HP Integrated Archive Platform Version 2.1

Administrator Guide

Part number: PDF
First edition: March 2010

Legal and notice information

Contents

Figures

Tables

About this guide

This guide provides information about administering the HP Integrated Archive Platform (IAP). For information on administering HP Email Archiving Software (EAs) for Exchange or Domino, see the respective administration guide included on the documentation CD in those products.

Intended audience

This guide is intended for HP Integrated Archive Platform administrators.

Related documentation

HP provides the following related documentation.

For administrators and installers:

- *HP Integrated Archive Platform Installation Guide* (available to HP personnel installing IAP)
- Online help for the Platform Control Center (PCC) (the help is a subset of the administrator guide)
- HP Email Archiving software for Microsoft Exchange or IBM Lotus Domino administration guide (located on the documentation CD included in those products)
- HP Email Archiving software for Microsoft Exchange or IBM Lotus Domino installation guide (available to HP personnel installing those products)

For users:

- *HP Integrated Archive Platform User Guide* (located on the documentation CD)
- Online help for the IAP Content Search and Retrieve Web Interface, the help is a subset of the user guide
- HP Email Archiving software for Microsoft Exchange or IBM Lotus Domino user guide (located on the documentation CD included in the those products)

For developers:

This release includes the following guides for developers, which are available at the HP Developer and Solution Partner Program web site at http://www.hp.com/go/ilmdspp/:

- *HP Integrated Archive Platform Web Service API Developer Guide*
- *HP Integrated Archive Platform ILM Object Storage API Developer Guide*

Document conventions and symbols

Table 1 Document conventions

Convention	Element
Medium blue text: Related documentation	Cross-reference links and email addresses

Convention	Element
Medium blue, underlined text (http://www.hp.com)	Web site addresses
Bold font	• Key names • Text typed into a graphical user interface (GUI) element, such as into a box • GUI elements that are clicked or selected, such as menu and list items, buttons, and check boxes
Italic font	Text emphasis
`Monospace font`	• File and directory names • System output • Code • Text typed at the command line
`Monospace, italic font`	• Code variables • Command-line variables
`Monospace, bold font`	Emphasis of file and directory names, system output, code, and text typed at the command line

⚠ WARNING!

Indicates that failure to follow directions could result in bodily harm or death.

△ CAUTION:

Indicates that failure to follow directions could result in damage to equipment or data.

IMPORTANT:

Provides clarifying information or specific instructions.

NOTE:

Provides additional information.

TIP:

Provides helpful hints and shortcuts.

Support

You can visit the HP Software Support web site at: http://www.hp.com/go/hpsoftwaresupport

HP Software Support Online provides an efficient way to access interactive technical support tools. As a valued support customer, you can benefit by using the support site to:

- Search for knowledge documents of interest
- Submit and track support cases and enhancement requests
- Download software patches
- Manage support contracts
- Look up HP support contacts
- Review information about available services
- Enter into discussions with other software customers
- Research and register for software training

Most of the support areas require that you register as an HP Passport user and sign in. Many also require a support contract.

To find more information about access levels, go to: http://support.openview.hp.com/new_access_levels.jsp

For continuous quality improvement, calls may be recorded or monitored.

Subscription service

HP strongly recommends that customers register online using the Subscriber's choice web site: http://www.hp.com/go/e-updates.

Subscribing to this service provides you with email updates on the latest product enhancements, newest driver versions, and firmware documentation updates as well as instant access to numerous other product resources.

1 IAP product overview

This chapter describes key concepts involving the HP Integrated Archive Platform.

Introduction to IAP

HP's Integrated Archive Platform (IAP) is a tightly integrated, fault-tolerant, secure system of hardware, software, and services for archiving and retention management. IAP provides scalable, long-term information retention combined with high-speed, interactive search and retrieval capabilities to help organizations meet their regulatory requirements.

Understanding the components

In order to create a framework for understanding the components of IAP and how they work together, this chapter covers a number of high level themes, such as a conceptual overview, key benefits and features, software components, hardware components, network architecture and data flow. Taking a look at the anatomy or architecture of IAP and its components will support the goal of operating, maintaining, optimizing and troubleshooting the IAP system more effectively.

Conceptual overview

In broad conceptual terms, the IAP can be thought of as having three aspects:

1. Enterprise content
2. Application archiving software
3. Working with archived Enterprise content effectively through the capabilities of IAP

Enterprise content

IAP serves as a platform for many types of enterprise content. IAP automatically and actively archives files, email (Exchange, Domino), documents and other types of data, see Appendix B on page 163.

Application archiving software

HP provides a variety of application archiving software that works with IAP to capture and store data. These software components are focused on ingesting data into IAP. The IAP can be precisely customized and configured to support an organization's archiving and retention policies.

IAP works with the following application archiving software:

HP Email Archiving software (EAs) for Exchange

HP Email Archiving software (EAs) for Exchange is mail administration software that archives messages and attachments, and stores the messages in IAP. Users can search for and retrieve their archived messages directly from Outlook. HP EAs for Exchange fully integrates with Microsoft Exchange and

Outlook. Outlook Integrated Archive Search enables search of the IAP repositories directly from Outlook by key words and date ranges.

HP Email Archiving software (EAs) for Domino

HP Email Archiving software (EAs) for Domino is mail administration software that archives messages and attachments, and stores the messages in IAP. Users can retrieve their archived messages directly from IBM Lotus Notes and iNotes, and EAs for Domino is fully integrated with Lotus Notes and iNotes.

HP File Archiving software

HP File Archiving software (FAs) is a policy-based file archiving solution that enables you to automatically archive files, web and application server data from a Windows environment for migration into IAP. (FAs was formerly known as FMA.)

Archiving methods

HP EAs for Exchange and HP EAs for Domino support the following archiving methods:

- Compliance Archiving. In EAs for Exchange, compliance archiving captures all inbound and outbound messages. EAs for Exchange provides two different mechanisms to implement compliance archiving. One technique copies each inbound and outbound message to a journal mailbox which the Archive Gateway scans periodically to archive the contents on the IAP. The other technique, available on Exchange 2007 and later, uses Simple Mail Transport Protocol (SMTP) to relay every message directly to the Archive Gateway which then archives them to the IAP. In EAs for Domino, compliance archiving can capture messages at an individual, group, and/or enterprise level when the Advanced Filtering feature is used.
- Selective archiving (mail database mining). Selective archiving helps administrators reduce the amount of primary storage used on the mail servers. It provides policies to mine messages from users' mailboxes. Administrators can centrally configure policies, including quota archive thresholds, message age or size, sender/recipient, specific folders, and keywords. Policies can be set at the individual, group, and/or enterprise-wide level. Messages are automatically archived when the policies are enabled. After archiving, links to the archived messages (called tombstones) can be placed in mail databases.
- Archiving of inactive mail databases. EAs for Exchange and EAs for Dominos also provide optional tombstoning of messages in inactive PST and NSF databases.

Working with enterprise content on IAP

An overarching focus of IAP is the ability to work effectively with stored enterprise content. The many abilities and benefits of working with stored enterprise content on IAP include:

- Storage
- Indexing
- Single instancing
- Retention policies
- Encryption
- Authentication
- Replication
- Monitoring
- Search
- Retrieval

- Backup/restore

Key benefits

Using IAP helps you:

- Mitigate the costs and business impact of e-discovery preparation, legal response, and regulatory compliance. When entering litigation, companies need to be able to show that their emails have not been tampered with and to be able to pull the information that pertains to the particular case.
- Effectively manage, retain, and retrieve information. IAP turns data into information by providing fast retrieval of archived content. Data retrieval from the archival system is accelerated using storage grid architecture that evenly distributes full text indexes and original content across many Smartcells (storage cells). Searches are executed across all Smartcells in parallel.
- Increase email, file, and application server performance by offloading documents from the server to the IAP. Storing mail and files on IAP provides a faster means for query and retrieval. IAP supports free text searching through any Web browser.
- Reduce the cost and complexity of managing e-mail, file system storage, and quotas. IAP contains standardized server hardware throughout the system, eliminates redundant servers, and automatically migrates email to lower cost storage according to value of data.
- Defer capital expenditures on new primary storage systems.

Features

IAP automatically and actively archives files, email messages, and other types of data. It provides interactive data querying to quickly and easily search for and retrieve archived data according to various criteria. Authorized users are granted access to archived information through IAP's Web user interface.

IAP combines HP server and grid storage technology, native content indexing, search, and policy management software onto a single, factory-assembled rack system for easy installation and simpler maintenance.

The following is a list of high-level features of IAP:

- Grid computing architecture. HP IAP's grid computing architecture provides scalability in performance and capacity. Searches are executed across all Smartcells in parallel.
- Integrated distributed content indexing and search tools. Comprehensive integrated search capabilities that leverage the IAP's distributed indexing and grid architecture technologies ensure fast access to consolidated and centralized multi-content information silos, such as email and files, for hundreds of thousands of users across hundreds of terabytes and billions of objects.
- Storage reduction algorithms. IAP uses storage reduction algorithms including data compression and object level single instancing. This enables IAP to provide an efficient means for identifying duplicate emails or documents and removing them from the archive.
- Search tool integrated with Exchange Outlook and Lotus Notes. Users can search for their archived messages directly from within their email application; they don't have to learn or bring up a separate utility.
- Supports business policy-driven automated archiving applications. Selectively archive emails and documents according to flexible policies.
- Reduced administration overhead for managed applications such as Microsoft Exchange servers. HP IAP Platform Control Center (PCC) software is included to monitor the status and administration of the IAP system.
- Tamper-resistant environment. Digital signatures, timestamps, and data retention policies are used to protect data against erasure and tampering. Data cannot be deleted from the archive until the end of the retention period.

- Data replication and mirroring. Data is mirrored synchronously across a group of Smartcells for increased data availability.
- Supports unstructured, semi-structured, and structured data such as email and files.
- Mitigates regulatory and legal risks and supports corporate governance initiatives. Data stored in HP IAP is managed by retention policies and is tamper-resistant in alignment with SEC and regulatory rules. Data can be searched with proper credentials. Litigation holds can be applied when required.
- Flexible retention management. Preservation of information through flexible retention policies ensures proper disposition of data at the end of its retention period whether it is to meet compliance, legal hold, or governance requirements.

Software components

There are two main types of software components for working with IAP – user applications and administrator applications.

Applications for users

The following application is available for all IAP users.

Table 2 IAP application for users

Application	Tasks
IAP Web Interface	Use a web browser to: • Search for documents archived on the system and save and reuse search criteria and results. • Send archived emails to the email client. • Download files to a local computer or network folder. • Export emails and files when the appropriate export tool is installed on the user's system. • Place a legal hold on emails so they are retained indefinitely. • (Compliance officers) View the AuditLog, the compliance system log.

To interact with the system, the following optional EAs and FAs applications are available for users. Outlook Plug-In, Notes Client Plug-In, Local Cache, Export Search, and IAP File Export are installed on users' computers. The Outlook Plug-In can also be installed on a server running Citrix Presentation Manager.

Table 3 EAs and FAs applications for users

Application	Tasks
Outlook Plug-In (customer option)	Search for, view, and work with archived emails using Outlook with a Microsoft Exchange mail server. Export email from the IAP to a PST file.
OWA Extension (customer option)	View and work with archived emails using Outlook Web Access with a Microsoft Exchange mail server.

Application	Tasks
Notes Client Plug-In Local Cache Export Search (customer options) (Windows users only)	View and work with archived emails using Lotus Notes with an IBM Domino mail server. Access archived email offline while traveling. Export email from the IAP to a Notes NSF file.
DWA Extension (customer option)	View and work with archived emails using iNotes (Domino Web Access) with an IBM Domino mail server.
IAP File Export (customer option)	Export files migrated using FAs to the local computer or a network folder.

Tools for administrators

The IAP provides the following troubleshooting and administrative features.

Table 4 IAP tools for administrators

Feature	Tasks
Platform Control Center (PCC)	Monitor and troubleshoot system status and performance, configure system options, and manage user accounts using the web-based administrator interface. See "Introduction to Platform Control Center (PCC)" on page 31.
Command-line interface	Allows HP authorized support personnel to configure and enable system options from the command line.
AuditLog	Enable the AuditLog for regulatory compliance. See "Enabling the AuditLog" on page 127.
External access	View and configure EAs for Exchange options using VNC access to the Archive Gateway. See "External access" on page 121.
Remote server management	Administer remote management of IAP servers using HP ProLiant Onboard Administrator Powered by Lights-Out 100 or Integrated Lights-Out 2. See Chapter 13 on page 143.

The following EAs, FAs, and backup applications are available for use with the IAP.

Table 5 Applications for administrators

Application	Tasks
HP Email Archiving software for Exchange	Archive journaled mail, configure archiving rules, and batch process multiple Outlook PST files for Microsoft Exchange. See the *HP Email Archiving software for Microsoft Exchange Administrator Guide* included with the EAs for Exchange product on the documentation CD.
HP Email Archiving software for Domino	Archive journaled mail, configure archiving rules, and batch process mail databases for IBM Lotus Domino. See the *HP Email Archiving software for IBM Lotus Domino Administrator Guide* included with the EAs for Domino product on the documentation CD.
HP File Archiving software	Archive files, web, and application server data from Windows systems.

Application	Tasks
HP Data Protector	Automates high-performance backup and recovery with scalable data protection.

Hardware components

IAP base system

The IAP base system consists of the following components:

- 42U rack
- (9x) management servers
- (1x) Universal Smartcell
- (1x) KVM/Mon
- (1x) 5406 switch
- (1x) 2510G-48 switch for iLO
- Total U used in base rack: 19U
- Available U for upgrade: 23U

The management servers consist of the following:

- Gateway (option)
- Firewall
- Load Balancer
- Query portal (HTTP)
- Store portal (SMTP)
- Metaserver
- Cloud router
- Database (DB2)
- PCC (Platform Control Center)
- Kickstart Server

The base system ships with DL360 G6 servers for all servers except the Universal Smartcell, which uses DL180 G6 servers. It includes one pair of Smartcells which provides 5 TB of storage. Two additional pairs of Smartcells can be ordered to fully load the base system. This increases the storage capacity from 5 TB to 15 TB.

Expansion rack

42U rack

(10x) Universal Smartcells

Total: 40U

Total capacity of one expansion rack: 50 TB

Up to 21 expansion racks per HP IAP system

Maximum capacity of an IAP is now 10650 TB (213 groups times 5 TB)

Performance upgrades (for example additional SMTP or HTTP portals) can optionally be housed in expansion racks

The expansion rack can hold an additional ten pairs of Smartcells providing 50 TB of storage capacity.

Additional options

The base system can accommodate up to three Smartcell groups. If you choose to grow or add Smartcells, you would be required to have an expansion rack. An expansion rack can accommodate up to ten pairs of Smartcells. It can also hold performance upgrades, such as an additional query portal or store portal performance upgrade. You can mix and match on the expansion rack.

Additional factory-integrated options

The following are available as factory-integrated options:

- 5 TB Smartcell storage capacity upgrade
- Query (HTTP) portal performance upgrade
- Store (SMTP) portal performance upgrade
- Replication performance upgrade
- IAP backup option (a dedicated tape backup solution)

Expansion Capacity Performance upgrade option

The Expansion Capacity Performance upgrade option allows for improved performance with a larger (expanded capacity) IAP.

Replication option

Another option is an IAP replication option for one system. The replication option is a software license that supports the replication of a base system, local or remote.

Network architecture

The following diagram shows the overall picture of the network architecture. It is a basic network diagram of IAP. The VLAN 500 is the backbone of the IAP network. Users connect through the firewall/NAT. Administrative functions connect through the PCC. With the PCC, you can monitor the IAP and its health, the state of the servers, and whether or not IAP is storing or not storing.

IAP 2.1 Hardware diagram

All systems are Linux-based, except for the optional HP Gateway server, which runs Windows Server software.

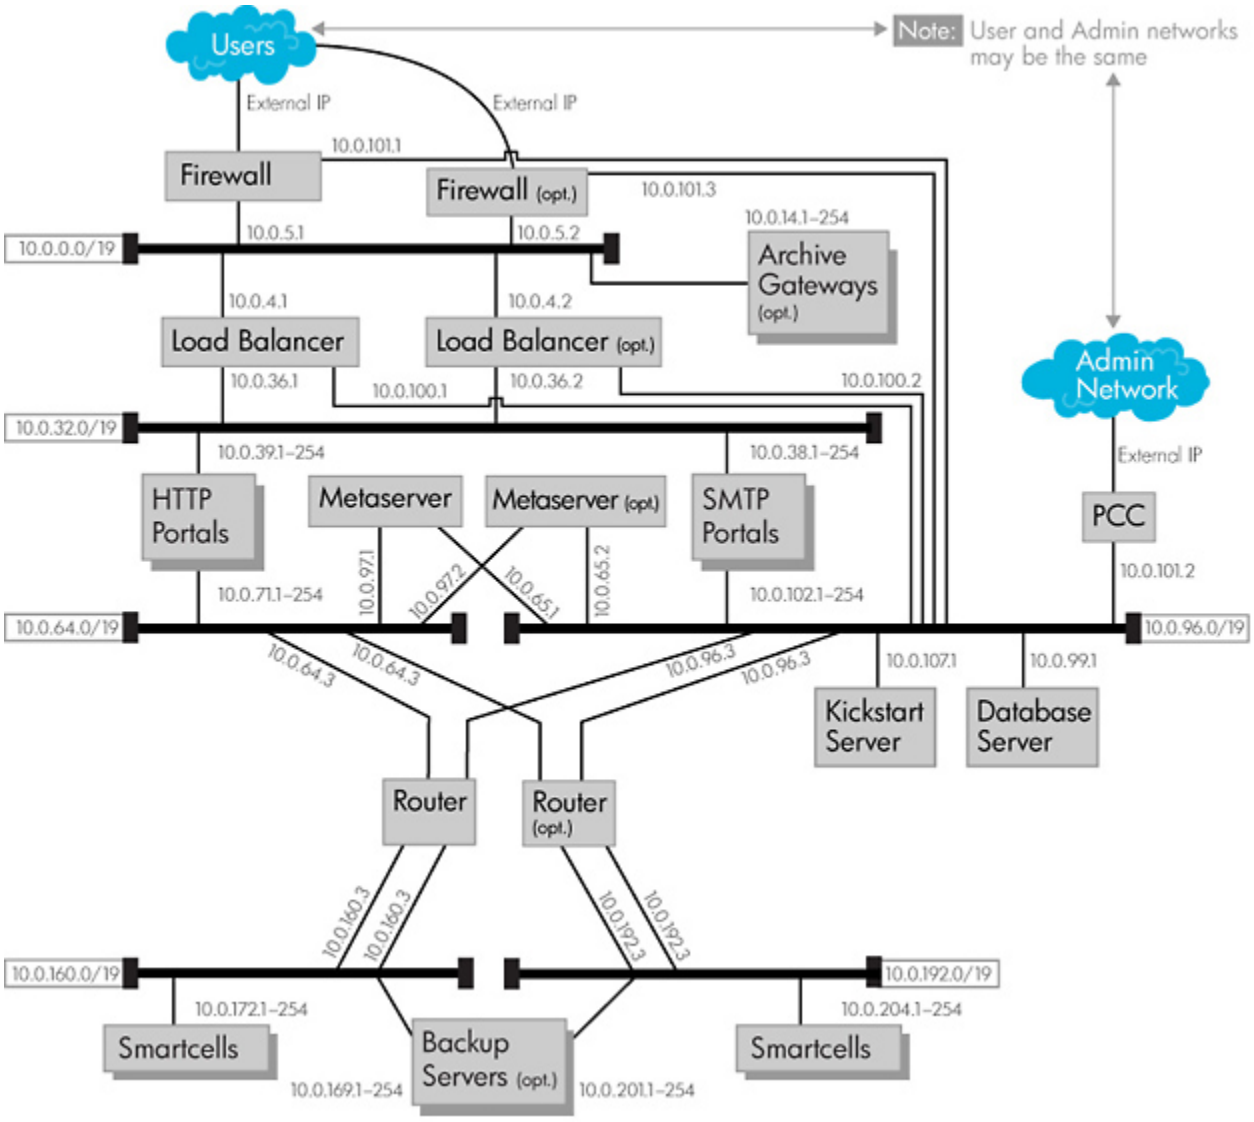

Figure 1 Architecture diagram

Virtual LANS (VLANS)

IAP uses Virtual LANS (VLANS), or Virtual Local Area Networks. The VLANS connect the management servers to one another. The following indicates how various components are connected through the VLANS:

- VLAN 200: Connects the Firewall/NAT with the Load Balancer and the Gateway servers
- VLAN 300: Connects the Load Balancer with the HTTP portal and the SMTP portal
- VLAN 400: HTTP portal that goes to VLAN 400 and can talk through the VLAN 300
- VLAN 500: The backbone of the IAP network. Connects with the PCC server, the database server, kickstart server, cloud router, and metaservers.
- VLAN 600: Connects to Cloud Router 1 and backup servers
- VLAN 700: Connects to Cloud Router 2 and backup servers

Grid architecture

IAP's grid storage implementation using Universal Smartcells is a distributed computer system of self-contained, all-inclusive data repositories. It enables intelligent distribution of hundreds of terabytes of content and billions of objects across a grid of Smartcells. It allows parallel searches, so that within three to five seconds, IAP can perform a quick search across terabytes of referenced information.

IAP's grid computing architecture helps manage explosive growth in information by consolidating and centralizing silos of information throughout the enterprise. It ensures the highest levels of performance as the number of users or document traffic into HP IAP increases.

Part of the grid architecture implementation involves domains and repositories. Domains are sub-groups of Smartcells within the larger grid architecture. For example, domain 1 may have four pairs of Smartcells, while domain 2 may have three pairs of Smartcells (both domains being in the same IAP grid). Each domain can have distinct policies, representing different divisions of a large organization. Repositories are logical entities spanning one or more Smartcells. Multiple repositories can exist within a single domain. A repository might represent a single user's mailbox, or it might represent the mailboxes of a group of users. Access control determines who has access to a particular domain or repository.

Management servers

The basic IAP system includes nine management servers. They are built using DL360 G6 servers. The following section briefly describes each management server's primary functions.

Firewall/NAT

The Firewall/NAT (Network Address Translation) restricts access to IAP. It leads out to users. When information is stored it goes through the firewall/NAT, which provides network address translation between internal and external networks. The firewall/NAT static IP address is 10.0.101.1. You can have a second firewall.

Load Balancer

The Load Balancer balances incoming traffic to the SMTP and HTTP portals. Its IP address is 10.0.100.1. You can have a second load balancer.

HTTP portal

The HTTP portal is used for querying and retrieving documents. Queries first pass through the firewall, then to the Load Balancer. The Load Balancer selects an available HTTP portal. The HTTP portal queries the metaserver to identify which Smartcell groups contain the document.

The HTTP portal's IP address is 10.0.71.1. Additional HTTP portals would use IP addresses that increase incrementally by the last digit. For example: 10.0.71.2, 10.0.71.3, and so forth.

SMTP portal

The SMTP portal is used for storing messages. It identifies which repositories have access to messages. Its IP address is 10.0.102.1. Additional SMTP portals would use IP addresses that increase incrementally by the last digit. For example: 10.0.102.2, 10.0.102.3, and so forth.

Metaserver

The metaserver manages Smartcell allocation, system state, and Bias Dimensional Indexing (BDI) mapping. It determines if additional Smartcells need to be allocated for storage. The metaserver continuously monitors the Smartcell states; for example, suspended state or stopped state.

The metaserver also tracks which IAP domains correspond to which Smartcells and users. For example, certain Smartcells belong to domain 1. Certain other Smartcells belong to domain 2. The metaserver communicates with the DB2 server and with the HTTP and SMTP portals to determine which Smartcell is applicable to each store or query request.

The metaserver's IP address is 10.0.97.1. It is possible to have a second metaserver. In that case, you would have a primary or master metaserver. The additional or secondary metaserver would use IP address 10.0.97.2.

Cloud router

The cloud router routes internal network traffic to and among Smartcells. In order for data to be stored or retrieved, it goes through the cloud router and it will either go to the Smartcells via the VLAN 600 or VLAN 700. Its static IP address is 10.0.96.3. It is possible to add a second cloud router.

Database server (DB2)

The Database server stores usernames, passwords, and other access-control information. It stores routing rules for messages entering the repository, reflecting who will have access to that repository and message. The database server also stores query results. Those query results can be pulled out at a later time. It stores Smartcell state information. Smartcell state information is also stored in the Smartcell itself which supersedes the state in the DB2 database. The database server also controls the replication flow between the primary IAP and the replica IAP in cases in which the replication option is utilized. The DB2 server's IP address is 10.0.99.1.

Kickstart server

The kickstart server acts as a hub for all the IAP management servers. It stores the operating systems, software applications and configuration files needed for each of the management servers. When one of the management servers performs a PXE boot, it goes through the kickstart server. Each management server has its own MAC address. For any given MAC address, the kickstart server knows the management server, the operating system (OS) of the server, the applications that the server runs, and the configuration files the server requires. The kickstart server's IP address is 10.0.107.1.

Platform Control Center (PCC)

The PCC provides a web-based interface for administrators to monitor and troubleshoot the IAP and manage user accounts. It also provides a command-line interface. Common administrator tasks include:

- Checking overall system health and performance
- Checking Smartcell health and performance
- Monitoring system status and RAID support
- Starting, stopping, and restarting system servers
- Synchronizing user accounts
- Managing user accounts
- Monitoring system alerts

The PCC's IP address is: 10.0.101.2.

Archive Gateway (option)

The Archive Gateway server is optional, based on customer requirements. It contains the software components necessary for archiving email. It is a connector to the IAP and drains email messages from production mail servers (Domino or Exchange). The Archive Gateway also creates message stubs or tombstones in the client mailboxes if selective archiving is implemented.

There are two types of Archive Gateways, one for Exchange and one for Domino. For Exchange, the Archive Gateway contains HP Email Archiving software for Exchange (HP EAs for Exchange). The IP address for the Exchange Archive Gateway is 10.0.14.x.

For Domino, the Archive Gateway contains HP Email Archiving software for Domino. It will not have the same IP address used in the Exchange environment because it sits outside of the IAP firewall.

Backup server (option)

Customers can choose to add the backup option, which is a backup server with HP Data Protector software. The backup server performs backups of the IAP's Smartcell and DB2 data and IAP configuration to tape. You can have one or more backup servers.

The backup server's IP addresses are 10.0.169.x and 10.0.201.x (one interface on each Smartcell network).

Universal Smartcell

The Universal Smartcell is the basic element of storage for IAP. It is configured on DL180 G6 servers. Two Smartcells make one Smartcell group; Smartcells are always configured as a pair – a primary and a secondary Smartcell. The Smartcells are "intelligent storage" servers that handle their own data operations, including capture, indexing, compression, storage, searching and retrieval. A Universal Smartcell group has up to 5 TB of storage capacity and it is in a RAID 6 configuration with ADG protection.

The IP address for the primary Smartcell can be 10.0.172.1 or 10.0.204.1, and it is accessed through VLAN 600. The IP address for the secondary Smartcell can be 10.0.172.1 or 10.0.204.1 (but it must not be the same as the primary), and it is accessed through VLAN 700. For situations in which customer needs require additional Smartcells, additional primary Smartcells will use IP addresses of 10.0.172.2, 10.0.172.3, and so forth. Additional secondary Smartcells will use IP addresses of 10.0.204.2, 10.0.204.3, and so forth.

Data flow

This section describes the general flow of data going into and coming out of IAP.

Store path

The store path is a part of the data flow in which data goes into IAP. It refers to email and files being archived. The actual unit of data that the store path works with is a bitfile. Before describing the data flow in more detail, we need to understand the bitfile.

Bitfile

A bitfile is a basic document storage unit used by the IAP. It is a four-part ZIP file consisting of these parts:

- Manifest
- Data
- Metadata
- Sig

Manifest

The manifest is a plain text date. It contains three comma-separated pieces of information. The first piece of information is a plain-text date (for example, Tue Jul 10 17:45:18 JST 2007); the second piece is a version number (for example, 1); the third is a key type number (for example, 1). An example of a manifest looks like this:

(Tue Jul 10 17:45:18 JST 2007,1,1)

Data

This is the original data. For email, messages are archived in the RFC822 format.

Metadata

IAP adds metadata to each bitfile. The metadata describes the bitfile's data; it consists of dates, repositories, an indexable structure of the document, and e-mail envelope information. The metadata is originally added by the SMTP portal, which scans the data being streamed into the bitfile. Metadata is primarily used for indexing and access control.

Sig

The Sig is a digital signature, used to ensure that the digital content has not been tampered with. The way it is computed ensures that the data content's digital signature is also “time locked.”

How the SMTP portal processes data

Now that we know something about the bitfile, we can discuss the store path in more detail.

In the store path, an object enters the system through the network. A mail message comes from the Exchange or Domino mail server via the Gateway, goes through the front firewall/NAT, and is load balanced. If there are multiple SMTP portals, the system determines which SMTP portal is available.

The SMTP portal processes the object by doing the following:

- It creates a digital signature of the object plus the date and time of receipt.
- For email, it performs address resolution.
- It locates the repository for the data. In a sense, it is asking the question, “Whose document is this?” It then locates the corresponding repository. If the data is an email message for Jan, for example, the message will be archived for Jan. The SMTP portal will find a repository to which Jan has access.
- It streams data through the cloud router to the primary and secondary Smartcell.

How the Smartcells process data

At this point in the store path, the Smartcells receive and process the data by doing the following:

- Confirm successful storing of data on both the primary and secondary Smartcells.
- Schedule the bitfile for indexing (parsing the data for keywords so that when searches are performed, the relevant data can be found).
- Transmit an acknowledgement to the sending application that the data was successfully stored, and then move on to the next item.

This store path process just described is the same for all types of archiving.

Query and retrieval path

The query and retrieval path refers to email and files being searched and retrieved, resulting in data flowing out of IAP to the user.

Query

Queries are performed using the system's HTTP portals, and are performed from the IAP Web Interface (via a web browser) or through the Outlook Integrated Archive Search.

The query is submitted through the front firewall. If there are multiple HTTP portals, the load balancer determines which HTTP portal is available.

The HTTP portal formats the query and uses information from the metaserver to determine which Smartcells will return data.

The Smartcells receive a multicast request in parallel; those with the data perform a local query.

Retrieval

After the multicast is performed, the Smartcells search for a match to the search criteria. Each Smartcell returns results from its search, along with results' digital signature.

The metaserver receives the results, places them in a time-sequential order, and passes them back to the HTTP portal.

The HTTP portal computes the digital signature and validates it against the stored digital signature. This is to make sure that nothing was tampered with before returning data to the querying application.

At that point, the system sends the data back through the firewall, and to the user's browser screen or Outlook client to show the results. For example, if a user had performed a search for "technical information" in messages archived during the past three months, the system returns email messages that contain the key words from that search query.

IAP power on/off

Below are instructions for turning the IAP on and off, and specific instructions to follow in case of a power failure.

Power off

To turn off the IAP, from the PCC enter:

```
# /opt/bin/stop
# /opt/bin/shutdown
```

Wait until the shutdown script completes before removing power from the IAP systems.

Power on

To power on the IAP:

1. Make sure the ProCurve switch(es) inside the IAP are powered on. Once power is restored, the ProCurve switch(es) should automatically come up.
2. Power on the kickstart server. Wait five minutes.
3. Power on everything else. Order is insignificant, unless there has been a power failure (see below).
4. Wait for all machines to start, then log in to the PCC console and issue the command: `/opt/bin/restart`

How to restart IAP after a power failure

After a power failure has occurred, a specific power on sequence is required:

1. Power off all systems.
2. Power on the kickstart server.
3. When the kickstart machine is running, log on and issue the commands:

   ```
   /etc/init.d/postgresql stop
   /etc/init.d/postgresql start
   ```

4. Wait for the start to complete successfully.
5. Power on DB2, routers, firewall, and load balancers.
6. Power on Smartcells.
7. Power on metaservers.
8. Power on remaining servers.
9. Wait for all machines to start, then log in to the PCC console and issue the command: `/opt/bin/restart`
10. Once IAP has restarted, verify (with the PCC web interface) that the IAP is running and monitoring is reporting system availabilities as expected.

NOTE:

In the event both routers go down, the system should be restarted with `/opt/bin/restart` once the routers are back up.

Maximum file size

The maximum file size that the IAP API allows for all file types is 1.47 GB.

VPN access

Your HP authorized support representative can discuss your VPN access options with you and help select and implement the option that will work best for you.

2 Introduction to Platform Control Center (PCC)

This chapter introduces the Platform Control Center (PCC) administration tool for monitoring and troubleshooting the IAP and user accounts.

It includes the following topics:

NOTE:

PCC pages must be fully loaded (as indicated by the progress bar or icon in the browser) before using the functions on the page.

Accessing PCC

To access the PCC, open a web browser, enter the PCC server's IP address, then log in using the administrative user name and password.

Administrator privileges are set up in the Account Manager. See "User account information" on page 70 for more information.

You can also log in as the super user, if directed to do so by HP technical support. The IAP super user login name and password are set up during system installation.

NOTE:

When opening PCC using Mozilla Firefox, accept the exception request if one appears.

User interface components

The PCC is an HTML-based application containing a menu on the left side of the page (referred to as the left menu). Use the left menu to access most pages in the PCC.

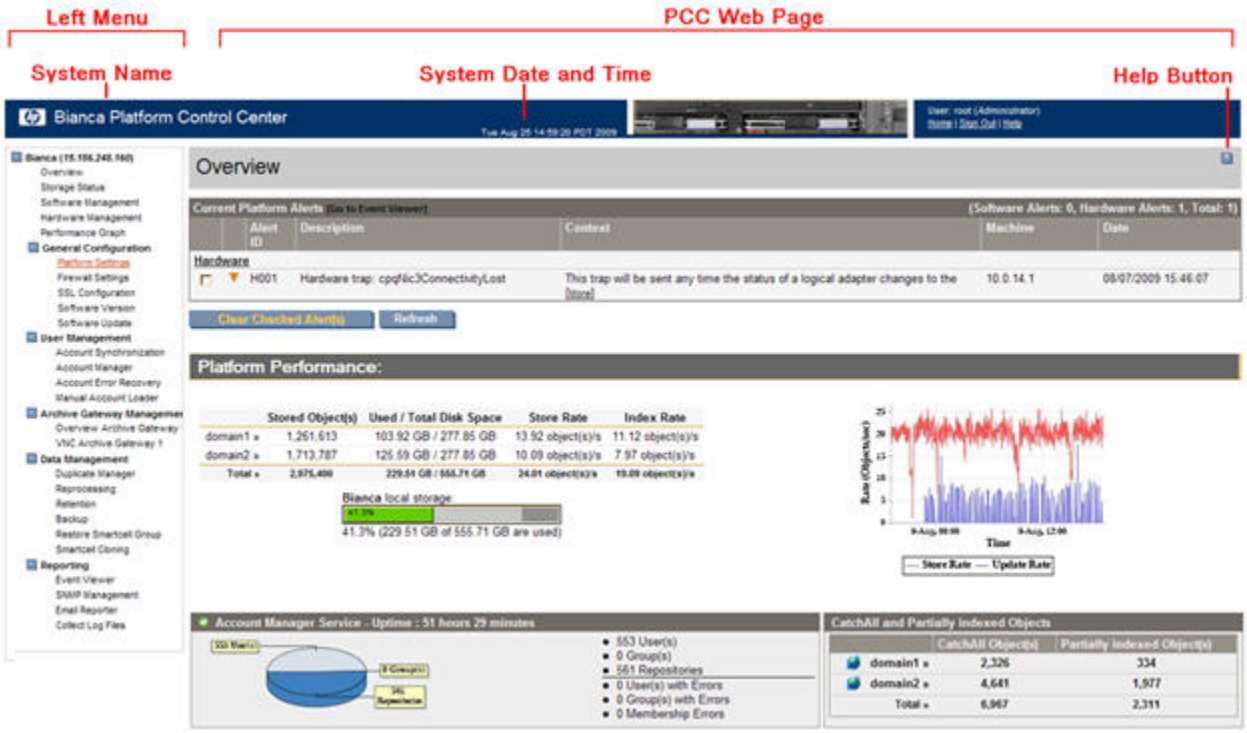

Figure 2 PCC user interface

User interface orientation tips

To orient yourself, pay attention to the different ways a PCC page is characterized.

- Link text: A navigation link leading to a page is a general description of the page. Most links to a page are from the left menu.
- HTML name: Each PCC page has a descriptive HTML name, which is displayed in the browser.

NOTE:

Be sure to wait until a PCC page is fully loaded before using any of the page functionality.

Pages for common tasks

Table 6 Pages for common system administration tasks

Task	Page
Check overall system health and performance	"Overview" on page 39
Start, stop, and restart system servers	"Software Management" on page 45
Monitor system hardware status or remotely manage servers	"Hardware Management" on page 46
Check the platform configuration	"Platform Settings" on page 53

Task	Page
Display firewalled ports enabled in the system	"Firewall Settings" on page 54
View IAP software version code installed	"Software Version" on page 57
Synchronize user accounts	"Account Synchronization" on page 59
Manage user accounts	"Account Manager (AM)" on page 67
Monitor, start, and stop replication for domains	"Replication" on page 86
Clone Smartcells (copy data)	"Smartcell Cloning" on page 107
Check database backup history	"Backup" on page 100
Monitor system events	"Event Viewer" on page 113
Activate SNMP traps and send email notifications	"SNMP Management" on page 115
Configure periodic email reports of system status and performance	"Email Reporter" on page 117
Link to Archive Gateway	"Archive Gateway Management" on page 121

Updating pages before printing

PCC pages displayed in the web interface are not automatically updated. To manually update the page, click **Refresh** (or **Reload**) in your browser.

If the browser caches web pages, the cached page displayed when you click the browser's **Back** button can be out of date. Refresh it manually.

Some browsers print from an updated version of the web page without refreshing the browser display. If the displayed page is out of date, the printout can be different from the displayed page. To ensure you print what is displayed, refresh the browser manually before printing.

Left menu

The left menu provides quick access to PCC views. The left menu varies depending on the way the system is configured. For example, systems not using replication do not have the Replication menu item available. Also, HP support personnel utilizes a Service Tools menu item which they may make available to you as needed.

Table 7 Pages accessible from left menu

Left menu item	Description
"Overview" on page 39	View summary of system health, storage status, Smartcell performance by domain, and system alerts.
"Storage Status" on page 44	View summary, by domain, of document storage rates and used/free disk space.
"Software Management" on page 45	View server software status, and start, stop, or restart one or more servers on the system.
"Hardware Management" on page 46	Monitor server hardware status and remotely manage servers.
"Performance Graph" on page 50	Graph system storage and indexing rates and system performance.
General Configuration pages	
"Platform Settings" on page 53	Display hardware and configuration information about the IAP.
"Firewall Settings" on page 54	Display each firewalled port that is enabled in the system.
"SSL Configuration" on page 54	Generate third party public certificate requests for the PCC and HTTP portals.
"Software Version" on page 57	View the IAP software version code used by the system.
"Software Update" on page 58	Install patches to the IAP software.
User Management pages	
"Account Synchronization" on page 59	Configure dynamic account synchronization (DAS) to automatically create and update IAP users with information obtained from LDAP servers.
"Account Manager (AM)" on page 67	Provision, update, and manage IAP user accounts.
"Account Error Recovery" on page 81	Repair account synchronization errors.
Archive Gateway Management pages	
"Overview Archive Gateway" on page 122	View status information about the Archive Gateway for each domain.
"VNC Archive Gateway" on page 121	Access the Archive Gateway using VNC.
Data Management pages	
"Duplicate Manager" on page 83	View status of duplicate merge jobs and schedule duplicate merge jobs.
"Reprocessing" on page 90	Schedule and enable reprocessing based on new routing rules.
"Retention" on page 92	Configure retention periods for domains and repositories.
"Backup" on page 100	View status information about database, configuration, and Smartcell backup. Access HP Data Protector user interface for backup system detailed information and configuration.

Left menu item	Description
"Replication" on page 86	Monitor and start or stop replication for IAP domains.
"Restoring a Smartcell group" on page 105	Restore a Smartcell group from tape.
"Smartcell Cloning" on page 107	Clone Smartcells (copy data) to give them a new, viable mirror cell.
Reporting pages	
"Event Viewer" on page 113	View events that have occurred in a system service or application or system hardware.
"SNMP Management" on page 115	Set SNMP traps for system monitoring and send email notifications.
"Email Reporter" on page 117	Configure system monitoring reports to be sent to email recipients.
"Collecting log files" on page 118	Collect log file reports and email or FTP them to HP technical support or other recipients.

Monitoring and reporting

PCC monitors the system and reports on its health and activity. PCC provides reports on:

- System health
- System performance
- Smartcell states

Servers in the system (and their services) are organized into groups of the same type, called host groups.

As long as services appear to be functioning correctly, the server is assumed to be healthy. If monitoring indicates a server is not functioning correctly, none of its services are available. If a service has critical status, but the host is up, the service probably needs to be restarted.

Statuses and states

Several PCC pages show current life cycle states of Smartcells or status values of particular hosts or services. Status values measure relative health, and can be associated with a status condition conveying a measure of confidence in the reported value.

For example, the health of a Smartcell in the SUSPENDED life cycle state can be reported with the HEALTHY host status value, which means the IAP operating system and applications are functioning as they should be.

PCC pages often use *status* and *state* loosely and interchangeably when referring to hosts and services. State is always used when referring to Smartcell life cycle states, but status and state are both used when reporting Smartcell health, since Smartcells are regarded as a host like any other. PCC pages also refer to status conditions as states or state types.

Smartcell life cycle states

Table 8 Smartcell life cycle states

Life cycle state	Definition	Importance
ASSIGNED 	The Smartcell is assigned to a domain. The Smartcell is available for document storage, search, and retrieval. If backup is enabled, Smartcell data can be backed up.	normal
CLOSED	The Smartcell is full. It is available for document search and retrieval, but not storage. If backup is enabled, all Smartcell data was backed up before the Smartcell entered this state.	normal
COMPLETE_PROCESSING 	Data indexing is being completed. The Smartcell is full. It is available for document search and retrieval, but not storage. If backup is enabled, Smartcell data can be backed up.	maintenance
RESTORE 	The Smartcell is a target for data restoration from another Smartcell. The Smartcell is not available for document storage, search, or retrieval.	maintenance
DISCOVERY	The metaserver and Smartcell are determining the Smartcell's start state (the state following DISCOVERY), based on expected states of the Smartcell and its mirror Smartcell. The Smartcell is not available for document storage, search, or retrieval.	maintenance (startup only)
RESET	The Smartcell is being recycled. Stored documents and corresponding management data, such as document indexes, are destroyed during recycling. The system administrator has determined that existing Smartcell data is no longer needed. The RESET state is only set manually. The Smartcell is not affiliated with any domain, so it is not available for document storage, search, or retrieval.	maintenance

Life cycle state	Definition	Importance
SUSPENDED 	Either of the following is true: • The Smartcell or its mirror Smartcell has one or more failed processes. • The mirror Smartcell is DEAD. NOTE: If the Smartcell's status is OK, only the mirror Smartcell has failed. The Smartcell is not available for storage. It is available for document search and retrieval (unless a failed process disabled the search engine). If backup is enabled, the Smartcell is backing up new data that has not yet been backed up.	failure
DEAD 	The Smartcell has failed. It is not available for document storage, search, or retrieval. If backup is enabled, some or all Smartcell data might *not* be backed up; if so, data will never be backed up.	failure
UNKNOWN 	The state of the Smartcell is unknown.	unknown
FREE	The Smartcell is free. (Shown in blue.) It can become ASSIGNED or become a target for data restoration. The Smartcell is not affiliated with any domain, so it is not available for document storage, search, or retrieval.	normal

3 System Status

This chapter discusses the information that is found in the system status pages.

It includes the following topics:

Overview

The Overview provides a high-level look at system health. It displays the following information:

- Alerts for problems that are currently occurring in the system.
- Information about document storage rates and capacity, for the system and for each domain.
- Platform information on objects stored or deleted.
- Summary of the number of IAP users, groups, and repositories in the system. (Groups are not currently supported, so the number is always 0.)
- Summary of the number of partially indexed objects and CatchAll repository objects.
- Information, by domain, about the status, health, and storage rate of each Smartcell.
- Information about the IAP software version code installed.
- The number of SMTP connections in each domain.

NOTE:

Typically, you should monitor the Overview every day.

To open the page, go to **Overview** in the left menu.

Current Platform Alerts

A Current Platform Alerts log at the top of the Overview displays alerts for problems that are currently occurring in system hardware, services, or applications.

NOTE:

Current Platform Alerts does not appear in the Overview if problems are not currently occurring in the system.

Alerts are grouped by features. For example, all hardware alerts are grouped under Hardware and all indexing alerts are grouped under Indexing. Each alert includes a description of the problem, the

context in which the problem is occurring, the server on which the problem is occurring, and the date and time the problem began. Clicking **More**, if it appears in the alert context, displays more details about the issue.

Whenever an alert is raised or cleared, a system event is logged. You can view all recent system events by clicking **Go to Event Viewer** or navigating to **Reporting** > **Event Viewer**. See "Event Viewer" on page 113 for more information.

Alert levels

There are three levels of alerts:

- *Critical*: The feature is completely down. It is critical for the entire system and requires immediate action from the IAP administrator.
- *Major*: The feature has major issues, but it is not critical for the entire system.
- *Minor*: The feature has some internal problems, but it is not critical for the feature health.

Application alerts

Application alerts are generated by the IAP software. They are cleared from the PCC Overview page automatically when the problem is fixed or has disappeared. (See Clearing alerts.) For a description of each application alert, see "IAP application-generated alerts" on page 155.

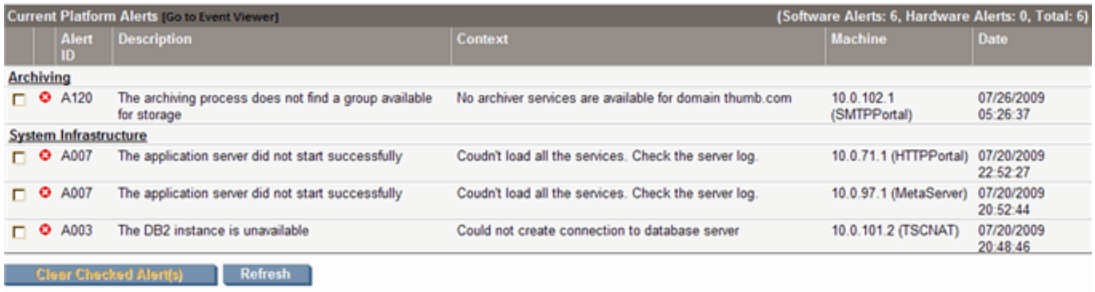

Figure 3 Sample application alerts

Application alerts can be forwarded as SNMP traps to external SNMP trap listeners or forwarded in email to designated recipients. For more information, see "SNMP Management" on page 115.

Hardware alerts

Hardware alerts are generated by the Proliant Service Packs installed on the machines in the IAP system. Generally, hardware alerts require your intervention. They are not cleared automatically, unless the server where the problem occurred is restarted. All hardware alerts have an alert ID of H001.

Each hardware alert provides a link to the server's HP System Management Homepage, where more information can be found. See "Hardware monitoring" on page 47.

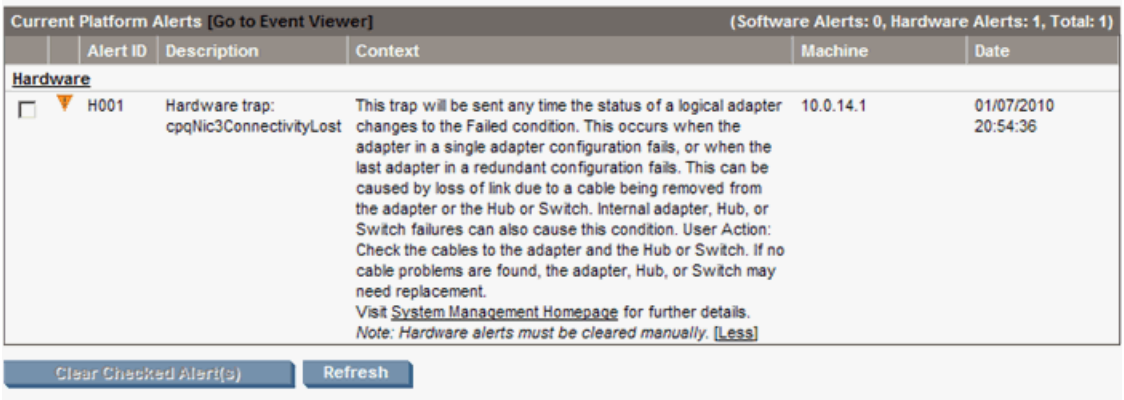

Figure 4 Sample hardware alerts

Clearing alerts

Alerts are maintained as runtime information. They are only generated once, when the problem first occurs. They can be cleared in one of three ways:

- The problem is resolved; the system observes that the problem is resolved; and the alert is cleared automatically from the PCC Overview page. (Applies to application alerts only.)
- A server is restarted; it notifies the PCC that it has restarted; and it requests the PCC to remove any alerts previously raised by it.
- The problem is resolved manually. You can resolve the problem and then manually clear the alert on the PCC Overview page by selecting the check box in front of the alert and clicking **Clear Checked Alert(s)**.

IMPORTANT:

Use caution when clearing an alert manually, as the system does not perform any verification that the problem is actually fixed. The assumption is that you know the problem does not exist anymore.

Platform Performance

This area of the overview provides information from the Storage Status page (see "Storage Status" on page 44). The table on the left displays the number of objects (documents) stored, the amount of used and total disk space, the document storage rate, and the document index rate for each domain. The bar graph displays the system's storage space ratio. The line graph on the right shows the messages per second that the system stored and updated in the past day, ending with the current time, as well as the amount of folder updates.

Account Manager Service

The Account Manager Service provides a brief summary of information from the PCC Account Manager (see "Account Manager (AM)" on page 67). This area displays the number of individual IAP users and groups, pending users and groups, and the number of IAP repositories. (Groups are not currently supported, so the number is always 0.) This area also displays the number of synchronization errors, if any, pertaining to IAP user accounts. Synchronization errors can be corrected in the Error Recovery page (see "Account Error Recovery" on page 81).

CatchAll and Partially indexed objects

This area of the Overview displays the following information:

- CatchAll Objects: The number of documents in the CatchAll repository. This repository contains documents that fail to route into the correct user repositories because they cannot be associated with:
  - Email addresses of registered IAP users imported by DAS
  - A domain routing rule

  Mailing-list messages cannot be routed if the recipient's name is not included in the message as a destination.

  The CatchAll repository can be viewed in the Account Manager (see "Account Manager (AM)" on page 67 by clicking the **Repository** radio button in AM, clicking the **Other** tab, then opening the CatchAll repository.
- Partially Indexed Objects: The number of documents in the failed indexing repository. This repository contains documents that failed to be indexed or were partially indexed. (For example, the system might not have been able to index the particular file type.) The failed indexing repository can be viewed in the Account Manager (see "Account Manager (AM)" on page 67. Click the **Repository** radio button in Account Manager, click the **Other** tab, then open the failed indexing repository.

The CatchAll and failed indexing repositories are automatically created when the system is started.

NOTE:

If the number of documents shown is `-1`, the values cannot be read by the system.

Platform Statistics

The Platform Statistics area provides status, health, and storage information about the IAP Smartcells. You can click a tab to view information about Smartcells in all domains or Smartcells in a particular domain. The Platform Statistics area also shows the IP addresses of free Smartcells in the system.

Each Smartcell's life cycle state is color-coded.

- A green table row indicates a Smartcell is ASSIGNED.
- A light green table row indicates Smartcell is CLOSED.
- A yellow-orange table row indicates a Smartcell is in the COMPLETE_PROCESSING or RESTORE state.
- A light yellow table row indicates a Smartcell is in the DISCOVERY or RESET state.
- A red table row indicates a Smartcell is SUSPENDED.
- A black table row indicates a Smartcell is DEAD.
- A gray table row indicates the state of a Smartcell is UNKNOWN.

See "Smartcell life cycle states" on page 36 for more information.

The icon at the beginning of the table row displays the JBoss status of the Smartcell.

- A green check icon indicates the Smartcell is healthy, and both machine and JBoss are up.
- A yellow icon indicates the Smartcell is unhealthy. The machine and JBoss are up, but there are some failed and/or unhealthy MBeans.
- A red X icon indicates the Smartcell is down. Either the whole machine or JBoss is down.
- A gray icon indicates that service monitoring failed. The server status is unknown.

TIP:

If you move your mouse over the icon, you will see more information including the host name, active thread count, and the Smartcell's MAC address.

Platform Statistics features

Table 9 Platform Statistics features

Feature	Description
Platform	The platform name, IP address, and document storage rate.
Domain	The name of the domain.
Group Name	A Smartcell group identifier generated automatically by IAP. This number is unique across all systems.
Smartcell IP	The IP address of the Smartcell.
Smartcell Role	A Smartcell can be Primary, Secondary, Replica-1, or Replica-2.
State	The current life cycle state of the Smartcell. See "Smartcell life cycle states" on page 36.
Stored Object(s)	The number of objects stored since the Smartcell was assigned. NOTE: When the system is actively storing objects, this count might be different than the stored object count in "Platform performance" on page 41. This number is the real-time count on the Smartcell, while the Platform Performance count (taken from the local database) is only updated every minute.
Indexed Object(s)	The number of objects indexed since the Smartcell was assigned.
Store Rate	The number of objects being stored per second.
Index Rate	The number of documents being indexed per second.
Index Queue	The number of documents waiting to be indexed.
Index Deletion Queue	The number of indexed documents waiting to be deleted.
Update Queue	The number of documents waiting to be updated. (For example, email folder updates.)
Other Smart Cells	IP addresses of Smartcells in the FREE state.

IAP Version

Near the bottom of the Overview, you will find information about the IAP software version code installed.

SMTP Flow Control

At the bottom of the Overview, the SMTP Flow Control area shows the following information, by domain:

- The maximum number of connections allowed
- The current number of connections
- The number of archiver connections

Storage Status

The Storage Status page provides detailed object storage information for each domain.

To open the page, go to **Storage Status** in the left menu.

Table 10 Storage Status features

Feature	Description
Platform Store	The number of objects and store rate per domain, and the allocated space on the system for storage and replication. The dark area on the right side of the example storage bar graph below shows the point at which storage space is 90 percent full. The dark area typically represents the amount of disk space reserved for work space for the Smartcell maintenance tasks. 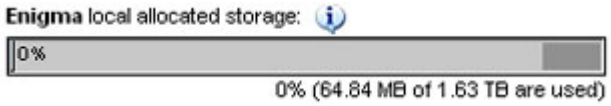 NOTE: The storage bar graph shows only assigned and allocated Smartcells for all active domains. The IAP might have free, unallocated hardware that is not represented in the bar graph.
Store Rate Graph	A line graph showing the number of messages per second that the system stored in the past day, ending with the current time. It also shows the amount of folder updates. Example: 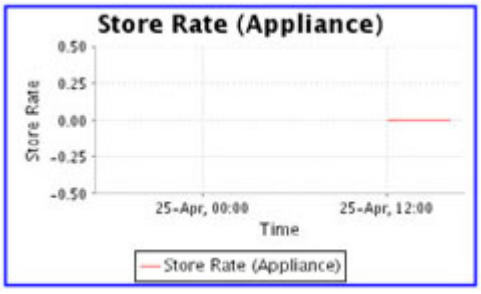

Feature	Description
Smart Cell Allocation	The number of Smartcells in each life cycle state. (See "Smartcell life cycle states" on page 36 for an explanation of each state.) Example: 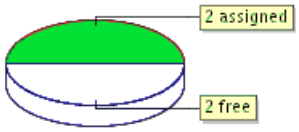
Domain Details	The domain group ID, the number of objects stored, the amount of used and total storage space, and the disk/space ratio, shown in bar graph format.

Software Management

Use the Software Management page to view the status of all IAP servers that are running JBoss, and to start, stop, or restart one or more of the IAP servers.

This page is useful to show the software status of a server. However, you should use the start/stop/restart functions only when necessary — for example, when you are upgrading a server or before a planned power outage.

To open the page, go to **Software Management** in the left menu.

Software Management features

The following information is displayed on the Software Management page.

Table 11 Software Management features

Feature	Description
Hostname	The type of host group, the servers in the group, and the full name of each server. Host group types include the following: • BACKUP Servers: Backup servers • CLOUD ROUTER Servers: Cloud router servers • DB Servers: DB2 database server • FIREWALL Servers: Firewall servers • HTTP Servers: HTTP portal servers • KICKSTART Servers: Server for loading and deploying system software/configuration files • LOAD BALANCER Servers: Load Balancer servers • PCC Servers: Platform Control Center server • META Servers: Metaservers • SMARTCELL Servers: Servers to store archived documents • SMTP Servers: SMTP portal servers The Archive Gateway does not use JBoss in IAP 2.1 and later versions, so is not shown on the Software Management page.

Feature	Description
Status icon	The status icon in front of the host name displays the JBoss status of the server. • A green check icon indicates the server is healthy, and both machine and JBoss are up. • A yellow icon indicates the server is unhealthy. The machine and JBoss are up, but there are some failed and/or unhealthy MBeans. • A red X icon indicates the server is down. Either the whole machine or JBoss is down. • A gray icon indicates that service monitoring failed. The server status is unknown.
IP Address	The server's IP address.
Status	The status of the server. Corresponds to the icon in front of the server in the Hostname column.

Starting, stopping, and restarting servers on the system

To stop, start, or restart one or more servers, follow these instructions:

1. Select the machine(s) on which the action will be performed:
   - To perform an action on all servers or all machines in a server group, click the **Machine Type** radio button and select the servers from the drop-down list.
   - To perform an action on specified machine(s), click the **Selected Machine(s) Below** radio button, and then select the check box in front of the machine(s).

NOTE:

Because the PCC server hosts the user interface, it is not listed in the Machine Type list and cannot be selected if you specify Selected Machine(s) below.

Stop or restart the PCC using the command line interface.

2. In the Action drop-down list, select the action to perform:
   - **Start**: Start a single machine, start all machines, or start all machines in a selected server group.
   - **Stop**: Stop a single machine, stop all machines, or stop all machines in a selected server group.
   - **Restart**: Stop and immediately start a single machine, or stop and immediately start all machines or all machines in a selected server group.
   - **Staggered Restart**: Restart all machines in sequence, or all server group machines in sequence, with a minimum of downtime.

     Staggered Restart can only be used with Smartcell, HTTP, and metaservers.
3. Click **Run Now**.

Hardware Management

Use the Hardware Management page to view the servers that are part of the IAP infrastructure and links to their corresponding HP System Management pages and Remote Management iLO pages (if

enabled). (See the IAP Administration Guide chapter “Remote management of servers” for more information about remote management.)

To open the page, go to **Hardware Management** in the left menu.

Hardware Management features

Table 12 Hardware Management features

Feature	Description
Hostname	The type of host group, the servers in the group, and the full name of each server. Move your mouse over the status icon to view a server's IP address. Host group types include the following: • ARCHIVEGATEWAY Servers: Archive Gateway (email mining) servers • BACKUP Servers: Backup servers • CLOUD ROUTER Servers: Cloud router servers • DB Servers: DB2 database server • FIREWALL Servers: Firewall servers • HTTP Servers: HTTP portal servers • KICKSTART Servers: Server for loading and deploying system software/configuration files • LOAD BALANCER Servers: Load Balancer servers • PCC Servers: Platform Control Center servers • META Servers: Metaservers • SMARTCELL Servers: Smartcell servers to store archived data • SMTP Servers: SMTP portal servers
Status icon	The icon in front of the host name displays the hardware status of each server. • A green check icon indicates there are no hardware alerts for the server. • A red X icon indicates there is at least one hardware alert for the server. Move your mouse over the status icon to display the number of alerts. To view the alerts, navigate to Current Platform Alerts on the PCC Overview page.
MAC Address	The server's MAC address.
System Management	Click **View** to link to the HP System Management Homepage on the server, which displays the server hardware status and other data. For more information, see Hardware monitoring below.
Remote Management	Shows whether iLO/Lights-Out 100 is set up on the server and configured in Proxy Access Mode. If remote management is enabled, a **Manage** link appears. Clicking the link displays the iLO or Lights-Out 100 Web interface and allows you to perform operations such as power cycling the server.

Hardware monitoring

You can monitor IAP hardware health from either the PCC user interface, the HP System Management Homepage on every server, or (if available) an HP System Management console. Hardware monitoring for IAP servers is provided by Proliant management agents. We include the standard HP Proliant

hardware management and system homepage agents as part of the IAP software. The management agents provide SNMP support for hardware monitoring by way of SNMP traps.

A standalone IAP Proxy Agent running on PCC aggregates SNMP hardware traps from the Proliant agents on the IAP servers. The agent is used for SIM discovery and trap forwarding

SIM discovery: The IAP Proxy agent responds to SNMP get requests from external SIM servers on the PCC external network IP address. An external SIM server needs to be configured to identify IAP appliances. The IAP appliance is displayed in the SIM console with the IP address as the PCC external network IP address.

Trap forwarding: The proxy agent listens for SNMP traps from the individual IAP servers.

- **SIM:** Once discovered in SIM, the proxy agent in the PCC server must be configured with the SIM server's IP address for trap forwarding. Any number of external SIM servers are supported. The SIM server displays traps received from the IAP Proxy agent as long as the source address is the IAP appliance address: the PCC's external IP address. The proxy agent modifies the traps received from the IAP servers by setting the source address to the PCC external address before forwarding the traps to SIM servers.
- **PCC:** The IAP proxy agent loads known MIB definitions used by the Proliant agents in order to parse the SNMP traps and retrieve source and description information. The IAP Proxy agent forwards the traps in the form of IAP alerts to the PCC, where they are displayed in the Current Platform Alerts log at the top of the Overview page. A link to the System Management Homepage for the IAP server is displayed in the alert context.

The Proliant management agents monitor hardware components such as storage (raid controllers and disks), server (CPU, memory, temperature, and power), and NICS. The full list of hardware traps is shown in the documentation for the Proliant Support Packs: http://bizsupport2.austin.hp.com/bc/docs/support/SupportManual/c00760188/c00760188.pdf.

Other information is included in the Proliant support pack user and installation guide: http://bizsupport2.austin.hp.com/bc/docs/support/SupportManual/c01527454/c01527454.pdf

Since hardware monitoring of IAP servers is dependent on the Proliant management agents, the hardware requirements are tied to the minimum requirements for the Proliant Support Packs. See the IAP Support Matrix for more information.

Log messages from the IAP Proxy agent are stored in the `proxyagent.log` files in the `/opt/tools/snmpProxyAgent/` directory on the PCC server.

SIM Console

If you use the HP SIM Central Management Server (CMS), it can recognize the IAP appliance via SNMP from the SIM console, using the PCC external IP address. Once configured to identify IAP appliances, the SIM console displays the IAP as a single managed system as shown below.

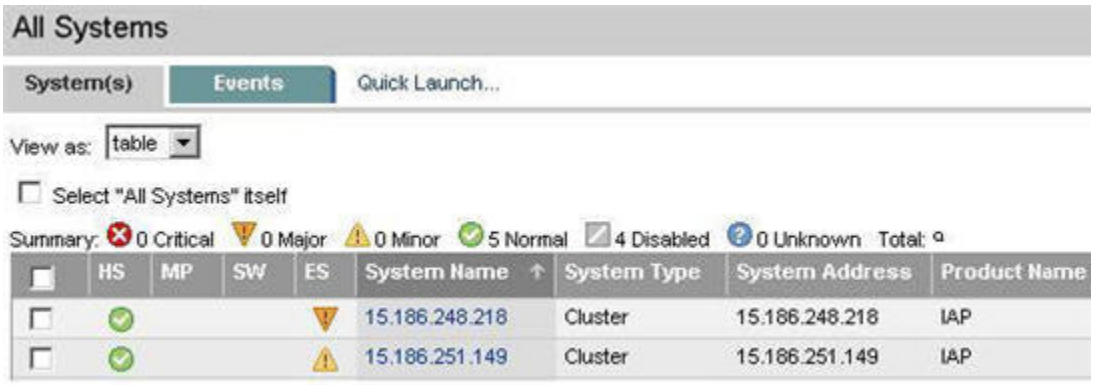

Figure 5 SIM Console

The ES column indicates the event status of the IAP, which is the most severe event or trap that has been received. The trap detail contains information about the actual IAP server that originated the trap.

IMPORTANT:

The addresses of SIM servers must be manually configured in a text file to enable trap forwarding from the proxy agent. For trap forwarding, create or use the file `/opt/tools/snmpProxyAgent/conf/ProxyAgent.properties` and add the following line: `trapsink=<sim_server_ip>` (for example 192.168.50.59).

Clicking the IAP System Name brings up the System Homepage as shown below.

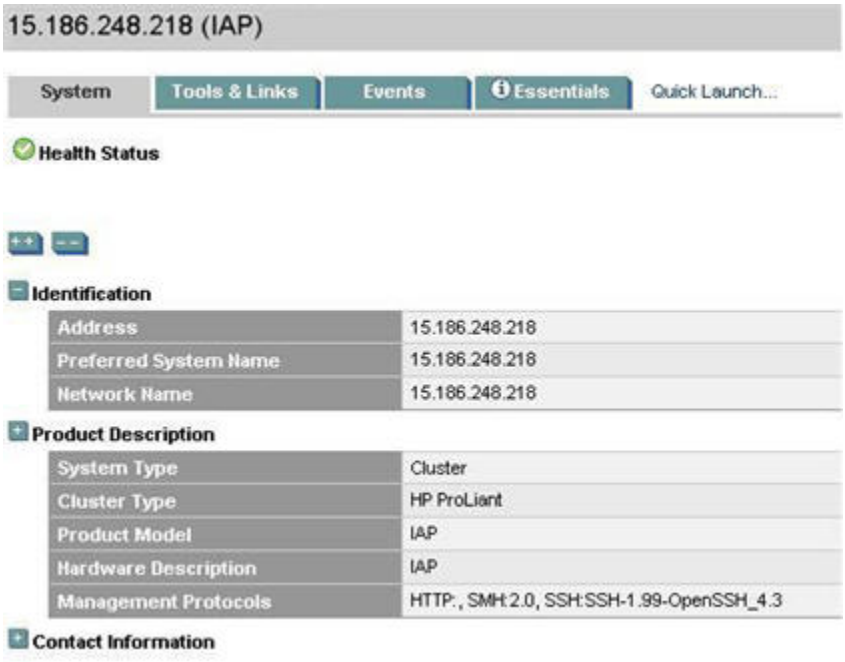

Figure 6 System Homepage

The events received from the IAP can be seen by clicking the Events tab. This brings up the events table. Trap detail can be obtained by clicking the Event Type. The detail includes the IAP server that was the source of the trap.

Figure 7 IAP events

The SIM documentation describes the use, configuration, and user interface of the HP SIM server apart from the specific use to monitor the IAP:

- HP SIM installation and configuration guide for Windows installations: http://docs.hp.com/en/418812-005/418812-005.pdf
- HP SIM installation and configuration guide for Linux installations: http://docs.hp.com/en/418811-004/418811-004.pdf

Performance Graph

Use this page to create graphs showing different types of system events over specified time periods. You can generate two categories of graphs:

- *System monitoring graphs* that show idle CPU usage, free memory usage, or the number of threads used.
- *Platform store and indexing graphs* that show the number of objects stored or indexed on the system, and the rate at which objects are stored or indexed per second.

Table 13 Performance Graph features

Feature	Description
Graph Type	The type of graph or the server name.
Assigned Smartcell or Domain	The platform, domain, or Smartcell to be graphed.
Time Frame	The time period being reported.

To open the page, go to **Performance Graph** in the left menu.

Example: Platform Store graph

An example of a platform store or indexing performance graph is shown below. This graph charts the Domain1 store rate for today at five minute intervals.

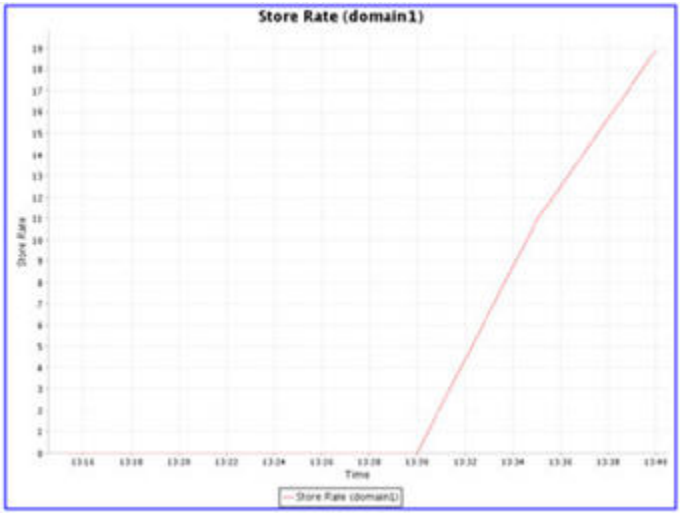

Figure 8 Performance Graph: Store Rate

Example: System Monitoring graph

An example of a system monitoring performance graph is shown below. This graph charts the free memory on the database server at hourly intervals over the past 24 hours.

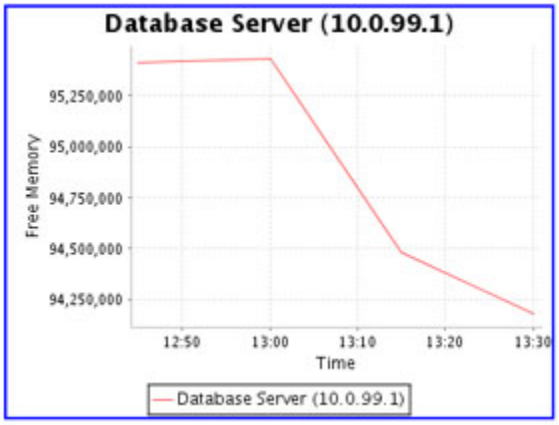

Figure 9 Performance Graph: Free Memory

Creating performance graphs

1. Click the **System Monitoring** tab or the **Platform Store and Indexing** tab for the category of graph that you want to create.

2. Select a graph type:
   - For Platform Store and Indexing, select one of the following:
     - Store Rate (Objects/Sec)
     - Store Rate (MB/sec)
     - Index Rate
     - Object Count
     - Index Count
     - Active Transactions
     - Update Rate (Objects/Sec)
   - For System Monitoring, select one of the following:
     - Idle CPU

       Percentage of the CPU that is idle.
     - Free Memory

       Percentage of the amount of free memory in the Java virtual machine vs. the total amount of memory in the Java virtual machine (same as JVM Heap Free / JVM Heap Total)
     - JVM Heap Total

       The total amount of memory in the Java virtual machine.
     - JVM Heap Used

       The total amount of memory in the Java virtual machine minus the amount of free memory in the Java virtual machine (same as JVM Heap Total – JVM Heap Free).
     - JVM Virtual Memory Used

       The amount of virtual memory the JVM processes were using, in bytes. (Includes the amount in RAM and in the swap memory.)
     - IO

       The percentage of disk resource (I/O) utilization for all main disks on the system.
     - Active Thread Count

       Number of Java threads that JBoss was running at the time.
3. Select one of the following options:
   - Machine Type (System Monitoring graphs): Select the type of machine (for example, PCC Servers or SMTP).
   - Assigned Smartcell/Domain (Platform Store and Indexing graphs): Select **Entire Platform**, the domain name, or the Smartcell IP address.
4. Select the time frame and reporting interval:
   - Time Frame: For a preselected time period, click **Select Time Frame**, and select a time frame from the drop-down list.
   - From Date, To Date: For a custom time period, click **Custom Time Range**, and select the start date and time and end date and time. The custom time range can be useful for troubleshooting.
   - Interval: Select the reporting interval from the drop-down list.
5. Click **Generate Graph**.

4 Configuration

This chapter contains the following information:

Platform Settings

The Platform Settings page is an administrative tool that displays hardware and configuration information about the IAP.

This page is divided into two parts:

- Information about the services enabled in each domain is in the upper portion of the page.
- Information about the IAP configuration is in the lower portion of the page.

To open the page, go to **General Configuration** > **Platform Settings** in the left menu.

Domain Configuration

The upper portion of the Platform Settings view shows information about each domain. It includes the domain type (Exchange or Domino), the services that are enabled (for example, indexing, replication, folder support, compliance, and data backup), supported object (document) types, and retention period(s).

The domain configuration information is drawn from the `Domain.jcml` file on the kickstart server.

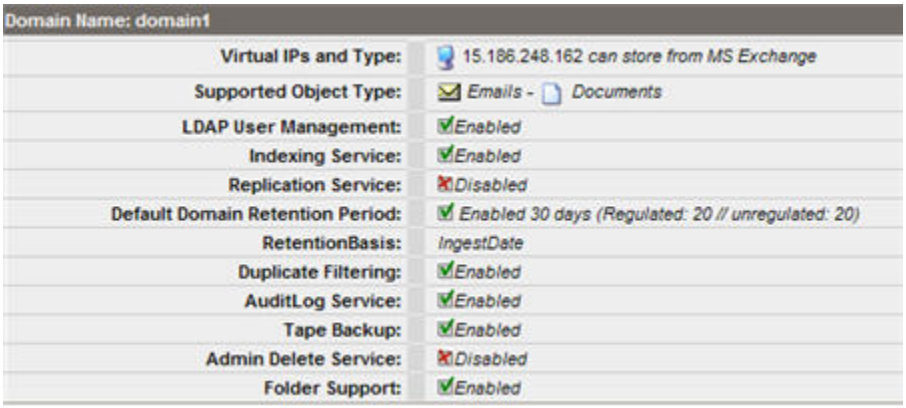

Figure 10 Domain Configuration

Platform Settings

The lower portion of the Platform Settings page displays the setup details for the IAP. This information is taken from the `BlackBoxConfig.bct` file on the kickstart server.

Firewall Settings

The Firewall Settings page shows the firewall status and settings for the PCC server and the IAP HTTP portals, and their virtual IP (VIP) addresses.

It includes the following information:

Table 14 Firewall ports

Feature	Description
Virtual IP	The virtual IP address.
Port	The port number.
Service	The service running on the port.
Type	The transfer protocol used on the port: TCP or UDP
Inbound/Outbound	Shows if the firewall is enabled or disabled on the port's inbound and outbound traffic. Outbound traffic is not filtered on the PCC server ports.

To open the page, go to **General Configuration** > **Firewall Settings** in the left menu.

SSL Configuration

SSL, or Secure Socket Layer, is a technology that allows web browsers and web servers to communicate over a secured connection. This means that the data being sent is encrypted by one side, transmitted, then decrypted by the other side before processing. It is a two-way process, meaning that both the server AND the browser encrypt all traffic before sending out data.

The SSL Configuration page lets you generate certificate signing requests (CSRs), and corresponding private keys, for secured connections on the IAP. You can create two types of CSRs: one for access to the PCC portal, and one for access to the HTTP portals.

After generating the CSRs, see "Installing a third party certificate on the PCC portal" on page 56 and "Installing a third party certificate on the HTTP portals" on page 56 for the steps needed to complete the process.

The SSL Configuration page contains two areas. The top area shows any current certificate signing requests in the system. The bottom area contains a form to complete to generate a certificate signing request (CSR) and RSA private key.

To open the page, go to **General Configuration** > **SSL Configuration** in the left menu.

Available certificate signing requests

Table 15 Available certificate signing requests (CSRs) in the IAP

Feature	Description
Machine Type	The host for which the CSR has been generated. The host is either a PCC or HTTP portal.
Virtual IP	The virtual IP address of the host.
Creation Date	The date the CSR was created.
Path	The path to the CSR. The CSR files are always placed on the PCC host.

Creating a certificate signing request

To create a certificate signing request (CSR):

1. Complete the form at the bottom of the SSL Configuration page.

   You can create only two types of CSR files: one for access to the PCC, and one for access to the HTTP portal machines.

2. Click **Generate CSR**.

   To install a certificate on the PCC portal, see
   "Installing a third party certificate on the PCC portal" on page 56.

   To install a certificate on each HTTP portal, see
   "Installing a third party certificate on the HTTP portals" on page 56.

NOTE:

If you enter incorrect information in the form and need to generate a new CSR, see "Deleting a certificate signing request" on page 55 for the procedure to follow.

Deleting a certificate signing request

After a certificate signing request (CSR) is generated for the PCC or HTTP portals, it cannot be regenerated. If you have entered incorrect information in the certificate request, the file must be manually deleted before you can create a new CSR in the SSL Configuration page.

To delete a CSR:

1. Log in to the PCC console.
2. Go to `/opt/keys` and delete the CSR with one of the following commands:

   `rm -f pccCert.pem` (to remove the PCC certificate request)

   `rm -f httpCert.pem` (to remove the HTTP certificate request)

3. Upon deleting `pccCert.pem` and/or `httpCert.pem` in directory `/opt/keys`, log off or close the PCC UI. If you do not, refreshing the PCC UI will re-create these files, and the <SSL Configuration> page will not allow new CSRs be created.

IMPORTANT:

Important! Do not delete the private key files (`pcckey.pem` or `httpkey.pem`).

NOTE:

After deleting `pccCert.pem` or `httpCert.pem` in `/opt/keys`, be sure to log off or close the PCC UI. If you don't and refresh, the PCC UI will re-create these files. (The SSL Configuration page will also not allow new CSRs be created.)

Installing and generating a certificate on the PCC portal

Follow these steps to generate and install a certificate for the IAP PCC portal.

1. Create a certificate signing request (CSR) for the PCC:
   a. Log in to the PCC Web interface and go **General Configuration** > **SSL Configuration**.
   b. Complete the CSR generation form.
   c. Log out of the PCC Web interface.

   This generates two files on the PCC:
   - `/opt/keys/pccCert.pem` (the certificate request)
   - `/opt/keys/pcckey.pem` (the RSA private key)
2. Manually copy the certificate request file to your local machine:

   ```
   scp root@[external ip address of PCC]:/opt/keys/pccCert.pem
   ```
3. Send the certificate request to a certificate authority (CA) such as VeriSign for signing.

   Follow the instructions provided by your CA.
4. Import the certificate you receive from the CA into the IAP PCC:
   a. Store the certificate from the CA on your local machine (for example, as `pccCertSigned.pem`).
   b. Copy the certificate to the PCC:

   ```
   scp pccCertSigned.pem root@[external ip address of PCC]:/opt/keys/
   pccCertSigned.pem
   ```
5. Import the certificate into the PCC's Apache server:

   ```
   /usr/local/bin/ssl_cert_update.pl --pcc --cert /opt/keys/
   pccCertSelfSigned.pem --key /opt/keys/pcckey.pem
   ```
6. Restart the PCC's Apache server by issuing the following command:

   ```
   /etc/init.d/httpd restart
   ```

Installing and generating a certificate on the HTTP portals

Follow these steps to install a certificate on the IAP HTTP portals.

1. Create a certificate signing request (CSR) for the HTTP portals:
   a. Log in to the PCC Web interface and go **General Configuration** > **SSL Configuration**.
   b. Complete the CSR generation form.
   c. Log out of the PCC Web interface.

   This generates two files on the PCC:
   - `/opt/keys/httpCert.pm` (the certificate request)
   - `/opt/keys/httpkey.pem` (the RSA private key)
2. Manually copy the certificate request file to your local machine:

   ```
   scp root@[external ip address of PCC]:/opt/keys/httpCert.pm
   ```
3. Send the certificate request to a certificate authority (CA) such as VeriSign for signing.

   Follow the instructions provided by your CA.
4. Import the certificate you receive from the CA into the IAP PCC:
   a. Store the certificate from the CA on your local machine (for example, as `httpCertSigned.pem`).
   b. Copy the certificate to the PCC:

      ```
      scp httpCertSigned.pem root@[external ip address of PCC]:/opt/keys/
      httpCertSigned.pem
      ```
5. From the PCC console:
   a. Import the certificate into the Apache server on each HTTP portal.

      ```
      /usr/local/bin/ssl_cert_update.pl --http --cert /opt/keys/
      httpCertSelfSigned.pem --key /opt/keys/httpkey.pem
      ```
   b. Restart all services on the HTTP portal by issuing the following command:

      ```
      /opt/bin/restarthttp
      ```

      (You can also restart the services using Platform Control in the PCC Web interface. See "Software Management" on page 45.)

Software Version

This page shows the software version that is used by hosts in the system, and any software updates that have been installed.

Table 16 Software Version features

Feature	Description
Software Version	IAP software version code installed.
Patches Applied	History of IAP software patches (updates) applied to the system.

To open the page, go to **General Configuration** > **Software Version** in the left menu.

Software Update

The Software Update link will connect you to the IAP tool which allows you to install patches. To open the page, go to **General Configuration > Software Update** in the left menu.

5 Account Synchronization

Use this page to configure dynamic account synchronization (DAS), which automatically creates and updates email user accounts on the IAP. You can define multiple configurations that track sets of users from one or more LDAP servers for specific IAP domains.

NOTE:

This chapter contains a DAS example for EAs for Exchange. EAs for Domino administrators will find information on DAS in the *EAs for Domino Administrator Guide*.

To open the page, go to **User Management** > **Account Synchronization** in the left menu.

Account Synchronization overview

The Account Synchronization page is divided into three areas.

- The DAS Available Jobs area lists all jobs created and assigned to an HTTP portal.
- The LDAP Server Connectors area lists available LDAP connections.
- The Jobs History Logs area shows the history of DAS job runs.

Creating and running DAS jobs

The basic steps for creating and running a DAS job are as follows:

1. Create an LDAP connection. See "Creating LDAP server connections" on page 59.
2. Create the job. When you create a new job, you assign the job a name and an LDAP connection, and set up the job query in the LDAP server. See "Creating jobs" on page 60.
3. Assign the job to an HTTP portal. See "Assigning HTTP portals" on page 63.
4. Run the job. See "Running DAS jobs" on page 63.

Creating LDAP server connections

To create an LDAP connection:

1. In the LDAP Server Connectors area, click **New LDAP**.

2. Complete the form to create an LDAP service connection by entering the following information:

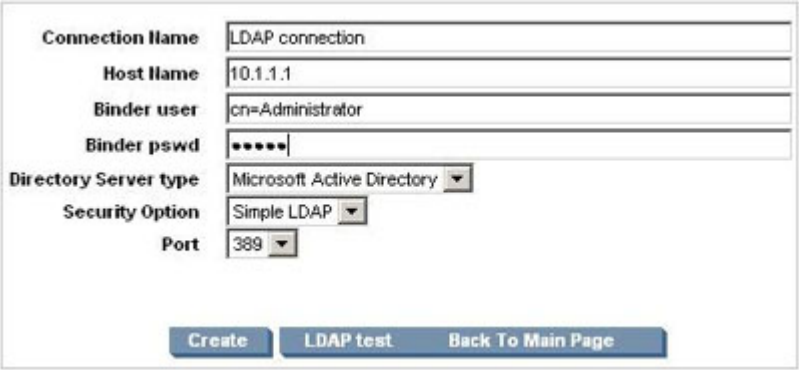

Figure 11 New LDAP connection

- Connection Name: Name used to identify the LDAP connection.
- Hostname: IP address of the LDAP server.
- Binder user: User in the LDAP directory tree that you want to bind to. At a minimum, the user must have read access to all users objects. For example: `cn=Administrator,cn=Users,dc=hostname,dc=com`.
- Binder pswd: Password of the Binder user.
- Directory Server type: Type of LDAP server to which you are connecting: Microsoft Active Directory
- Security Option: Type of LDAP security: simple authentication or SSL.
- Port: Open LDAP port of the LDAP server. Use the port 389 for simple authentication. Use port 636 for SSL support.

3. To test the LDAP server connection before creating it, click **LDAP test**.

   A content pane displays the status of the LDAP connection. It tells you whether the connection and bind are successful and the authentication types that are supported by the LDAP server. Errors are displayed in red.

4. Click **Create**.
5. Return to the Account Synchronization page and verify that the new LDAP server connection is listed under LDAP Server Connectors.

Creating jobs

To create a DAS job:

1. In the DAS Available Jobs area, click **New JOB**.
2. Name the job you are creating by entering it in the Job Name box. (In addition to the characters noted below, the name can contain no spaces.) Click **Next Step.**

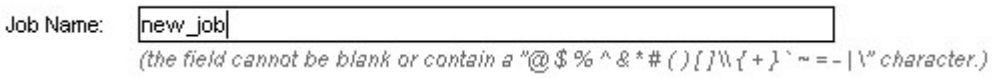

Figure 12 Create DAS job

3. From the drop-down list, select the LDAP connection you want to use with the job.

   If you need to create the LDAP connection, click **Create New LDAP Connection** and see "Creating LDAP server connections" on page 59 for further instructions.

4. Click **Next Step**.

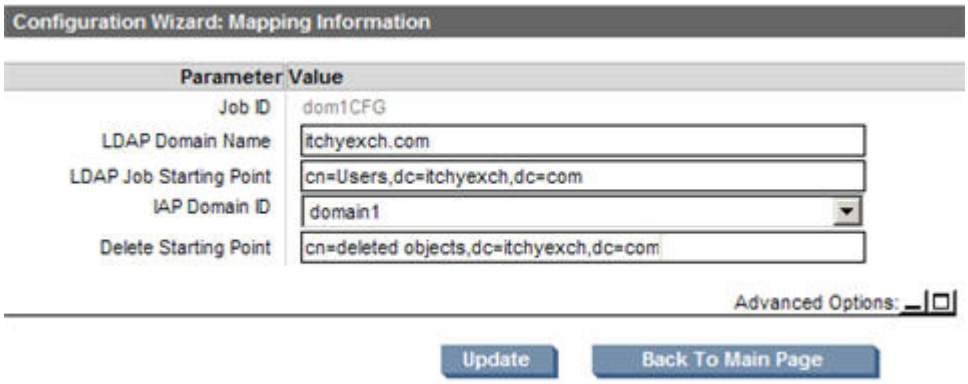

Figure 13 Mapping information

5. Complete the form by entering the following:
   - LDAP Domain name: Domain to which the users belong. For example: `ldaptest.com`.
   - LDAP Job Starting Point: Root node where the user accounts are stored. Example:
     For Exchange, enter `cn=Users,dc=ldaptest,dc=com` for node `Users` in domain `ldaptest.com`. The value must specify the relative location in the LDAP tree, including parent nodes and domain name. (Another example is `ou=stest,dc=sandbox2k_12,dc=com` where ou is the organization name in the exchange.)
   - IAP DomainID: The IAP domain ID (not the domain name) where users are synchronized with users on the LDAP server. This is the same as the domainID set in `Domain.jcml`.
   - Delete Starting Point: The root node where deleted user objects are stored on the LDAP server.
     Example: For Exchange, enter `cn=deletedObjects,dc=ldaptest,dc=com` for node `deletedObjects` in domain `ldaptest.com`. The value must specify the relative location in the LDAP tree, including parent nodes and domain name.
6. The console displays "`Do you want to attach the job to an HTTP portal`" and a "Back to main page" button.

   To complete the advanced options, see "Mapping advanced options" on page 61.

Mapping advanced options

Click the Advanced Options icon (⊡) to display the advanced options.

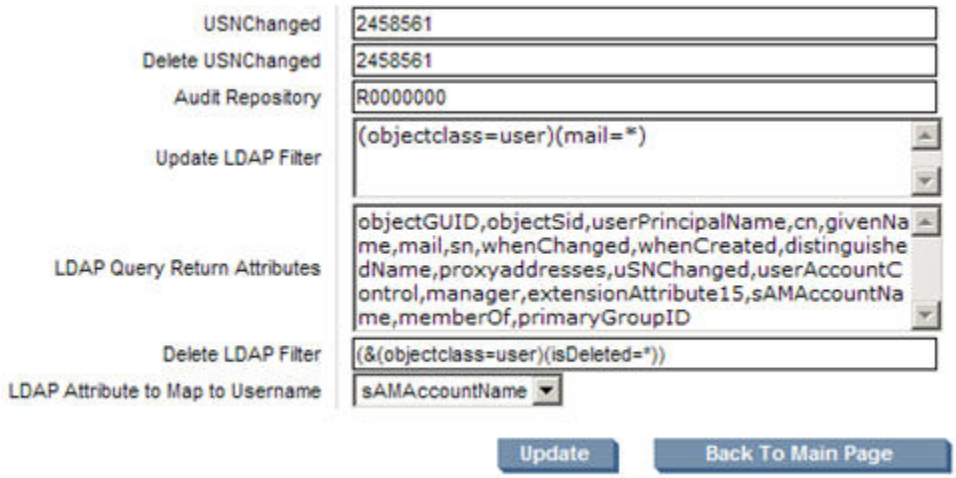

Figure 14 Advanced options

1. Complete the advanced options form by entering the following information:
   - USNChanged: Active Directory's unique sequence number (USN). Active Directory increments the USN for each change in any of its user accounts. When DAS finds a larger USN, it extracts new information. For initial setup, set USNChanged to **1** so DAS extracts all users. Thereafter, do not change this value.
   - Delete USNChanged: USN in deleted users directory. This is the same as USNChanged but for deleted objects. For initial setup, set this value to **0**. Thereafter, do not change this value. This value must be set to 0 the first time the mapping is set up.
   - NextRepID: IAP repository ID assigned to new users. DAS retrieves this value from the database, but you can use this feature to specify the repository ID for the first user inserted when DAS runs. For example, if you enter **67**, the repository created and assigned for the first imported user is R0000067. You can use this feature if users already exist in the system or to reserve repository IDs for any reason. If you specify a value that is lower than an existing repository ID, DAS automatically changes the value to the next higher number. The default is set to 5.
   - Audit Repository: Do not change.
   - Update LDAP filter: Criteria to include or exclude specific users. For example:

     For Exchange, use at least the default, `(objectclass=user)(mail=*)`, which excludes users who do not have email accounts.
   - LDAP Query return attributes: List of return attributes. Example:

     For Exchange, use the default attributes unless your LDAP schema requires mapping changes.
   - Delete LDAP Filter: Criteria to include or exclude specific users in the LDAP deleted users directory. Add to the default for special cases.
   - LDAP Attribute to Map to Username: Active Directory attribute that will be mapped to users' username. The attribute that is selected will be the webui username login name. You may use Active Directory fields sAMAccountName, userPrincipalName or mail as the IAP username mapping.

NOTE:

It is no longer necessary to modify the `LoadChanges.dse` file on the HTTP portal to perform the username mapping.

2. Click **Next Step** to create the job.

3. To assign this job to an HTTP portal, click **Assign Job**.

   See Assigning HTTP portals for further instructions.

Assigning HTTP portals

Before running a DAS job, assign an HTTP portal on which to run the job. Only one job can be assigned to an HTTP portal.

If all HTTP portals are being used, reassign the HTTP portal from another job. (To reassign an HTTP portal, select an existing DAS job and click the **Unassign HTTP** button.)

To assign an HTTP portal to a job:

1. Click **Assign Job**.

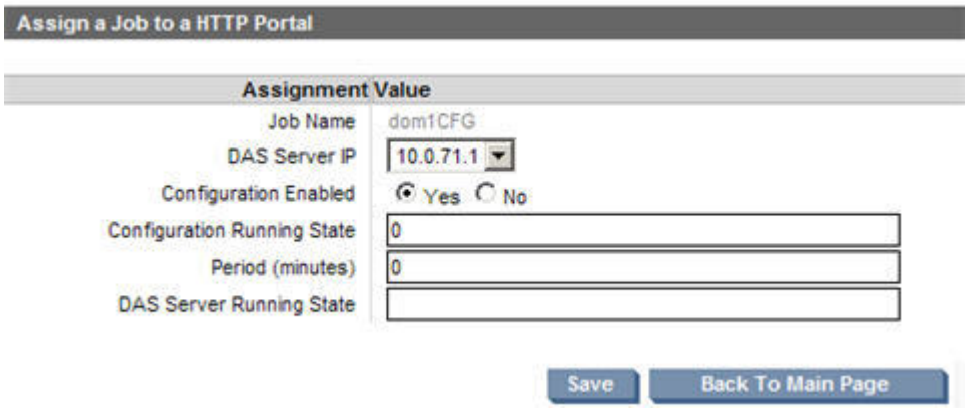

Figure 15 Assign a job to a portal

2. Complete the form by entering the following information:
   - DAS server IP: IP of the DAS HTTP portal where DAS runs the configuration.
   - Configuration Enabled: Select **Yes** to enable. If not enabled, the job cannot be scheduled or started with this console.
   - Configuration running state: Do not change.
   - Period: Number of minutes between job runs. Enter **0** to run the job once.
   - DAS server running state: Do not change.
3. Click **Save**.

Starting, scheduling, and stopping DAS jobs

1. In DAS Available Jobs, select the job.
2. To start the job without making changes to the schedule, click **Start**.
3. To make a change to the schedule before starting the job:
   a. Click **Schedule**.
   b. Enter the number of minutes between job runs. To run a job once, enter **0**. The DAS schedule should not be set to any interval less than 60 minutes.
   c. Click **Confirm Schedule** to save your changes.
   d. To launch the selected job, click **Start**.

   A notice appears stating that the job started successfully.

4. Click **Back to Main Page** to return to the DAS page.

To stop a DAS job from running:

1. In the DAS Available Jobs page, select the job and click **Start/Stop**.

   A notice appears stating that the job has stopped. When the DAS job is stopped, DAS will unschedule itself and not run any more of the DAS job. (A currently running DAS job will not stop.)
2. Click **Back to Main Page** to return to the DAS page.

NOTE:

If you are running DAS for the first time or have moved DAS to a different HTTP portal, you must run DAS with the initialize option set to true.

Editing or deleting jobs

To edit or delete an active or configured job:

1. In the DAS Available jobs area, click the name of the job you want to edit or delete.
2. To edit the job, click **Edit** and edit the job mapping.

   Click the Advanced Options icon (⊟) to edit advanced options. See "Creating jobs" on page 60 for information about job options.
3. To delete the job, click **Delete**.

Managing available HTTP portals

To start, stop, or restart the HTTP portals from the Account Synchronization page:

1. In the DAS Available Jobs area, click the name of the job.
2. Click **Assign HTTP Server** or **Unassign HTTP** for a particular server.

Editing or deleting available LDAP connections

To edit or delete an LDAP connection:

1. In the LDAP Server Connectors area, click the name of the connection you want to edit or delete.
2. To edit the connection, click **Edit**, complete the form, and click **Save**.
3. To delete the connection, click **Delete**.

Viewing DAS history logs

The DAS jobs history log provides a list of job runs for each configured active job. The log includes the job name; the number of IAP users that were added, deleted, and updated; the time between job runs; the job status; and the date and time the job was completed.

The history also displays the number of IAP groups that were added, deleted, and updated.

To display the history log, scroll down to the Jobs History Log area at the bottom of the Account Synchronization page.

To display a history of the previous runs for a specific DAS job, click the job name.

To check the status of a job, view the process icon in the State column. For example, the  icon indicates the job is *Complete*.

You can also point to the process icon to display the status. If a red X icon is displayed, issues were found during the DAS run, or errors occurred in the job during its previous runs which have not yet been repaired using Account Error Recovery. For more information about the issues, check the tooltip on the status icon, Account Error Recovery Tool, and DAS logs on the HTTP portal.

6 Account Manager (AM)

Use the Account Manager (AM) page to provision and update user accounts.

This chapter contains the following topics:

To open the page, go to **User Management** > **Account Manager** in the left menu.

Account Manager overview

Use Account Manager (AM) to view and update user accounts and repositories for unusual circumstances on the IAP. Initial setup, and routine addition and deletion of user accounts, occur automatically through synchronization with Active Directory or the Domino Directory.

IAP archives emails and other documents in repositories — virtual collections of documents associated with a given user by routing rules (storing) and access lists (retrieving).

Use AM to update accounts as follows:

- Add users not imported through dynamic account synchronization (DAS).
- Change user permissions.
- Give users access to other repositories.
- Add special-purpose repositories such as audit repositories.
- Add or modify routing rules.
- Edit the retention period for some repositories.

NOTE:

You can view accounts on replica domains, but you cannot use AM to add or edit them. Fields cannot be edited and action buttons are unavailable when you select a domain that is a replica of another domain.

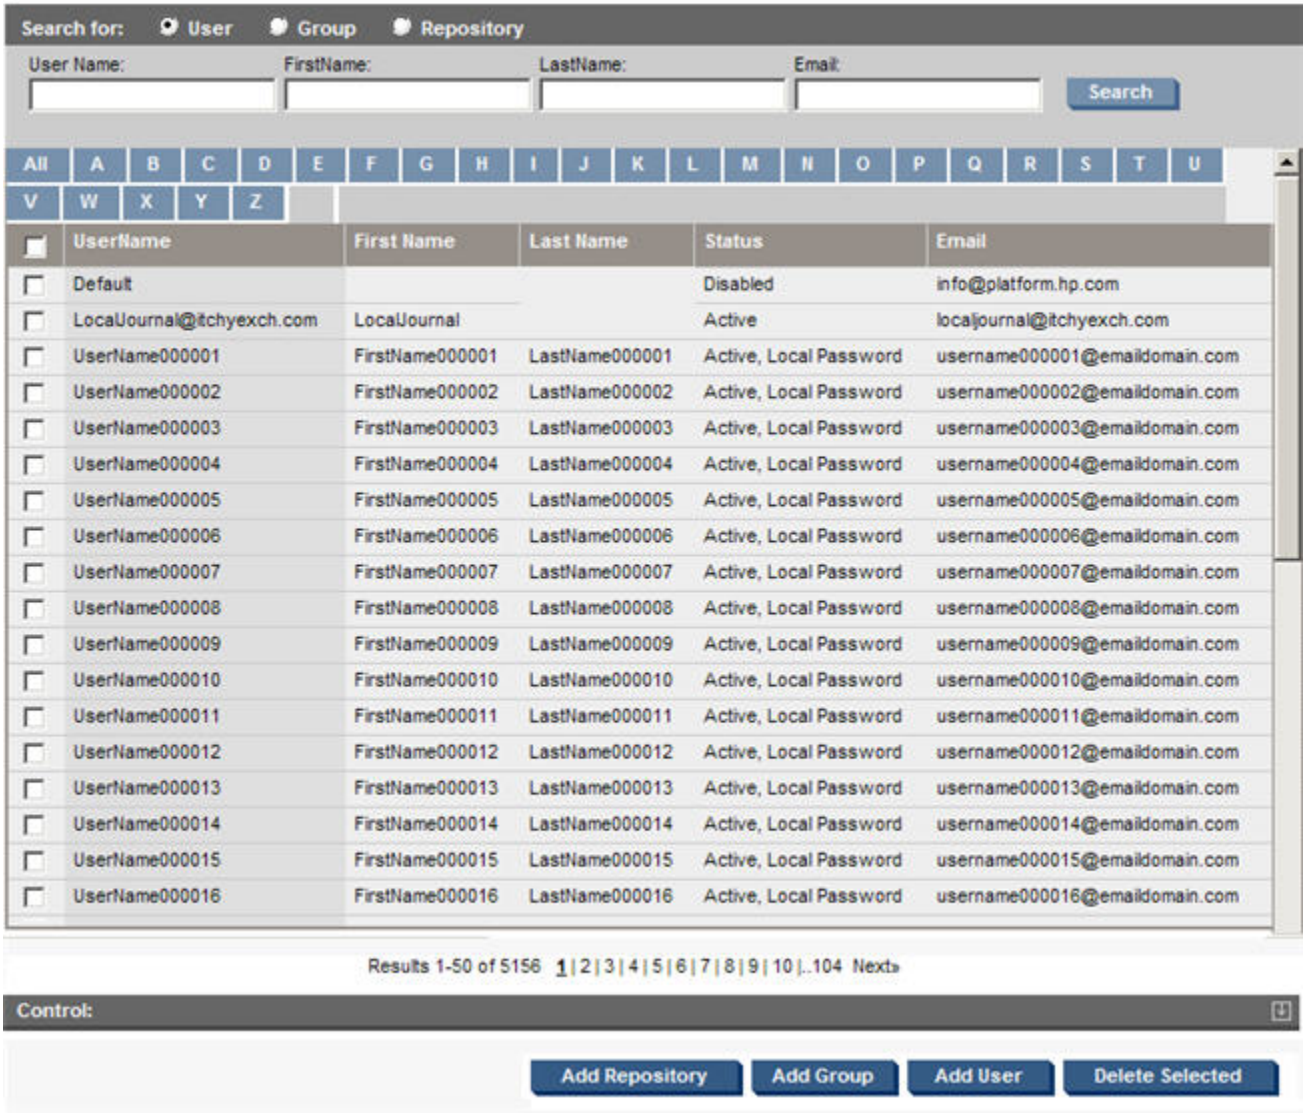

Figure 16 Account Manager page

The total number of users and groups for the domain(s) is shown in the upper right corner of the page.

Account Manager features

Table 17 Account Manager features

Feature	Description
Search button	Use the search feature to find users, groups, or repositories. The search function uses the "Like" SQL database capability. For example, in the User panel, you could enter *jack* to match users jackdoe or jacksmith. Entering *%doe* would match users jackdoe, janedoe, or maryjanedoe. Entering *%ja%* would match users jadams, jackdoe, janedoe, jacksmith, or maryjanedoe. Searches are not case sensitive. For users and groups, you can search by email account name.

Feature	Description
All check box	Select this check box to display all objects of the type indicated by filter name (user, group, or repository). For example, select the **User** filter and **All** button to show all users.
A-to-Z buttons	Use to display only names starting with the button letter. Names correspond to the objects of the current filtered panel. You can use the search function within the filtered panel.
Panels	Select the user, group, or repository radio button to view the corresponding panel. • Users (see "Managing user accounts" on page 69) - You can add users, view and edit user information, change user status and privileges, and remove users from the system. • Groups - Groups are not currently supported in the IAP. • Repositories (see "Managing repositories" on page 72). You cannot delete existing repositories.
Navigation buttons	Click the navigation buttons at the bottom of the page to: • Add a repository. • Add a user. • Delete a user from the system. Grant or revoke access to a repository

Managing user accounts

Open the Users panel to view, add, remove, or change individual user accounts on the IAP.

Adding a user in Account Manager automatically creates a primary repository for the user. It also adds routing rules and filters for the user's contact email address, and gives the user access to his or her personal repository.

When the IAP is first installed, a single user (Default) is listed in the Users panel. You can delete this entry after users are imported into the system using DAS.

Adding a new user

Users can be added directly to the system using the Account Manager. For example, it is likely that you will create an IAP administrator directly in the system, rather than granting administrative privileges to specific users imported by DAS.

Users created in the system and given a local password are known as *local* users. Users imported by DAS are known as *remote* users.

To add a new user:

1. Click **Add User** at the bottom of the Account Manager page.
2. Complete the Integrated Archive Platform Account and LDAP Information form that appears.

   See "User account information" on page 70 for details on completing this form.
3. Click **Save Now!** to save the user information.

Editing user information

To edit user information:

1. Click the **User** radio button, locate the user account that you want to edit, and then click the user name.

   See "Account Manager features" on page 68 for information on searching for a user.

2. In the Integrated Archive Platform Account and LDAP Information form that appears, clear the check box labeled **Deactivate this check box and then edit the user entries**.

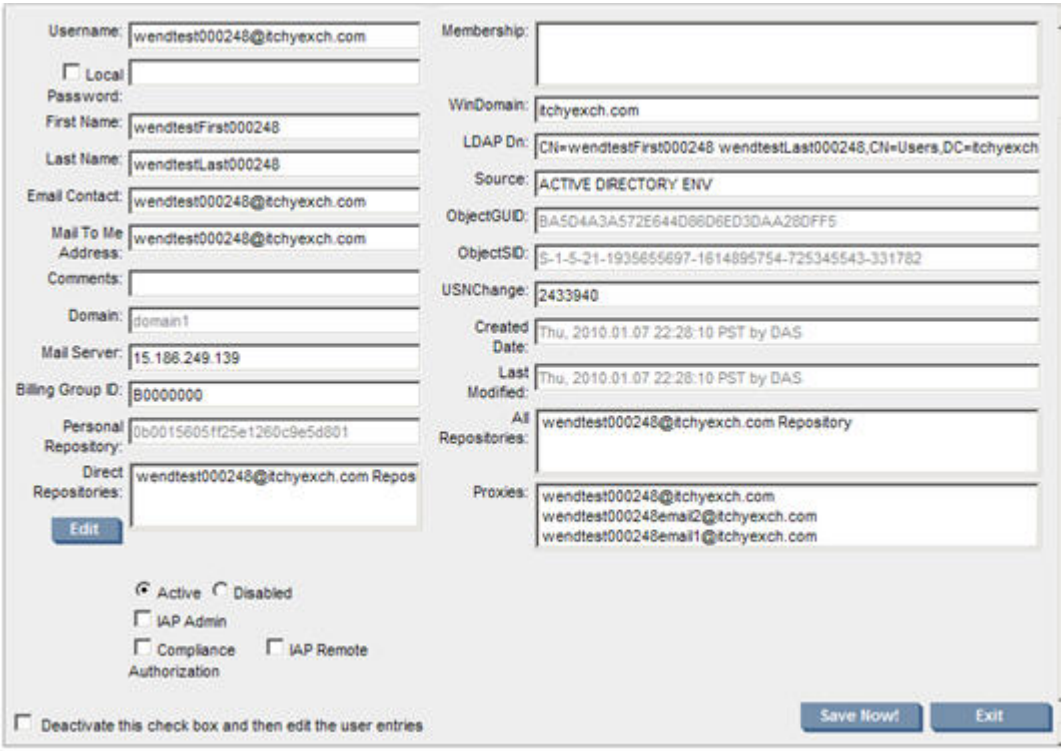

Figure 17 Editing user account information

3. Edit the relevant user entries.
4. Click **Save Now!** when you are finished.

NOTE:

To change a user's repository retention period, edit the repository. See "Editing repository information" on page 76.

User account information

Table 18 User account information

Feature	Description
Integrated Archive Platform Account Information	

Feature	Description
Username	(Required.) The system login name for the selected user. Usernames must be unique, but they can be the same except for their domains. For example, johnkdoe@company.com could be an Active Directory user imported into the system through dynamic synchronization (DAS); at the same time, user johnkdoe could be created as a local user in the IAP.
Local Password	The password in the IAP for a selected user. Active Directory users who are imported into the system through DAS are not required to have a local password. However, local users must have one.
First Name	First name of the selected user.
Last Name	Last name of the selected user.
Email Contact	(Required.) Email address for the selected user.
Mail To Me Address	(Required.) Email destination when the selected user clicks *Send results* to send a copy of an archived document from the Web Interface.
Comments	The system administrators' comments on the selected user account.
Domain	(Required.) The IAP domain to which the selected user belongs. Select the domain from the drop-down list; the selection limits the scope of the Search and A-to-Z filter buttons.
Mail Server	The IP address of the mail server for the selected user. (Optional, unless querying within Microsoft Outlook is required.)
Billing Group ID	In RISS versions prior to 1.5, this number identifies the account to be billed for services. It is not used in later RISS or IAP versions.
Personal Repository	The ID of the repository created for the selected user.
Direct Repositories	The IDs of the repositories to which the user has access. The user automatically has access to his or her own repository. If the user is a member of a group, he or she also has automatic access to the group's primary repository. You can give the user access to another repository by clicking **Edit**, selecting the check box for the repository, clicking **Add**, and then clicking **Exit**. Repositories (except for the user's personal repository) can be removed from the list by selecting the repository name and clicking **Remove**.
LDAP information	
Membership	Any groups to which a remote user belongs.
WinDomain	The Windows domain to which a remote user belongs.
LDAP Dn	The remote user's object name, relative location in the LDAP tree, and domain. CN (Common Name) is the object name, cn is its branch in the tree, and dc values are domain names. DAS supplies this value.
Source	An account synchronization maintained field that indicates which DAS job was used to create or manage the information.
Object GUID	The Global Unique Identifier of the user account on the LDAP server. This is account synchronization's key to the correct account on the LDAP server.

Feature	Description
Object SID	An identifier for the user, automatically generated and maintained by Active Directory. (Cannot be edited.)
USNChange	The USN (Active Directory unique sequence number) imported from the corresponding account on the LDAP server the last time account synchronization was run.
Created Date	The date the user account was created.
Last Modified	The date the user account was last modified.
All Repositories	All repositories to which a user has access — either through direct access or through group membership.
Proxies	Displays the email addresses that will route to the user's primary repository.
Check boxes	
Active/Disabled	Select or clear this check box to enable or disable the user account on the IAP.
IAP Admin	Select this check box to grant the user administrative privileges on the PCC. It's best to create a new, local account for the administrator. Be sure to also define a local password.
Compliance	Select this check box if the user is authorized to view all recipients (including BCC addressees) in the repositories to which the user has access. This function is generally limited to compliance officers. For a compliance officer to see any BCC information, either BCC or envelope journaling **must** be enabled on the mail server and the corresponding setting must be enabled on the gateway server. Compliance Archiving (archiving of journal mailboxes) must be enabled, as this feature is not available with Selective Archiving.
IAP Remote Authorization	An IAP remote authorization user must be in the system for replication statistics to be displayed properly in the Replication page. A user account is automatically created for this purpose, with the username of AuthorizationUser and a randomly-generated password. You can change the password for this account. You can also create a new IAP remote authorization user account and password if you want. If you create a new account, be sure to clear this check box in the AuthorizationUser account, or delete AuthorizationUser from the system.
Admin Delete	This user is authorized to delete messages from the IAP before their retention period has expired. The check box only appears when the Admin Delete service is enabled in `Domain.jcml`. See "Administrative Delete" on page 110.

Managing groups

Groups are not currently supported in the IAP.

Managing repositories

Use the Repositories pane to change the routing rules for a repository, change the retention period for some repositories, or add repositories. You could, for example, add a repository to provide a container for special routings (email addresses or email domains).

You cannot delete a repository in Account Manager.

Repository overview

The IAP archives emails and other documents into repositories. A repository is a virtual collection of documents associated with a given user by routing rules (for storing documents) and access control lists (for retrieving documents).

Repositories are used to manage retention and access to documents stored on the IAP. They define the scope of a query; users can search for documents only in the repositories to which they have access.

Email messages that are stored on the IAP can be associated with several repositories. For example, an email might be sent to several recipients and archived into each recipient's repository. The message itself is only stored once in the system, but contains references to all the associated repositories. The message remains on the IAP until the retention period of all repositories associated with the message has been exceeded.

There are several types of repositories that reside on an IAP: user repositories, access only repositories, quarantine repositories, AuditLog repositories, and system repositories. They are described in the following sections.

User repository types

The IAP supports the following types of user repositories:

- Unregulated repositories
- Regulated repositories

Unregulated repositories

Unregulated repositories are automatically created for users that are imported by DAS from Active Directory or the Domino Directory. An unregulated repository can also be created directly in Account Manager, for example to set up a local user account for the IAP administrator. (Users created in the system and given a local password are known as *local* users. Users imported by DAS are known as *remote* users.)

An email that is sent to a user's email address and ingested into the IAP is associated with the user's repository. When the user is deleted from the Active Directory or the Domino Directory and a DAS job is run, the user's repository is disabled.

The document retention period for an unregulated repository is often the same as the retention period for its IAP domain, but it can be greater. This value is set in `Domain.jcml`. For more information, see "Retention periods" on page 93.

Regulated repositories

Regulated repositories can be created for users whose activities require a different retention period from users in unregulated repositories. The retention period is the only difference between regulated and unregulated repositories. Regulated repositories can be created in Account Manager, or an existing unregulated repository can be converted to a regulated repository by editing the repository information in Account Manager. (See "Repository information" on page 77.)

The retention period for regulated repositories must be set in `Domain.jcml`, even if your organization does not use regulated repositories. For more information, see "Retention periods" on page 93.

Access only (audit) repositories

Documents sent to the IAP can be routed to an audit repository in addition to their respective user repositories, to enable compliance or legal discovery searches.

Audit repositories are special purpose repositories that allow compliance officers to execute searches within the archived documents of more than one user. This can be all employees of an organization or a subset of employees in a specific email domain.

Before IAP 2.0, audit repositories could be created using either regulated or unregulated repository types. IAP 2.0 introduced a new type of repository, the *access only* repository, which is specifically intended for audit purposes. The access only repository was created to avoid two retention issues that sometimes occurred when regulated or unregulated repositories were used as audit repositories: the retention period for the repository was too long, causing storage capacity problems on the IAP; or the retention period was too short, leading to email being deleted from the audit repository and impacting compliance searches.

The access only repository addresses this situation by keeping a document for as long as it is retained in the corresponding user or quarantine repositories. After the retention period has expired in the last user repository that is associated with the document, and any quarantine hold has been removed, the document is removed from the access only repository and deleted. This type of repository does not have its own retention period.

If your system has been upgraded from an earlier IAP version and you currently have one or more repositories used for audit purposes, they can be converted from a regulated or unregulated repository type to an access only repository type. This conversion is not required. You should only convert an existing audit repository into an access only repository if the current arrangement does not allow you to implement useful retention policies. Any conversion must be handled by HP technical support.

For the steps required to set up a new audit repository, see "Setting up an audit repository" on page 78.

Quarantine repositories

Documents that are placed in a quarantine repository are not subject to automatic deletion. As long as items are in this repository, they are preserved in a legal hold and cannot be deleted when their retention period expires.

(The exception to this rule is the deletion of quarantined messages by an Administrative Delete user. See "Administrative Delete" on page 110 for more information.)

When the hold is no longer necessary, the contents of the quarantine repository can be deleted. This action does not delete the items from the IAP. It releases documents from the legal hold and makes them subject once again to the retention period that was set for them.

Compliance officers can create one or more quarantine repositories directly in the Web Interface. In the PCC Account Manager, these repositories appear in the user's Account Information form, and are shown under the Quarantine tab in the Repositories panel.

There is no limit to the number of active quarantine repositories that a user can have.

Because documents in a quarantine repository are removed from retention processing, the retention period for this repository type is shown as -1 in the IAP Repository form.

The steps involved in creating and deleting quarantine repositories are described in "Using quarantine repositories" in the IAP user guide.

AuditLog repository

The AuditLog provides a surveillance system log for companies that are required to prove they are adhering to compliance processes. Logs are created for each user session in the Web Interface, each Duplicate Manager job, and the granting of Administrative Delete authority in the PCC. The logs are stored in the domain's AuditLog repository (*<domain>.auditlog*). This repository is created automatically if the AuditLog feature is enabled in the `Domain.jcml` file. The default retention period for the repository is seven years.

NOTE:

The minimum retention period allowed, as for all repositories, is 30 days.

For information on setting up the AuditLog and granting compliance users access to the AuditLog repository, see "Enabling the AuditLog" on page 127.

For information on performing AuditLog searches, see "Querying the AuditLog" in the IAP User Guide.

System repository types

Each of the following repositories is automatically created for an IAP domain:

- *Recycle Bin repository (<domain>.recyclebin)*: If the folder capture and/or End User Delete features are enabled for a domain, messages that no longer have email folders or user repositories associated with them are placed in the Recycle Bin repository. When the retention period for the Recycle Bin repository expires, the documents are automatically removed from IAP during retention processing. The default retention period for this repository is 30 days and can be extended.
- *CatchAll repository (<domain>.catch all)*: The CatchAll repository holds documents that fail to route into any user or audit repositories because they cannot be associated with:
  - Email addresses of registered IAP users imported by DAS
  - A domain routing rule

  The default retention period for this repository is 30 days and can be extended.
- *Failed indexing repository (<domain>.failed.indexing.repository)*: This repository holds documents that cannot be indexed or are partially indexed. (For example, the system might not be able to index the particular document type of an email attachment.)

Repository grouping on PCC pages

In the Account Manager and Retention pages, repositories are listed under the following tabs:

- Unregulated
- Regulated
- Quarantine
- Other or Other Type

  Includes the AuditLog, failed indexing, CatchAll, Recycle Bin, and access only repositories for a domain.

Adding repositories

User accounts and repositories are created automatically for each user that is imported by the DAS process. You can also manually create the following types of repositories in Account Manager:

- Unregulated repositories
- Regulated repositories
- Access only repositories

See "Repository overview" on page 73 for information about these repository types.

To manually add a repository from the Account Manager:

1. Click **Add repository** on the Account Manager page.
2. In the IAP Repository form that appears:
   - Enter a unique name for the repository in the Name box.
   - In the Domain box, select the IAP domain in which the repository will reside.
   - In the Retention box, enter the number of days that documents should be retained in the repository. The field will be populated with the default value for the specified repository type as configured in `Domain.jcml`. This value can be modified but must be equal to or greater than the number of days specified for the domain and the repository type in the `Domain.jcml` file.

     For an access only (audit) repository, this field cannot be modified.
   - In the Type box, select the repository type from the drop-down menu.

     You can create repositories of the following types: Unregulated, Regulated, or Access Only.

     If you select Access Only, a dialog box appears stating that this repository type should only be used for audit repositories. Click **OK** to close the dialog box.
   - In the Add Email box, specify the user's email address(es). For example, user@company.com.

     Do not complete this field for an access only (audit) repository.
   - In the Add Email Domain box, specify the email routing domain to be associated with the repository. For example, company.com.

     You do not need to complete this field for a user repository.
3. Click **Save Now!**

For more information on the fields in the IAP Repository form, see "Repository information" on page 77.

For information on creating access only (audit) repositories, see "Access only (audit) repositories" on page 74.

Editing repository information

To edit an IAP Repository form:

1. Click the **Repository** radio button on the Account Manager page.
2. Click the tab for the type of repository: Unregulated, Regulated, or Other (for access only repositories).
3. Locate and click the name of the repository you want to edit.

4. In the IAP Repository form, clear the check box labeled **Deactivate this check box and then edit the user entries**.

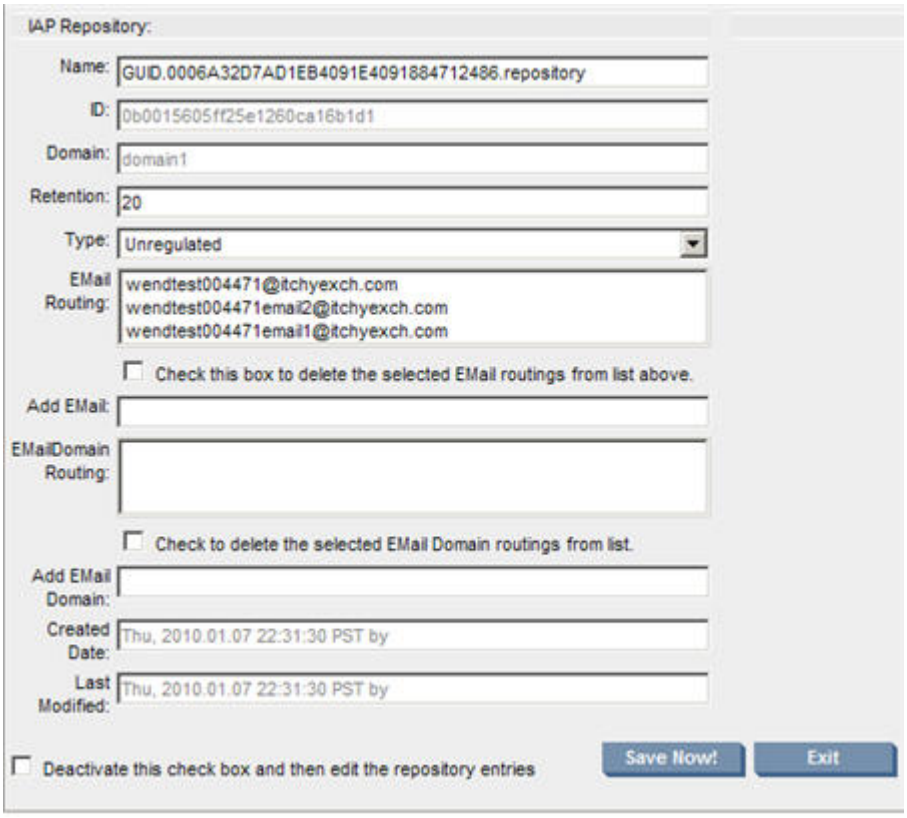

Figure 18 Editing repository information

5. Edit the relevant entries, using the descriptions in Repository information.
6. Click **Save Now!** to save the repository information.

Repository information

Table 19 Repository information

Feature	Description
Name	(Required.) Enter a name if adding a repository.
Domain	The IAP domain to which the selected repository belongs. Select the domain from the drop-down list; the selection limits the scope of the Search and A-to-Z filter buttons.
Retention	Enter the number of days that documents are retained in the repository. (For example, 30=30 days.) For regulated or unregulated repositories, this value cannot be lower than two values set in the `Domain.jcml` file: • The retention period specified for the domain • The retention period specified for the repository type (See "Retention overview" on page 92 for more information on these values.) Do not complete this field if you are creating an access only repository.

Feature	Description
Type	If adding a repository, you can select one of the following types: • Regulated • Unregulated • Access Only If changing the type of repository, you can: • Change a regulated repository to an unregulated repository. • Change an unregulated repository to a regulated repository. Access only repositories cannot be changed to another repository type.
Email routing	Mailing address(es) for email routing. To delete one or more addresses, select the relevant address(es) in the list, and then select the check box beneath the field.
Add Email	Specify the email addresses associated with this repository (for example, user@company.com). This entry appears in the EMail routing box after the repository information is saved.
Email Domain Routing	The email domain associated with the repository.
Add Email Domain	Specify the email routing domain to be associated with this repository (for example, company.com). This entry appears in the EMail Domain Routing box after the repository information is saved.
Created Date	The date the repository was created. (Cannot be edited.)
Last Modified	The date the repository was last modified. (Cannot be edited.)

Setting up an audit repository

Audit repositories are special purpose repositories that allow compliance officers to execute searches within the archived documents of more than one user. This can be all employees of an organization or a subset, such as users with significant legal exposure.

Perform the following steps to set up a repository for audit purposes:

1. Create the repository.
   a. Follow the instructions in "Adding repositories" on page 75.
   b. Specify the email routing domain to be associated with the repository in the **Add Email Domain** box.

      See Configuring routing rules for an audit repository for details.
2. Give one or more compliance officers access to the audit repository.
   - To give a user compliance privileges, select the **Compliance** check box in the user's account form.
   - To give a compliance user access to the audit repository, select the repository in the **Direct Repositories** box in the user account form.

   See "Editing user information" on page 70.

Configuring routing rules for an audit repository

Audit repositories provide access to email messages based on the email routing domain. Specify the email routing domain to be associated with the repository in the Add Email Domain box in the IAP Repository form. When the form is saved, the domain is saved in the Email Domain Routing field. (See "Repository information" on page 77.)

In the following example, messages containing at least one email address with the @company.com email domain will be routed into the repository named Audit. A compliance officer with access to this repository will be able to search for messages sent to any user that matches the email address pattern: `<user>@company.com`.

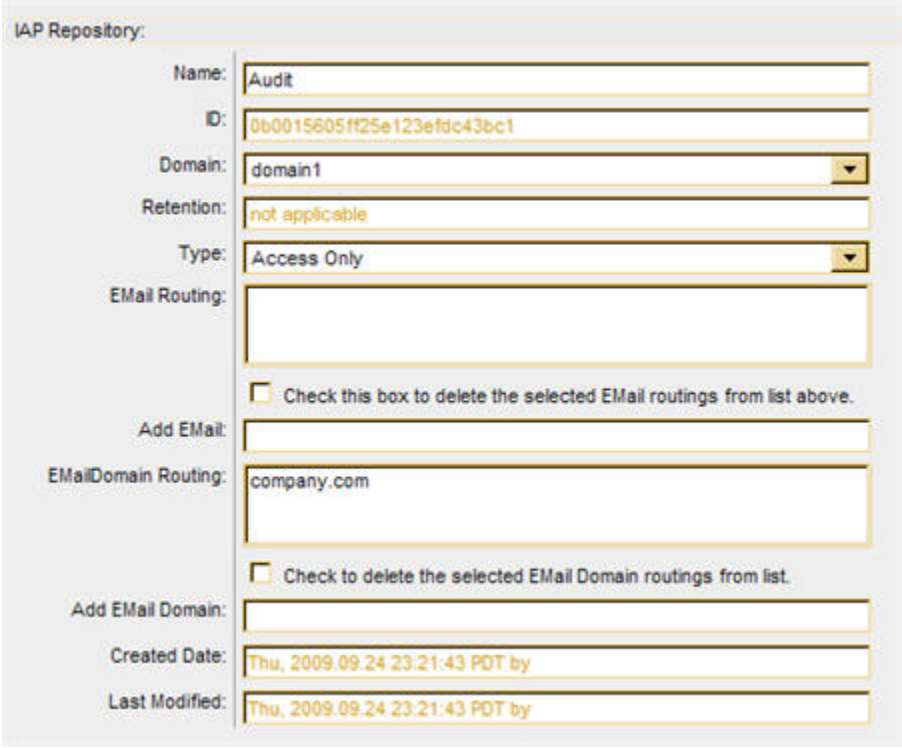

Figure 19 Repository form for an audit repository

Audit repositories can also be created to segregate messages based on email domain. For example, if your organization has several email domains, you can create individual audit repositories for each email domain.

You could also create a master audit repository that was associated with all email domains for the organization, then create individual audit repositories for each specific email domain. This would allow compliance officers to limit or widen a search as needed.

7 Account Error Recovery

The Account Error Recovery page displays account synchronization activities that have not been performed successfully.

To open the page, go to **User Management** > **Account Error Recovery** in the left menu.

Error Recovery overview

Unsuccessful account synchronization activities can include:

- Not being able to add a new user because the user name or object GUID (Global Unique Identifier) already exists
- Not being able to update a user because the user's entry cannot be found
- Not being able to remove a user because the user's entry cannot be found
- Not being able to add an alias or email routing rule for a user's repository
- Not being able to add a membership to a group

Activity errors are shown grouped by object GUID, with the most recent error at the top of the list. You can select an activity to attempt it again, or you can delete it.

Activities can be filtered by clicking **User**, **Group**, or **Membership** at the top of the page.

Error Recovery features

Table 20 Error Recovery features

Feature	Description
Date	Date the error occurred.
Error	Type of activity error. For example, not being able to add an alias or email routing rule for a user's repository.
USN	Active Directory unique sequence number. (User or group filters only.)
UserName	The name of the user account. (User filter only.)
GroupName	The name of the group account. (Group filter only.)
Group	Automatically generated identifier for the selected group; unique to the system. (Membership filter only.)
Member	The name of the repository. (Membership filter only.)

Feature	Description
MemberType	The type of repository: • Regulated • Unregulated • Quarantine • Other (includes the AuditLog, failed indexing, CatchAll, Recycle Bin, and access only repositories for a domain) (Membership filter only.)
Object GUID LDAP DN	The Global Unique Identifier of the user account on the LDAP server; its corresponding object name, relative location in the LDAP tree, and domain. CN (Common Name) is the object name, cn is its branch in the tree, and dc values are domain names. Account synchronization (DAS) supplies these values. (User filter only.)

Repairing synchronization errors

To repair synchronization activity errors, identify the activities that you want to reattempt or delete. Click an entry to display more information about an activity, including its Java User Management Services (UMS) database entry (only shown if a UMS database entry matches the activity).

You will need to decide the order in which to reattempt or delete the activities.

To repair or delete synchronization errors:

1. Identify the activities that you want to reattempt and select the check box in front of those entries.

   You can click the entry in the UserName or GroupName column to find out more about the error.
2. Click **Apply Selected**.

   If the reattempt is successful, the entry is removed from the list.
3. Identify the activities that you want to delete and select the check box in front of those entries.

   You can click the entry in the UserName or GroupName column to find out more about the error.
4. Click **Delete Selected**.

   When an activity is deleted, it is immediately removed from the list and cannot be recovered.

8 Data management

This chapter discusses the following topics:

Duplicate Manager

The Duplicate Manager page lets you schedule duplicate merge jobs and view the status of duplicate merge jobs.

NOTE:

Duplicate Merge is only supported for Exchange email.

To open the page, go to **Data Management** > **Duplicate Manager** in the left menu.

Duplicate Manager overview

Duplicate Manager is intended for organizations that are storing with single instancing enabled. Single instancing is the ability to store a single copy of an email that exists in multiple user repositories. Emails are uniquely identified by a cryptographic checksum called a hash. If an email client attempts to store the same email twice (as determined by the message hash) only one copy is physically stored and the same reference URI is returned for both store requests.

Storing duplicate emails uses unnecessary space on the Smartcells and clutters search results by listing the same email multiple times. Once the email file has been stored there is additional data attached to the file called metadata. Over the lifecycle of the email, the metadata will change as user access permissions and folder information change. The purpose of the Duplicate Manager tool is to eliminate duplicate emails while retaining all associated metadata. To that end, the Duplicate Manager will take the union of all metadata and attach it to a single copy of the email and remove all redundant copies.

If a user previously ran a successful query that returned duplicate emails and saved the query results, after running this tool all but one of the duplicate emails should be accessible. If a user has quarantined the files in the saved query result set after running this tool, the single remaining copy of the email will remain quarantined.

Smartcells have a storage threshold that is lower than the total capacity of the hard disk. This is because the services that run on the Smartcells require a percentage of the disk space in order to operate. Once the Smartcell reaches this threshold it will transition into a closed state. Because duplicate emails can be stored across Smartcell groups, there is a possibility that the merge job will not be able to work on a set of duplicates because the Smartcells have exceeded the storage threshold and are no longer able to accept additional metadata. These duplicates cannot be merged until additional space has been made available on the closed Smartcell by either migrating some of the emails to a new group or waiting for retention to delete emails and free additional storage space.

The amount of time for a merge job to complete is dependent on the total number of files stored, the percentage of those files that are duplicates, and the average number of duplicates stored per hash. The first merge job can take many days to complete, while subsequent merge jobs can complete much faster.

If no errors occur during the duplicate merge, duplicates are significantly reduced in the system. However, there is no guarantee that all duplicates are removed. Some emails might not have an associated hash, and cannot be merged into a single copy. The duplicate merge procedure might also miss a few duplicate emails on the first pass, since duplicate merge speed was traded for an exhaustive duplicate elimination. Duplicate merge does not attempt to merge all emails in a single pass. It merges all emails in the system in manageable batches according to default or user configurable batch size. As a result of this batching, unmerged duplicate emails often occur at the beginning or end of the batches. A subsequent duplicate merge pass will remove the remaining duplicate emails.

Duplicate Manager job schedules

The Duplicate Manager Schedule shown on the top portion of the Duplicate Manager page shows each domain, domain portal IP address, days of the week that a duplicate merge job is scheduled, time for which it is scheduled, and whether the job status is enabled.

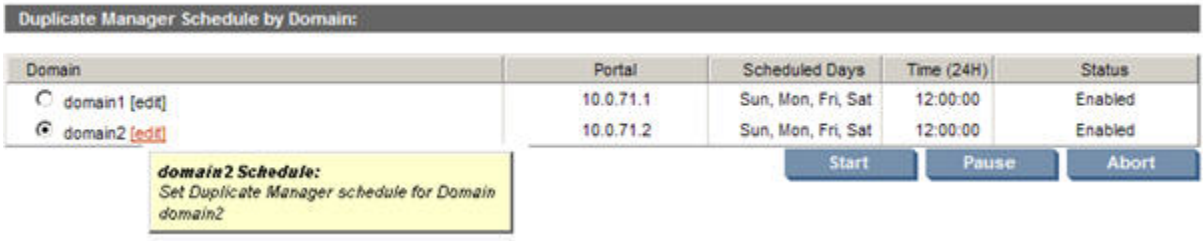

Figure 20 Duplicate Manager job schedules

Scheduling a job

To schedule a duplicate merge job:

1. In the Duplicate Manager Schedule by Domain section, click the **edit** link next to the domain for which you want to schedule a job.
2. On the resulting schedule page, use the check box to enable the job, select the HTTP portal where the job will run, the days of the week to run the job, the time to run the job, and then click **Save Schedule Now!**

TIP:

For optimal performance this tool should run on a different HTTP portal than DAS and replication.

Enabling or disabling a job

To enable or disable a duplicate merge job:

1. In the Duplicate Manager Schedule by Domain section, click the **edit** link next to the domain for which you want to enable or disable a job.
2. On the resulting schedule page, select or deselect the **Enable** check box and then click **Save Schedule Now!**.

Starting, pausing, or aborting a job

To start, pause, or abort a duplicate merge job:

Use the radio button to the left of the desired domain, and click **Start**, **Pause**, or **Abort** as appropriate. Start will start a duplicate merge job on the selected domain immediately. Pause will pause a duplicate merge job currently running on the selected domain. Abort will abort a currently running or paused duplicate merge job on the selected domain.

Duplicate Manager job histories

The Duplicate Manager History Logs displayed on the bottom portion of the Duplicate Manager page shows information for the five latest duplicate merge job runs. You may also click a domain to see all (up to 50) duplicate merge job history logs for that domain.

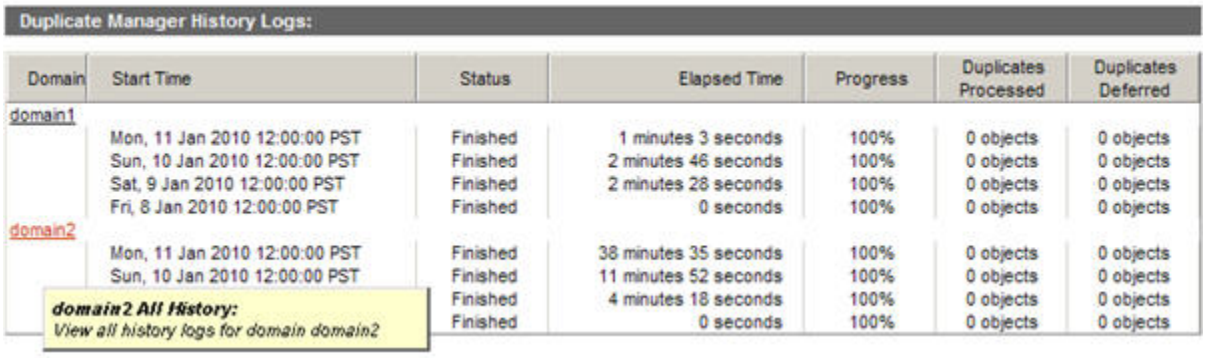

Figure 21 Duplicate Manager job history

The Duplicate Manager History Log provides the following job history information.

Feature	Description
Domain	Domain name
Start Time	Time stamp of when Duplicate Manager job started
Status	Initializing, Started, Not Available, Failed, or Finished
Elapsed Time	Time since a job was started until finished, aborted, or paused
Progress	Percentage of Duplicate Manager processed sets compared to the total
Duplicated Processed	Number of duplicates processed

Feature	Description
Duplicates Deferred	Number of duplicates that could not be processed by the Duplicate Manager in the job.

Replication

NOTE:

This page is available only if a replicated system is configured.

Use the Replication page to monitor and to start or stop replication for a domain on a remote system.

To open the page, go to **Data Management** > **Replication** in the left menu.

Replication overview

Replication status is updated after each polling cycle. Therefore, it could be up to five minutes after you start replication before you see results on the graph of replication rates. Errors and warnings, however, are displayed as soon as they happen on the system. (The Replication page is unavailable until at least one domain on the PCC host system is configured for replication.)

The behavior of the primary Smartcell groups on the replication process is the following: The primary Smartcell of a primary group by default is the one responsible of replicating the email. If this primary Smartcell fails a failover occurs. A failover is the action performed by the secondary Smartcell of taking over the responsibility to replicate email for its group. When the primary Smartcell is restarted it will regain control of replicating email, this is known as failback. The intention of the failover and failback mechanisms is to be able to provide fault tolerance by taking advantage of the Smartcell redundancy that IAP provides and continue replicating even if the primary Smartcell is down and to replicate from the primary Smartcell whenever it is up and healthy.

The behavior of the replica Smartcell groups in an IAP installation is basically to ingest messages from its primary Smartcell group. In order for the ingestion to happen both replica 1 and replica 2 Smartcells need to be available and healthy; or, in the case of a IAP installation configured to only have one replica Smartcell, only replica 1 needs to be available and healthy. Once a replica group is allocated no further allocation will occur in the case where the replica Smartcell(s) are unavailable for any reason. In the case where the replica Smartcell(s) are unavailable for any reason, no allocation for replica Smartcells will occur, and the replication will fail until the replica Smartcell(s) become available again.

If the primary site is not available for queries, users must enter the address of the replica site, instead of the primary site, in their browsers.

NOTE:

In order for the Replication page to display statistics from the remote system, an IAP remote authorization user must be set up in Account Manager. See IAP Remote Authorization in “User account information” on page 70.

Database Replication

The top part of the Replication page describes the database replication.

Table 21 Database Replication features

Feature	Description
Local Server/ Source Server	The replication and primary systems if you are logged into the replicated system.
Local Server/ Target Server	The primary and replication systems if you are logged into the primary system.
Replication Set	A set is made of one or more DB2 tables that replicated.
Inbound Sync Time/ Outbound Sync Time	The time at which the database tables in the replication set got last synchronized. The synchronization time is displayed in two colors: green and red. Green means that the database tables on the primary and replica DB2 have been synchronized successfully within the last 30 minutes. Red means that the synchronization has not happened for 30 minutes or more; this could indicate a DB2 replication issue and you should contact HP technical support for advice.
Inbound Status/ Outbound Status	The replication status on the replicated and primary systems.

Re-initializing DB2 replication

Follow these steps to remove and re-initialize DB2 replication on the primary and replica systems. Be sure to wait for each step to complete before proceeding to the next step.

1. Stop the IAP applications on the replica system:

   PCC server on replica: `/opt/bin/stop`
2. Restart DB2 on the replica system:

   DB2 server on replica: `/etc/init.d/db2_app start`
3. Stop the IAP applications on the primary system:

   PCC server on primary: `/opt/bin/stop`
4. Restart DB2 on the primary system:

   DB2 server on primary: `/etc/init.d/db2_app start`
5. Unconfigure DB2 replication on the replica system:

   DB2 server on replica: `/opt/bin/repl/dbRepl rm`
6. Unconfigure DB2 replication on the primary system:

   DB2 server on primary: `/opt/bin/repl/dbRepl rm`
7. Re-initialize DB2 replication on the primary system:

   DB2 server on primary: `/opt/bin/repl/dbRepl init`
8. Re-initialize DB2 replication on the replica system:

   DB2 server on replica: `/opt/bin/repl/dbRepl init`

9. Wait for the replication process to complete.

   On average, it should take less than 30 minutes for the DB2 databases to synchronize.

   Use the following steps to monitor DB2 replication and determine when the DB2 databases on the primary and replica systems are synchronized:

   a. Log on to the DB2 server of the replica IAP and enter the following commands:

```
ssh db2
su - db2udb
db2 connect to ark01db
while true
do
   date
   db2 "select set_name,status,synchpoint,synchtime from asn.ibmsnap_subs_set"
   sleep 10
done
```

   b. Wait until:

   - All items listed in STATUS are 0.

   NOTE:

   If one or more items in STATUS is -1, replication is unlikely to succeed. You might want to start over and repeat steps 1–8.

   - All items listed in SYNCHPOINT have values.
   - All items listed in SYNCHTIME have values and are within a few minutes of the current time.

   This is an example of success:

```
SET_NAME            STATUS SYNCHPOINT               SYNCHTIME
------------------- ------ ---------------------- --------------------------
DB2UDB.base              0 x'4A9DA0D8000000010000' 2009-09-04-16.08.39.244175
DB2UDB.base              0 x'4A9DA113000000010000' 2009-09-04-16.08.51.000000
HAPPY.UBQ                0 x'4A9DA122000000010000' 2009-09-04-16.09.51.955488
HAPPY.UBQ                0 x'4A9DA0D8000000010000' 2009-09-04-16.08.39.244175
GRUMPY.UBQ               0 x'4A9DA0D8000000010000' 2009-09-04-16.08.39.244175
GRUMPY.UBQ               0 x'4A9DA122000000010000' 2009-09-04-16.09.51.955488
HAPPY.D2D                0 x'4A9DA122000000010000' 2009-09-04-16.09.16.904729
HAPPY.D2D                0 x'4A9DA0D8000000010000' 2009-09-04-16.08.04.184687
GRUMPY.D2D               0 x'4A9DA0D8000000010000' 2009-09-04-16.08.04.184687
GRUMPY.D2D               0 x'4A9DA122000000010000' 2009-09-04-16.09.51.955488
CSFR                     0 x'4A9DA0D8000000010000' 2009-09-04-16.08.39.244175
CSFR                     0 x'4A9DA122000000010000' 2009-09-04-16.09.16.904729
```

   This is an example of failure:

```
SET_NAME            STATUS SYNCHPOINT               SYNCHTIME
------------------- ------ ---------------------- --------------------------
DB2UDB.base              0 x'4A9DA0D8000000010000' 2009-09-04-16.08.39.244175
DB2UDB.base             -1 x'4A9DA113000000010000' 2009-09-01-15.32.51.000000
HAPPY.UBQ                0 x'4A9DA122000000010000' 2009-09-04-16.09.51.955488
HAPPY.UBQ                0 x'4A9DA0D8000000010000' 2009-09-04-16.08.39.244175
GRUMPY.UBQ               0 x'4A9DA0D8000000010000' 2009-09-04-16.08.39.244175
GRUMPY.UBQ               0 x'4A9DA122000000010000' 2009-09-04-16.09.51.955488
HAPPY.D2D                0 x'4A9DA122000000010000' 2009-09-04-16.09.16.904729
HAPPY.D2D                0 x'4A9DA0D8000000010000' 2009-09-04-16.08.04.184687
```

```
GRUMPY.D2D              0 x'4A9DA0D8000000010000' 2009-09-04-16.08.04.184687
GRUMPY.D2D              0 x'4A9DA1220000000010000' 2009-09-04-16.09.51.955488
CSFR                    0 x'4A9DA0D8000000010000' 2009-09-04-16.08.39.244175
CSFR                    0 x'4A9DA1220000000010000' 2009-09-04-16.09.16.904729
```

10. After replication is complete, restart the applications on the primary system:

    PCC server on primary: `/opt/bin/start`

11. Restart the applications on the replica system:

    PCC server on replica: `/opt/bin/start`

Data Replication Flow and Detail Information

The middle part of the Replication page displays the flow of replicated data. You can also suspend or resume the replication process per domain from this area.

Table 22 Data Replication Flow

Feature	Description
Domain	The domain name and group ID of the domain being replicated, and the following information: • The IP addresses of the Smartcells being duplicated, the status of the replication, and the number of stored and indexed documents being copied. A check icon indicates normal operation. An X icon indicates that replication failed the last time it was tried. An ! icon indicates replication is being retried. • An arrow showing the direction of data. The direction is left to right from the primary system to the replicated system. The direction is right to left on cross replica configurations, reflecting that a primary Smartcell group in the remote IAP installation is replicating to a replica Smartcell group on this IAP. • The number of stored and indexed documents that have been copied to the replicated domain. • The percentage of data that has been copied to the replicated domain.
Suspend/Resume	To stop replication, click **Suspend**. Replication will stop after the current batch is replicated. To start replication, click **Resume**. The batch that would have been replicated next when replication stopped will be replicated.

The Data Replication Flow and Detail Information section also includes a More Details link. Click **More Details** to open the Data Replication Detail Information page to see additional primary and replica system statistics. Statistics are provided by system group name and include Stored Objects, Indexed Objects, Active Transactions, Active Entries, Stored Resources, Stored Chunks, Stored References, and Used Diskspace.

Replication Status

The bottom part of the Replication page displays the status of the replication.

Table 23 Replication Service General Status

Feature	Description
Domain	The domain name of the domain being replicated.
Group ID	Group ID of the domain being replicated.
State	Shows whether or not replication is in progress.
File Batch Count	The number of file batches that have been iterated to replicate. A batch is composed of 100 messages; therefore, file batch count is the number of replication attempts for a batch of 100 messages. (Under certain conditions a batch might get partially replicated, meaning only an *x* number of files out of the 100 will be successfully replicated. In such a case, replication attempts for the batch with the remaining 100-*x* also increases the file batch count.
Update Time	The time of the last replication update.
Next Retry Time	The date and time of the next replication update.
Message	Any messages about the last replication update.

Reprocessing

The Reprocessing page displays each domain and shows whether reprocessing is enabled, when reprocessing is scheduled, and a history log report. Reprocessing is an engine service that scans recently stored data and reallocates messages to the proper repository.

Most reprocessing occurs as an automated and transparent background process that does not require any operational effort other than passive monitoring.

In most cases, messages are reprocessed to reapply routing rules that have been recently added and are applied to a short time period before the change has occurred. The Reprocessing view allows you to schedule and enable reprocessing based on new routing rules.

NOTE:

When you make a change using the Reprocessing Utility, the change does not take place immediately. It is put in the job queue, and the reprocessing itself occurs 24 hours later.

To open the page, go to **Data Management** > **Reprocessing** in the left menu.

Rescheduling all reprocessing schedules

To reschedule all the domains to the same reprocessing schedule:

1. In the Reprocessing page, click **Reschedule All**.
2. Complete the form to set the status and schedule.
3. Click **Reschedule All**.

4. To ensure that all schedules are enabled, verify that all Reprocessing Status check boxes are selected.

   If selected, the text *Enabled* appears next to the check box.

Editing reprocessing schedules

To edit a domain's reprocessing schedule:

1. In the Reprocessing page, click the **edit** link next to the domain you want to edit.

Figure 22 Editing reprocessing schedules

2. Complete the form to set the status and schedule.
3. Click **Save Schedule**.
4. To ensure the schedule is enabled, verify that the corresponding Reprocessing Status check box is selected.

   If selected, the text *Enabled* appears next to the check box.

Changing the reprocessing status

You can enable or disable a reprocessing schedule of a specific domain or all the domains listed.

To change a reprocessing schedule:

1. In the Reprocessing page, locate the domain whose schedule you want to enable or disable.
2. Complete any of the following tasks to change the status:

Figure 23 Change the reprocessing status

- To enable reprocessing for a specific domain, select the corresponding Reprocessing Status check box.
- To disable reprocessing for a specific domain, clear the corresponding Reprocessing Status check box.
- To enable reprocessing for all domains, click **Resume All**.
- To disable reprocessing for all domains, click **Suspend All**.

3. Verify the status by noting the text *Disabled* or *Enabled* next to the Reprocessing Status check box.

Using the Reprocessing Utility

You can reprocess a user repository if you think a user hasn't been included in the processing.

In the Reprocessing Utility area of the page:

Figure 24 Reprocessing utility

1. Enter the following information in the text boxes:
   - The user email address
   - The user repository ID

   Users can only be reprocessed if their email address and repository are in the IAP. If so, they are listed in the Account Manager.
2. In the Domain Name drop-down list, select the domain in which the user repository resides.
3. Specify a date range, and click **Add in Reprocessing Queue**.

   The job is sent to the queue, and the reprocessing takes place 24 hours later.

Viewing reprocessing history logs

The reprocessing history log at the bottom of the Reprocessing page shows a list of the last successful reprocessing runs for each configured domain. The log includes each domain group, when a domain was reprocessed last, the number of processed files, and the elapsed time of the run. This report is based on data from Smartcell groups, which is averaged from the individual Smartcells of each group.

Retention

Retention periods are configured for an IAP domain and its repositories when the system is installed. Use the Retention page to change the amount of time that documents are kept before they are deleted from the system.

Retention is useful in meeting compliance requirements, when your company must retain emails or files for a specified number of years.

To open the page, go to **Data Management** > **Retention** in the left menu.

Retention overview

Retention determines how long a document resides on the IAP before it is subject to automatic deletion and permanently deleted from the system. Retention periods can be applied to IAP domains and

repositories. Documents remain physically on the IAP until all retention periods that are associated with the item are exceeded.

This section describes how documents are retained in the system and includes the following topics:

- Retention basis
- Retention periods
- End User Delete
- When the retention period expires

Retention basis

Documents are retained based on either the send date or the archive date (also known as the ingest date).

The send date for an email represents the date on the email client when the email was sent or received. The send date for a file, stored via the HP File Archiving software or the Object Storage API interface, represents the "last modified timestamp" of the file.

The ingest date (archive date) of an email or file is the date it was stored in the IAP, using the local time on the SMTP server.

The retention basis is set in the `Domain.jcml` file on the kickstart server. One of the following options should be listed in the file:

```
RetentionBasis=IngestDate
RetentionBasis=SendDate
```

If neither of the options is specified in the `Domain.jcml` file, the configuration defaults to the archive date.

If the retention basis is changed from the archive date to the send date, email with an older send date but recently archived (for example via the HP EAs Exchange PST Importer) might be deleted the next time a retention job runs, if the send date is outside the retention period for the domain and repository.

The PCC Domain Configuration page and the Retention page show the retention basis.

NOTE:

Whenever you configure any retention settings in the `Domain.jcml` file, run the following commands:

`regloader.pl –cv –clearallConfirm=<IAP name>` on the kickstart server and `/opt/bin/restart` on the PCC to restart the IAP.

Retention periods

When the system is set up, three retention periods are configured in `Domain.jcml`:

- The retention period for the domain:

  `DomainRetentionPeriodDays=[number of days]`

  This is the minimum number of days that documents are stored in an IAP domain's repositories before they are automatically deleted from the system. Values can be -1 (which disables retention) or can range from 30 (days) to 2,147,483,647 (days).

  After the system is put into production, domain retention can only be altered using the PCC Retention page. You can enable or disable retention or increase the domain retention period. The domain retention period cannot be decreased after the IAP starts archiving documents.

When retention is enabled on the Retention page, the -1 setting in `Domain.jcml` has no effect. Configuration changes made via the PCC retention page will override the settings made in the `Domain.jcml` file.

- The retention period for unregulated repositories:

  `UnregulatedRetentionPeriodDays=[number of days]`

  Unregulated repositories are created for users imported by DAS from the Active Directory or the Domino Directory. The unregulated retention period is used as the default retention period when the repositories are created. If necessary, individual repository retention periods can later be changed via the PCC Account Manager page.

  Setting the value to -1 one does not disable retention.

  You can change the unregulated retention period by increasing the value in the `Domain.jcml` file.

- The retention period for regulated repositories:

  `RegulatedRetentionPeriodDays=[number of days]`

  Regulated repositories can be created for users whose activities require a different retention period than users in unregulated repositories. These repositories can either be created in Account Manager, or an existing unregulated repository can be converted to a regulated repository by editing the repository information in Account Manager.

  If your organization uses regulated repositories, the retention period is usually greater than the period for domain retention. A value must be configured for this field, even if you do not use regulated repositories. Setting the value to -1 one does not disable retention.

  You can change the regulated retention period by increasing the value in the `Domain.jcml` file.

All three retention periods must be set for retention to work. The retention periods are displayed on the PCC Domain Configuration page and the Retention page.

See "Editing domain retention periods" on page 96 and "Editing repository retention periods" on page 99.

Exceptions

The configurable retention periods do not apply in the following cases:

- Quarantine repositories: Documents placed in a quarantine repository are preserved in a legal hold and cannot be deleted when their retention period expires. See "Quarantine repositories" on page 74 for more information on these repositories.
- Admin Delete: A user with administrative delete privileges can manually delete documents from a quarantine repository, or a regulated or unregulated repository before its minimum retention period has expired. See "Administrative Delete" on page 110 for more information.

End User Delete

What is End User Delete?

If users delete an archived email from an Outlook folder, the email is not automatically deleted from the IAP. Users are deleting a "tombstone," a pointer to the actual email that is now stored on the IAP. (When users select the message in their email client, they are viewing a retrieved copy of the message.)

When a tombstone is deleted, and retention and End User Delete are enabled in the `Domain.jcml` file, the reference to the user's repository is removed from the archived message. If folder capture is enabled, references to the Outlook folders where the tombstone was located are also removed from the archived item. The message is deleted from the user's repository, and the user can no longer

access it—either from Outlook, the Web Interface, or Outlook Integrated Archive Search. (This does not affect other users who access the message via their own repositories.)

The End User Delete function applies to items in regulated or unregulated repositories, and is used in conjunction with HP EAs Exchange Outlook clients. An archiving event, Synchronize Deleted Items, must be scheduled in EAs Exchange so that removal of a tombstone is synchronized with removal of the message in the user repository.

End User Delete removes a message from a user repository. The archived message itself can be removed from the IAP only when all repository (and folder) references are removed from the message. Messages that no longer have email folders or user repositories associated with them are placed in the Recycle Bin repository when folder support and/or End User Delete are enabled for a domain. After the retention period for the Recycle Bin repository expires (the default is 30 days), the messages are automatically removed from the index and the Smartcell during the next retention job.

NOTE:

When messages are placed in the Recycle Bin repository, they are re-indexed. The items can be searched for in the Recycle Bin repository but cannot be searched for in the user repository from which they were removed.

Configuring End User Delete

Use the following fields in `Domain.jcml` to configure End User Delete:

- `MinUnregulatedRetentionPeriodDays`
- `MinRegulatedRetentionPeriodDays`

The value in each field determines the amount of time an email remains in a user repository after the tombstone is deleted from Outlook. If set to `-1`, End User Delete is disabled. If set to `>=0`, End User Delete is enabled for the repository type (regulated or unregulated).

When the value is set to `0`, messages can be deleted from the user's repository immediately. When the value is the same as the retention period for the repository type (`RegulatedRetentionPeriodDays` or `UnregulatedRetentionPeriodDays`), messages cannot be deleted via the Outlook client. (This is common in compliance environments.)

Archived messages are removed automatically from the IAP when the retention period has expired for each repository in which the message was located.

The following diagram displays how the process works.

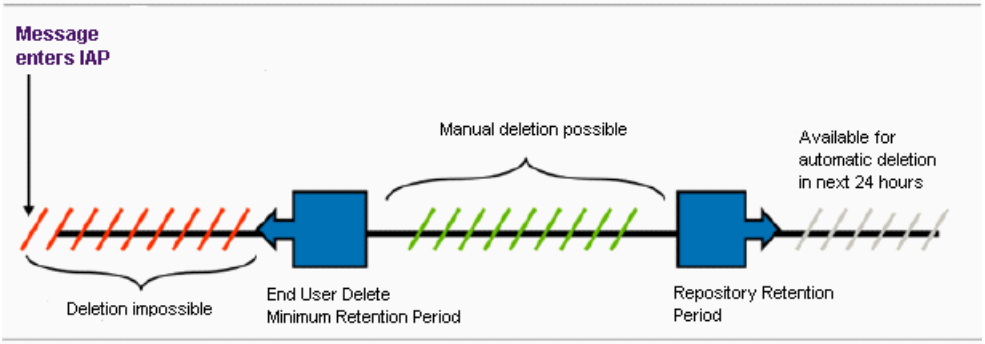

Figure 25 End User Delete and retention

Example

An example of a configuration with End User Delete could be:

```
DomainRetentionPeriodsDays=30
UnregulatedRetentionPeriodsDays=30
RegulatedRetentionPeriodDays=365
MinUnregulatedRetentionPeriodDays=0
MinRegulatedRetentionPeriodDays=365
```

This configuration means that:

- Users can delete messages from their unregulated repository immediately, using their Outlook client.
- From the minimum retention period (`MinUnregulatedRetentionPeriodDays`) to the maximum retention period (`UnregulatedRetentionPeriodsDays`), manual deletion is possible in unregulated repositories.
- Users cannot delete messages from their regulated repository (if one exists) since, for regulated repositories, the minimum retention period is the same as the maximum retention period.
- After the retention period has expired for all repositories associated with a message, the message is available for automatic deletion during the next retention job.

When the retention period expires

When the retention feature is enabled, a retention job runs every night at midnight. The Smartcell retention manager first checks the domain retention period for the minimum number of days and verifies that retention processing is enabled. If retention is enabled, it retrieves the information for a repository, including the repository retention period, and locates documents that were ingested or sent before the age limit specified for the repository. The reference to the expired repository, and any corresponding folder references, are removed from the document. If a document is also located in other repositories and is still within their retention period, it is re-indexed.

Documents that no longer contain repository references are physically deleted from the Smartcell and removed from the index.

Editing domain retention periods

The Setting and Configuration area of the Retention page shows the retention basis and retention period for each IAP domain. It also shows whether retention has been enabled for the domain.

You can increase the retention period and/or enable retention processing for a domain by following these steps:

1. In the Retention page, click **Edit** under Set Domain Retention.

2. Set a new retention period by selecting a preset time period from the drop-down list, or by entering a specific time period in the text box.

   The time period must be greater than the Minimum Domain Retention Period.

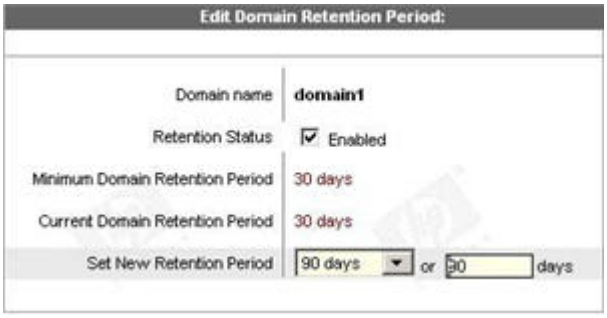

Figure 26 Editing the domain retention period

3. If retention processing should be enabled for the domain, select the **Retention Status** check box.

   (Clear the check box to disable retention processing.)

4. Click **Save Retention Now** to save the revised settings.

You can enable retention processing for all domains by clicking **Resume All**. Retention processing can be disabled by clicking **Suspend All**.

NOTE:

After the system starts archiving documents, retention processing can only be disabled or enabled using the PCC Retention page. It cannot be disabled by editing the retention fields in the `Domain.jcml` file.

Viewing deletion statistics

Documents can be deleted automatically from the IAP when their retention period expires, or they can be deleted manually. The Retention page displays an Auto Delete Statistics table and a Manual Delete Statistics table with information collected during the most recent retention job run.

These tables, located at the bottom of the Retention page, show aggregate data for domains and Smartcells.

Automatic deletion statistics

The Auto Delete Statistics table shows information for the current (or last) retention job for each Smartcell. You can view the history of retention jobs for each Smartcell by clicking the Smartcell link.

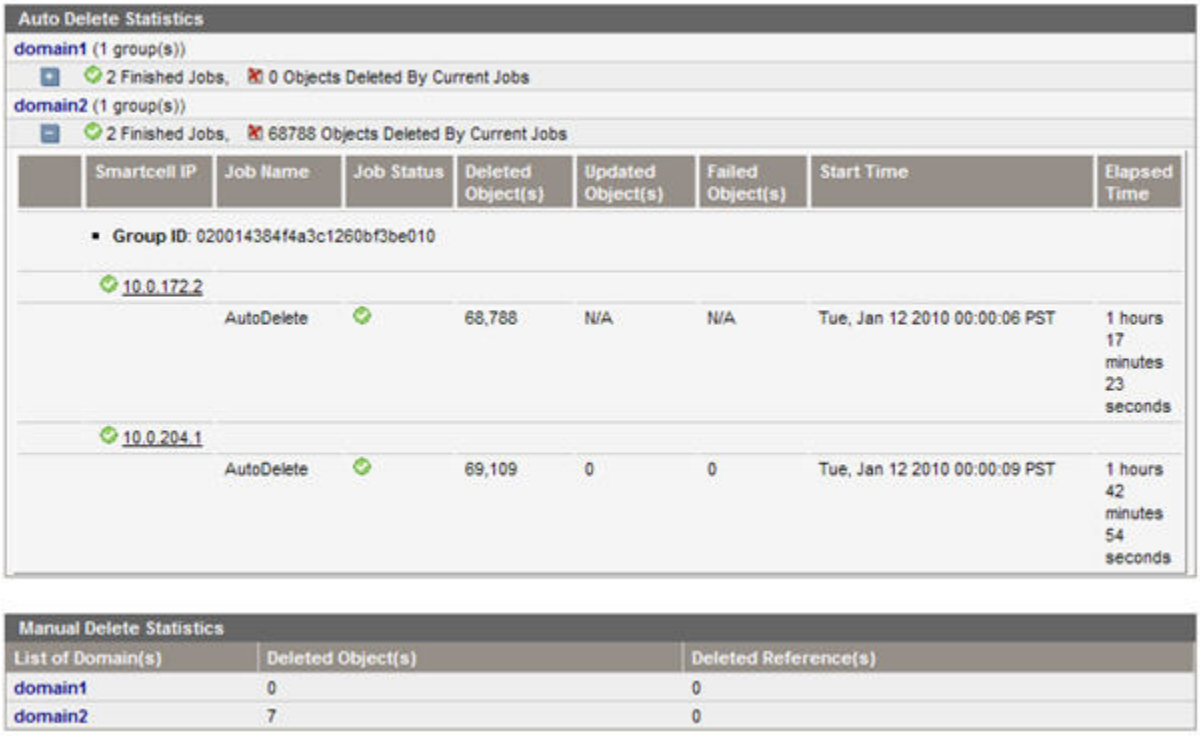

Figure 27 Auto Delete Statistics

Click the "+" icon below the domain name to see the detailed activity statistics by Smartcell. When a job is running, the statistics are updated every five minutes.

Table 24 Auto Delete Statistics

Field	Description
Smartcell IP	The IP address of each Smartcell in the domain. Move the cursor (mouse) over the icon in front of the Smartcell IP address to view the Smartcell server name, the type of Smartcell (primary or secondary), the Smartcell state and status, active thread number, and Mac address. You can click the Smartcell IP address to see the Platform Auto Delete History table, which displays the retention history for that Smartcell.
Job Name	Always **AutoDelete**.
Job Status	A retention job is either: • Finished • Failed • >>>> Running Move the cursor over the icon in the Job Status column to find the progress (percentage complete) of the job. When a job is running, the status is updated every five minutes.
Deleted Objects	The number of documents that were deleted during the run.

Field	Description
Updated Objects	The number of documents for which a repository reference was removed, because the repository retention period had expired. Documents with no remaining repository references are marked for deletion. Documents that still contain repository references are re-indexed.
Failed Objects	The number of documents that failed to be deleted plus the number of documents that failed to be updated.
Start Time	Start date and time of the job run. Retention jobs run automatically every night at midnight (12:00 a.m.).
Elapsed Time	The amount of time the job took to complete.

Manual deletion statistics

The Manual Delete Statistics table shows domain level statistics for the End User Delete function, if this function is enabled in `Domain.jcml`. (See "End User Delete" on page 94.) End User Delete occurs when users manually delete message tombstones (pointers to archived messages) from their Outlook folders. If the minimum retention period for End User Delete has expired, the reference to the user's repository is removed from the message. When folder capture is enabled, the reference to the Outlook folder where the tombstone was located is also removed from the message. The number of removed repository references are displayed in the Deleted References column.

Messages can be deleted from the IAP only when all references have been removed from the message.

Editing repository retention periods

Follow the instructions below to modify the retention period for an unregulated or regulated repository. You can change the retention period as long as it is above the default retention period for the repository type.

You can also edit the retention periods for the AuditLog, CatchAll, and Recycle Bin repositories in a domain. Documents in a quarantine repository are removed from retention processing, so there is no retention period for this repository type. (For more information about IAP repositories, see "Repository overview" on page 73.)

1. In the Configuration and Setting area of the Retention page, click the name of the domain.
2. Click the repository tab (Regulated, Unregulated, Other Type) to locate the repository.

   The AuditLog, CatchAll, and Recycle Bin repositories are located in Other Type.

   You can search for a user repository by entering the user name in the Edit User Repositories box and clicking **Search**.

   The search function uses the "Like" SQL database capability. For example, you could enter *jack* to match users jackdoe or jacksmith. Entering *%doe* would match users jackdoe, janedoe, or maryjanedoe. Entering *%ja%* would match users jadams, jackdoe, janedoe, jacksmith, or maryjanedoe.

   Searches are not case sensitive.
3. Click the repository name to edit its retention period.

4. In the form that appears, click the drop-down list in Set New Retention Period or enter the number of days in the text box.

   Retention periods are shown in days. The number of days must be the same or greater than the domain retention period. The system ignores any value that is less than the domain retention period.

5. Click **Save Retention Now**.

NOTE:

You can view user repositories on replica domains, but you cannot edit them. Retention periods and other fields cannot be edited and action buttons are unavailable when you select a domain that is a replica of another domain.

Backup

The PCC Backup pages provide status information about the standard and optional backup services running on the platform and provide access to HP Data Protector, which runs on optional backup server(s).

NOTE:

If none of the IAP domains is configured with tape backup (`DataBackupEnabled=false` in `Domain.jcml`), tape backup is disabled in the IAP. If tape backup is disabled, tape backup status will not be included on the backup pages.

IAP backup data:

- Configuration files such as IAP configuration from the kickstart, SSL certificates, and SSO keys are backed up locally (the standard backup). IAP configuration files from the kickstart server are backed up locally to the PCC when the converter script is run. Configuration files can also be backed up to tape (the optional backup system) by full and incremental backups.
- Database data is backed up locally. A local backup of the DB2 database is made daily on the platform's local disk. A copy of that local backup is kept on the database server and the kickstart server for added security. Database data files can also be backed up to tape (optional backup system) by full and incremental backups.
- Smartcell data: This data can be backed up to tape (optional backup system) by full and incremental backups. The content indexes are also backed up to tape.

To open the Backup pages, go to **Data Management** > **Backup** in the left menu.

For detailed information about the optional backup system which runs Data Protector, see "Optional Backup System" on page 129.

Backup Overview

The Backup Overview page provides backup status information for configuration, database, and Smartcell data. From this page you can also disable or enable tape backup to run as scheduled if you have the optional backup server.

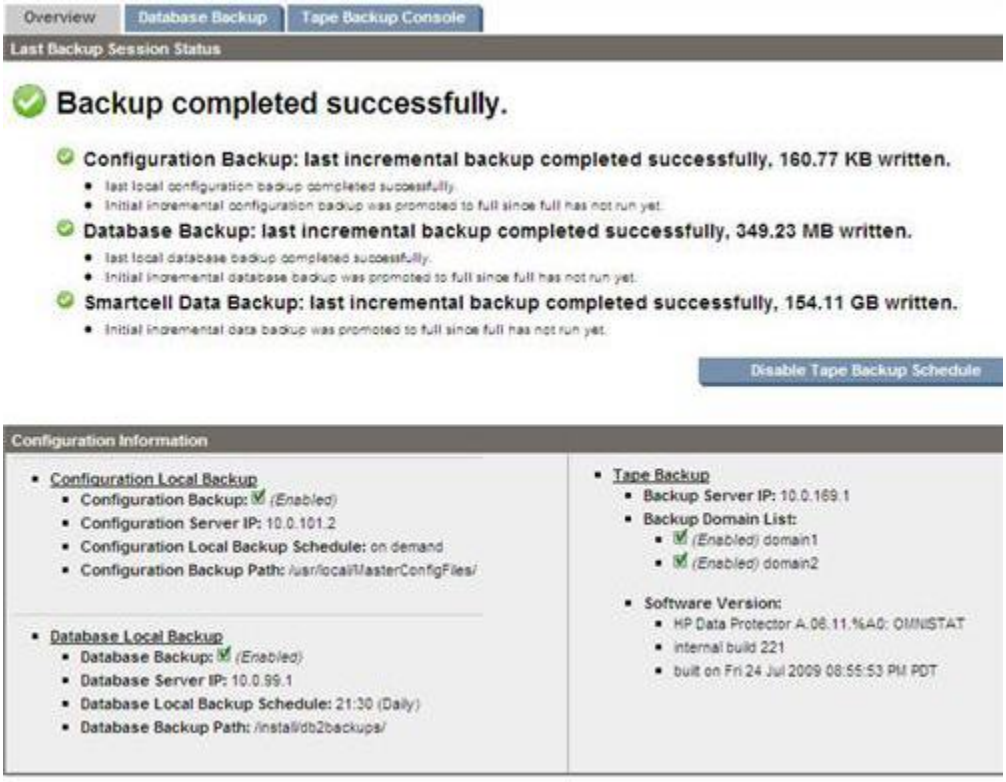

Figure 28 Backup Overview

Overall backup service summary

An overall backup service summary is at the top of the page.

A check icon indicates backup service completed successfully. This status is displayed if the last backup of all three sets of data completed successfully without error.

An ! icon indicates backup service completed with warnings. This status is displayed if the last backup of all three sets of data completed but at least one backup completed with warnings.

An X icon indicates backup service failed. This status is displayed if the last backup of any of the three sets of data failed.

An information icon indicates backup service information is being provided. This status is displayed in non-error conditions. When backup service is initially set up, the information icon signifies that no backup session has run yet.

Individual backup service summaries

The Backup Overview page also shows the backup status of each set of data: for configuration, database, and Smartcells.

The success check icon is displayed if there is at least one full backup in the last 28 days and one incremental backup in a 24 hour period. These time intervals are default but they can be customized. Otherwise, the page shows a warning icon or an error icon . Warnings are displayed when the last full or incremental backup completed with warnings. Errors are shown when the last backup failed or completed with errors, or the last full backup is older than 28 days and the last incremental backup is older than a day.

An information icon indicates backup service information is being provided for the backup type. This status is displayed in non-error conditions. (When backup service is initially set up, the information icon signifies that no backup session has run yet.)

Conditional information will inform you when:

- The given backup type has a backup schedule enabled or disabled.
- There is a current backup session running and its status (if any).
- A full backup is older than a certain threshold (by default 28 days).
- An incremental backup is older than a certain threshold (by default one day).

For the configuration and database data types (which have local backups performed on the IAP) you will see bulleted information regarding the last backup session status.

Configuration data backup status

Configuration data to be backed up resides on various servers (for example, PCC, HTTP, and kickstart servers). These servers are called the backup clients.

A check icon is displayed if the last backups of the configuration data on all backup clients completed successfully without error and are not older than a certain number of days (one day for incremental backup, 28 days for full backup).

An ! icon is displayed if the last backups of the configuration data on all backup clients completed, but incremental backup on some backup clients is older than two days (for example, “Last completed incremental backup on kickstart is older than 1 day”).

An X icon is displayed if the last FULL backup of the configuration data on any client is older than 28 days (for example, “Last completed full backup on PCC is older than 28 days”), or if the backup status on any backup client cannot be found (for example, “No backup found on backup server”).

Database data backup status

Database data is backed up on the IAP local disk (kickstart and database servers) and also backed up to tape if tape backup is enabled. The database data that is backed up to tape resides on the kickstart server.

A check icon is displayed if the last local backup and tape backup of the database data completed successfully without error, and tape backups are not older than a certain number of days (one day for incremental backup, 28 days for full backup).

An ! icon is displayed if the last backup of database data completed but the local database backup has a warning (for example, “Local backup 1 completed with warnings”), or an incremental tape backup is older than one day (for example, “Last completed incremental backup on kickstart is older than 1 day”).

An X icon is displayed if the last full tape backup is older than 28 days (for example, “Last completed full backup on kickstart is older than 28 days”), if the tape backup status cannot be found (for example, “No backup found on tape”), or if the local backup status cannot be found (for example, “No backup found on DB2 server”).

If tape backup is disabled, the status of the database backup to tape is not displayed on the dashboard.

Smartcell data backup status

Smartcell data to be backed up resides on assigned Smartcells. These Smartcells are also backup clients. The name of the Smartcell backup client consists of *gid* and the group ID that the Smartcell belongs to and the domain name of the IAP.

A check icon is displayed if the last backup of the data on all assigned Smartcells completed successfully without error and is not older than a certain number of days (one days for incremental backup, 28 days for full backup).

An ! icon is displayed if the last backup of the data on all assigned Smartcells completed but incremental backup on some Smartcells is older than one day (for example, “Last completed incremental backup on gidxxx1 is older than 1 day”).

An X icon is displayed if the last full backup of the data on any assigned Smartcell is older than 28 days (for example, “Last completed full backup on PCC is older than 28 days”), or if the backup status on any assigned Smartcell cannot be found (for example, “No backup found on gidxxx1.”).

If tape backup is disabled, Smartcell data backup is not displayed on the dashboard.

Disable/enable tape backup schedule

If the tape backup feature is configured on the Backup Overview page, a button to disable/enable scheduled tape backup is displayed. If the tape backup schedule is disabled, tape backup will no longer occur. The data on previously backed-up tape will not be erased. When the tape backup schedule is disabled, a bullet under each of the three backup types is displayed stating that the backup is disabled.

Configuration Information

The lower portion of the Backup Overview page displays configuration information for the local and tape backups.

The Configuration Information area contains the internal IAP IP address of the configuration and database server, the run time schedule for the local backup process, and the location of the local backup files.

The Configuration Information area also contains the internal IAP IP address of the backup server, the version of the Data Protector backup software being used, and a list of the IAP domains that have tape backup enabled.

NOTE:

If tape backup is disabled, this section displays “Tape Backup is disabled.” If tape backup is enabled but the backup server is not available, the Backup Server IP section displays “Backup application Server is down, cannot retrieve information.”

Database Backup

Database data is backed up on the IAP local disk (kickstart and database servers) and also backed up to tape if tape backup is enabled. The Database Backup page provides detailed information on the last two local backups (one on the local database server and the other on the kickstart server). The Database Backup page also provides detailed information on the last tape full and incremental database backups (for optional backup servers). It includes completion time, size, type (full or incremental) information, and informational messages.

The entry in the Last Database Backup table for the local database backup contains information about the database backup job, which extracts information from the DB2 database and stores it in a recoverable format.

The PCC displays statistics about the database backup process itself. In particular, the Completion Time column shows the time at which the database extraction process completed. Other job logs

might show the local backup job completing at a later time, but those later timestamps include the overhead tasks that are managing the disk storage space for the database backup archive, and are not the specific database backup process.

The Local Completion Time is useful for determining the data that was included in the backup, for situations where recovery of backed-up data is dependent on the time the data was extracted.

Tape Backup Console

The Tape Backup Console page provides VNC access to the HP Data Protector graphical user interface (GUI), which runs on the optional backup server(s). Use the HP Data Protector GUI for detailed optional backup server status information and configuration.

Local backup file locations

The contents of the `install/config/primary` directory are backed up.

- Database backup files: The database backup files and DB2 transaction log files are stored on the database server (partition `/db2/VolBack`). The files are mirrored to a directory on the kickstart machine (`/install/db2backups`). The DB2 database is backed up once a day and the last two backup files are kept on disk.
- Master configuration files: There are three master configuration files: `BlackBoxConfig.jcml`, `Domain.jcml`, and the `UserKeyStore` file in the `/install/configs/primary` directory on the kickstart server. Whenever one of these files is changed and the converter and Registry Loader are run, the modified files are copied to the PCC machine, to the directory `/usr/local/MasterConfigFiles`.

  (To run the RegistryLoader after changing a configuration file, run the following script on the kickstart server: `regloader.pl -cv`.)

NOTE:

The file `BlackBoxConfig.bct` is not included in the local backup, because it can be generated from `BlackBoxConfig.jcml` (`/install/tools/converter/dumpBCT`). The script `createBBC.zr` generates the `.jcml` version that is backed up.

Restoring DB2 local backup files

The DB2 database can be restored from the backup files on the `/db2/VolBack` partition on the database server. When data is lost because the database machine was reinstalled, restore the backup files first. Restore the backup files from either the copy on the kickstart server, or tape backups if application backup is enabled.

The functionality to restore the database backup partition and the database itself is implemented in the `/opt/bin/backup/db2BackupUtility` script on the database server. Obtain detailed usage information by running the command without any arguments.

Follow these steps to restore the database:

1. Before restoring the database, stop all IAP applications by issuing `/opt/bin/stop` on the PCC server.
2. Log on to the database system.

3. If the database server was reinstalled, run the following command to restore the `/db2/VolBack` partition from the version stored on the kickstart server:

   ```
   /opt/bin/backup/db2BackupUtility -restore_from_kickstart
   ```

NOTE:

This will delete the content of the VolBack directory.

4. To restore the database, make sure the database is up by issuing the command:

   ```
   su – db2udb –c db2start
   ```

5. To restore DB2, run the following command. You will be asked for a timestamp of the backup to restore.

   ```
   /opt/bin/backup/db2BackupUtility -restore_db2
   ```

   Normally, the last two backups of the database are kept.

6. After the installation is finished, run `/opt/bin/start` on the PCC server to start the IAP and validate that it operates correctly.

Restoring the master configuration files

Copies of the master configuration files are kept in the `/usr/local/MasterConfigFiles` directory on the PCC server. Whenever one of the files is changed on the kickstart server and the converter is run, the master configuration file is copied to the PCC server, and renamed to `YYYY-MM-DD-hh.mm.ss_FileName`.

Restoring a Smartcell group

The Restore Smartcell Group page lists broken groups that can be restored from tape, groups that are currently being restored (including progress information), as well as past completed or aborted restore operations.

If you lose one Smartcell of a group, use the Smartcell Cloning page to clone the Smartcell to rebuild the group. Cloning is much faster than tape restore and the healthy Smartcell may have recent data not on the tape backup.

If you experience the loss of both Smartcells in a group, use the Restore Smartcell Group page to restore one Smartcell from tape. After one Smartcell is restored from tape, use cloning to get the second Smartcell thus rebuilding the group.

NOTE:

You should only restore from tape when cloning is not possible. Cloning a Smartcell is faster than a tape restore and also guarantees that the data is complete (data on tape is only as current as the time the last backup took place).

To open the Restore Smartcell Group page, go to **Data Management** > **Restore Smartcell Group** in the left menu.

Smartcell Group Restore candidates

The Smartcell Group Restore candidates section appears on the Restore Smartcell Group page when there are one or more broken groups that can be restored from tape. Only groups for which both the primary and secondary Smartcells are unavailable are listed here. If only one of the Smartcells in a group has to be restored, cloning is the preferred option to recover the group.

Each entry in the group lists the domain name, the group ID, and (if available) the IP address of primary and secondary Smartcell.

The **Start Restore** button initiates the tape restore for the given group. When you click the button, tape restore for the particular group begins and an entry for the operation appears in the list of current restore jobs.

Current Smartcell Group Restore Jobs

The Current Smartcell Group Restore Jobs section lists running or failed restore jobs along with their current progress. Each entry provides the following information:

Table 25 Current Smartcell Group Restore Jobs information

Field	Description
Group ID	The group ID of the group being restored from tape.
Status	Either Running or Failed.
Smartcell IP	IP address of the Smartcell to which the data is being restored. This field is populated as soon as a free Smartcell has been assigned for restore.
Started	Date and time when this operation began.
Completed Step	Information about the restore step currently being executed.
Action	One or more possible actions to control the job.

You can perform the following actions:

- **Reset**

  Can be performed as soon as the Smartcell has been assigned. This action aborts the current job and resets the Smartcell that was assigned as a restore target. All data that was potentially restored is already removed and the cell is moved back into the pool of free Smartcells.

- **Retry**

  Only available if the current step fails. This action resumes the operation, retrying from the last failed step.

- **Clone**

  Only available once the Smartcell is restored completely and is ready to be cloned. This action takes you to the Cloning PCC page and marks the restore operation for this Smartcell as complete.

Unregistering a configured Smartcell

If a Smartcell that is already configured in an IAP needs to be “rekickstarted” for some reason, you must unregister the Smartcell in the Tape Backup Console to be able to use it for restore. (If this is a new Smartcell added to IAP, this step is not necessary.)

To unregister a Smartcell:

1. Open the Tape Backup Console using the tab on the Backup page in the IAP PCC UI.

2. In the Tape Backup Console, connect to the HP Data Protector Cell Manager, by clicking on **File > Connect to Cell Manager**. Select the primary backup server to connect. If you are already connected you can skip the connection step.
3. Once connected, select the Smartcell hostname from the "Clients" drop down section. Right click and select **Delete**. In the figure below, the Smartcell being selected for delete is sc-s1-172-1.persistcorp.com.

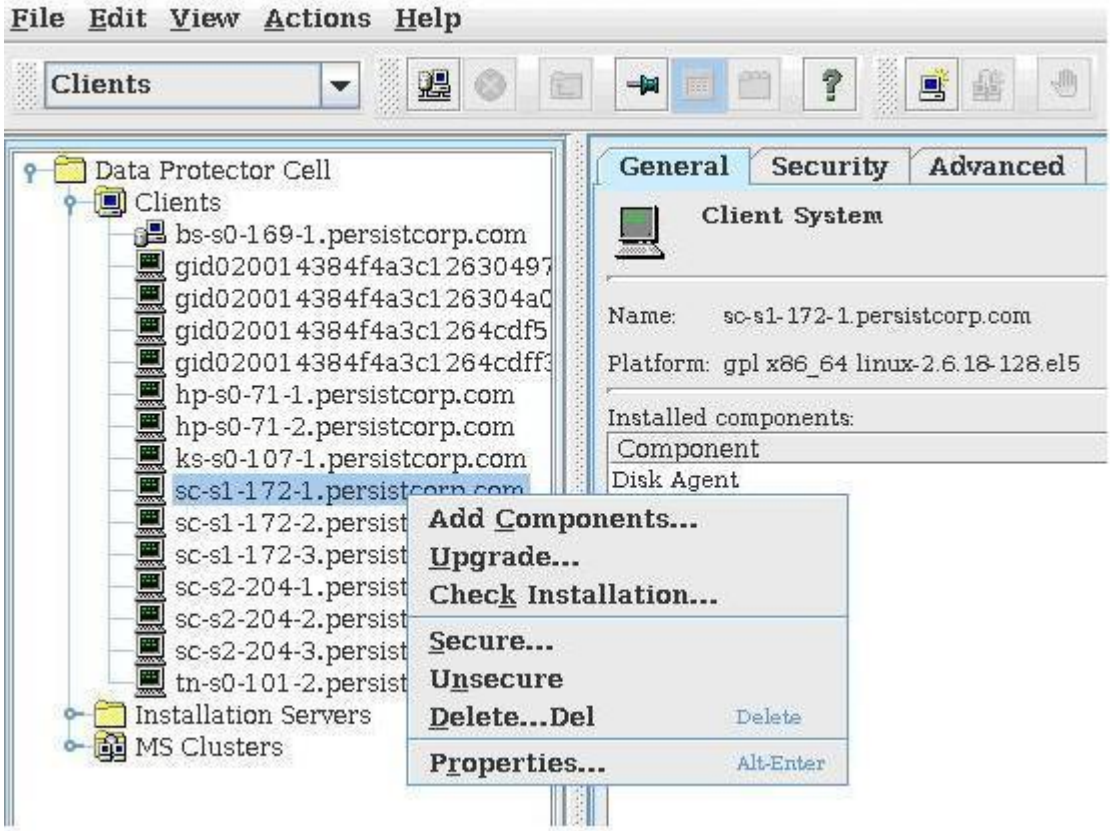

Figure 29 Unregistering a Smartcell

4. When asked if you want to remove Data Protector software, click **No**. Once this step finishes, the Smartcell is ready to be rekicked and available for restore.

Smartcell Cloning

Use the Smartcell Cloning page to show the status of current and past cloning operations, and to clone a Smartcell.

To open the page, go to **Data Management** > **Smartcell Cloning** in the left menu.

Cloning overview

You can clone a Smartcell if its mirror Smartcell is SUSPENDED, DEAD, or FAILED. (See "Smartcell life cycle states" on page 36.) Cloning a Smartcell copies all its information to another Smartcell that is in the FREE state to give the Smartcell a new, viable mirror. Cloning operations can take a long time (as much as a day), depending on the amount of information cloned.

To place a source Smartcell in the free pool, contact your HP authorized support representative.

When you access the Cloning page, PCC searches for ongoing cloning operations and loads the current data. Only one Smartcell can be cloned at a time, so you see the progress of any ongoing cloning operation.

Cloning features

Table 26 Cloning features

Feature	Description
Source	The IP address of a Smartcell without a viable mirror. If all Smartcells have viable mirrors, the Source displays **No Broken Groups Found**. If more than one Smartcell needs a mirror, a Change Source button appears below the automatically selected IP address.
Free Cells	The number of Smartcells currently in the FREE state. This number is decreased by one after a cloning operation starts.
Assigned/Free	A graphical representation of the assigned and free Smartcells.
Clone Cell	Starts the cloning operation. See "Cloning Smartcells (copying data)" on page 108.
Cloning Area	This area provides the following information about an ongoing cloning operation: • Source selected: IP address of the Smartcell being duplicated. • Target selected: IP address of the Smartcell receiving duplicate data. • Current Step Percentage: Dynamic bar showing how much source data has been duplicated. • Overall Percentage: Current step in the cloning operation. Cloning operation steps are as follows: • Initializing • Assigning target host • Transferring data • Transferring indexes • Waiting for indexer to complete • Updating history log • Completed or Failed Some steps happen so quickly you might not see them. Steps such as Transferring data and Transferring indexes can take a long time if there is a great deal of data to clone.
History Logs	The following information about each cloning operation that has occurred since startup: • Source • Target • Time Elapsed • Status • Date

Cloning Smartcells (copying data)

To clone a Smartcell:

1. Perform one of the following actions:
   - Select the Smartcell from the Source field.
   - If the option is present, click **Change Source** to select a different Smartcell for cloning. When the selection box appears, select the Smartcell you want from the drop-down list, and click **Select**.
2. Click **Clone Cell**.

   This button is unavailable if there are no Smartcells to clone.

   When cloning is successful a pop-up panel appears on the PCC.
3. Check the Status Area to see results of the cloning operation.

   A check icon indicates the cloning operation is proceeding normally.

   An X icon appears if the cloning operation has failed.

Folder capture support

IAP 2.x supports folder capture for Email Archiving software for Exchange (HP EAs Exchange) version 2.x.

This feature captures the original Outlook location of archived messages and updates the corresponding documents in the IAP. The folder location (for example, `/Inbox/project`) is stored as metadata in the archived email.

For folder capture to occur, it must be enabled on both the IAP and HP EAs Exchange. In HP EAs Exchange, folder capture is enabled by default. On the IAP, it is disabled by default and must be enabled in `Domain.jcml`.

When folder information is indexed, the folder name can be used to search for and retrieve messages stored on the IAP. To limit the number of archived messages that are re-indexed when folder information is added, a cut off date is specified in `Domain.jcml`. Messages stored before the cut off date have only their metadata updated, while messages stored after the cut off date have both the metadata and indexes updated with folder information.

NOTE:

If the IAP contains a number of closed Smartcells, available disk space can be quickly depleted when indexed folder information is added.

To prevent this error, we strongly recommend setting the cut off date to the date the system is upgraded so that only new messages are indexed.

For more detail on folder support, see the "Working with folder capture" chapter in *HP Email Archiving for Exchange Administration Guide version 2.x*.

Enabling folder capture

The FolderSupportEnabled option in `Domain.jcml` indicates whether folder support is enabled in the IAP for a domain. If enabled, the IAP allows users from this domain to search for email by folder and display the folder in which the email is stored. To enable folder support, this option must be set to true: `FolderSupportEnabled=true`

After configuring `Domain.jcml`, run the following commands: `regloader.pl -cv -clearallConfirm=<IAP name>` on the kickstart server and `/opt/bin/restart` on PCC to restart the IAP.

The PCC Domain Configuration shows whether folder support is enabled.

Use the date of the IAP 2.x installation or upgrade as the date, for example 07/30/2008.

Administrative Delete

The Administrative Delete function gives designated users the ability to manually delete an email that has been stored in the archive. This function is typically used in the Federal area for dealing with classified data; for example, to remove inadvertently distributed messages from the archive. Email can be manually deleted from regulated, unregulated, and quarantine repositories.

IMPORTANT:

It is the responsibility of customers to verify that using Administrative Delete does not interfere with their organization's record retention policies. HP will not be held liable for the irresponsible use of this feature. HP has made every effort to put the proper controls and logging in place to prevent unintentional removal of data from the system.

Enabling Administrative Delete

Administrative Delete is enabled or disabled at the domain level. If enabled, the system allows a user who has the appropriate privilege to delete email.

1. Open the `Domain.jcml` file:
   a. On the console, press **Ctrl** twice. In the menu that appears, double-click **01** (on IAP 2.1 hardware) to select the kickstart server.
   b. At the prompt enter:
      ```
      cd /install/configs/primary
      ```
   c. Open `Domain.jcml` in a text editor.
      For example:
      ```
      vi Domain.jcml
      ```
2. Enable Administrative Delete by setting the `AdminDeleteEnabled` field to true:
   ```
   AdminDeleteEnabled=true
   ```
   Add this field if it does not exist in the file.

   NOTE:

   If Administrative Delete is enabled, the AuditLog must also be enabled, as described in "Enabling the AuditLog" on page 127.

   (Administrative Delete can be disabled by setting the `AdminDeleteEnabled` field to `false`.)
3. Save and close the file.

4. Run the RegistryLoader by entering:

   `regloader.pl –cv –clearallConfirm=<BlackBoxName>`

   Enter y at prompt to confirm.

5. After the RegistryLoader has completed, go to the PCC. On the console, press **Ctrl** twice. In the menu that appears, double-click **02** (on IAP 2.1 hardware) to select the PCC NAT.

6. Restart the IAP from the PCC:

   `# /opt/bin/restart`

7. Log into PCC Web administration.

   The General Configuration Platform Settings page should show the Admin Delete Service as enabled.

Granting the Delete Admin privilege

All IAP Admin users, with the exception of super users, can grant or revoke the Delete Administration privilege. Only a remote user (a user existing in the Active Directory and imported by DAS) can be granted Delete Administration privilege.

To grant this privilege to a user:

1. Log in to PCC as the super user.
2. Open the PCC Account Management page and grant the IAP Admin privilege to a user.
3. Log in to PCC as this IAP Admin user and open the PCC Account Management page.
4. Select the user who needs to execute Administrative Delete and edit this user by selecting the **Delete Admin** check box.

   The domain that this user belongs to should have Administrative Delete enabled; otherwise, the Delete Admin check box will not appear.

5. Save the change.

To revoke the Admin Delete privilege:

1. Log in to PCC as the remote IAP Admin user.
2. Open the PCC Account Management page and select the user who has the Delete Admin privilege and edit this user by deselecting the **Delete Admin** check box.
3. Save the change.

Executing Administrative Delete

After obtaining the Delete Admin privilege, a user is allowed to delete any messages stored in IAP. To delete messages:

1. Log in to the IAP Web Interface as the user with the Delete Admin privilege.
2. Search for and select the message to be deleted.
3. Click **More Options**.
4. Select **Delete Checked Items**.
5. To confirm the deletion, click **Confirm Delete**.

6. After confirming deletion, a status page appears, which contains a report on the success or failure of the deletion submittal.

   After the message is submitted to delete, it may take up to two hours for the message to be removed from the index. During that time, the message contents will be still searchable, but not retrievable.

When a message is deleted by Administrative Delete:

- The deletion physically removes the message and all references of the message from IAP.
- The message is deleted if it is quarantined.
- A message in Saved Results is deleted from the index and the message body is no longer searchable. However, the message is still listed in Saved Results.
- The deletion is done on both primary and secondary Smartcells. If the IAP is running in a replication environment, the message is also deleted on the replica Smartcells.
- If the IAP is running in a replication environment, Administrative Delete can only be executed to delete a message that is stored on the primary IAP; the deletion is triggered on the replica IAP by the replication process. Administrative Delete can not delete any message stored on the replica IAP directly.

Logging in AuditLog

Activities in Administrative Delete are logged in the AuditLog and stored in the AuditLog repository.

- When a user obtains or loses the Delete Admin privilege, the change is logged in the AuditLog.
- When a message is deleted by Administrative Delete, the operation is logged in AuditLog.

Current limitations

The following are current limitations of Administrative Delete:

- In IAP 2.x, Administrative Delete only deletes emails. It does not delete documents that were stored with the Object Store API (BIBO) or the HP File Archiving software (formerly FMA), and it does not delete an audit trail.
- If email has been backed up with the IAP backup feature, the email can be deleted from the IAP using Administrative Delete. However, the email and any related information on the tape will not be deleted. Because of this, all administrative delete operations should be re-executed on the IAP when any Smartcell recovery has been performed. (This is a low-probability event.) The AuditLog provides the necessary information to determine which delete operations need to be re-executed.
- At this time, deleted messages in a saved query result appear when reloading the saved result in the Web Interface. However, if you click on the deleted messages, you will see `Error displaying message`. The saved queries keep the reference to the file until the saved query is cleaned up.
- This feature does not currently satisfy DoD or NSA data destruction standards.

9 Reporting

This chapter includes information about the following topics:

Event Viewer

The Event Viewer shows the events that are occurring or have occurred on system servers and in system services and applications.

NOTE:

All alerts are logged as events in the PCC, and are displayed among the other events on the Event Viewer page. See "Current Platform Alerts" on page 39.

There are two ways to open the Event Viewer:

- Go to **Reporting** > **Event Viewer** in the left menu.
- Select **Overview** from the left menu, and then click **Go to Event Viewer** in Current Platform Alerts.

Event Viewer overview

The following information is displayed in the Event Viewer.

Table 27 Event Viewer features

Feature	Description
Description	A short description of the event. The description is preceded by an icon showing the event level: • *Critical*: The issue is critical for the entire system. • *Major*: A server or application feature has a major problem, but the issue is not critical for the entire system. • *Minor*: A server or application feature has some internal problems, but they are not critical for the server or feature health. • *Resolved/Manually resolved*: The issue has been automatically or manually resolved. • *Unknown*: The cause of the issue was unknown. • *Information*: The event was part of the normal operation of the system.
Context	The context in which the event occurred.

Feature	Description
Machine	The type of server on which the event occurred and the IP address of the machine.
Date	The date and time that the event first occurred or was resolved.

Searching the Event Viewer

You can choose to display only events that match the search type and criteria specified in the Search area of the viewer.

To search the Event Viewer:

1. Select the number of events to be shown on each page.

   The default is 20 events per page.

2. Select a Search Type in the drop-down list.

   You can choose from the following types of searches:

   - Show All Events
   - Level (of event)
   - Machine Type (of the machine on which the event occurred)
   - IP Address (of the machine on which the event occurred)
   - Host Name (of the machine on which the event occurred)

3. In the Search Criteria text box, enter the criteria for the search.

   You must enter criteria for all searches except Show All Events. Searches are not case sensitive.

   - For Level, see the legend at the bottom of the Event Viewer.
   - For Machine Type, see the Machine column in the Event Viewer for examples.
   - For IP Address, see the Machine column of the Event Viewer or the Software Management page.
   - For Host Name, see the Software Management page.

   The search function uses the "Like" SQL database capability. For example, you could use Search Type: Host Name and enter *sc%* in Search Criteria to match host names sc-s1-172-1.company.com or sc-s2-204-1.company.com. Entering *%204%* would match hosts sc-s2-204-1.company.com, or sc-s2-204-2.company.com.

   Or, to search for resolved events, you could select Search Type: Level and enter *%resolved* in Search Criteria to view both resolved and manually resolved events.

4. Click **Submit**.

Deleting events

Delete events from the Event Viewer in one of the following ways:

1. Choose one of the following:

   - Select the check box in front of one or more events.
   - Select a cutoff time in the drop-down list at the bottom of the page. For example, events older than two weeks.
   - Select All Events in the drop-down list at the bottom of the page to remove all events in the viewer.

2. Click **Delete** after making your selection.

NOTE:

The system automatically clears events (oldest first) when the number of records reaches 20,000. Ten percent (2,000) of the maximum number of entries are removed. Currently, event retention is not configurable.

SNMP Management

Application-generated alerts that occur in the IAP system are shown in Current Platform Alerts on the PCC Overview page. These alerts can be forwarded as SNMP traps to external monitors in your network and/or to email recipients.

Use the SNMP Management page to perform the following functions:

- Download the IAP MIB file to be imported into your monitoring management server(s).
- Select the SNMP traps to be monitored.
- Set the monitoring servers and/or email recipients to be notified when an alert is received.
- Set SNMP community.
- Send a test trap to verify that the configuration has been set up properly.

NOTE:

Only application-generated alerts can be selected as SNMP traps on this page. Hardware monitoring for IAP servers is provided by Proliant management agents. See "Hardware monitoring" on page 47 for more information.

To open the page, go to **Reporting** > **SNMP Management** in the left menu.

Downloading the IAP MIB

The IAP MIB (management information base) allows monitoring applications such as HP SIM to recognize SNMP events coming from the IAP.

IAP SNMP traps are placed in an SNMP XML file, which is used by the IAP system to generate a MIB file, `iap.mib`. The generated `iap.mib` uses SNMP version 2, which is compliant with the transfer protocol SNMP version 1.

Whenever a new version of the IAP software is installed on your system, download `\opt\mibs\iap.mib` from the PCC server and import it into your monitoring management software so the monitor can recognize the IAP SNMP traps. In the MIB Manager Service area, click **Download** to copy the IAP MIB file to your local machine.

To import the MIB to the monitoring servers, refer to the documentation for your monitoring management software.

Selecting SNMP traps

In the SNMP Trap Manager Service area, select the SNMP traps to be monitored. When one of the alerts is generated, a corresponding trap is sent to the SNMP monitoring server and/or email recipient that is entered on the SNMP Management page.

To set the SNMP traps:

1. Select the system for which the traps will be activated.

   This is always IAP.

2. Select the traps to be sent:

   Traps are grouped by features. You can activate all traps, activate traps by feature, or activate traps individually.

   The trap IDs correspond to the alert IDs that are listed in "IAP application-generated alerts" on page 155.

3. Click **Save** when you are finished.

Setting the SNMP server

In the SNMP Trap Server Monitors area, enter the IP address of the server where you wish to send SNMP Traps, then click **Add**.

You can enter more than one IP address, since the SNMP information can be broadcast to multiple servers, or receivers.

NOTE:

Always enter the IP address of the server; do not enter the hostname.

HP SIM servers can be used as SNMP receivers. To be compliant with SIM servers, the system uses SNMP protocol version 1. All SNMP receivers should be compatible with this version.

Use of a server can be enabled or disabled by clicking the radio button in front of the server entry, then clicking **Toggle Enabled/Disabled**.

A server can be deleted from the list by clicking the radio button in front of the server entry, then clicking **Delete**.

If you do not have an SNMP trap monitoring server, you can receive SNMP event notifications via email. See Sending SNMP events by email.

Sending SNMP events by email

If you do not have an SNMP trap monitoring server, you can receive SNMP notifications via email. This option is also useful if you are away from the office.

To enable the service, enter your email address or the email address of the person who will receive the notifications, and click **Add**.

Repeat this step if you want to add more recipients.

To enable/disable an email recipient, click the radio button in front of the person's name and click **Toggle Enabled/Disabled**.

To delete an email recipient, click the radio button in front of the person's name, and then click **Delete**.

IMPORTANT:

The mail host's IP address must be entered in the Email Reporter page for SNMP email notifications to be sent. See "Creating and scheduling email reports" on page 117 for more information.

Setting SNMP Community

An SNMP community string is a text string that acts as a password. It is used to authenticate messages that are sent between the management station (the SNMP manager) and the device (the SNMP agent). The community string is included in every packet that is transmitted between the SNMP manager and the SNMP agent.

Public is the default choice for the SNMP Community.

Testing the configuration

After configuring or updating the settings, you can click **Send Test Trap** to verify that the configuration has been set up properly.

Email Reporter

Use the Email Reporter to configure summary monitoring reports that are sent periodically to chosen email recipients.

You can select the information to be sent and the frequency of the reports— from one time only to every seven days. These reports can include system status, storage rates, node health, and other information from the PCC pages.

To open the page, go to **Reporting** > **Email Reporter** in the left menu.

Creating and scheduling email reports

To generate and schedule email reports:

1. Configure the fields in Mail Configuration if this is the first time you are using email reports:
   a. Enter the IP address of the mail host.
   b. Enter an email address for the report's sender, for example IAPAdmin@mycompany.com, or iap-no-reply@mycompany.com. This does not have to be a real email address.
   c. Click **Save** to save the mail configuration.

2. In Generate Report With, select one of the following:
   - **All Contents** to send all the information listed in Custom Contents.

   NOTE:

   Always use All Contents when you are sending email to HP technical support.

   - **Custom Contents** to select specific information to send.

Report on	Report taken from
Alerts	**Overview** > **Current Platform Alerts**
Archive Gateway (EAs Exchange only)	**Archive Gateway Management** > **Overview Archive Gateway**
Auto Delete Statistics	**Data Management** > **Retention Information** > **Auto Delete Statistics**
Backup	**Data Management** > **Backup**
DAS Statistics	**User Management** > **Account Synchronization** > **DAS Available Job(s)**
Domain Settings	**General Configuration** > **Platform Settings**
Node Health	**Software Management**
Platform Performance	**Overview** > **Platform Performance**
Smartcell Metrics	**Overview** > **Platform Statistics**
Version	**Software Version**

3. Enter one or more email addresses in the **Send to Recipient(s)** text box.

   When entering several email addresses in the box, separate them with a space, comma, or semicolon. For example: recipient1@mycompany.com,recipient2@mycompany.com.
4. In the **Frequency** list, select the frequency that the report should be sent.

   The frequency ranges from one time only to every seven days.
5. Click **Schedule**.

The Email Report information appears in the Current Active Schedules area of the page, unless the frequency is set to one time only. You can click **Send Now** anytime to immediately receive a copy of the defined report rather than wait for the scheduled time.

To delete a recipient, select the radio button in front of the recipient's email address, then click **Delete**.

Collecting log files

The Collect Log Files page lets you collect logs detailing system runtime events and errors. These logs can be gathered from specific servers or from all servers on the IAP. For example, you can collect logs from the HTTP servers and email them to HP technical support for help in troubleshooting.

The logging level can be set to INFO, to collect announcements about the normal operation of the system, or changed to DEBUG or TRACE to collect more detailed information and to help diagnose runtime errors. This change is made by HP technical support, either on the Archive Gateway (to log errors on the gateway server), or in the PCC Service Tools (to log errors on other system servers). You might be asked to change the logging level under HP support's guidance.

To open the Collect Log Files page, go to **Reporting** > **Collect Log Files** in the left menu.

Collecting the logs

Follow these instructions to collect a log file:

1. Select the type of server from which logs should be collected.

   You can select one of the following:

   - SMARTCELL servers
   - META servers
   - DB server
   - SMTP servers
   - HTTP servers
   - SMTP and HTTP servers
   - PCC server
   - ARCHIVEGATEWAY servers (Exchange only)
   - BACKUP servers
   - All Servers (except Exchange ARCHIVEGATEWAY)

2. Determine how the log will be collected by selecting one of the following methods:

   - Download the log to a .tar file.
   - Email the log.

     Enter the recipient's email address in the Email Logs field. To send to more than one email address, separate the addresses by a space, comma, or semicolon.

     The nms-mail-relay in `BlackBoxConfig.bct` on the kickstart server must be configured properly for the logs to be sent via email.

     NOTE:

     The size of the logs may be huge. Depending on the email server rules regarding the attachment sizes in the customer environment, the email might not go through. There is no check to see whether the email has been correctly sent to the recipient.

   - FTP the log.

     Enter the destination host address and the username, password, and destination directory.

     (Optional) Enter an encryption password to encrypt the logs. You will need to provide the password to HP support engineers so they can open the logs.

NOTE:

Log files are compressed into .zip format for logs from the Archive Gateway or .tgz format for logs from the other system servers. If the logs are downloaded, the .zip or .tgz files are bundled into a .tar file.

3. Click **Run Now**.

10 External access

The PCC left menu contains an Archive Gateway Management section. The Archive Gateway Management section has a link to an overview of the Archive Gateway status for each EAs for Exchange domain, and a link for VNC access to the Archive Gateway.

Archive Gateway Management

For information on the Overview Archive Gateway page, see "Overview Archive Gateway" on page 122.

To open the page, go to **Archive Gateway Management** in the left menu.

VNC Archive Gateway

To access an Archive Gateway for an EAs for Exchange domain:

1. In the left menu, go to Archive Gateway Management and select the **VNC Archive Gateway** link.

   The VNC Viewer: Connection Details dialog box appears.
2. Click **OK**.

   The VNC Authentication dialog box appears.
3. Enter the password that was provided by your HP service representative, and press **Enter**.

   The Windows Welcome screen appears.
4. Place your mouse cursor anywhere in the VNC viewer and press **F8**. In the menu that is displayed, select **Send Ctrl-Alt-Del** to display the Windows login screen.

   You can also call the Windows login screen by pressing **Shift**+**Ctrl**+**Alt**+**Delete** or **AltGr**+**Delete**.
5. Log in to VNC using the password provided by HP.

NOTE:

Important! Do not use **Ctrl**+**Alt**+**Delete**. This key sequence does not work when logging in to VNC from the PCC.

Overview Archive Gateway

This page provides information about the Archive Gateway system for each domain, including service status and health.

Table 28 Overview Archive Gateway features

Feature	Description
Header	The header shows the name of the Archive Gateway (for example, EM-S0-110-1), and the version of the Email Archiving software (for example, 1.05.0000).
Statistics Collected	The date and time that the statistics were collected (for example, 5/11/2009 2:10:17 PM).
Sections	Each section in the overview describes a major area inside the Archive Gateway, together with an overall health status for that area. The sections include information on: • System services • Configured tasks • Journal mining • Selective archiving • Synchronized deleted items • Tombstone maintenance If a disabled icon is displayed, the section is either disabled or data was not available at the time statistics were collected.

System Services

The System Services section displays the health of the specific system services running on the Archive Gateway.

Table 29 System Services features

Feature	Description
Name	The name of the service.
Status	The following status types can be displayed: • The service is running and the Service Account is not set to LocalSystem. (VNC Server is expected to have LocalSystem as its Service Account.) • The service is stopped or the Service Account is set to LocalSystem. (VNC Server is expected to have LocalSystem as its Service Account.) The Service Information summary displays the red icon if any service is stopped or the Service Account is set to LocalSystem. Otherwise, the summary displays the green icon.
Startup Type	Possible values for Startup Type are Auto, Manual, or Disabled.
Service Account	The name of the service account.

Configured Tasks

The Configured Tasks section displays configuration information about tasks that have been created and modified using the Scheduler program.

Table 30 Configured Tasks features

Feature	Description
Task Name	The name of the task.
Status	The following status types can be displayed: • The task is enabled. • The task is disabled. The Task Information summary displays the yellow icon if all configured tasks are disabled. Otherwise, the summary displays the green icon.
Task Type	The type of task, for example Selective Archiving or Journal Mining.
Process Count	Journal Mining only: Displays the number of processes that are configured to run concurrently.
Schedule Frequency	Journal Mining only: Displays the frequency of the processes that are configured to run concurrently.

Journal Mining

The Journal Mining section contains information on each Journal Mailbox being mined. It includes three areas: Exchange Statistics, Process Information, and Message Statistics. If any of these areas has minor or major error conditions, then a fourth section is displayed. The fourth area, Processes in Failed or Retry State, contains information about the particular processes that have errors.

Table 31 Journal Mining features

Feature	Description
Exchange Statistics	• Connection Status: The Exchange server is available and communicating with the email miner. A connection to the Exchange server cannot be completed. • Exchange Server: The Exchange server on which the journal mailbox resides. • Journal Mailbox Name: The name of the journal mailbox being mined. • Journal Message Count: The number of messages in the journal mailbox at the time statistics were collected. • Journal Folder Size: The size (in KB) of the journal mailbox at the time statistics were collected. • Failed To Archive Count: The number of messages in the FailedToArchive folder of the journal mailbox at the time statistics were collected. • Failed To Archive Size: The size (in KB) of the FailedToArchive folder of the journal mailbox at the time statistics were collected.

Feature	Description
Process Information	• Status: No error conditions for any journal mining process on the journal mailbox. A journal mining process is in a retry state. A journal mining process is in a failed state. • Active: The number of journal mining processes that were executing at the time statistics were collected. • Failed: The number of journal mining processes that were in a failed state at the time statistics were collected. • Retry: The number of journal mining processes that were in a retry state at the time statistics were collected. • Completed: The number of journal mining processes that had completed processing at the time statistics were collected. • Queued: The number of journal mining processes that were queued for execution at the time statistics were collected.
Message Statistics	• Processed: The number of messages that have been processed. Note that *processed* means different things depending on the context in which it is displayed. • Submitted: The number of messages that were submitted to the IAP for archiving. • Tombstoned: The number of messages that were replaced with "stubs" on the Exchange server. • Ignored: The number of messages that were ignored. Note that *ignored* means different things depending on the context in which it is displayed. • Rejected: The number of messages that were rejected because there was something wrong with the original MAPI message, possibly message corruption on the Exchange server. • Average Processing Rate (msg/sec): The average processing rate.
Processes in Failed or Retry State	• Entry: This information helps identify the process that has entered a failed or retry state. An icon identifies the state: A minor error has occurred. These errors will be retried. A major error has occurred. These errors will not be retried. • Status: The problem that caused the process to enter a failed or retry state. • Proc: The number of messages that were processed before the condition occurred. • Tomb: The number of messages that were tombstoned before the condition occurred. • Sub: The number of messages that were submitted to the IAP before the condition occurred. • Rej: The number of messages that were rejected before the condition occurred. • Ign: The number of messages that were ignored before the condition occurred. • Elapsed: The number of elapsed seconds of processing before the condition occurred. • Last Run: The date and time that the condition occurred.

Selective Archiving

The Selective Archiving section contains the following areas described in Journal Mining:

- Process Information
- Message Statistics

- Processes in Failed Or Retry State

See "Journal Mining" on page 123 for information on these areas.

Synchronize Deleted Items

The Synchronize Deleted Items section contains the following areas described in Journal Mining:

- Process Information
- Message Statistics
- Processes in Failed Or Retry State

See "Journal Mining" on page 123 for information on these areas.

Tombstone Maintenance

The Tombstone Maintenance section contains the following areas described in Journal Mining:

- Process Information
- Message Statistics
- Processes in Failed Or Retry State

See "Journal Mining" on page 123 for information on these areas.

11 Enabling the AuditLog

The AuditLog provides a surveillance system log for companies that are required to prove they are adhering to compliance processes. The AuditLog records details about the following types of actions:

- Actions that users perform from the Web Interface, such as searching for an email, downloading a file, or viewing a message.
- Actions that are performed to manually delete emails from the IAP. Logged actions include granting Administrative Delete privileges to a user from the IAP's Platform Control Center (PCC), and deleting emails from the IAP via the Administrative Delete feature in the Web Interface.
- Actions that are performed to merge duplicate copies of an email in IAP's Duplicate Manager.

For a complete list of the actions that are logged, and for information on performing AuditLog repository queries, see "Querying the AuditLog" in the IAP user guide.

This chapter describes how to enable the AuditLog feature, set the retention period, monitor status, and grant user access to the repository.

Enabling the AuditLog feature

The AuditLog feature is enabled per domain. To enable the AuditLog feature for a domain, add the following attributes in the `/install/configs/primary/Domain.jcml` file for each domain in which the AuditLog feature needs to be enabled:

- `AuditLogEnabled=true`
- `domainLog_IP=<VIP for this domain>`

  If the IAP is part of a replicated site, the value of the domainLog_IP attribute needs to be the VIP of the primary site.
- `DocClass=email,windoc`

Granting user access to the AuditLog repository

By default, no user has access to the AuditLog repository. To grant access to a compliance officer or other user:

1. Log in to the Platform Control Center (PCC).
2. Go to the Account Manager page (**User Management** > **Account Manager**).
3. Click the **User** radio button at the top of the page.

4. Find the user's name in the list.

   You can search for the user with the letter buttons, or by using the search function.

   See "Account Manager features" on page 68 for information about user name searches.

5. In the Integrated Archive Platform Account and LDAP Information form that appears, clear the check box labeled **Deactivate this check box and then edit the user entries**.
6. Click any entry in the Direct Repositories box.

   There should be at least one entry in the box, for the user's personal repository.

7. In the Adding Repository Access form that appears, click the **Other** tab.
8. Select the check box for the AuditLog repository, for example `<domainname>.userauditlog.repository`.

NOTE:

The user's personal repository must be in the same domain as the AuditLog repository.

9. Click **Save to List**.
10. Click **Exit**.

    The repository is added to the user's direct repositories.

Monitoring status

Use the PCC Platform Settings page (**General Configuration** > **Platform Settings**) to check whether the AuditLog feature (AuditLog Service) is enabled for a specific domain.

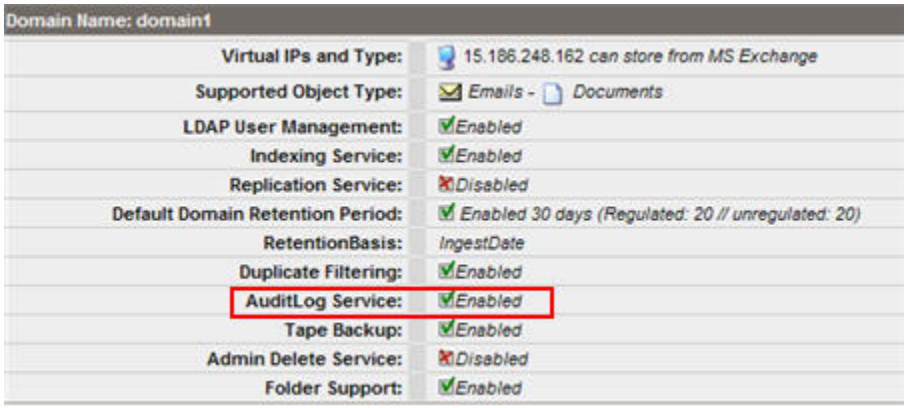

Figure 30 AuditLog enabled

AuditLog repository retention period

The default retention period for the AuditLog repository is seven years, and can be changed. For instructions on changing the retention period, go to "Editing repository retention periods" on page 99.

12 Optional Backup System

For customers who require a tape backup of IAP configuration, DB2 database, and Smartcell data, HP offers an optional backup system which runs HP Data Protector backup service software. An HP authorized representative can add the optional backup system by adding one or more backup servers to IAP and pre-configuring the backup system. The HP representative enables the hardware and application backup (APP) in the `BlackBoxConfig.bct` file and sets the `Domain.jcml` DataBackupEnabled flag to true: `DataBackupEnabled=true`

Once properly configured, the backup server won't require manual interaction (beyond managing tapes, e.g. making sure there are enough free tapes in the library). Administrators need to scan and format the tapes – tapes can only be used for backup after they are formatted. For more information, consult the Data Protector documentation.

Administrators are also expected to connect the media to the backup server and configure the device for single or additional backup servers. For more information, consult the Data Protector documentation.

The PCC Backup pages provide general backup status information for database, configuration, and Smartcell data.

Topics in this chapter include:

For more information and best practices, see the *Data Protector 6.1 Concepts Guide* available at http://www.hp.com/support/manuals and Data Protector online help.

Default configuration

Data Protector (as a backup solution) is automatically installed and pre-configured on IAP. This configuration includes:

- Auto-configuring of your attached tape library by Data Protector
- Adding the IAP Smartcells, DB2 database, and configuration data to the backup specification
- Setting up a default backup schedule

After the HP representative has set up backup, backups are done on a regular basis using the attached tape library. However, you may want to change this configuration to customize the setup to certain needs or to introduce vaulting capabilities. This chapter includes a sample scenario which explains how to use Data Protector to make these changes.

Tape library

Data Protector auto-configures the attached tape library when the HP service representative or administrator uses the Autoconfigure Devices option to auto detect the library. (To use this option, use VNC to open the Data Protector UI, select the **Device & Media** context, right click **Device**, and then click **Autoconfigure Devices**. When a the tape library is attached or added after the installation, this step can be performed by the administrator.)

You can see the configuration by selecting the **Devices & Media** context from the Data Protector UI. To open the Data Protector UI, initiate a VNC connection from the PCC or your local desktop to the backup server.

After expanding the Devices section, you can see the drives, robotic paths and slots. In this example, the tape library has two tape drives and several slots for tapes partly loaded with tapes. Data Protector uses its standard algorithm for tape rotation to use the tapes equally available in the tape library.

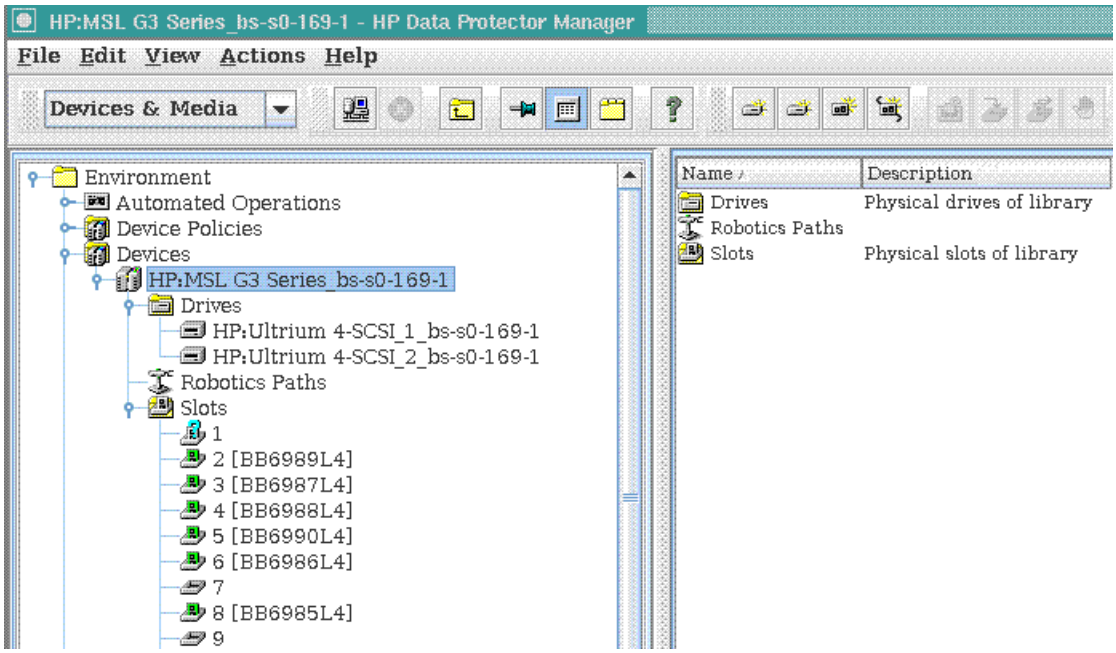

Figure 31 Data Protector Devices & Media

NOTE:

The tape library drive is specified in the backup specification by going to the Data Protector Backup context, expanding the tree, and selecting one backup specification, e.g. IAP_DATA, selecting the Destination tab, and selecting the tape library drive, and clicking **Apply**. After doing the same steps for the other backup specifications IAP_CONFIG, and IAP_DB2, backup runs on the schedule.

Backup specification and scheduling

The auto-configuration also sets up the backup specification including the data sources, shown in the figure below. You can access it by selecting **Backup** from the context drop down box. The left pane lists all backup specifications needed to properly back up your IAP, i.e. IAP_CONFIG, IAP_DATA, and IAP_DB2. The right pane shows the properties for each individual backup specification, in the data sources. The Smartcells that are included are automatically updated whenever the Smartcells

setup of the IAP changes. You should not make any changes here. If you experience misconfigurations here, contact HP.

NOTE:

If an IAP configuration was changed to remove some existing Smartcell groups and add new groups, the tape backup console will show the old groups that are removed as registered backup clients. Although, these groups are still available in the tape backup console, these groups are not selected as a source for backup in the specification. Only valid groups are selected for backup.

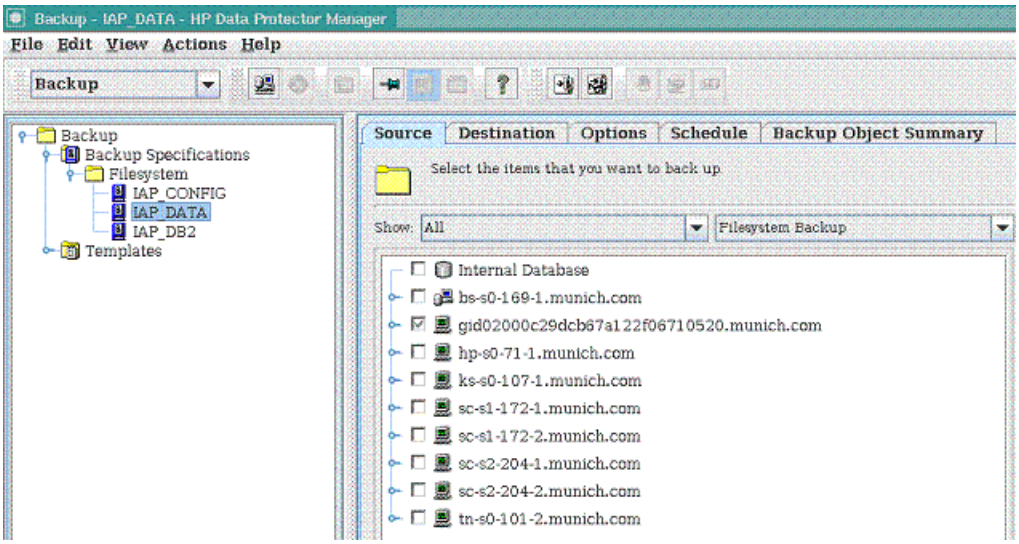

Figure 32 Backup Specifications and Source

If you click **Schedule**, you will see the predefined schedule. As you can see, the schedule is automatically set up to do a full backup every four weeks and an incremental backup once a day.

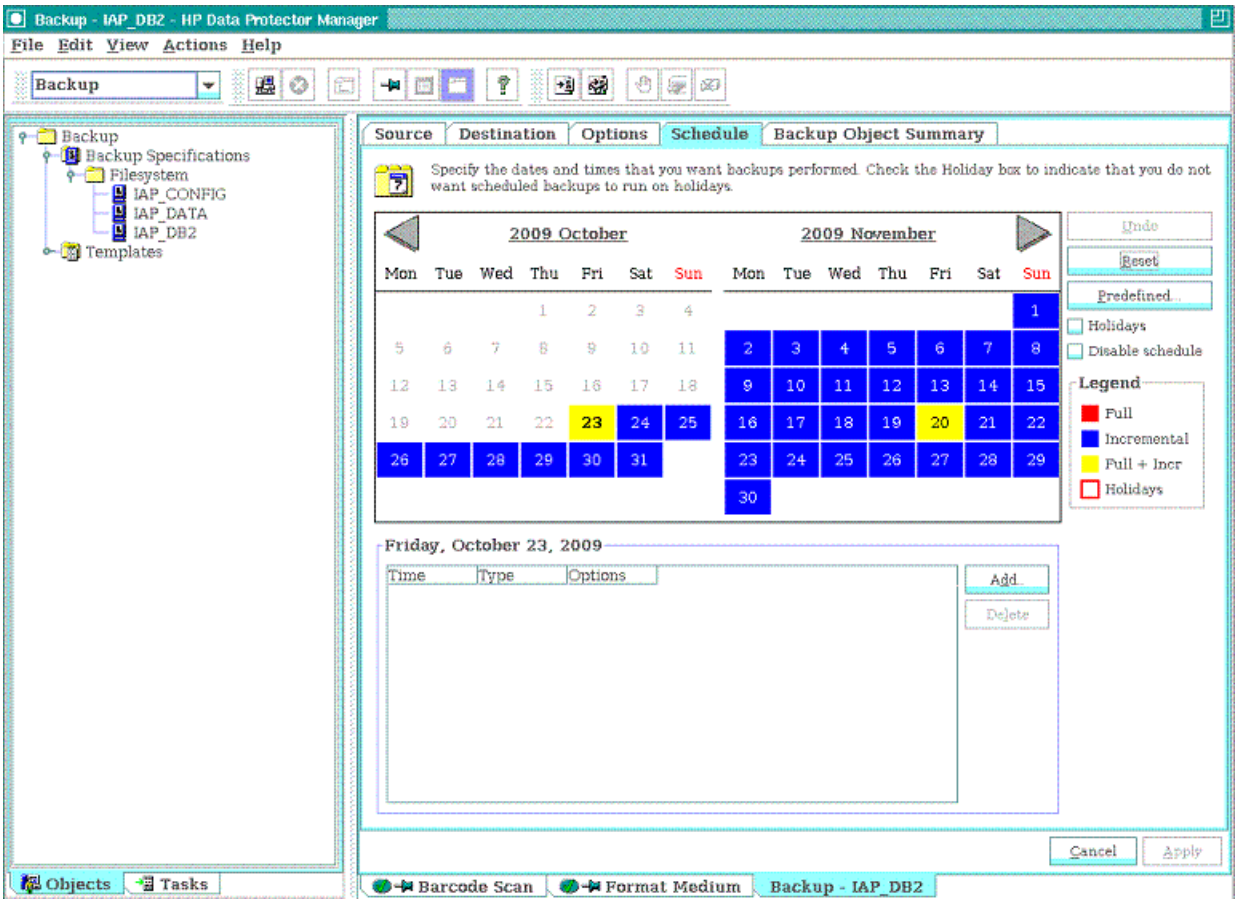

Figure 33 Backup Schedule

Changing the configuration

The section above introduced Data Protector and explained the default backup configuration that IAP uses. Let's look at what you would need to do to deviate from this default configuration and customize your backup strategy to a full backup every two weeks and an incremental backup once a day.

NOTE:

It is best to have the full backups of IAP_CONFIG, IAP_DATA, and IAP_DB2 at the same point in time in case you would like to use the vaulting strategy (described later in this chapter).

The first step is to change the backup schedule for each backup specification. To do this, select the **Schedule** tab and click **Reset**. This will remove all pre-configured schedules. Then click **Add** and set up the schedule for the full backup as shown below. After you click **OK**, every other Friday in the calendar should be marked red.

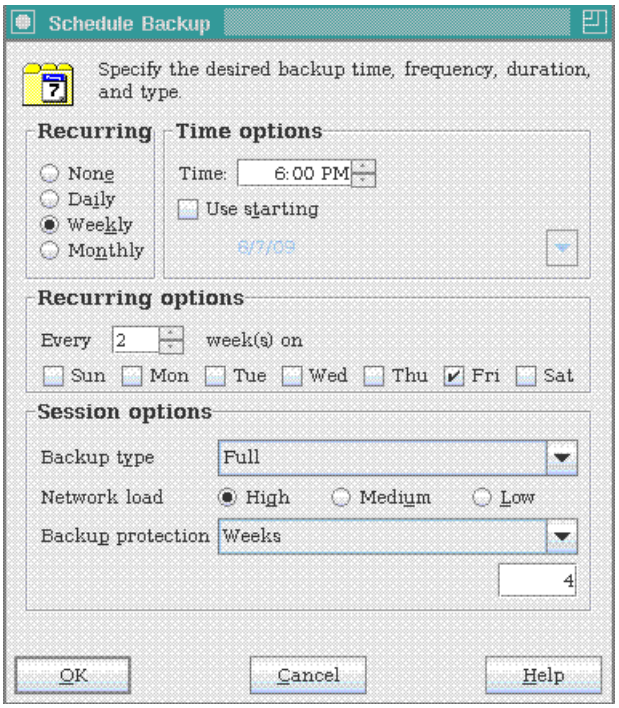

Figure 34 Full backup

The second step is to do the same to set up the incremental backups. Note that for both schedules it is very important that you specify backup protection to four weeks, so that Data Protector can reuse the tapes after four weeks.

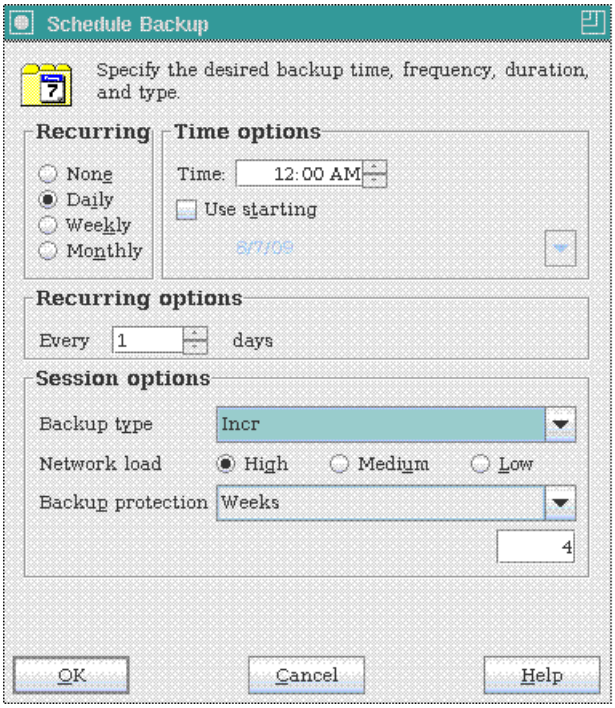

Figure 35 Incremental backup

Having changed the configuration, your backup works already. If you do not need any vaulting you can stop here; Data Protector will automatically use the available tapes for storage and make sure

that your data is protected. However, you may want to introduce some vaulting capability and store sets of tapes offsite. This will be explained in the next section.

Operational workflow for vaulting

One recommended way to perform vaulting is to have three different sets of tapes and cycle through them. This means:

- For the first backup cycle have a set of tapes in the tape library storing the data of the first full backup followed by the frequently applied incremental backups
- After the first backup cycle, the complete set of tapes used for the backup is replaced by a second set of tapes. The first set can be brought to an offsite location for vaulting purposes.
- After another backup cycle, the second set of tapes is replaced by a third set of tapes. Now the second set of tapes can be brought to the offsite location as well.
- After yet another backup cycle, the third set of tapes is replaced by the first set of tapes, overwriting the old backup.

For the sample scenario explained in the previous section, the duration of a cycle would be four weeks. A backup cycle always begins with a full backup followed by a number of (optional) incremental backups and ends with the last incremental backup before a new full backup.

To enhance data protection even more, the scenario could be extended by making copies of the sets of tapes before moving them to one offsite location. The copies could then be moved to another offsite location.

Estimating needed number of tapes

Let's assume that there is only one pair of Smartcells in the IAP, i.e. you have 5 TB of data to back up. Therefore, for the full backup you need roughly a maximum of seven tapes (assuming LTO-4). Depending on how many changes there are during one backup cycle, you need an appropriate additional amount of tapes for the incremental backups, say one more tape. For three different set of tapes (for vaulting) you need 24 tapes.

In the same scenario for LTO-3 you would need twice as many tapes.

This is a rough estimate and it can deviate depending on your use case. The amount of changes can also vary over time within one setup, so you should always make sure that there are enough tapes in the library. For a quick overview, however, the following table shows the numbers of tapes needed for one backup cycle (note that you will need three times as many tapes for implementing vaulting).

Smartcells	LTO-3 tapes	LTO-4 tapes
1 (5 TB)	16	8
2 (10 TB)	32	16
3 (15 TB)	48	24

Refer to *Data Protector Concepts Guide*, Chapter 3 “Media management and devices” for more information.

Replacing tapes

Let's assume that you have the first set of tapes already in the library and the backup is working. After each backup cycle, operational work is required.

You can determine the end of a complete backup cycle by looking at the backup schedule. A backup cycle ends with the last incremental backup before a new full backup.

If you would like to make sure that your operational work does not conflict with the next full backup, navigate to the PCC Web Interface and disable all backup schedules. Also make sure that no backup is currently running using the Data Protector graphical user interface. Don't forget to enable the backup schedule once you are done with operational work.

You can label tapes, eject them and replace them with new ones from within the **Devices & Media** context in Data Protector. Navigate from Devices via your tape library to the Slots node and expand it. You should see all your tapes.

To check which tapes were used for backup and need to be ejected, select each tape and click the **Objects** tab. This shows you the backup object it contains. It also shows the backup times (Start Time/End Time), which you can use to determine if a tape was used in the current backup cycle or not.

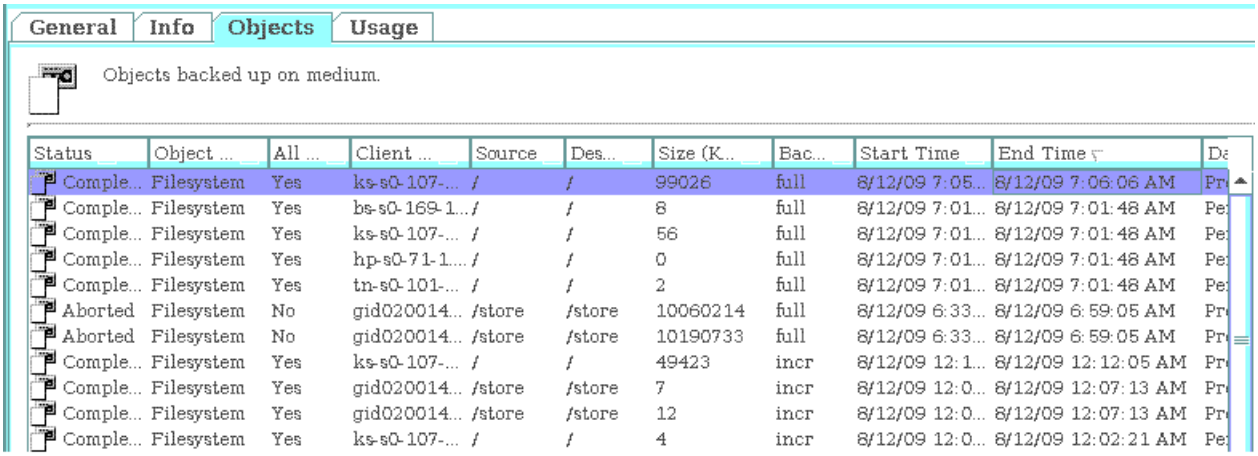

Figure 36 Backup Object

NOTE:

Data Protector respects retention times set in the backup schedule. This means that no backup object will be overwritten on a tape until it expires. It also means that you do not have to delete certain backup objects, since Data Protector will automatically do that if it needs the space on the tape.

Before ejecting a tape:

1. Go to Data Protector UI **Device & Media > Device > Slots** and select the tape which you want to eject for vaulting and right click and select **Recycle**.
2. Again select the tape and click on **Export**.
3. Select the tape and click on **Eject**.
4. The tape will be ejected and can be taken to an offsite location for vaulting.

Every tape you eject should be replaced by a tape from the next set of tapes.

TIP:

If you have many tapes to eject, select them all while holding down the **Ctrl** key to select them individually or while holding down the **Shift** key to select a complete range. Then open the context menu and select **Eject**.

When you hold down the **Ctrl+eject** keys you see a page that lets you enter a vault location for a tape — for multiple tapes it will ask one by one the location. After you enter the location for all the tapes, they will eject.

IMPORTANT:

You should not eject a tape directly from the library panel. If you do, Data Protector cannot keep track of the tapes. Always use Data Protector to eject backup tapes. To enter a tape, insert it into the tape library and select **Enter** on an empty slot from within Data Protector.

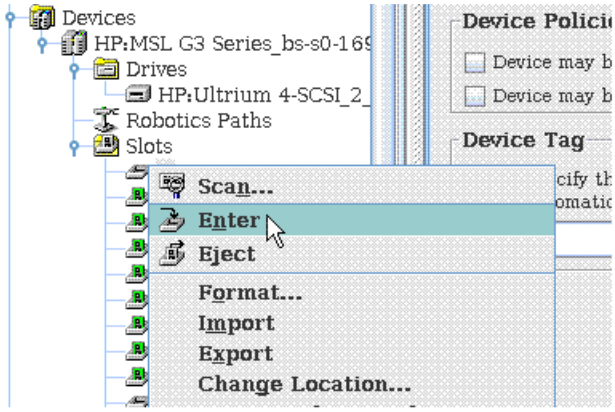

Figure 37 Entering a tape

After replacing all used tapes of the set by a new set, you should move the ejected set of tapes to an offsite location for vaulting purposes. Do not forget to enable your backup schedule if you disabled it for replacing the tapes. The operational work is done. After the next backup cycle, perform the operational work again.

Restoration of a Smartcell group from a vaulted tape

If you experience the loss of both Smartcells in a group, you can use the last backups to restore the Smartcell group. (If you lose only one Smartcell of a group, use the PCC Smartcell Cloning menu function to recover the lost Smartcell.) To recover both Smartcells from tape, go to the PCC Data Management Backup page and make sure that backup scheduling is disabled.

With the backup scheduling disabled, verify that there are no backups in progress using the Data Protector console's **Monitor** context. For example, the Current Sessions view should display the sentence "No running sessions."

Once you have verified that no backups are in progress, eject the current set of tapes from the tape library and replace them with the set of tapes from the last full backup of the Smartcell group.

When the tapes are in the tape library, use the Data Protector console's **Devices & Media** context to scan the barcodes of the tapes. If the tapes were created on a different backup server, or if the backup

server has had its Data Protector's internal database (IDB) cleared of information about the vaulted tapes, then the tapes will be listed as either "DataProtector Foreign" or "unknown" tape formats. If either of these formats is displayed after the scan barcode function, then highlight the tapes and select the **Import Catalog** function.

Importing a medium writes detailed information about backed up data that is 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. Refer to the *HP OpenView Storage Data Protector Administrator's Guide* for more information on media management.

After inserting vaulted tapes:

1. Scan the device.
2. Right click on the slot/tape and click on **Import**. The sessions will be imported.

NOTE:

Attribute information such as object or media size is not 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.

When finished importing the catalog, verify that the correct tapes have been mounted and cataloged by using the Data Protector console's **Devices & Media** context. Also examine each slot of the tape library under the **Objects** tab to verify that the Smartcell group's data is on tape in one or more of the slots, i.e. under the heading **Client** you should see the groupid (gid0200...) of the Smartcell group that was backed up. Verify that the tapes are loaded by looking at each of the tapes, searching for the Smartcell group ID.

With the tapes from the last full backup placed in the tape library, use the PCC Restore Smartcell Group menu function to start the Smartcell group restore. The PCC will allocate a Smartcell to hold the restored data and Data Protector will begin reading data from the tape restoring it to the newly allocated Smartcell.

During the restore process, you may find that the restore requires additional tapes from subsequent incremental backups. Use the Data Protector's **Monitor** context to keep track of the progress of the restore, and respond to the additional media mount requests as Data Protector issues them. Pay close attention to the device and slot information that is displayed in the mount message and make sure to place the correct tape into the requested device and slot. For more information about Data Protector devices and media management, refer to the Data Protector documentation.

When all the data has been restored, the PCC will manage the post data restore processing. It will require user interaction to start cloning to complete the restore process.

Restoration in case of non-Smartcell group data loss

If you experience data loss of backed up information other than Smartcell groups, use Data Protector for the restoration process. Change to the **Restore** context and expand the **Filesystem** node. Select the client and backup object you would like to restore and click **Restore**. Data Protector will

automatically know which tape it needs. If the tape is not in the library Data Protector will tell you. From the location label you will know where the tape is located.

Restoring IAP configuration information

The tape backup system stores selected IAP configuration information to tape. The Data Protector backup specification named IAP_CONFIG defines the specific information that is stored to tape.

The IAP_CONFIG backup specification stores information from the PCC, kickstart, and backup servers:

- The PCC server information includes server certificates and host keys.
- The kickstart server information includes IAP configuration and Archive Gateway backup data.
- The backup server information includes backup specifications and schedules.

In the event that the configuration information is lost, overwritten or corrupted, it can be recovered from the tape backup. The information should be restored only on an as needed basis as it may not be as current as information stored on servers themselves. The information stored on the tape backup is intended for disaster recovery and under most operating conditions; it should not need to be restored.

The IAP_CONFIG information is restored using the Tape Backup Console (the Data Protector graphical user interface). To bring up the console:

1. Open the tape backup console by clicking the tab on the backup page in the IAP PCC UI.
2. In the tape backup console, connect to the HP Data Protector Cell Manager, by selecting **File** > **Connect to Cell Manager**. Select the primary backup server to connect. If you are already connected you can skip the connection step.
3. Once connected, you can use the Data Protector **Restore** context to select the information to restore. The information is restored by selecting the machine and then selecting the files to restore from tape using the **Source** tab. In the example below, the PCC server (tn-s0-101-2.viap11.com) certificate and host key files are being restored.

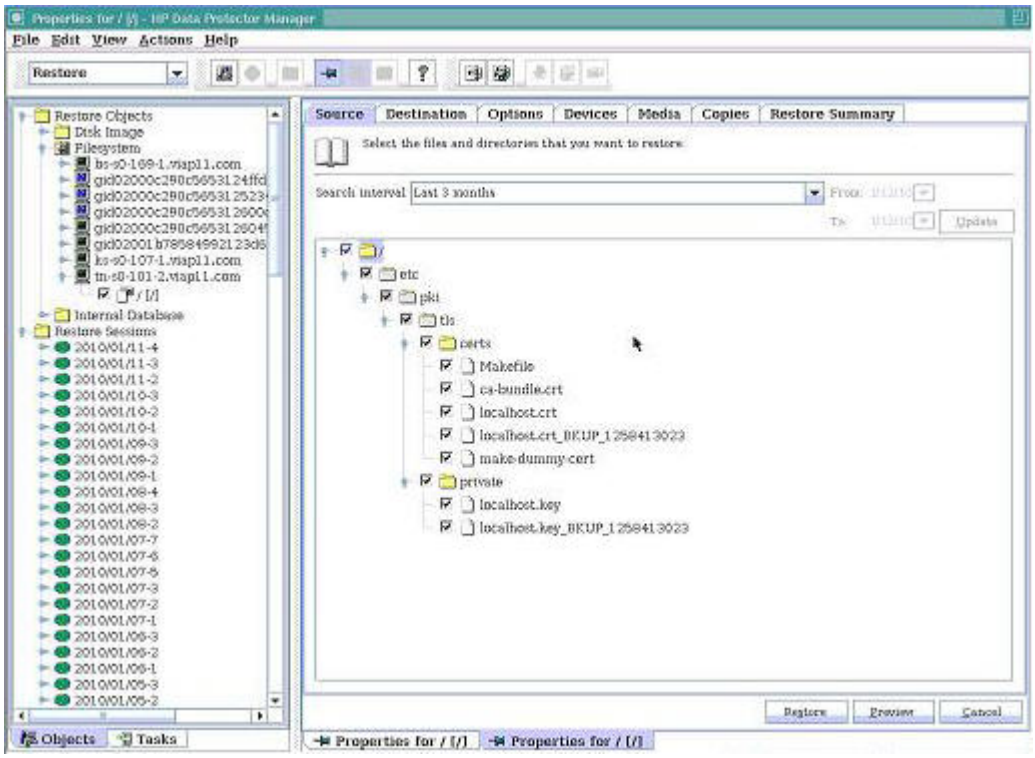

Figure 38 Restoring information

4. Click the **Destination** tab and verify that the files are being restored to their original locations and that the most recent file is being used.
5. Click the button labeled **Restore** (located at the bottom of the page) when all the options have been selected. The **Preview** button can be used to perform the restore without actually storing the files to the destination. After clicking **Restore** or **Preview** you may be asked to choose multiple or single objects to restore. If you are not asked this question, then skip to step 7. Otherwise, choose single and click **Next**.
6. If asked to select host (object), make sure you select the host chosen earlier (step 3) from the drop-down box. In this example, we are restoring the PCC (tn-s0-101.2.viap11.com).
7. A summary of the proposed restore will be displayed. If it is correct, click **Finish**. Otherwise, use the **Back** button to navigate back and make your changes. Once **Finish** is clicked, the restore is started as a Data Protector restore session and progress messages will be displayed. When the restore is finished, a completion message is displayed.

Configuring multiple backup servers

IAP 2.1 supports adding multiple backup servers to shorten full backup cycles. We recommend that you add one backup server per 50 TB of data. Up to three backup servers are supported. This translates to 25 groups of 2 TB Smartcells or 10 groups of 5 TB Smartcells per backup server. If more than three backup servers are required, replication should be used.

The installation of additional backup servers is the same as installing the primary backup server.

The HP Data Protector Cell Manager is installed and configured to work only from the primary backup sever. (This is the first backup server configured in the `BlackboxConfig.bct` file, typically 10.0.169.1). The tape backup console (the Data Protector graphical user interface) runs on the primary backup server and all the backup specifications and media configurations are configured

and controlled from the primary backup server. A default specification is provided and can be customized based on the customer needs.

The second and third backup servers act as media agents only, providing additional backup devices to increase the amount of data that can be concurrently backed up. Additional tape devices are connected to these backup servers. Once the device and media is connected to the additional backup servers, the backup devices can be imported into the primary backup server as devices from the additional backup servers. The importing of devices is done using the tape backup console running on the primary backup server.

If the primary backup server is down for any reason, tape backup will not be run. If the primary backup server cannot be brought back up, the IAP must be reconfigured to use one of the additional backup servers as the primary backup server.

Automatic failover of the primary backup server to one of the additional backup servers is not supported in IAP 2.1. A fresh kick of the backup servers is used to perform the reconfiguration to a different primary backup server.

To reconfigure backup server configurations, comment or remove the "down" backup server from the `BlackboxConfig.bct` file, then kick all the backup servers. Note that the `regloader.pl -cv` command needs to be run after changing the `Blackbox.bct` file.

To configure devices attached to additional backup servers:

1. Open the tape backup console using the tab on the backup page in the IAP PCC UI.
2. In the tape backup console, connect to the HP Data Protector Cell Manager, by selecting **File** > **Connect to Cell Manager**. Select the primary backup server to connect. If you are already connected you can skip the connection step.
3. Once connected, configure the devices connected to the additional backup servers using the **Devices & Media** drop down menu. In the figure example below, bs-s0-169-2.paris.com is the second backup server:

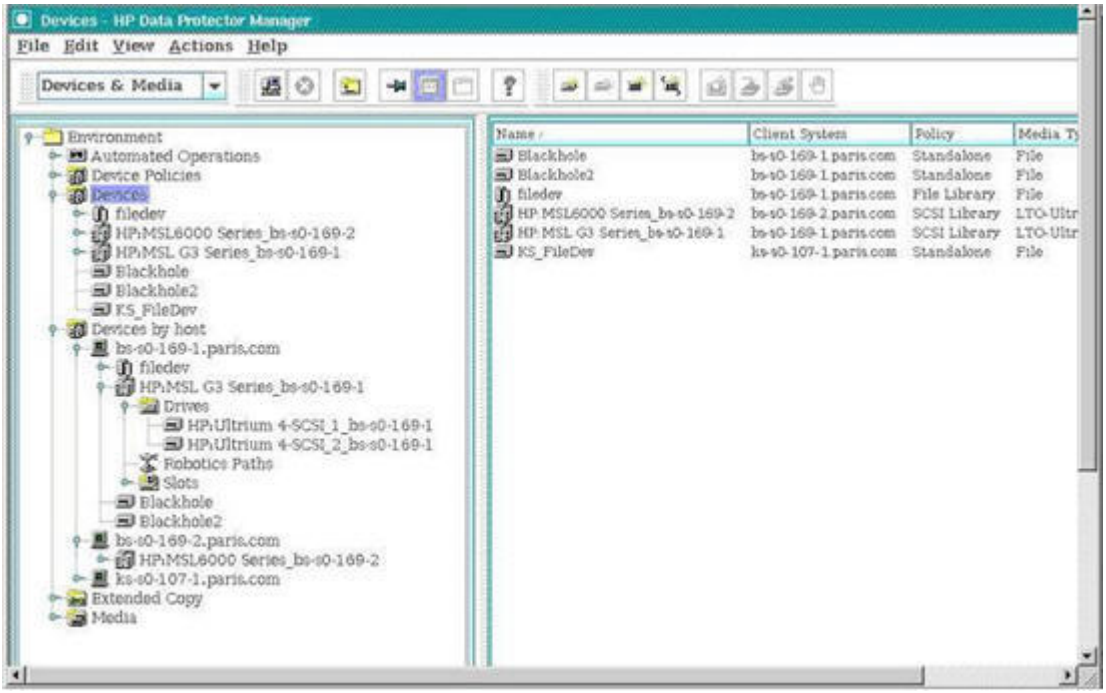

Figure 39 Configuring devices

4. To add an additional backup device, right click the **Devices by host** menu item and select **Add Device** from the drop down menu. In the **Client** field, choose the additional backup server. As we mentioned, for our example this would be bs-s0-169-2.paris.com.
5. Select the device and finish the configuration of the device.
6. Once the backup device on the additional backup server is configured and its media prepared for backup, change to the Data Protector **Backup** context and modify the IAP_DATA backup specification in the **Backup** drop down menu to add this newly configured device as the destination for backup.
7. The backup will run based on the schedule and data is backed up to all the selected media at the same time. The tapes required for restore are automatically identified by Data Protector and no additional setup is required.

Troubleshooting

For detailed backup server status information, use the Data Protector graphical user interface running on the backup server (connecting via VNC). The backup service on the Smartcells, the kickstart server, and the DB2 server keeps a log file of recent events.

Data Protector keeps a log file on the backup server. This log file is also be included in the IAP log collection script.

If a problem with the IAP backup occurs, determine if the root cause is on the IAP or on the Data Protector side. Check the PCC status page and the alerts sent through the alert framework. For example, an alert notifying you that the library is out of free tapes would clearly indicate that this is a Data Protector issue in the error message. In addition, the log files both on the IAP and the Data Protector side help determine on which side the problem originates.

NOTE:

There is currently a Data Protector limitation in which the Data Protector **Devices & Media Usage** tab information for tapes/slots may show zero percent usage even though the **Objects** tab correctly shows that data has been imported.

13 Remote management of servers

Server remote management capabilities allow you to perform administrative activities on IAP servers without having to be physically present at the console. Remote management capabilities allow you to use remote power management, remote console, and other features that iLO and Lights-Out 100 support. Every server within IAP can be remotely managed using iLO or LO100.

This chapter describes common tasks for administering remote management for IAP servers using HP ProLiant Onboard Administrator Powered by Lights-Out 100 or Integrated Lights-Out 2. Throughout this chapter we refer to the HP ProLiant Onboard Administrator Powered by Lights-Out 100 or Integrated Lights-Out 2 as LO100, iLO, or Remote Processor.

Topics in this chapter include:

IAP 2.1 hardware (DL180 G6, DL360 G6) and newer will ship with the iLO/LO100 Advanced Pack license preinstalled on the servers. Customers with previous hardware releases who want the Advanced Pack features will need to obtain and install the licenses through their standard channels. This chapter assumes that the Advanced Pack license has been installed on the server. Licenses can be installed on the servers either through the iLO/LO100 web interface or using the command line tools `hponcfg` and `lo100cfg`. Refer to the iLO/LO100 user guides listed below for details.

Refer to the HP ProLiant Onboard Administrator Powered by Lights-Out 100 or Integrated Lights-Out 2 product documentation for comprehensive instructions for configuring and using remote management.

- HP ProLiant Lights-Out 100 Remote Management User Guide for HP ProLiant DL140 G2, ML110 G3, and ML150 G2 Servers at http://bizsupport1.austin.hp.com/bc/docs/support/SupportManual/c00857454/c00857454.pdf
- HP ProLiant Onboard Administrator Powered by Lights-Out 100 User Guide for HP ProLiant ML150 G6, DL160 G6, and DL180 G6 Servers at http://bizsupport1.austin.hp.com/bc/docs/support/SupportManual/c01716486/c01716486.pdf
- HP Integrated Lights-Out User Guide for HP Integrated Lights-Out firmware 1.91 at http://bizsupport1.austin.hp.com/bc/docs/support/SupportManual/c00209014/c00209014.pdf
- HP Integrated Lights-Out 2 User Guide for Firmware 1.75 and 1.77 at http://bizsupport1.austin.hp.com/bc/docs/support/SupportManual/c00553302/c00553302.pdf

Remote management solution overview

The following sections provide an overview of the IAP remote management solution.

Goals

The primary goals of the IAP remote management solution are to:

- Allow you to choose between:
  - Direct access to iLO remote management processors via your management network for usability and higher availability (no single point of failure)
  - Proxy access to iLO remote management processors for manageability and greater security
- For any IAP server, allow you to:
  - Establish a secure connection to the iLO remote management processors
  - Power cycle or turn the remote server on or off
  - Monitor or interact with the boot sequence of the remote server
  - Access local media or ISO images on the server to remotely upgrade firmware

Network topologies

The IAP remote management network keeps the iLO remote management processors isolated, both physically and logically. The administrator can choose whether the IAP remote management network is connected to the customer management network directly or by proxy. For both direct and proxy access, **the PCC remote management processor must always** be connected to the customer management network.

Direct mode

In Direct mode the iLO interfaces of the IAP servers are connected to the customer management network directly. Since the IAP remote management network is self-contained, connecting it directly to the customer network does not pose the risk of accidentally exposing internal IAP network. In this network topology, a port on one of the dedicated remote management network switches is uplinked to a switch in the customer's management network. The administrator is then responsible for configuring and managing access to the remote processors in the IAP. For example, the administrator has to allocate IP addresses on his network for each remote management processor in the IAP. This association between a given IAP server and its remote management processor has to be maintained by the administrator, e.g. using a spreadsheet that needs to be updated each time physical servers are added or replaced in the IAP. To help with the setup task, HP provides a PCC command to use to set up the network settings and user management of the remote management processors. This script and the required configurations steps are discussed below.

Advantages:

- Usability
  - Can access all remote management processors concurrently
  - Can be integrated with customer network management solution for access from their management console
- Availability: Direct access to remote management processors means no IAP Single Point of Failure

Disadvantages or limitations:

- Security: All internal IAP servers can be accessed via their remote management processor, bypassing the IAP firewalls
  - Potential for unauthorized access to IAP servers
  - The administrator is responsible for restricting access to the IAP remote management processors only to authorized personnel
- Manageability: No automated mapping of a given IAP server IP address to its remote management processor IP address
  - Requires maintaining manually a mapping of IAP server IP to remote management processor IP

- Such a mapping would have to be communicated to HP Support personnel needing access to the IAP remote management processors

- More IP addresses required on customer network for IAP: Need one IP address for each IAP server's remote management processor

Proxy mode

In Proxy mode all IAP server iLO interfaces are connected through a separate network to the PCC. A dedicated NIC on the PCC acts as the gateway to the IAP remote management network. Access to the iLO network is through an additional IP address on the PCC that acts as the proxy. IAP assigns internal IP addresses to all iLOs and the PCC gateway. The PCC iLO is connected to the customer management network. Additionally, a Virtual IP (VIP) address on the same network as the PCC IP address is used for the proxy access from the UI.

When it is connected by proxy, the network is contained inside the IAP. The remote processors are not accessible from the outside unless an action is initiated to request access to them. In this network topology, a user who is successfully authenticated on the PCC or the PCC Web UI has to request access to a remote processor 1) via the PCC command line interface using HP provided commands, or 2) via the PCC Web UI. The IAP firewall dynamically grants access to one remote management processor at a time. All access requests to remote processors are logged in an audit log.

Advantages:

- Security
  - Direct access to the remote management processors of all IAP servers (except PCC) is not available.
  - There is an audit log of requests to enable proxy access to remote management processors
- Manageability

  The PCC UI displays a list of IAP servers and redirects the user transparently to the remote management processor of the selected server.
- You need fewer IP addresses for IAP on the customer network

  You need one IP address for the PCC remote management processor and one VIP for proxy access from the user interface.

Disadvantages or limitations:

- Usability

  Can only enable proxy access to one remote management processor at a time
- Availability: PCC is an IAP Single Point of Failure
  - Access to internal remote management processors is via proxy on the PCC, which requires at a minimum that the PCC operating system and networking are working.
  - Though access to remote management is through the PCC UI, JBoss does not need to be up and running to gain access to the remote management processor. Users can bypass the UI and use a direct SSH from the PCC to the iLO.
  - You can enable proxy from the command line by running the "iloproxy" script.
- The PCC must be a ProLiant DL320 G5 or newer server. (The DL360 G4/G4p are lacking an eth2 network interface that is required to connect the Remote Management network to the PCC.)
- The DL140 G2 LO100 web user interface can not be accessed via PCC in Proxy mode. The LO100 command line interface can be accessed by logging on to the PCC and telneting to the IP address of the LO100 remote processor.

Hardware wiring

All LO100s/iLOs for IAPs hardware generation 2.1 or above are pre-wired at the factory in the proxy access mode.

At deployment time, the PCC iLO is wired to the customer management network. Should the customer prefer a direct access topology, the wiring can be altered to remove the connection to the PCC eth2 interface and uplink the switches to the customer management network.

Checking the iLO mode

To determine the iLO access mode of the IAP, refer to the "REMOTE MANAGEMENT" section of the `/install/configs/primary/BlackBoxConfig.bct` configuration file on the kickstart machine. There is a field called "iLO Access Mode" with the value of either direct or proxy. When this section is missing from the configuration file, remote management is disabled for the IAP.

Disabling remote management

The IAP is shipped with remote management wired for Proxy mode, but is not enabled by default. If remote management has been enabled by an administrator, it can be disabled with the following steps:

1. Edit the file `/install/configs/primary/BlackBoxConfig.bct` on the kickstart machine to remove the section "REMOTE MANAGEMENT" from the file.
2. Run `createBBC.zr`.
3. Run `converter.zr`
4. Run `regloader.pl –cv`

Using Proxy mode

Prerequisites

You must supply the following:

- iLO Subnet – The default subnet that iLO is configured with is 10.1.0.0/16. However, if this conflicts with your network, the subnet can be updated to X.X.0.0/16 as you desire.

  The value of 10.1.0.0 for "iLO Subnet" is a default and can be changed to one that is suitable for your environment. In Proxy mode, the static IP addresses for the iLOs are derived automatically from the server IP address and assigned on the configured iLO Subnet. For example, when the server is at IP address 10.0.172.1 and the iLO Subnet is 14.0.0.0/16, then the iLO is assigned the IP address 14.0.172.1, and the PCC server's eth2 interface is assigned a static IP address 14.0.0.1 on the iLO network. When the default iLO Subnet of 10.1.0.0 is used, then the iLO of the server whose IP address is 10.0.172.1 is assigned the IP address 10.1.172.1 and the PCC server's eth2 interface is assigned the IP address 10.1.0.1.
- iLO Proxy VIP – The iLO Proxy VIP is an IP on your network to allow access to the iLO proxy. Enter the IP in the BCT with the mask binary shorthand. This IP must be on the same subnet as the PCC external IP.
- IP of PCC iLO – The PCC iLO IP is the IP of the PCC's remote management processor. This IP does not have to be on the same subnet as the PCC external IP. In Proxy mode, the PCC is the only

server for which its iLO device sits on your network. The rest of the servers use internal IP's. This is required to be able to access the PCC iLO device when the PCC server is unreachable.

Configuring remote management

To configure remote management in Proxy mode:

1. Edit the `/install/configs/primary/BlackBoxConfig.bct` file on the kickstart machine to add the following section:

   ```
   ########################

   ### REMOTE MANAGEMENT ###

   ########################

   iLO Access Mode: proxy

   iLO Subnet: 10.1.0.0/16

   iLO Proxy VIP: <proxy_VIP>/??
   ```

2. Run `createBBC.zr`.
3. Run `converter.zr`
4. Run `regloader.pl –cv`
5. When run on the PCC server, the following command configures iLO on all the IAP servers, creates the Administrator user for iLO, and allows you to set the administrator supplied password for that user: `/opt/tools/ilo/configure_ilo –v –set –f /opt/tools/ilo/ilo_ips.csv`

   The file `ilo_ips.csv` is a text file that must be in the directory from which the command is run. The file must contain the contents:

   <Internal IP of PCC>,<IP of PCC iLO>,<iLO netmask>,<Gateway IP for iLO Subnet>

   Here is an example of the contents: `10.0.101.2,15.1.2.30,255.255.0.0,15.1.2.1`
6. When prompted, enter a password for the Administrator account.

Accessing the remote management web interface

To access the remote management web interface of all servers except the PCC:

1. Log in to the PCC using the PCC external VIP address.
2. Select **Hardware Management** from the left menu.
3. In the Hardware Management view, click on the **Manage** link for the server for which remote management is required. For iLO, this opens the Remote Management login page for the login credentials. (For LO100, an intermediate access page will be opened first which provides access links for the LO100 Web UI and the LO100 Virtual KVM/Media applet. Virtual KVM/Media features are only available if you are licensed for ILO/LO100 Advanced Pack.)
4. Enter the user ID and Password login credentials to gain access to the remote management web interface.

To access the remote management web interface of the PCC: Connect to the URL *https://<PCC iLO IP address>* directly.

NOTE:

When Proxy mode is configured, access iLO/LO100 remote management only via the links in the Hardware Management page of the PCC because each link requests specific iLO/LO100 access for the server clicked on (since all servers have the same URL for iLO/LO100 in Proxy mode). Also, logging out (or timing out) from the PCC page will cause the iLO/LO100 access to be revoked. Therefore, connecting to the URL *https://<PCC iLO IP address>* directly will not work unless the remote management web interface has already been accessed at least once via the Hardware Management page link, and the PCC webpage stays logged in.

NOTE:

When accessing iLO/LO100 using a web browser, the browser may complain that the certificate presented matches another server. For IE, click **Continue** to ignore this warning. For Firefox, remove the previously stored certificate to avoid the warning.

Accessing the command line interface

To access the remote management command line interface:

1. Using `ssh`, login to the PCC server as `root`.
2. Enter the command `ssh Administrator@<iLO IP Address>` for the server to which iLO access is required.

   For example, use `ssh Administrator@10.1.172.1` to connect to the iLO IP 10.1.172.1.
3. When prompted, enter the password to gain access to the iLO command line interface of that server.

Using the command line interface to power cycle a server

To use the command line interface to power cycle a server:

1. Using `ssh`, login to the PCC server as `root`.

   NOTE:

   To access the LO100 of DL140 G2 servers, use `telnet` not `ssh`.
2. Enter the command `ssh Administrator@<iLO IP Address>` for the server to which iLO access is required.

   For example, use `ssh Administrator@10.1.172.1` to connect to the iLO IP 10.1.172.1.
3. When prompted, enter the password to gain access to the iLO command line interface of that server.
4. If you have iLO or iLO2, at the `hpiLO->` prompt, enter: `power reset`
5. If you have LO100, press Enter at the `EMS->` prompt to get the `Lights-Out>` prompt and enter the command: `reset [WARM|COLD]`

Using the web interface to power cycle a server

To use the web interface to power cycle a server:

1. Log in to the remote management web interface as described above in Accessing the remote management web interface.
2. If you are using iLO, click on the **Virtual Devices** tab, and then **Virtual Power**.
3. If you are using iLO2, click on the **Power Management** tab at the top.
4. If you are using LO100, in the left menu click on the Power Control Options link, and then select one of the "Power Control Options" from the pull-down menu and click the **Apply** button.
5. For the iLO type, select one of the power cycle options.

Accessing the KVM console

To access the KVM console:

1. Log in to the remote management web interface as described above in Accessing the remote management web interface.
2. If you are using LO100, click on the KVM link in the "intermediate" page.
3. If you are using iLO or iL02, click on the **Remote Console** tab at the top.
4. If you are using iLO, click on the Remote Console link.
5. If you are using iLO2, click on the Integrated Remote Console link.

The server console will open in a new window.

NOTE:

Virtual KVM features are only available if you are licensed for ILO/LO100 Advanced Pack.

Monitoring a server reboot

To monitor a server reboot:

1. Log in to the remote management web interface as described above in Accessing the remote management web interface.
2. Click on the **Remote Console** tab at the top.
3. If you are using iLO, click on the Remote Console link.
4. If you are using iLO2, click on the Integrated Remote Console link.
5. When the server console opens in a new window, log in to the IAP server as `root`.
6. Enter the command: `reboot`
7. Monitor the boot sequence of the server in the console window.

Turning on a server UID light

To turn on a server UID light:

1. Log in to the remote management web interface as described above in Accessing the remote management web interface
2. If you are using iLO, click on the **Virtual Devices** tab at the top, click on the Virtual Indicators link, and click on the **Turn Unit ID On** button.

3. If you are using iLO2, click on the **System Status** tab at the top and click on the **Turn UID On** button.
4. If you are using LO100, click the Power Control Options link in the left menu, select one of the Chassis Locator options from the pull-down menu, and click on the **Identify** button.

Power cycling an unresponsive proxy server (PCC)

To power cycle an unresponsive proxy server (PCC):

1. Connect to the URL `https://<iLO IP address>` to access proxy server (PCC) iLO.
2. Log in to the remote management web interface using the user ID and password.
3. If you are using iLO, click on the **Virtual Devices** tab, and then **Virtual Power**.
4. If you are using iLO2, click on the **Power Management** tab at the top.
5. If you are using LO100, click the Power Control Options link in the left menu, select one of the Power Control Options from the pull-down menu, and click on the **Apply** button.
6. To power cycle the proxy server (PCC), select one of the power cycle options available for the iLO type.

Changing the administrative password on all iLOs

To change the administrative password on all iLOs:

1. Log in to the PCC server as `root` using SSH.
2. Enter the command:

   ```
   /opt/tools/ilo/configure_ilo -changepassword all
   ```

3. When prompted, enter the new password for the administrator account. The new password will be applied to all configured iLO Administrator accounts.

Updating iLO firmware

To update iLO firmware:

1. If you are using iLO and iLO2 devices, the firmware is automatically updated with the latest firmware shipped with the IAP whenever the iLO devices are configured. To update to a newer firmware than what the IAP shipped with:
   a. Extract the iLO bin file and place it under /opt/tools/ilo/firmware/<ilo|ilo2>
   b. Enter the command:

      ```
      /opt/tools/ilo/configure_ilo –query all
      ```

2. If you are using Lights-Out 100 devices, you must manually perform firmware updates:
   a. Create the `bin` file using the firmware package.
   b. Place the bin file on the PCC under `/tftpboot/firmware/lo100/<version>` omitting any periods in the version number.
   c. From the PCC command line, enter the command:

      ```
      /opt/tools/ilo/flashLO100 <iLO IP> <admin name> <admin password>
      ```

      For example, to update LO100 firmware for IP 10.1.172.1, enter the command:

      ```
      /opt/tools/ilo/flashLO100 10.1.172.1 Administrator <admin password>
      ```

      The iLO IP is generated by taking the first two octets of the iLO subnet (set in the REMOTE MANAGEMENT section in the `BlackBoxConfig.bct` file) and the last two octets of the host IP.

Configuring the iLO processor of a new or replacement server

Use the `configure_ilo` command from the PCC server to configure iLO on a new or replaced IAP server, create the Administrator user for iLO, and set the customer supplied password for that user:

1. Run `/opt/tools/ilo/configure_ilo –v –set –f ilo_ips.csv <IP Address>`
   <IP Address> is the new or replaced server's IAP IP address. For example, use `/opt/tools/ilo/configure_ilo –v –set –f ilo_ips.csv 10.0.172.11` to configure the iLO of the server whose IP address is 10.0.172.11.
2. Enter a password for the Administrator account when prompted.

The file `ilo_ips.csv` is a text file that must be current directory with the following contents:

<Internal IP of IAP server>,<IP of PCC iLO>,<iLO netmask>,<Gateway IP for iLO Subnet>

The following line is an example of the contents of the `ilo_ips.csv` file:

```
10.0.101.2,15.1.2.30,255.255.0.0,15.1.2.1
```

Using Direct mode

Prequisites

You must supply one IP address for each server to be configured with iLO access.

Configuring remote management

To configure remote management in Direct mode:

1. Edit the `/install/configs/primary/BlackBoxConfig.bct` file on the kickstart machine to add the following section:

   ```
   ########################

   ### REMOTE MANAGEMENT ###

   ########################

   iLO Access Mode: direct
   ```
2. Run `createBBC.zr`.
3. Run `converter.zr`
4. Run `regloader.pl –cv`
5. When run on the PCC server, the following command configures iLO on all the IAP servers, creates the Administrator user for iLO, and allows you to set the administrator supplied password for that user: `/opt/tools/ilo/configure_ilo –v –set –f ilo_ips.csv`

   The file `ilo_ips.csv` is a text file that must be in the directory from which the command is run. The file must contain the following entry for each server to be configured with iLO access, one per line:

   <Internal IP of IAP server>,<IP of iLO>,<iLO netmask>,<Gateway IP for iLO Subnet>

   Here is an example of the contents: `10.0.204.2,16.1.2.40,255.255.254.0,16.1.2.1`
6. When prompted, enter a password for the Administrator account.

NOTE:

The IPs of iLO (16.1.2.40) are addresses on the customer management network. One such IP address must be included in the `ilo_ips.csv` file for each server whose iLO needs to be configured.

Accessing the direct mode remote management web interface

To connect to an IAP server for remote management:

1. Connect to the URL `https://<Server iLO IP address>`
2. This opens the Remote Management login page for the login credentials. To gain access to the remote management web interface, enter the user ID and Password for remote management.

NOTE:

On the PCC Hardware Management page, the links in the Remote Management column display "Not Enabled" in Direct mode.

Accessing the command line interface

To access the remote management command line interface:

1. Enter the command `ssh Administrator@<iLO IP Address>` for the server to which iLO access is required.

   For example, use `ssh Administrator@16.1.2.40` to connect to the iLO IP 16.1.2.40.
2. When prompted, enter the password to gain access to the iLO command line interface of that server.

Using the command line interface to power cycle a server

To use the command line interface to power cycle a server:

1. Enter the command `ssh Administrator@<iLO IP Address>` for the server to which iLO access is required.
2. When prompted, enter the password to gain access to the iLO command line interface of that server.
3. If you have iLO or iLO2, at the `hpiLO->` prompt, enter: `power reset`
4. If you have LO100, press Enter at the `EMS->` prompt to get the `Lights-Out>` prompt and enter the command: `reset [WARM|COLD]`

Using the web interface to power cycle a server

To use the web interface to power cycle a server:

1. Log in to the remote management web interface as described above in Accessing the direct mode remote management web interface
2. If you are using iLO, click on the **Virtual Devices** tab, and then **Virtual Power**.
3. If you are using iLO2, click on the **Power Management** tab at the top.

4. If you are using LO100, in the left menu click on the Power Control Options link, and then select one of the "Power Control Options" from the pull-down menu and click the **Apply** button.
5. For the iLO type, select one of the power cycle options.

Accessing the KVM console

To access the KVM console:

1. Log in to the remote management web interface as described above in Accessing the direct mode remote management web interface
2. Click on the **Remote Console** tab at the top.
3. If you are using iLO, click on the Remote Console link.
4. If you are using iLO2, click on the Integrated Remote Console link.

The server console will open in a new window.

NOTE:

Virtual KVM features are only available if you are licensed for ILO/LO100 Advanced Pack.

Monitoring a server reboot

To monitor a server reboot:

1. Log in to the remote management web interface as described above in Accessing the direct mode remote management web interface
2. Click on the **Remote Console** tab at the top.
3. If you are using iLO, click on the Remote Console link.
4. If you are using iLO2, click on the Integrated Remote Console link.
5. When the server console opens in a new window, log in to the IAP server as `root`.
6. Enter the command: `reboot`
7. Monitor the boot sequence of the server in the console window.

Turning on a server UID light

To turn on a server UID light:

1. Log in to the remote management web interface as described above in Accessing the direct mode remote management web interface
2. If you are using iLO, click on the **Virtual Devices** tab at the top, click on the Virtual Indicators link, and click on the **Turn Unit ID On** button.
3. If you are using iLO2, click on the **System Status** tab at the top and click on the **Turn UID On** button.
4. If you are using LO100, click the Power Control Options link in the left menu, select one of the Chassis Locator options from the pull-down menu, and click on the **Identify** button.

Changing the administrative password on all iLOs

To change the administrative password on all iLOs:

1. Log in to the PCC server as `root` using SSH.

2. Enter the command:

   ```
   /opt/tools/ilo/configure_ilo -changepassword all
   ```

3. When prompted, enter the new password for the administrator account. The new password will be applied to all configured iLO Administrator accounts.

Updating iLO firmware

To update iLO firmware:

1. If you are using iLO and iLO2 devices, the firmware is automatically updated with the latest firmware shipped with the IAP whenever the iLO devices are configured. To update to a newer firmware than what the IAP shipped with:
   a. Extract the iLO bin file and place it under /opt/tools/ilo/firmware/<ilo|ilo2>
   b. Enter the command:

      ```
      /opt/tools/ilo/configure_ilo –query all
      ```

2. If you are using Lights-Out 100 devices, you must manually perform firmware updates:
   a. Create the `bin` file using the firmware package.
   b. Place the bin file on the PCC under `/tftpboot/firmware/lo100/<version>` omitting any periods in the version number.
   c. From the PCC command line, enter the command: `/opt/tools/ilo/flashLO100 <iLO IP> <admin name> <admin password>`

      For example, to update LO100 firmware for IP 16.1.2.40, enter the command: `/opt/tools/ilo/flashLO100 16.1.2.40 Administrator <administrator password>`

Configuring the iLO processor of a new or replacement server

Use the `configure_ilo` command from the PCC server to configure iLO on a new or replaced IAP server, create the Administrator user for iLO, and set the customer supplied password for that user:

1. Modify the file `/opt/tools/ilo/ilo_ips.csv` and add or edit the entry for the new server. Replacement servers do not need modification if they are using the same IP address.
2. Run `/opt/tools/ilo/configure_ilo –v –set –f ilo_ips.csv`
3. Enter a password for the Administrator account when prompted.

The file `ilo_ips.csv` is a text file that must be current directory with the following contents:

<Internal IP of IAP server>,<IP of iLO>,<iLO netmask>,<Gateway IP for iLO Subnet>

The following line is an example of the contents of the `ilo_ips.csv` file:

```
10.0.204.11,16.1.2.50,255.255.254.0,16.1.2.1
```

NOTE:

The IP of iLO (16.1.2.50) is an address on the customer management network.

A IAP application-generated alerts

The table below shows the application-generated alerts that appear in the IAP system.

These alerts notify the IAP administrator about significant problems that affect or can affect the functional behavior of the IAP.

Alerts appear on the PCC Overview page, in the Current Platform Alerts area at the top of the page. They can also be sent to external monitors or email recipients. (See "SNMP Management" on page 115.)

Alerts are logged as events in the persistent storage (postgres) of the PCC, and displayed among the other events on the Event Viewer page. (See "Event Viewer" on page 113.)

NOTE:

Some alerts are generated or cleared with a time delay of up to 30 minutes. (The exact delay depends on the alert specifics.)

Table 32 IAP application-generated alerts

ID	Description	Context	After PCC/ server restart	After manual removal from PCC Overview	Alert level
		System Infrastructure			
A010-01	The Smartcell is suspended because of a service failure	The reason is specified by various features	***After PCC restart:*** Alert is shown ***After server restart:*** Alert usually disappears, or appears again if Smartcell is suspended	Alert disappears and is not shown even if the problem still exists	Major
A010-03	The DB2 instance is unavailable	Could not create connection to DB2	***After PCC restart:*** PCC cannot start unless DB2 is up and running ***After server restart:*** After DB2 restarts the alert disappears	Alert disappears and is not shown even if the problem still exists	Critical

ID	Description	Context	After PCC/ server restart	After manual removal from PCC Overview	Alert level
A010-06	A DB2 transaction log file is missing	Found gap in DB2 transaction log: Missing log file <missing_tx_log>	***After PCC restart:*** Alert is shown ***After server restart:*** Alert is shown	Alert is shown	Critical
A010-07	The application server did not start successfully	JBoss failed to load all the services	***After PCC restart:*** Alert is shown ***After server restart:*** If the problem is not fixed, the alert appears again	Alert disappears and is not shown even if the problem still exists	Critical
A010-09	The MySQL instance on the Smartcell is unavailable	Could not create connection to MySQL	***After PCC restart:*** Alert is shown ***After server restart:*** After MySQL restarts the alert disappears	Alert disappears and is not shown even if the problem still exists	Critical
A010-11	The runtime environment detected an out-of-memory problem	Out of memory error on server <ip_address>	***After PCC restart:*** Alert is not shown ***After server restart:*** This is the recovery condition for Out of Memory alert. Alert is cleared when machine restarts.	Alert disappears	Critical
A010-12	The machine cannot be reached	The machine is not responding to '<command>' requests, where <command> is 'ping' or 'netcat at port 22'	***After PCC restart:*** Alert is shown ***After server restart:*** Alert disappears	Alert disappears and is not shown even if the problem still exists	Critical
A010-13	The application server on this machine is down	JBoss is down	***After PCC restart:*** Alert is shown ***After server restart:*** Alert disappears	Alert disappears and is not shown even if the problem still exists	Critical
A010-14	The Postgres instance on the Kickstart is unavailable	Could not create connection to Postgres (PCC persistent storage)	***After PCC restart:*** Alert is shown ***After server restart:*** Alert is shown	Alert disappears and is not shown even if the problem still exists	Critical

ID	Description	Context	After PCC/ server restart	After manual removal from PCC Overview	Alert level
A010-15	Free swap space is less than threshold limit	Free swap space is <free_swap>% (the threshold limit is: <threshold_free_swap>%)	***After PCC restart:*** Alert is shown ***After server restart:*** Alert is shown	Alert is shown	Minor
A010-16	Disk partition size has reached the low threshold limit (95%)	Disk Partition <disk_partition_ name> is <used_percentage> full. <used_absolute> is used of <total_size>	***After PCC restart:*** Alert is shown ***After server restart:*** Alert is shown again when next check for disk sizes is made and the disk size is still above the threshold	Alert is shown again when next check for disk size is made and the disk size is still above the threshold	Minor
A010-17	Disk partition size has reached the high threshold limit (99%)	Disk Partition <disk_partition_ name> is <used_percentage> full. <used_absolute> is used of <total_size>	***After PCC restart:*** Alert is shown ***After server restart:*** The alert is shown again when next check for disk sizes is made and the disk size is still above the threshold	Alert is shown again when next check for disk size is made and the disk size is still above the threshold	Major
		Archiving			
A020-01	The archiving process could not find any available Smartcell group	No archiver services are available for domain <domain>	***After PCC restart:*** Alert is shown ***After server restart:*** The alert disappears for a short period of time, and appears again	Alert disappears and is not shown even if the problem still exists	Critical
		Indexing			
A030-01	An index is corrupted	Index <index> is unavailable and needs repair. Search results might be incomplete.	***After PCC restart:*** Alert is shown ***After server restart:*** The alert disappears if the problem is recovered during Smartcell start-up	Alert disappears and is not shown even if the problem still exists	Major

ID	Description	Context	After PCC/ server restart	After manual removal from PCC Overview	Alert level
			Query Service		
A040-01	The query service cannot reach a SmartCell group	Group <group> is unavailable for query. Search results might be incomplete.	***After PCC restart:*** Alert is shown ***After server restart:*** Alert disappears if the problem is recovered during Smartcell restart	Alert disappears and is not shown even if the problem still exists	Major
			Replication		
A050-01	The replica SMTP server cannot be reached	Failed to replicate data because the connection to <smtp_server> failed	***After PCC restart:*** Alert is shown ***After server restart:*** Alert is shown	Alert is shown	Critical
A050-02	The replica SMTP server is busy	Failed to replicate data for group <group> in domain <domain> because the connection was closed by the replica SMTP server at <smtp_server>	***After PCC restart:*** Alert is shown ***After server restart:*** Alert is shown	Alert is shown	Major
A050-03	A connection to the replica SMTP server timed out	Failed to replicate data because connection to the replica SMTP server timed out at <smtp_server>	***After PCC restart:*** Alert is shown ***After server restart:*** Alert is shown	Alert is shown	Critical
A050-04	The replication process failed to send data	Failed to replicate data for group <group> in domain <domain>	***After PCC restart:*** Alert is shown ***After server restart:*** Alert is shown	Alert is shown	Critical
A050-05	The replica archiving service for a group on the replica system is unavailable	The replica archiving service for group <group> in domain <domain> is unavailable	***After PCC restart:*** Alert is shown ***After server restart:*** Alert is shown	Alert is shown	Critical
A050-06	The replication process failed to list files for a group	Could not get list of files to replicate for group <group> in domain <domain>	***After PCC restart:*** Alert is shown ***After server restart:*** Alert is shown	Alert is shown	Critical

ID	Description	Context	After PCC/ server restart	After manual removal from PCC Overview	Alert level
A050-07	The replication process cannot access data from the source group	Could not retrieve data to replicate for group <group> in domain <domain>	***After PCC restart:*** Alert is shown ***After server restart:*** Alert is shown	Alert is shown	Critical
A050-08	The replication process failed for a group	Could not replicate data for group <group> in domain <domain>	***After PCC restart:*** Alert is shown ***After server restart:*** Alert is shown	Alert is shown	Critical
A050-09	The replica archiving service for a group is unavailable	The replica archiving service for group <group> in domain <domain> is unavailable	***After PCC restart:*** Alert is shown ***After server restart:*** Alert is shown	Alert is shown	Critical
		Database Replication			
A051-01	A DB2 transaction log required by database replication is missing	A DB2 transaction log <missing_tx_log> required by DB2 replication could not be found. DB2 replication has been stopped and will have to be manually re-initialized.	***After PCC restart:*** Alert is shown ***After server restart:*** Alert is shown	Alert is shown	Critical
A051-02	An IPC file is missing	An IPC file <missing_ipc_ queue_file> required by DB2 replication (<program>) could not be found. This may cause DB2 replication to fail to stop when requested to do so.	***After PCC restart:*** Alert is shown ***After server restart:*** Alert is shown	Alert is shown	Minor
A051-03	DB2 replication has errors	DB2 replication is in error state. Details: <error_msg>	***After PCC restart:*** Alert is shown ***After server restart:*** Alert is shown	Alert is shown	Critical

ID	Description	Context	After PCC/ server restart	After manual removal from PCC Overview	Alert level
A051-04	The primary and replica databases are not synchronized	The primary and replica DB2 databases are not synchronized for more than <threshold_ latency> secs (Actual Latency: <actual_latency> secs). Details: <error_msg>	***After PCC restart:*** Alert is shown ***After server restart:*** Alert is shown	Alert is shown	Major
A051-05	DB2 replication has not captured updates	DB2 replication has not captured updates to the DB2 database for more than <threshold_ latency> secs (Actual Latency: <actual_latency> secs). Details: <error_msg>	***After PCC restart:*** Alert is shown ***After server restart:*** Alert is shown	Alert is shown	Major
User Management and User Synchronization					
A060-01	The dynamic account synchronization job for users failed	DAS job <das_job> failed to import users	***After PCC restart:*** Alert is shown ***After server restart:*** Alert is shown when the job is run again	Alert is shown when the job is run again	Critical
A060-02	The dynamic account synchronization job for groups failed	DAS job <das_job> failed to import groups	***After PCC restart:*** Alert is shown ***After server restart:*** Alert is shown when the job is run again	Alert is shown when the job is run again	Critical
A060-03	The dynamic account synchronization job for deleted objects failed	DAS job <das_job> failed to disable users and groups	***After PCC restart:*** Alert is shown ***After server restart:*** Alert is shown when the job is run again	Alert is shown when the job is run again	Critical
A060-04	The dynamic account synchronization job completed with errors	Check Account Error Recovery for any importing issues for starting point=<ldapdn>, Server=<ldap_ server>	***After PCC restart:*** Alert is shown ***After server restart:*** Alert is shown when the job is run again	Alert is shown when the job is run again	Major

ID	Description	Context	After PCC/ server restart	After manual removal from PCC Overview	Alert level
			Database Backup		
A070-01	The DB2 online backups are too old	DB2 online backups under <online_backup_ dir> are older than <max_days> days	***After PCC restart:*** Alert is shown ***After server restart:*** Alert is shown	Alert is shown	Major
A070-02	The DB2 online backup cannot be found.	No DB2 online backups could be found under <online_ backup_dir>	***After PCC restart:*** Alert is shown ***After server restart:*** Alert is shown	Alert is shown	Major
			Tape Backup		
A071-01	Tape restore failed. Manual intervention required.	Error allocating smart cell to restore to. Group ID: <group_id>	***After PCC restart:*** Alert is shown if the problem persists when the job is run again ***After server restart:*** Alert is shown if the problem persists when the job is run again	Alert is shown if the problem persists when the job is run again	Critical
A071-02	Tape restore failed. Manual intervention required.	Tape restore interrupted on smart cell.	***After PCC restart:*** Alert is shown ***After server restart:*** Alert is shown if the problem persists when the job is run again	Alert is shown if the problem persists when the job is run again	Critical
A071-03	Tape restore failed. Manual intervention required.	Tape restore step failed on smart cell.	***After PCC restart:*** Alert is shown ***After server restart:*** Alert is shown if the problem persists when the job is run again	Alert is shown if the problem persists when the job is run again	Critical

ID	Description	Context	After PCC/ server restart	After manual removal from PCC Overview	Alert level
			Administration		
A080-01	The Postgres instance on the PCC is unavailable	Could not create connection to Postgres (PCC persistent storage).	***After PCC restart:*** Alert is shown ***After server restart:*** After Postgres restarts the alert disappears	Alert disappears and is not shown even if the problem still exists	Critical
			Retention		
A090-01	Retention is misconfigured and will not run	The domain retention period of <domain_ret_ period> is below the minimum of <min_ret_period> for Domain <domain>	***After PCC restart:*** Alert is shown ***After server restart:*** Alert is shown if problem persists when the machine is restarted	Alert disappears and is not shown even if the problem still exists	Critical

B Indexed document types

In addition to email messages, the file types shown in the table below are indexed by the IAP. These files can be email attachments or files that are migrated to the IAP using HP File Archiving software.

Microsoft Access, Project, and Visio files are not indexed. If these files are archived, users can search for them only by using external identifying information, such as the file name or file extension.

Beginning with IAP 2.1, we have added file type and content-type support for the following applications: WordPerfect Office, WordPerfect Presentations, WordPerfect for DOS, Quattro Pro for Windows, and Quattro Pro for DOS.

See the "IAP overview" chapter in the IAP user guide for more detailed information on document indexing.

Table 33 IAP indexed document MIME types

File extension	File type	MIME content-type
.xml	XML document	text/xml
.txt	Plain text file; treated as ISO-8859-1 unless otherwise specified	text/plain
.htm, .html, .stm	HTML document	text/html, rtf/html
.rtf	Rich text format	rtf/text, application/rtf
.dat	TNEF (for Microsoft Exchange)	ms/tnef
.mht, .mhtml, .nws, .eml	Email message	message/RFC 822
.doc, .dot	Microsoft Word 97-2003 document	application/msword
.xla, .xlc, .xlm, .xls, .xlt, .xlw	Microsoft Excel 97-2003 document	application/vnd.ms-excel, application/ms-excel
.pot, .pps, .ppt	Microsoft PowerPoint 97-2003 document	application/vnd.ms-powerpoint, application/vnd.msppt
.pdf	Adobe Portable Document format	application/pdf
.zip	ZIP archive	application/zip
.docx	Microsoft Word 2007 document	application/vnd.openxmlformats-officedocument.wordprocessingml.document
.docm	Microsoft Word 2007 macro-enabled document	application/vnd.ms-word.document.macroEnabled.12

File extension	File type	MIME content-type
.dotx	Microsoft Word 2007 template	application/vnd.openxmlformats-officedocument.wordprocessingml.template
.dotm	Microsoft Word 2007 macro-enabled document template	application/vnd.ms-word.template.macroEnabled.12
.xlsx	Microsoft Excel 2007 workbook	application/vnd.openxmlformats-officedocument.spreadsheetml.sheet
.xlsm	Microsoft Excel 2007 macro-enabled workbook	application/vnd.ms-excel.sheet.macroEnabled.12
.xltx	Microsoft Excel 2007 template	application/vnd.openxmlformats-officedocument.spreadsheetml.template
.xltm	Microsoft Excel 2007 macro-enabled workbook template	application/vnd.ms-excel.template.macroEnabled.12
.xlam	Microsoft Excel 2007 add-in	application/vnd.ms-excel.addin.macroEnabled.12
.pptx	Microsoft PowerPoint 2007 presentation	application/vnd.openxmlformats-officedocument.presentationml.presentation
.pptm	Microsoft PowerPoint 2007 macro-enabled presentation	application/vnd.ms-powerpoint.presentation.macroEnabled.12
.ppsx	Microsoft PowerPoint 2007 slide show	application/vnd.openxmlformats-officedocument.presentationml.slideshow
.ppsm	Microsoft PowerPoint 2007 macro-enabled slide show	application/vnd.ms-powerpoint.slideshow.macroEnabled.12
.potx	Microsoft PowerPoint 2007 template	application/vnd.openxmlformats-officedocument.presentationml.template
.potm	Microsoft PowerPoint 2007 macro-enabled presentation template	application/vnd.ms-powerpoint.template.macroEnabled.12
.wpd	Corel WordPerfect for Windows – versions through Version 12.0, X3	application/wordperfect, application/wpd
.qpw, .wb1, .wb2, .wb3	Corel Quattro Pro for Windows – versions through Version 12.0, X3	application/qpw, application/wb1, application/wb2, application/wb3
.shw	Corel Presentations – versions through Version 12.0, X3	application/presentations

Index

Q

R

S

T

U

V

W