

# HP Integrated Archive Platform Administrator Guide

Version 2.0



P D F

Part number: PDF  
First edition: April 2008



**Legal and notice information**

© Copyright 2004-2008 Hewlett-Packard Development Company, L.P.

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

This document contains proprietary information, which is protected by copyright. No part of this document may be photocopied, reproduced, or translated into another language without the prior written consent of Hewlett-Packard. The information contained in this document is subject to change without notice. The only warranties for HP products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. HP shall not be liable for technical or editorial errors or omissions contained herein.

Microsoft® and Windows® are US registered trademarks of Microsoft Corporation. Outlook™ is a trademark of Microsoft Corporation.

Hewlett-Packard Company shall not be liable for technical or editorial errors or omissions contained herein. The information is provided "as is" without warranty of any kind and is subject to change without notice. The warranties for Hewlett-Packard Company products are set forth in the express limited warranty statements for such products. Nothing herein should be construed as constituting an additional warranty.

---

# Contents

<b>About this guide</b>	<b>11</b>
Intended audience	11
Related documentation	11
Document conventions and symbols	12
HP technical support	12
Subscription service	13
Other web sites	13
<b>1 IAP overview</b>	<b>15</b>
IAP power on/off	15
Power off	16
Power on	16
How to restart IAP after a power failure	16
Maximum file size	16
<b>2 Introduction to Platform Control Center (PCC)</b>	<b>17</b>
Accessing PCC	17
User interface components	17
User interface orientation tips	18
Views for common tasks	19
Updating views before printing	19
Left menu views	19
Monitoring and reporting	21
Statuses and states	21
Smart cell life cycle states	21
<b>3 System Status</b>	<b>23</b>
Overview	23
Events	23
Events features	24
Platform Performance	24
Account Manager Service	24
Partially indexed and CatchAll	24
Platform Statistics	25
Platform Statistics features	26
IAP Version	26
SMTP Flow Control	26
Storage Status	26
System Status	27
Platform Control	29
Platform Control view features	29
Starting, stopping, and restarting servers on the system	30
Performance Graph	30
Example: Platform Store graph	30
Example: System Monitoring graph	31
Creating performance graphs	31
API Configuration and Statistics	32
API Configuration	32
API Connections	33

<b>4 Configuration</b>	<b>35</b>
Platform Settings	35
Domain Configuration	35
Platform Settings	36
Firewall Settings	36
SSL Configuration	37
Available certificate signing requests	37
API port configuration	37
Creating a certificate signing request	37
Deleting a certificate signing request	38
Installing and generating a certificate on the PCC portal	38
Installing and generating a certificate on the HTTP portals	39
Software Version	40
Displaying Software Version	40
<b>5 Account Synchronization</b>	<b>41</b>
Account Synchronization overview	41
Creating and running DAS jobs	41
Creating LDAP server connections	41
Creating jobs	42
Mapping advanced options	43
Assigning HTTP portals	44
Starting, scheduling, and stopping DAS jobs	45
Editing or deleting jobs	45
Managing available HTTP portals	46
Editing or deleting available LDAP connections	46
Viewing DAS history logs	46
<b>6 Account Manager (AM)</b>	<b>47</b>
Account Manager overview	47
Account Manager view features	49
Managing user accounts	49
Adding a new user	49
Editing user information	50
User account information	50
Administrative delete	52
Enabling Administrative Delete	52
Granting "Delete Administration" privilege	53
Executing Administrative Deletion	53
Logging in Auditlog	53
Current Limitations	54
Managing groups	54
Managing repositories	54
Adding repositories	54
Editing repository information	54
Repository information	56
<b>7 Other user management features</b>	<b>57</b>
Manual Account Loader	57
Exporting user account information	57
Loading user account information	57
Error Recovery	58
Error Recovery features	58
Repairing synchronization errors	59
<b>8 Data management</b>	<b>61</b>
Replication	61

Database Replication . . . . .	62
(Re-)Initializing db2 replication . . . . .	62
Replication Status . . . . .	63
Data Replication Flow . . . . .	63
Cloning . . . . .	63
Cloning view features . . . . .	64
Cloning smart cells (copying data) . . . . .	65
Reprocessing . . . . .	65
Rescheduling all reprocessing schedules . . . . .	65
Editing reprocessing schedules . . . . .	66
Changing the reprocessing status . . . . .	66
Using the Reprocessing Utility . . . . .	66
Viewing reprocessing history logs . . . . .	67
Retention . . . . .	67
Searching for and editing a repository retention period . . . . .	68
Editing domain retention periods . . . . .	69
Changing the retention processing status . . . . .	69
Viewing retention history logs . . . . .	70
Setting the retention basis . . . . .	70
Database and data backup . . . . .	70
Backup file locations . . . . .	70
Restoring DB2 and master configuration file backups . . . . .	71
Restoring the master configuration files . . . . .	71
Duplicate Manager . . . . .	72
Duplicate Manager job schedules . . . . .	73
Scheduling a job . . . . .	73
Enabling or disabling a job . . . . .	73
Starting, pausing, or aborting a job . . . . .	73
Duplicate Manager job histories . . . . .	73
Folder Support . . . . .	74

## 9 Reporting . . . . . 75

Event Viewer . . . . .	75
Searching the Event Viewer . . . . .	75
Other Event Viewer features . . . . .	76
SNMP Management . . . . .	76
Downloading the IAP MIB . . . . .	76
Setting the SNMP server . . . . .	76
Selecting SNMP traps . . . . .	77
Receiving SNMP events by email . . . . .	78
Setting SNMP Community . . . . .	79
Email Reporter . . . . .	79
Detailed email reports . . . . .	79
Creating and scheduling email reports . . . . .	80
Logfile Sender . . . . .	80

## 10 External access . . . . . 81

Archive Gateway Management . . . . .	81
VNC Archive Gateway . . . . .	81
Overview Archive Gateway . . . . .	81
System Services . . . . .	82
Configured Tasks . . . . .	82
Journal Mining . . . . .	83
Selective Archiving . . . . .	84
Synchronize Deleted Items . . . . .	84
Tombstone Maintenance . . . . .	84
Troubleshooting . . . . .	85

<b>11 Audit Log</b>	<b>87</b>
Enabling the Audit Log feature	87
Granting user access to the Audit Log repository	87
Monitoring status	88
Setting Audit Log repository retention periods	88
<b>12 Backup system administration</b>	<b>91</b>
IAP backup strategy	91
Tivoli Storage Manager	91
Gaining access to the IAP backup server	91
Smart cell data backups	92
Separate Group Volumes	92
TSM backup terms	92
How IAP configures TSM	93
Adding and labeling new media (Web interface)	94
Adding and labeling new media (command line)	97
Restoring a smart cell	98
Preparing the backup server for disaster recovery	100
Things to back up	100
To perform a backup	100
Recovering the backup server	101
<b>Index</b>	<b>103</b>

---

# Figures

1	PCC user interface . . . . .	18
2	Performance Graph: Store Rate . . . . .	31
3	Performance Graph: Free Memory . . . . .	31
4	Domain Configuration . . . . .	36
5	New LDAP connection . . . . .	42
6	Create DAS job . . . . .	42
7	Mapping information . . . . .	43
8	Advanced options . . . . .	44
9	Assign a job to a portal . . . . .	45
10	Account Manager view . . . . .	48
11	Editing user account information . . . . .	50
12	Editing repository information . . . . .	55
13	Editing reprocessing schedules . . . . .	66
14	Change the reprocessing status . . . . .	66
15	Reprocessing utility . . . . .	67
16	User Repository search results . . . . .	68
17	Edit repository retention period . . . . .	68
18	Edit domain retention period . . . . .	69
19	Audit Log enabled . . . . .	88
20	Accessing domain repository . . . . .	88
21	Accessing Audit Log repository . . . . .	89
22	Policy domain structure . . . . .	93
23	Library properties . . . . .	95
24	Label and check in volumes . . . . .	96
25	Server process list . . . . .	97
26	Provisioner status . . . . .	99

---

# Tables

1	Document conventions . . . . .	12
2	Applications for users . . . . .	15
3	Applications for administrators . . . . .	15
4	Views for common system administration tasks . . . . .	19
5	Views accessible from left menu . . . . .	20
6	Smart cell life cycle states . . . . .	21
7	Link to Overview . . . . .	23
8	Events features . . . . .	24
9	Platform Statistics features . . . . .	26
10	Storage Status view features . . . . .	27
11	Link to Storage Status view . . . . .	27
12	System Status view features . . . . .	28
13	Link to System Status view . . . . .	29
14	Link to Platform Control view . . . . .	29
15	Platform Control view features . . . . .	29
16	Performance Graph features . . . . .	30
17	Link to Performance Graph view . . . . .	30
18	Link to API Configuration and Statistics . . . . .	32
19	Link to Platform Settings view . . . . .	35
20	Firewall ports . . . . .	36
21	Link to Firewall Settings view . . . . .	37
22	Link to SSL Configuration view . . . . .	37
23	Available certificate signing requests (CSRs) in the IAP . . . . .	37
24	Software Version view features . . . . .	40
25	Link to Software Version view . . . . .	40
26	Link to Account Synchronization view . . . . .	41
27	Link to Account Manager view . . . . .	47
28	Account Manager view features . . . . .	49
29	User account information . . . . .	50
30	Repository information . . . . .	56
31	Links to the Manual Account Loader view . . . . .	57
32	Links to Error Recovery view . . . . .	58
33	Error Recovery features . . . . .	58
34	Link to Replication view . . . . .	61
35	Database Replication features . . . . .	62



36	Replication Service General Status . . . . .	63
37	Data Replication Flow . . . . .	63
38	Link to Cloning view . . . . .	64
39	Cloning view features . . . . .	64
40	Link to Reprocessing view . . . . .	65
41	Link to Retention view . . . . .	67
42	DB Backup History . . . . .	70
43	Link to DB and Data Backup view . . . . .	70
44	Link to Duplicate Manager view . . . . .	72
45	Event Viewer features . . . . .	75
46	Links to Event Viewer . . . . .	75
47	Links to SNMP Management view . . . . .	76
48	Links to Email Reporter view . . . . .	79
49	Detailed Email Reports . . . . .	79
50	Links to LogFile Sender view . . . . .	80
51	Link to Archive Gateway Management view . . . . .	81
52	Overview Archive Gateway view features . . . . .	82
53	System Services features . . . . .	82
54	Configured Tasks features . . . . .	83
55	Journal Mining features . . . . .	83



---

# 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 EAs for Exchange or Domino administration guide (located on the documentation CD included in those products)
- HP EAs for Exchange or 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 Web Interface, also included in the above user guide
- HP EAs for Exchange or 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: <a href="#">Related documentation</a>	Cross-reference links and email addresses
Medium blue, underlined text ( <a href="http://www.hp.com">http://www.hp.com</a> )	Web site addresses
<b>Bold font</b>	<ul style="list-style-type: none"><li>• Key names</li><li>• Text typed into a GUI element, such as into a box</li><li>• GUI elements that are clicked or selected, such as menu and list items, buttons, and check boxes</li></ul>
<i>Italic font</i>	Text emphasis
Monospace font	<ul style="list-style-type: none"><li>• File and directory names</li><li>• System output</li><li>• Code</li><li>• Text typed at the command line</li></ul>
<i>Monospace, italic font</i>	<ul style="list-style-type: none"><li>• Code variables</li><li>• Command-line variables</li></ul>
<b>Monospace, bold font</b>	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.

---

## HP technical support

Telephone numbers for worldwide technical support are listed on the HP support web site: <http://www.hp.com/support/>.

Collect the following information before calling:

- Technical support registration number (if applicable)
- Product serial numbers
- Product model names and numbers
- Applicable error messages
- Operating system type and revision level
- Detailed, specific questions

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.

After subscribing, locate your products by selecting **Business support** and then **Storage** under Product Category.

## Other web sites

For other product information, see the following HP web sites:

- <http://www.hp.com>
- <http://www.hp.com/go/storage>
- [http://www.hp.com/service\\_locator](http://www.hp.com/service_locator)
- <http://www.hp.com/support/manuals>



# 1 IAP overview

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

IAP is a fault-tolerant, secure system of hardware and software that archives files and email messages for your organization, and lets you search for archived documents. IAP provides the following main functions:

- Automatic, active data archiving (email and specific document types) that helps your organization meet regulatory requirements.
- Interactive data querying to search for and retrieve archived data according to various criteria.

EAs is management software for IAP. To interact with the system, users can use the following applications:

**Table 2 Applications for users**

Application	Tasks
EAs for Exchange (customer option)	Search for emails using Microsoft Outlook with a Microsoft Exchange mail server. View and work with archived emails.
EAs for Domino (customer option)	Search for emails using IBM Lotus Notes with an IBM Domino mail server. View and work with archived emails.

IAP and EAs provide the following troubleshooting and administrative tools:

**Table 3 Applications for administrators**

Application	Tasks
Platform Control Center (PCC)	Monitor and troubleshoot system status and performance, and manage user accounts. See " <a href="#">Introduction to Platform Control Center (PCC)</a> " on page 17.
PST Importer	Batch process multiple PST files. See the <i>HP Email Archiving software for Exchange Administrator Guide</i> included with the EAs for Exchange product on the documentation CD.
Audit Log	Enable the Audit Log for regulatory compliance. See " <a href="#">Audit log</a> " on page 87.
EAs for Exchange	Create selective archiving rules for Microsoft Exchange and Outlook. See the <i>HP Email Archiving software for Exchange Administrator Guide</i> included with the EAs for Exchange product on the documentation CD.
EAs for Domino	Create selective archiving rules for Domino. See the <i>HP EAs for Domino Administrator Guide</i> included with the EAs for Domino product on the documentation CD.
Web UI	The Web UI allows administrators and users to use their web browser to search for documents archived on the system, and save and reuse your search-query definitions and results. See the <i>IAP User Guide</i> .

## 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 PCC enter:

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

Wait a few minutes until the PCC console shutdown is complete before removing power from the IAP systems.

## Power on

To power on the IAP:

1. Make sure the IAP switch(es) are powered up. Once power is restored, the 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).

## 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, and loadbalancers.
6. Power on smart cells.
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.



---

# 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:

- [Accessing PCC](#), page 17
- [User interface components](#), page 17
- [User interface orientation tips](#), page 18
- [Views for common tasks](#), page 19
- [Updating views before printing](#), page 19
- [Left menu views](#), page 19
- [Monitoring and reporting](#), page 21
- [Statuses and states](#), page 21

## 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 50 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.

## User interface components

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 views in the PCC.

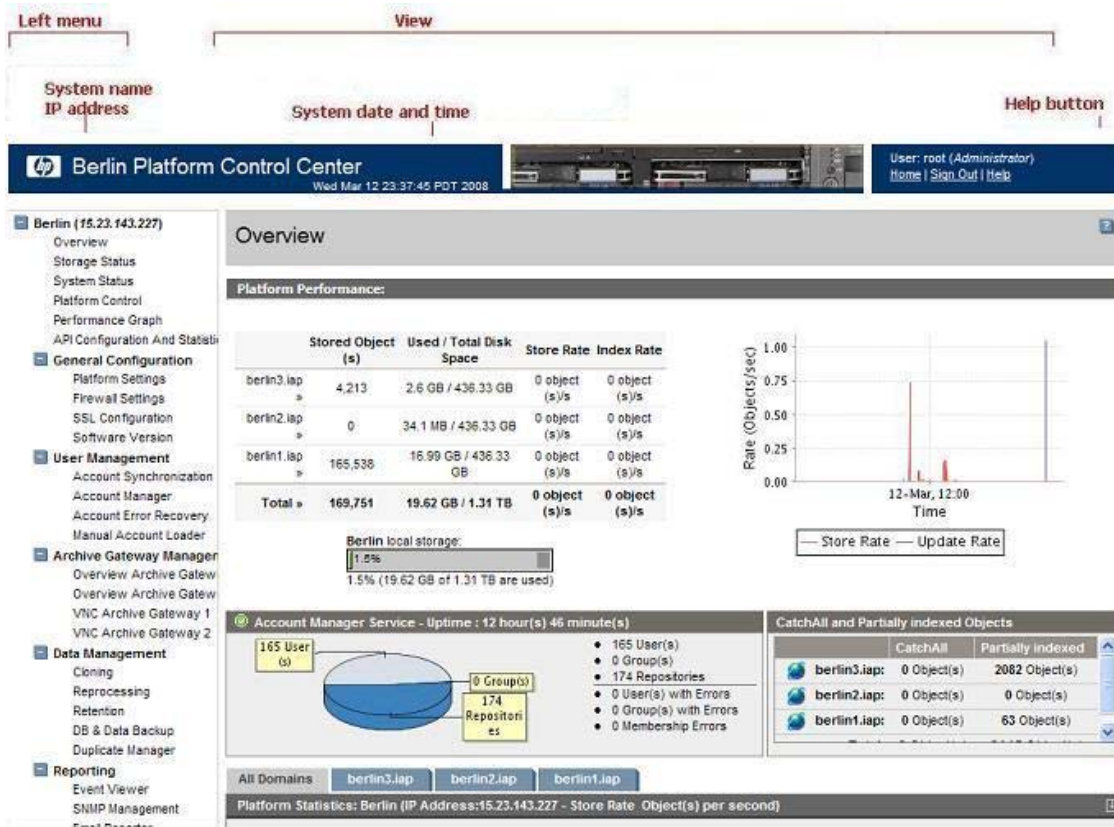


Figure 1 PCC user interface

## User interface orientation tips

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

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

# Views for common tasks

**Table 4 Views for common system administration tasks**

Task	View
Check overall system health and performance	" <a href="#">Overview</a> " on page 23
Check smart cell health and performance	" <a href="#">Platform Statistics</a> " on page 25
Monitor system status and RAID support	" <a href="#">System Status</a> " on page 27
Start, stop, and restart system servers	" <a href="#">Platform Control</a> " on page 29
Check the platform configuration	" <a href="#">Platform Settings</a> " on page 35
Display firewalled ports enabled in the system	" <a href="#">Firewall Settings</a> " on page 36
View software versions used by system hosts	" <a href="#">Software Version</a> " on page 40
Synchronize user accounts	" <a href="#">Account Synchronization</a> " on page 41
Manage user accounts	" <a href="#">Account Manager (AM)</a> " on page 47
Monitor, start, and stop replication for domains	" <a href="#">Replication</a> " on page 61
Clone smart cells (copy data)	" <a href="#">Cloning</a> " on page 63
Check database backup history	" <a href="#">Database and data backup</a> " on page 70
Monitor system alerts	" <a href="#">Event Viewer</a> " on page 75
Activate SNMP traps and send email notifications	" <a href="#">SNMP Management</a> " on page 76
Configure periodic email reports of system status and performance	" <a href="#">Email Reporter</a> " on page 79
Link to archive gateway services	" <a href="#">Archive Gateway Management</a> " on page 81

## Updating views before printing

PCC views displayed in the web browser are not automatically updated. To manually update the view, click **Refresh** (or **Reload**) in the browser.

If the browser caches web pages, the cached view 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 view is out of date, the printout can appear different from the displayed view. To ensure you print what is displayed, refresh the browser manually before printing.

## Left menu views

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.

**Table 5 Views accessible from left menu**

Left menu item	Description
"Overview" on page 23	View summary of system health, storage status, smart cell performance by domain, and system alerts and warnings.
"Storage Status" on page 26	View summary, by domain, of document storage rates and used/free disk space.
"System Status" on page 27	View summary, by server, of system capacity and performance.
"Platform Control" on page 29	Start, stop, or restart one or more servers on the system.
"Performance Graph" on page 30	Graph system storage and indexing rates and system performance.
"API Configuration and Statistics" on page 32	View API configuration and connections.
<b>General Configuration views</b>	
"Platform Settings" on page 35	Display hardware and configuration information about the IAP.
"Firewall Settings" on page 36	Display each firewalled port that is enabled in the system.
"SSL Configuration" on page 37	Generate third party public certificate requests for the PCC and HTTP portals.
"Software Version" on page 40	View software versions used by system hosts.
<b>User Management views</b>	
"Account Synchronization" on page 41	Configure dynamic account synchronization (DAS) to automatically create and update IAP users with information obtained from LDAP servers.
"Account Manager (AM)" on page 47	Provision, update, and manage IAP user accounts.
"Account Error Recovery" on page 58	Repair account synchronization errors.
"Manual Account Loader" on page 57	Create and update IAP users if the server is not using LDAP.
<b>Archive Gateway Management views</b>	
"Overview Archive Gateway" on page 81	View status information about the archive gateway for each domain.
"VNC Archive Gateway" on page 81	Access the archive gateway using VNC.
<b>Data Management views</b>	
"Replication" on page 61	Monitor and start or stop replication for domains.
"Cloning" on page 63	Clone smart cells (copy data) to give them a new, viable mirror cell.
"Reprocessing" on page 65	Schedule and enable reprocessing based on new routing rules.
"Retention" on page 67	Configure retention periods for domains and repositories.
"Database and data backup" on page 70	View status information about database and configuration file backup, including results of the last two backups.
"Duplicate Manager" on page 72	View status of duplicate merge jobs and schedule duplicate merge jobs.
<b>Reporting views</b>	
"Event Viewer" on page 75	View events with a critical or recovery status that have occurred in a system service or application.
"SNMP Management" on page 76	Set SNMP traps for system monitoring and send email notifications.

Left menu item	Description
"Email Reporter" on page 79	Configure system monitoring reports to be sent to email recipients.
"LogFile Sender" on page 80	Send output and error log file reports, by machine type, to email recipients.

## Monitoring and reporting

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

- system health
- system performance
- smart cell states

Hosts in the system (and their services) are organized into groups of the same type, called host groups. For example, to view all smart cell hosts, display the status of the host group SMARTCELL Servers in the System Status view.

As long as services appear to be functioning correctly (OK), the host is assumed to be healthy (UP). If monitoring indicates a host is not functioning correctly (DOWN), none of its services are available (they can have any status except OK). If a service has CRITICAL status, but the host is UP, the service probably needs to be restarted.

## Statuses and states




Several PCC views show current life cycle states of smart cells 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 smart cell 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 views often use *status* and *state* loosely and interchangeably when referring to hosts and services. State is always used when referring to smart cell life cycle states, but status and state are both used when reporting smart cell health, since smart cells are regarded as a host like any other. PCC views also refer to status conditions as states or state types.

## Smart cell life cycle states

**Table 6 Smart cell life cycle states**

Life cycle state	Definition	Importance
ASSIGNED 	The cell is assigned to a domain. The cell is available for document storage, search, and retrieval. If backup is enabled, cell data can be backed up.	normal
CLOSED 	The cell is full. It is available for document search and retrieval, but not storage. If backup is enabled, all cell data was backed up before the cell entered this state.	normal
COMPLETE_PROCESSING 	Data indexing is being completed. The cell is full. It is available for document search and retrieval, but not storage. If backup is enabled, cell data can be backed up.	maintenance

Life cycle state	Definition	Importance
BACKING_UP 	The cell is available for document search and retrieval. If backup is enabled, the cell is backing up all its indexes and new data that has not yet been backed up.	maintenance
SYNC_WAIT 	The cell is available for document search and retrieval.	maintenance
RESTORE 	The cell is a target for data restoration from another smart cell. The cell is not available for document storage, search, or retrieval.	maintenance
DISCOVERY 	The meta server and smart cell are determining the cell's start state (the state following DISCOVERY), based on expected states of the cell and its mirror smart cell. The cell is not available for document storage, search, or retrieval.	maintenance (startup only)
RESET 	The cell is being recycled. Stored documents and corresponding management data, such as document indexes, are destroyed during recycling. The system administrator has determined that existing cell data is no longer needed. The RESET state is only set manually. The cell is not affiliated with any domain, so it is not available for document storage, search, or retrieval.	maintenance
SUSPENDED 	Either of the following is true: <ul style="list-style-type: none"> <li>The cell or its mirror cell has one or more failed processes.</li> <li>The mirror cell is DEAD.</li> </ul> <b>NOTE:</b> If the cell's status is OK, only the mirror cell has failed. The cell 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 cell is backing up new data that has not yet been backed up.	failure
DEAD 	The cell has failed. It is not available for document storage, search, or retrieval. If backup is enabled, some or all cell data might <i>not</i> be backed up; if so, data will never be backed up.	failure
UNKNOWN 	The state of the smart cell is unknown.	unknown
FREE	The cell is free. (Shown in blue.) It can become ASSIGNED or become a target for data restoration. The cell 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 views.

It includes the following topics:

- [Overview](#), page 23
- [Storage Status](#), page 26
- [System Status](#), page 27
- [Platform Control](#), page 29
- [Performance Graph](#), page 30
- [API Configuration and Statistics](#), page 32

## Overview

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

- Critical events that are occurring in a system service or application.
- Information about document storage rates and capacity, for the system and for each domain.
- A summary of the number of IAP users, groups, and repositories in the system.
- A summary of the number of partially indexed objects and CatchAll repository objects.
- Information, by domain, about the status, health, and storage rate of each smart cell.
- Information about the IAP software version.
- The number of SMTP connections in each domain.



### NOTE:

Typically, you would monitor the Overview every day.

---

**Table 7 Link to Overview**

Origin	Link
left menu	Overview

## Events

The Events log at the top of the Overview displays the critical events that are currently occurring in system services or applications.

Clicking **More Details** takes you to the Event Viewer, where you can search for events by type and time period. See "[Event Viewer](#)" on page 75 for more information.




### NOTE:

The Events log does not appear in the Overview if critical events are not currently occurring. You can view previous system events by navigating to **Reporting > Event Viewer**.

---

## Events features

**Table 8 Events features**

Feature	Description
Event	Information describing the event or error, including the service or application name.
Machine	The name of the server on which the event is occurring.
IP	The IP address of the server on which the event is occurring.
Date	The date of the event.
Level	The status of the event. In Events, the only status shown is  <i>critical</i> .

## Platform Performance

This area of the overview provides information from the Storage Status view (see “[Storage Status](#)” on page 26). 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 47). This area displays the number of individual IAP users and groups, pending users and groups, and the number of IAP repositories. It also displays the number of synchronization errors, if any, pertaining to IAP user accounts. Synchronization errors can be corrected in the Error Recovery view (see “[Account Error Recovery](#)” on page 58).

## Partially indexed and CatchAll

This area of the Overview displays the following information:

- **Partially indexed:** The number of documents that were partially indexed. (For example, the system might not have been able to index the particular file type.) These documents are in the failed indexing repository, which can be viewed in the Account Manager (see “[Account Manager \(AM\)](#)” on page 47). Click the **Repository** radio button in AM, click the **Other** tab, then open the failed indexing repository.
- **CatchAll:** The number of messages in the system’s catchall repository, including messages too large to be indexed, messages that cannot be parsed, and messages that cannot be routed to a registered IAP user. The number includes messages that went into catchall from the SMTP portals as well as those from the smart cell indexers.  
Messages that cannot be parsed have malformed message structures (MIME) or unsupported character sets.  
Messages that cannot be routed do not correspond to any system routing rule. They are not recognized as destined for a registered IAP user. 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 47 by clicking the **Repository** radio button in AM, clicking the **Other** tab, then opening the catch-all repository.

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



**NOTE:**








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

---

## Platform Statistics





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

Each smart cell's life cycle state is color-coded.

-  A green table row indicates a smart cell is ASSIGNED.
-  A light green table row indicates smart cell is CLOSED.
-  A yellow-orange table row indicates a smart cell is in the COMPLETE\_PROCESSING, SYNC-WAIT, BACKING\_UP, or RESTORE state.
-  A light yellow table row indicates a smart cell is in the DISCOVERY or RESET state.
-  A red table row indicates a smart cell is SUSPENDED.
-  A black table row indicates a smart cell is DEAD.
-  A gray table row indicates the state of a smart cell is UNKNOWN.

See "[Smart cell life cycle states](#)" on page 21 for more information.

The icon at the beginning of the table row displays the health of the IAP operating system and applications on the smart cell:

-  A green check icon indicates the smart cell server is started and healthy.
-  A gray icon indicates that JBoss and the IAP applications have stopped.
-  A yellow icon indicates that JBoss is running, but one or more IAP applications are unhealthy.
-  A red X icon indicates that JBoss is running, but one or more IAP applications have failed.

**TIP:**

If you move your mouse over the icon, you will see more information including node status, active thread count, and the smart cells's MAC address and IP 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 smart-cell group identifier generated automatically by IAP. This number is unique across all systems.
Smartcell IP	The IP address of the smart cell.
Smartcell Role	A smart cell can be Primary, Secondary, Replica-1, or Replica-2.
State	The current life cycle state of the smart cell. See <a href="#">“Smart cell life cycle states”</a> on page 21.
Stored Object(s)	The number of objects stored since the smart cell was assigned. <b>NOTE:</b> When the system is actively storing objects, this count might be different than the stored object count in <a href="#">“Platform performance”</a> on page 24. This number is the real-time count on the smart cell, 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 smart cell was assigned.
Store Rate	The number of objects being stored per second.
Index Rate	The number of documents being indexed per second.
Index Deletion Queue	The number of documents scheduled for deletion.
Update Queue	The number of documents scheduled for update.
Other Smart Cells	IP addresses of smart cells in the FREE state.

## IAP Version

Near the bottom of the Overview, you will find information about the IAP base version (also known as L2) and software version (also known as L3).

## SMTP Flow Control


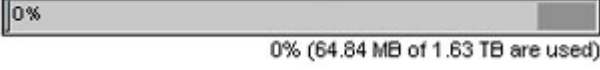
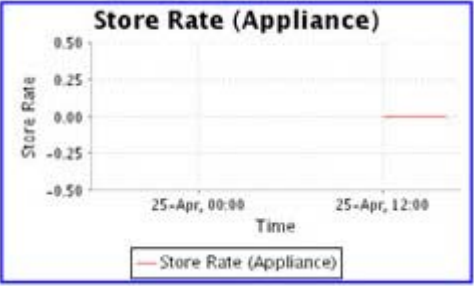
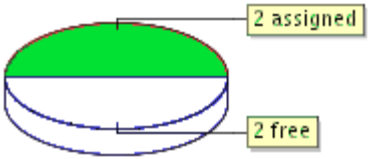
At the bottom of the Overview, the SMTP Flow Control area shows the following information, by domain:

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

## Storage Status

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

**Table 10 Storage Status view features**

Feature	Description
Platform Store	<p>The number of objects and store rate per domain, and the allocated space on the system for storage and replication.</p> <p>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.</p> <p>Enigma local allocated storage: </p>  <p><b>NOTE:</b></p> <p>The storage bar graph shows only assigned and allocated smart cells for all active domains. The IAP might have free, unallocated hardware that is not represented in the bar graph.</p>
Store Rate Graph	<p>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.</p> <p>Example:</p> 
Smart Cell Allocation	<p>The number of smart cells in each life cycle state. (See <a href="#">“Smart cell life cycle states”</a> on page 21 for an explanation of each state.)</p> <p>Example:</p> 
Domain Details	<p>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.</p>










**Table 11 Link to Storage Status view**

Origin	Link
left menu	Storage Status

## System Status

The System Status view provides hardware and system information for all hosts in the system.

**Table 12 System Status view features**

Feature	Description
Status	<p>The icon in front of the host name displays the status of the host machine.</p> <ul style="list-style-type: none"> <li>•  A green check icon indicates the server is started and healthy.</li> <li>•  A gray icon indicates that JBoss and the IAP applications on the server have stopped.</li> <li>•  A yellow icon indicates that JBoss is running, but one or more IAP applications on the server are unhealthy.</li> <li>•  A red X icon indicates that the host is either down, or that JBoss is running, but one or more IAP applications on the server have failed.</li> <li>•  A blue question mark icon indicates that the state of the server is unknown.</li> </ul>
Hostname	<p>The type of host group, the servers in the group, the full name of each server, and the status of each server.</p> <p>Host group types include the following:</p> <ul style="list-style-type: none"> <li>• SMARTCELLS Servers: Smart cell servers</li> <li>• META Servers: Meta servers</li> <li>• SMTP Servers: SMTP portal servers</li> <li>• DB Servers: DB2 database servers</li> <li>• ARCHIVEGATEWAY Servers: Email mining servers</li> <li>• HTTP Servers: HTTP portal servers</li> <li>• PCC Servers: Platform Control Center servers</li> <li>• KICKSTART Servers: Kickstart servers.</li> <li>• CLOUDROUTER Servers: Cloud router servers</li> <li>• LOADBALANCER Servers: Load balancer servers</li> <li>• FIREWALL Servers: Firewall servers</li> </ul>
Battery Backed Write Cache	<p>Battery Backed Write Cache support and status for the machine.</p>
RAID Controller	<p>RAID support and status for the machine.</p> <p>The following type of messages and status may be displayed:</p> <ul style="list-style-type: none"> <li>•  FAILED: LD 1 (136 GB, RAID 1+0, Interim Recovery Mode)</li> <li>•  REBUILD: LD 1 (136 GB, RAID 1+0, Rebuilding 1%). (Indicates a drive is rebuilding.)</li> <li>•  OK: No SmartArray RAID controller to check.</li> <li>•  OK: LD 1 (136 GB, RAID 1+0, OK)</li> </ul>
Memory	<p>The machine's memory use.</p> <p><b>TIP:</b></p> <p>If you move your mouse over the entry's performance bar, you will see more information including maximum memory (the physical memory available), total memory (the memory allocated to the IAP applications), and the free and used memory (for the memory allocated to the IAP applications).</p>
Processor	<p>The machine's CPU use.</p> <p><b>TIP:</b></p> <p>If you move your mouse over the entry's performance bar, you will see the actual percentage of CPU use.</p>
Used Threads	<p>The number of threads that are being used by the machine.</p>

**Table 13 Link to System Status view**

Origin	Link
left menu	System Status

## Platform Control

Use the Platform Control view to start, stop, or restart one or more servers on the system.

This view is useful to show the start, stop, and pending status of a server. However, you should use it only when necessary — for example, when you are upgrading a host or before a planned power outage. The Platform Control view should be used only by system administrators or HP service representatives.






**Table 14 Link to Platform Control view**

Origin	Link
left menu	Platform Control

## Platform Control view features

The following information is displayed in the Platform Control view.

**Table 15 Platform Control view features**

Feature	Description
Hostname	<p>The type of host group, the servers in the group, and the full name of each server. Host group types include the following:</p> <ul style="list-style-type: none"> <li>• LOAD BALANCER Servers: Load balancer servers</li> <li>• CLOUD ROUTER Servers: Cloud router servers</li> <li>• SMARTCELL Servers: Smart cell servers</li> <li>• META Servers: Meta servers</li> <li>• SMTP Servers: SMTP portal servers</li> <li>• DB Servers: DB2 database servers</li> <li>• ARCHIVEGATEWAY Servers: Email mining servers</li> <li>• HTTP Servers: HTTP portal servers</li> <li>• PCC Servers: Platform Control Center servers</li> </ul>
Status	<p>The icon in front of the host name displays the status of the host machine.</p> <ul style="list-style-type: none"> <li>•  A green check icon indicates the server is started and healthy.</li> <li>•  A gray icon indicates the server is stopped.</li> <li>•  A yellow icon indicates the server is starting or stopping.</li> <li>•  A red X icon indicates the server is down or an action has failed.</li> <li>•  A blue question mark icon indicates the state of the server is unknown.</li> </ul>
Host IP Address	The host's IP address.
MAC Address	The host's MAC address.
Other Details	Other details about the host.

## Starting, stopping, and restarting servers on the system

1. 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.

---

 **NOTE:**

Staggered Restart can only be used with smart cell, HTTP, and meta servers.

---

2. To perform an action on all machines in a server group:
  - a. Click the **Machine Type** radio button.
  - b. Select the type of server from the Machine Type drop-down list.
3. To perform an action on a particular machine:
  - a. Click the **Machine** radio button.
  - b. Select the machine in the Hostname column.
4. Click **Run Now!**

## Performance Graph

Use this view 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 16 Performance Graph features**


Feature	Description
Heading	The type of graph or the server name.
Event History	The graphed event history over the reported time period.
Time	The time period being reported.

**Table 17 Link to Performance Graph view**

Origin	Link
left menu	Performance Graph

### 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 2 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 3 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 System Monitoring, select for example **Idle CPU Usage**, **Free Memory**, or **Active Thread Count**.
  - For Appliance Store and Indexing, select for example **Store Rate**, **Index Rate**, **Object Count**, or **Index Count**.

3. Select one of the following options:
  - Machine Type (System Monitoring graphs): Select the type of machine (for example, PCC Servers or Smart Cell Servers).
  - Assigned Smartcell/Domain (Platform Store and Indexing graphs): Select **Entire Platform**, **Domain name**, or **Smart Cell 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**.

## API Configuration and Statistics

The API Configuration and Statistics view gives an overview of how API access to the IAP is configured. It also reports how many clients are currently connected to each portal.

**Table 18 Link to API Configuration and Statistics**

Origin	Link
left menu	API Configuration and Statistics

## API Configuration

Information about the API configuration attributes is in the upper portion of the API Configuration and Statistics view. The attributes are described in the following table:

### API Configuration

Attribute	Value
Secure Port	Indicates whether the encrypted (SSL) port is enabled or disabled. Also shows the port number.
Clear Text Port	Indicates whether clear text port is enabled or disabled. Also shows the port number.
Maximum connections per portal	Shows the maximum number of API client connections supported per portal. For the number of API portals, refer to the table "API Connections" below.
API Server Version	Version information for the API server component.
API Wire Protocol Version	Information for the highest API client/server protocol version supported by the server component.

### NOTE:

Only one of the two ports can be enabled at a time. To select secure or clear text connection, refer to the PCC **General Configuration > SSL Configuration** section "API Port Configuration."



## API Connections

Information about the API connections is displayed in the lower portion of the API Configuration and Statistics view. The API Connections section shows the API portals and clients connected to them. The attributes are described in the following table:

### API Connections

Attribute	Value
Domain	The IAP domain identifier for the domain to which the client is connected.
Username	The API client user name used to establish the connection.
Connection Time	The time the API client established the connection.
Client Version	API client version information.
Protocol Version	API client/server protocol version used for active connection.



---

# 4 Configuration

This chapter contains the following information:

- [Platform Settings](#), page 35
- [Firewall Settings](#), page 36
- [SSL Configuration](#), page 37
- [Software Version](#), page 40

## Platform Settings

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

This view is divided into two parts:

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

**Table 19** [Link to Platform Settings view](#)

Origin	Link
left menu	General Configuration > Platform Settings

## 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.jcm1` file on the kickstart server. See the Domain Configuration figure below.

Domain Name: thumb.com	
Virtual IPs and Type:	15.186.249.66 can store from MS Exchange
Supported Object Type:	Emails -  Documents
LDAP User Management:	<input checked="" type="checkbox"/> Enabled
Indexing Service:	<input checked="" type="checkbox"/> Enabled
Replication Service:	<input checked="" type="checkbox"/> Disabled
Default Domain Retention Period:	—
RetentionBasis:	IngestDate
Duplicate Filtering:	<input checked="" type="checkbox"/> Enabled
AuditLog Service:	<input checked="" type="checkbox"/> Enabled
Data Backup:	<input checked="" type="checkbox"/> Disabled
Admin Delete Service:	<input checked="" type="checkbox"/> Enabled
Folder Support:	<input checked="" type="checkbox"/> Enabled

**Figure 4 Domain Configuration**

## Platform Settings

The lower portion of the Platform Settings view displays the setup details for the IAP. This information is taken from the BlackBoxConfig.bct file.

## Firewall Settings

The Firewall Settings view 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 20 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 <input checked="" type="checkbox"/> enabled or <input checked="" type="checkbox"/> disabled on the port's inbound and outbound traffic. Outbound traffic is not filtered on the PCC server ports.

**Table 21 Link to Firewall Settings view**

Origin	Link
left menu	General Configuration > Firewall Settings

## 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 view 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 38 and [“Installing a third party certificate on the HTTP portals”](#) on page 39 for the steps needed to complete the process.

The SSL Configuration view 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.

**Table 22 Link to SSL Configuration view**

Origin	Link
left menu	General Configuration > SSL Configuration

## Available certificate signing requests

**Table 23 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.

## API port configuration

There is also an API Port Configuration section which allows you to select one of two connections for accessing the IAP API. The IAP API can be accessed: either via an encrypted **Secure Connection (SSL)** or by using a **Clear Text Connection** (unencrypted). To change between the two options, select the preferred configuration and click the **Apply** button. Please note that using the clear text connection could impose a security risk as the network communication between the IAP and client applications is not protected this way.

## Creating a certificate signing request

To create a certificate signing request (CSR):

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

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 38.

To install a certificate on each HTTP portal, see [“Installing a third party certificate on the HTTP portals”](#) on page 39.

---

 **NOTE:**

If you enter incorrect information in the form and need to generate a new CSR, see [“Deleting a certificate signing request”](#) on page 38 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 view.

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.

---

 **NOTE:**

**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/pccCertSigned.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. Import the certificate into the Apache server on each HTTP portal:

```
usr/local/bin/ssl_cert_update.pl -http -cert /opt/keys/httpCert-Signed.pem -key /opt/keys/httpkey.pem
```

6. From the PCC console, 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 “[Platform Control](#)” on page 29.

## Software Version

This view shows the software versions that are used by hosts in the system, and any patches that have been installed.

**Table 24 Software Version view features**

Feature	Description
Base Version	IAP and operating system software on Linux servers, also called L2 or spine.
Software Version	Third-party software package, also called L3.
Installer Version	IAP installation program.
Patches Applied	History of IAP software patches applied to the system.

**Table 25 Link to Software Version view**

Origin	Link
left menu	General Configuration > Software Version

## Displaying Software Version

You can display the software versions by host type:

1. Select the type of host in the HostType drop-down list.  
To display all software versions of machines in the system, select **System** as HostType.
2. Click **Update**.



# 5 Account Synchronization

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

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

- [Account Synchronization overview](#), page 41
- [Creating and running DAS jobs](#), page 41
- [Editing or deleting jobs](#), page 45
- [Managing available HTTP portals](#), page 46
- [Editing or deleting available LDAP connections](#), page 46
- [Viewing DAS history logs](#), page 46

**Table 26 Link to Account Synchronization view**

Origin	Link
left menu	User Management > Account Synchronization

## Account Synchronization overview

The Account Synchronization view is divided into three sections.

- The DAS Available Jobs section lists all jobs created and assigned to an HTTP portal.
- The LDAP Server Connectors section lists available LDAP connections.
- The Jobs History Logs section 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 41.
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 42.
3. Assign the job to an HTTP portal. See [“Assigning HTTP portals”](#) on page 44.
4. Run the job. See [“Running DAS jobs”](#) on page 45.

## 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:

Connection Name	LDAP connection
Host Name	10.1.1.1
Binder user	cn=Administrator
Binder pswd	•••••
Directory Server type	Microsoft Active Directory
Security Option	Simple LDAP
Port	389

**Figure 5 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 view 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**.

Job Name:   
*(the field cannot be blank or contain a "@ \$ % ^ & \* # ( ) [ ] \ { + } ` ~ = - | \" character.)*

**Figure 6 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 41 for further instructions.

4. Click **Next Step**.

User Management / Dynamic Account Synchronization

Configuration Wizard: Mapping Information

Parameter	Value
Job ID	berlin3_Domino
LDAP Domain Name	Germany
LDAP Job Starting Point	
IAP Domain ID	berlin3.iap
Delete Starting Point	N/A

Advanced Options: 

**Figure 7 Mapping information**

5. Complete the form by entering the following:
  - LDAP Domain name: Domain to which the users belong. For example: `ldaptest.com`.
  - LDAP 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`.
  - Deletion 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 43.

## Mapping advanced options

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

Group Name	<input type="text"/>
USNChanged	10000000
Delete USNChanged	10000000
NextRepID	5
Audit Repository	R0000000
Update LDAP filter	(objectclass=user)(mail=*)
LDAP Query return attributes	objectGUID,userPrincipalName,cn,givenName,mail,s n,whenChanged,whenCreated,distinguishedName,pr oxyaddresses,uSNChanged,userAccountControl,ma nager,extensionAttribute15,sAMAccountName
Delete LDAP Filter	(&!(objectclass=group)(objectclass=user))(isDeleted=*)

**Figure 8 Advanced options**

- Complete the advanced options form by entering the following information:
  - Group Name: Not used at this time.
  - 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.
- Click **Next Step** to create the job.
- To assign this job to an HTTP portal, click **Assign Job**.  
See "[Assigning HTTP portals](#)" on page 44 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**.

Job Name	New_Job
DAS server IP	10.0.71.2
Configuration Enabled	<input checked="" type="radio"/> Yes <input type="radio"/> No
Configuration running state	0
Period (minutes)	0
DAS server running state	

**Figure 9 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 the DAS Available Jobs view, 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 anything 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 view.

To stop a DAS job:

1. In the DAS Available Jobs view, select the job and click **Start/Stop**.  
A notice appears stating that the job has stopped.
2. Click **Back to Main Page** to return to the DAS view.

## 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 42 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 view:

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 view.

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 icon in the State column. For example, the  icon indicates the job is *Complete*. You can also point to the icon to display the status.

---

# 6 Account Manager (AM)

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

This chapter contains the following topics:

- [AM overview](#), page 47
- [Managing user accounts](#), page 49
- [Managing groups](#), page 54
- [Managing repositories](#), page 54

**Table 27 Link to Account Manager view**

Origin	Link
left menu	User Management > Account Manager

## 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 Windows Active 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:

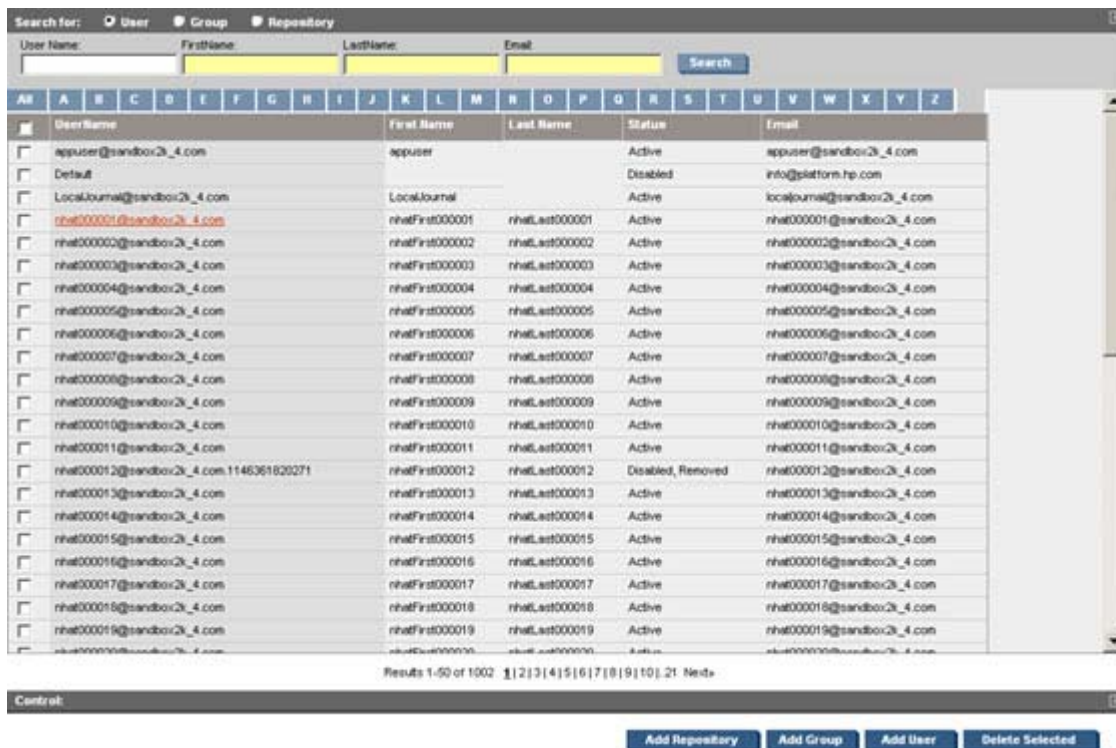
- Add users not imported through dynamic account synchronization (DAS).
- Change user permissions.
- Give users access to other repositories.
- Add special-purpose repositories.
- Add or modify routing rules.
- Edit the retention period for repositories.

---

 **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 10 Account Manager view**

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



## Account Manager view features

**Table 28 Account Manager view 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 <i>jack</i> to match users jackdoe or jacksmith. Entering <i>%doe</i> would match users jackdoe, janedoe, or maryjanedoe. Entering <i>%ja%</i> 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.
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 <b>User</b> filter and <b>All</b> 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. <ul style="list-style-type: none"><li>• Users (see “<a href="#">Managing user accounts</a>” on page 49) - You can add users, view and edit user information, change user status and privileges, and remove users from the system.</li><li>• Groups (see “<a href="#">Managing groups</a>” on page 54) - You can view group information.</li><li>• Repositories (see “<a href="#">Managing repositories</a>” on page 54) - You can create repositories and edit repository information. You cannot delete existing repositories.</li></ul>
Navigation buttons	Click the navigation buttons at the bottom of the view to: <ul style="list-style-type: none"><li>• Add a repository.</li><li>• Add a user.</li><li>• Delete a user from the system.</li></ul>

## 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.

### Adding a new user

To add a new user:

1. Click the Add User button at the bottom of the Account Manager view.
2. Complete the Integrated Archive Platform Account and LDAP Information form that appears.  
See “[Editing user information](#)” on page 50 for more details.
3. Click the Save Now! button.

## Editing user information

To edit user information:

1. Click the **User** radio button, then click the user name that you want to edit.  
See “[Account Manager view features](#)” on page 49 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**.

Integrated Archive Platform Account Information:

LDAP Information:

Username:

Local Password:

First Name:

Last Name:

Email Contact:

Mail To Me Address:

Comments:

Domain:

Mail Server:

Billing Group ID:

Personal Repository:

Direct Repositories:

Active  Disabled

IAP Admin

Compliance  IAP Remote Authorization

Deactivate this check box and then edit the user entries

LDAP Information:

Membership:

WinDomain:

LDAP Dn:

Source:

ObjectGUID:

ObjectSID:

USNChange:

Created Date:

Last Modified:

All Repositories:

Proxies:

**Figure 11 Editing user account information**

3. Edit the relevant user entries.
4. Click the **Save Now!** button.

### NOTE:

You can change a user’s repository retention period by editing the repository. See “[Editing repository information](#)” on page 54.

## User account information

**Table 29 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 <i>Mail To Me</i> for an archived document.
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 higher 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 <b>Edit</b> , selecting the check box for the repository, clicking <b>Add</b> , and then clicking <b>Exit</b> . Repositories (except for the user's personal repository) can be removed from the list by selecting the repository name and clicking <b>Remove</b> .
<b>LDAP Information</b>	
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.
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.

Feature	Description
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.
<b>Check boxes</b>	
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 <b>MUST</b> be enabled on the Exchange Server and the corresponding setting must be enabled on the email miner. The system must be running with Compliance archiving, 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 view. 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.

## Administrative delete

The Administrative Delete functionality gives a user, who is granted with a designated privilege, the ability to physically delete all references of a message that has been stored in the archive. This gives the designated user the ability to remove inadvertently distributed messages from the archive. This feature will typically be used in the Federal space for dealing with classified data; however, it does have application outside this area. *It is the responsibility of the customer to verify that using this feature does not interfere with their corporation'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.

1. In the `Domain.jcml` file, use the `AdminDeleteEnabled` field to indicate whether a Domain has Administrative Delete enabled. If enabled, the system will allow a user who has appropriate privilege to delete emails.

To enable Administrative Delete, `AdminDeleteEnabled` has to be set to true:  
`AdminDeleteEnabled=true`

If the field doesn't exist in the `Domain.jcml` file, add it.

If Administrative delete is enabled, Audit log must also be enabled as described in [Audit log](#).

To disable Administrative Delete, set the `AdminDeleteEnabled` field to false.

2. After configuring `Domain.jcml`, run `regloader.pl -cv -clearallConfirm=xxx` on the Kickstart server and `/opt/bin/restart` on PCC to restart the whole appliance.

The PCC General IAP Configuration page shows the enable status of the Administrative Delete Service.

## Granting “Delete Administration” privilege

All IAP Admin users, with the exception of “root” users, can grant or revoke 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 a user with this privilege:

1. Log in to PCC as `root`.
2. Open the PCC Account Management page and grant “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 checking 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 a user with this 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 unchecking the **Delete Admin** check box.
3. Save the change.

## Executing Administrative Deletion

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

1. Log in to the Web UI as the user with Delete Admin privilege.
2. Search for and select the message to be deleted.
3. Click the **More Options** button.
4. Click the **Delete Checked Items** button.
5. To confirm the deletion, click the **Confirm Delete** button.
6. After confirming deletion, you will see a status page which will report on the success or failure of the deletion submittal.  
After the message is submitted to delete, it will 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 deletion also deletes a quarantined message.
- The deletion is done on both primary and secondary smart cells. If the IAP is running in a replication environment, the message is also deleted on the replica smart cells.
- If the IAP is running in a replication environment, Administrative Delete can only be executed to delete a message which is stored on the primary IAP – the deletion is triggered on the replica IAP by 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. Auditlog messages are stored in the Auditlog repository.

- When a user obtains or loses 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

Current limitations of Administrative Delete:

- In IAP 2.0, Administrative Delete only deletes emails; it doesn't delete documents that were stored with the ObjectStoreAPI (BIBO) or the HP File Migration Agent (FMA), or delete the audit trail.
- If email has been backed up with the IAP backup feature, the email can be deleted from RISS using Administrative Delete; however, the email and any related information on the tape will not be deleted. Because of this, all admin delete operations should be re-executed on the RISS when any smart cell recovery has been executed. (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 result appear when reloading the saved result. 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 gets cleaned up.
- This feature does not currently satisfy DoD or NSA data destruction standards.

## Managing groups

**IAP 1.6 and 2.0 do not support group management in Account Manager.**

## Managing repositories

Use the Repositories panel to change the routing rules or retention periods for a repository, or to 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.

## Adding repositories

To add a repository:

1. Click Add repository.
2. Complete the IAP Repository information form that appears.  
See "[Editing repository information](#)" on page 54.
3. Click the Save Now! button.

## Editing repository information

To edit repository information:

1. Click the **Repository** radio button, then click the repository you want to edit.

2. In the form that appears, clear the check box labeled **Deactivate this check box and then edit the user entries.**

The screenshot shows a web-based configuration form for an IAP Repository. The form is titled "IAP Repository:" and contains several input fields and checkboxes. The fields are: Name (GUID.021BFB521F2444A91E667711265A099.repository), ID (0b00110a2ebfa2118865ab88601), Domain (berlin1.iap), Retention (31), Type (Unregulated), EMail Routing (ztu1@roseilm11.com), Add EMail, EMailDomain Routing, Add EMail Domain, Created Date (Thu, 2008.03.06 15:08:55 PST by), and Last Modified (Thu, 2008.03.06 15:08:55 PST by). There are two checkboxes: one for deleting selected EMail routings and one for deleting selected EMail Domain routings. At the bottom left, there is a checkbox labeled "Deactivate this check box and then edit the repository entries". At the bottom right, there are two buttons: "Save Now!" and "Exit".

IAP Repository:

Name: GUID.021BFB521F2444A91E667711265A099.repository

ID: 0b00110a2ebfa2118865ab88601

Domain: berlin1.iap

Retention: 31

Type: Unregulated

EMail Routing: ztu1@roseilm11.com

Check this box to delete the selected EMail routings from list above.

Add EMail:

EMailDomain Routing:

Check to delete the selected EMail Domain routings from list.

Add EMail Domain:

Created Date: Thu, 2008.03.06 15:08:55 PST by

Last Modified: Thu, 2008.03.06 15:08:55 PST by

Deactivate this check box and then edit the repository entries

Save Now! Exit

**Figure 12 Editing repository information**

3. Edit the relevant entries.
4. Click the Save Now! button.

## Repository information

**Table 30 Repository information**

Feature	Description
Name	(Required.) The name of the selected 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	The amount of time that messages and documents are retained in the repository. The time period is shown in days. For example, 2556 days is 7 years. Note that quarantine repositories have an infinite retention period.
Type	The type of repository: <ul style="list-style-type: none"><li>• Regulated</li><li>• Unregulated</li><li>• Quarantine</li><li>• Other (includes automatically-created repositories for non-indexed documents, catch-all messages, and audit logs)</li></ul>
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	Enter a new email routing mailing address.
Email Domain Routing	The email domain associated with the repository.
Add Email Domain	Enter a new email domain to be associated with the repository.
Created Date	The date the repository was created.
Last Modified	The date the repository was last modified.



# 7 Other user management features

This chapter explains how to use the Manual Account Loader and Error Recovery user management features.

The chapter contains the following topics:

- [Manual Account Loader](#), page 57
- [Account Error Recovery](#), page 58

## Manual Account Loader

Manual Account Loader (MAL) is a batch tool used to load users into the IAP when the Exchange or mail server is not using LDAP. If the Exchange server is using LDAP, use dynamic account synchronization (DAS) instead. (See "[Account Synchronization](#)" on page 41.)

Do not use this tool with Microsoft Exchange 2000 (or later) servers.

Loading user accounts into the system is a three step process:

1. Export a CSV file containing the user accounts you are loading into the IAP. See "[Exporting user account information from Exchange](#)" on page 57.
2. Copy the exported CSV file to the PCC server.
3. Load the CSV file into the IAP. See "[Loading user account information](#)" on page 57.

**Table 31** Links to the Manual Account Loader view

Origin	Link
left menu	User Management > Manual Account Loader

## Exporting user account information

Export user account information by domain. You will need one CSV file for each domain.

To export user account information from an Exchange 5.5 server:

1. Open the Exchange Admin tool.
2. Select **Tools > Export Data**.
3. Copy the exported CSV file to the PCC server.

## Loading user account information

1. Open the Manual Account Loader view in PCC.
2. In the CSV File Uploader box, browse to the CSV file you created in "[Exporting user account information from Exchange](#)" on page 57.
3. Click **Upload CSV File**.
4. In the Domain Name drop-down list, select the domain associated with the user accounts. Only one domain can be selected.
5. Click **Load Selected File**.

The loading process takes about one hour for 50,000 accounts.

- To verify the results, open the Account Manager.

## Error Recovery

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

**Table 32 Links to Error Recovery view**

Origin	Link
left menu	User Management > Account Error Recovery

## Error Recovery features

**Table 33 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 and 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.)
MemberType	The type of repository: <ul style="list-style-type: none"> <li>• Regulated</li> <li>• Unregulated</li> <li>• Quarantine</li> <li>• Other (includes automatically-generated repositories for non-indexed documents, catch-all messages, and audit logs)</li> </ul> (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 and Group filters 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:

- [Replication](#), page 61
- [Cloning](#), page 63
- [Reprocessing](#), page 65
- [Retention](#), page 67
- [Database and data backup](#), page 70
- [Duplicate Manager](#), page 72
- [Folder Support](#), page 74

## Replication

---

 **NOTE:**

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

Use the Replication view to monitor and to start or stop replication for a domain on a remote system. 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 view is unavailable until at least one domain on the PCC host system is configured for replication.)

A failed replicated smart cell triggers allocation of new primary and secondary smart cells on the replica site. When the original site returns to service, it automatically becomes the replica site for new primary and secondary smart cells. Depending on configuration, some administration might be necessary for continued data storage. Specifically, if the system uses Selective Archiving, the replica site mining servers must be configured and enabled to start mining to the replica site.

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 view 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 50.

---

**Table 34 Link to Replication view**

Origin	Link
left menu	Data Management > Replication

## Database Replication

### IMPORTANT:

When installing a replica IAP, finish db2 replication before starting JBoss on all other servers. If you do not some IAP applications will not be able to access db2 because db2 tables are locked by the db2 replication process.

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

**Table 35 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 the steps below to initialize or re-initialize db2 replication:

1. Stop the IAP applications on the primary and replica:  
PCC on primary: `/opt/bin/stop`  
PCC on replica: `/opt/bin/stop`
2. Unconfigure db2 replication on primary/replica:  
DB2 on replica: `/opt/bin/repl/dbRepl rm`  
DB2 on primary: `/opt/bin/repl/dbRepl rm`
3. Restart db2 on primary/replica:  
DB2 on primary: `/etc/init.d/db2_app start`  
DB2 on replica: `/etc/init.d/db2_app start`
4. Re-initialize db2 replication on primary/replica:  
DB2 on primary: `/opt/bin/repl/dbRepl init`  
DB2 on replica: `/opt/bin/repl/dbRepl init`
5. Wait for the replication process to complete (this can take anywhere between a few minutes to an hour or more depending on the size of the db2 database and the bandwidth of the network between the two IAPs):  
DB2 on replica: `/opt/bin/repl/dbRepl status -brief -force`  
Look for the message "Replication is active." Re-issue the command every few minutes until this message is returned, or while you see the message `Refresh is in-progress`.
6. Once replication is successfully initialized, restart the applications on primary/replica:  
PCC on primary: `/opt/bin/start`  
PCC on replica: `/opt/bin/start`

## Replication Status

The middle part of the Replication view displays the status of the replication.

**Table 36 Replication Service General Status**

Feature	Description
Domain	The domain name and group ID of the domain being replicated.
State	Shows whether or not replication is in progress.
File Batch Count	The number of batches being replicated.
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.
Remote Authorization Service	Shows if the IAP remote authorization user has been set up in Account Manager. This user must be set up in the system for the Replication view to display statistics from the remote system. See IAP Remote Authorization in <a href="#">“User account information”</a> on page 50.

## Data Replication Flow

The bottom part of the Replication view displays the flow of replicated data. You can also suspend or resume the replication process from this area.

**Table 37 Data Replication Flow**

Feature	Description
Domain	<p>The domain name and group ID of the domain being replicated, and the following information:</p> <ul style="list-style-type: none"> <li>The IP addresses of the smart cells being duplicated, the status of the replication, and the number of stored and indexed documents being copied. <ul style="list-style-type: none"> <li> A check icon indicates normal operation.</li> <li> An X icon indicates that replication failed the last time it was tried.</li> <li> An ! icon indicates replication is being retried.</li> </ul> </li> <li>An arrow showing the direction of data. Normally, the direction is left to right, from the primary system to the replicated system. In a failover situation, the direction is right to left, as the data replicated on the replicated system is being used to restore the group on the primary system.</li> <li>The number of stored and indexed documents that have been copied to the replicated domain.</li> <li>The percentage of data that has been copied to the replicated domain.</li> </ul>
Suspend/Resume	<p>To stop replication, click <b>Suspend</b>. Replication will stop after the current batch is replicated.</p> <p>To start replication, click <b>Resume</b>. The batch that would have been replicated next when replication stopped will be replicated.</p>

## Cloning

Use the Cloning view to show the status of current and past cloning operations, and to clone a smart cell. You can clone a smart cell if its mirror smart cell is SUSPENDED, DEAD, or FAILED. (See [“Smart cell life cycle states”](#) on page 21.)

Cloning a smart cell copies all its information to another smart cell that is in the FREE state to give the smart cell 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 smart cell in the free pool, contact HP support.

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

**Table 38 Link to Cloning view**

Origin	Link
left menu	Data Management > Cloning

## Cloning view features

**Table 39 Cloning view features**

Feature	Description
Source	The IP address of a smart cell without a viable mirror. If all smart cells have viable mirrors, the Source displays <b>No Broken Groups Found</b> . If more than one smart cell needs a mirror, a Change Source button appears below the automatically selected IP address.
Free Cells	The number of smart cells 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 smart cells.
Clone Cell	Starts the cloning operation. See " <a href="#">Cloning smart cells (copying data)</a> " on page 65.
Cloning Area	<p>This area provides the following information about an ongoing cloning operation:</p> <ul style="list-style-type: none"> <li>• Source selected: IP address of the smart cell being duplicated.</li> <li>• Target selected: IP address of the smart cell receiving duplicate data.</li> <li>• Current Step Percentage: Dynamic bar showing how much source data has been duplicated.</li> <li>• Overall Percentage: Current step in the cloning operation.</li> </ul> <p>Cloning operation steps are as follows:</p> <ul style="list-style-type: none"> <li>• Initializing</li> <li>• Assigning target host</li> <li>• Transferring data</li> <li>• Transferring indexes</li> <li>• Waiting for indexer to complete</li> <li>• Updating history log</li> <li>• Completed or Failed</li> </ul> <p>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.</p>
History Logs	<p>The following information about each cloning operation that has occurred since startup:</p> <ul style="list-style-type: none"> <li>• Source</li> <li>• Target</li> <li>• Time Elapsed</li> <li>• Status</li> <li>• Date</li> </ul>



## Cloning smart cells (copying data)

To clone a smart cell:

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

This button is unavailable if there are no smart cells 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.

## Reprocessing

The Reprocessing view 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.

**Table 40 Link to Reprocessing view**

Origin	Link
left menu	Data Management > Reprocessing

## Rescheduling all reprocessing schedules

To reschedule all the domains to the same reprocessing schedule:

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

Domain:	Reprocessing Status	Scheduled Days	Time (24H)
domain1 [edit]	<input checked="" type="checkbox"/> Enabled	Sun, Mon, Tues, Wed, Thurs, Fri, Sat	02:00:00
domain2 [edit]	<input checked="" type="checkbox"/> Enabled	Sun, Mon, Tues, Wed, Thurs, Fri, Sat	02:00:00

Suspend All   Resume All   Reschedule All   Refresh

**Figure 13** 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 view, locate the domain whose schedule you want to enable or disable.
2. Complete any of the following tasks to change the status:

Domain:	Reprocessing Status	Scheduled Days	Time (24H)
domain1 [edit]	<input checked="" type="checkbox"/> Enabled	Sun, Mon, Tues, Wed, Thurs, Fri, Sat	02:00:00
domain2 [edit]	<input checked="" type="checkbox"/> Enabled	Sun, Mon, Tues, Wed, Thurs, Fri, Sat	02:00:00

Suspend All   Resume All   Reschedule All   Refresh

**Figure 14** 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 view:

**Figure 15 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 view 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, and the number of processed files. This report is based on data from smart cell groups, which is averaged from the individual smart cells of each group.

## Retention

IAP archives emails and other documents in one or more repositories. A repository is a virtual collection of documents associated with a given user by routing rules (storing) and access control lists (ACLs) (retrieving).

You can use the Retention view to configure document retention periods for domains and repositories. For example, you can remove expired stored data or remove access to stored data for specific repositories based on user-defined retention policies. When the retention rules locate expired data, references to the data are removed from the index. The data itself is removed from the file system after all references to it have been removed. This is useful in meeting compliance requirements, when your company must retain all emails for a specified number of years.

In addition to a history log report, the Retention view displays a domain's retention period and status. Each repository within the domain also has a retention period. If the domain's retention period is greater than a specific repository's retention period, the domain's retention period is applied to that repository. Otherwise, the repository's retention period is applied.



### NOTE:

Quarantine repositories are the exception to this rule because they have an infinite retention period.

**Table 41 Link to Retention view**

Origin	Link
left menu	Data Management > Retention

## Searching for and editing a repository retention period

Changes to user repository retention periods are handled by the Account Manager. You can change the retention period in the Account Manager, or you can edit the same form from the Retention view.


To search for a specific repository and change its retention period in the Retention view:

1. Enter all or part of the user name in the Edit User Repositories box and click **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.

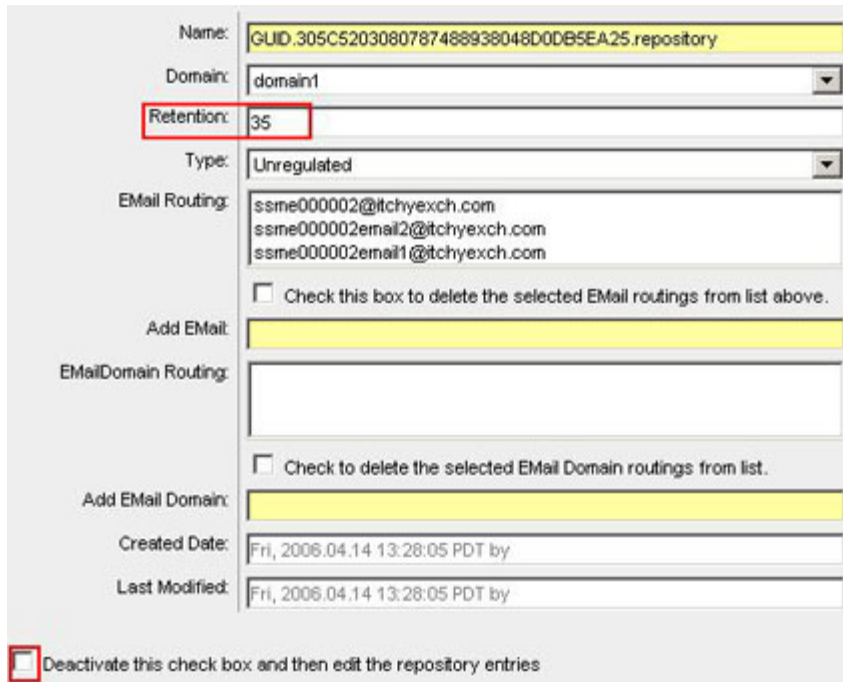
Searches are not case sensitive.

2. If necessary, click one of the repository tabs (Regulated, Unregulated, Quarantine, Other) to locate the repository.
3. Click the repository name to edit its retention period.



**Figure 16 User Repository search results**

4. In the form that appears, clear the check box at the bottom of the form.



The screenshot shows a form for editing a repository. The 'Name' field contains 'GUID.305C5203080787488938048D0DB5EA25.repository'. The 'Domain' dropdown is set to 'domain1'. The 'Retention' field is highlighted with a red box and contains the value '35'. The 'Type' dropdown is set to 'Unregulated'. The 'Email Routing' section lists three email addresses: 'ssme000002@tchyexch.com', 'ssme000002email2@tchyexch.com', and 'ssme000002email1@tchyexch.com'. There is a checkbox to 'Check this box to delete the selected Email routings from list above.' The 'Add Email' field is empty. The 'EMailDomain Routing' section is empty. There is a checkbox to 'Check to delete the selected Email Domain routings from list.' The 'Add Email Domain' field is empty. The 'Created Date' and 'Last Modified' fields both show 'Fri, 2006.04.14 13:28:05 PDT by'. At the bottom, there is a checkbox to 'Deactivate this check box and then edit the repository entries'.

**Figure 17 Edit repository retention period**

5. Change the retention period in the Retention text box.  
Retention periods are shown in days.

6. Click **Save Now!**.

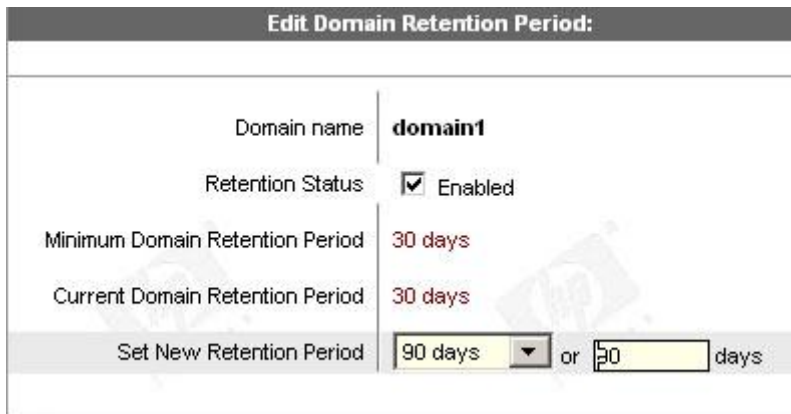
 **NOTE:**

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.

## Editing domain retention periods

To view or edit a domain's retention period:

1. In the Retention view, click **Edit** next to the domain you want to view or edit.
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.



Edit Domain Retention Period:	
Domain name	domain1
Retention Status	<input checked="" type="checkbox"/> Enabled
Minimum Domain Retention Period	30 days
Current Domain Retention Period	30 days
Set New Retention Period	90 days ▼ or 30 days

**Figure 18 Edit domain retention period**

3. Click **Save Retention Now!**.
4. To ensure that retention is enabled, verify that the domain's **Retention Status** check box is selected. If selected, the text *Enabled* appears next to the check box.

## Changing the retention processing status

You can enable or disable retention processing for a specific domain or all the IAP domains.

To change the retention processing status:

1. Locate the domain whose retention processing you want to enable or disable.
2. Complete any of the following tasks to change the status:
  - To enable retention processing for a specific domain, select the corresponding **Retention Status** check box.
  - To disable retention processing for a specific domain, clear the corresponding **Retention Status** check box.
  - To enable retention processing for all domains, click **Resume All**.
  - To disable retention processing for all domains, click **Suspend All**.
3. Verify the status by noting the text *Disabled* or *Enabled* next to the Retention Status check box.

## Viewing retention history logs

The Statistics and Logs area at the bottom of the Retention view shows the last data optimization runs for each configured domain. The log includes each domain group, when a domain was reprocessed last, and file activity.

## Setting the retention basis

Documents can be retained based on either send date or archive date. The send date for an email represents the date when the email was sent by the email client. The send date for a document, stored via the File Migration Agent or via the Object Storage API interface, represents the "last modified timestamp" of the document.

The archive date of an email or document is the date it was stored within the IAP.

If none of the options below are explicitly specified, the default configuration is the archive date for the document retention.

The option can be changed by changing the Domain.jcml file. The two options are:

RetentionBasis=IngestDate

RetentionBasis=SendDate



### NOTE:

When the send date option is configured, email recently stored via the PST Importer disappears the next time retention runs if their send date is outside the retention period specified for the domain and repository.

The PCC [Domain Configuration](#) shows the status of the retention basis option.

## Database and data backup

DB2 database files and the master configuration files are backed up automatically every night. This view shows the status of the last two local backups, the last backup to the kickstart server, and the tape backup if it is enabled.

**Table 42 DB Backup History**

Feature	Description
Backup Status	General status of the backup server and each of its services. <ul style="list-style-type: none"><li>•  A check icon indicates normal operation.</li><li>•  A ! icon indicates a problem or inactive service.</li></ul>
Configuration Information	Configuration information such as backup service (enabled or disabled), job schedule (time) database backup file path location, Tivoli version.
Backup Events	Information about the backup.

**Table 43 Link to DB and Data Backup view**

Origin	Link
left menu	Data Management > DB and Data backup

## Backup file locations

- Database backup files: The database backup files and DB2 transaction log files are stored on the database server (partition: /db2/Vo1Back). The files are mirrored to a directory on the

kickstart machine (/install/db2backups). Database files are backed up once a day, and the local backups are kept for two days before they are deleted.

- 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 changed files are copied to the PCC machine, to the directory /usr/local/MasterConfigFiles.  
(To run the RegistryLoader script 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 backup, because it can be generated from BlackBoxConfig.jcml (/install/tools/converter/dumpBCT). The script createBBC.zr generates the .jcml version that is then backed up.

---

If the backup server is installed and application backup (APP) is enabled in the BlackBoxConfig.bct file, the database backup files and copies of the master configuration files are also backed up to tape.

## Restoring DB2 and master configuration file backups

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 script: /opt/bin/backup/db2BackupUtility 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 one of the following commands to restore the /db2/VolBack partition from either the version stored on the kickstart server or tape backups.

---

 **NOTE:**

This will delete the content of the VolBack directory.

---

- /opt/bin/backup/db2BackupUtility -restore\_from\_kickstart
  - /opt/bin/backup/db2BackupUtility -restore\_from\_tape
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. It will ask for a timestamp of a backup to restore.  
Normally, the last two backups of the database are kept.  
/opt/bin/backup/db2BackupUtility -restore\_db2
  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. It is also backed up to tape, if application backup is enabled.

To restore this directory from the tape backup, run the following command on the PCC server:

```
/usr/local/tsmBackup/rotateMasterConfigBackup -restore
```

## Duplicate Manager

The Duplicate Manager view allows the administrator to schedule duplicate merge jobs and view the status of duplicate merge jobs.

This tool is meant for customers who are storing with Single Instancing enabled. Under certain circumstances the single instancing mechanism allows duplicate copies of the same email to be stored. Storing duplicate emails uses unnecessary space on the smart cells 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 meta data. Over the lifecycle of the email, the meta data 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 meta data. To that end, the Duplicate Manager will take the union of all meta data and attach it to a single copy of the email and remove all redundant copies.

Single instancing is the ability to store a single copy of an email that exists in multiple user repositories. Users can access the email through a reference pointer. Emails are uniquely identified by a cryptographic checksum called a hash. If a client attempts to store the same email twice, as determined by the hash, only one copy is physically stored and the same reference URI is returned for both store requests. In the event that the IAP cannot determine if an email has already been stored, it will store another copy. This can lead to multiple copies of the same email being stored.

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.

Smart cells have a storage threshold that is lower than the total capacity of the hard disk. This is because the services that run on the smart cells require a percentage of the disk space in order to operate. Once the smart cell reaches this threshold it will transition into a closed state. Because duplicate emails may be stored across smart cell groups, there is a possibility that the merge job will not be able to work on a set of duplicates because the smart cells have exceeded the storage threshold and are no longer able to accept additional meta data. These duplicates cannot be merged until additional space has been made available on the closed smart cell 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 may take many days to complete, while subsequent merge jobs may 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. As mentioned above some emails may not have an associated hash. Such emails cannot be merged into a single copy. The duplicate merge procedure may 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, the unmerged duplicate emails often occur for emails at the beginning or end of these batches. A subsequent duplicate merge pass will remove these remaining duplicate emails.

**Table 44** [Link to Duplicate Manager view](#)

Origin	Link
left menu	Data Management > Duplicate Manager



## Duplicate Manager job schedules

The Duplicate Manager Schedule shown on the top portion of the Duplicate Manager view 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.

Duplicate Manager Schedule by Domain:				
Domain	Portal	Scheduled Days	Time (24H)	Duplicate Manager Status
<input type="radio"/> berlin3.iap <a href="#">[edit]</a>				
<input type="radio"/> berlin2.iap <a href="#">[edit]</a>				
<input type="radio"/> berlin1.iap <a href="#">[edit]</a>				

**berlin3.iap Schedule:**  
Set Duplicate Manager schedule for Domain: berlin3.iap

### 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 checkbox 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 the **Save Schedule Now!** button.



#### 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, check or uncheck the **Enable** checkbox and click the **Save Schedule Now!** button.

### 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 the **Start**, **Pause**, or **Abort** button 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 view 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.

#### Duplicate Manager History Logs:

Domain	Start Time	Duplicate Manager Status	Elapsed Time	Progress	Duplicates Processed	Duplicates Deferred
<a href="#">berlin3.iap</a>						
<a href="#">berlin2.iap</a>						
<a href="#">berlin1.iap</a>						

**berlin1.iap All History:**  
View all history logs for domain berlin1.iap

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

Feature	Description
Domain	Domain name
Start Time	Time stamp of when Duplicate Manager job started
Duplicate Manager 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
Duplicates Deferred	Number of duplicates that could not be processed by the Duplicate Manager in the job.

## Folder Support

IAP 2.0 supports Folder Capture support for Email Archiving software for Exchange 2.0. When folder support is enabled, IAP 2.0 captures the folder information of any new email stored.

The FolderSupportEnabled option indicates whether the Domain is folder support enabled. The FolderSupportEnabled option is set in `Domain.jcm1`. 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`

The PCC [Domain Configuration](#) shows whether folder support is enabled.

See the *HP Email Archiving for Exchange Administration Guide version 2.0* for more detail on folder support.

# 9 Reporting

This chapter includes information about the following topics:



- [Event Viewer](#), page 75
- [SNMP Management](#), page 76
- [Email Reporter](#), page 79
- [LogFile Sender](#), page 80

## Event Viewer

The Event Viewer shows the critical and recovery events that have occurred in system services or applications. You can also use the Event Viewer to search for events by type.

The following information is displayed in the Event Viewer.

**Table 45 Event Viewer features**

Feature	Description
Event	Information describing the event, including the service or application name.
Machine Type	The type of server on which the event occurred.
IP Address	The IP address of the machine on which the event occurred.
Date	The date and time of the event.
Level	Status of the event: <ul style="list-style-type: none"><li>•  Critical</li><li>•  Recovery (Recovery represents a transition from a critical status value to the normal status value.)</li></ul>

**Table 46 Links to Event Viewer**

Origin	Link
left menu	Reporting > Event Viewer
left menu	Select <b>Overview</b> , and then click <b>More Details</b> in Platform Events

## Searching the Event Viewer

To search the Event Viewer:

1. Select a Search Type in the drop-down list.  
You can choose from the following types of searches:
  - Show All Alerts
  - Level
  - IP address
  - Host Name
  - Class Name
  - Machine Type

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

You must enter criteria for all searches except Show All Alerts.

The search function uses the “Like” SQL database capability. For example, you could enter `sc` 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 `ms-s0-204-1.company.com`.

Searches are not case sensitive.

3. Click **Submit**.

## Other Event Viewer features

You can change the number of events displayed on each Event Viewer page by following these steps:

1. Go to the Events Per Page drop-down list in the bottom right corner of the view.
2. Select the number of events to show on the page, and click **First**.

You can delete events from the Event Viewer by following these steps:

1. Select an event or go to the time period drop-down list in the bottom left corner of the view.
2. Click **Delete** in the bottom left corner of the view.

## SNMP Management

Use the SNMP Management view to perform the following actions:

- After installing or upgrading to IAP 2.0, download the MIB file `iap.mib`. Import it into your monitoring management software so that the monitoring management software can recognize the IAP SNMP traps.
- Set the SNMP monitoring management server.
- Select SNMP traps for monitoring.
- Receive SNMP events via email.
- Set SNMP community.

**Table 47 Links to SNMP Management view**

Origin	Link
left menu	Reporting > SNMP Management

## Downloading the IAP MIB

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

In the MIB Manager Service area, click the **Download** button to copy the IAP MIB file to your local machine. To import the MIB, refer to the documentation for your monitoring management software.

## Setting the SNMP server

In the SNMP Server Monitors area, enter the IP address of the monitoring management server, then click **Add Server**.

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.

---

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 Server**.

---

**NOTE:**

If you do not have a monitoring management server, you can receive SNMP event notifications via email. See ["Receiving SNMP events by email"](#) on page 78.

---

## Selecting SNMP traps

In the Trap Manager Service area, select the SNMP events, or traps, to be monitored. When one of the traps fails, a notification is sent to the SNMP monitoring management server and/or email recipient.

To set the SNMP traps:

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

2. Select the traps to be activated, and click **Set Traps**.

The following traps can be activated:

- `smartcell_dead`: A smart cell has failed. This trap is generated when a smartcell in the IAP goes in the dead state.
- `reprocessing_trap`: Reprocessing has failed. This trap is generated when the reprocessing service encounters a problem during execution.
- `retention_trap`: Retention has failed. This trap is generated when the retention service encounters a problem during execution.
- `quarantine_trap`: One or more quarantine repositories has failed. This trap is generated when the quarantine service encounters a problem during execution.
- `raidstatus_trap`: There is a disk hardware failure. This trap indicates an abnormal condition on an IAP server's RAID controller.
- `bbwcstatus_trap`: This trap indicates an abnormal condition on a RAID controller's battery-backed write cache.
- `smartcell_suspend`: A smart cell is in the suspended state and cannot store data. This trap indicates that a suspend event has been sent to the SC.
- `smartcell_unsuspend`: This trap indicates that an unsuspend event has been sent to the SC.
- `smartcell_assign`: This trap indicates that an assign event has been sent to the SC.
- `smartcell_assign_space`: This trap indicates that an `assign_space` event has been sent to the SC.
- `smartcell_assign_restore`: This trap indicates that an `assign_restore` (cloning) event has been sent to the SC.
- `smartcell_found`: This trap indicates that a previously lost SC has been started. An attempt will be made to recover it.
- `smartcell_lost`: This trap indicates that an SC has been shutdown or stopped. Its pair will be suspended.
- `pingFailure`: This trap indicates that a specific host has stopped responding to the ping service. The ping service down means either the machine is getting rebooted or there is a network failure. If a reboot has not been issued by the administrator then there is problem with the machine.
- `partition_full`: This trap indicates that one partition has reached a certain critical usage.
- `fs_full_trap`: A filesystem on an IAP has reached 99% full.
- `machine_unresponsive`: The IAP OS and applications on a machine are down. Cannot ping the machine.
- `Assign_space`: A trap that is sent when a smartcell group has run out of room and a new smartcell group needs to be assigned.

## Receiving SNMP events by email

If you do not have a monitoring management 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.

You can disable or delete an email recipient by clicking the radio button in front of the person's name, then clicking **Disable** or **Delete**.

---

 **NOTE:**

The mail host's IP address must be entered in the Email Reports view for SNMP email notifications to be sent. See "[Detailed email reports](#)" on page 79 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.

## Email Reporter

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

You can select the information to be sent and the frequency of the reports— from every 2 hours to every 24 hours. These reports can include system status, storage rates, node health, and other information from the PCC views.

**Table 48 Links to Email Reporter view**

Origin	Link
left menu	Reporting > Email Reporter

## Detailed email reports

The following information is provided in detailed email reports.

**Table 49 Detailed Email Reports**

Feature	Description
VERSION	Provides information about the IAP software version from the Overview. See <a href="#">"IAP Version"</a> on page 26.
EMAILMINER_METRICS	Provides statistics collected by Email Archiving software. See <a href="#">"Overview Archive Gateway"</a> on page 81.
DAS_STATS	Provides user account information from the Account Manager Service section of the Overview. See <a href="#">"Account Manager Service"</a> on page 24. Also provides names of the scheduled account synchronization (DAS) jobs.
NODE_HEALTH	Provides host status and other information from the Platform Control view. See <a href="#">"Platform Control"</a> on page 29.
EVENTS	Provides information on current critical events from the Events section of the Overview. See <a href="#">"Events"</a> on page 23.
SMARTCELL_METRICS	Provides smart cell information from the Platform Statistics section of the Overview. See <a href="#">"Platform Statistics"</a> on page 25.
CONFIG	Provides domain configuration information from the Platform Settings view. See <a href="#">"Domain Configuration"</a> on page 35.
CORE	Provides storage information from the Platform Performance section of the Overview. See <a href="#">"Platform performance"</a> on page 24.

## Creating and scheduling email reports

1. Select one of the following:
  - **Default Selection** to send all information listed in Features.

---

 **NOTE:**

Always use the Default Selection when you are sending emails to HP technical support.

---

- **Custom Selection** to pick specific information to send.
2. Enter one or more email addresses in the Recipient text box.

When entering several email addresses in the box, separate them with a comma or semicolon. For example: recipient1@mycompany.com,recipient2@mycompany.com.
  3. In Features, select the information to be sent.

This function is only available if you have made a custom selection.
  4. In the Frequency drop-down list, select the frequency that the report should be sent.

The frequency ranges from two hours to 24 hours.
  5. Click the maximize box in More Details and enter the email address of the report's sender.
  6. If this is the first time you are using email reports, enter the IP address of the mail host.
  7. Click **Schedule**.

The Email Report information appears in the Current Active Schedules portion of the view. If you want to delete a recipient, select the recipient, then click **Delete**.

## Logfile Sender

The Logfile Sender view lets you select log files to email. For example, you would send log files to HP technical support to assist in troubleshooting.

To send a log file:

1. Enter the recipient's email address in the Email Address field.

You can enter only one email address.
2. Select a log file to send.

You can select only one log file to send.
3. Click **Collect Logs Now!**

### Table 50 Links to LogFile Sender view

Origin	Link
left menu	Reporting > Logfile Sender



---

# 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 view, see “[Overview Archive Gateway](#)” on page 81.

**Table 51** Link to Archive Gateway Management view

Origin	Link
left menu	Archive Gateway Management

## 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 view provides information about the archive gateway system for each domain, including service status and health.



**Table 52 Overview Archive Gateway view 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, 5/11/2006 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: <ul style="list-style-type: none"> <li>• System services</li> <li>• Configured tasks</li> <li>• Journal mining</li> <li>• Selective archiving</li> <li>• Synchronized deleted items</li> <li>• Tombstone maintenance</li> </ul> <input type="checkbox"/> 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 53 System Services features**

Feature	Description
Name	The name of the service.
Status	The following status types can be displayed: <ul style="list-style-type: none"> <li>•  The service is running and the Service Account is not set to LocalSystem. (VNC Server is expected to have LocalSystem as its Service Account.)</li> <li>•  The service is stopped or the Service Account is set to LocalSystem. (VNC Server is expected to have LocalSystem as its Service Account.)</li> </ul> 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 54 Configured Tasks features**

Feature	Description
Task Name	The name of the task.
Status	<p>The following status types can be displayed:</p> <ul style="list-style-type: none"> <li>•  The task is enabled.</li> <li>•  The task is disabled.</li> </ul> <p>The Task Information summary displays the yellow icon if all configured tasks are disabled. Otherwise, the summary displays the green icon.</p>
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 55 Journal Mining features**

Feature	Description
Exchange Statistics	<ul style="list-style-type: none"> <li>• Connection Status: <ul style="list-style-type: none"> <li>•  The Exchange server is available and communicating with the email miner.</li> <li>•  A connection to the Exchange server cannot be completed.</li> </ul> </li> <li>• Exchange Server: The Exchange server on which the journal mailbox resides.</li> <li>• Journal Mailbox Name: The name of the journal mailbox being mined.</li> <li>• Journal Message Count: The number of messages in the journal mailbox at the time statistics were collected.</li> <li>• Journal Folder Size: The size (in KB) of the journal mailbox at the time statistics were collected.</li> <li>• Failed To Archive Count: The number of messages in the FailedToArchive folder of the journal mailbox at the time statistics were collected.</li> <li>• Failed To Archive Size: The size (in KB) of the FailedToArchive folder of the journal mailbox at the time statistics were collected.</li> </ul>
Process Information	<ul style="list-style-type: none"> <li>• Status: <ul style="list-style-type: none"> <li>•  No error conditions for any journal mining process on the journal mailbox.</li> <li>•  A journal mining process is in a retry state.</li> <li>•  A journal mining process is in a failed state.</li> </ul> </li> <li>• Active: The number of journal mining processes that were executing at the time statistics were collected.</li> <li>• Failed: The number of journal mining processes that were in a failed state at the time statistics were collected.</li> <li>• Retry: The number of journal mining processes that were in a retry state at the time statistics were collected.</li> <li>• Completed: The number of journal mining processes that had completed processing at the time statistics were collected.</li> <li>• Queued: The number of journal mining processes that were queued for execution at the time statistics were collected.</li> </ul>

Feature	Description
Message Statistics	<ul style="list-style-type: none"> <li>• Processed: The number of messages that have been processed. Note that <i>processed</i> means different things depending on the context in which it is displayed.</li> <li>• Submitted: The number of messages that were submitted to the IAP for archiving.</li> <li>• Tombstoned: The number of messages that were replaced with “stubs” on the Exchange server.</li> <li>• Ignored: The number of messages that were ignored. Note that <i>ignored</i> means different things depending on the context in which it is displayed.</li> <li>• 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.</li> <li>• Average Processing Rate (msg/sec): The average processing rate.</li> </ul>
Processes in Failed or Retry State	<ul style="list-style-type: none"> <li>• Entry: This information helps identify the process that has entered a failed or retry state. An icon identifies the state: <ul style="list-style-type: none"> <li>⚠ A minor error has occurred. These errors will be retried.</li> <li>🚫 A major error has occurred. These errors will not be retried.</li> </ul> </li> <li>• Status: The problem that caused the process to enter a failed or retry state.</li> <li>• Proc: The number of messages that were processed before the condition occurred.</li> <li>• Tomb: The number of messages that were tombstoned before the condition occurred.</li> <li>• Sub: The number of messages that were submitted to the IAP before the condition occurred.</li> <li>• Rej: The number of messages that were rejected before the condition occurred.</li> <li>• Ign: The number of messages that were ignored before the condition occurred.</li> <li>• Elapsed: The number of elapsed seconds of processing before the condition occurred.</li> <li>• Last Run: The date and time that the condition occurred.</li> </ul>

## 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 83 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 83 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 83 for information on these areas.

## Troubleshooting

If a "DiskSpaceBuffer Threshold has been reached" error is reported during Selective Archiving and folder capture is enabled, the IAP is attempting to update folder information on messages that are in a smart cell that has run out of disk space. Messages impacted by the error cannot be processed by the HP EAs Exchange software.

An example of a folder update is a user moving archived messages from one Outlook folder to another. The next time the folders are archived, the IAP attempts to update the folder name on the previously-archived messages.

Folder changes can be made on messages in a closed smart cell until the available disk space drops below an established watermark. At that point, the IAP rejects the changes with a "DiskSpaceBuffer Threshold has been reached" error. An override to the disk full behavior can be configured on the IAP, which will force the smart cell to accept the update regardless of the disk space check results.

This override is accomplished by setting the 'Application:name=MetaDataRetriever' MBean's attribute 'MetadataChangeQueueThrottleOnDiskFull' to false (from its default setting of true). This can be accomplished by using the ViewCellspace menu option on the PCC for each smart cell, changing the MBean value through MBean attribute and method screen. Alternatively, all the smart cell's can be changed by editing the kickstart file 'SmartCell.xml' setting the value to false and then restarting all the smart cells (or the entire IAP).

When the disk capacity is reached, the folder data submit fails with an error code=12. This should log the errant condition as an Alert and just continue to the next message. The log file may contain many such alert conditions, after it has occurred once.

The errant message is not tombstoned. If the message is actually in the IAP, then its DocRef will be set into the message Exchange properties. But the message class is not changed to PERSISTMailItem.



---

# 11 Audit Log

The Audit Log feature provides a surveillance system log for companies that are required to prove they are adhering to surveillance processes. This chapter describes how to enable the Audit Log feature, set retention periods, monitor status, and grant user access to the repository. For information on performing Audit Log repository queries, see the Audit Log section in the *HP IAP User Guide*.

The chapter includes the following topics:

- [Enabling the Audit Log feature](#), page 87
- [Granting user access to Audit Log repository](#), page 87
- [Monitoring status](#), page 88
- [Setting Audit Log repository retention periods](#), page 88

## Enabling the Audit Log feature

The Audit Log feature is enabled per domain. To enable the Audit Log feature for a domain, add the following attributes in the `/install/configs/primary/Domain.jcml` file for each domain in which the Audit Log 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.

To turn on windoc support:

- `DocClass=email,windoc`

## Granting user access to the Audit Log repository

By default, no user has access to the Audit Log 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 view (**User Management > Account Manager**).
3. Click the **User** radio button at the top of the view.
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 view features](#)" on page 49 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 Information form that appears, click the **Other** tab.

8. Select the check box for the Audit Log repository, for example <domainname>.useraudit-log.repository.



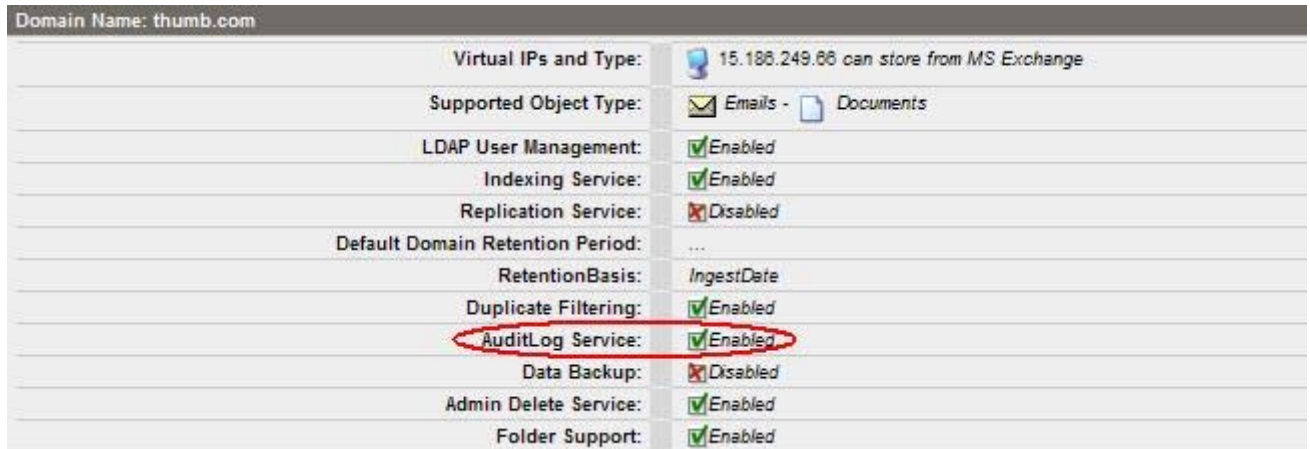
**NOTE:**

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

9. Click **Add**  
The repository is added to the user's direct repositories.
10. Click the Save Now! button.

## Monitoring status

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



**Figure 19 Audit Log enabled**

## Setting Audit Log repository retention periods

The Audit Log repository behaves like a regulated repository. The same rules for changing retention settings apply for the Audit Log repositories.

To view the Audit Log repository for a domain, click the domain name in the PCC Retention view (**Data Management > Retention**).

List of Domain(s)	Retention Period	Retention Status	Set Domain Retention
domain1	2556 days	Enabled	Edit
domain2	30 days	Disabled	Edit

**Figure 20 Accessing domain repository**

Click the **Other Type** tab, and locate the entry for the Audit Log repository.





**Figure 21 Accessing Audit Log repository**

To change the retention period:

1. Click the entry for the domain's userauditlog repository.  
The retention period form appears.
2. Select the retention period from the drop-down list, and click **Save Retention**.



---

# 12 Backup system administration

The optional IAP backup system is the final line of defense in the integrated IAP data-protection strategy. The IAP backup system uses Tivoli Storage Manager (TSM) to create backups of IAP data.

This chapter describes the processes involved in accessing the IAP backup server, configuring TSM, and managing smart cells, and contains detailed procedures for maintaining and labeling backup files and media. It includes the following topics:

- [IAP backup strategy](#), page 91
- [Tivoli Storage Manager](#), page 91
- [Gaining access to the IAP backup server](#), page 91
- [Smart cell data backups](#), page 92
- [Separate Group Volumes](#), page 92
- [TSM backup terms](#), page 92
- [How IAP configures TSM](#), page 93
- [Adding and labeling new media \(Web interface\)](#), page 94
- [Adding and labeling new media \(command line\)](#), page 97
- [Restoring a smart cell](#), page 98
- [Preparing the backup server for disaster recovery](#), page 100
- [Recovering the backup server](#), page 101

## IAP backup strategy

The IAP backup option is the final line of defense in the integrated IAP backup strategy. The first line of defense is that each smart cell group contains two mirrored smart cells. If either of these cells fails, then the survivor can be cloned to recover the mirrored pair. To offer further protection, when a IAP system is replicated, the replicated system also maintains two additional cells that can be used to recover the originals. The IAP backup system is utilized only in the event that all mirrors and replicas are lost.

After the IAP backup option is installed, administrators have the choice of which IAP store domains they want to back up. All store domains can be backed up if desired.

## Tivoli Storage Manager

Tivoli Storage Manager 5.5 (TSM) is used to perform backups to external media. You should have a thorough understanding of TSM before attempting backup procedures. This guide explains the most important TSM concepts involved in IAP management.

For more details about TSM, see <http://www.ibm.com/software/tivoli/products/storage-mgr>. The links provided later in this chapter also provide access to online HTML and PDF versions of the Tivoli Storage Manager Installation and Administrator's Guide and the Administrator's Reference.

## Gaining access to the IAP backup server

A IAP system that has been configured with the backup option has one central backup server machine that runs Tivoli Storage Manager software. This backup server is named INTERNAL and is available at IP 10.0.105.1. Use any of the following methods to access this backup server:

- Use the IAP KVM console to access the backup server directly.
- Use the IAP PCC machine for remote access:

- Use a remote console program like VNC. (Use the PCC's IP address.)
- Use the command line shell from the Tivoli administrative console program (dsmadm, with *admin* as the user name and *admin* as the password).

The availability of these options depends on the access mode that is configured for your IAP. Discuss this with your HP service representative.

## Smart cell data backups

After a IAP storage domain has been configured for backup, each smart cell group in that domain performs a backup every hour. Any new data files that have been stored in the group since the previous backup are combined into a single aggregate file and backed up.

The secondary smart cell usually performs this backup. However, the primary cell in the group performs the backup if the secondary cell is unavailable.

The smart cell itself initiates the backup. Because this is automatic and does not involve Tivoli's internal scheduler, the TSM configuration does not include a schedule for smart cell backups.

## Separate Group Volumes

An option for Separate Group Volumes is available in the Backup section of the IAP master configuration file. If selected, this option ensures that data backed up from one smart cell group is never stored on the same media as another smart cell group.

Only select this option if you specifically require this functionality, because it complicates the creation of offline copies and leads to a less efficient use of backup tapes.

If you select the Separate Group Volumes option, a separate TSM node object and accompanying objects for each smart cell group are created. After you have selected this option, it cannot be changed without first removing all existing smart cell backups.

## TSM backup terms

The following section contains short descriptions of important TSM terms and concepts and how they relate to each other.

The three basic TSM elements are *library*, *device*, and *path*:

A *library* is a set of one or more drives that have similar media mounting requirements. A library can also include robotic devices.

A *device* represents a tape or optical drive.

A *path* describes a one-to-one relationship between a source and a destination. Data can flow from source to destination and back. The path is the element that connects the server, library, and device.

A *device class* is a logical collection of devices for similar media. Every device that is attached to a library has to be a member of a device class. The device class contains the library name and the device type.

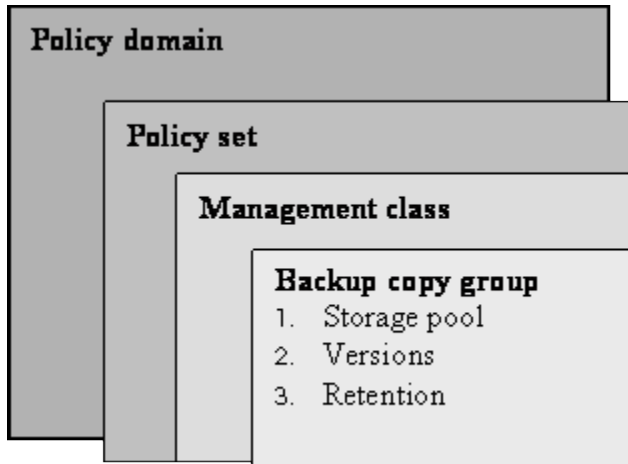
A *storage volume* is the basic unit of storage media (for example, a tape). Storage volumes are grouped into storage pools, which can be arranged in a storage hierarchy. The term server storage represents all storage pools in the system.

The following terms define how client data is handled:

- A *policy domain* is a logical construct that keeps policy sets, management classes, and copy groups together.
- A *policy set* is a set of rules that defines how to handle client data within TSM. (A policy is a rule.)
- A *management class* determines how client files are managed and where they are initially stored. It also contains information about how to handle files that clients back up, archive, or migrate.
- Each management class can contain up to two *copy groups*: a backup copy group and an archive copy group. Each of them points to a destination, which is a storage pool where files

are first stored after they are backed up or archived. A copy group is used to define how many versions of a file are kept and how long client data is retained.

The figure below shows how these categories are related.



**Figure 22 Policy domain structure**

## How IAP configures TSM

IAP's implementation of TSM backup uses only a primary storage pool. No copy storage pools are created. All volumes in the primary storage pool must remain in the library at all times. A full backup is taken when backups are first initialized. After that, only incremental backups are taken.

If a tape volume becomes corrupted or has data loss, backups must be re-initialized. This means that all existing data on tape is removed and a full backup is taken for all smart cell groups in the IAP.

The IAP backup server is initially configured with the script `/usr/local/tsmBackup/configTiv init` from the PCC machine.

This script uses the TSM administrative command line interface to connect to TSM, and sends TSM commands that create the TSM objects and configurations that are necessary for backing up IAP data.

The TSM objects that are created are always preceded with an appropriate prefix to make it easier to remember and recognize the object names. The association of prefixes and objects is as follows:

- LIB: Libraries
- DC: Device Classes
- DR: Drives
- SP: Storage Pools

The details of the `configTiv` script are not required for day-to-day operation, but might be useful for TSM administrators who are interested in customizing their systems.

The process of setting up the initial Tivoli configuration is divided into three phases:

1. Establishing general backup server settings.
  - The default password expiration time is the maximum available (27 years).
  - The password for the administrator is set to never expire.
  - To prevent clients from creating unexpected nodes, the client registration is set to closed. Only the administrator can register client nodes. (See the Tivoli Administration Guide on "Closed Registration" at <http://publib.boulder.ibm.com/infocenter/tivihelp/v1r1/index.jsp?topic=/com.ibm.itsmcw.doc/anrwqd55375.htm>.)
  - An administrator account is created with the name `INTERNAL_ADMIN`. The granted privilege class is `system`.

- A database volume of 1000 MB is allocated on the backup server, if one does not already exist under C:\WINNT\SYSTEM32\DB.2.
2. Initializing the library, storage paths, drives, device classes, and labeling of all of the media. The following steps represent an example configuration and can vary from system to system.
    - Libraries are defined. The name for the libraries is taken from the master configuration file (bct) and the string "LIB" is added as a prefix (for example, LIB.TAPE).
    - A path is defined for each library. In the path, the device parameter specifies the device alias name of the library's robotic mechanism.
    - The devices that belong to each library are defined. For each device a path is defined that contains the device alias and the previously defined library name. A path is the connecting element for the server, library, and device. (See the Tivoli Administration Guide on "Defining Devices and Paths" at <http://publib.boulder.ibm.com/infocenter/tivihelp/v1r1/index.jsp?topic=/com.ibm.itsmcw.doc/anrwqd55169.htm>.)
    - A new device class is defined. The device class name is the same as the library to which it will be connected, except that the string DC is added as a prefix (for example, DC.TAPE). (See the Tivoli Administration Guide on "Defining Device Classes" at <http://publib.boulder.ibm.com/infocenter/tivihelp/v1r1/index.jsp?topic=/com.ibm.itsmcw.doc/anrwqd55272.htm>.)
    - All volumes in all libraries are labeled.
  3. Defining a storage pool, domain, policy set, management class, and copy group for primary group data.
    - A storage pool is created using the device class that was previously defined. (See the Tivoli Administration Guide on "Managing Storage Pools and Volumes" at <http://publib.boulder.ibm.com/infocenter/tivihelp/v1r1/index.jsp?topic=/com.ibm.itsmcw.doc/anrwqd55289.htm>.)
    - A policy domain, a policy set for that domain, and a management class that is assigned to that domain are created.
    - The copy group is defined and the storage pool is activated. The copy group defines that the server will keep one version of a backed up file for an indefinite time. (See the Tivoli Administration Guide on "Defining and Updating a Backup Copy Group" at <http://publib.boulder.ibm.com/infocenter/tivihelp/v1r1/index.jsp?topic=/com.ibm.itsmcw.doc/anrwqd55459.htm>.)
    - Client nodes are registered during the backup process that runs on the smart cells. Therefore, no node is registered during the initial configuration.

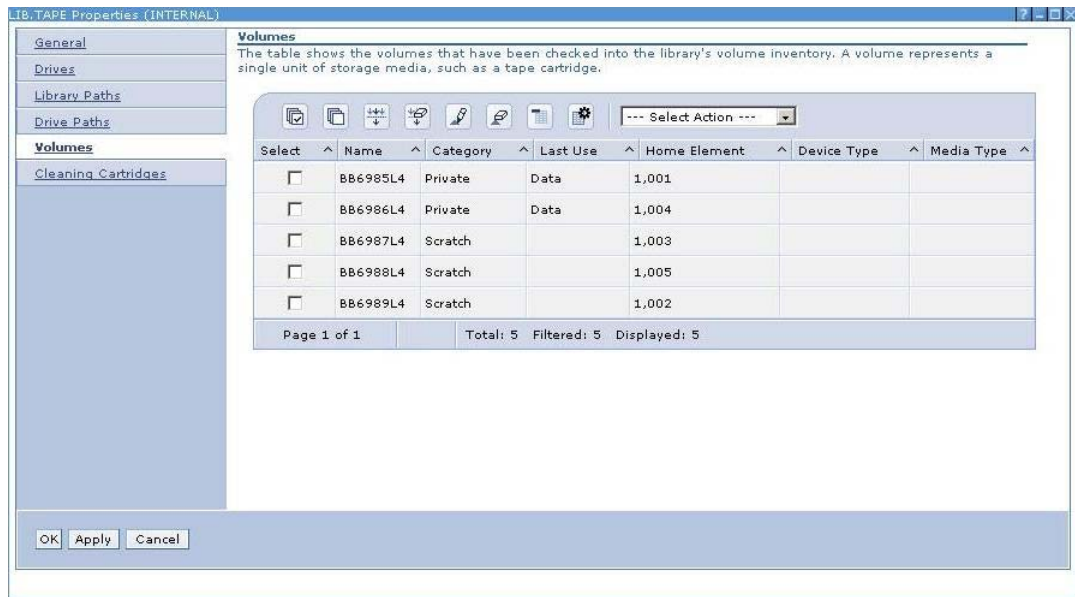
## Adding and labeling new media (Web interface)

To add, label, and check in new media, use the Add Volumes wizard that can be started from the Web administration console. In the following example, two new empty volumes were (physically) added to the library before the procedure was initiated.

1. Log on to the Tivoli Integrated Solutions Console by clicking **Tivoli Web Console** on the backup server desktop, or by going to `http://<PCC-SYSTEM-NAME>:8421/ibm/console`.
2. On the left, open the **Tivoli Storage Manager > Storage Devices** page.

You should see one entry for the backup server and one for the library. If no library is defined, verify that the backup has been initialized with the command `configTiv init` and that the server instance has been added to the Administration Center, as documented in the *IAP Installation Guide version 2.0*.

3. Select the library to which you want to add the media. (In the example, the library is LIB.TAPE.) A properties page appears, as shown in the figure below.



**Figure 23 Library properties**

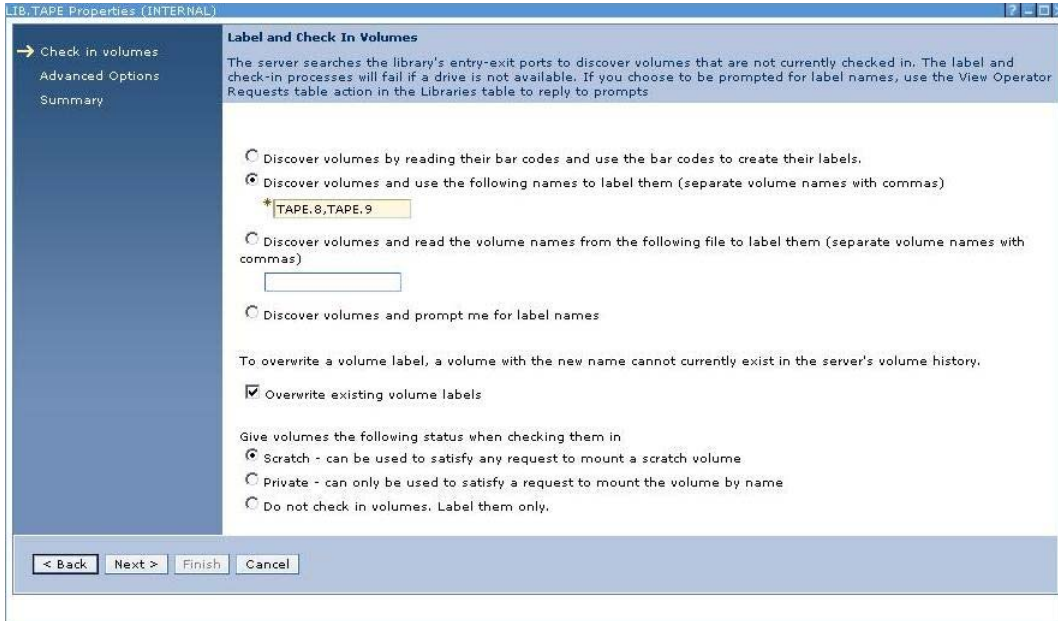


**NOTE:**

In the example, barcodes are available and used as labels.

4. Click the Volumes link to see a list of volumes that are assigned to that library.
5. In the Select Action list, select Add Volumes.
6. Select Not all of the volumes are labeled, and then click Next.
7. Select Search for all eligible volumes in the library's regular slots, and then click Next.

8. Depending on whether or not your library has a barcode reader, select the appropriate option. (See the figure below.)



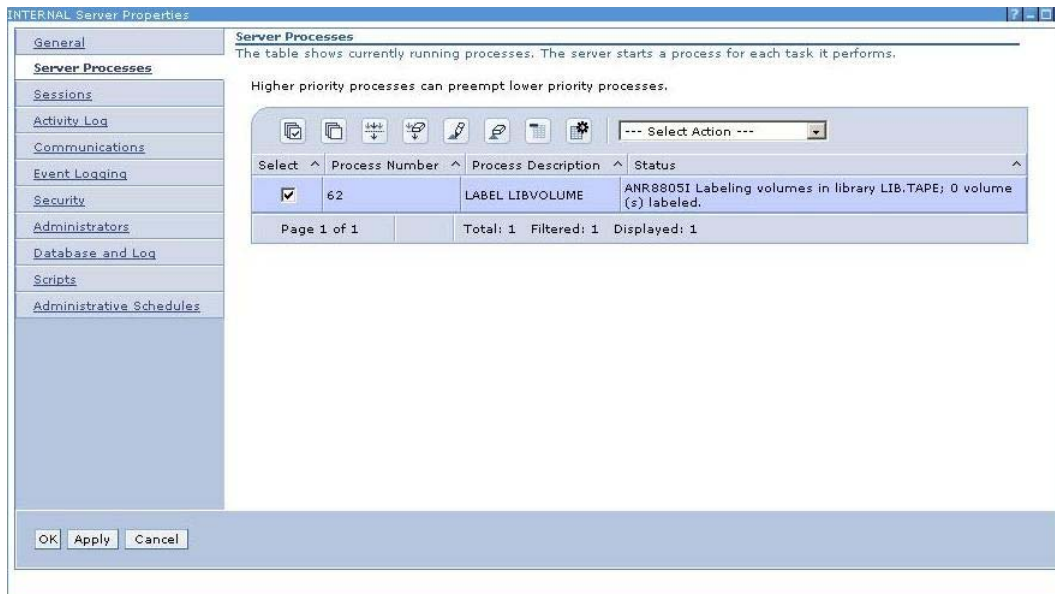
**Figure 24 Label and check in volumes**

In the example, the labels for the new media are specified manually. Also select the check box to overwrite any existing labels.

9. Verify that the volume status is Scratch, and then click Next.
10. Select the default of 60 minutes for mount requests, and then click Next.

A search begins for new volumes in the library. You can monitor the status of the search by selecting the backup sever properties and clicking Server Processes as shown in the figure below. When the search ends, the new media appears in the library volumes list.





**Figure 25 Server process list**

## Adding and labeling new media (command line)

The Add Volumes wizard mentioned in [“Adding and labeling new media \(Web interface\)”](#) on page 94 gathers information to build the necessary Tivoli command. The command can be issued on the command line manually (for example, if the Web interface is not available).

The following command (which includes no line break) is equivalent to the procedure that was described in the previous section:

```
LABEL libvol LIB.TAPE search=yes labelsource=vollist
vollist=TAPE.8,TAPE.9 overwrite=YES checkin=SCRATCH WAITTIME=60
```

This command causes Tivoli Storage Manager to search for media and label them according to the names specified as vollist. It also checks in the new volumes. Use the following procedure to issue the command:

1. Open two ssh sessions to the PCC machine. Use one to observe the command.
2. In the first shell, type the command: `dsmdm -con`.  
The Tivoli monitoring console opens.
3. In the second shell, begin the Tivoli command line interface by typing `dsmdm` and issue the `libvol` command as described.

Note that the volume labels that are specified as a list are separated by commas with no spaces between.

4. You should now see the same command and additional information in the monitoring console. When the process is finished you should see output that looks something like this:

```
ANR8810I Volume TAPE.8 has been labeled in library LIB.TAPE.
```

```
ANR8810I Volume TAPE.9 has been labeled in library LIB.TAPE.
```

```
ANR8801I LABEL LIBVOLUME process 96 for library LIB.TAPE completed;
```

```
2 volume(s) labeled, 2 volume(s) checked-in.
```

If the library has a barcode reader, change the command to use the barcodes as volume labels:

```
LABEL libvolume LIB.TAPE search=yes labelsource=barcode overwrite=YES
```

checkin=SCRATCH

More information about this topic is available at the Tivoli Online Information Center:

<http://publib.boulder.ibm.com/infocenter/tivihelp/v1r1/index.jsp?topic=/com.ibm.itsmcw.doc/anrwgd55242.htm>

## Restoring a smart cell

To restore data on a failed smart cell, you need at least one free smart cell. If only one smart cell in a group failed, then clone a smart cell instead. See “Cloning” on page 63 for more information about cloning.

To restore a smart cell, the following conditions must be met:

- The smart cell domain has `DataBackupEnabled=true` set in the `Domain.jcml` configuration file.
- The smart cell group must have been assigned in this system at some time.
- The smart cell must be marked as missing in the View Cell State table, or no record of the cell's existence appears in the View Cell State table.  
View Cell State is located under Service Tools in the PCC's left menu. It can only be accessed if you log in to the PCC as the super user.

Rebuilding the index on the smart cell takes a significant amount of time, depending on the amount of data that is being restored.

If both smart cells in a group failed, perform the following procedure:

1. In the View Cell Space view, scroll to the Other Smart Cells area, and verify that at least one free smart cell exists.

2. Locate the primary controller:
  - a. Click one of the links in the MetaServer table.  
The Agent view appears.
  - b. From the MBean list, click ProvisionerMBean.  
The MBean view appears.
  - c. Verify that the ProvisionerMasterBackupStatus attribute is set to Master Provisioner.  
If the attribute is set to Backup Provisioner, return to the View Cell Space view, click the link to the other meta server, and repeat the previous steps.

**List of MBean attributes:**

Name	Type	Access	Value	Description
RunningState	int	R	3	MBean Attribute.
ProvisionerPingPeriod	long	RW	60000	MBean Attribute.
StateString	java.lang.String	R	Started	MBean Attribute.
State	int	R	3	MBean Attribute.
Message	java.lang.String	R		MBean Attribute.
Name	java.lang.String	R	ProvisionerWrapper	MBean Attribute.
ProvisionerPingDelay	long	RW	45000	MBean Attribute.
MBeanUptime	long	R	233	MBean Attribute.
DebugLogging	boolean	RW	<input checked="" type="radio"/> True <input type="radio"/> False	MBean Attribute.
ProvisionerMasterBackupStatus	java.lang.String	R	Master Provisioner	MBean Attribute.

Apply Changes

**Figure 26 Provisioner status**

3. Determine which smart cells failed and can be restored:
  - a. In the MBean view for the primary controller, click the button next to ListBrokenGroups.  
Groups that failed and have not been recovered are listed.
  - b. Copy GroupIDs and Roles of broken groups.
  - c. Click **Back to MBean view**.

4. Restore the smart cell:

After you determine which smart cell you want to restore, begin the other Provisioner operation, RestoreSmartCellUsingGroupIDAndRole. A smart cell from the Free Pool (the group of unassigned and empty smart cells) is assigned as a restore target. The smart cell begins restoring its data. If the data indexes were not backed up, then the smart cell will rebuild them before coming back online. (This generally happens if the smart cell was still open for storage when it failed.)

- a. Under the **RestoreCmartCellUsingGroupIDAndRole** operation, enter the group ID in the first field and the role in the second field.
- b. Click **Invoke**.

If the assignment of a cell is successful, a confirmation message appears. This can take up to two minutes. Note the restore target information.

If an error occurs, verify that the group ID and role are correct. If the information is correct, verify that the smart cell needs to be restored and that a free smart cell is available.

5. Verify that the restore process is running correctly or has completed:

The smart cell should restore its data and indexes and then restart itself to open in the correct state. (It does not reboot.) During the restore, *restore* appears next to the cell state on the View Cell Space page. Afterward, *assigned* should appear next to both the primary and secondary cell. If the restore has failed, the BackupSystem MBean on that restore target will be listed as *failed* on the Overview page of the PCC.

- a. In the View Cell Space view, locate the restored smart cell that is listed in the confirmation message.
- b. If the assigned smart cell is not in the restore state, the restore process has completed and no further action is required. Otherwise, click the link to the smart cell in the restore state.
- c. From the MBean list, click **BackupSystem/MBean**.  
The MBean view appears.
- d. Verify that the RunningState attribute is **3** and the FailureReason attribute is **No Failure**.

This indicates that the restore process is occurring normally. If the RunningState attribute is 5, then the restore process failed. Check the FailureReason attribute to discover the cause.

## Preparing the backup server for disaster recovery

The following steps and procedures describe how to restore the backup server (not the entire IAP system) in case the server itself is lost or destroyed. All volumes in the library must remain intact. Restoring the backup server involves performing backups of the backup server database and configuration files. At this time, taking tapes offsite is not supported.

### Things to back up

In order to protect the backup server, the following information has to be backed up:

1. A backup of the TSM database and recovery log. The database contains information about the client data that is stored in storage pools. The recovery log contains changes that have been made to the database.
2. A backup of the TSM configuration, including all of these components:
  - A copy of the Volume History File: Volume history information is stored in the database, but is not available during a database restore. Among other things, this file contains information about database backups and volumes that have been added, removed, or reused.
  - A copy of the Device Configuration File: This file contains information that is required to read backup data, such as library, path, device, and device class definitions.
  - A copy of the Server Options File: This file contains general settings for the server.

We recommend taking a backup of the TSM database at the end of each business day.

### To perform a backup

The following steps show how to perform a backup.

---

 **NOTE:**

The commands that begin with *tsm: INTERNAL>* were issued from the Tivoli command line interface (dsmadm) from the PCC machine. The Administrative Command Line is also installed on the backup server. To start it, click **Start > Programs > Tivoli Storage Manager > Administrative Command Line**.

1. Perform a full backup of the database. Note that TAPE is the library name set in the BCT file. To find your device class name, run 'query devclass' from the Administrative Command Line:

```
tsm: INTERNAL>backup db type=full devclass=DC.TAPE
```

2. Back up the volume history file:  
tsm: INTERNAL>backup volhistory
3. Back up the device configuration:  
tsm: INTERNAL>backup devconfig
4. Copy the following files off of the backup server: volhist.out, devcnfg.out, dsmserv,opt, dsmserv.dsk. These files are located under C:\Program Files\Tivoli\tsm\server1. Back up these files by copying them onto a USB key or scp them to the kickstart server.

## Recovering the backup server

To restore the TSM database, contact HP support.



# Index

## Symbols

IAP, definition, 15

## A

Account Error Recovery view, 58  
Account Manager, 47  
    *See also* AM  
    Account Manager Service, 24  
account synchronization  
    *See* DAS  
Account Synchronization view, 41  
accounts, user, 49  
administrative privileges, IAP, 52  
AM, 47  
    *See also* Account Manager  
    about AM, 47  
    Account Manager window, 47  
    adding repositories, 54  
    adding users, 49  
    definition, 47  
    group panel, 54  
    user accounts, 49  
API Configuration and Statistics view, 32  
Archive Gateway management  
    logging in, 81  
Archive Gateway Management view, 81  
ASSIGNED smart cell state, 21  
audit log, 15, 87

## B

BACKING\_UP smart cell state, 22  
backup files  
    accessing IAP backup server, 91  
    backup file history, 70  
    backup file locations, 70  
    backup strategy, 91  
    configuration files, 71  
    database files, 70  
    disaster recovery, 100  
    restoring configuration files, 71  
    restoring DB2 backups, 71  
    restoring smart cell, 98  
    using Tivoli Storage Manager, 91  
backup system administration, 91

## C

catchall repository, 24  
certificate signing request (CSR)  
    deleting, 37  
    generating, 37

certificates  
    installing, 37  
cloning smart cell, 63  
CLOSED smart cell state, 21  
COMPLETE\_PROCESSING smart cell state, 21  
compliance, regulatory, 52, 67, 87  
configuration  
    domain, 35  
    IAP, 36  
configuration file backup, 70  
conventions  
    text symbols, 12  
CPU use, 25, 27  
critical events, 23, 75

## D

DAS  
    configuring, 41  
    creating jobs, 41  
    editing or deleting jobs, 45  
    history logs, 46  
    repairing synchronization errors, 58  
    scheduling jobs, 45  
    starting jobs, 45  
    stopping jobs, 45  
data  
    archiving, 15  
    querying, 15  
    reprocessing, 65  
    retention, 67  
Database and data backup view, 70  
database backup, 70, 91  
DEAD smart cell state, 22  
disaster recovery, 100  
DISCOVERY smart cell state, 22  
document retention, 67

## E

EAs  
    application programs for users, 15  
email log files, 80  
Email Reporter, 79  
email SNMP notifications, 78  
event log, 23, 75  
Event Viewer, 75  
events  
    critical, 23, 75  
    recovery, 75

## F

failed indexed repository, 24  
firewall settings, 36  
FREE smart cell state, 22

## H

health  
  checking system, 23  
  machine, 27  
  smart cell, 25  
help, obtaining, 12  
host groups  
  definition, 21  
  types, 27, 29  
host machines  
  starting, stopping, or restarting, 29  
  status, 21, 27, 29  
HP  
  storage web site, 13  
  Subscriber's choice web site, 13  
  technical support, 12  
HP OpenView, 76  
HTTP servers, starting, stopping, or restarting, 29

## I

IAP  
  applications, 15  
IAP administrator, 52  
IAP Authorization User, 52, 61, 63  
index rate, 24, 26, 30

## J

JBoss, 25, 27  
journal mining, 83

## L

L2, L3 software, 40  
LDAP servers  
  configuring DAS, 41

left menu, PCC, 19

life cycle states  
  definition, 21

life cycle states, types, 21

log in

  email mining system, 81

  PCC, 17

  VNC, 81

Logfile Sender, 80

logs

  email, 80

Lotus Notes Interface

  description, 15

## M

Manual Account Loader view, 57

memory

  host machines, 27

MIB,, 76

Mining

  Archive Gateway Management, 81

monitoring reports, 78, 79

monitoring, PCC, 21

## N

notifications

  detailed email reports, 79

  SNMP events, 78

## O

Outlook Interface

  description, 15

Overview Archive Gateway, 81

Overview, PCC, 23

## P

patch history, 40



## PCC

- Cloning view, [63](#)
  - about, [17](#)
  - accessing, [17](#)
  - Account Error Recovery view, [58](#)
  - Account Manager view, [47](#)
  - Account Synchronization view, [41](#)
  - API Configuration and Statistics view, [32](#)
  - Archive Gateway Management view, [81](#)
  - Archive Gateway overview, [81](#)
  - common administration tasks, [19](#)
  - Database and data backup view, [70](#)
  - description, [15](#)
  - detailed email reports, [79](#)
  - Email Reporter, [79](#)
  - Event Viewer, [75](#)
  - health, checking system, [23](#)
  - left menu, [17](#), [19](#)
  - log in, [17](#)
  - Logfile Sender, [80](#)
  - Manual Account Loader view, [57](#)
  - monitoring tools, [21](#)
  - Overview, [23](#)
  - patch history, [40](#)
  - Performance Graph view, [30](#)
  - Platform Control view, [29](#)
  - Platform Settings view, [35](#)
  - printing, [19](#)
  - refreshing views, [19](#)
  - Replication view, [61](#)
  - Reprocessing view, [65](#)
  - Retention view, [67](#)
  - smart cell life cycle states, [21](#)
  - SNMP Management view, [76](#)
  - Software Version view, [40](#)
  - SSL Configuration view, [37](#)
  - states, [21](#)
  - status conditions, [21](#)
  - Storage Status view, [26](#)
  - System Status view, [27](#)
  - updating views, [19](#)
  - user interface, [18](#), [18](#)
  - views, [17](#), [18](#)
- performance graphs
  - appliance storage and indexing, [30](#)
  - system monitoring, [30](#)
- Platform Control Center
  - See PCC
- Platform Control view, [29](#)
- platform performance, [24](#)
- Platform Settings view, [35](#)
- power on/off, [15](#)
- printing PCC views, [19](#)
- private keys (certificates), [37](#)
- PST Importer
  - description, [15](#)

## Q

- querying
  - data, [15](#)

## R

- RAID controller, [27](#)
- regulatory compliance, [52](#), [67](#), [87](#)
- related documentation, [11](#)
- Remote Authorization, IAP, [52](#)
- replicating smart cells, [61](#)
- replication
  - suspending and resuming, [63](#)
- Replication view, [61](#)
- reports
  - detailed email, [79](#)
  - SNMP events, [78](#)
- repositories
  - adding, [54](#)
  - catchall, [24](#)
  - creating, modifying, [54](#)
  - definition, [47](#)
  - domain retention period, [69](#)
  - failed indexed, [24](#)
  - quarantine, [67](#)
  - retention period, [67](#), [68](#)
- reprocessing data, [65](#)
- Reprocessing view, [65](#)
- RESET smart cell state, [22](#)
- restarting servers, [29](#)
- RESTORE smart cell state, [22](#)
- Retention view, [67](#)
- retention, document, [67](#)
- routing
  - email, [56](#)
- RSA private key, [37](#)

## S

- search function, [49](#)
- selective archiving, [81](#), [81](#)
- servers
  - CPU use, [27](#)
  - DAS, configuring, [41](#)
  - MAC address, [29](#)
  - memory, [27](#)
  - starting, stopping, or restarting, [29](#)
  - status, [29](#)
  - thread count, [27](#)
- smart cells
  - allocation, [26](#)
  - MAC address, [25](#)
  - machine health, [25](#)
  - overview information, [25](#)
  - replicating, [61](#)
  - restoring, [98](#)
  - states, life cycle, [21](#), [25](#)
  - System Status view, [27](#)
  - thread count, [25](#)
- SMTP Flow Control, [26](#)

SMTP Servers, starting, stopping, or restarting, [29](#)

SNMP Management view, [76](#)

SNMP traps

notifications, [76, 78](#)

selecting, [77](#)

setting SNMP server, [76](#)

Software Version view, [40](#)

software versions, [26, 40](#)

SSL Configuration view, [37](#)

states

definition, [21](#)

types of, [21](#)

stopping servers, [29](#)

storage rate, [24, 26, 30](#)

Storage Status view, [26](#)

Subscriber's choice, HP, [13](#)

SUSPENDED smart cell state, [22](#)

symbols in text, [12](#)

SYNC\_WAIT smart cell state, [22](#)

synchronization, [41](#)

synchronization errors, [24](#)

repairing, [58](#)

System Status view, [27](#)

## T

technical support

HP, [12](#)

text symbols, [12](#)

thread count, [25, 27](#)

## U

UNKNOWN smart cell state, [22](#)

users

accounts, [49](#)

adding new, [49](#)

PCC management, [41](#)

Users panel, [49](#)

## V

version identifier, viewing, [26, 40](#)

VNC, [81](#)

## W

W2 software, [40](#)

web sites

HP documentation, [13](#)

HP storage, [13](#)

HP Subscriber's choice, [13](#)