

A. The Cisco 2600, 3600, and 3700 Series IDS Network Module receives copies of packets directly from the router's backplane in a passive or promiscuous mode. The packets are passed through the internal monitoring interface for classification and processing. The Cisco 2600, 3600, and 3700 Series IDS Network Module analyzes the captured packets and compares them against a rule set of typical intrusion activity. If the captured packets match a defined intrusion pattern in the rule set, the IDS Network Module can take one of two actions. It can send a command to the router to either shut down the interface or it can send a TCP reset packet to the sender to stop the TCP session causing the attack.

Q. What is the rated performance of the Cisco 2600, 3600, and 3700 Series IDS Network Module?

A. The Cisco 2600, 3600, and 3700 Series IDS Network Module provides up to 10 Mbps for the 2600XM Series and up to 45 Mbps for the 3700 Series, depending on the platform in which the network module is inserted.

Q. Can the Cisco Network IDS Sensors monitor trunked traffic?

A. Both IDSM and the appliance sensors can monitior 802.1q traffic and, hence, are VLAN aware.

Q. What type of interfaces are supported on the appliance sensors?

A. Copper interfaces are supported on the IDS 4215 and IDS 4235. Both Copper and Fiber interfaces are supported on the IDS 4250 Sensor. The 4250-XL supports dual fiber interfaces with MTRJ connectors.

Q. Does Cisco IDS provide multi-interface support?

A. Yes. Dual sniffing interfaces are supported on the IDS 4250-XL. Up to 5 interfaces are supported on the IDS 4215, 4235, and 4250 Sensor appliances.

A configurable four Fast Ethernet interface card is provided for other models of the Cisco 4200 Series sensors to deliver a total of five sniffing interfaces for each sensor—one onboard sniffing interface plus four Fast Ethernet configurable interfaces.

The Cisco IDSM-2 can be used to monitor traffic from multiple interfaces. The network module for the Cisco access routers can monitor traffic from any of the router interfaces.

Q. Is the user notified when the sniffing interface of a sensor is disconnected?

A. Sensors are equipped with a monitoring interface for data packet capture and a command and control interface for transmitting alarms to the management console and receiving configuration information from the management console. When the sniffing interface is disconnected, an alarm is triggered. This setup provides the user with an alert mechanism when the interface is tampered with and, hence, assures persistent operation.

18.900

Q. Does the Cisco IDS Sensor provide an indication of when it is oversubscribed?

A. Yes. The sensor, IDSM, and network module issue an alarm when their respective performance ratings are exceeded.

Q. Where can I find more details on the IDS signature algorithms?

A. For more information on the signature algorithms, please refer to a white paper that may be downloaded at:

http://www.cisco.com/warp/public/cc/pd/sqsw/sqidsz/prodlit/idssa_wp.htm

Q. Does Cisco IDS protect against common IDS evasion techniques?

A. Cisco IDS also includes protection from a number of advanced, anti-IDS evasion techniques including:

- IP fragmentation reassembly
- TCP streams reassembly
- Unicode Web deobfuscation

Q. Can the sensor itself be attacked and compromised?

A. A properly configured and installed sensor cannot be compromised. The monitoring interface (connected to the production network) cannot be detected, and packets cannot be directed at it. The interface is in promiscuous mode, and has neither a protocol stack nor an IP address bound to it. It is not susceptible to "antisniff" detection techniques. The separate management interface does have an IP address, but Cisco recommends that a separate, isolated management subnet be used to provide connectivity from the management interface on the sensor to the IDS management console. In addition, only a very limited number of services are available from the management interface, and access controls can be configured to allow only designated management systems to connect to the sensor.

Q. What is Cisco Countermeasures Research Team (C-CRT)?

A. The core of the Cisco IDS solution—the advanced protection capabilities—is developed and maintained by C-CRT. This team of elite security professionals is dedicated to:

- Advancing countermeasures research
- Identifying and responding to new threats
- Distributing proactive signature files and signature micro-engines
- Maintaining our network security database (NSDB)
- Contributing research to the Cisco Security Encyclopedia (CSEC)
- Improving the state of threat mitigation science

3899

The C-CRT is comprised of elite "white hat" personnel. C-CRT's esteemed credentials include:

- *Heritage and tenure*—most joined Cisco through the WheelGroup acquisition.
- *Government clearance*—greater than 65 percent have held Secret and Top Secret Department of Defense security clearances.
- *Military backgrounds*—experience logged from USAF Information Warfare Center, Department of Defense (DOD), Department of Energy (DOE), National Security Agency (NSA), Central Intelligence Agency (CIA), or other notable government organizations.
- *Security experience*—average member of the C-CRT has over six years of computer security experience, allowing Cisco to deliver the most mature, accurate, and industry-proven intrusion protection solutions.

Q. Can I create my own signatures?

A. Because the security objectives for each IDS deployment are unique, Cisco IDS adds granularity to the way in which sensors may be tuned to specifically suit the environment in which they are deployed. Using our innovative TAME (Threat Analysis Micro Engine) policy language, users can create new policies or modify existing policies to meet their unique security objectives. Since the TAME policies are decoupled from the sensing application, changes can be made without affecting sensor performance or reliability.

Q. Is it possible to record and replay the IP session of the source IP address that triggered an IDS alarm?

A. IP session logging provides extensive logging that is important for system troubleshooting as well as for reconstructing system events before and after attacks. Cisco IDS augments this capability by converting these logs to a standard TCP dump format that allows them to be viewed and replayed using public domain utilities, such as Ethereal and TCPReplay.

Q. Can the sensor detect attacks if the traffic is encrypted, for example IPsecurity (IPsec) or Secure Sockets Layer (SSL)?

A. The Cisco IDS Sensor analyzes both packet header information (context data) and packet data information (content data) to determine if suspicious activity is occurring. Encryption algorithms encrypt the data portion of the packet for confidentiality. Because it can process only what it can "see," the Sensor cannot detect attacks that require inspection of the payload or data fields within a packet. It will, however, still alarm and respond to attacks, which are detected from the unencrypted packet header information. All network-based IDSs suffer this problem. Therefore, in networks carrying encrypted traffic, Sensor placement is critical. To take advantage of their full intrusion-detection capability, the Cisco IDS Sensors should be installed where the traffic has already been decrypted. Otherwise, the Sensor can be placed on an encrypted segment and will detect all but the packet data or payload-based attacks.

Q. What techniques does Cisco use for mitigating threats?

A. Several techniques provide comprehensive protection against the latest cyber threats, including simple pattern matching, stateful pattern matching, protocol anomaly detection, heuristic-based detection, and anomaly detection.

Q. Does Cisco IDS deliver Peer to Peer signatures?

A. Yes. Cisco IDS delivers protection against file-sharing threats with support for advanced P2P attack mitigation techniques.

18.558

Q. How are the Cisco IDS signatures updated?

A. Cisco posts signature updates on Cisco Connection Online (CCO) approximately every 14 days. Cisco IDS provides a facility to automatically distribute new signature files and application upgrades to sensors without operator involvement. Utilizing a secure staging technique, new signature files are placed on a central server and passed to the sensor at scheduled intervals. After verifying the integrity of the package, the sensor automatically installs the update. This new capability significantly streamlines the process of regularly updating remote sensors, thereby lowering the recurring operational costs associated with this task. Additionally, users can subscribe to Cisco Active Update notification services to stay informed about breaking vulnerability news and posted countermeasures at:

http://www.cisco.com/warp/public/779/largeent/issues/security/idsnws/archive.html

Users may refer to the following site for a chronological listing of the Cisco IDS Active Update Notification Bulletins:

http://www.cisco.com/warp/public/779/largeent/issues/security/idsnws/archive.html

Q. How do the Cisco IDS Sensors and management consoles communicate with each other?

A. Communication between the Cisco IDS 4.0 Sensors and management consoles is provided by a secure (SSL) XML based messaging format. All alarm transmissions from the sensor to the management console are acknowledged.

If connectivity from the sensor to the management console is disrupted, the sensor will continue to monitor the network, and will queue alarms and retransmit until successful.

Q. How much additional network traffic does the Cisco IDS generate?

A. Because each alarm and acknowledgment is contained in a single UDP packet, there is negligible impact on network traffic.

Q. Is there a site that lists all the supported Cisco IDS signatures?

A. Yes. Users may access the latest Cisco IDS signatures at the Cisco Secure Encyclopedia site at:

http://www.cisco.com/pcgi-bin/front.x/csec/idsHome.pl

Q. Does Cisco support a centralized site that contains a compiled listing of the latest vulnerabilities?

A. Yes. Cisco's Security Encyclopedia is a one-of-a-kind clearinghouse of security and vulnerability information. Unlike other security databases that simply consolidate vulnerability information published on a number of public-source Web sites, the CSEC contains statistics on the vulnerabilities by industry or by sector. These statistics are compiled from over 400 actual Security Posture Assessments (SPA) performed by the Cisco Security Consulting team. CSEC is developed and maintained by the elite C-CRT. You may visit the CSEC site at:

http://www.cisco.com/go/csec

Q. Where can I download the latest IDS software?

A. Both current and archived IDS Sensor software can be downloaded at the Software Center on CCO (CCO login required):

http://www.cisco.com/public/sw-center/ciscosecure/ids/crypto/

Q. Where can I access documentation on the Cisco IDS Sensor Software?

A. Documentation for sensor software updates are available at:

http://www.cisco.com/univercd/cc/td/doc/product/iaabu/csids/index.htm

18.897

Threat Response Technology

Q. Is Threat Response an event correlator?

A. Threat Response is not an event correlator. Event correlation involves analyzing data from NIDS sensor, firewalls, routers, and other sources. Instead of correlating this type of data, Threat Response investigates the actual target of an attack. This is the same process an expert network security specialist would use, and is the best way to determine if a system has been compromised.

Q. Does Threat Response conduct network vulnerability scans?

A. No, Threat Response does not conduct enterprise wide scans of your environment, nor does Threat Response inventory your network. The designers of Threat Response have been network administrators for large mission-critical enterprises such as the U.S. Pentagon and know that downtime is unacceptable. Because of this, Threat Response conducts a low-impact investigation of targeted systems only when needed. Threat Response is able to work in dynamic network environments (including DHCP and wireless) without the need to run regular vulnerability scans that can disrupt your enterprise.

Q. Does Threat Response require deployment of software across the enterprise?

A. No, Threat Response does not require the deployment of software across the enterprise. Threat Response accesses these systems in the same way a security network administrator would—with read access privileges.

Q. How does Threat Response stay up to date with the latest attacks?

A. Cisco releases updates to keep the Threat Response IDS signature database up to date with the IDS vendors, as well as corresponding forensic signature updates to investigate IDS events. When an update is available, the administrator will be notified via the Threat Response GUI and can use the integrated auto-update feature to keep the product current.

Q. What does Threat Response do once real attacks are identified?

A. Threat Response will provide the user with detailed information on how the event was investigated, as well as any forensic data gathered showing details on the actual attack. This information can then be used by an administrator to quickly remediate an intrusion.

Q. What type of systems can Threat Response investigate?

A. Threat Response can conduct a full active investigation of systems running Windows-based operating systems. For systems running Linux, Solaris, and other forms of UNIX, Threat Response will perform passive checks as a first line of investigation. Based on this initial analysis, Threat Response can eliminate many alarms that are not targeted to those specific platform types.

Q. How do I get Threat Response technology?

A. Threat Response technology is currently available as a full featured, 90-day free trial software solution. The trial version ships with every Cisco IDS sensor. Once the trial has expired, customers have the choice of:

- switching to a reduced capabilities free version (only conducts basic level investigation of the targeted system)
- purchasing the full featured version, which will be offered as part of a VMS bundle

Management

Q. What management console options are available for the Cisco IDS?

A. See Table 1 and Table 2.

Table 1 Event Management

	IEV	IEV with CTR Technology	VMS SecMon	VMS SecMon with CTR Technology
Deployment method	Dedicated system required	Dedicated system required	Dedicated system required	Dedicated system required
GUI type	Java desktop application	Browser-based GUI	Browser-based GUI	Browser-based GUI
# sensors	5	5	Unlimited	Unlimited
Event types	IDS	IDS	IDS, Firewall, Router	IDS, Firewall, Router

Table 2 Device Management

	IDM	CLI	VMS Management Center
Deployment method	Integrated on sensor	Integrated on sensor	Dedicated system required
GUI type	Browser-based GUI	Browser-based GUI	Browser-based GUI
# sensors	Unlimited, by sensor	Unlimited, by sensor	Unlimited, by sensor groups
Event types	IDS	IDS	IDS, Firewall, Router

IEV = Cisco IDS Event Viewer, included free of charge with IDS sensor
VMS SecMon = CiscoWorks Monitoring Center for IDS, part of the CiscoWorks VMS bundle
IDM = IDS Device Manager, included free of charge with the IDS sensor
CLI = Command-line interface, included free of charge with the IDS sensor
VMS Mgmt Center = CiscoWorks Management Center for IDS, part of the CiscoWorks VMS bundle

For More Information

More information on Cisco's VMS solutions can be found at:

http://www.cisco.com/go/vms

In addition, augmentation to the alarm viewing, analysis, and reporting capabilities of the Cisco IDS Management solution are provided through third-party applications that are available from Cisco Security Associates partners.

For more details see:

http://www.cisco.com/warp/public/778/security/sap/management.html

Q. Are Cisco IDS communications encrypted?

A. IPsec functionality is included on the appliance Sensors to allow customers to encrypt traffic to management consoles with IPsec capabilities.

Q. If I lose connectivity to a remote sensor, can I tell from the management console?

A. The Cisco IDS management consoles monitor the health of a Sensor via a continuous heartbeat. If communications with the Sensor are lost for more than one minute (by default), a visual indicator is displayed on the management console, indicating a communications failure with the Sensor. If it is determined that a sensor has failed, it can be quickly replaced with another sensor and the configuration, stored on the management console, and can be quickly pushed to the new sensor.

Q. How many sensors can one Cisco IDS Management console manage?

A. Although the technical limit is very large (greater than 1000), Cisco typically recommends a ratio of 20 to 25 sensors per management console for practical reasons. With ratios greater than this, operators can be easily overwhelmed with the volume of information that they may be required to analyze, thereby diminishing the overall effectiveness of the IDS. For deployments larger than 25 sensors, multiple management consoles can be installed to scale the number of sensors.

Q. Can I have multiple Cisco IDS management consoles?

A. The Cisco IDS architecture supports the deployment of multiple management platforms. Sensors can send alarms to multiple management consoles simultaneously, and management consoles can forward alarms to other management consoles, allowing customers to build large, hierarchical management infrastructures.

CISCO SYSTEMS

Corporate Headquarters
Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA
www.cisco.com
Tel: 408 526-4000
800 553-NETS (6387)
Fax: 408 526-4100

European Headquarters
Cisco Systems International BV
Haarlerbergpark
Haarlerbergweg 13-19
1101 CH Amsterdam
The Netherlands
www-europe.cisco.com
Tel: 31 0 20 357 1000
Fax: 31 0 20 357 1100

Americas Headquarters
Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA
www.cisco.com
Tel: 408 526-7660
Fax: 408 527-0883

Asia Pacific Headquarters
Cisco Systems, Inc.
Capital Tower
168 Robinson Road
#22-01 to #29-01
Singapore 068912
www.cisco.com
Tel: +65 6317 7777
Fax: +65 6317 7799

Cisco Systems has more than 200 offices in the following countries and regions. Addresses, phone numbers, and fax numbers are listed on the **Cisco Web site at www.cisco.com/go/offices**

Argentina • Australia • Austria • Belgium • Brazil • Bulgaria • Canada • Chile • China PRC • Colombia • Costa Rica • Croatia
Czech Republic • Denmark • Dubai, UAE • Finland • France • Germany • Greece • Hong Kong SAR • Hungary • India • Indonesia • Ireland
Israel • Italy • Japan • Korea • Luxembourg • Malaysia • Mexico • The Netherlands • New Zealand • Norway • Peru • Philippines • Poland
Portugal • Puerto Rico • Romania • Russia • Saudi Arabia • Scotland • Singapore • Slovakia • Slovenia • South Africa • Spain • Sweden
Switzerland • Taiwan • Thailand • Turkey • Ukraine • United Kingdom • United States • Venezuela • Vietnam • Zimbabwe

ANEXO (26H)

18.894

CHAPTER 5

Monitoring Device Status and Events

From the Monitor tab, you can monitor device status and events. To monitor device status, select Monitor > Device Status. To view real-time and historical events, select Monitor > Event Viewer.

This chapter contains the following sections:

- Monitoring Device Status, page 5-1
- Using Event Viewer, page 5-2

Monitoring Device Status

From the Device Status page, which you access by selecting Monitor > Device Status, you can view a list of monitored postoffice devices, their device types, and connection status.

Displaying the Status of Monitored Devices

You can view the connection status of the postoffice devices you are monitoring with Security Monitor.

18.893

To display the status of the postoffice devices, follow these steps:

Step 1 Select **Monitor > Device Status.**

The Device Status page appears.

Step 2 To refresh the display, click **Refresh**.

The device list and connection status for each is updated.

Using Event Viewer

You can use Event Viewer to view real-time and historical events. Events include IDS alerts (generated by network-based and host-based sensors, IOS devices, and PIX devices), syslog messages, and audit logs. This section contains the following topics:

- Learn More About Event Viewer, page 5-2
- Event Viewer Display, page 5-4
- Event Viewer Terminology and Conventions, page 5-5
- Event Viewer Task List, page 5-13
- Defining Event Viewer Preferences, page 5-22

Learn More About Event Viewer

Sensors and other network devices can continually forward events to Monitoring Center for Security (Security Monitor). These events are stored in the Security Monitor database. Event Viewer allows you to view the events stored in the Security Monitor database. You can view real-time events as they are forwarded to Security Monitor, and you can also view historical events stored in the database.

Note Event Viewer is not the same as the Windows Administrative Tool also known as Event Viewer.

18.892

The following list contains examples of events that can be viewed in Event Viewer:

- An attempt by a hacker to break into a computerxd5 s IDS Messages (IDS Alerts)
- General security-related messages (Security Summaries)
- A status message from a program or a computer (Audit Logs)

Event Viewer queries the database at regular intervals to extract the latest events. You can specify the query interval on the Event Viewer Preferences panel. For more information, see Specifying Event Viewer Preferences, page 5-17.

You can access Event Viewer by selecting Monitor > Event Viewer.

Starting Event Viewer

Before you start Event Viewer, you must specify which events you want to display.

Note The event start and stop times are calculated using the time at which events are stored in the database, not the time that the events were generated by the sensor. Usually, the two times are close, if not identical. Storage and generation times differ greatly only if there are communications problems that postpone sending events from the sensor to the database.

To start Event Viewer, follow these steps:

Step 1 Select **Monitor > Event Viewer**.

The Launch Event Viewer page appears.

Step 2 To select which event type appears in Event Viewer, select an option from the **Event Type** list box.

Step 3 Select an option in the **Event Start Time** section to specify the oldest events that appear in Event Viewer.

- Select **At Earliest** to view events starting with the oldest stored in the database.
- Select **At Time** to specify an exact date and time from which you want to start displaying events.

Step 4 Select an option in the **Event Stop Time** section to specify the most recent events that appear in Event Viewer.

- Select **Donxd5 t Stop** for real-time event analysis.
- Select **At Time** to specify an exact date and time up to which you want to display events.

Step 5 To start Event Viewer, click **Launch Event Viewer**.

Event Viewer appears.

Tip To start another Event Viewer window from the current Event Viewer window, choose File > New > Window in Windows Explorer or File > New > Navigator Window in Netscape Navigator.

Event Viewer Display

Event Viewer combines the functionality of a spreadsheet (such as Lotus 1-2-3 or Microsoft Excel) with that of a hierarchical, drill-down directory (such as Windows Explorer) to create a collection of event records called a *drillsheet* (a *drill*down spread*sheet*). The drillsheet displays groups of similar event records on a single row of a grid, enabling you to detect patterns in the data.

Event Viewer contains a grid pane that organizes and displays event records. Event Viewer can read real-time events and historical events from the database. You can configure the grid pane in a variety of ways to display information about alarms detected by the sensor. For example, you can delete unwanted columns and expand and collapse cells.

gray background but has data displayed, rather than a "+". This means that this cell has not been expanded, but there is only one data element to be displayed, so it is displayed anyway.

Figure 5-1 Event Viewer Drillsheet

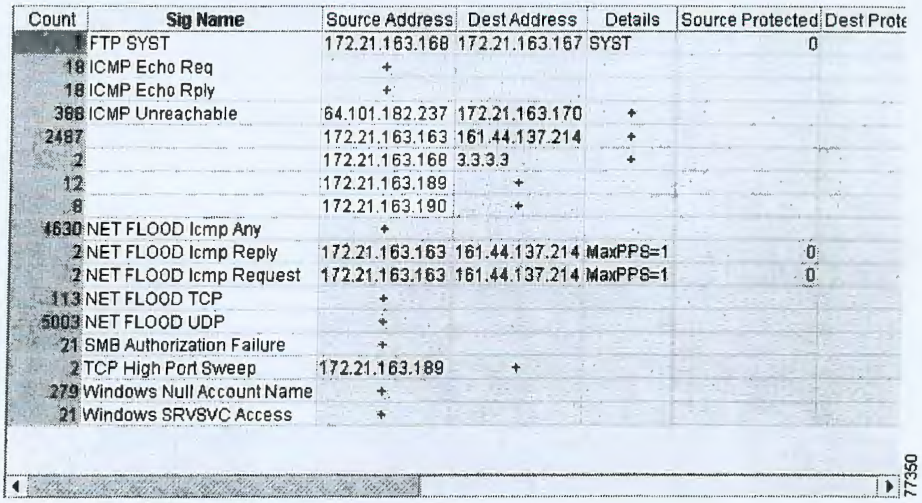

Count	Sig Name	Source Address	Dest Address	Details	Source Protected	Dest Prote
1	FTP SYST	172.21.163.168	172.21.163.167	SYST	0	
18	ICMP Echo Req	+				
18	ICMP Echo Rply	+				
388	ICMP Unreachable	64.101.182.237	172.21.163.170	+		
2487		172.21.163.163	161.44.137.214	+		
2		172.21.163.168	3.3.3.3	+		
12		172.21.163.189	+			
8		172.21.163.190	+			
4630	NET FLOOD Icmp Any	+				
2	NET FLOOD Icmp Reply	172.21.163.163	161.44.137.214	MaxPPS=1	0	
2	NET FLOOD Icmp Request	172.21.163.163	161.44.137.214	MaxPPS=1	0	
113	NET FLOOD TCP	+				
5003	NET FLOOD UDP	+				
21	SMB Authorization Failure	+				
2	TCP High Port Sweep	172.21.163.189	+			
279	Windows Null Account Name	+				
21	Windows SRVSVC Access	+				

77350

Selecting Cells

Many of the functions performed by Event Viewer require you to select cells in the drillsheet. Typically, you select a cell by clicking it. It is important to understand what it means to select a cell in the drillsheet.

When you select a cell in the drillsheet you are actually selecting a node in the event tree. When you perform an operation against a selected cell, you are actually performing an operation on all the branches of nodes that pass through the selected cell. For example, in Figure 5-2, if you select the "ICMP Unreachable" cell, any operation that you run on that cell is performed for all events that have the name "ICMP Unreachable". In this case, that would be all elements in rows 4 through 8. If you intend to execute an operation against only row 4, you must select, in Figure 5-2, either the "64.101.182.237" cell or a cell to its right.

Figure 5-2 Event Viewer Drillsheet

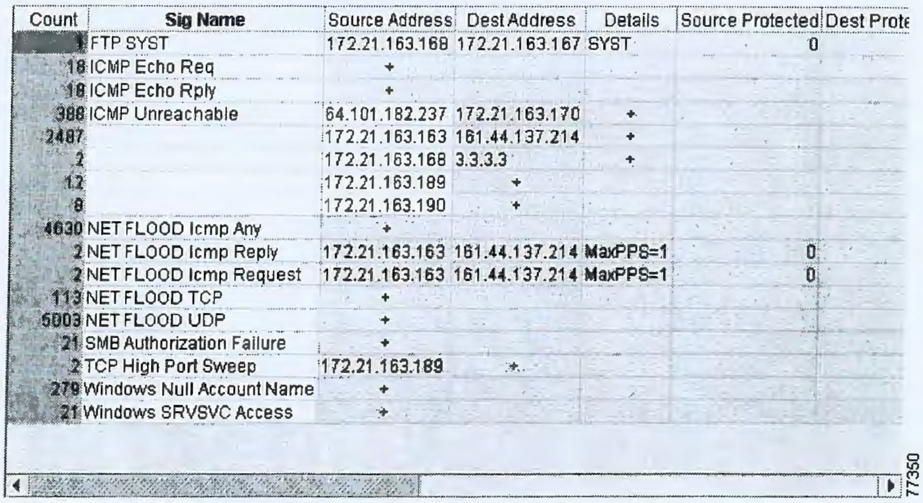

Count	Sig Name	Source Address	Dest Address	Details	Source Protected	Dest Prote
1	FTP SYST	172.21.163.168	172.21.163.167	SYST	0	
18	ICMP Echo Req	+				
18	ICMP Echo Rply	+				
388	ICMP Unreachable	64.101.182.237	172.21.163.170	+		
2487		172.21.163.163	161.44.137.214	+		
2		172.21.163.168	3.3.3.3	+		
12		172.21.163.189	+			
8		172.21.163.190	+			
4630	NET FLOOD Icmp Any	+				
2	NET FLOOD Icmp Reply	172.21.163.163	161.44.137.214	MaxPPS=1	0	
2	NET FLOOD Icmp Request	172.21.163.163	161.44.137.214	MaxPPS=1	0	
113	NET FLOOD TCP	+				
5003	NET FLOOD UDP	+				
21	SMB Authorization Failure	+				
2	TCP High Port Sweep	172.21.163.189	+			
279	Windows Null Account Name	+				
21	Windows SRVSVC Access	+				

Furthermore, if you select a cell that is blank because its value is implied by the cell above it (for example, the cell just below the "ICMP Unreachable" cell), the branch of the node that is operated on is the branch that is defined by the first cell that is filled in to the right of the blank cell that you selected. For example, in Figure 5-2, if you select the blank cell just below the "ICMP Unreachable" cell, when you perform an action, Event Viewer behaves as though you selected the "172.21.163.163" cell.

Note You can use the Preferences panel to change the behavior described in the above paragraph. For more information, see Specifying Event Viewer Preferences, page 5-17.

The Count Column and the Event Count Tool-Tip

Event Viewer provides two mechanisms for displaying the number of events in a group: the Count column and the event count tool-tip.

- **Count Column**—The Count column is the first column in the drillsheet; you cannot move, collapse, or delete it. In the Count column, a cell for a given row displays the number of events represented by that row. For example, the drillsheet in Figure 5-3 indicates that there are 18 "ICMP Echo Req" events.

18.888
)

However, the count of 7 in the fourth row does not mean that there are 7 "ICMP Unreachable" events; it means that there are 7 "ICMP unreachable" events with a source address of "172.21.163.190".

- **Event Count Tool-Tip**—You can find out how many events are represented in a branch that spans more than one row by resting the mouse pointer on the cell you are interested in. A tool-tip indicates how many events pass through that branch. The tool-tip also displays a *child count.* The child count is the number of unique data elements to the right of the cell you have selected.

In Figure 5-3, when you rest the mouse pointer on the source address 172.21.163.190, you see a count of 8 and a child count of 2. This means that there are 8 "ICMP Unreachable" events with a source address of 172.21.163.190. The values in the Count column confirm this. The Count column indicates that there are 7 events with the fields "ICMP Unreachable", 172.21.163.190 and 64.101.128.56 and 1 event with the fields "ICMP Unreachable", 172.21.163.190 and 171.70.168.183. The sum of 7 and 1 is 8.

Figure 5-3 Event Count Tool-Tip

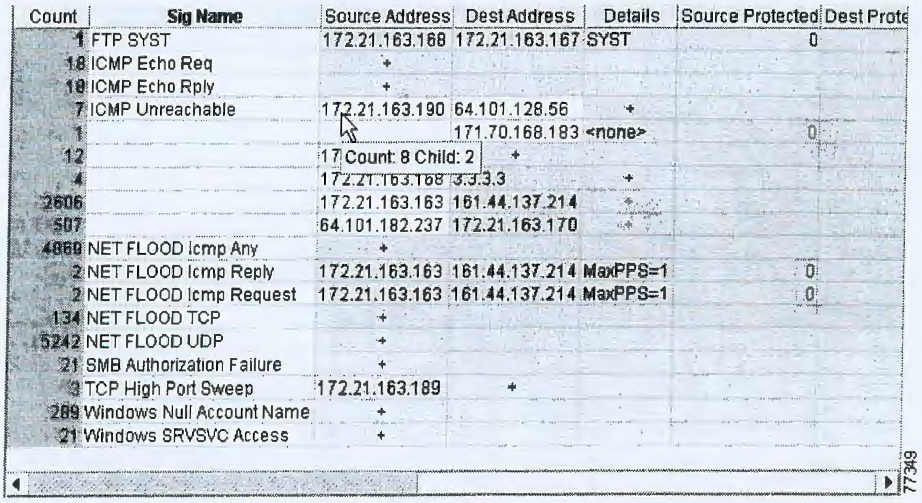

Count	Sig Name	Source Address	Dest Address	Details	Source Protected	Dest Prote
1	FTP SYST	172.21.163.168	172.21.163.167	SYST	0	
18	ICMP Echo Req	+				
18	ICMP Echo Rply	+				
7	ICMP Unreachable	172.21.163.190	64.101.128.56	+		
1			171.70.168.183	<none>	0	
12		17	+			
4		172.21.163.168	3.3.3.3	+		
2606		172.21.163.163	161.44.137.214	+		
507		64.101.182.237	172.21.163.170	+		
4869	NET FLOOD Icmp Any	+				
2	NET FLOOD Icmp Reply	172.21.163.163	161.44.137.214	MaxPPS=1	0	
2	NET FLOOD Icmp Request	172.21.163.163	161.44.137.214	MaxPPS=1	0	
134	NET FLOOD TCP	+				
5242	NET FLOOD UDP	+				
21	SMB Authorization Failure	+				
3	TCP High Port Sweep	172.21.163.189	+			
289	Windows Null Account Name	+				
21	Windows SRVSVC Access	+				

19.887

Status Propagation

This section describes how Event Viewer determines the severity for individual events and groups of events.

- **Individual Events**—Some events are more severe than others. Some events represent unmistakable and devastating actions, while others might represent occurrences that are either less damaging or more ambiguous, or both. To indicate the severity of an alarm, a sensor associates a severity level with each alarm that is generated. In general, those severity levels are Informational, Low, Medium, and High, and the colors associated with those levels are blue, green, yellow, and red, respectively.
- **Event Groups**—Event Viewer uses a "propagate most severe" status propagation scheme. This means that in a group of events, the severity of the group is the severity of the most severe event in the group. For example, if an event group contains one *High* event and 17 *Low* events, the severity of the group is *High.*

  The background color of the event groupxd5 s Count column cell is the color associated with the event groupxd5 s severity. For example, if row number 17 represents 200 events, and if one of those 200 events is High, the event group itself is considered High, and the background color of the Count column cell at row number 17 is red.

  As rows are modified (either through the addition or deletion of events or through the manipulation of rows by the user), the status of the rows is modified in real time.

  In addition to being shown in the Count column, the severity of an event group is reflected in the Severity column. For more information about how you can manipulate drillsheets to group events by severity, see Sorting Data and Shifting Columns, page 5-11.

Context Buffer

Some alarms have additional data associated with them. This data segment, called a *context buffer*, records exactly what traffic was traversing the network at the time the alarmxd5 s signature was detected. The context buffer contains up to 256 bytes of incoming traffic and 256 bytes of outgoing traffic.

Not all events have context buffers. The following is a partial list of alarms that have context buffers:

- 3100 Smail Attack
- 3101 Sendmail Invalid Recipient
- 3102 Sendmail Invalid Sender
- 3103 Sendmail Reconnaissance
- 3104 Archaic Sendmail Attacks
- 3200 WWW Phf Attack
- 3201 WWW General cgi-bin Attack
- 6251 Telnet Authorization Failure
- 8000 String Match

  The 8000 signature contains the following subsignatures:

  - 2101 FTP Retrieve Password File
  - 2302 Telnet-/etc/shadow Match
  - 2303 Telnet-++
  - 51301 Rlogin-IFS Match
  - 51302 Rlogin-/etc/shadow Match
  - 51303 Rlogin-++

For more information about signatures, see the Network Security Database (NSDB). You can access the NSDB at https://*hostname*/vms/nsdb/html/all_sigs_index.html, where *hostname* is the name of the computer on which Security Monitor is installed. For information about viewing the NSDB entry for an event in Event Viewer, see Viewing the Network Security Database, page 5-19.

If even one event represented by a row has a context buffer, the value in the Count column is **bold**. To view the context buffer(s) associated with an event group, select a cell, and then select **View > Context Buffer** in the TOC. For more information, see Viewing the Context Buffer, page 5-18.

18.885

Sorting Data and Shifting Columns

This section describes two functions that help you find data: the ability to sort data within a column, and the ability to change the order of columns within a drillsheet.

Sorting Data

By default, all columns except time-related columns are displayed in ascending order. This means that, from top to bottom, numbers are displayed from least to greatest, and words are displayed from A to Z. To change the sorting scheme of a column from ascending to descending (or vice versa), click the column header. To change it back, click the column header again.

Note By default, time-related columns (times, dates, and timestamps) are displayed in descending order. The most recent dates are displayed at the top of the list, ensuring that recent events are displayed at the top of lists. To change to ascending order, click the column header.

Sorting within a drillsheet is different from sorting in a spreadsheet in one significant way: In a drillsheet, sorting data elements in a particular column is constrained by the nature of the data in the columns to the left.

For example, Table 5-1 shows two columns. The first column has last names, and the second column has first names.

Table 5-1 First Names Sorted in Ascending Order

Last Name	First Name
Baker	Alan
	Wanda
Jones	Bob
	Xena
Smith	Charles
	Yvonne

18.824

The Last Name column and the First Name column are ascending. First names are associated with last names, so any sorting of first names must be within last names. If you click the First Name header to change the sorting scheme to descending, you obtain the results shown in Table 5-2.

Table 5-2 First Names Sorted in Descending Order

Last Name	First Name
Baker	Wanda
	Alan
Jones	Xena
	Bob
Smith	Yvonne
	Charles

The data in the first column did not change when you changed the sorting scheme of the second column.

Shifting Columns

The order of the columns in a drillsheet determines how events are grouped together. For example, if your first three columns (excluding the Count column) are, in order, "Name, Source Address, Dest Address", all events are grouped by name, and then each of those name groups is divided into subgroups by source address, and then each of those subgroups is divided into even smaller groups by destination address.

To change the way events are grouped, you must change the order of the columns.

To change column order, click and hold the cursor over the header of the column you want to move, and then drag the header to the desired location and release the mouse button. The window is redrawn.

In most cases, redrawing after a column shift is nearly instantaneous. However, with large numbers of events (tens of thousands or more), a slight delay may occur during redrawing.

18.883

Event Viewer Terminology and Conventions

The following sections describe the terminology and common usage conventions surrounding Event Viewer and its drillsheets.

Event Display

A drillsheet has rows and columns, and the intersection of a row and a column is called a *cell.*

The background color of a cell gives some information about the cell:

- If a cell is white, only one data element is associated with that cell.
- If a cell is gray, that cell may represent more than one data element.
  - If a cell is gray and displays the + symbol, that cell represents more than one data element. You can see all the data elements by expanding this cell.
  - If a cell is gray but displays a single data element (for example, 172.21.172.6), that cell has not been expanded, but it contains only a single data element, so that element is displayed anyway.

Note You can use the Preferences panel to modify some of the behavior described in this section. For more information, see Specifying Event Viewer Preferences, page 5-17.

Note The conventions governing the background colors of cells in the Count column are different and are described in Status Propagation, page 5-9.

For example, in Figure 5-1, there is more than one source address associated with the events that have the name "ICMP Echo Req". Therefore, the Source Address cell in the ICMP echo request row is gray and displays "+". We also see that Source Address column has been expanded for the "ICMP Unreachable" events. Therefore, the cells in the Source Address column for the ICMP Unreachable rows are white. Finally, note that the destination address 172.21.163.170 has a

18.882

The Count column is always the first column in the display. You cannot drag the Count column to another position, and you cannot drag another column to the left of the Count column. If you try, a dialog box appears, and the columns revert to their original positions.

When columns are shifted, the entire window is redrawn, meaning that all rows are expanded to the Event Expansion Boundary for that window. To reduce the number of rows that are drawn with each column shift, consider making one of the first few columns the Event Expansion Boundary. For more information, see Setting the Event Expansion Boundary, page 5-15.

Deleting Columns

You can delete a column from the Event Viewer display. Deleting a column affects only the Event Viewer display that you are viewing. It does not change the default column arrangement for other existing or future Event Viewer displays.

To delete a particular column from the current Event Viewer display, select any cell in the column that you want to delete. Then, select **Delete > One Column** in the TOC.

Event Viewer Task List

This section describes the tasks that you can perform from the TOC in Event Viewer.

- Deleting Events, page 5-14
- Collapsing Cells, page 5-15
- Setting the Event Expansion Boundary, page 5-15
- Expanding Cells, page 5-16
- Suspending and Resuming New Events, page 5-17
- Specifying Event Viewer Preferences, page 5-17
- Viewing the Context Buffer, page 5-18
- Viewing Host Names, page 5-19
- Viewing the Network Security Database, page 5-19
- Viewing Event Statistics, page 5-20

18.881

- Refreshing Events, page 5-20
- Blocking a Host or a Network, page 5-21
- Removing a Block, page 5-22

Deleting Events

You can delete events from the current Event Viewer display, or you can use Event Viewer to mark events for deletion from the database. You can also delete a column, which deletes the column from the current Event Viewer display. Deleting a column does not mark those events for deletion from the database.

Tip To delete events that you have marked for deletion from the database, run the PruneMarkedForDeletion.pl script or the Alarm Export Utility. You can also use the Alarm Export Utility to "unmark" events marked for deletion. You can run the PruneMarkedForDeletion.pl script using database rules. For more information about database rules, see Using Database Rules, page 7-2. For more information about the Alarm Export Utility, see Using the Alarm Export Utility, page A-1.

To delete events or a column from Event Viewer, follow these steps:

Step 1 Select a cell in the Event Viewer.

The selected cell is highlighted and outlined in gray.

Step 2 To delete the selected event from the current Event Viewer display, select **Delete From this Grid** from the TOC.

Step 3 To mark selected event for deletion from the database, select **Delete From Database** from the TOC.

Step 4 To delete one column, select **Delete One Column** from the TOC.

Note You cannot delete the Count column.

18.880

Collapsing Cells

When a cell is collapsed, all branches that pass through the selected cell provide less detail. For each branch, the background color of the cells in the newly hidden column(s) changes from white to gray. Also, rows are removed as necessary to conceal the appropriate data.

Note Collapsing does not delete anything; it merely hides data from view.

Events can be collapsed by one column, by first group, or all the way (all columns). If a cell is collapsed by one column, each branch through the selected cell gives one less column of detail. If a cell is collapsed by first group, Event Viewer traverses the tree from the selected node and collapses all nodes up the branch until a node with multiple child nodes is collapsed. If a cell is collapsed all the way, all branches through the selected cell are condensed into the selected cell.

To collapse a cell, follow these steps:

Step 1 Select a cell in the Event Viewer.

The selected cell is highlighted and outlined in gray.

Step 2 To collapse a cell by one column, select **Collapse > One Column** from the TOC.

Step 3 To collapse a cell by first group, select **Collapse > First Group** from the TOC.

Step 4 To collapse a cell all the way, select **Collapse > All Columns** from the TOC.

Setting the Event Expansion Boundary

The Event Expansion Boundary dictates the number of a new eventxd5 s columns that will be expanded if the new event does not match an existing event group. The cells in an event are expanded as long as the event matches an existing event group. After there are no matches, a new row is created for the event, and the cells in the new event are expanded until the Event Expansion Boundary is reached.

The default value for the Event Expansion Boundary is one column. You can change the default value in the Preferences dialog box. For more information, see Specifying Event Viewer Preferences, page 5-17.

88.879

To set the Event Expansion Boundary, follow these steps:

Step 1 To establish a column as the Event Expansion Boundary, select a cell in that column.

The selected cell is highlighted and outlined in gray.

Step 2 Select **Set Event Expansion Boundary** from the TOC.

The Event Expansion Boundary is set. The column heading is bold.

Expanding Cells

When a cell is expanded, all *branches* that pass through the selected cell provide more detail. For each branch, the background color of the cells in the newly filled-in column(s) changes from gray to white. Also, rows are created as necessary to display the exposed data.

Event rows can be expanded by one column, and they can be expanded all the way (all columns). If a cell is expanded by one column, each branch through the selected cell gives one more column of detail. If a cell is expanded by first group, Event Viewer traverses the tree from the selected node and expands all nodes down the branch until a node with multiple children is reached. If a cell is expanded all the way, all branches through the selected cell are fully expanded.

Note Sometimes expanding events can cause many rows to be created. If the number of new rows exceeds a certain maximum, a popup window asks you to confirm that you want to continue.

To expand a cell, follow these steps:

Step 1 Select a cell in the Event Viewer.

The selected cell is highlighted and outlined in gray.

Step 2 To expand a cell by one column, select **Expand > One Column** from the TOC.

Step 3 To expand a cell by first group, select **Expand > First Group** from the TOC.

Step 4 To expand a cell all the way, select **Expand > All Columns** from the TOC.

Suspending and Resuming New Events

You can suspend new events from being added to the current Event Viewer display. You can resume receiving new events when you are ready.

To suspend or resume events, follow these steps:

Step 1 To suspend receiving new events, select **Suspend New Events** from the TOC.

Event Viewer stops querying the database for new events.

Step 2 To resume receiving new events, select **Resume New Events** from the TOC.

Event Viewer resumes querying the database for new events.

Specifying Event Viewer Preferences

Use the options in the Preferences dialog box to specify Event Viewer settings for the current Event Viewer display. To modify preferences for all Event Viewer displays, see Defining Default Event Viewer Preferences, page 5-23 and Defining Custom Event Viewer Preferences, page 5-24.

To specify the Event Viewer preferences, follow these steps:

Step 1 Select **Preferences** from the TOC.

The Preferences dialog box appears.

Step 2 To determine how long, in seconds, Event Viewer will wait for a response from the remote sensor or host before concluding that the remote sensor or host is not connected, enter a value in Command Timeout field. The default timeout value is 10 seconds.

Step 3 To specify how long, in minutes, a sensor blocks traffic from a specified source when you issue a Block command from Event Viewer, enter a value in the Time to Block field. The default value is 1440 minutes.

18.877

Step 4 Specify the subnet mask in the Subnet Mask field.

Step 5 To enable the Blank Left feature, select the **Blank Left** check box.

Step 6 To enable the Blank Right feature, select the **Blank Right** check box.

Step 7 Specify whether events are sorted by count or content.

- **a.** To sort events based on the number of events per row from highest to lowest, select the **Count** radio button.
- **b.** To sort events alphabetically based on the column to the right of the Count column, select the **Content** radio button.

Step 8 Specify the default Event Expansion Boundary in the Default Expansion Boundary field.

Step 9 To specify the maximum number of events that can be displayed in a single grid, enter a value in the Maximum Events per Grid field.

Step 10 To enable automatic collapsing of cell, select the **Auto Collapse Enabled** check box.

Note This option has no effect in the current product.

Step 11 Specify whether Event Viewer uses colors or icons to indicate event severity.

- **a.** To use colors to display event severity, select the **Color** radio button.
- **b.** To use icons to display event severity, select the **Icon** radio button.

Step 12 To enable automatic queries of the database for new events, select the **Auto Query Enabled** check box.

Step 13 To specify how often, in minutes, that Event Viewer queries the database for new events, enter a value in the Query Interval (minutes) field.

Step 14 To save your changes, click **OK**.

Viewing the Context Buffer

A context buffer records exactly what traffic was traversing the network at the time the alarmxd5 s signature was detected. Not all signatures contain context buffers. For more information, see Context Buffer, page 5-9.

To view the context buffer, follow these steps:

Step 1 Select a cell in Event Viewer.

The selected cell is highlighted and outlined in gray.

Step 2 Select **View > Context Buffer** from the TOC.

A dialog box appears. If the signature has a context buffer, the dialog box displays the context buffer information. Otherwise the dialog box displays the following message: "No context buffer data for the selected cell."

Viewing Host Names

You can view the host names that correspond to the source and destination addresses. If a host name cannot be resolved, you receive a message stating that the name cannot be resolved.

To view the host names, follow these steps:

Step 1 Select a cell in Event Viewer.

The selected cell is highlighted and outlined in gray.

Step 2 Select **View > Host Names** from the TOC.

The Hostname Resolution dialog box appears. The dialog box displays the addresses and corresponding host names, if available.

Viewing the Network Security Database

The Network Security Database (NSDB) provides detailed information about signatures, including descriptions, versions, benign triggers, and related vulnerabilities. You can access the NSDB information for a signature directly from Event Viewer.

18.875

To access the NSDB, follow these steps:

Step 1 Select a cell in Event Viewer.

The selected cell is highlighted and outlined in gray.

Step 2 Select **View > Network Security Database** from the TOC.

If there is an NSDB entry for the event you selected, the NSDB opens in a new window. Otherwise, a dialog box notifies you that there is not an NSDB entry for the event you selected and the NSDB index page opens.

Viewing Event Statistics

You can view event statistics for a cell in Event Viewer. The statistics can include the following:

- The number of events represented by the cell.
- The severity level.
- The number of child cells.
- The percentage of total events that the selected cell and its child cells represent in the current Event Viewer display.

To view event statistics, follow these steps:

Step 1 Select a cell in Event Viewer.

The selected cell is highlighted and outlined in gray.

Step 2 Select **View > Statistics** from the TOC.

The Event Statistics dialog box appears and displays the event statistics.

Refreshing Events

Based on the settings you specified in the Preferences dialog box, Event Viewer queries your monitored devices at regular intervals for new events. If you want to check for new events between intervals or if you have automatic queries disabled,

18.874

you can use the Refresh Events option to query your monitored devices for new events manually. For more information about specifying Event Viewer preferences, see Specifying Event Viewer Preferences, page 5-17.

To refresh the Event Viewer events, follow these steps:

Step 1 Select **Refresh Events** from the TOC.

The Event Viewer display refreshes to include any new events.

Step 2 Repeat Step 1 as often as you would like to check for new events.

Blocking a Host or a Network

Blocking a host causes a sensor to block all traffic emanating from the source IP address associated with the selected event. In a similar way, blocking a network causes the sensor to block all traffic emanating from the network that contains the source IP address of the selected event. Blocking is accomplished through a properly configured Cisco router. For information about removing a block, see Removing a Block, page 5-22.

To block a host or a network, follow these steps:

Step 1 To select an event whose source (a host or a network) you want to block, click the corresponding cell in Event Viewer.

The selected cell is highlighted and outlined in gray.

Step 2 To block a host, select **Block > Host** from the TOC.

The traffic is blocked for the number of minutes specified in the Preferences dialog box.

18.873

Step 3 To block a network, select **Block** > **Network** from the TOC.

The traffic is blocked for the number of minutes specified in the Preferences dialog box.

Note The network address of a blocked network is calculated by applying the network mask in the Preferences panel to the source IP address of the selected event.

Removing a Block

You can remove any blocks that you have added in Event Viewer.

To remove a block, follow these steps:

Step 1 To select the event from which you want to remove the block, select the corresponding cell in Event Viewer.

The selected cell is highlighted and outlined in gray.

Step 2 To remove a sensorxd5 s block from a host, select **Remove Block** > **Host** from the TOC.

Step 3 To remove a sensorxd5 s block from a network, select **Remove Block** > **Network** from the TOC.

Step 4 To remove all blocks, select **Remove Block** > **All** from the TOC.

Defining Event Viewer Preferences

This section describes how to define Event Viewer preferences. It also describes how to administer the preferences of Event Viewer users. This section contains the following procedures:

- Defining Default Event Viewer Preferences, page 5-23
- Defining Custom Event Viewer Preferences, page 5-24

3872

- Viewing Event Viewer Users, page 5-26
- Deleting Users from the Event Viewer Database, page 5-26

Defining Default Event Viewer Preferences

If you have administrative privileges, you can define the default Event Viewer preferences. These preferences are the default preferences used by all users. However, users can define custom preferences to reconfigure their views. For more information, see Defining Custom Event Viewer Preferences, page 5-24.

To define the default Event Viewer preferences, follow these steps:

Step 1 Select **Admin > Event Viewer**.

Step 2 Select **Default Preferences** from the TOC.

The Default Preferences page appears.

Step 3 To determine how long, in seconds, Event Viewer will wait for a response from the remote sensor or host before concluding that the remote sensor or host is not connected, enter a value in Command Timeout field. The default timeout value is 10 seconds.

Step 4 To specify how long, in minutes, a sensor blocks traffic from a specified source when you issue a Block command from Event Viewer, enter a value in the Time to Block field. The default value is 1440 minutes.

Step 5 Specify the subnet mask in the Subnet Mask field.

Step 6 Specify the default Event Expansion Boundary in the Default Expansion Boundary field.

Step 7 Enter a value in the Maximum Events per Grid field to specify the maximum number of events that can be displayed in a single grid.

Step 8 Select the **Auto Collapse Enabled** check box to enable automatic collapsing of cell.

Note This option has no effect in the current product.

Step 9 To specify how often, in minutes, Event Viewer queries the database for new events, enter a value in the Query Interval (minutes) field.

Step 10 To enable automatic queries of the database for new events, select the **Auto Query Enabled** check box.

Step 11 Specify whether Event Viewer uses colors or icons to indicate event severity.

a. To use colors to display event severity, select the **Color** radio button.

b. To use icons to display event severity, select the **Icon** radio button.

Step 12 To enable the Blank Left feature, select the **Blank Left** check box.

Step 13 To enable the Blank Right feature, select the **Blank Right** check box.

Step 14 Specify whether events are sorted by count or content.

a. To sort events based on the number of events per row from highest to lowest, select the **Count** radio button.

b. To sort events alphabetically based on the column to the right of the Count column, select the **Content** radio button.

Step 15 Click **Apply**.

The preferences you specified are the default preferences used by all Event Viewer users.

Defining Custom Event Viewer Preferences

You can define custom Event Viewer preferences that override the default Event Viewer preferences. Custom Event Viewer preferences affect only the Event Viewer displays opened by the user for whom the preferences were defined.

To define custom Event Viewer preferences, follow these steps:

Step 1 Select **Admin > Event Viewer**.

Step 2 Select **Your Preferences** from the TOC.

The Your Preferences page appears.

Step 3 To determine how long, in seconds, Event Viewer will wait for a response from the remote sensor or host before concluding that the remote sensor or host is not connected, enter a value in Command Timeout field. The default timeout value is 10 seconds.

Step 4 To specify how long, in minutes, that a sensor blocks traffic from a specified source when you issue a Block command from Event Viewer, enter a value in the Time to Block field. The default value is 1440 minutes.

Step 5 Specify the subnet mask in the Subnet Mask field.

Step 6 Specify the default Event Expansion Boundary in the Default Expansion Boundary field.

Step 7 Enter a value in the Maximum Events per Grid field to specify the maximum number of events that can be displayed in a single grid.

Step 8 Select the Auto Collapse Enabled check box to enable automatic collapsing of cell.

Note This option has no effect in the current product.

Step 9 To specify how often, in minutes, that Event Viewer queries the database for new events, enter a value in the Query Interval (minutes) field.

Step 10 To enable automatic queries of the database for new events, select the Auto Query Enabled check box.

Step 11 Specify whether Event Viewer uses colors or icons to indicate event severity.

- **a.** To use colors to display event severity, select the Color radio button.
- **b.** To use icons to display event severity, select the Icon radio button.

Step 12 To enable the Blank Left feature, select the Blank Left check box.

Step 13 To enable the Blank Right feature, select the Blank Right check box.

Step 14 Specify whether events are sorted by count or content.

- **a.** To sort events based on the number of events per row from highest to lowest, select the Count radio button.
- **b.** To sort events alphabetically based on the column to the right of the Count column, select the Content radio button.

Step 15 To save your changes, click **Apply**.

Your Event Viewer displays will use the preferences you defined.

Step 16 To revert to the default Event Viewer preferences, click **Reset to Defaults**.

Your custom preferences are overwritten by the default preferences used by all Event Viewer users.

Viewing Event Viewer Users

You can view a list of users that have custom Event Viewer preferences stored in the database.

To view a list of Event Viewer users, follow these steps:

Step 1 Select **Admin** > **Event Viewer**.

Step 2 Select **Users** from the TOC.

The Users page appears. The users are listed in a table on this page.

Deleting Users from the Event Viewer Database

To clean up your database, you can delete preferences for users who no longer view events. Only the event viewing preferences for that user are deleted from the database.

Note You must have administrative privileges to delete a user from the database.

Tip Security Monitor administers only Event Viewer user records. To administer user permissions, you must use IDS MC. For more information, refer to *Using Management Center for IDS Sensors*.

18.868

To delete a user from the Event Viewer database, follow these steps:

Step 1 Select **Admin** > **Event Viewer**.

Step 2 Select **Users** from the TOC.

The Users page appears.

Step 3 To select which user to delete, select the check box next to the user ID.

Note You can select all users by clicking Select All.

A check mark appears next to the user ID that you selected.

Step 4 To delete the selected user, click **Delete**.

The event viewing preferences for the selected user are deleted from the Event Viewer database.

18.867

CISCO SYSTEMS

Home | Logged In | Profile | Contacts &

18.866

TechnicalSupport GO

Select a L

TECHNICAL SUPPORT
Cisco Secure Encyclopedia

TECHNICAL SUPPORT

Search:

Search A

Toolkit:

Feedback

Related T
TAC Case
TAC Case
TAC Case
Dynamic (

Cisco Secure Encyclopedia

Terms of Use:

1. These terms of use state Cisco's requirements for use of the Website. Use of these or any other interactive areas of Cisco's Website constitutes acceptance of these legal terms and conditions.

2. You agree that the information you include in your user profile for the Website (which will be accessible to other users of the Website) will be accurate and not misleading. You agree not to impersonate any person or entity [including any Cisco or third party moderator of this forum] or to misrepresent your affiliation with any person or entity.

3. In providing this Website, it is Cisco's intention to maintain an informative and valuable service that meets the needs of its users and is free from verbal abuse, violations of privacy and other harmful activities. Cisco prohibits posting of threatening, libelous, defamatory, obscene, pornographic, or other material that would violate any law. You agree not to post or submit any information that may infringe upon or violate any third party's intellectual property rights, including, without limitation, copyright, patent, trademark, trade secret, moral right, proprietary rights or rights of publicity or privacy, or contain obscenity, pornography or be otherwise inflammatory, or which contains any viruses, Trojan horses, worms, time bombs, cancelbots or other computer programming routines that are intended to damage, detrimentally interfere with, surreptitiously intercept or expropriate any system, data or personal information, or which is materially false, misleading or inaccurate.

4. We welcome your participation in this Website. However, in so participating, please be aware that you are solely responsible for your content or other information you distribute, post, include, link to or otherwise upload to the Website. You agree that Cisco is only acting as a venue and that Cisco has no liability related to the content of any posting or any other user-created content on its Website, whether or not arising under the laws of copyright, libel, privacy, obscenity, or otherwise. By its very nature, the information available in this Website is changed frequently, may be offensive, harmful or inaccurate and in some cases may be mislabeled or deceptively labeled. We do not make any representation or warranty, express or implied, as to the accuracy, timeliness or completeness of such information nor do we undertake to update or correct such information. We recommend that you confirm all information you obtain from this Website.

5. Messages posted to open forums are not pre-screened. Cisco, or a third party organization Cisco may appoint to moderate a particular forum, each reserves the right, although it does not have the obligation, to monitor content throughout the Website (or a specific forum, as applicable) and to remove any content which Cisco or such third party organization, in its sole discretion, determines to be harmful, offensive, defamatory, libelous, illegally threatening or harassing or which contains obscenity, pornography or is otherwise inflammatory; which may create liability for Cisco; or is unrelated to the intention of the community or otherwise in violation of Cisco's policies of acceptable use.

6. You agree not to use this Website to violate any applicable local, state, national or international law, including, but not limited to, regulations promulgated by the U.S. Securities and Exchange Commission, any rules of any national or other securities exchange, including, without limitation, the New York Stock Exchange, the American Stock Exchange or the NASDAQ, the export control laws of the United States or your country of residence regarding the transmission of technical data, and any regulations having the force of law.

7. You agree that you will not interfere with or disrupt this Website or any servers or networks connected to this Website or disobey any

rules, policies or procedures of any networks connected to this Website.

8. You should assume that everything hosted on this Website is the intellectual property of Cisco and our suppliers. All rights to this information are reserved. All reproduction, total or partial, of the contents of this Website by any means whatsoever and without the written consent of Cisco is illegal.

18.865

9. Cisco and our suppliers provide this Website "as is." WE AND OUR SUPPLIERS DISCLAIM ALL WARRANTIES OF ANY KIND EITHER EXPRESS OR IMPLIED, INCLUDING WITHOUT LIMITATION, THE WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, TITLE AND NON-INFRINGEMENT. WE DO NOT GUARANTEE CONTINUOUS OR UNINTERRUPTED ACCESS TO THE WEBSITE OR ANY FORUM AND MAKE NO WARRANTY AS TO THE OPERATION, FUNCTIONALITY, OR AVAILABILITY OF THE WEBSITE, OR THAT THE WEBSITE WILL BE ERROR-FREE, OR THAT DEFECTS WILL BE CORRECTED. ANY MATERIAL DOWNLOADED OR OTHERWISE OBTAINED THROUGH THE USE OF THE WEBSITE IS DONE AT YOUR OWN DISCRETION AND RISK AND THAT YOU WILL BE SOLELY RESPONSIBLE FOR ANY DAMAGE TO YOUR COMPUTER SYSTEM OR LOSS OF DATA THAT RESULTS FROM THE DOWNLOAD OF ANY SUCH MATERIAL. NO ADVICE OR INFORMATION, WHETHER ORAL OR WRITTEN, OBTAINED BY YOU FROM CISCO OR THROUGH OR FROM THIS WEBSITE OR ANY FORUM SHALL CREATE ANY WARRANTY NOT EXPRESSLY STATED IN THESE TERMS OF USE.

10. These terms of use shall be governed in all respects by the laws of the State of California as such laws are applied to agreements entered into and to be performed entirely within California between California residents. Any action or proceeding arising from or relating to these terms and conditions must be brought in a federal court in the Northern District of California or in state court in San Francisco County, California, and you irrevocably submit to the jurisdiction and venue of any such court in any such action or proceeding. If any provision of these terms of use is held to be invalid or unenforceable, such provision shall be struck and the remaining provisions shall be enforced. Our failure to act with respect to a breach by you or others does not waive our right to act with respect to subsequent or similar breaches. These terms of use may be amended or changed by us at our discretion. You agree that your continued use of this Website after such amendment or change shall signify your acceptance of any such change.

11. The services offered under these use policies are provided by Cisco Systems, Inc. 170 W. Tasman Drive, San Jose, CA 95134. If you are a California resident, you may have this same information emailed to you by sending a letter to the address above with your email address and a request for the information. If you notice any user that is violating these terms of use, please contact us at csec-feedback@cisco.com.

12. Please report any violations of these terms of use to: csec-feedback@cisco.com.

13. Cisco reserves the right at any time and from time to time to modify or discontinue, temporarily or permanently, this Website (or any part thereof) with or without notice. Cisco shall not be liable to you or to any third party for any modification, suspension or discontinuance of this Website.

14.If applicable, you will receive a password and account designation upon completing this Website's registration process. You are responsible for maintaining the confidentiality of the password and account, and are fully responsible for all activities that occur under your password or account. You agree to immediately notify Cisco of any unauthorized use of your password or account or any other breach of security. Cisco will not be liable for any loss or damage arising from your failure to comply with these provisions. You agree that Cisco, in its sole discretion, may terminate your password, account (or any part thereof) or use of the Website, and remove and discard any content on the Website, for any reason, including, without limitation, for lack of use or if Cisco believes that you have violated or acted inconsistently with these terms of use.

15. By submitting your content to the Website (including, but not limited to sending messages through the Website or posting messages to message boards, forums, and chat rooms), you hereby grant to Cisco a worldwide, perpetual, irrevocable, royalty-free, sublicenseable (through multiple tiers) right to exercise all copyright and publicity rights, in any

media now known or not currently known, with respect to any content you provide to us in any public space on this Website or in any email that you voluntarily provide to us through this Website.

16. We will treat any feedback or suggestions you provide us as non-confidential and non-proprietary. Thus, please do not submit confidential or proprietary information to on this Website.

18.864

17. You agree to indemnify and hold Cisco and its subsidiaries, affiliates, officers, agents and employees, harmless from any losses, damages, expenses or liability, including reasonable attorneys' fees, arising out of any claim by a third party based on or arising from any content or information you submit, post, transmit or make available through this Website, your violation of the these terms of use, or your violation of any rights of another.

I Accept | I Decline

BUSINESS STRATEGIES & SOLUTIONS | NETWORKING SOLUTIONS & PROVISIONED SERVICES | F
| TECHNOLOGIES | ORDERING | TECHNICAL SUPPORT | LEARNING & EVENTS | PARTNERS & RESI
Home | Logged In | Profile | Contacts & Feedback | Site Help