HP Medical Archive solution

Software version: 8.1.0

user guide

Document release date: May 2010
Software release date: November 2009

Legal notices

Warranty

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.

The information contained herein is subject to change without notice.

Restricted rights legend

Confidential computer software. Valid license from HP required for possession, use or copying. Consistent with FAR 12.211 and 12.212, Commercial Computer Software, Computer Software Documentation, and Technical Data for Commercial Items are licensed to the U.S. Government under vendor's standard commercial license.

Licensing

The use of HP products is governed by the terms and conditions of the applicable End User License Agreement (EULA).

Copyright notices

Trademark notices

Microsoft, Windows, Windows NT, and Windows XP are U.S. registered trademarks of Microsoft Corporation.

UNIX® is a registered trademark of The Open Group.

Tivoli® Storage Manager (TSM) server

Contents

About this document

The *user guide* provides an introduction to:

- an overview of the features and components of the HP MAS product
- an introduction to the NMS web interface
- a description of the data flow when objects are ingested, replicated, retrieved, modified, and deleted
- monitoring tips
- an overview of grid configuration
- a general description of grid servers
- an introduction to ILM (information lifecycle management)
- procedures to execute grid tasks
- an overview of Server Manager
- troubleshooting information for NMS alarms

Intended audience

This guide is intended for:

- administrators
- PACS administrators
- technical support staff responsible for monitoring the HP Medical Archive solution system

The document assumes general understanding of the grid's components and functionality. A fairly high level of computer literacy is assumed, including knowledge of file systems, tree-structured hierarchies, and network connectivity. You should also be familiar with using a web browser.

This document assumes familiarity with many terms related to computer operations and programming, network communications, and operating system file operations. There is wide use of acronyms.

Prerequisites

Prerequisites for using this product include:

- Operating system knowledge

Related documentation

In addition to this guide, please refer to other documents for this product:

- *HP Medical Archive solution Release Notes* for 8.1.1
- *HP Medical Archive solution audit message reference*
- *HP Medical Archive solution DICOM Integration Guide*
- *HP Medical Archive solution DICOM Conformance Statement*
- *HP Medical Archive solution IHE Integration Statement*
- *HP Medical Archive solution Siemens Integration Guide*

These and other HP documents can be found on the HP documents web site:

http://www.hp.com/support/

Document conventions and symbols

Convention	Element
Medium blue text: Figure 1 http://www.hp.com	• Cross-reference links • E-mail addresses • Web site addresses
Bold	• Key names or key sequence • GUI elements that are clicked or selected, such as menu and list items, buttons, and check boxes • Text typed into a GUI element, such as into a box
Italics	• Document titles • Text emphasis • You must supply a value for a variable in a GUI element.
`Monospace`	• File and directory names • Text displayed on the screen, such as system output and application messages • Command or reserved keyword in a CLI, API, program language, or operating system • Script or code example
`Italic monospace`	You must supply a value on the command line.
`Bold monospace`	• Text typed at the command line • Emphasis of file and directory names, system output, and code

NOTE Provides additional information.

RECOMMENDATION Provides guidance from HP for a best practice or for optimum performance.

△ CAUTION Caution messages appear before procedures which, if not observed, could result in loss of data or damage to equipment.

Documentation updates

The title page of this document contains the following identifying information:

- Software version number

  Indicates the software version.

- Document release date

  Changes each time the document is updated.

- Software release date

  Indicates the release date of this version of the software.

Subscription service

HP strongly recommends that customers sign up online using the Subscriber's choice web site:

http://www.hp.com/go/e-updates

- Subscribing to this service provides you with e-mail updates on the latest product enhancements, newest versions of drivers, and firmware documentation updates as well as instant access to numerous other product resources.
- After signing up, you can quickly locate your products under Product Category.

Support

You can visit the HP Software Support web site at:

http://www.hp.com/go/hpsoftwaresupport

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

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

- Research and register for software training

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

To find more information about access levels, go to:

http://h20230.www2.hp.com/new_access_levels.jsp

For more information about HP Passport, go to:

http://h20229.www2.hp.com/passport-registration.html

1 Introduction to the HP MAS system

The HP MAS System

The HP MAS system is a storage management solution that stores, protects, and preserves fixed-content over long periods of time (tens of years). Through the use of sophisticated Information Lifecycle Management (ILM) rules, when client applications store objects to the grid they are securely stored and replicated to protect from loss. Wide Area Network (WAN) links extend the HP MAS system, enabling off-site replication of content for business continuity and increased content availability. In systems with multiple sites, this replication means that if one site is lost, data is not lost, and clients are still able to seamlessly retrieve data from copies stored at other sites.

The HP MAS system employs a grid architecture of interconnected servers. This grid architecture enables object replication and data protection across multiple servers or sites, while creating a continuously available and highly reliable system. If one part of the grid goes down another takes over.

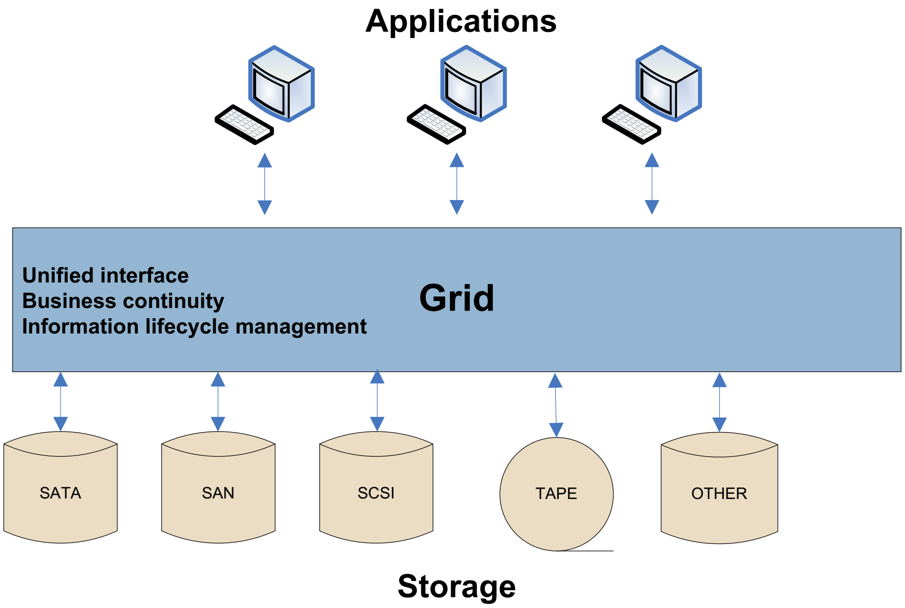

Figure 1 Storage Management System for Massive Volumes of Fixed-Content Data

The HP MAS system:

- employs standardized network file system protocols (NFS & CIFS) to exchange data with external systems, including PACS, viewing workstations, and modalities
- relies on open standards for interoperability with external systems and is installed on Linux based servers
- addresses hardware obsolescence via transparent migration of data from previous to current generations of hardware
- supports multi-vendor hardware and policy driven tiered storage to reduce overall storage costs
- virtualizes storage across sites and tiers, insulating applications from changes to underlying storage infrastructure
- monitors and verifies data integrity proactively as data is stored, replicated, and retrieved

The HP MAS system's linked servers each host at least one grid node. A grid node is a collection of one or more grid services. A grid service is a software module providing a set of capabilities to the HP MAS system. Each grid node within the HP MAS system can be upgraded, decommissioned, replaced, or temporarily disconnected without disruption to client applications.

Deployment Topologies

The HP MAS system can be deployed in a number of topology configurations, including the following common configurations:

- Data Center
- Data Center + Disaster Recovery

The deployments described in this section are simplified examples and do not necessarily represent a complete grid deployment.

NOTE Any topology configuration can also include a Tape Node that manages storage of data to a nearline system.

Data Center

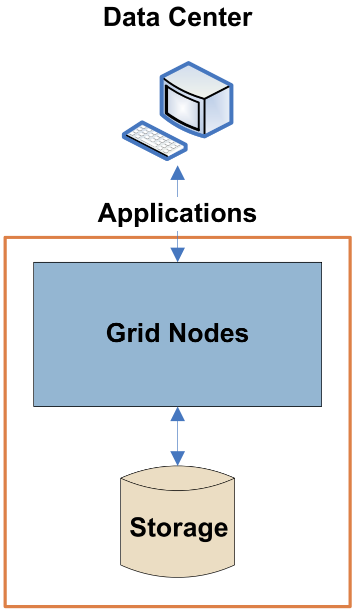

Figure 2 Data Center Deployment

In a Data Center (DC) deployment, the infrastructure and operations of the HP MAS system are centralized in a single site. There is no off-site Disaster Recovery facility.

The following is an example of a DC deployment that includes a Tape Node to manage the storage of data to a nearline system.

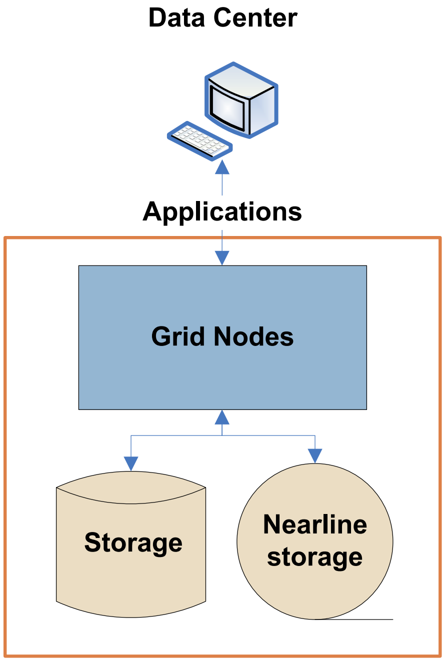

Figure 3 Data Center + Archive Deployment

Data Center + Disaster Recovery

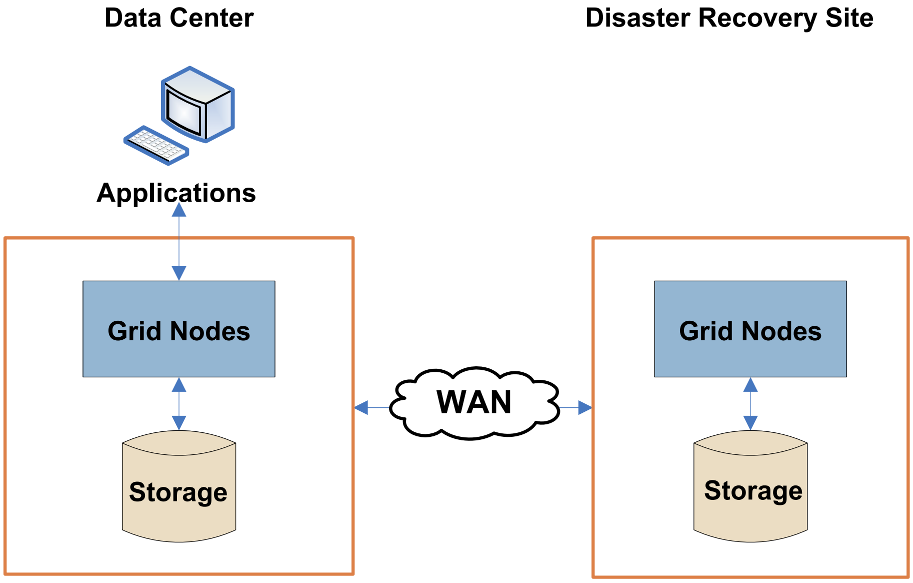

Figure 4 Data Center + Disaster Recovery Deployment

In a Data Center + Disaster Recovery (DC + DR) deployment, the infrastructure is consolidated at the DC site and replicated to the DR site. Generally, the DR site is located in a geographically different location than the DC site.

Grid Nodes

The basic building blocks of an HP MAS system are grid nodes. A grid node consists of one or more grid services running on a server. A grid service is a software component that performs a specific function. For more information on grid services, see Grid Services (page 19).

The basic grid node types are:

- Gateway Node—provides the interface to the grid through which applications, or clients, communicate with the grid
- Storage Node—manages data storage on spinning disks
- Control Node—stores and manages content metadata
- Admin Node—provides grid management services such as grid monitoring, logging, and grid configuration
- Tape Node—for environments where content is also stored on archive media, provides an interface to the middleware that manages the archive media storage device such as a tape library

NOTE Grid nodes can be combined to reduce the number of servers deployed in a grid.

Figure 5 displays how grid nodes and services are commonly arranged within the grid.

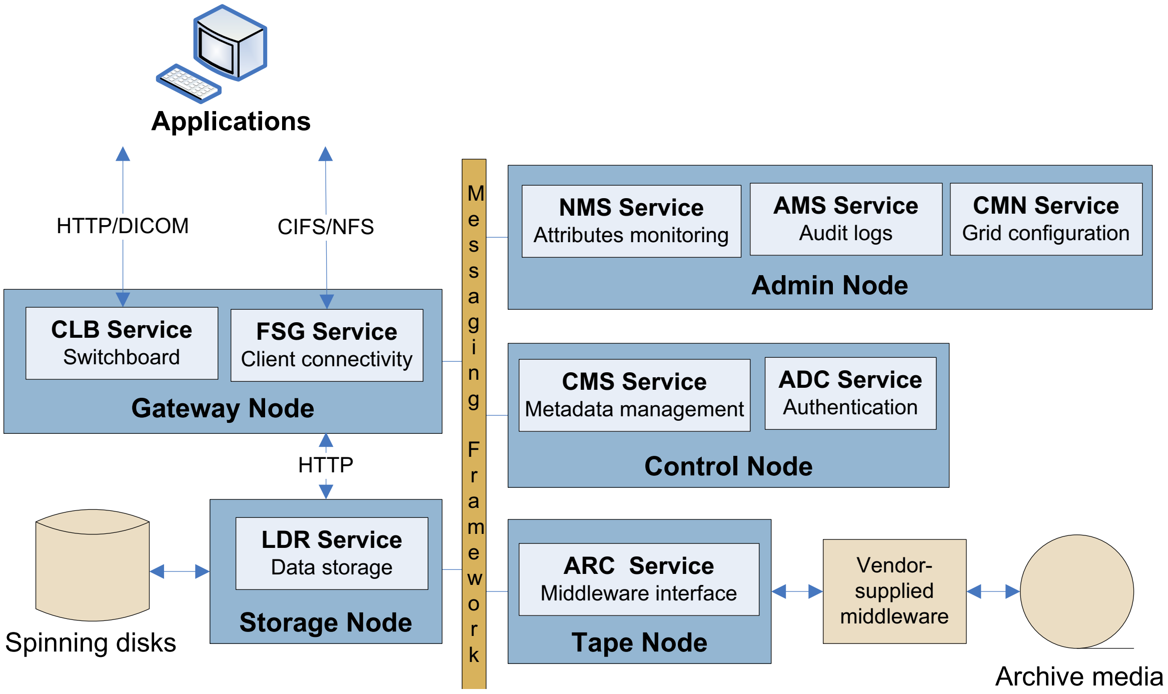

Figure 5 Common HP MAS System Deployment

Table 1 lists the grid nodes and what grid services they host.

Table 1 Grid Nodes

Grid Node	Function	Grid Services	Notes
Admin Node	Grid management	NMS CMN AMS SSM	The Admin Node provides tools for grid administration. For purposes of redundancy, grids can have more than one Admin Node. The primary Admin Node hosts the CMN service. (a grid only has one CMN service). The optional High Capacity Admin Cluster (HCAC) allows the NMS service to support a greater number of grid nodes. The HCAC is made up of a reporting Admin Node and a processing Admin Node. The reporting Admin Node hosts the CMN service and is thus the primary Admin Node.
Tape Node	Data storage on archive media	ARC SSM	The Tape Node manages storage of data to nearline (neither “online”—instantly available—nor “offline”) data storage devices such as tape libraries (via IBM Tivoli® Storage Manager)
Control Node	Content management	CMS ADC SSM	The Control Node manages content metadata and content replication. Each grid has at least two Control Nodes for redundancy of metadata storage. Note that best practices suggest that a grid should have at least three Control Nodes.
Gateway Node	Client connectivity	CLB FSG SSM	The Gateway Node provides an interface between client applications and the grid. Each grid must have at least two Gateway Nodes for redundancy. Depending on what client services are required, a Gateway Node may include an FSG service (for CIFS/NFS), a CLB service (for HTTP/DICOM), or both.
Storage Node	Data storage on spinning disks	LDR SSM	The Storage Node manages data storage on spinning disks. Each grid must have at least two Storage Nodes for redundancy.

Table 1 Grid Nodes *(continued)*

Grid Node	Function	Grid Services	Notes
Admin/ Gateway Node	Grid management and client connectivity	NMS CMN AMS CLB FSG SSM	Combined Admin Node and Gateway Node functionality on a single server. The HCAC configuration is not supported in combined Admin/Gateway Nodes.
Control/ Storage Node	Content management and Data storage	CMS ADC LDR SSM	Combined Control Node and Storage Node functionality on a single server.
Custom nodes	Your system may also include custom grid nodes, for example, an Admin/Gateway/Control/Storage Node or a Gateway/Control/Storage Node that combines the functionality of multiple grid nodes on a single server. The HCAC configuration is not supported for custom nodes.		

Grid Services

A grid service is a software module providing a set of capabilities to the HP MAS system. For an explanation of how these grid services work together during object ingest, retrieval and delete, see Chapter 3, Data Flow.

Table 2 Grid Services

Acronym	Name	Description	Grid Node
ADC 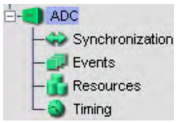	Administrative Domain Controller	Maintains topology information and provides authentication services.	Control Node
AMS 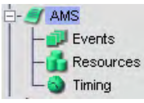	Audit Management System	Keeps logs of grid activity and events.	Admin Node (reporting in HCAC)
ARC ARC Replication Store Retrieve Middleware Events Resources Timing	Archive	Communicates with archiving middleware to store and retrieve data to and from archive media.	Tape Node

Table 2 Grid Services ***(continued)***

Acronym	Name	Description	Grid Node
CLB 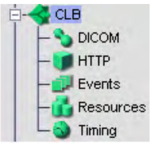 	Connection Load Balancer	Acts as switchboard for connecting remote entities to the most efficient LDR. Primary connection point for remote entities using the DICOM or HTTP protocols.	Gateway Node
CMN 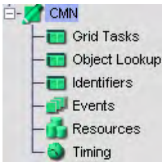 CMN Grid Tasks Object Lookup Identifiers Events Resources Timing	Configuration Management Node	Manages grid-wide configurations: connection profiles, FSG replication groups, grid tasks, grid options.	Primary Admin Node (reporting in HCAC)
CMS 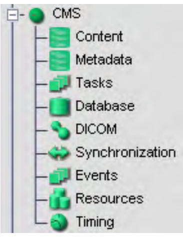 CMS Content Metadata Tasks Database DICOM Synchronization Events Resources Timing	Content Management System	Keeps track of what data is stored on the grid. Stores content metadata and manages content replication based on ILM rules.	Control Node
FSG 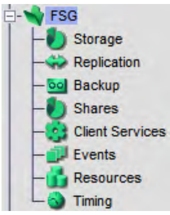 FSG Storage Replication Backup Shares Client Services Events Resources Timing	File System Gateway	Allows connections to the grid via a standard file system (NFS or CIFS).	Gateway Node

Table 2 Grid Services *(continued)*

Acronym	Name	Description	Grid Node
LDR 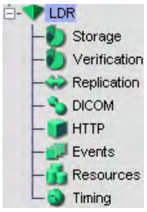 LDR Storage Verification Replication DICOM HTTP Events Resources Timing	Local Distribution Router	Stores, moves, verifies, and retrieves object data stored on disks.	Storage Node
NMS 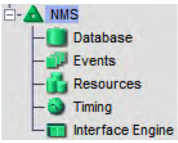 NMS Database Events Resources Timing Interface Engine	Network Management System	Provides a window into the grid. Used to monitor grid status and to configure the grid. Large grids may be configured with an HCAC to increase the grid's grid services and thus grid node capacity.	Admin Node (both reporting and processing in an HCAC)
SSM 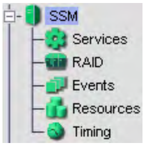	Server Status Monitor	Monitors hardware performance such as key operating system metrics and network metrics. The current operating system installed on the node is listed (SSM > Services).	Present on all nodes

Grid Options

The HP MAS system may include a number of options that effect how the grid operates. The following table lists the key options that may be deployed with a HP MAS system:

Table 3 HP MAS Options

Option	Description
DICOM	Enables the HP MAS system to communicate with medical imaging devices and applications using the DICOM (Digital Imaging and Communications in Medicine) protocol.
Audit	Provides user access to the audit logs which contain a complete record of grid activity.
Compression	Compresses objects saved to the grid, reducing file size by roughly 50% for content that is not already in a compressed format.
High Capacity Admin Cluster	Increases the grid's grid services and thus grid node capacity by configuring the grid with a reporting Admin Node and a processing Admin Node.
Deduplication	Deletes unnecessary identical copies of an object from the grid. Deduplication is unavailable on a grid configured with metadata replication. If deduplication is enabled, metadata replication cannot be enabled.

Table 3 HP MAS Options *(continued)*

Option	Description
Deletion protection	Prevents clients from deleting any content that has been stored to the grid.
Metadata synchronization	Content metadata is synchronized among all CMSs in the grid. After a capacity expansion, metadata is synchronized among all CMSs of the same generation. Required in grids that use deduplication, or non-configurable ILMs. All other grids use metadata replication.
Encryption	Enables encrypted storage of all ingested data. Content is encrypted during ingest and objects are stored in an encrypted form so that if a server is compromised no data can be retrieved in any readable form. Note that after it is enabled, encryption cannot be disabled.
Parallel loading	Enables Gateway Nodes to preload in their cache all the files in a directory upon an initial file request. This enhances performance for systems that store related files in a single directory.
Secondary preloading	Enables caching on all Gateway Nodes in the Gateway Node replication group as files are ingested or retrieved to speed access via the secondary Gateway Node.
Security partitions	Provides the ability to isolate content ingested from different FSG replication groups or HTTP API clients, such that each client only has access to its own data.

Grid Configuration Information

SAID Package

The Software Activation and Integration Data (SAID) package contains customer-specific files and software needed to install, expand, perform maintenance on, or upgrade a grid. The SAID package contains grid specific configuration and integration information including server hostnames and IP addresses. It is generated during the provisioning phase of installation and saved to the Provisioning USB flash drive.

NOTE The SAID package contains highly confidential passwords and encryption keys needed during system maintenance, updates, and expansion. Only trained and authorized service personnel should have access to the SAID package. Store the SAID package in a secure location.

Grid Configuration HTML Pages

The SAID package includes a `\Doc` directory that contains html pages documenting the configuration of the grid. Click the `index.html` file to open these grid configuration web pages.

NOTE You must use Microsoft's Internet Explorer web browser to access these html pages.

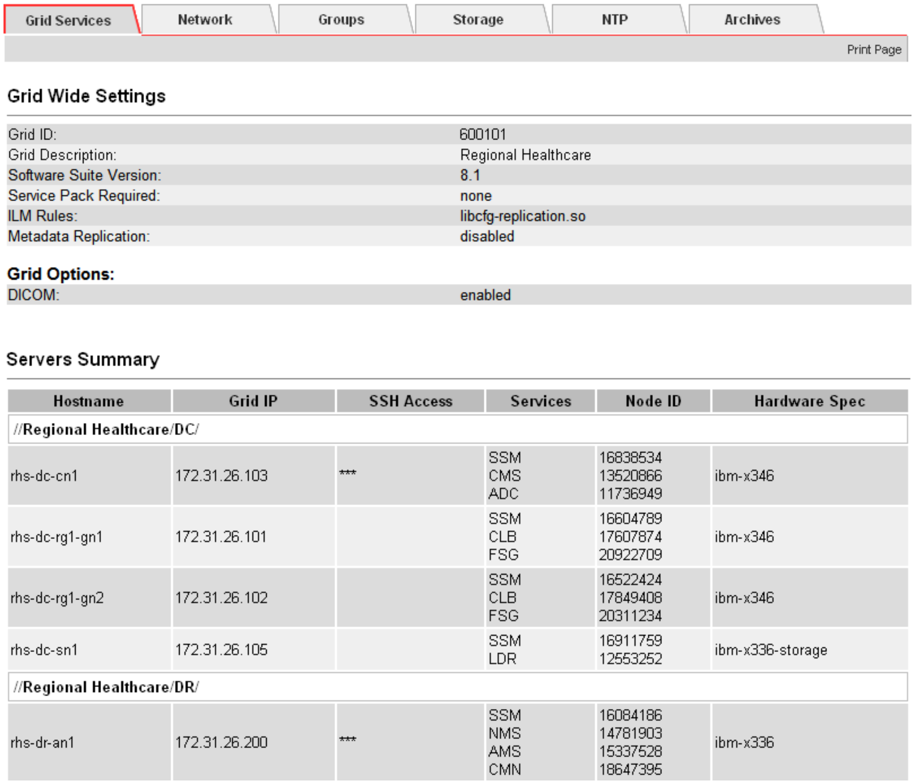

Grid Wide Settings

Grid ID:	600101
Grid Description:	Regional Healthcare
Software Suite Version:	8.1
Service Pack Required:	none
ILM Rules:	libcfg-replication.so
Metadata Replication:	disabled

Grid Options:

DICOM:	enabled

Servers Summary

Hostname	Grid IP	SSH Access	Services	Node ID	Hardware Spec
//Regional Healthcare/DC/					
rhs-dc-cn1	172.31.26.103	***	SSM CMS ADC	16838534 13520866 11736949	ibm-x346
rhs-dc-rg1-gn1	172.31.26.101		SSM CLB FSG	16604789 17607874 20922709	ibm-x346
rhs-dc-rg1-gn2	172.31.26.102		SSM CLB FSG	16522424 17849408 20311234	ibm-x346
rhs-dc-sn1	172.31.26.105		SSM LDR	16911759 12553252	ibm-x336-storage
//Regional Healthcare/DR/					
rhs-dr-an1	172.31.26.200	***	SSM NMS AMS CMN	16084186 14781903 15337528 18647395	ibm-x336

Figure 6 Grid Configuration HTML Pages

Passwords

Passwords used to access the grid can be found in the `Passwords.txt` file.

License Agreement

The text of the HP license agreement is located in a file stored in `/var/local/install/` for HP MAS 7.5 or earlier and `/var/local` for HP MAS 8.0 or later.

Server Consoles

During regular day-to-day operations, you do not need to access the HP MAS system's server consoles. However, occasionally, you may be required to run commands directly from the server console in order to troubleshoot problems or execute maintenance procedures.

Server Manager

Each server in a HP MAS system runs the Server Manager application. Server Manager is used to supervise the starting and stopping of services on the server, ensuring services gracefully join and leave the grid. Server Manager also monitors services on the server and attempts to restart any that report errors.

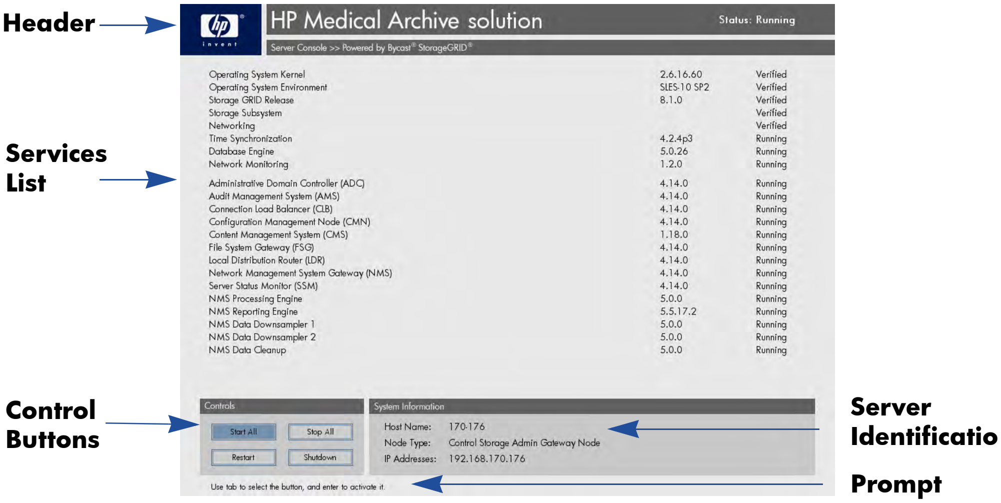

HP Medical Archive solution Status: Running

Server Console >> Powered by Bycast® StorageGRID®

Operating System Kernel	2.6.16.60	Verified
Operating System Environment	SLES-10 SP2	Verified
Storage GRID Release	8.1.0	Verified
Storage Subsystem		Verified
Networking		Verified
Time Synchronization	4.2.4p3	Running
Database Engine	5.0.26	Running
Network Monitoring	1.2.0	Running
Administrative Domain Controller (ADC)	4.14.0	Running
Audit Management System (AMS)	4.14.0	Running
Connection Load Balancer (CLB)	4.14.0	Running
Configuration Management Node (CMN)	4.14.0	Running
Content Management System (CMS)	1.18.0	Running
File System Gateway (FSG)	4.14.0	Running
Local Distribution Router (LDR)	4.14.0	Running
Network Management System Gateway (NMS)	4.14.0	Running
Server Status Monitor (SSM)	4.14.0	Running
NMS Processing Engine	5.0.0	Running
NMS Reporting Engine	5.5.17.2	Running
NMS Data Downsampler 1	5.0.0	Running
NMS Data Downsampler 2	5.0.0	Running
NMS Data Cleanup	5.0.0	Running

Controls

System Information

Host Name: 170-176
Node Type: Control Storage Admin Gateway Node
IP Addresses: 192.168.170.176

Use tab to select the button, and enter to activate it.

Figure 7 Server Manager

The body of the display is the list of services being monitored by the Server Manager on this server.

For more information on Server Manager, see Chapter 9, Server Manager.

Command Shell Access

Occasionally, you may be asked to run commands directly from a command shell on the server console.

Log In

To log in and access a command shell:

At the server, press **<Alt>+<F1>** to access a command shell and log in as `root` using the password listed in the `Passwords.txt` file.

Log Out

To log out of a command shell session:

Close the current command shell session. Enter: **`exit`**

Press **<Alt>+<F7>** to return to the Server Manager GUI.

2 NMS Management Interface

This chapter contains an overview of the browser-based NMS Management Interface (MI) that you use to monitor the grid. The chapter explains how to log in and out, describes the interface elements, and contains procedures to configure your account, monitor alarms, and create reports.

Browser Requirements

The only supported browser is Microsoft Internet Explorer v6.0 SP2 or v7.0. JavaScript and cookies must be enabled.

Configure Internet Options Settings

Internet Explorer settings for temporary internet files, security, and privacy must be set correctly.

To verify the Internet Explorer settings:

1 Go to **Tools > Internet Options > General**.

2 In the Temporary Internet files box, click **Settings**.

3 In the Check for newer versions of stored pages section, verify that Automatically is selected.

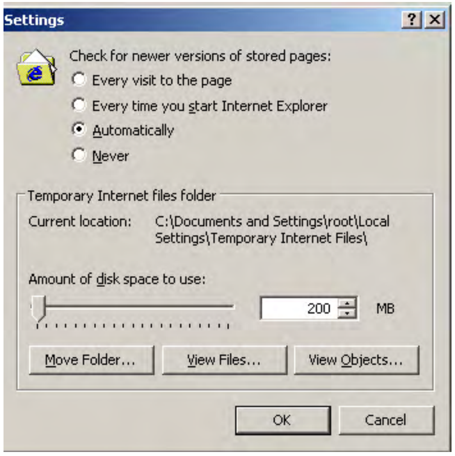

Settings

Check for newer versions of stored pages:

○ Every visit to the page

○ Every time you start Internet Explorer

● Automatically

○ Never

Temporary Internet files folder

Current location: C:\Documents and Settings\root\Local Settings\Temporary Internet Files\

Amount of disk space to use: 200 MB

Figure 8 Temporary Files Setting

4 Go to **Tools > Internet Options > Security > Custom Level** and ensure that the Active scripting setting is Enable.

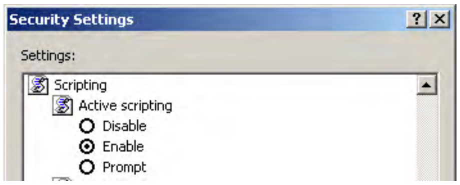

Security Settings

Settings:

Scripting

Active scripting

○ Disable

● Enable

○ Prompt

Figure 9 Active Scripting Setting

5. Go to **Tools > Internet Options > Privacy** and ensure that the privacy setting is Medium or lower (cookies must be enabled).

Enable Pop-ups

NOTE To make changes to passwords, you must ensure that Internet Explorer has the Pop-up Blocker turned off.

To enable pop-ups in Internet Explorer:

1. Select **Tools > Pop-up Blocker > Turn off Pop-up Blocker** from the Internet Explorer main menu.
2. The menu option is a toggle. If the pop-up blocker is already disabled, the menu option is Turn On Pop-up Blocker.

Log In and Out

Accessing the NMS MI requires a web browser (Internet Explorer) with grid access to a known address defined by your system administrator. You will need a user name and password to access the system. Each system user is assigned a user name and password when first introduced to the NMS MI.

NOTE The only supported browser is Microsoft Internet Explorer v6.0 SP2 or v7.0. Cookies must be enabled.

Log In

To log in to the NMS MI:

1. Launch a web browser.
2. Enter the address: https://<IP_Address>
3. If you are prompted with a Security Alert dialog, do one of the following:

   — Proceed with this session. The alert will appear again the next time you access this URL.

   —or—

   — View and install the certificate using the installation wizard so that you no longer receive the alert.
4. Enter your user name and password in the Login window and click log in.

Both user name and password are case sensitive. Keystrokes appear as bullet characters (•) to protect your password.

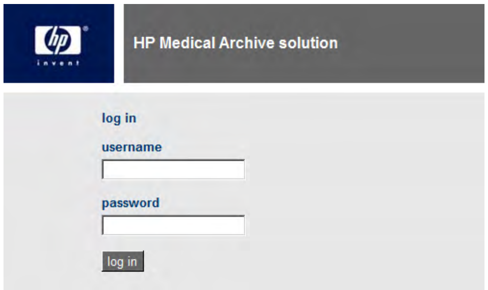

Figure 10 Login Window

If the user name or password you entered cannot be validated, the following message appears in red under the password field: Invalid credentials. Repeat the login process to correct your entry. If you have forgotten your password, contact HP Support to have your password reset.

Log Out

When you have finished your NMS MI session, log out to keep the system secure.

To log out:

1 Click Logout Logout ➔ located at the top right corner of the screen.

2 The logging out message appears.

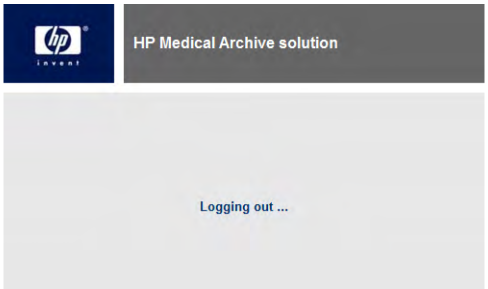

Figure 11 Logout Window

3 You may safely close the browser or continue using other applications.

4 Failure to log out may give unauthorized users access to your NMS MI session. Simply closing your browser is not sufficient to log out of the session.

NMS Management Interface

The NMS Management Interface (MI) provides basic operational data, alarm status, reporting functionality, and configuration options for each grid node, service, and component.

The main elements of the NMS MI are:

- Header
- Grid Management Menu
- Grid Topology Tree
- Content Tabs

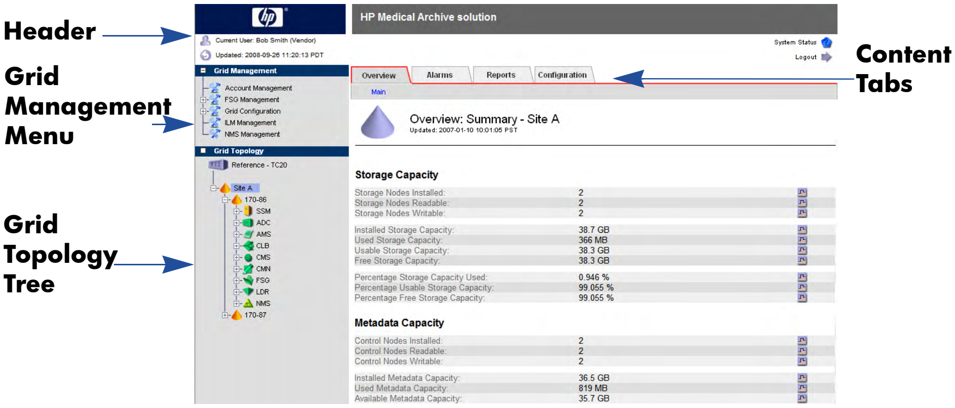

Overview: Summary - Site A
Updated: 2007-01-10 10:01:05 PST

Storage Capacity

Storage Nodes Installed:	2
Storage Nodes Readable:	2
Storage Nodes Writable:	2
Installed Storage Capacity:	38.7 GB
Used Storage Capacity:	366 MB
Usable Storage Capacity:	38.3 GB
Free Storage Capacity:	38.3 GB
Percentage Storage Capacity Used:	0.946 %
Percentage Usable Storage Capacity:	99.055 %
Percentage Free Storage Capacity:	99.055 %

Metadata Capacity

Control Nodes Installed:	2
Control Nodes Readable:	2
Control Nodes Writable:	2
Installed Metadata Capacity:	36.5 GB
Used Metadata Capacity:	819 MB
Available Metadata Capacity:	35.7 GB

Figure 12 NMS Management Interface Elements

Header

The header contains high-level grid status information. The currently logged in user's name and User Group, and the latest browser refresh time are displayed on the left. Clicking the user's name opens the Account Management page. The Hewlett-Packard Product Name logo on the left serves as a button to access NMS MI version information. The right side contains the System Status indicator and the Logout button.

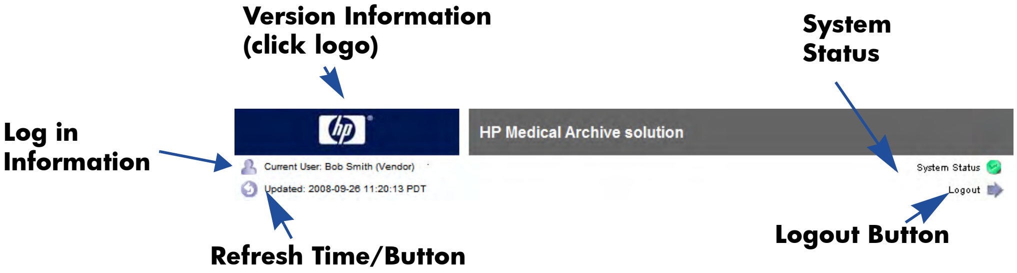
Version Information (click logo)
System Status
Log in Information
HP Medical Archive solution
Current User: Bob Smith (Vendor)
Updated: 2008-09-26 11:20:13 PDT
System Status
Logout
Logout Button
Refresh Time/Button

Figure 13 NMS Management Interface — Header

Refreshing the Display

Information presented in the NMS MI is time-sensitive. Since the interface is delivered as HTML pages, the content shown is static. The Updated timestamp indicates when the data shown was collected, that is, the time at which the last grid status “snapshot” was taken. Local time is shown as determined from the preferences set in the user account. The information is refreshed automatically at set intervals (the default is 15 seconds).

To refresh the display manually, do one of the following:

- Click Refresh Page ↻ .
- Click the timestamp.
- Click the refresh button for the browser (for Microsoft Internet Explorer, press **<F5>**).

After the page has finished reloading, the content and the timestamp are updated.

If you leave the NMS MI open with no activity, the session expires after the configurable timeout period and returns you to the Login Window.

If you have configured the NMS GUI Timeout period to 0 and then minimize your browser for an extended period of time (for example, greater than one week), when you expand the browser, the NMS MI may not correctly display attribute values and the current time.

Displaying NMS MI Version Information

To display NMS MI version information:

- Click the HP logo in the header.

  The About window displays the interface version number, the software build number, and copyright information.

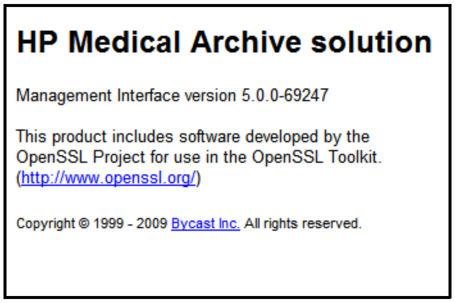

HP Medical Archive solution

Management Interface version 5.0.0-69247

This product includes software developed by the OpenSSL Project for use in the OpenSSL Toolkit. (http://www.openssl.org/)

Copyright © 1999 - 2009 Bycast Inc. All rights reserved.

Figure 14 About Window—Version Information

NMS Interface Engine Status

You can view the status of the NMS MI at any time through the **<Admin_Node> > NMS > Overview** page. The NMS Interface Engine Status attribute displays the current status of the NMS MI for the selected Admin Node.

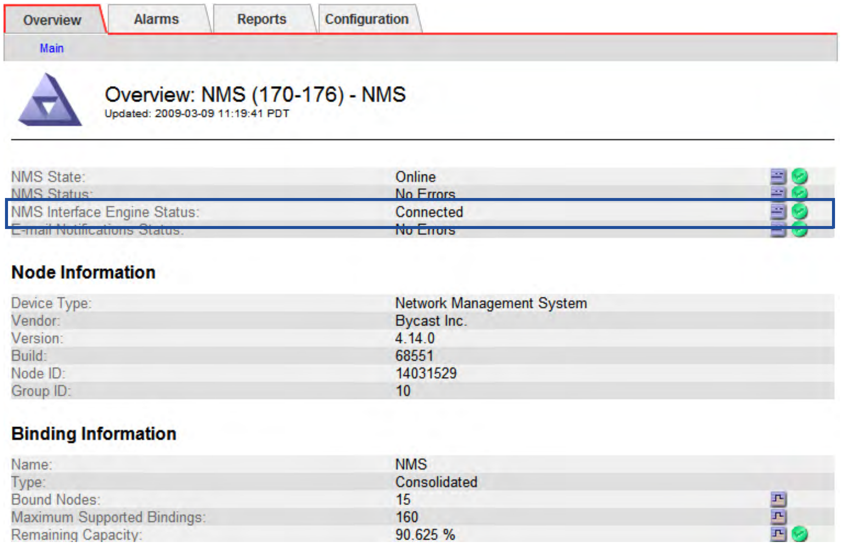

Overview: NMS (170-176) - NMS
Updated: 2009-03-09 11:19:41 PDT

NMS State:	Online
NMS Status:	No Errors
NMS Interface Engine Status:	Connected
E-mail Notifications Status:	No Errors

Node Information

Device Type:	Network Management System
Vendor:	Bycast Inc.
Version:	4.14.0
Build:	68551
Node ID:	14031529
Group ID:	10

Binding Information

Name:	NMS
Type:	Consolidated
Bound Nodes:	15
Maximum Supported Bindings:	160
Remaining Capacity:	90.625 %

Figure 15 NMS Interface Engine Status

This interface engine status information is useful if your grid is configured with multiple Admin Nodes and thus multiple NMS services. You can monitor the status of the grid's other NMS MIs to which web clients are connected. Monitoring the NMS interface engine status can tell you if there is a connectivity problem with any of your Admin Nodes.

Grid Management Menu

The Grid Management menu provides access to a number of configuration pages.

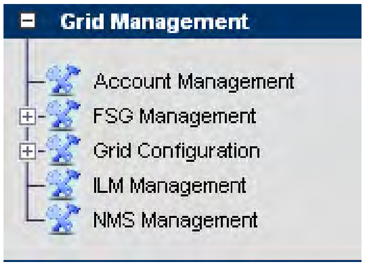

Figure 16 Grid Management Menu

Table 4 Grid Management Menu

Item	Used to Configure
Account Management	User Accounts. For more information on accounts and user groups, see User Accounts (page 36).
FSG Management	FSG settings such as content protection options, cache space, and backups. Configuration is restricted to user accounts that have Maintenance permissions such as the Admin and Vendor accounts.
Grid Configuration	Grid options such as Audit levels, HTTP and DICOM profiles, link cost groups, and storage grades. Configuration is restricted to user accounts that have Grid Management permissions such as the Vendor account.
ILM Management	ILM (Information Lifecycle Management) policies. Configuration is restricted to user accounts that have Grid Management permissions such as the Vendor account.
NMS Management	NMS overview, custom alarms, alarm notifications, Admin Node name, and GUI timeout period. Configuration is restricted to user accounts that have Maintenance permissions such as the Admin and Vendor accounts.

Grid Topology Tree

The grid topology tree provides quick access to grid elements.

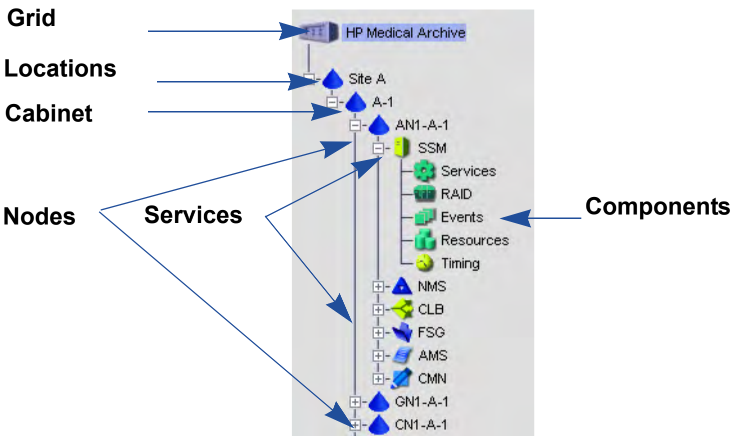

Figure 17 Grid Topology Tree

Occasionally the NMS MI fails to display the entire grid topology tree. The tree will be restored when the navigation frame is automatically refreshed. To prevent this, ensure the browser settings are correct. For more information, see Configure Internet Options Settings *(page 25).*

The highest level of detail is the grid as a whole, shown in the NMS MI as the root of the grid topology tree. Together, all of the other elements shown below constitute the grid. From the highest (big picture) to the lowest (most granular detail) the elements are:

- Grid
- Location
- Grid nodes
- Grid services
- Service components

The HP MAS can be deployed as a Single Site (Site A) or Single Site + DR (Site A and Site B) in a standard appliance configuration. In an Enterprise deployment, locations may be cities throughout the country, buildings within a city, or any other grouping. Each cabinet can be expanded to reveal one or more grid nodes, a grid node being a server hosting a collection of one or more grid services. A grid service consists of software components that deliver a particular capability.

Expand or Collapse the Grid Topology Tree

To expand and collapse the grid topology tree:

- In the grid topology tree, click ⊞ and ⊟.

  <Ctrl> clicking at either the location or nodes level opens (or closes) all items in the grid topology tree at the level clicked.

View Grid Elements

To view detailed information about a grid element:

- In the grid topology tree, click an element's name.

Naming Conventions

Names in the grid topology tree use the following convention:

- Locations:
  - Site A—The primary site for the HP MAS.
  - Site B —An optional Disaster Recovery (DR) site.
- Cabinets:
  - A-1 through A-8—Cabinets at the primary site.
  - B-1 through B-4—Cabinets at an optional DR site.
- Nodes—Named using the following elements:
  - Two or three-letter code for the type of node (such as: AN for Admin Node, GN for Gateway Node, or CSN for Control/Storage Node)
  - A sequence number (1 through 4) within the cabinet
  - The cabinet identifier (such as B-1, A-3, and so on)

Example: GN1-A-1 is the primary Gateway Node in cabinet A-1.

- Services—Named by the software with a three letter acronym, such as ADC for an Administrative Domain Controller service.
- Components—Named by the software

Content Tabs

Content on each page is organized under four tabs: Overview, Alarms, Reports, Configuration.

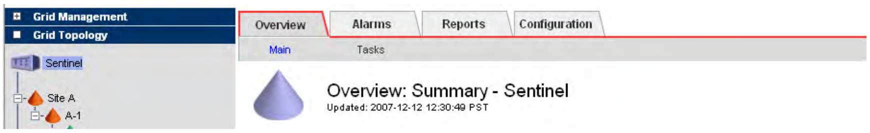

Content Tabs

- **Overview**—The Overview tab is used to monitor grid attributes. Each attribute represents a property, for example the number of managed objects, free storage space, backup size, or service state. These attributes are used to monitor normal grid operation and to detect and troubleshoot abnormal conditions. While there are hundreds of attributes, most of them are used for troubleshooting and only a small number must be monitored on a regular basis to ensure smooth operation. For examples on how to work with attributes, see Operations (page 103).
- **Alarms**—The Alarms tab is used to view and acknowledge alarms. For more information, see Alarms and State Indicators (page 39).
- **Reports**—The Reports tab is used to create charts and text reports. For more information, see Reports (page 51).
- **Configuration**—The Configuration tab is used to change configuration settings at the location, grid node, grid service, or component level. Configuration is restricted to user accounts that have Maintenance permissions such as the Admin and Vendor accounts.

Some tabs contain multiple pages. Click the page selector to access the content. The page currently selected is shown in blue and the other pages in black.

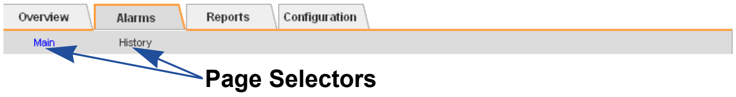

Figure 18 Page Selector

Attribute Description

The NMS MI contains a description of each attribute.

To find out more about each attribute:

- Click the attribute name to display its description. Click ☒ to close the description.

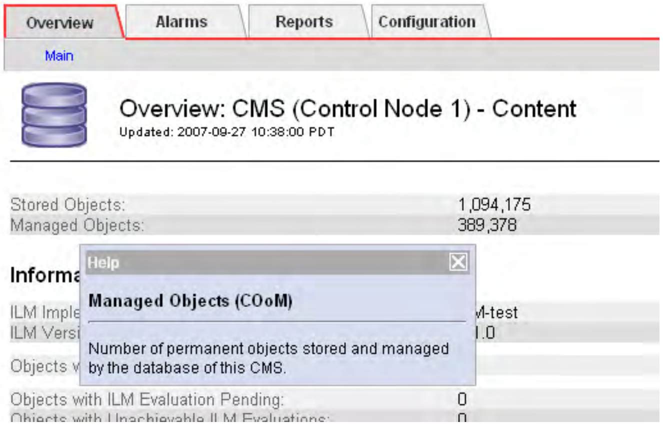

Overview: CMS (Control Node 1) - Content
Updated: 2007-09-27 10:38:00 PDT

Stored Objects:	1,094,175
Managed Objects:	389,378

Informa

ILM Imple	M-test
ILM Versi	1.0
Objects v	
Objects with ILM Evaluation Pending:	0
Objects with Unachievable ILM Evaluations:	0

Help ☒

Managed Objects (COoM)

Number of permanent objects stored and managed by the database of this CMS.

Figure 19 Attribute Description

Attribute Value Updates

The reporting of attributes is subject to propagation delays within the grid. Updated values for most attributes, except for state attributes, are sent to the NMS MI at fixed intervals. Therefore, it may take a few moments before an update is visible in the NMS MI, and two attributes that change more or less simultaneously may be reported at slightly different times.

Valid Characters

The NMS MI accepts only valid UTF-8 characters as user input in text fields.

Units of Measure

For units of “Seconds” or “Bytes”, the values displayed in the NMS MI are scaled to a suitable unit. For example, durations scale to microseconds, milliseconds, seconds, minutes, hours, or days; bytes scale to kilobytes, megabytes, or gigabytes.

NOTE The scale of bytes displayed by the NMS MI uses the “natural” measure of powers of 10. For example 3 MB = 3 x 106 = 3,000,000 bytes. This is not the same as powers of 2 normally used for computing, where 3 MiB = 3 x 220 = 3,145,728 bytes.

Apply Changes Button

To commit changes, for example to acknowledge alarms or change configuration settings, you must click the Apply Changes button at the bottom of the page. After you click the button, the button dims until changes are complete. Changes may take time to process. Do not click Apply Changes more than once. Wait for the page to refresh.

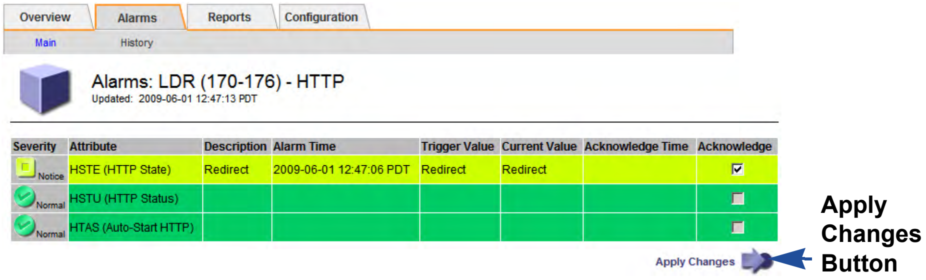

Alarms: LDR (170-176) - HTTP
Updated: 2009-06-01 12:47:13 PDT

Severity	Attribute	Description	Alarm Time	Trigger Value	Current Value	Acknowledge Time	Acknowledge
Notice	HSTE (HTTP State)	Redirect	2009-06-01 12:47:06 PDT	Redirect	Redirect		☑
Normal	HSTU (HTTP Status)						□
Normal	HTAS (Auto-Start HTTP)						□

Apply Changes Button

Figure 20 Apply Changes Button

To abort changes prior to clicking Apply Changes, simply refresh the page using the Refresh Page button at the top left in the header or the browser's refresh button.

Node Path Links

Clicking an underlined grid service path takes you to the node's Overview page. For example, in Figure 21, clicking Site A/170-41/FSG opens the FSG > Overview page for the primary Gateway Node.

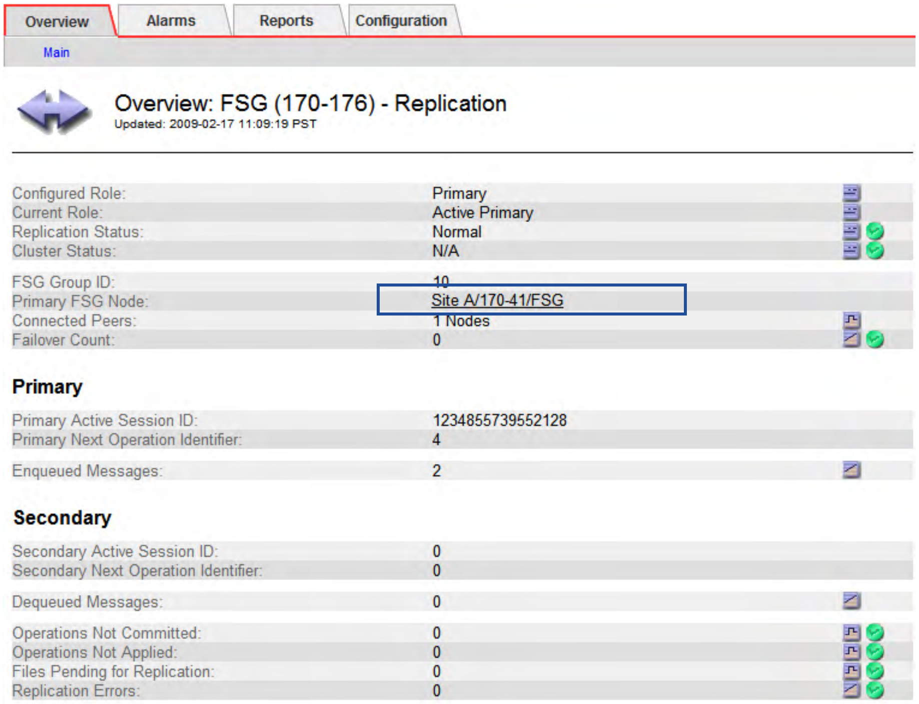

Overview: FSG (170-176) - Replication
Updated: 2009-02-17 11:09:19 PST

Configured Role:	Primary
Current Role:	Active Primary
Replication Status:	Normal
Cluster Status:	N/A
FSG Group ID:	10
Primary FSG Node:	Site A/170-41/FSG
Connected Peers:	1 Nodes
Failover Count:	0

Primary

Primary Active Session ID:	1234855739552128
Primary Next Operation Identifier:	4
Enqueued Messages:	2

Secondary

Secondary Active Session ID:	0
Secondary Next Operation Identifier:	0
Dequeued Messages:	0
Operations Not Committed:	0
Operations Not Applied:	0
Files Pending for Replication:	0
Replication Errors:	0

Figure 21 Node Path Links

User Accounts

The HP MAS system has two built-in user accounts:

Admin	Responsible for grid maintenance. The admin account can configure services and components but cannot make grid-wide changes.
Vendor	Responsible for grid configuration. The vendor account has full permissions.

The built-in accounts cannot be deleted. Additional accounts may exist on your grid depending on how it is configured. For example, the grid could have a read-only access intended for people who simply monitor the grid.

Permissions

Three built-in user groups (group accounts) have been configured for the HP MAS system: Vendor, Admin, and User.

Built-in groups are granted a collection of permissions. There are four types of permissions:

- Grid management
- Maintenance
- Alarm acknowledgement
- Accounts

Table 5 describes the allowable tasks for each set of permissions.

Table 5 User Groups Permissions

Permission Set	Allowable Tasks
Grid Management	• Configure grid-wide options • Configure ILM
Maintenance	• Configure FSGs • Configure the NMS MI (customize alarms, configure e-mail notifications, and configure GUI time-out) • Configure services and components
Alarm acknowledgement	• Acknowledge alarms
Accounts	• Create new accounts, configure existing accounts, and delete accounts • Create new user groups, configure existing user groups, and delete user groups

Figure 22 shows the built-in user accounts and group accounts. Depending on the Accounts permission on your account, you may not see all the accounts. See Figure 23 for a comparison.

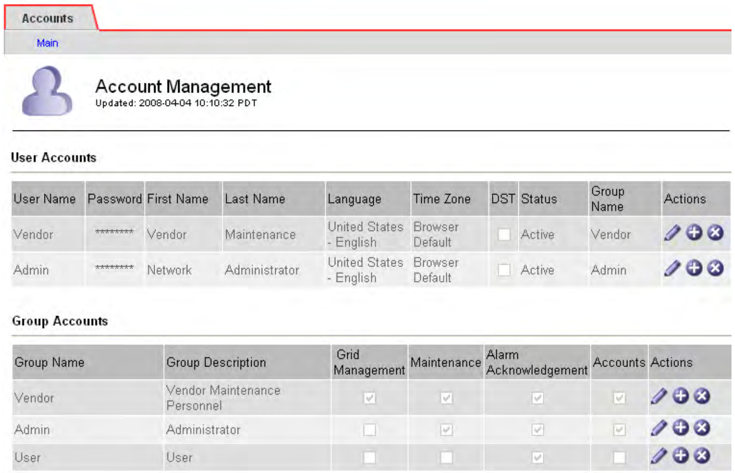

Account Management
Updated: 2008-04-04 10:10:32 PDT

User Accounts

User Name	Password	First Name	Last Name	Language	Time Zone	DST	Status	Group Name	Actions
Vendor	********	Vendor	Maintenance	United States - English	Browser Default	□	Active	Vendor	✎ ⊕ ⊗
Admin	********	Network	Administrator	United States - English	Browser Default	□	Active	Admin	✎ ⊕ ⊗

Group Accounts

Group Name	Group Description	Grid Management	Maintenance	Alarm Acknowledgement	Accounts	Actions
Vendor	Vendor Maintenance Personnel	☑	☑	☑	☑	✎ ⊕ ⊗
Admin	Administrator	□	☑	☑	☑	✎ ⊕ ⊗
User	User	□	□	☑	□	✎ ⊕ ⊗

Figure 22 Account Management for Vendor Account

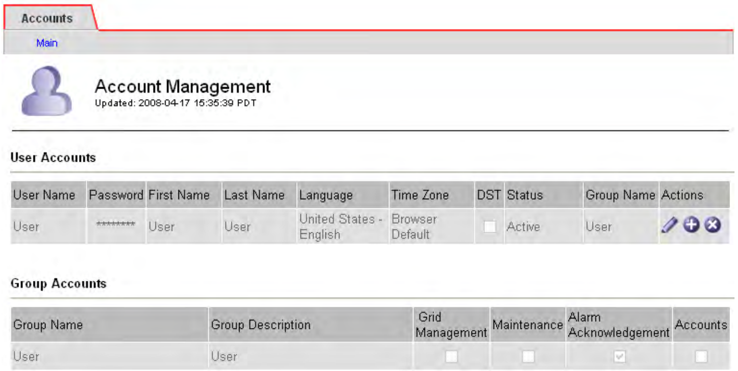

Account Management
Updated: 2008-04-17 15:35:39 PDT

User Accounts

User Name	Password	First Name	Last Name	Language	Time Zone	DST	Status	Group Name	Actions
User	********	User	User	United States - English	Browser Default	□	Active	User	✎ ⊕ ⊗

Group Accounts

Group Name	Group Description	Grid Management	Maintenance	Alarm Acknowledgement	Accounts
User	User	□	□	☑	□

Figure 23 Account Management for Account Without Accounts Permissions

Configuring Your Account

You can configure your account to change your password, first name, last name, and time zone.

Creating accounts and modifying accounts of other users is restricted to accounts that have "Accounts" permission such as the Admin and Vendor accounts.

To configure your account:

1. If you intend to change the password, make sure pop-ups are enabled. For more information, see Enable Pop-ups (page 26).
2. Go to **Grid Management > Account Management**.

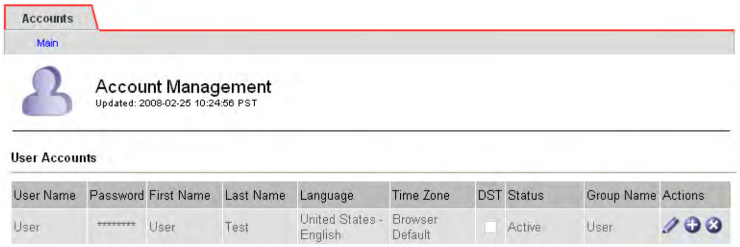

Account Management
Updated: 2008-02-25 10:24:56 PST

User Accounts

User Name	Password	First Name	Last Name	Language	Time Zone	DST	Status	Group Name	Actions
User	********	User	Test	United States - English	Browser Default	□	Active	User	✎ ⊕ ⊗

Figure 24 Account Management

3 Click Edit ✎ and update entries as needed:

- User Name—The user name entered at login. Read-only.
- Password—Masked password for the account; shown as a string of asterisks.
- First Name—User’s first name.
- Last Name—User’s last name.
- Language—The default language to be used for this user. At this time, only English is supported.
- Time Zone—Time zone of the NMS MI. By default, this is Browser Default which is the time of the computer from which the web browser is accessing the NMS MI. The time zone can be changed so that the time zone of the NMS MI is the same as the HP MAS system—particularly useful if the HP MAS system and the NMS MI are in different time zones.

  If you change the time zone settings of your computer, the change is effective the next time you log in to the NMS interface.
- DST—Click to set to Daylight Saving Time.

  You can modify the DST field only if Time Zone is not set to Browser Default. When Time Zone is set to Browser Default, daylight saving time is determined by the browser settings.

  Note that in order to update the DST box, Time Zone must be set to a time zone that supports Daylight Savings Time.
- Status:
  - Active—the user can log in and use the NMS MI.
  - Disabled — the user is prevented from logging in.

  Depending on your permissions, this field may be read-only.
- Group Name—Profile that governs the permitted activities for this user.

  Depending on your permissions, this field may be read-only.

4 To change your password:

a Double-click the Password box to select the complete field.

b Type a new password. Your password must contain between 8 and 32 characters and is case-sensitive.

c Press **<Tab>**. A confirmation pop-up window appears.

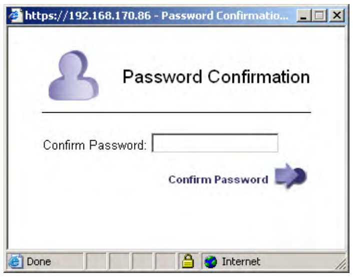

Figure 25 Password Confirmation Pop-up Window

d Re-enter the password and click Confirm Password. If the password fails to match, re-enter the password as prompted.

5 To change your name, edit the First Name and Last Name boxes.

6 To change the time zone, select a new time zone from the list.

7 Click **Apply Changes**.

Alarms and State Indicators

The color of the icon next to each location, cabinet, grid node, grid service, and service component in the grid topology tree reflects the overall status of that part of the grid.

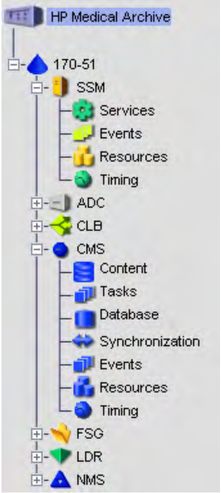

Figure 26 Grid Topology Tree with State and Alarm Colors

If there are no alarms and all services are connected, the icon appears in the normal (green) state. If there is an alarm, the color of the icon reflects the most severe alarm currently active on that branch of the tree. Locations display the highest alarm level of the grid nodes on that branch. Grid nodes display the color of the most severe state or alarm among their hosted services. Each individual service reflects the highest alarm severity of its components.

Service State Indicators

A service can have one of three states: Unknown, Connected, or Administratively Down. A service that is Unknown is problematic and must be investigated. A service that is Connected is operating normally and displays the color of its highest alarm severity—either itself or its components. A service that is Administratively Down has been deliberately shut down for maintenance by a grid administrator.

Table 6 Service States

Icon	Color	State	Meaning
?	Blue	Unknown	An unknown condition exists that has stopped normal operation. Requires immediate attention. The "Unknown" state is considered the most severe. It is typically used to indicate loss of connection between the NMS and a service.
✓	Green	Connected	All services are working normally.
?	Gray	Administratively Down	A service has been purposefully stopped. All alarms on the stopped service including acknowledged alarms are removed.

Alarm Indicators

A change in the value of an attribute can trigger an alarm. A change in the state of a service does not trigger an alarm.

An alarm is triggered when the value of an attribute reaches the alarm threshold value. When an alarm is triggered, the alarm information is displayed in the NMS MI and an e-mail notification is automatically sent to designated personnel.

Alarms are generated at the attribute level. There are five alarm severity levels displayed in the NMS MI. Each alarm level has an associated color and icon (see Table 7).

Table 7 Alarm Severity and Indicators

Severity Level	Icon	Color	State	Severity	Meaning
Lowest	✓	Green	Connected	Normal	All functions are working normally.
	🟨	Yellow	Connected	Notice	An unusual condition exists that does not affect normal operation.
	◆	Light Orange	Connected	Minor	An abnormal condition exists that could affect operation in the future; should be investigated to prevent escalation.
	⚠	Dark Orange	Connected	Major	An abnormal condition exists that currently affects operation; requires prompt attention to prevent escalation.
Highest	⊗	Red	Connected	Critical	A critical alert of an abnormal condition that has stopped normal operation; should be addressed immediately.

Propagation

Alarm Indicators

Alarms are generated at the attribute level. When an issue is detected, the alarm is propagated up through the grid topology tree. The associated attribute, component, service, node, and location information displayed in the NMS MI all change to reflect the alarm's severity. The color displayed reflects the most severe alarm currently active on that branch of the grid topology tree. As a result, you can view the general alarm severity level at the grid level, then drill down through the service components to locate the specific details.

For example, in Figure 27, the SSM service has at least two alarms: the Events component has at least one alarm with a severity of Notice and the Resources component has at least one alarm with a severity of Minor Alert. Minor Alert is the more severe of the two alarms and therefore it propagates up the grid topology tree so that the SSM service takes on the Minor Alert alarm color, light orange.

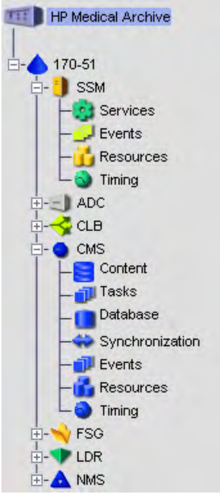

Figure 27 Propagation of Alarm and State Indicators

State Indicators

State indicators are displayed at the services level and above and take priority over alarm indicators. When a service enters either an Administratively Down or Unknown state, the state color is propagated down through the service to its components. This overrides any component alarm indicators displayed in the grid. For example, a service state of "Unknown" supersedes an alarm severity of Critical and results in the service displaying the Unknown state color and not the Critical alarm severity color. The state of the service also propagates up to the node level.

For example, in Figure 27, the CMS service has a state of "Unknown." This state overrides any alarms that may have been raised on any of it components. "Unknown" is a more severe state than any alarm and therefore propagates down the grid topology tree so that all CMS components take on the "Unknown" state color, blue. The grid node also displays the "Unknown" state color of blue as this is the most critical state of any of its services. The state condition of a node's services propagates up the grid topology tree in the same manner that alarms do.

E-mail Notifications

E-mail notifications are automatically sent to designated personnel to alert recipients that an alarm has been triggered or a service state has changed. Managing e-mail notifications is restricted to user accounts that have Maintenance permissions such as the Admin and Vendor accounts.

E-mail Notification Status

You can view the status of e-mail notifications at any time via the **<Admin_Node> > NMS > Overview** page. The E-mail Notifications Status attribute displays the selected Admin Node's current ability to send e-mail notifications to the mail server.

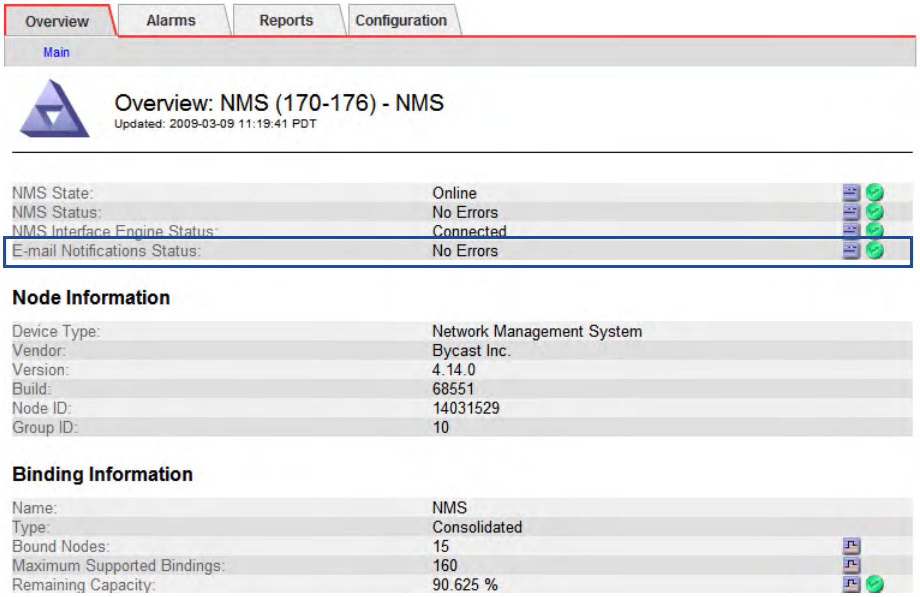

Overview: NMS (170-176) - NMS

Updated: 2009-03-09 11:19:41 PDT

NMS State:	Online
NMS Status:	No Errors
NMS Interface Engine Status:	Connected
E-mail Notifications Status:	No Errors

Node Information

Device Type:	Network Management System
Vendor:	Bycast Inc.
Version:	4.14.0
Build:	68551
Node ID:	14031529
Group ID:	10

Binding Information

Name:	NMS
Type:	Consolidated
Bound Nodes:	15
Maximum Supported Bindings:	160
Remaining Capacity:	90.625 %

Figure 28 NMS Notification Status

If there is an error, the selected NMS service cannot send e-mail notifications to the mail server. Depending on grid configuration, this may mean that the NMS service is not sending notifications and that switch-over to another NMS service has occurred (if the grid has multiple Admin Nodes).

Alarm Customization

The NMS MI is configured with a set of default alarms. In addition, it is possible to create custom alarms at the service or component level, or at the grid level.

The **Configuration > Alarms** page of each service or component is used to view configured Default alarms and Global Custom alarms and to create Custom alarms for a service. Access to this page is restricted to user accounts that have Maintenance permissions such as the Admin and Vendor accounts.

The **Grid Management > NMS Management** area is used to create Global Custom alarms and to enable or disable Default alarms globally. Access to this page is restricted to user accounts that have Grid Management permissions such as the Vendor account.

Reviewing Alarms

Table 8 summarizes how to review information on the current status of alarms. This information is located on the:

- System Status page
- Alarms tab for each component and service

- Overview tab for each component and service

Table 8 Reviewing Alarms

<table>
<tr><th>To:</th><th>Do this:</th></tr>
<tr><td>Get a list of all current alarms in the grid</td><td>

- Click **System Status** System Status ⚠ in the header to display alarms. Alarms are sorted by severity.

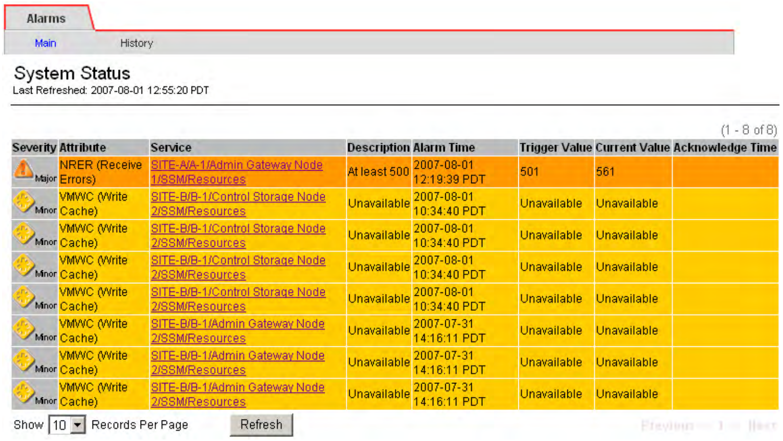

Alarms

Main History

System Status

Last Refreshed: 2007-08-01 12:55:20 PDT

(1 - 8 of 8)

<table>
<tr><th>Severity</th><th>Attribute</th><th>Service</th><th>Description</th><th>Alarm Time</th><th>Trigger Value</th><th>Current Value</th><th>Acknowledge Time</th></tr>
<tr><td>⚠ Major</td><td>NRER (Receive Errors)</td><td>SITE-A/A-1/Admin Gateway Node 1/SSM/Resources</td><td>At least 500</td><td>2007-08-01 12:19:39 PDT</td><td>501</td><td>561</td><td></td></tr>
<tr><td>Minor</td><td>VMWC (Write Cache)</td><td>SITE-B/B-1/Control Storage Node 2/SSM/Resources</td><td>Unavailable</td><td>2007-08-01 10:34:40 PDT</td><td>Unavailable</td><td>Unavailable</td><td></td></tr>
<tr><td>Minor</td><td>VMWC (Write Cache)</td><td>SITE-B/B-1/Control Storage Node 2/SSM/Resources</td><td>Unavailable</td><td>2007-08-01 10:34:40 PDT</td><td>Unavailable</td><td>Unavailable</td><td></td></tr>
<tr><td>Minor</td><td>VMWC (Write Cache)</td><td>SITE-B/B-1/Control Storage Node 2/SSM/Resources</td><td>Unavailable</td><td>2007-08-01 10:34:40 PDT</td><td>Unavailable</td><td>Unavailable</td><td></td></tr>
<tr><td>Minor</td><td>VMWC (Write Cache)</td><td>SITE-B/B-1/Control Storage Node 2/SSM/Resources</td><td>Unavailable</td><td>2007-08-01 10:34:40 PDT</td><td>Unavailable</td><td>Unavailable</td><td></td></tr>
<tr><td>Minor</td><td>VMWC (Write Cache)</td><td>SITE-B/B-1/Admin Gateway Node 2/SSM/Resources</td><td>Unavailable</td><td>2007-07-31 14:16:11 PDT</td><td>Unavailable</td><td>Unavailable</td><td></td></tr>
<tr><td>Minor</td><td>VMWC (Write Cache)</td><td>SITE-B/B-1/Admin Gateway Node 2/SSM/Resources</td><td>Unavailable</td><td>2007-07-31 14:16:11 PDT</td><td>Unavailable</td><td>Unavailable</td><td></td></tr>
<tr><td>Minor</td><td>VMWC (Write Cache)</td><td>SITE-B/B-1/Admin Gateway Node 2/SSM/Resources</td><td>Unavailable</td><td>2007-07-31 14:16:11 PDT</td><td>Unavailable</td><td>Unavailable</td><td></td></tr>
</table>

Show 10 Records Per Page Refresh Previous 1 Next

NOTE The first thing to look at in the NMS MI is the System Status indicator. It immediately tells you the most serious status (state or alarm) of the grid.

</td></tr>
</table>

Table 8 Reviewing Alarms *(continued)*

To:	Do this:
Get a list of all alarms triggered over a period of time	1 Click **System Status** in the header. 2 Click the **History** page. 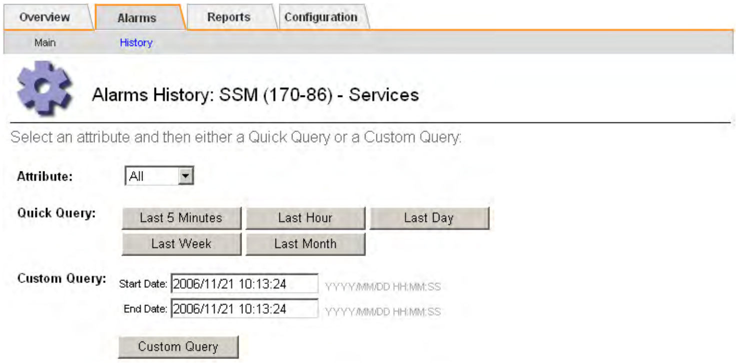 Overview Alarms Reports Configuration Main History Last 5 Minutes Last Hour Last Day Last Week Last Month Custom Query Alarms History: SSM (170-86) - Services Select an attribute and then either a Quick Query or a Custom Query. Attribute: All Quick Query: Custom Query: Start Date: 2006/11/21 10:13:24 YYYY/MM/DD HH:MM:SS End Date: 2006/11/21 10:13:24 YYYY/MM/DD HH:MM:SS 3 Do one of the following: — Click one of the time periods. — Enter a custom range and click **Custom Query**.

Table 8 Reviewing Alarms *(continued)*

To:	Do this:
Find out more about an alarm	1 Click the service path in the System Status table to go to the Alarms tab for the selected alarm (or use the grid topology tree). The Alarms tab displays the current status of the attributes for the selected service or component. The colors and icons reflect the severity levels of the alarms. After an alarm is resolved, the alarm returns to the green "Normal" severity level.

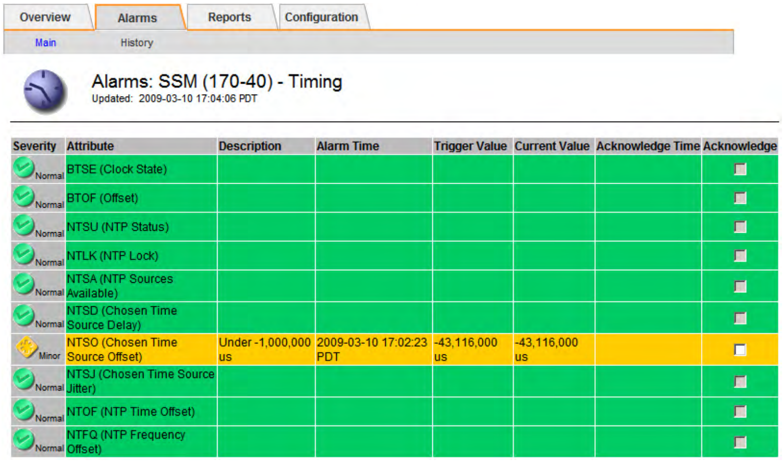

Alarms: SSM (170-40) - Timing
Updated: 2009-03-10 17:04:06 PDT

Severity	Attribute	Description	Alarm Time	Trigger Value	Current Value	Acknowledge Time	Acknowledge
Normal	BTSE (Clock State)						□
Normal	BTOF (Offset)						□
Normal	NTSU (NTP Status)						□
Normal	NTLK (NTP Lock)						□
Normal	NTSA (NTP Sources Available)						□
Normal	NTSD (Chosen Time Source Delay)						□
Minor	NTSO (Chosen Time Source Offset)	Under -1,000,000 us	2009-03-10 17:02:23 PDT	-43,116,000 us	-43,116,000 us		□
Normal	NTSJ (Chosen Time Source Jitter)						□
Normal	NTOF (NTP Time Offset)						□
Normal	NTFQ (NTP Frequency Offset)						□

See Table 9 (page 48) for a description of the fields.

2 Click the alarm to display a description of the attribute.

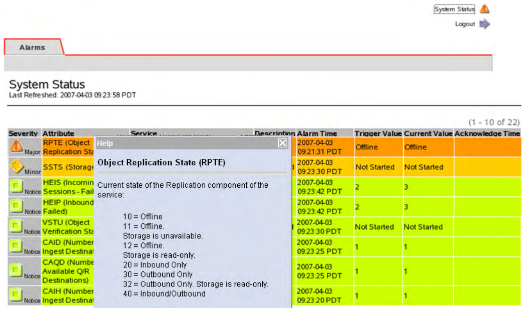

System Status
Last Refreshed: 2007-04-03 09:23:58 PDT

(1 - 10 of 22)

Severity	Attribute	Service	Description	Alarm Time	Trigger Value	Current Value	Acknowledge Time
Major	RPTE (Object Replication Sta			2007-04-03 09:21:31 PDT	Offline	Offline	
Minor	SSTS (Storage			2007-04-03 09:23:30 PDT	Not Started	Not Started	
Notice	HEIS (Incomin Sessions - Fail			2007-04-03 09:23:42 PDT	2	3	
Notice	HEIP (Inbound Failed)			2007-04-03 09:23:42 PDT	2	3	
Notice	VSTU (Object Verification Sta			2007-04-03 09:23:30 PDT	Not Started	Not Started	
Notice	CAID (Number Ingest Destinat			2007-04-03 09:23:25 PDT	1	1	
Notice	CAQD (Numbe Available Q/R Destinations)			2007-04-03 09:23:25 PDT	1	1	
Notice	CAIH (Number Ingest Destinat			2007-04-03 09:23:20 PDT	1	1	

Object Replication State (RPTE)

Current state of the Replication component of the service:

10 = Offline
11 = Offline.
Storage is unavailable.
12 = Offline.
Storage is read-only.
20 = Inbound Only
30 = Outbound Only
32 = Outbound Only. Storage is read-only.
40 = Inbound/Outbound

3 Look up the four-character attribute in the reference table in Chapter 10, Alarm Reference.

4 Click the **Overview** tab.

5 Locate the alarm, and if it can be charted, click the chart button to view a trend of the attribute over the last hour. Adjust the time period as required. For more information, see Displaying Charts (page 53).

Table 8 Reviewing Alarms *(continued)*

<table>
<tr><th>To:</th><th>Do this:</th></tr>
<tr><td>Find out how often alarms have been triggered for a particular attribute</td><td>
1 Go to the service or component that has the attribute.
2 Click the **Alarms** tab and then the **History** page.
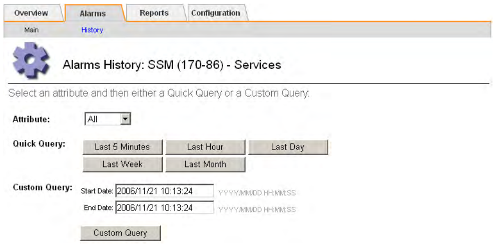

Alarms History: SSM (170-86) - Services
Select an attribute and then either a Quick Query or a Custom Query:
Attribute: All
Quick Query:
Custom Query: Start Date: 2006/11/21 10:13:24 YYYY/MM/DD HH:MM:SS
End Date: 2006/11/21 10:13:24 YYYY/MM/DD HH:MM:SS
3 Select the attribute from the list.
4 Do one of the following:
— Click one of the time periods.
— Enter a custom range and click **Custom Query**.
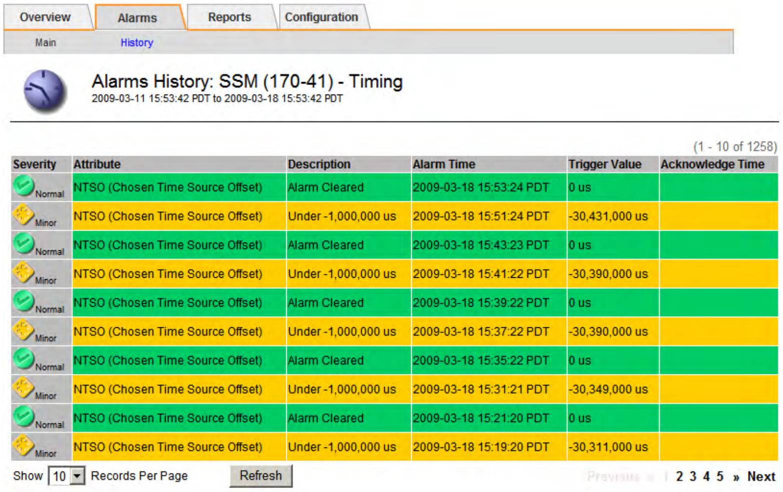

Alarms History: SSM (170-41) - Timing
2009-03-11 15:53:42 PDT to 2009-03-18 15:53:42 PDT
(1 - 10 of 1258)
<table>
<tr><th>Severity</th><th>Attribute</th><th>Description</th><th>Alarm Time</th><th>Trigger Value</th><th>Acknowledge Time</th></tr>
<tr><td>Normal</td><td>NTSO (Chosen Time Source Offset)</td><td>Alarm Cleared</td><td>2009-03-18 15:53:24 PDT</td><td>0 us</td><td></td></tr>
<tr><td>Minor</td><td>NTSO (Chosen Time Source Offset)</td><td>Under -1,000,000 us</td><td>2009-03-18 15:51:24 PDT</td><td>-30,431,000 us</td><td></td></tr>
<tr><td>Normal</td><td>NTSO (Chosen Time Source Offset)</td><td>Alarm Cleared</td><td>2009-03-18 15:43:23 PDT</td><td>0 us</td><td></td></tr>
<tr><td>Minor</td><td>NTSO (Chosen Time Source Offset)</td><td>Under -1,000,000 us</td><td>2009-03-18 15:41:22 PDT</td><td>-30,390,000 us</td><td></td></tr>
<tr><td>Normal</td><td>NTSO (Chosen Time Source Offset)</td><td>Alarm Cleared</td><td>2009-03-18 15:39:22 PDT</td><td>0 us</td><td></td></tr>
<tr><td>Minor</td><td>NTSO (Chosen Time Source Offset)</td><td>Under -1,000,000 us</td><td>2009-03-18 15:37:22 PDT</td><td>-30,390,000 us</td><td></td></tr>
<tr><td>Normal</td><td>NTSO (Chosen Time Source Offset)</td><td>Alarm Cleared</td><td>2009-03-18 15:35:22 PDT</td><td>0 us</td><td></td></tr>
<tr><td>Minor</td><td>NTSO (Chosen Time Source Offset)</td><td>Under -1,000,000 us</td><td>2009-03-18 15:31:21 PDT</td><td>-30,349,000 us</td><td></td></tr>
<tr><td>Normal</td><td>NTSO (Chosen Time Source Offset)</td><td>Alarm Cleared</td><td>2009-03-18 15:21:20 PDT</td><td>0 us</td><td></td></tr>
<tr><td>Minor</td><td>NTSO (Chosen Time Source Offset)</td><td>Under -1,000,000 us</td><td>2009-03-18 15:19:20 PDT</td><td>-30,311,000 us</td><td></td></tr>
</table>
Show 10 Records Per Page
The alarms are listed in reverse chronological order. See Table 9 for a description of the fields.
5 To return to the alarms history request form, click **History**.
</td></tr>
</table>

Table 9 Alarms Table Fields

Field	Description
Severity	Color icon indicating the alarm severity level.
Attribute	Code that identifies the attribute and issue being monitored. See Chapter 10, Alarm Reference for an alphabetical listing of the attributes and reference information on each alarm.
Description	Brief details about the cause of the alarm.

Table 9 Alarms Table Fields *(continued)*

Field	Description
Alarm Time	The date and time in your local time zone at which the trigger value was reported. This field is blank if the alarm has been acknowledged.
Trigger Value	Value that triggered the alarm.
Current Value	Value of the attribute as last reported by the service or component.
Acknowledge Time	Date and time the alarm was acknowledged.
Acknowledge	Selecting the **Acknowledge** check box acknowledges the alarm. See the procedure on page 50.

Acknowledging Alarms

Depending on the situation, you may choose to acknowledge alarms while you are trying to resolve the underlying issue.

Acknowledging alarms is restricted to user accounts that have Alarm Acknowledgement permissions such as the Admin and Vendor accounts.

An acknowledged alarm continues to display as an alarm at the component level on the System Status page. However, after an alarm has been acknowledged, it no longer propagates up the grid topology tree. The grid topology tree is displayed as "Normal" (green) or the color of the next most severe unacknowledged alarm or more severe service state (see Figure 29).

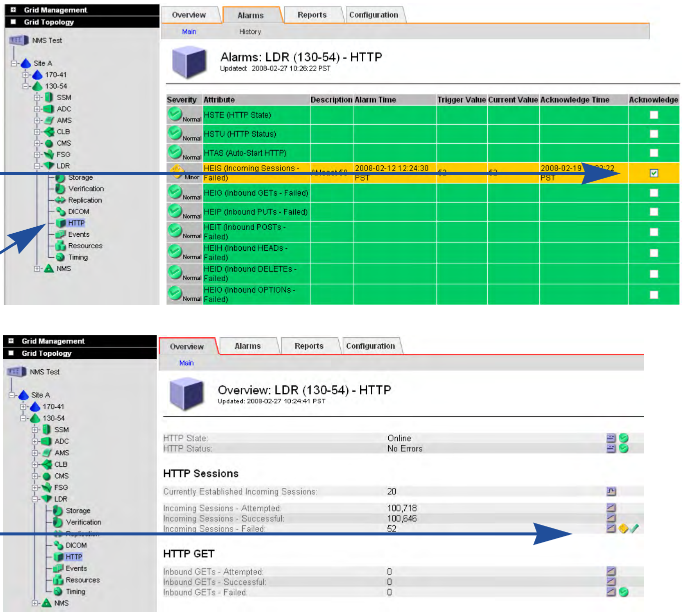

Alarms: LDR (130-54) - HTTP
Updated: 2008-02-27 10:26:22 PST

Severity	Attribute	Description	Alarm Time	Trigger Value	Current Value	Acknowledge Time	Acknowledge
Normal	HSTE (HTTP State)						□
Normal	HSTU (HTTP Status)						□
Normal	HTAS (Auto-Start HTTP)						□
Minor	HEIS (Incoming Sessions - Failed)	At least 50	2008-02-12 12:24:30 PST	52	52	2008-02-19 [illegible]:22 PST	☑
Normal	HEIG (Inbound GETs - Failed)						□
Normal	HEIP (Inbound PUTs - Failed)						□
Normal	HEIT (Inbound POSTs - Failed)						□
Normal	HEIH (Inbound HEADs - Failed)						□
Normal	HEID (Inbound DELETEs - Failed)						□
Normal	HEIO (Inbound OPTIONs - Failed)						□

Overview: LDR (130-54) - HTTP
Updated: 2008-02-27 10:24:41 PST

HTTP State:	Online
HTTP Status:	No Errors

HTTP Sessions

Currently Established Incoming Sessions:	20
Incoming Sessions - Attempted:	100,718
Incoming Sessions - Successful:	100,646
Incoming Sessions - Failed:	52

HTTP GET

Inbound GETs - Attempted:	0
Inbound GETs - Successful:	0
Inbound GETs - Failed:	0

Minor alarm has been acknowledged on the Alarms tab

Component state is "Normal" even though a minor alarm has been triggered

Minor alarm with the Acknowledge check icon appears on the Overview tab

Figure 29 Acknowledging Alarms

There are many reasons why you may want to acknowledge an alarm. For instance, while testing or troubleshooting the grid, you may want to hide (by acknowledging) alarms that you are aware of in order to better track unknown issues. Or, you may, because of time constraints, want to acknowledge an alarm that you can more effectively attend to later.

When a grid is restarted, all unacknowledged alarms are reset to “Normal.” Acknowledged alarms remain unchanged unless the condition that has triggered the alarm has changed and the alarm has reached a new severity level. If the previously acknowledged alarm is triggered, it is considered a new alarm and can be re-acknowledged.

To acknowledge an alarm:

1 Go to the service or component that you are interested in.

2 Click the **Alarms** tab.

3 Select Acknowledge next to the alarm.

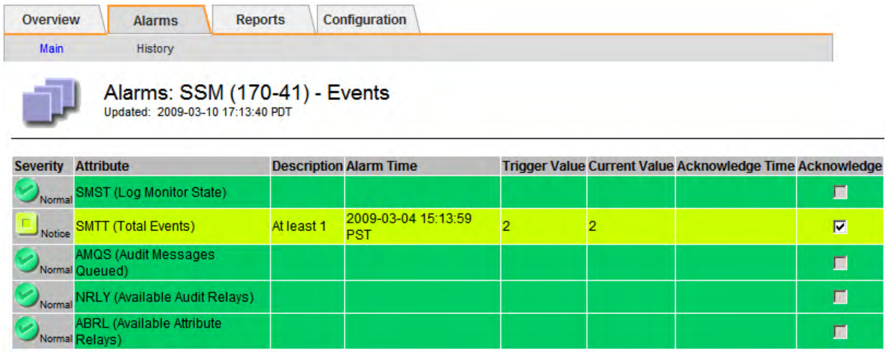

Alarms: SSM (170-41) - Events
Updated: 2009-03-10 17:13:40 PDT

Severity	Attribute	Description	Alarm Time	Trigger Value	Current Value	Acknowledge Time	Acknowledge
Normal	SMST (Log Monitor State)						□
Notice	SMTT (Total Events)	At least 1	2009-03-04 15:13:59 PST	2	2		☑
Normal	AMQS (Audit Messages Queued)						□
Normal	NRLY (Available Audit Relays)						□
Normal	ABRL (Available Attribute Relays)						□

Figure 30 Sample Acknowledged Alarm

4 Click **Apply Changes**.

The alarm is acknowledged and a notification is sent to designated personnel.

To unacknowledge an alarm:

1 Go to the service or component that you are interested in.

2 Click the **Alarms** tab.

3 Clear Acknowledge next to the alarm.

4 Click **Apply Changes**.

The alarm is unacknowledged and a notification is sent to designated personnel.

Reports

Reports are an invaluable tool you can use to monitor the state of the grid and to troubleshoot problems. There are two types of reports: chart reports and text reports.

Charts

Chart reports present the data with the attribute value (vertical axis) over a specified time span (horizontal axis).

Chart Types

There are three types of charts:

- line graph
- area graph
- state graph

Line Graph

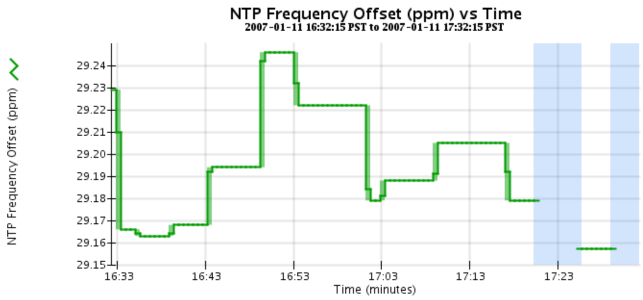

Figure 31 Line Graph

Line graphs are used to plot the values of an attribute that has a “unit” value (such as NTP Frequency Offset, in ppm). The changes in the value are plotted in bins at regular intervals over time.

Area Graph

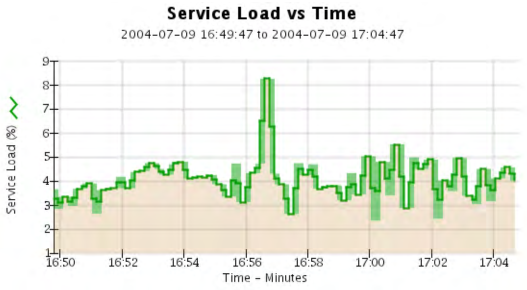

Figure 32 Area Graph

Area graphs are used to plot volumetric quantities, file count or service load values for instance. Area graphs are similar to line graphs but include a light brown shading below the line. The changes in the value are plotted in bins at regular intervals over time.

State Graph

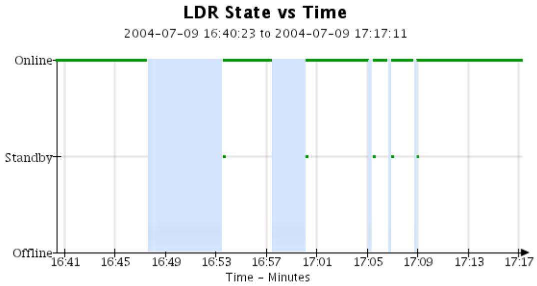

Figure 33 State Graph

State graphs are used to plot values that represent distinct states such as a service state that can be online, standby, or offline. State graphs are similar to line graphs but the transition is discontinuous, that is, the value jumps from one state value to another.

Interpreting Chart Colors

The chart colors have a specific meaning. Table 10 describes how to interpret the various colors and line types.

Table 10 Chart Colors and Shading

Sample	Meaning
▬	Reported attribute values are plotted using dark green lines.
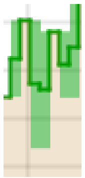	Light green shading around dark green lines indicates that the actual values in that time range vary and have been “binned” for faster plotting. The dark line represents the weighted average. The range in light green indicates the maximum and minimum values within the bin. Light brown shading is used for area graphs to indicate volumetric data.
	Blank areas (no data plotted) indicate that the attribute values were unavailable. The background may be blue, gray, or a mixture of gray and blue depending on the state of the service reporting the attribute.

Table 10 Chart Colors and Shading *(continued)*

Sample	Meaning
🟦	Light blue shading indicates that some or all of the attribute values at that time were indeterminate; the attribute was not reporting because the service was in an unknown state.
⬜	Gray shading indicates that some or all of the attribute values at that time were not known because the service reporting the attributes was administratively down.
▦	A mixture of gray and blue shading indicates that some of the attribute values at the time were indeterminate (because the service was in an unknown state), while others were not known because the service reporting the attributes was administratively down.

Displaying Charts

In most cases, the fastest way to create a chart is to go to an Overview page and click the chart button next to the attribute. Clicking this chart icon will immediately take you to the Reports > Charts page and display a chart for the attribute. This is known as an immediate chart.

You can also manually create reports from the Report tab.

NOTE There are some attributes for which it is not possible to create charts, for example, text attributes such as Node ID, version number, and build number.

Create an Immediate Chart Report

To create an immediate chart from the Overview page:

1 Go to the component or service that has the attribute you are interested in.

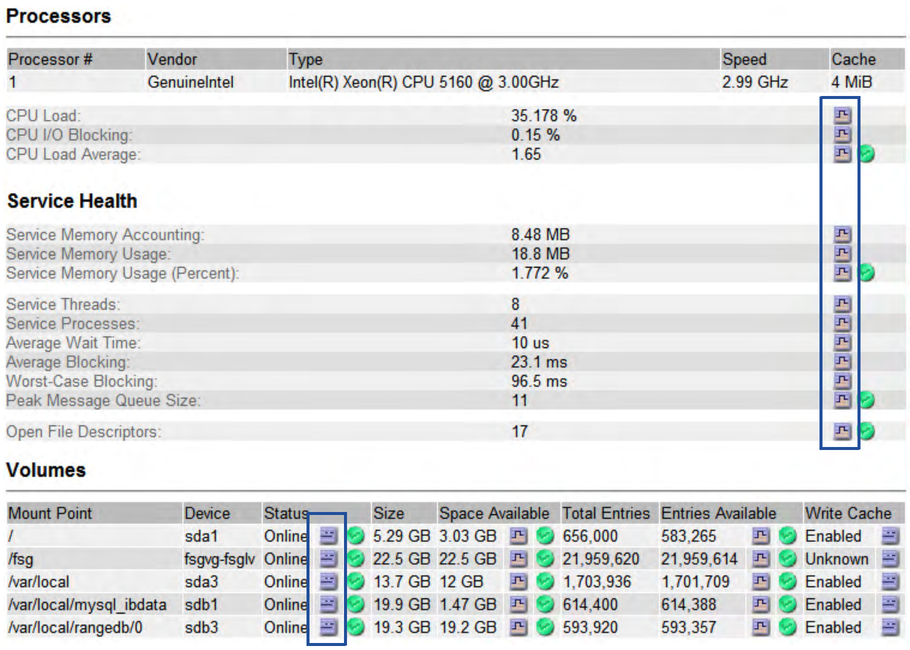

Processors

Processor #	Vendor	Type	Speed	Cache
1	GenuineIntel	Intel(R) Xeon(R) CPU 5160 @ 3.00GHz	2.99 GHz	4 MiB

CPU Load:	35.178 %
CPU I/O Blocking:	0.15 %
CPU Load Average:	1.65

Service Health

Service Memory Accounting:	8.48 MB
Service Memory Usage:	18.8 MB
Service Memory Usage (Percent):	1.772 %
Service Threads:	8
Service Processes:	41
Average Wait Time:	10 us
Average Blocking:	23.1 ms
Worst-Case Blocking:	96.5 ms
Peak Message Queue Size:	11
Open File Descriptors:	17

Volumes

Mount Point	Device	Status	Size	Space Available	Total Entries	Entries Available	Write Cache
/	sda1	Online	5.29 GB	3.03 GB	656,000	583,265	Enabled
/fsg	fsgvg-fsglv	Online	22.5 GB	22.5 GB	21,959,620	21,959,614	Unknown
/var/local	sda3	Online	13.7 GB	12 GB	1,703,936	1,701,709	Enabled
/var/local/mysql_ibdata	sdb1	Online	19.9 GB	1.47 GB	614,400	614,388	Enabled
/var/local/rangedb/0	sdb3	Online	19.3 GB	19.2 GB	593,920	593,357	Enabled

Figure 34 Immediate Chart Buttons

2 Click the chart button next to the attribute to display a chart.

The display automatically changes to the Reports > Charts page. The chart displays the attribute's data over the past hour. To view other time ranges, follow the procedure described in Manually Create a Chart Report.

Manually Create a Chart Report

To manually create a chart from the Reports tab:

1 Click the Reports tab of the service or component you are interested in.

2 From the **Attribute** pull-down menu, select an attribute.

3 To force the Y-axis to start at zero, clear Vertical Scaling.

4 To show values at full precision, select Raw Values. To round values to a maximum of three decimal places (for example, for attributes reported as percentages), clear Raw Data.

5 From the **Quick Query** pull-down menu, select a time period.

The chart appears after a few moments. Allow several minutes for long time ranges.

To display a chart for a custom time period:

a From the Quick Query pull-down menu, select Custom Query.

b Enter the Start Date and End Date.

Use the format YYYY/MM/DD HH:MM:SS in local time. Leading zeros are required to match the format. For example, 2007/4/8 7:30:00 fails validation; the correct format is 2007/04/08 07:30:00.

c Click **Update**.

Displaying Charts in a New Window

When you generate a chart report, it is often useful to compare it to another chart. The NMS MI provides the ability to view chart data in a new window. Multiple windows can be opened.

To open a new chart window:

- Click ⧉ to display the current view in a new window.
- Click ☒ to close the chart windows.

Text Reports

A text report displays a textual representation of attribute data values that have been processed by the NMS service. For attribute data that is expected to be continuously changing, this attribute data is sampled by the NMS service (at the source) at regular intervals. For attribute data that changes infrequently (for example, data based on events such as state or status changes) an attribute value is sent to the NMS service when the value changes.

There are two types of text reports: raw and aggregate. The type of report displayed depends on the configured time period. By default, aggregate text reports are generated for time periods longer than one week.

Grey text indicates the service was administratively down during the time it was sampled. Blue text indicates the service was in an unknown state.

Raw Text Report

A raw text report displays the following information:

- **Time Received**—Local date and time that a sample value of an attribute’s data was processed by the NMS service.
- **Sample Time**—Local date and time that an attribute value was sampled or changed at the source.
- **Value**—Attribute value at sample time.

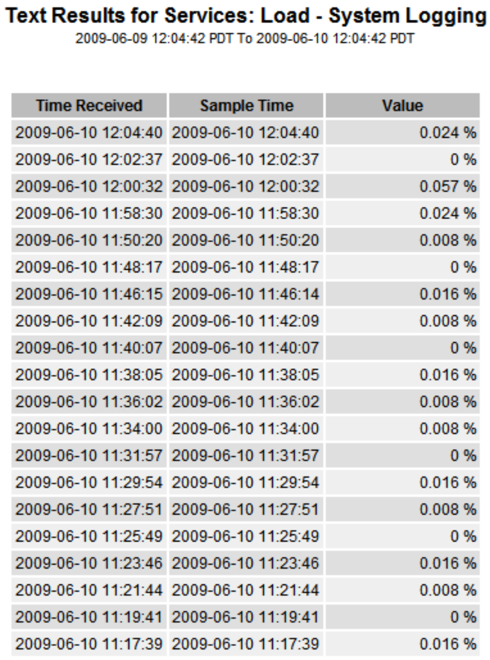

Text Results for Services: Load - System Logging

2009-06-09 12:04:42 PDT To 2009-06-10 12:04:42 PDT

Time Received	Sample Time	Value
2009-06-10 12:04:40	2009-06-10 12:04:40	0.024 %
2009-06-10 12:02:37	2009-06-10 12:02:37	0 %
2009-06-10 12:00:32	2009-06-10 12:00:32	0.057 %
2009-06-10 11:58:30	2009-06-10 11:58:30	0.024 %
2009-06-10 11:50:20	2009-06-10 11:50:20	0.008 %
2009-06-10 11:48:17	2009-06-10 11:48:17	0 %
2009-06-10 11:46:15	2009-06-10 11:46:14	0.016 %
2009-06-10 11:42:09	2009-06-10 11:42:09	0.008 %
2009-06-10 11:40:07	2009-06-10 11:40:07	0 %
2009-06-10 11:38:05	2009-06-10 11:38:05	0.016 %
2009-06-10 11:36:02	2009-06-10 11:36:02	0.008 %
2009-06-10 11:34:00	2009-06-10 11:34:00	0.008 %
2009-06-10 11:31:57	2009-06-10 11:31:57	0 %
2009-06-10 11:29:54	2009-06-10 11:29:54	0.016 %
2009-06-10 11:27:51	2009-06-10 11:27:51	0.008 %
2009-06-10 11:25:49	2009-06-10 11:25:49	0 %
2009-06-10 11:23:46	2009-06-10 11:23:46	0.016 %
2009-06-10 11:21:44	2009-06-10 11:21:44	0.008 %
2009-06-10 11:19:41	2009-06-10 11:19:41	0 %
2009-06-10 11:17:39	2009-06-10 11:17:39	0.016 %

Figure 35 Raw Text Report

Aggregate Text Report

An aggregate text report displays data over a longer period of time (usually a week) than a raw text report. Each entry is the result of summarizing multiple attribute values (an aggregate of attribute values) by the NMS service over time into a single entry with an average, maximum and minimum value that is derived from the aggregation.

Each entry displays the following information:

- **Aggregate Time**—Last local date and time that the NMS service aggregated (collected) a set of changed attribute values.
- **Average Value**—The average of the attribute’s value over the aggregated time period.
- **Minimum Value**—The minimum value processed over the aggregated time period.
- **Maximum Value**—The maximum value processed over the aggregated time period.

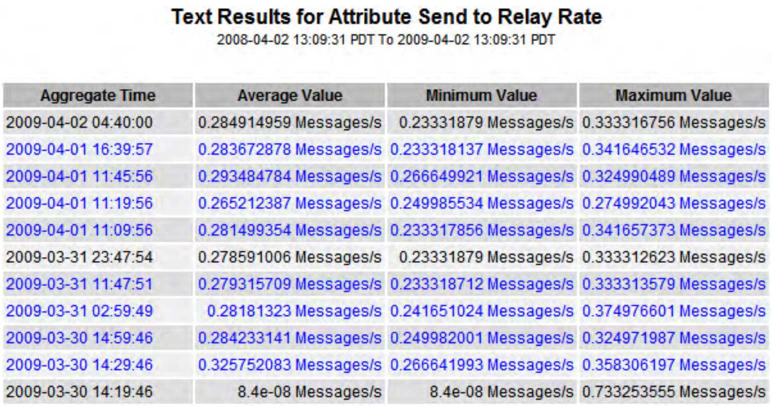

Text Results for Attribute Send to Relay Rate

2008-04-02 13:09:31 PDT To 2009-04-02 13:09:31 PDT

Aggregate Time	Average Value	Minimum Value	Maximum Value
2009-04-02 04:40:00	0.284914959 Messages/s	0.23331879 Messages/s	0.333316756 Messages/s
2009-04-01 16:39:57	0.283672878 Messages/s	0.233318137 Messages/s	0.341646532 Messages/s
2009-04-01 11:45:56	0.293484784 Messages/s	0.266649921 Messages/s	0.324990489 Messages/s
2009-04-01 11:19:56	0.265212387 Messages/s	0.249985534 Messages/s	0.274992043 Messages/s
2009-04-01 11:09:56	0.281499354 Messages/s	0.233317856 Messages/s	0.341657373 Messages/s
2009-03-31 23:47:54	0.278591006 Messages/s	0.23331879 Messages/s	0.333312623 Messages/s
2009-03-31 11:47:51	0.279315709 Messages/s	0.233318712 Messages/s	0.333313579 Messages/s
2009-03-31 02:59:49	0.28181323 Messages/s	0.241651024 Messages/s	0.374976601 Messages/s
2009-03-30 14:59:46	0.284233141 Messages/s	0.249982001 Messages/s	0.324971987 Messages/s
2009-03-30 14:29:46	0.325752083 Messages/s	0.266641993 Messages/s	0.358306197 Messages/s
2009-03-30 14:19:46	8.4e-08 Messages/s	8.4e-08 Messages/s	0.733253555 Messages/s

Figure 36 Aggregate Text Report

Create a Text Report

To create a text report:

1. Click the **Reports** tab of the service or component you are interested in.
2. Click **Text**.
3. From the **Attribute** pull-down menu, select an attribute.
4. From the **Results per Page** pull-down menu, select the number of values reported per page.
5. To round values to a maximum of three decimal places (for example, for attributes reported as percentages), clear Raw Data.
6. From the Quick Query pull-down menu, select the time period for the report.

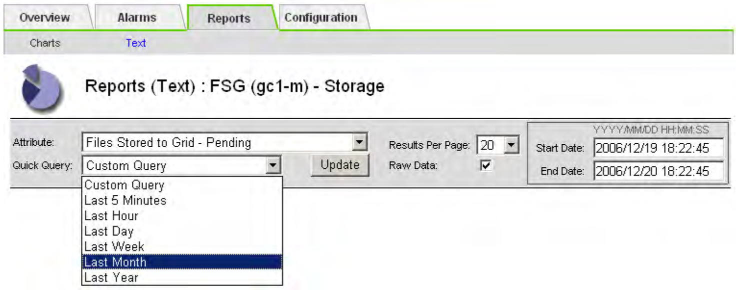

Reports (Text) : FSG (gc1-m) - Storage

Attribute: Files Stored to Grid - Pending

Quick Query: Custom Query

Results Per Page: 20

Raw Data: ☑

YYYY/MM/DD HH:MM:SS

Start Date: 2006/12/19 18:22:45

End Date: 2006/12/20 18:22:45

Figure 37 Report Request Form—Quick Text Reports

The report appears after a few moments. Allow several minutes for tabulation of long time ranges.

7. To display a report for a custom time period:
   a. From the Quick Query pull-down menu, select Custom Query.
   b. Enter the Start Date and End Date.

   Use the format YYYY/MM/DD HH:MM:SS in local time. Leading zeros are required to match the format. For example,
   2007/4/8 7:30:00 fails validation. The correct format is:
   2007/04/08 07:30:00.

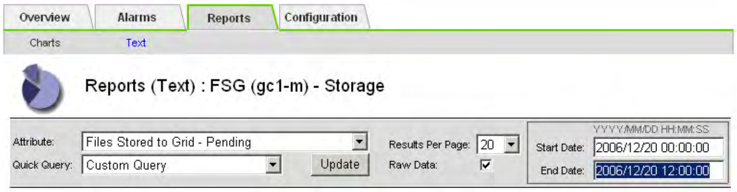

Reports (Text) : FSG (gc1-m) - Storage

Attribute: Files Stored to Grid - Pending

Quick Query: Custom Query

Results Per Page: 20

Raw Data: ☑

YYYY/MM/DD HH:MM:SS

Start Date: 2006/12/20 00:00:00

End Date: 2006/12/20 12:00:00

Figure 38 Report Request Form—Custom Query Report

c Click Update.

A text report is generated. Depending on the length of time set for the query, either a raw text report or downsampled text report is displayed.

Export a Text Report

Exporting text reports opens a new window which allows you to select and copy data. This copied data can then be saved into a new document (for example, a spreadsheet) and used to analyze the performance of the grid.

To export a text report:

1 Create a text report. For more information, see Create a Text Report (page 57).

2 Click Export .

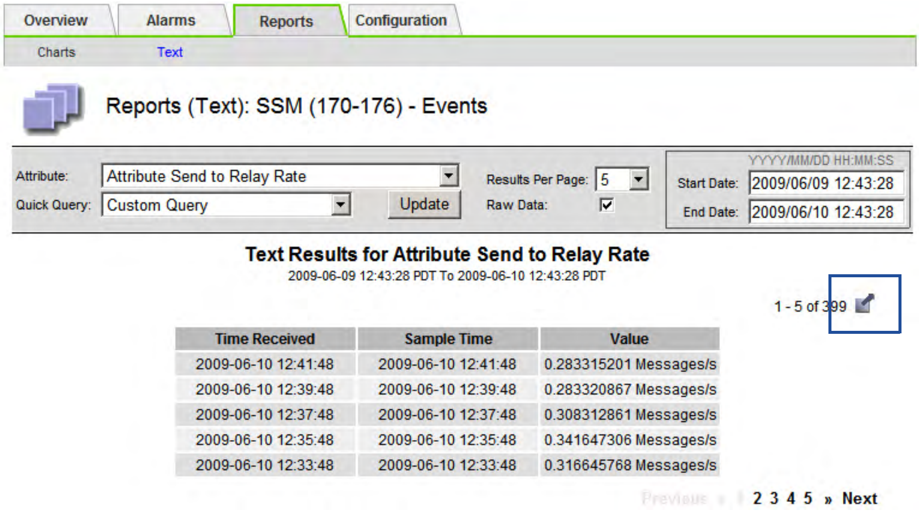

Reports (Text): SSM (170-176) - Events

Attribute: Attribute Send to Relay Rate

Quick Query: Custom Query

Results Per Page: 5

Raw Data: ☑

YYYY/MM/DD HH:MM:SS

Start Date: 2009/06/09 12:43:28

End Date: 2009/06/10 12:43:28

Text Results for Attribute Send to Relay Rate

2009-06-09 12:43:28 PDT To 2009-06-10 12:43:28 PDT

1 - 5 of 399

Time Received	Sample Time	Value
2009-06-10 12:41:48	2009-06-10 12:41:48	0.283315201 Messages/s
2009-06-10 12:39:48	2009-06-10 12:39:48	0.283320867 Messages/s
2009-06-10 12:37:48	2009-06-10 12:37:48	0.308312861 Messages/s
2009-06-10 12:35:48	2009-06-10 12:35:48	0.341647306 Messages/s
2009-06-10 12:33:48	2009-06-10 12:33:48	0.316645768 Messages/s

Figure 39 Export Text Report

The Export Text Report window opens displaying all results for the report.

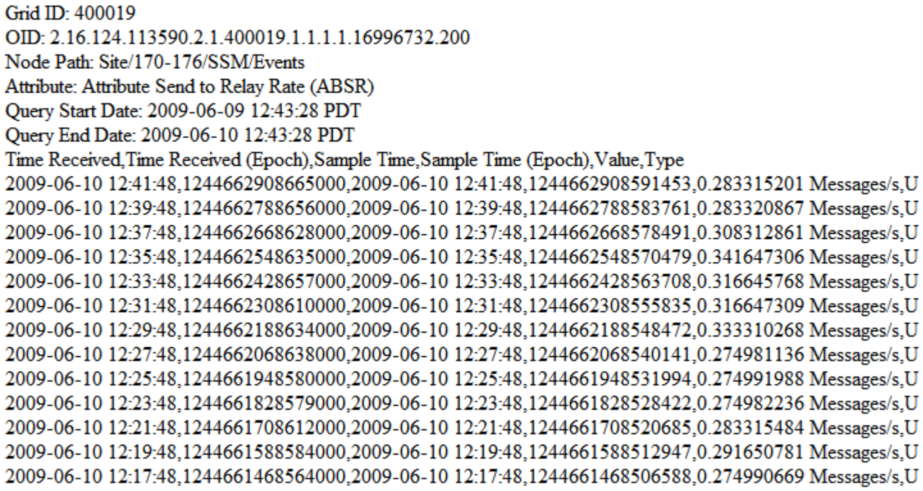

```
Grid ID: 400019
OID: 2.16.124.113590.2.1.400019.1.1.1.1.16996732.200
Node Path: Site/170-176/SSM/Events
Attribute: Attribute Send to Relay Rate (ABSR)
Query Start Date: 2009-06-09 12:43:28 PDT
Query End Date: 2009-06-10 12:43:28 PDT
Time Received,Time Received (Epoch),Sample Time,Sample Time (Epoch),Value,Type
2009-06-10 12:41:48,1244662908665000,2009-06-10 12:41:48,1244662908591453,0.283315201 Messages/s,U
2009-06-10 12:39:48,1244662788656000,2009-06-10 12:39:48,1244662788583761,0.283320867 Messages/s,U
2009-06-10 12:37:48,1244662668628000,2009-06-10 12:37:48,1244662668578491,0.308312861 Messages/s,U
2009-06-10 12:35:48,1244662548635000,2009-06-10 12:35:48,1244662548570479,0.341647306 Messages/s,U
2009-06-10 12:33:48,1244662428657000,2009-06-10 12:33:48,1244662428563708,0.316645768 Messages/s,U
2009-06-10 12:31:48,1244662308610000,2009-06-10 12:31:48,1244662308555835,0.316647309 Messages/s,U
2009-06-10 12:29:48,1244662188634000,2009-06-10 12:29:48,1244662188548472,0.333310268 Messages/s,U
2009-06-10 12:27:48,1244662068638000,2009-06-10 12:27:48,1244662068540141,0.274981136 Messages/s,U
2009-06-10 12:25:48,1244661948580000,2009-06-10 12:25:48,1244661948531994,0.274991988 Messages/s,U
2009-06-10 12:23:48,1244661828579000,2009-06-10 12:23:48,1244661828528422,0.274982236 Messages/s,U
2009-06-10 12:21:48,1244661708612000,2009-06-10 12:21:48,1244661708520685,0.283315484 Messages/s,U
2009-06-10 12:19:48,1244661588584000,2009-06-10 12:19:48,1244661588512947,0.291650781 Messages/s,U
2009-06-10 12:17:48,1244661468564000,2009-06-10 12:17:48,1244661468506588,0.274990669 Messages/s,U
```

Figure 40 Export Text Report Window

3 Select and copy the contents of the Export Text Report window.

This data can now be pasted into a third-party document such as a spreadsheet.

Printing Reports

To print a chart or text report:

1 Create a chart or text report. For more information, see Reports (page 51).

2 Right-click the chart or report to display the Context menu.

3 From the Context menu, select **Print** (for text reports) or Print Picture (for chart reports).

4 The Print dialog box opens.

5 In the Print dialog box, select printing options and then click **Print**.

The report is printed.

3 Data Flow

This chapter describes the grid activities that take place as objects are ingested, replicated, retrieved, modified, and purged.

Key Concepts

To follow objects through the grid as they are processed, you need to understand these concepts:

- Client shares and FSG managed file system
- FSG replication groups
- Topology queries
- HTTP protocol commands
- ILM policy
- Owner CMS and metadata replication
- Object content handle

Client Shares and FSG Managed File System

In order to communicate with the FSG service, client applications map a network drive to the FSG file share, for example, `/fsg/myDirectory`. The FSG service supports CIFS and NFS file share protocols.

Applications interface with the grid via the FSG's "managed file system". There is one managed file system per FSG and it is mounted at /fsg. The managed file system contains file pointers that point to the file's location in the grid. The FSG managed file system shown in Figure 41 contains a file pointer for the object /fsg/myDirectory/image.jpg. The FSG uses the object's unique identifier to communicate with the CMS service (via the LDR service) which tracks where the object is stored on the Storage Nodes (LDR service) and on the Tape Nodes (ARC service).

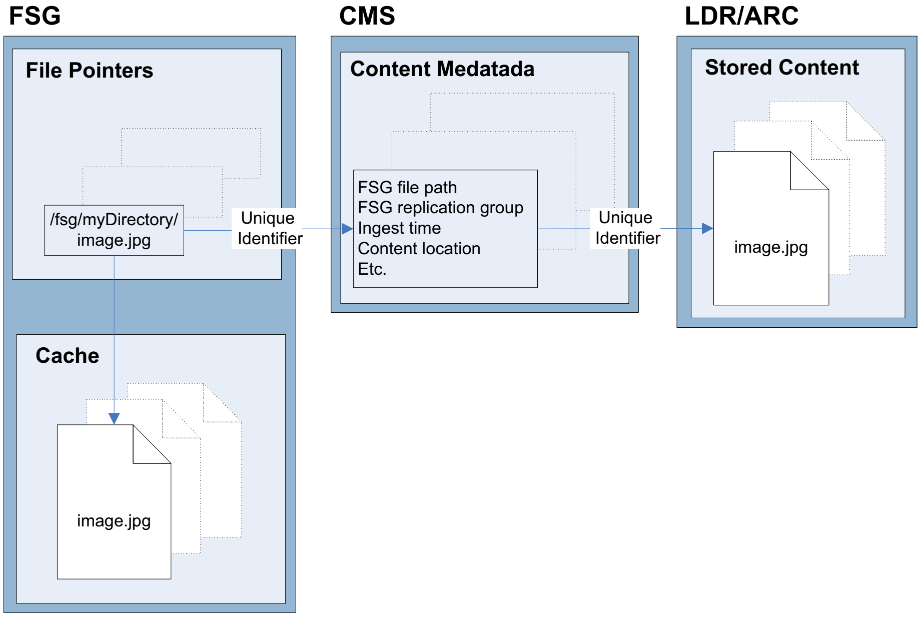

Figure 41 FSG File Pointers

The FSG cache contains recently stored and retrieved files. The purpose of the cache is to speed up data access for clients. Unaccessed files in the FSG cache are swapped out to make room for new files, leaving a file pointer that allows the file data to be re-cached if the client later retrieves the file. Files are swapped out of the FSG cache in approximately least-recently accessed order.

FSG Replication Groups

All FSGs belong to replication groups. A replication group contains a primary FSG and one or more secondary FSGs. The primary FSG provides read and write access to clients. The secondary FSG provides a mirror of the primary FSG's managed file system for redundancy in case the primary FSG fails or must be taken out of service.

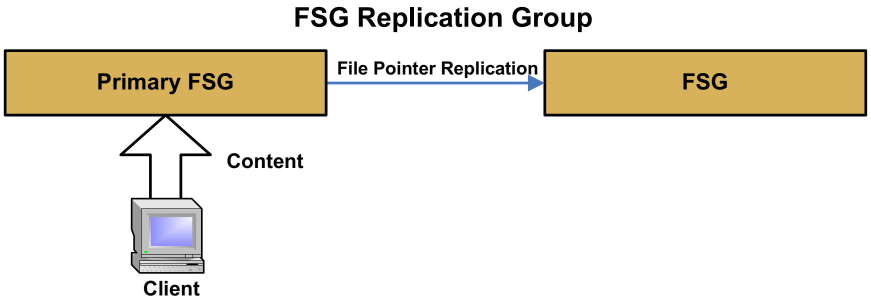

Figure 42 FSG Replication Group

The primary FSG replicates its file pointers (for instance Figure 43 shows the pointer for `/fsg/myDirectory/image.jpg`) to the secondary FSG.

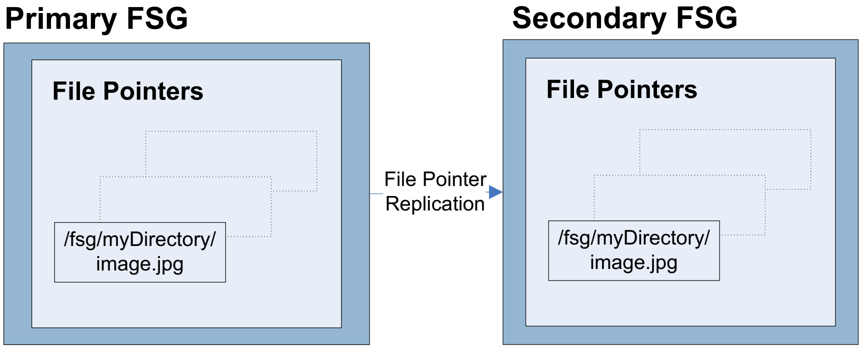

Figure 43 FSG File Pointer Replication

Depending on the configuration, the secondary FSG may also provide read-only client access and may perform backups of the file system into the grid for additional redundancy.

There are two types of FSG replication groups:

- Basic Gateway replication group
- High Availability Gateway replication group

For a detailed review of each type of FSG replication group, see FSG—File System Gateway (page 224).

Topology Queries

When a grid service needs information from another service or needs an action to be performed by another service, it contacts the ADC service to find the best service to process the request. This is called a topology query. The ADC service responds to each query with the latest information received from the grid. The information maintained by the ADC service includes CPU load, amount of available disk space, supported services, and location.

HTTP Protocol Commands

Internally, FSGs use standard HTTP protocol commands to communicate with LDRs to store, retrieve, and purge objects. For example, the FSG issues a PUT command to the LDR service to store an object into the grid, a GET command to retrieve the object, and a DELETE command to purge the object.

ILM Policy

The ILM policy defines how and where objects are stored in the grid. At a high level, ILM policies dictate:

- geography—the location of the files
- storage grade—what type of storage to use
- replication—the number of copies to make

Location, storage grade, and number of copies can vary over time. In the ILM example shown in Figure 44, a file is ingested into the grid via an FSG. At ingest, two copies of the object are stored in the Data Center site on Fibre Channel disks and one copy is stored in the Disaster Recovery site on SATA disks. One year after ingest, one copy at the DC site is deleted and one copy is created on archive media at the DR site.

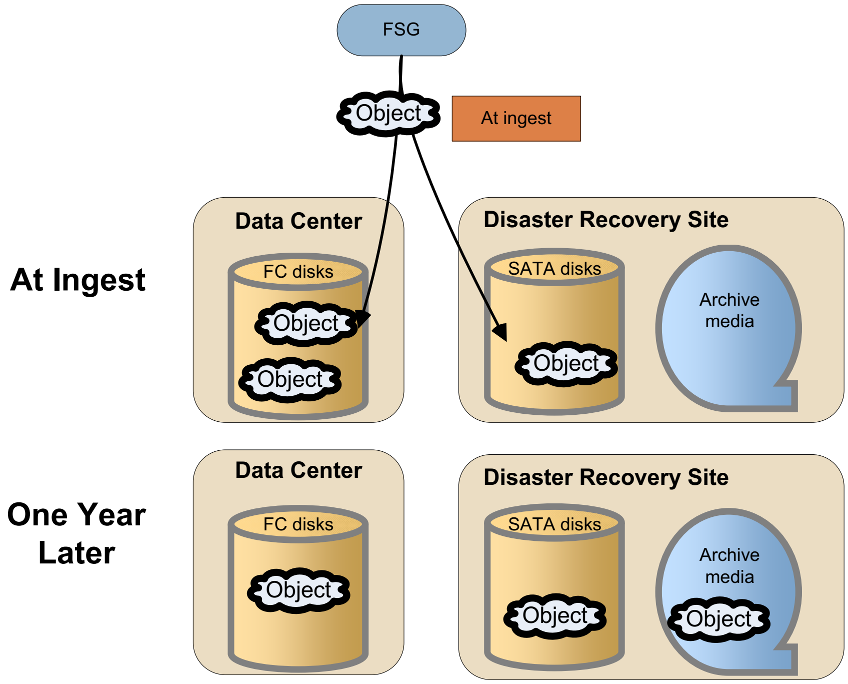

Figure 44 Information Lifecycle Management (ILM) Example

Owner CMS and Reproduction of Metadata

Each object ingested into the grid has a set of associated metadata. Metadata is information related to the object, for instance the ingest file path or the file ingest time. Content metadata is managed by the CMS service. The first CMS service to get the object metadata becomes the owner CMS. The owner CMS then copies the metadata to other CMS services.

A grid has either CMSs that use metadata replication or CMSs that use metadata synchronization. If the CMSs use metadata replication, the metadata is copied to all CMSs in the same CMS replication group as the owner. Metadata also follows content: for each copy of the data that is made, one copy of metadata is stored on a CMS in the same location as the data. If the CMSs use metadata synchronization, the owner CMS synchronizes metadata with all other read-write CMSs in the grid. After a grid that uses synchronized metadata has been expanded to add metadata storage capacity, the metadata is synchronized to all CMSs of the same generation as the owner CMS.

For a detailed review of the two kinds of CMS and their operation, see CMS—Content Management System (page 193).

Object Content Handle

The grid assigns a unique identifier to each object ingested into the grid. This identifier is called a “content handle”. The grid uses the content handle to refer to the object. As long as the object is referenced by the client application, it is said to have a content handle. When the client application deletes the object, the object’s content handle is said to be released.

Object Lifecycle

Figure 45 follows an object as it is ingested, retrieved, becomes inactive, and is finally deleted:

- The client application creates the file over CIFS/NFS. The file is ingested into the grid in the background from the FSG to an LDR over HTTP and is replicated according to the ILM policy. The file is also stored in the FSG cache.
- When the client reads the file, the file is retrieved either from the FSG cache over CIFS/NFS or from the LDR to the FSG and out to the client over CIFS/NFS.
- An inactive file may be swapped out of the FSG cache to make room for more recent or more active files
- When the client deletes the file, removal from the FSG triggers a removal notification to the LDR.

NOTE When an object is ingested, replicated, retrieved, modified, or purged, it may take several minutes for the NMS MI to display updated attribute values.

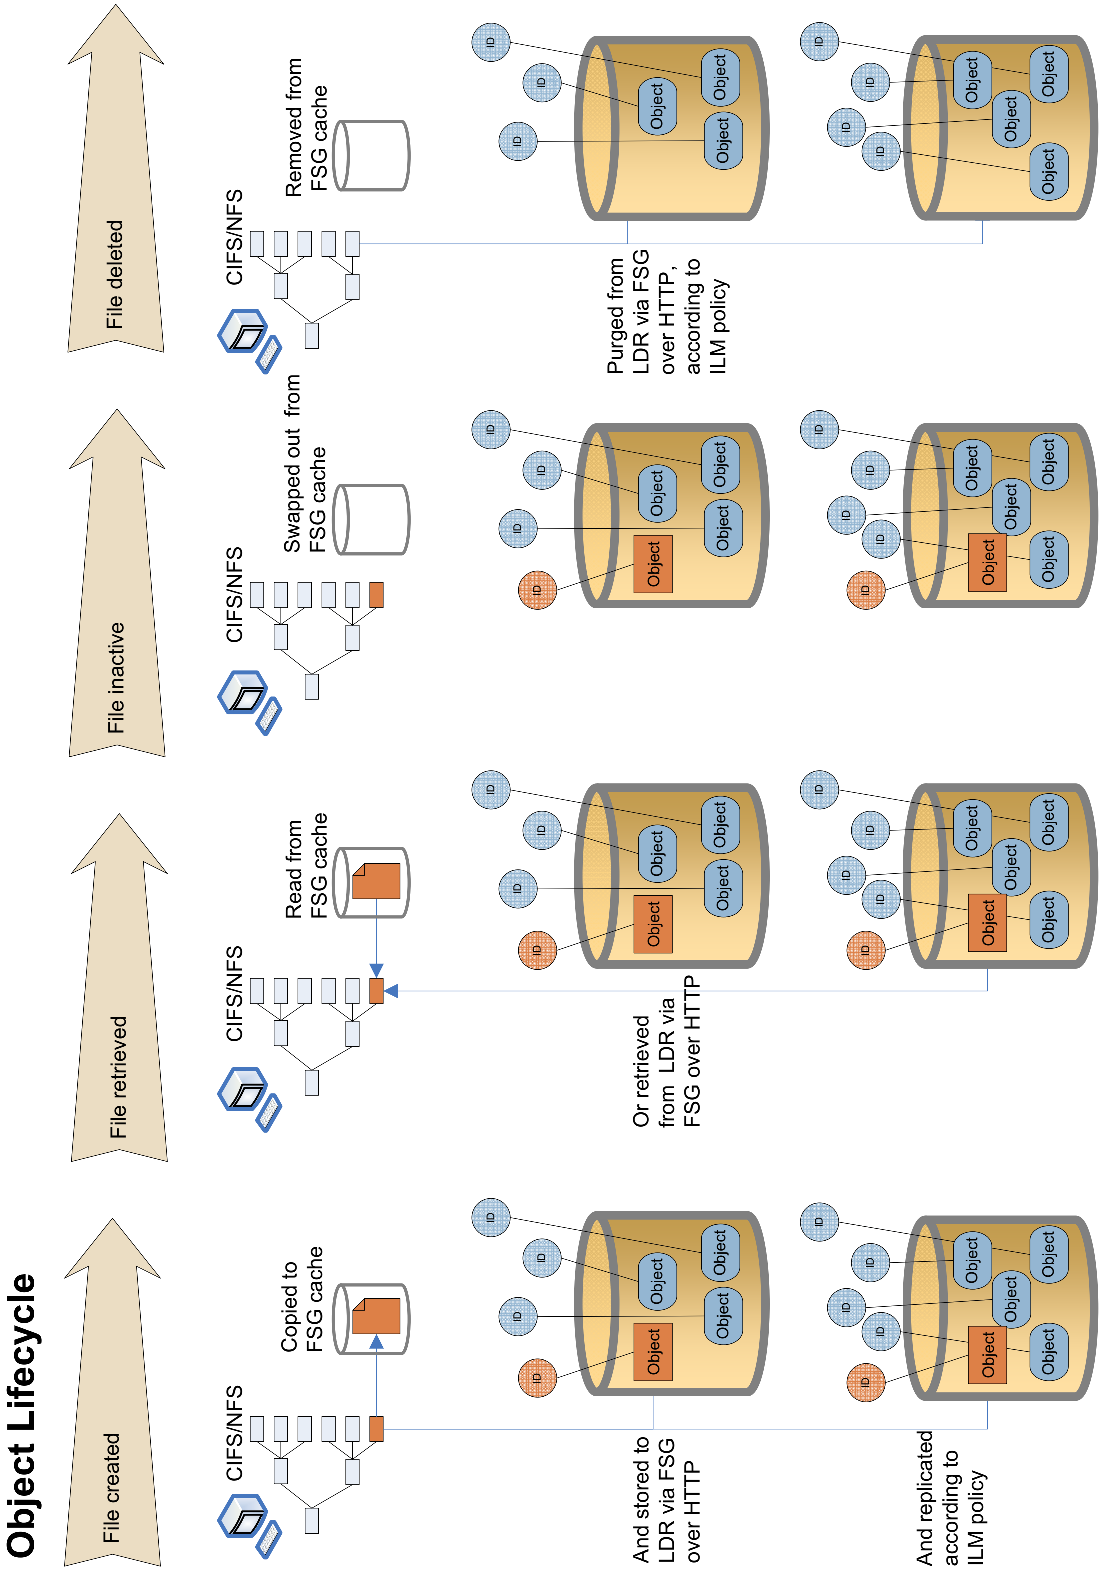
Object Lifecycle
File created
CIFS/NFS
Copied to FSG cache
And stored to LDR via FSG over HTTP
And replicated according to ILM policy
File retrieved
CIFS/NFS
Read from FSG cache
Or retrieved from LDR via FSG over HTTP
File inactive
CIFS/NFS
Swapped out from FSG cache
File deleted
CIFS/NFS
Removed from FSG cache
Purged from LDR via FSG over HTTP, according to ILM policy
ID
Object

Figure 45 Object Lifecycle

Ingest

Object ingest refers to the process of a client application storing content into the grid.

Data Flow

See Figure 46 for a simplified step-by-step description of what happens when objects are ingested into the grid.

When a client saves a file to the FSG file system (for example to the mapped network drive `/fsg/myDirectory`), the FSG stores a local copy in its cache and streams the file to permanent storage via an LDR. The LDR assigns a unique identifier (a content handle) to the file and transmits this information to the FSG. The LDR also transmits the object metadata to the CMS. The primary FSG replicates the content handle and file pointers to the other FSGs in its replication group, and the CMS replicates the metadata to other CMSs in the grid.

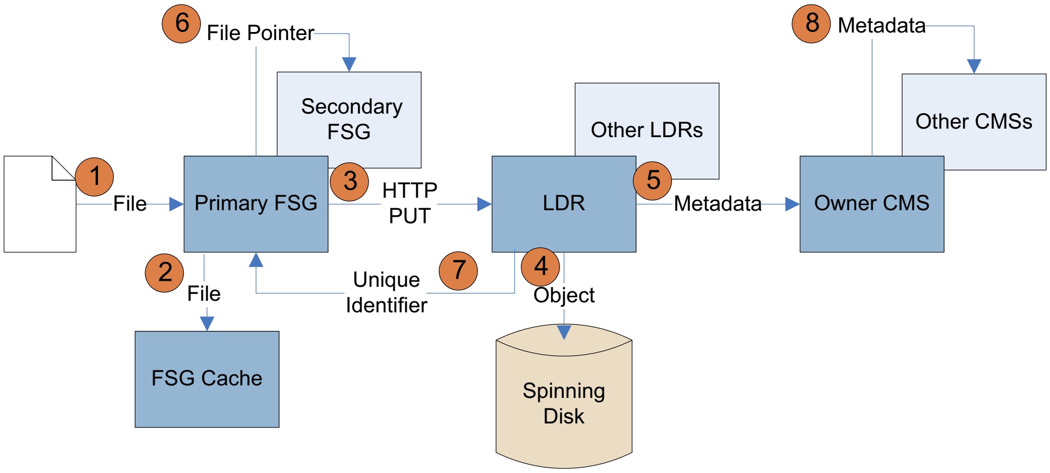

Figure 46 Ingest Data Flow

1 The client application saves a file to the primary FSG managed file system via its mapped network share, for example `/fsg/myDirectory`. This triggers an FSG “create operation” and as a result the FSG creates a file pointer for this object.

2 The primary FSG saves a copy of the object in its cache.

3 The primary FSG stores the file to an LDR via an HTTP PUT command. What LDR is chosen depends on the result of an ADC topology query.

4 The LDR saves the object to spinning disk and allocates a “content handle”, that is, a unique identifier (UUID), to the object.

5	The LDR notifies a CMS that a new piece of content has been ingested and sends the object metadata, which includes the unique identifier, to the CMS. This CMS becomes the owner CMS. What CMS is chosen depends on the result of a topology query.
6	The primary FSG replicates the file pointer to the other FSGs in its replication group.
7	The LDR sends the unique identifier (UUID) to the primary FSG. This UUID is used by the FSG to uniquely identify the object. Thus, the FSG uses the UUID to retrieve, query, and delete the correct object from the grid.
8	The owner CMS replicates the metadata and content location to the other CMSs. If the grid uses metadata replication, metadata is replicated to a subset of the CMSs; otherwise, content metadata is synchronized across all read-write CMSs.

Related Attributes

The actual shape of the trends varies with each grid and depends on ingest, replication, and purging rates.

Table 11 lists some of the NMS MI attributes used to track what happens when a single object is ingested into the grid.

Table 11 Object Ingest Attributes

Component	Attribute Changes
Primary FSG > Storage 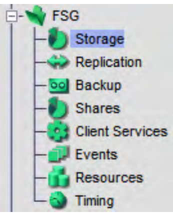 FSG Storage Replication Backup Shares Client Services Events Resources Timing 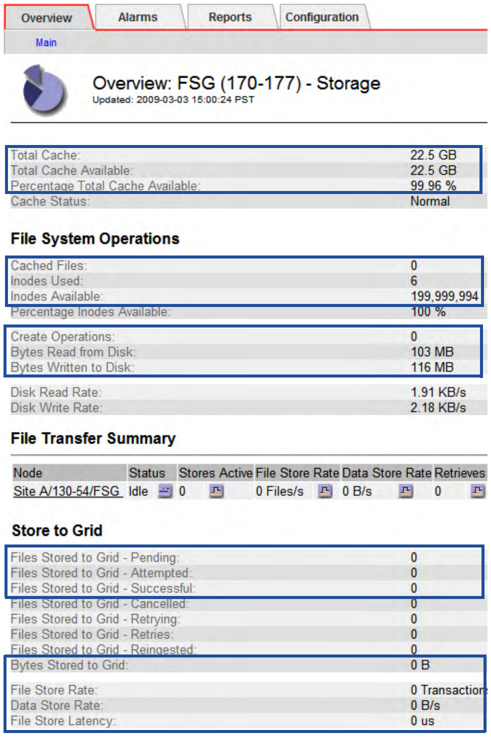 Overview: FSG (170-177) - Storage Updated: 2009-03-03 15:00:24 PST Total Cache: 22.5 GB Total Cache Available: 22.5 GB Percentage Total Cache Available: 99.96 % Cache Status: Normal **File System Operations** Cached Files: 0 Inodes Used: 6 Inodes Available: 199,999,994 Percentage Inodes Available: 100 % Create Operations: 0 Bytes Read from Disk: 103 MB Bytes Written to Disk: 116 MB Disk Read Rate: 1.91 KB/s Disk Write Rate: 2.18 KB/s **File Transfer Summary** Node | Status | Stores Active | File Store Rate | Data Store Rate | Retrieves Site A/130-54/FSG | Idle | 0 | 0 Files/s | 0 B/s | 0 **Store to Grid** Files Stored to Grid - Pending: 0 Files Stored to Grid - Attempted: 0 Files Stored to Grid - Successful: 0 Files Stored to Grid - Cancelled: 0 Files Stored to Grid - Retrying: 0 Files Stored to Grid - Retries: 0 Files Stored to Grid - Reingested: 0 Bytes Stored to Grid: 0 B File Store Rate: 0 Transaction Data Store Rate: 0 B/s File Store Latency: 0 us	Create Operations (FCRO): The number of new files or folders created in the file system increases by 1. 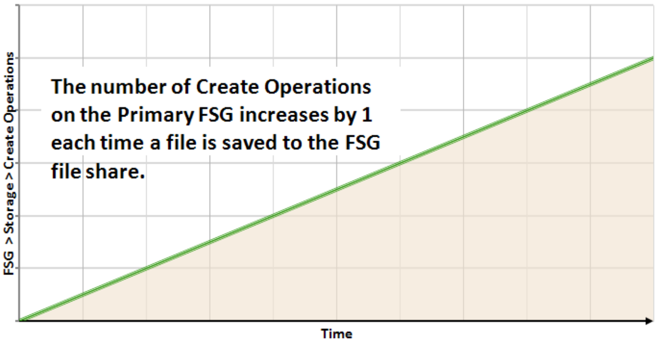 The number of Create Operations on the Primary FSG increases by 1 each time a file is saved to the FSG file share. Inodes Used (FSIU): The number of inodes (files and directories) used on the FSG file system increases by 1. Inodes Used and Create Operations can differ if objects are deleted and their inodes are later re-used for new objects. Inodes Available (FSIA): The number of inodes (files and directories) available on the FSG file system decreases by 1. Files Stored to Grid - Successful (FSGC): The number of files stored persistently on the grid increases by 1. The value increases after the FSG receives an acknowledgment from the LDR that the file has been successfully stored into the grid (which is at the same time as it receives the unique identifier from the LDR). 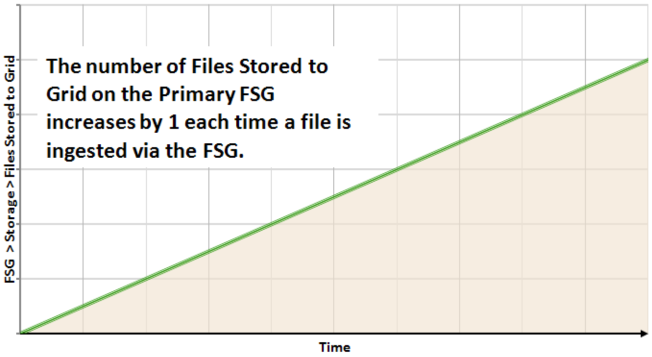 The number of Files Stored to Grid on the Primary FSG increases by 1 each time a file is ingested via the FSG. Files Stored to Grid - Attempted (FSGA): The number of initiated file transfers to the grid increases by 1.

Table 11 Object Ingest Attributes *(continued)*

Component	Attribute Changes
	Files Stored to Grid - Pending (FSGP): The number of new files cached locally and waiting for transfer to the grid for persistent storage increases by 1 temporarily, and decreases again after the file has been stored. Bytes Stored to Grid (FSGB): The number of bytes ingested successfully into the grid increases by an amount equivalent to the file size. Bytes Read from Disk (FSRB): The number of bytes read from disk increases. File ingest generates multiple read operations as the FSG reads the file and stores it to the grid. Bytes Written to Disk (FSWB): The number of bytes written to disk increases. File ingest generates multiple write operations to disk for the creation of the file and the saving of the file content and metadata. Total Cache Available (FSTA): The total local cache space on the primary FSG still available for use decreases unless the minimum value has been reached. After that, the primary FSG swaps files out of the cache to make room for the file that has been ingested. The cache may dip below the minimum value temporarily if files are being created faster than existing files can be swapped out. 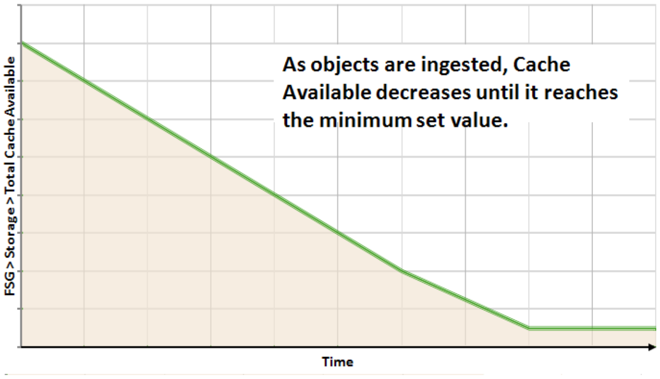 As objects are ingested, Cache Available decreases until it reaches the minimum set value.

Table 11 Object Ingest Attributes *(continued)*

<table>
<tr><th>Component</th><th>Attribute Changes</th></tr>
<tr><td></td><td>
Cached Files (FSCF): The number of cached files on the primary FSG increases by 1 unless the cache is full. The number of cached files will generally stabilize after the Total Cache Available reaches its minimum, but may fluctuate as files are added and removed from the cache depending on file size.

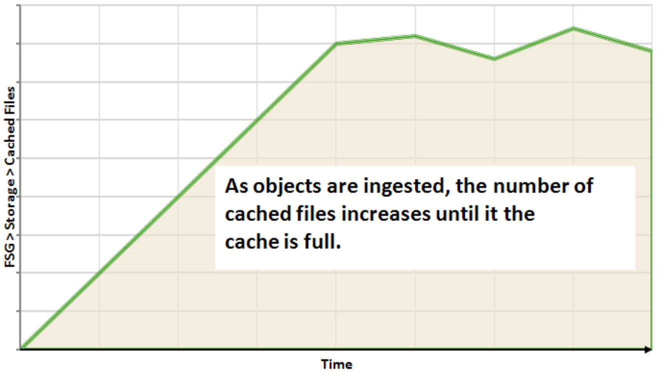

</td></tr>
<tr><td>
FSG >Shares

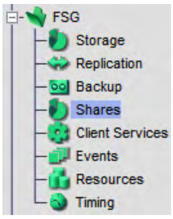

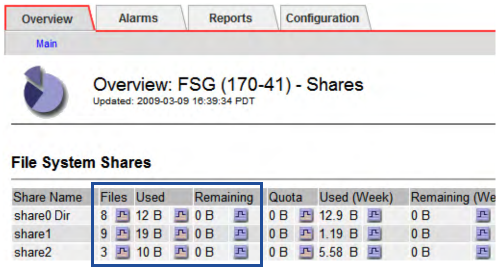

Overview: FSG (170-41) - Shares
Updated: 2009-03-09 16:39:34 PDT

File System Shares

<table>
<tr><th>Share Name</th><th>Files</th><th>Used</th><th>Remaining</th><th>Quota</th><th>Used (Week)</th><th>Remaining (We</th></tr>
<tr><td>share0 Dir</td><td>8</td><td>12 B</td><td>0 B</td><td>0 B</td><td>12.9 B</td><td>0 B</td></tr>
<tr><td>share1</td><td>9</td><td>19 B</td><td>0 B</td><td>0 B</td><td>1.19 B</td><td>0 B</td></tr>
<tr><td>share2</td><td>3</td><td>10 B</td><td>0 B</td><td>0 B</td><td>5.58 B</td><td>0 B</td></tr>
</table>
</td><td>
Files (FSSF): The number of files ingested on this share will increase by 1 the next time a backup takes place.

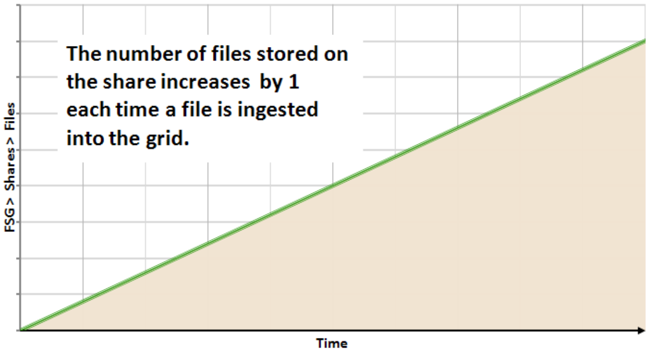

Used (FSSB) and Used (Week) (FSSR): The total size of all files referenced by this share will increase by an amount equivalent to the file size the next time a backup takes place.

Remaining (FSSA) and Remaining (Week) (FSSL): If a quota has been configured for the share, the total amount of space remaining on the share will decrease by the size of the file.
</td></tr>
</table>

Table 11 Object Ingest Attributes *(continued)*

Component	Attribute Changes
LDR > HTTP 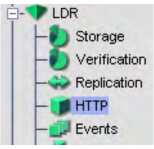 LDR Storage Verification Replication HTTP Events 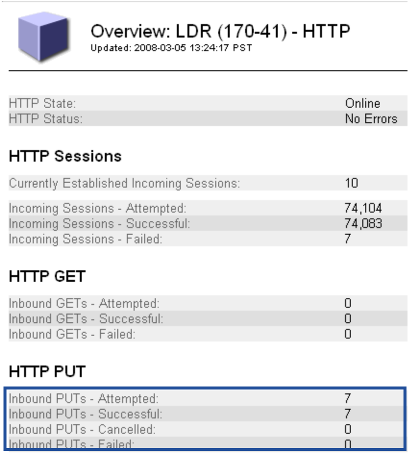 Overview: LDR (170-41) - HTTP Updated: 2008-03-05 13:24:17 PST HTTP State: Online HTTP Status: No Errors **HTTP Sessions** Currently Established Incoming Sessions: 10 Incoming Sessions - Attempted: 74,104 Incoming Sessions - Successful: 74,083 Incoming Sessions - Failed: 7 **HTTP GET** Inbound GETs - Attempted: 0 Inbound GETs - Successful: 0 Inbound GETs - Failed: 0 **HTTP PUT** Inbound PUTs - Attempted: 7 Inbound PUTs - Successful: 7 Inbound PUTs - Cancelled: 0 Inbound PUTs - Failed: 0	Inbound PUTs - Successful (HIPC): The total number of HTTP PUT (“content store”) requests that have been completed successfully by the LDR increases by 1. 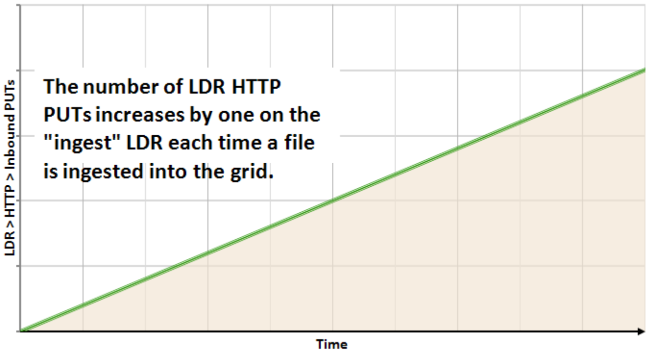 The number of LDR HTTP PUTs increases by one on the "ingest" LDR each time a file is ingested into the grid. Inbound PUTs - Attempted (HAIP): The total number of HTTP PUT (content store) requests that have been received by the LDR also increases by 1.

Table 11 Object Ingest Attributes *(continued)*

Component	Attribute Changes
LDR > Storage 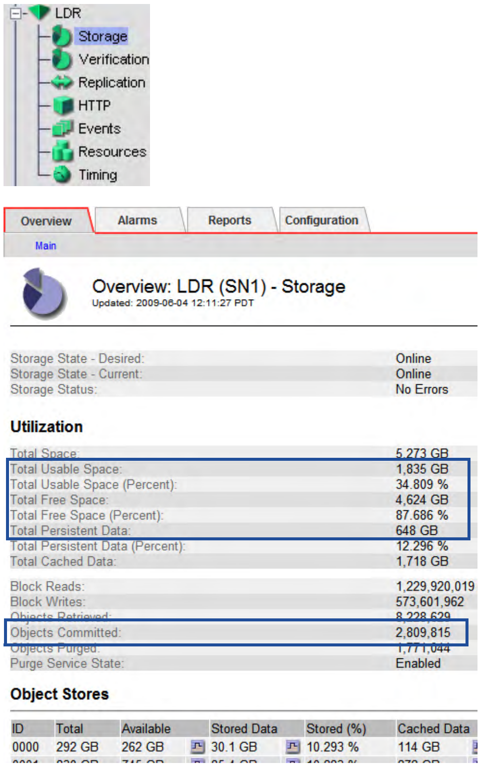 LDR Storage Verification Replication HTTP Events Resources Timing Overview Alarms Reports Configuration Main Overview: LDR (SN1) - Storage Updated: 2009-06-04 12:11:27 PDT Storage State - Desired: Online Storage State - Current: Online Storage Status: No Errors **Utilization** Total Space: 5,273 GB Total Usable Space: 1,835 GB Total Usable Space (Percent): 34.809 % Total Free Space: 4,624 GB Total Free Space (Percent): 87.686 % Total Persistent Data: 648 GB Total Persistent Data (Percent): 12.296 % Total Cached Data: 1,718 GB Block Reads: 1,229,920,019 Block Writes: 573,601,962 Objects Retrieved: 8,228,629 Objects Committed: 2,809,815 Objects Purged: 1,771,044 Purge Service State: Enabled **Object Stores** ID · Total · Available · Stored Data · Stored (%) · Cached Data 0000 · 292 GB · 262 GB · 30.1 GB · 10.293 % · 114 GB	Total Usable Space (STAS): The total amount of object storage space that is currently available to be used to store objects decreases by an amount roughly equivalent to the file size. 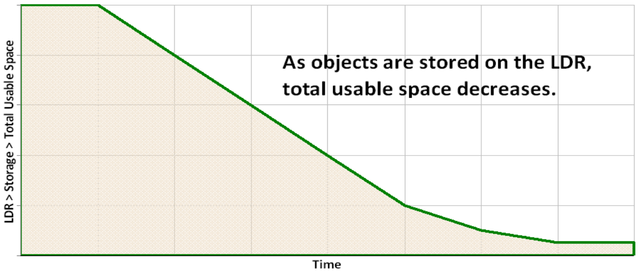 Total Usable Space (Percent) (SAVP): The total amount of object storage space (displayed as a percentage) that is currently available to be used to store objects decreases by an amount roughly equivalent to the file size. Total Free Space (SUSA): The total amount of all free space—available to be used or not—on all object stores decreases by an amount roughly equivalent to the file size. Total Free Space (Percent) (SUSP): The total amount (displayed as a percentage) of all free space—available to be used or not—on all object stores decreases by an amount roughly equivalent to the file size. Total Persistent Data (SPSD): The estimate of the size of the persistently stored data increases by an amount roughly equivalent to the file size. Objects Committed (OCOM): The number of object store operations that have been processed by the LDR increases by 1. 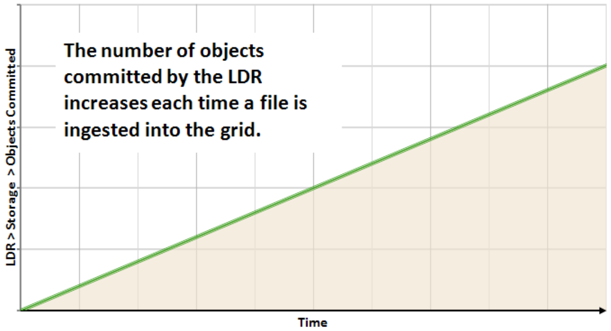

Table 11 Object Ingest Attributes *(continued)*

Component	Attribute Changes
CMS > Metadata 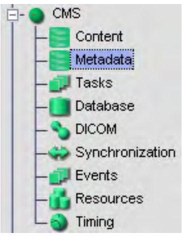 CMS Content Metadata Tasks Database DICOM Synchronization Events Resources Timing 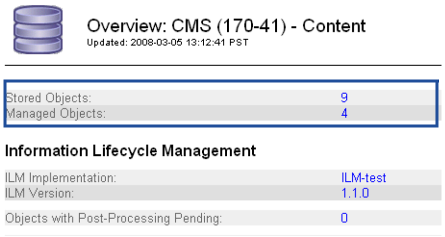 Overview: CMS (170-41) - Content Updated: 2008-03-05 13:12:41 PST Stored Objects: 9 Managed Objects: 4 Information Lifecycle Management ILM Implementation: ILM-test ILM Version: 1.1.0 Objects with Post-Processing Pending: 0 For grids that use synchronized metadata, the attributes Stored Objects and Managed Objects are shown under the Content component. For grids that use metadata replication, these two attributes appear under the Metadata component.	Stored Objects (COoT): The number of objects in the CMS metadata database increases by 1 on all CMSs in the same CMS replication group in a grid that uses metadata replication, and on all read-write CMSs in a grid that uses metadata synchronization. 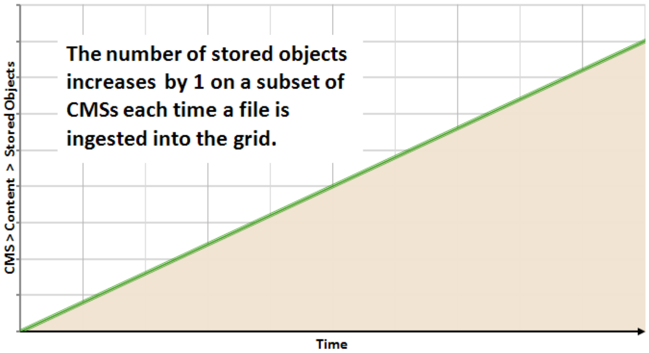 The number of stored objects increases by 1 on a subset of CMSs each time a file is ingested into the grid. Managed Objects (COoM): The number of objects owned by the owner CMS increases by 1. 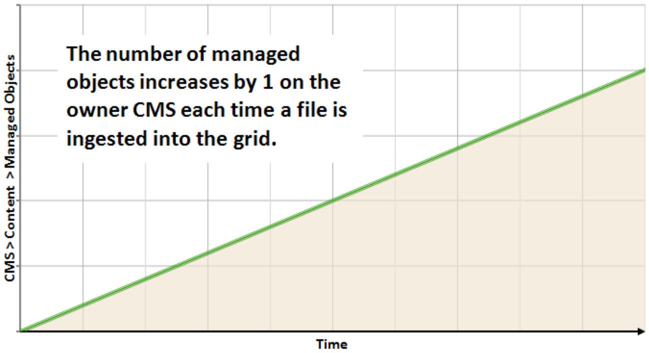 The number of managed objects increases by 1 on the owner CMS each time a file is ingested into the grid.

Table 11 Object Ingest Attributes *(continued)*

Component	Attribute Changes
CMS > Database 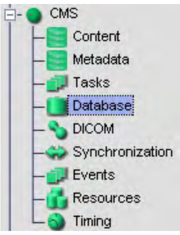 CMS Content Metadata Tasks Database DICOM Synchronization Events Resources Timing 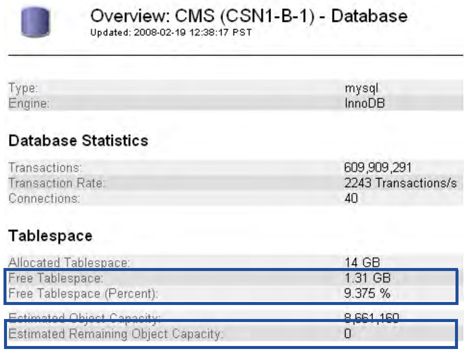 Overview: CMS (CSN1-B-1) - Database Updated: 2008-02-19 12:38:17 PST Type: mysql Engine: InnoDB Database Statistics Transactions: 609,909,291 Transaction Rate: 2243 Transactions/s Connections: 40 Tablespace Allocated Tablespace: 14 GB Free Tablespace: 1.31 GB Free Tablespace (Percent): 9.375 % Estimated Object Capacity: 8,661,160 Estimated Remaining Object Capacity: 0	Free Tablespace Percent (DBSP): The amount of space remaining in the metadata database decreases. MySQL manages free tablespace in chunks. Ingesting (or deleting) a single object may not change the reported free table space. Ingesting (or deleting) many objects will eventually change the free table space. 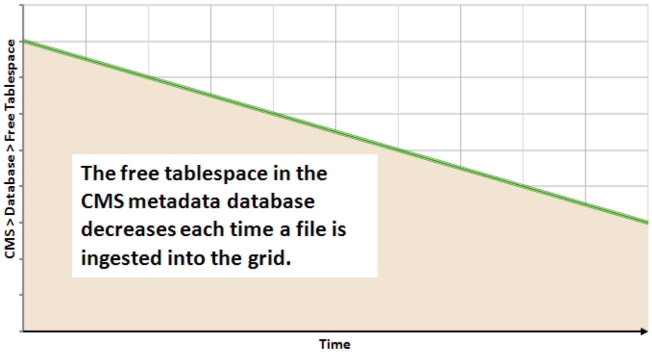 The free tablespace in the CMS metadata database decreases each time a file is ingested into the grid. Estimated Remaining Object Capacity (CORS): The estimate of how many more objects can be tracked in the CMS metadata database decreases by 1. 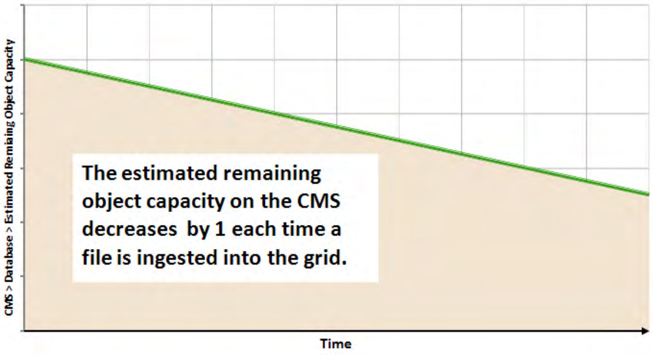 The estimated remaining object capacity on the CMS decreases by 1 each time a file is ingested into the grid.

Content Replication

Following ingest, the object is replicated according to the grid's ILM policy. Content replication refers to the process of making copies of the object in order to satisfy the ILM policy.

Data Flow

The owner CMS, which is the first CMS to receive the object metadata from the LDR, controls the replication, that is, ensures that the correct number of copies are stored in the correct locations for the duration specified by the ILM policy. See Figure 47 for a simplified step-by-step description of what happens as objects are replicated.

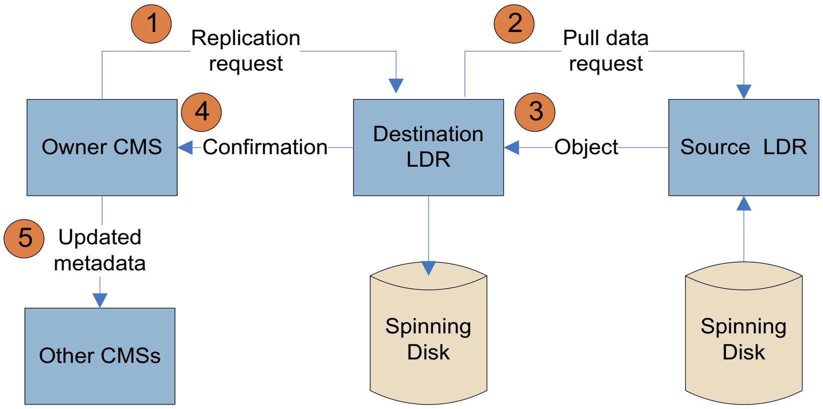

Figure 47 Replication Data Flow

1	The owner CMS queries the ADC to determine the best destination LDR within the storage pool defined by the ILM policy, and sends that LDR a command to initiate replication.
2	The destination LDR queries the ADC for the best source location and sends a replication request to the source LDR.
3	The source LDR sends a copy of the object to the destination LDR.
4	The destination LDR notifies the CMS that the object has been stored.
5	The owner CMS updates the location information and distributes that information to the other CMSs that store metadata for this object.

Related Attributes

Table 12 lists some of the NMS MI attributes used to track what happens when a single object is replicated.

Table 12 Object Replication Attributes

Component	Attribute Changes
CMS > Content 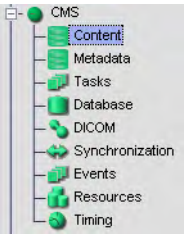 CMS Content Metadata Tasks Database DICOM Synchronization Events Resources Timing 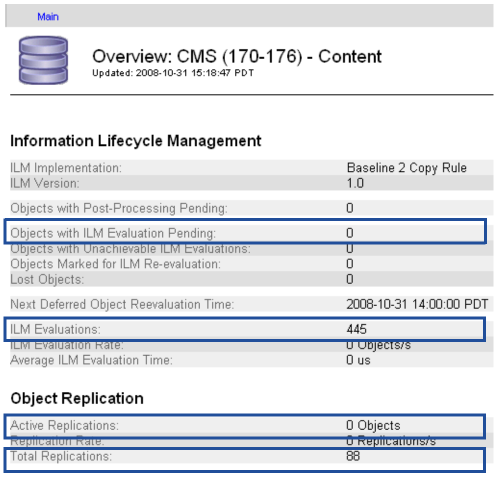 Overview: CMS (170-176) - Content Updated: 2008-10-31 15:18:47 PDT **Information Lifecycle Management** ILM Implementation: Baseline 2 Copy Rule ILM Version: 1.0 Objects with Post-Processing Pending: 0 Objects with ILM Evaluation Pending: 0 Objects with Unachievable ILM Evaluations: 0 Objects Marked for ILM Re-evaluation: 0 Lost Objects: 0 Next Deferred Object Reevaluation Time: 2008-10-31 14:00:00 PDT ILM Evaluations: 445 ILM Evaluation Rate: 0 Objects/s Average ILM Evaluation Time: 0 us **Object Replication** Active Replications: 0 Objects Replication Rate: 0 Replications/s Total Replications: 88	Objects with ILM Evaluation Pending (ORpe): The number of objects waiting to be processed through the business rules for replication increases by 1 when the object is ingested into the grid and decreases by 1 after replication is complete. Active Replications (DCdA): The number of objects in the process of being replicated increases by 1 when replication starts and decreases by 1 when replication is complete. 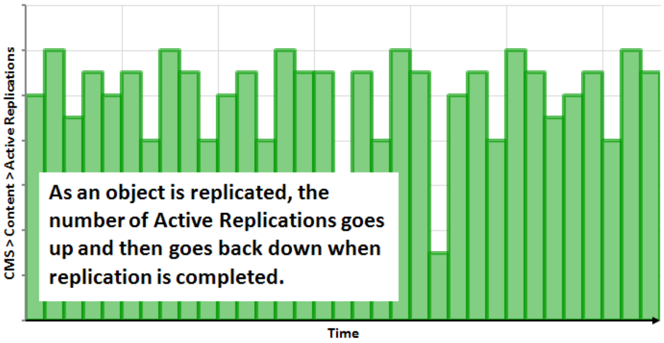 As an object is replicated, the number of Active Replications goes up and then goes back down when replication is completed. ILM Evaluations (ILev): The total number of ILM evaluations that have been performed to date increases when the object is evaluated after ingest and again when the object has been replicated. The number of evaluations is dependent on the ILM rules for the grid. Total Replications (DCdT): The total number of object replications performed by the owner CMS since grid startup increases by 1 each time a copy is made. 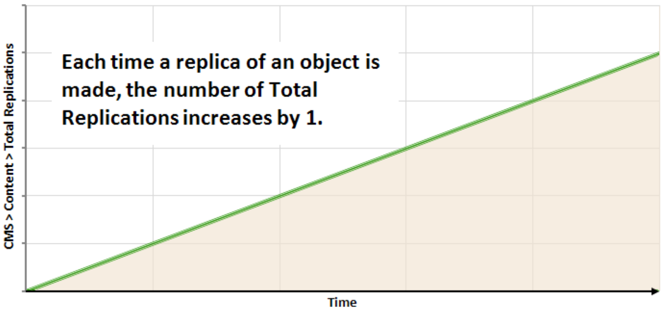 Each time a replica of an object is made, the number of Total Replications increases by 1.

Table 12 Object Replication Attributes *(continued)*

<table>
<tr><th>Component</th><th>Attribute Changes</th></tr>
<tr><td>

LDR > Storage

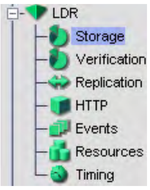

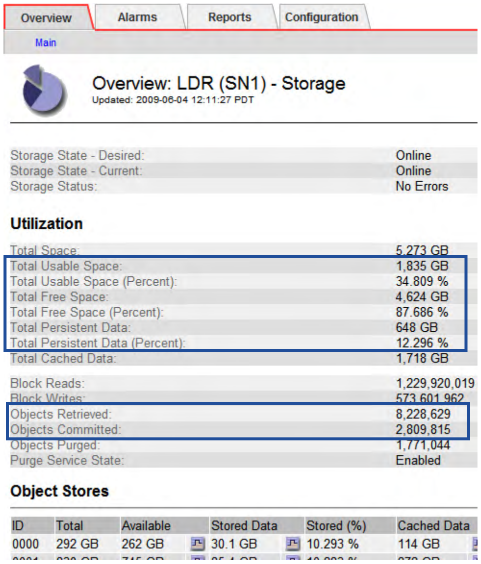

Overview: LDR (SN1) - Storage
Updated: 2009-06-04 12:11:27 PDT

Storage State - Desired: Online
Storage State - Current: Online
Storage Status: No Errors

Utilization

Total Space: 5,273 GB
Total Usable Space: 1,835 GB
Total Usable Space (Percent): 34.809 %
Total Free Space: 4,624 GB
Total Free Space (Percent): 87.686 %
Total Persistent Data: 648 GB
Total Persistent Data (Percent): 12.296 %
Total Cached Data: 1,718 GB

Block Reads: 1,229,920,019
Block Writes: 573,601,962
Objects Retrieved: 8,228,629
Objects Committed: 2,809,815
Objects Purged: 1,771,044
Purge Service State: Enabled

Object Stores

<table>
<tr><th>ID</th><th>Total</th><th>Available</th><th>Stored Data</th><th>Stored (%)</th><th>Cached Data</th></tr>
<tr><td>0000</td><td>292 GB</td><td>262 GB</td><td>30.1 GB</td><td>10.293 %</td><td>114 GB</td></tr>
</table>

</td><td>

Total Usable Space (STAS): The total amount of object storage space that is currently available to be used to store objects decreases by an amount roughly equivalent to the file size.

Total Usable Space (Percent) (SAVP): The total amount of object storage space (displayed as a percentage) that is currently available to be used to store objects decreases by an amount roughly equivalent to the file size.

Total Free Space (SUSA): The total amount of all free space—available to be used or not—on all object stores decreases by an amount roughly equivalent to the file size.

Total Free Space (Percent) (SUSP): The total amount (displayed as a percentage) of all free space—available to be used or not—on all object stores decreases by an amount roughly equivalent to the file size.

Total Persistent Data (SPSD): The estimate of the size of the persistently stored data increases by an amount roughly equivalent to the file size.

Total Persistent Data (Percent) (SPDP): The percentage of the total storage space used by persistent data on each destination LDR increases by an amount roughly equivalent to the size of the replicated object.

Objects Retrieved (ORET): The number of persistent objects retrieved from the source LDR increases by 1.

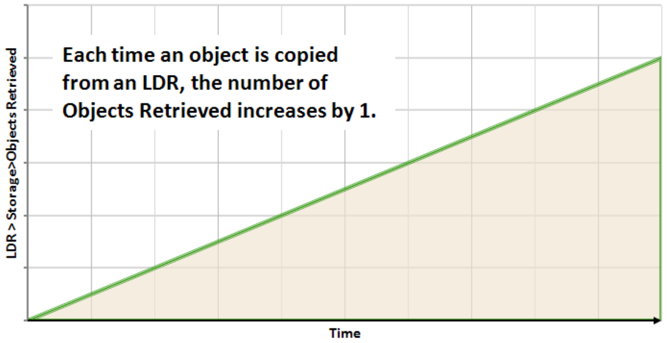

Each time an object is copied from an LDR, the number of Objects Retrieved increases by 1.

Objects Committed (OCOM): The number of persistent objects stored on each destination LDR increases by 1.

</td></tr>
</table>

Table 12 Object Replication Attributes *(continued)*

Component	Attribute Changes
LDR > Replication 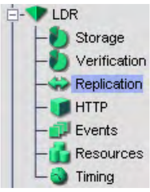 LDR Storage Verification Replication HTTP Events Resources Timing 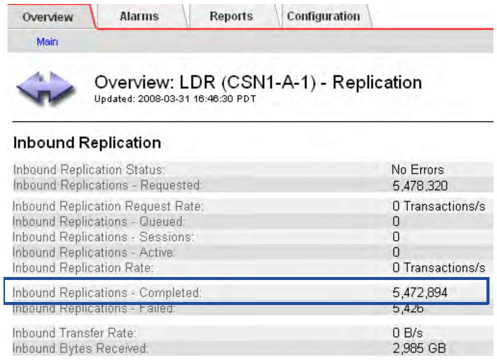 Overview: LDR (CSN1-A-1) - Replication Updated: 2008-03-31 16:46:30 PDT **Inbound Replication** Inbound Replication Status: No Errors Inbound Replications - Requested: 5,478,320 Inbound Replication Request Rate: 0 Transactions/s Inbound Replications - Queued: 0 Inbound Replications - Sessions: 0 Inbound Replications - Active: 0 Inbound Replication Rate: 0 Transactions/s Inbound Replications - Completed: 5,472,894 Inbound Replications - Failed: 5,426 Inbound Transfer Rate: 0 B/s Inbound Bytes Received: 2,985 GB	Inbound Replications Completed (RIRC): The total number of objects replicated to the destination LDR increases by 1. 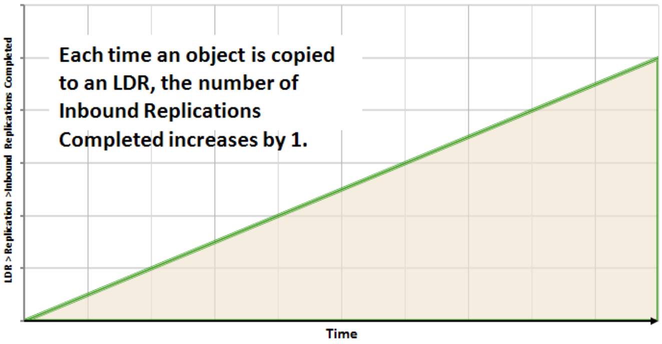 Each time an object is copied to an LDR, the number of Inbound Replications Completed increases by 1.
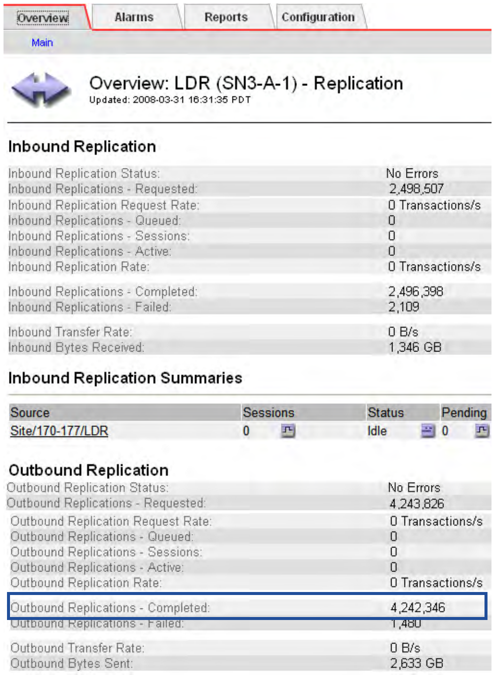 Overview: LDR (SN3-A-1) - Replication Updated: 2008-03-31 16:31:35 PDT **Inbound Replication** Inbound Replication Status: No Errors Inbound Replications - Requested: 2,498,507 Inbound Replication Request Rate: 0 Transactions/s Inbound Replications - Queued: 0 Inbound Replications - Sessions: 0 Inbound Replications - Active: 0 Inbound Replication Rate: 0 Transactions/s Inbound Replications - Completed: 2,496,398 Inbound Replications - Failed: 2,109 Inbound Transfer Rate: 0 B/s Inbound Bytes Received: 1,346 GB **Inbound Replication Summaries** Source / Sessions / Status / Pending Site/170-177/LDR / 0 / Idle / 0 **Outbound Replication** Outbound Replication Status: No Errors Outbound Replications - Requested: 4,243,826 Outbound Replication Request Rate: 0 Transactions/s Outbound Replications - Queued: 0 Outbound Replications - Sessions: 0 Outbound Replications - Active: 0 Outbound Replication Rate: 0 Transactions/s Outbound Replications - Completed: 4,242,346 Outbound Replications - Failed: 1,480 Outbound Transfer Rate: 0 B/s Outbound Bytes Sent: 2,633 GB	Outbound Replications Completed (RORC): The total number of objects replicated from the source LDR increases by 1. 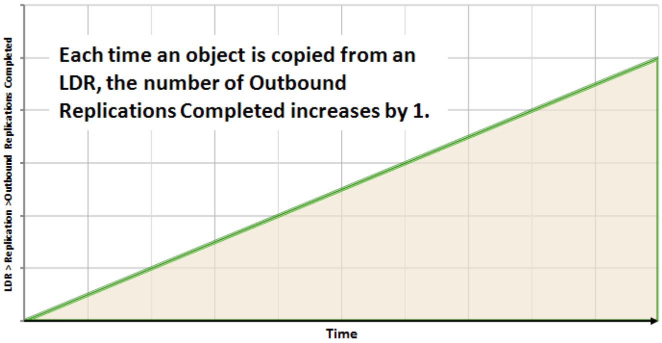 Each time an object is copied from an LDR, the number of Outbound Replications Completed increases by 1.

Content Replication to Archive Media

The Tape Node provides an interface between the grid and an archival media device which is external to the grid. The Tape Node communicates with a middleware layer that manages access to the physical storage device. Currently supported archive devices include any storage device managed by Tivoli Storage Manager (such as a tape library).

Data Flow

If the ILM policy requires an object to be stored on archive media, the CMS sends a request to the Tape Node which in turn sends the object to the middleware (see Figure 48).

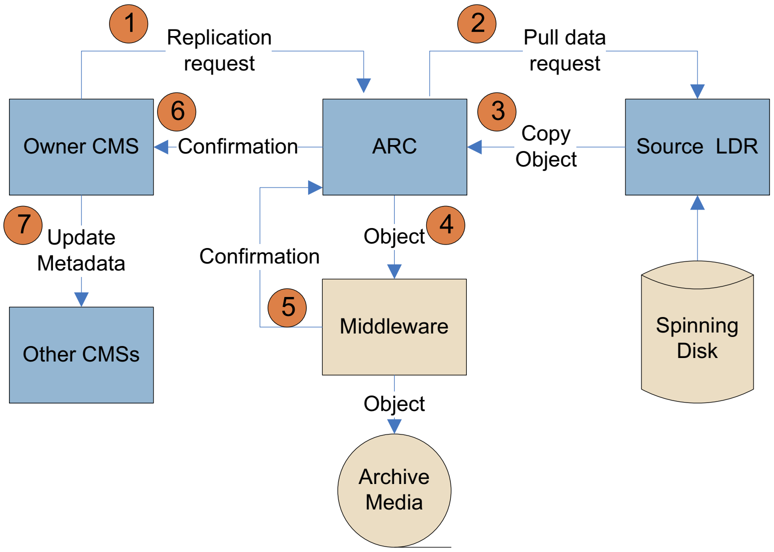

Figure 48 Archiving Data Flow

1	The owner CMS sends a request to the ARC to store a copy of the object on archive media.
2	The ARC queries the ADC for the best source location and sends a request to the source LDR.
3	The ARC retrieves the object from the LDR.
4	The ARC sends the object to the archiving middleware which in turn copies it to the archive media.

5	The middleware notifies the ARC that the object has been stored.
6	The ARC notifies the CMS that the object has been stored.
7	The owner CMS updates the location information and distributes that information to the other CMSs that store metadata for this object.

Related Attributes

Table 13 lists some of the NMS attributes used to track object replication to archive media.

Table 13 Object Replication to Tape Node Attributes

Component	Attribute Changes
ARC > Store 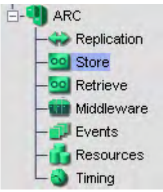 ARC Replication Store Retrieve Middleware Events Resources Timing 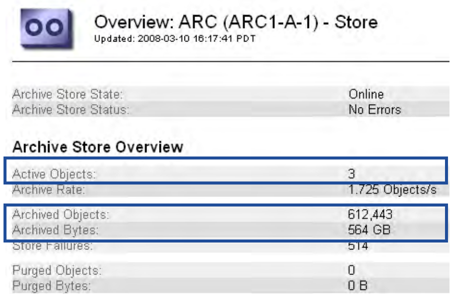 Overview: ARC (ARC1-A-1) - Store Updated: 2008-03-10 16:17:41 PDT Archive Store State: Online Archive Store Status: No Errors Archive Store Overview Active Objects: 3 Archive Rate: 1.725 Objects/s Archived Objects: 612,443 Archived Bytes: 564 GB Store Failures: 514 Purged Objects: 0 Purged Bytes: 0 B	Active Objects (AROP): The number of objects in the process of being written to archiving media increases by 1 as the object is being archived and decreases by 1 after the object is archived. 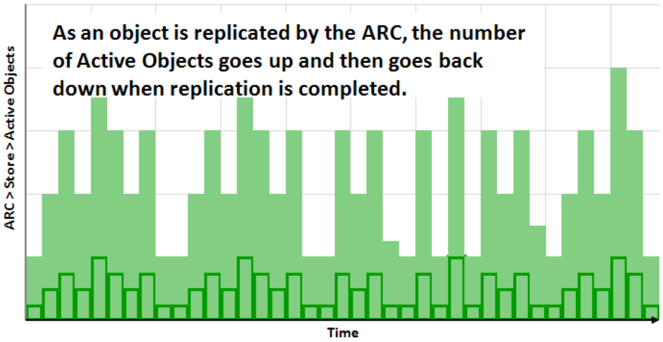 Archived Objects (AROA): The total number of objects written to archive media by this ARC increases by 1. 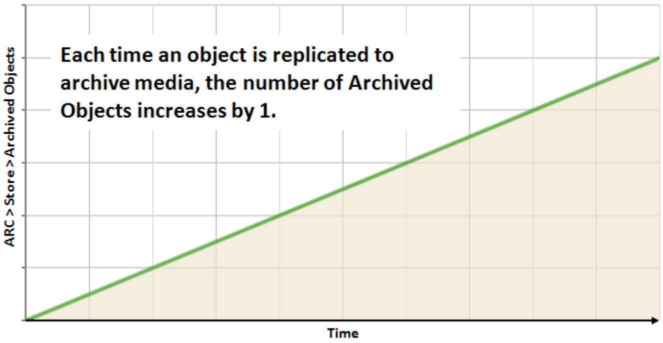 Archived Bytes (ARBA): The total amount of content written to archive media increases by an amount equivalent to the file size. The grid does not know how much installed and available storage is on the archival media device attached to the Tape Node.

Table 13 Object Replication to Tape Node Attributes *(continued)*

Component	Attribute Changes
ARC > Replication 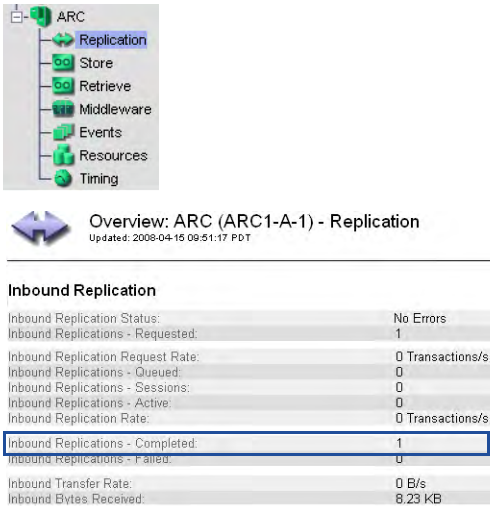 ARC Replication Store Retrieve Middleware Events Resources Timing Overview: ARC (ARC1-A-1) - Replication Updated: 2008-04-15 09:51:17 PDT **Inbound Replication** Inbound Replication Status: No Errors Inbound Replications - Requested: 1 Inbound Replication Request Rate: 0 Transactions/s Inbound Replications - Queued: 0 Inbound Replications - Sessions: 0 Inbound Replications - Active: 0 Inbound Replication Rate: 0 Transactions/s Inbound Replications - Completed: 1 Inbound Replications - Failed: 0 Inbound Transfer Rate: 0 B/s Inbound Bytes Received: 8.23 KB	Inbound Replications Completed (RIRC): The total number of objects replicated to the destination ARC increases by 1.
LDR > Storage 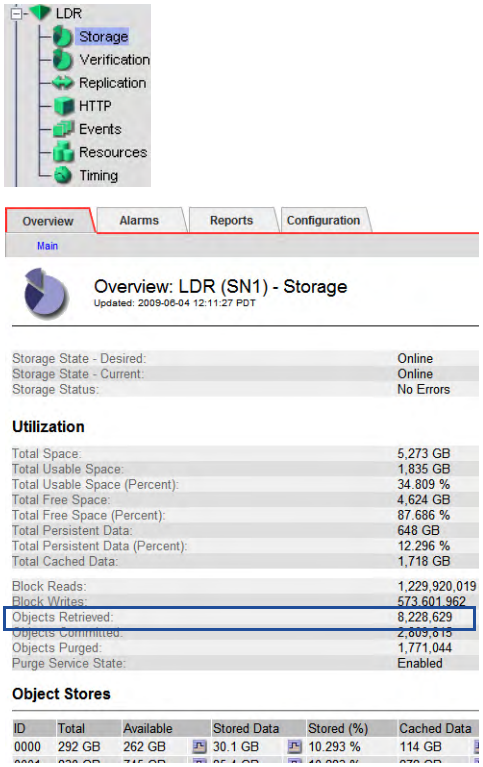 LDR Storage Verification Replication HTTP Events Resources Timing Overview Alarms Reports Configuration Main Overview: LDR (SN1) - Storage Updated: 2009-06-04 12:11:27 PDT Storage State - Desired: Online Storage State - Current: Online Storage Status: No Errors **Utilization** Total Space: 5,273 GB Total Usable Space: 1,835 GB Total Usable Space (Percent): 34.809 % Total Free Space: 4,624 GB Total Free Space (Percent): 87.686 % Total Persistent Data: 648 GB Total Persistent Data (Percent): 12.296 % Total Cached Data: 1,718 GB Block Reads: 1,229,920,019 Block Writes: 573,601,962 Objects Retrieved: 8,228,629 Objects Committed: 2,809,815 Objects Purged: 1,771,044 Purge Service State: Enabled **Object Stores** ID | Total | Available | Stored Data | Stored (%) | Cached Data 0000 | 292 GB | 262 GB | 30.1 GB | 10.293 % | 114 GB	Objects Retrieved (ORET): The number of persistent objects retrieved from the storage system of the source LDR increases by 1 each time the object is replicated from this source LDR to the Tape Node. 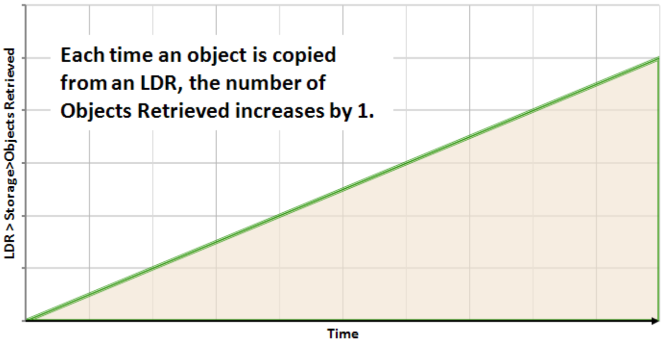 Each time an object is copied from an LDR, the number of Objects Retrieved increases by 1.

Table 13 Object Replication to Tape Node Attributes *(continued)*

Component	Attribute Changes
LDR > Replication 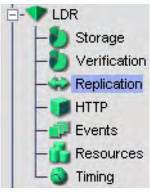 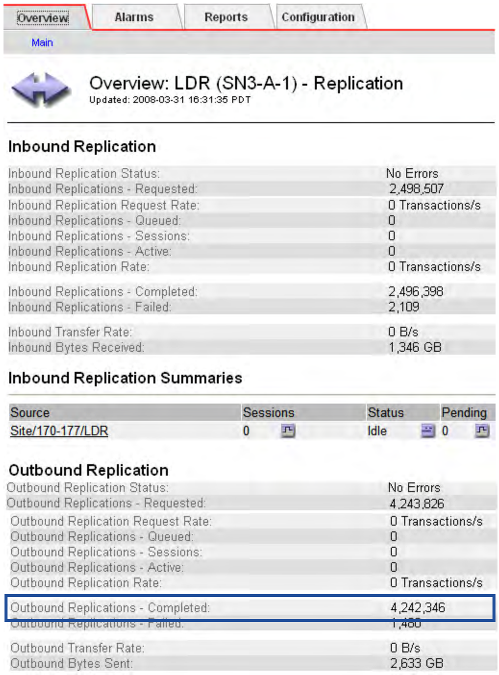	Outbound Replications - Completed (RORC): The total number of objects replicated from the source LDR increases by 1 each time the object is replicated from this source LDR to the Tape Node. 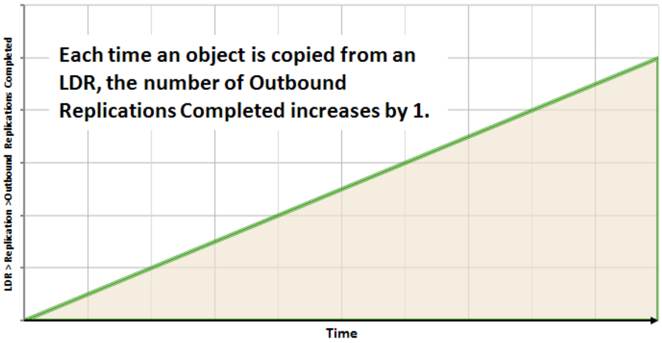

LDR
- Storage
- Verification
- Replication
- HTTP
- Events
- Resources
- Timing

Overview | Alarms | Reports | Configuration

Main

Overview: LDR (SN3-A-1) - Replication
Updated: 2008-03-31 16:31:35 PDT

Inbound Replication

Inbound Replication Status:	No Errors
Inbound Replications - Requested:	2,498,507
Inbound Replication Request Rate:	0 Transactions/s
Inbound Replications - Queued:	0
Inbound Replications - Sessions:	0
Inbound Replications - Active:	0
Inbound Replication Rate:	0 Transactions/s
Inbound Replications - Completed:	2,496,398
Inbound Replications - Failed:	2,109
Inbound Transfer Rate:	0 B/s
Inbound Bytes Received:	1,346 GB

Inbound Replication Summaries

Source	Sessions	Status	Pending
Site/170-177/LDR	0	Idle	0

Outbound Replication

Outbound Replication Status:	No Errors
Outbound Replications - Requested:	4,243,826
Outbound Replication Request Rate:	0 Transactions/s
Outbound Replications - Queued:	0
Outbound Replications - Sessions:	0
Outbound Replications - Active:	0
Outbound Replication Rate:	0 Transactions/s
Outbound Replications - Completed:	4,242,346
Outbound Replications - Failed:	1,480
Outbound Transfer Rate:	0 B/s
Outbound Bytes Sent:	2,633 GB

Retrieval

Retrieval refers to what happens when a client application accesses a file stored in the grid. There are two scenarios:

- The file is in the FSG cache.
- The file is not in the FSG cache.

Data Flow

File in FSG Cache

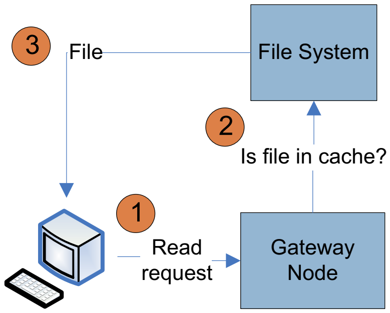

Figure 49 Retrieval Data Flow – File in Cache

1 The client requests the file.

2 The file system (CIFS/NFS) finds the file in the FSG cache.

3 The client reads the file.

File Not in FSG Cache

If the file is not in the FSG cache, the FSG sends a request to an LDR. The LDR returns the file if it has it. Otherwise, the LDR retrieves it from another LDR or an ARC after getting the file location from the CMS (see Figure 50). Retrieval preferentially goes to spinning media under normal performance load balancing. When there is no higher grade copy accessible, the retrieval request is directed to the Tape Node.

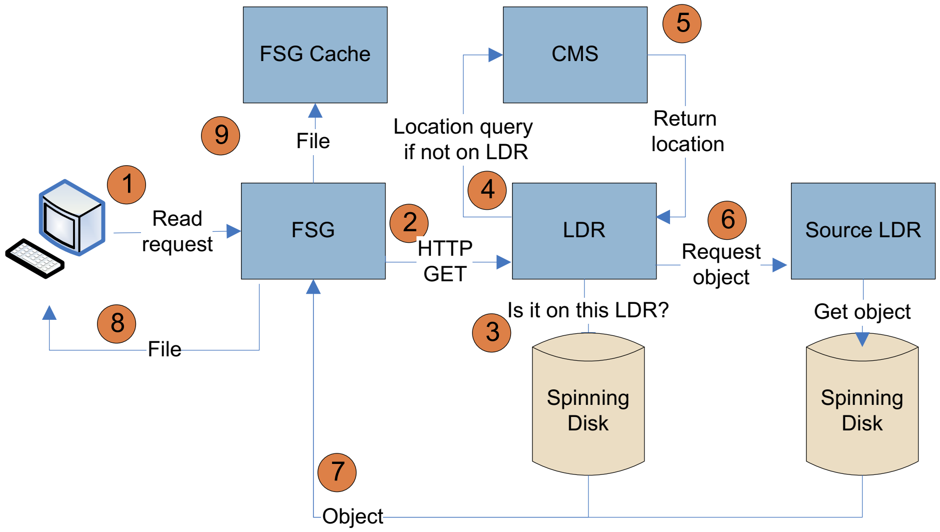

Figure 50 Retrieval Data Flow – File Not in Cache

1	The FSG receives a read request from a client.
2	Since the FSG cannot find the file in its cache, the FSG submits an HTTP GET command to an LDR. The LDR is chosen based on the result of a topology query.
3	The LDR checks if it has the object. If yes, it sends it to the FSG.
4	If the LDR does not find the object, it requests the location from a CMS. The CMS is chosen based on the result of a topology query.
5	The CMS returns the object location to the LDR.
6	The LDR retrieves the content from the LDR or ARC that has it and sends it to the FSG.
7	The LDR starts streaming the object to the FSG. Note that the object does not persist on the LDR.
8	The FSG sends the file to the client as soon as it has retrieved enough data from the LDR (it does not wait until it has retrieved the entire file).
9	The FSG stores the file in its cache for future use by the client.

Related Attributes

Table 14 lists some of the attributes used to track what happens when a client retrieves an object stored in the grid.

Table 14 Object Retrieval Attributes

Component	Attribute Changes
FSG > Storage 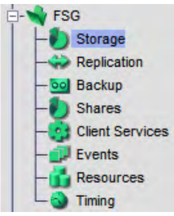 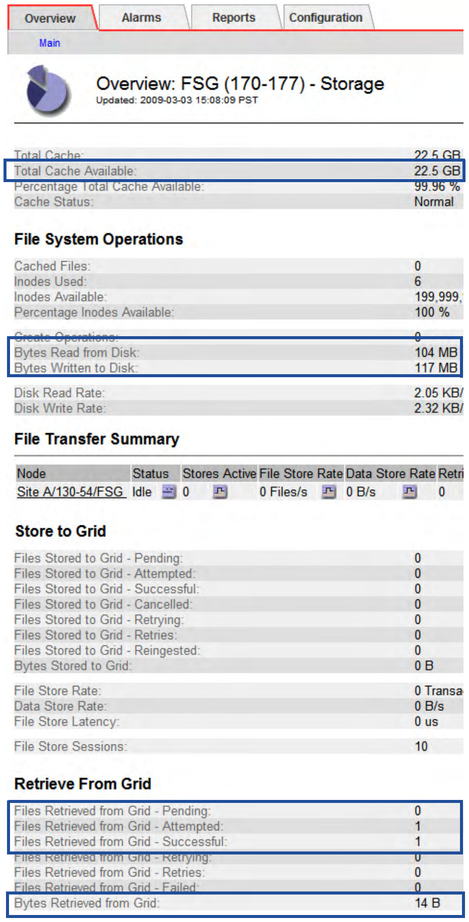	Total Cache Available (FSTA): The total local cache space that is still available for use does not change if the file was already in cache. Otherwise, total cache available decreases unless the minimum value has been reached. Bytes Read from Disk (FSRB): If the file is not in the FSG cache, the number of bytes read from disk increases. File retrieval generates multiple read operations as the FSG retrieves the file from the grid and the client application accesses the file. Bytes Written to Disk (FSWB): If the file is not in the FSG cache, the number of bytes written to disk increases. File retrieval generates multiple write operations as the FSG retrieves the file from the grid. Files Retrieved from Grid - Pending (FRGP) also increases by 1 when the transfer is requested and then decreases after it is completed. Files Retrieved from Grid - Attempted (FRGA): If the file is not in the FSG cache, the number of file retrieval requests waiting for a response from the grid increases by 1. Files Retrieved from Grid - Successful (FRGC): If the file is not in the FSG cache, the number of file transfer requests completed successfully increases by 1. 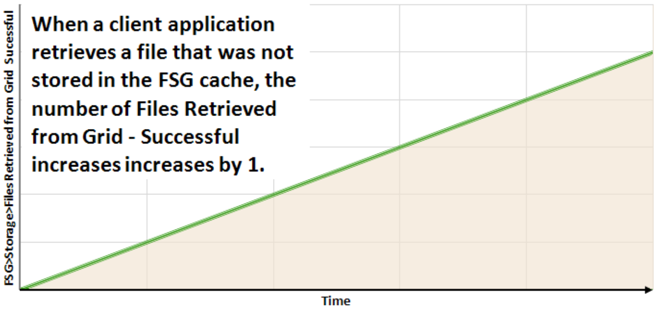 Bytes Retrieved from Grid (FRGB): If the file is not in the FSG cache, the number of bytes retrieved successfully from the grid increases by an amount equivalent to the file size.

FSG
- Storage
- Replication
- Backup
- Shares
- Client Services
- Events
- Resources
- Timing

Overview | Alarms | Reports | Configuration

Main

Overview: FSG (170-177) - Storage
Updated: 2009-03-03 15:08:09 PST

Total Cache:	22.5 GB
Total Cache Available:	22.5 GB
Percentage Total Cache Available:	99.96 %
Cache Status:	Normal

File System Operations

Cached Files:	0
Inodes Used:	6
Inodes Available:	199,999,
Percentage Inodes Available:	100 %
Create Operations:	0
Bytes Read from Disk:	104 MB
Bytes Written to Disk:	117 MB
Disk Read Rate:	2.05 KB/
Disk Write Rate:	2.32 KB/

File Transfer Summary

Node	Status	Stores Active	File Store Rate	Data Store Rate	Retri
Site A/130-54/FSG	Idle	0	0 Files/s	0 B/s	0

Store to Grid

Files Stored to Grid - Pending:	0
Files Stored to Grid - Attempted:	0
Files Stored to Grid - Successful:	0
Files Stored to Grid - Cancelled:	0
Files Stored to Grid - Retrying:	0
Files Stored to Grid - Retries:	0
Files Stored to Grid - Reingested:	0
Bytes Stored to Grid:	0 B
File Store Rate:	0 Transa
Data Store Rate:	0 B/s
File Store Latency:	0 us
File Store Sessions:	10

Retrieve From Grid

Files Retrieved from Grid - Pending:	0
Files Retrieved from Grid - Attempted:	1
Files Retrieved from Grid - Successful:	1
Files Retrieved from Grid - Retrying:	0
Files Retrieved from Grid - Retries:	0
Files Retrieved from Grid - Failed:	0
Bytes Retrieved from Grid:	14 B

Table 14 Object Retrieval Attributes *(continued)*

<table>
<tr><th>Component</th><th>Attribute Changes</th></tr>
<tr><td>LDR &gt; HTTP
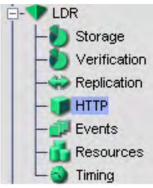

LDR
Storage
Verification
Replication
HTTP
Events
Resources
Timing

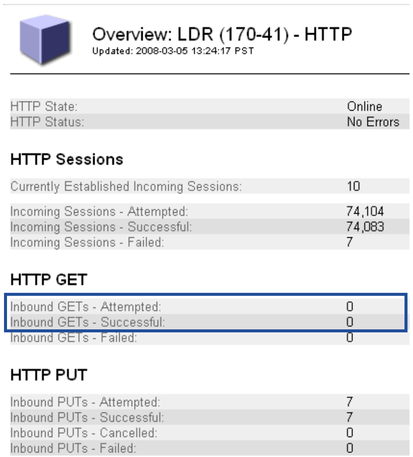
Overview: LDR (170-41) - HTTP
Updated: 2008-03-05 13:24:17 PST
<table>
<tr><td>HTTP State:</td><td>Online</td></tr>
<tr><td>HTTP Status:</td><td>No Errors</td></tr>
</table>
HTTP Sessions
<table>
<tr><td>Currently Established Incoming Sessions:</td><td>10</td></tr>
<tr><td>Incoming Sessions - Attempted:</td><td>74,104</td></tr>
<tr><td>Incoming Sessions - Successful:</td><td>74,083</td></tr>
<tr><td>Incoming Sessions - Failed:</td><td>7</td></tr>
</table>
HTTP GET
<table>
<tr><td>Inbound GETs - Attempted:</td><td>0</td></tr>
<tr><td>Inbound GETs - Successful:</td><td>0</td></tr>
<tr><td>Inbound GETs - Failed:</td><td>0</td></tr>
</table>
HTTP PUT
<table>
<tr><td>Inbound PUTs - Attempted:</td><td>7</td></tr>
<tr><td>Inbound PUTs - Successful:</td><td>7</td></tr>
<tr><td>Inbound PUTs - Cancelled:</td><td>0</td></tr>
<tr><td>Inbound PUTs - Failed:</td><td>0</td></tr>
</table>
</td><td>Inbound GETs - Successful (HIGC): If the file is not in the FSG cache, the total number of HTTP GET ("content retrieve") requests that have completed successfully increases by 1.
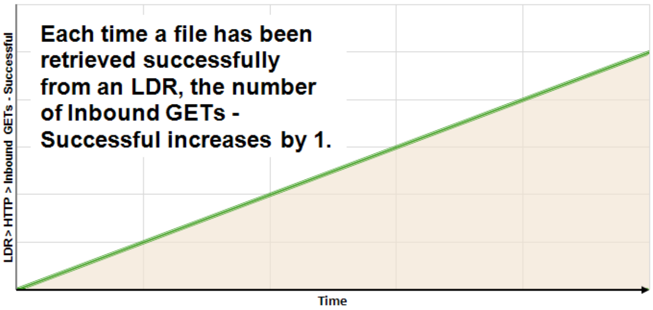
Each time a file has been retrieved successfully from an LDR, the number of Inbound GETs - Successful increases by 1.
Inbound GETs - Attempted (HAIG): The total number of HTTP GET (content retrieve) requests that have been received by the LDR also increases by 1.</td></tr>
<tr><td>LDR &gt; Storage
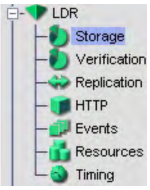

LDR
Storage
Verification
Replication
HTTP
Events
Resources
Timing

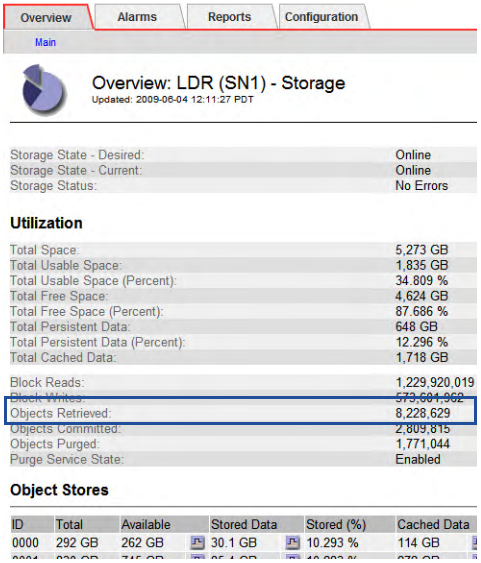
Overview: LDR (SN1) - Storage
Updated: 2009-06-04 12:11:27 PDT
<table>
<tr><td>Storage State - Desired:</td><td>Online</td></tr>
<tr><td>Storage State - Current:</td><td>Online</td></tr>
<tr><td>Storage Status:</td><td>No Errors</td></tr>
</table>
Utilization
<table>
<tr><td>Total Space:</td><td>5,273 GB</td></tr>
<tr><td>Total Usable Space:</td><td>1,835 GB</td></tr>
<tr><td>Total Usable Space (Percent):</td><td>34.809 %</td></tr>
<tr><td>Total Free Space:</td><td>4,624 GB</td></tr>
<tr><td>Total Free Space (Percent):</td><td>87.686 %</td></tr>
<tr><td>Total Persistent Data:</td><td>648 GB</td></tr>
<tr><td>Total Persistent Data (Percent):</td><td>12.296 %</td></tr>
<tr><td>Total Cached Data:</td><td>1,718 GB</td></tr>
<tr><td>Block Reads:</td><td>1,229,920,019</td></tr>
<tr><td>Block Writes:</td><td>573,601,962</td></tr>
<tr><td>Objects Retrieved:</td><td>8,228,629</td></tr>
<tr><td>Objects Committed:</td><td>2,809,815</td></tr>
<tr><td>Objects Purged:</td><td>1,771,044</td></tr>
<tr><td>Purge Service State:</td><td>Enabled</td></tr>
</table>
Object Stores
<table>
<tr><th>ID</th><th>Total</th><th>Available</th><th>Stored Data</th><th>Stored (%)</th><th>Cached Data</th></tr>
<tr><td>0000</td><td>292 GB</td><td>262 GB</td><td>30.1 GB</td><td>10.293 %</td><td>114 GB</td></tr>
</table>
</td><td>Objects Retrieved (ORET): If the file is not in the FSG cache, the number of persistent objects retrieved from the source LDR increases by 1.</td></tr>
</table>

Table 14 Object Retrieval Attributes *(continued)*

Component	Attribute Changes
ARC > Retrieve 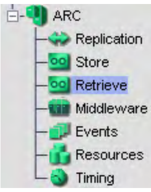 ARC Replication Store Retrieve Middleware Events Resources Timing 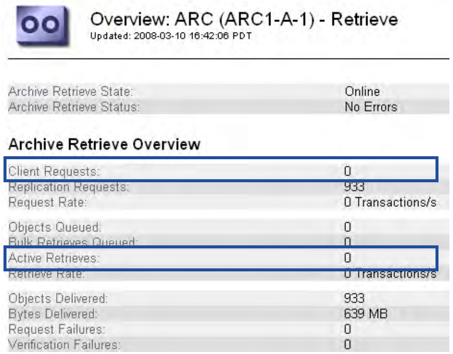 Overview: ARC (ARC1-A-1) - Retrieve Updated: 2008-03-10 16:42:06 PDT Archive Retrieve State: Online Archive Retrieve Status: No Errors Archive Retrieve Overview Client Requests: 0 Replication Requests: 933 Request Rate: 0 Transactions/s Objects Queued: 0 Bulk Retrieves Queued: 0 Active Retrieves: 0 Retrieve Rate: 0 Transactions/s Objects Delivered: 933 Bytes Delivered: 639 MB Request Failures: 0 Verification Failures: 0	Client Requests (ARCR): The total number of requests received from clients for objects stored on the ARC increases by 1 each time the grid attempts to retrieve an object from this Tape Node. This happens only if the file is not in the FSG cache or on an LDR. 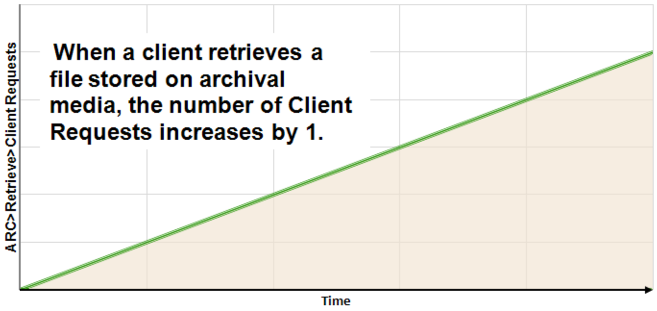 Active Retrieves (ARAR): The number of object retrievals in progress increases by 1 while retrieval is taking place on this Tape Node and then decreases by 1 after retrieval is complete. 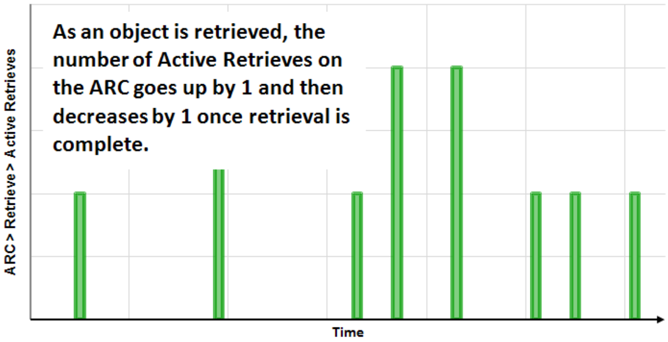

Purging

Removing an object from the grid's Storage Nodes and Tape Nodes is called "purging".

Data Flow

See Figure 51 for a simplified step-by-step description of what happens when the FSG receives a client request to delete a file.

NOTE By default, in order to protect against accidental or malicious object deletions, the ILM policy prevents content from being purged from the grid even if the client deletes the file on the FSG.

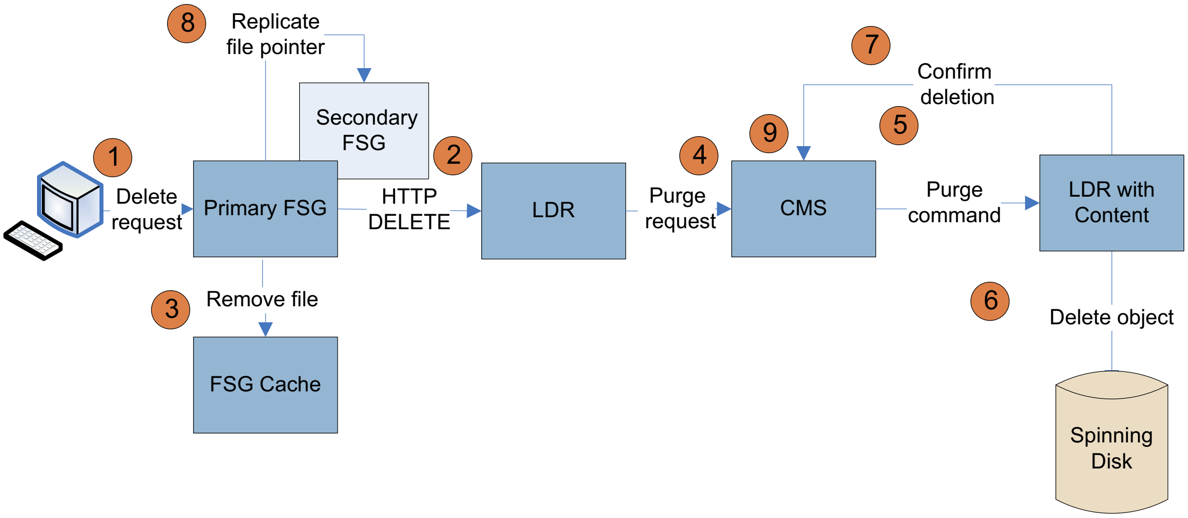

Figure 51 Purging Data Flow

1 The client sends a file delete request to the FSG which removes the file pointer from the FSG file system.

2 The FSG sends an HTTP DELETE command to an LDR. (This releases the content handle.) What LDR is chosen depends on the result of a topology query.

3 The file is removed from the FSG cache (that is, the file pointer and data are deleted).

4 The LDR notifies a CMS that the content handle has been released.

5 A purge command is sent to an LDR that has a copy of the object.

6 The LDR receives the purge command and deletes the content.

7	The LDR notifies the CMS that the content has been purged.
8	The primary FSG replicates the file pointer information to the other FSGs in its replication group.
9	In a grid that uses metadata replication, after all content locations are purged the metadata is purged. Metadata is never purged in a grid that uses metadata synchronization.

Related Attributes

Table 15 lists some of the attributes used to track what happens when a client deletes an object stored in the grid.

Table 15 Object Purging Attributes

Component	Attribute Changes
FSG > Storage 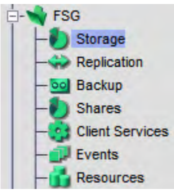 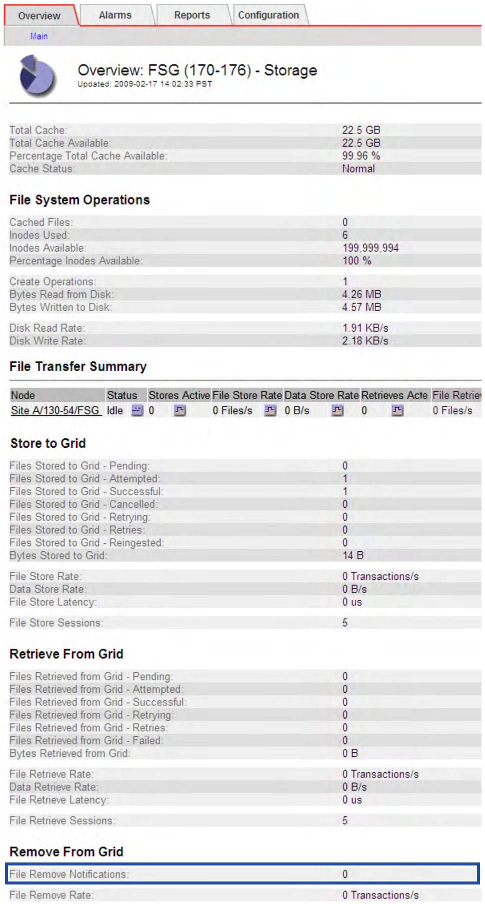	File Remove Notifications (FRGN): The number of content handle release notifications sent by the FSG to LDRs increases by 1 when the FSG receives a request to delete an object. 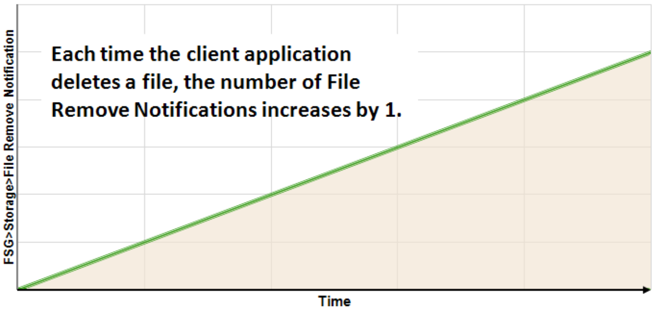

Overview: FSG (170-176) - Storage
Updated: 2009-02-17 14:02:33 PST

Total Cache:	22.5 GB
Total Cache Available:	22.5 GB
Percentage Total Cache Available:	99.96 %
Cache Status:	Normal

File System Operations

Cached Files:	0
Inodes Used:	6
Inodes Available:	199,999,994
Percentage Inodes Available:	100 %
Create Operations:	1
Bytes Read from Disk:	4.26 MB
Bytes Written to Disk:	4.57 MB
Disk Read Rate:	1.91 KB/s
Disk Write Rate:	2.18 KB/s

File Transfer Summary

Node	Status	Stores Active	File Store Rate	Data Store Rate	Retrieves Acte	File Retrie
Site A/130-54/FSG	Idle	0	0 Files/s	0 B/s	0	0 Files/s

Store to Grid

Files Stored to Grid - Pending:	0
Files Stored to Grid - Attempted:	1
Files Stored to Grid - Successful:	1
Files Stored to Grid - Cancelled:	0
Files Stored to Grid - Retrying:	0
Files Stored to Grid - Retries:	0
Files Stored to Grid - Reingested:	0
Bytes Stored to Grid:	14 B
File Store Rate:	0 Transactions/s
Data Store Rate:	0 B/s
File Store Latency:	0 us
File Store Sessions:	5

Retrieve From Grid

Files Retrieved from Grid - Pending:	0
Files Retrieved from Grid - Attempted:	0
Files Retrieved from Grid - Successful:	0
Files Retrieved from Grid - Retrying:	0
Files Retrieved from Grid - Retries:	0
Files Retrieved from Grid - Failed:	0
Bytes Retrieved from Grid:	0 B
File Retrieve Rate:	0 Transactions/s
Data Retrieve Rate:	0 B/s
File Retrieve Latency:	0 us
File Retrieve Sessions:	5

Remove From Grid

File Remove Notifications:	0
File Remove Rate:	0 Transactions/s

Each time the client application deletes a file, the number of File Remove Notifications increases by 1.

Table 15 Object Purging Attributes *(continued)*

<table>
<tr><th>Component</th><th>Attribute Changes</th></tr>
<tr><td>

FSG >Shares

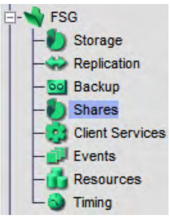

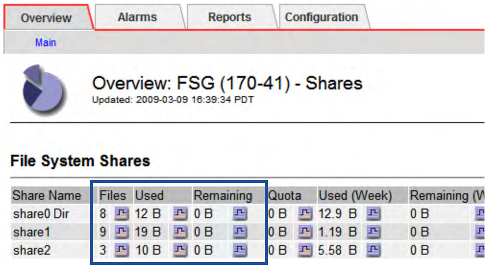

Overview: FSG (170-41) - Shares
Updated: 2009-03-09 16:39:34 PDT

File System Shares

<table>
<tr><th>Share Name</th><th>Files</th><th>Used</th><th>Remaining</th><th>Quota</th><th>Used (Week)</th><th>Remaining (W</th></tr>
<tr><td>share0 Dir</td><td>8</td><td>12 B</td><td>0 B</td><td>0 B</td><td>12.9 B</td><td>0 B</td></tr>
<tr><td>share1</td><td>9</td><td>19 B</td><td>0 B</td><td>0 B</td><td>1.19 B</td><td>0 B</td></tr>
<tr><td>share2</td><td>3</td><td>10 B</td><td>0 B</td><td>0 B</td><td>5.58 B</td><td>0 B</td></tr>
</table>

</td><td>

Files (FSSF): The number of files for the share where the object was ingested decreases by 1.

Used (FSSB) and Used (Week) (FSSR): The total size of all files referenced by the share decreases by an amount equivalent to the file size.

Remaining (FSSA) and Remaining (Week) (FSSL): If a quota has been configured for the share, the total amount of space remaining on the share increases by the size of the file.

</td></tr>
<tr><td>

LDR > HTTP

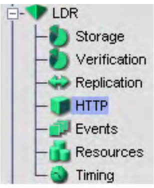

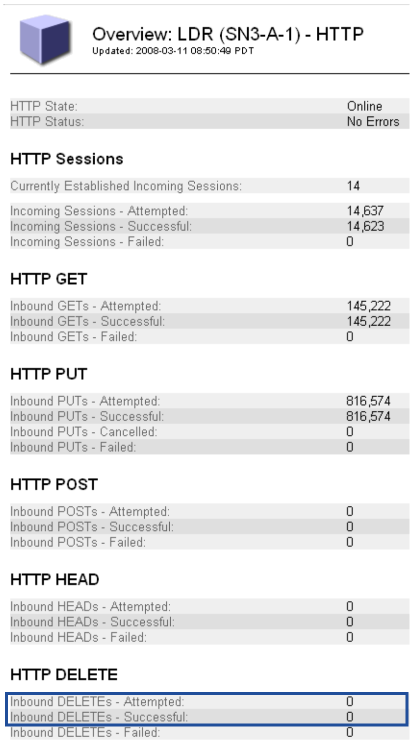

Overview: LDR (SN3-A-1) - HTTP
Updated: 2008-03-11 08:50:49 PDT

<table>
<tr><td>HTTP State:</td><td>Online</td></tr>
<tr><td>HTTP Status:</td><td>No Errors</td></tr>
</table>

HTTP Sessions

<table>
<tr><td>Currently Established Incoming Sessions:</td><td>14</td></tr>
<tr><td>Incoming Sessions - Attempted:</td><td>14,637</td></tr>
<tr><td>Incoming Sessions - Successful:</td><td>14,623</td></tr>
<tr><td>Incoming Sessions - Failed:</td><td>0</td></tr>
</table>

HTTP GET

<table>
<tr><td>Inbound GETs - Attempted:</td><td>145,222</td></tr>
<tr><td>Inbound GETs - Successful:</td><td>145,222</td></tr>
<tr><td>Inbound GETs - Failed:</td><td>0</td></tr>
</table>

HTTP PUT

<table>
<tr><td>Inbound PUTs - Attempted:</td><td>816,574</td></tr>
<tr><td>Inbound PUTs - Successful:</td><td>816,574</td></tr>
<tr><td>Inbound PUTs - Cancelled:</td><td>0</td></tr>
<tr><td>Inbound PUTs - Failed:</td><td>0</td></tr>
</table>

HTTP POST

<table>
<tr><td>Inbound POSTs - Attempted:</td><td>0</td></tr>
<tr><td>Inbound POSTs - Successful:</td><td>0</td></tr>
<tr><td>Inbound POSTs - Failed:</td><td>0</td></tr>
</table>

HTTP HEAD

<table>
<tr><td>Inbound HEADs - Attempted:</td><td>0</td></tr>
<tr><td>Inbound HEADs - Successful:</td><td>0</td></tr>
<tr><td>Inbound HEADs - Failed:</td><td>0</td></tr>
</table>

HTTP DELETE

<table>
<tr><td>Inbound DELETEs - Attempted:</td><td>0</td></tr>
<tr><td>Inbound DELETEs - Successful:</td><td>0</td></tr>
<tr><td>Inbound DELETEs - Failed:</td><td>0</td></tr>
</table>

</td><td>

Inbound DELETEs - Successful (HIDC): The total number of objects for which the content handle has been released increases by 1 after the HTTP DELETE command has been completed successfully.

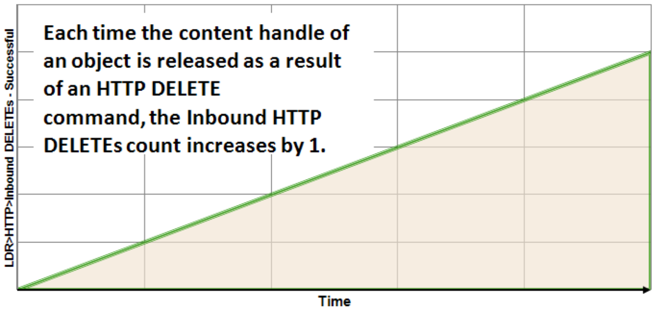

Inbound DELETEs - Attempted (HAID): The total number of HTTP DELETE (content handle release) requests that have been received by the LDR also increases by 1.

</td></tr>
</table>

Table 15 Object Purging Attributes *(continued)*

Component	Attribute Changes
LDR > Storage 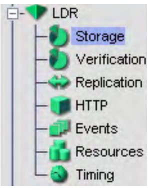 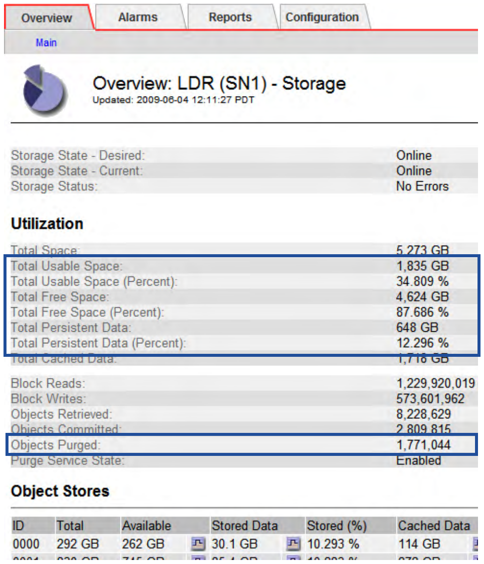	Objects Purged (OPUR): The number of persistent objects purged from this LDR increases by 1 for each copy of the object purged. 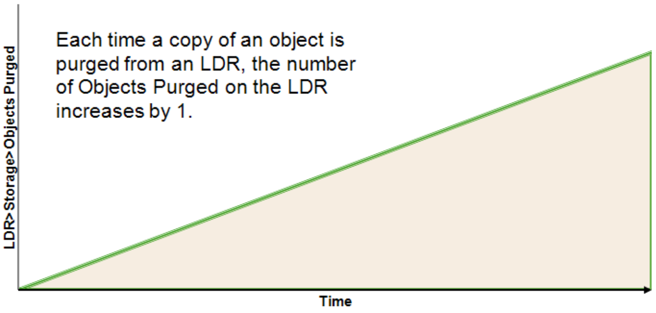 Total Usable Space (Percent) (SAVP): The percentage of object storage space available for use increases by an amount roughly equivalent to the size of the purged object. Total Persistent Data (Percent) (SPSD): The percentage of the total storage space used by persistent data decreases by an amount roughly equivalent to the size of the purged object.

Overview: LDR (SN1) - Storage
Updated: 2009-06-04 12:11:27 PDT

Storage State - Desired:	Online
Storage State - Current:	Online
Storage Status:	No Errors

Utilization

Total Space:	5,273 GB
Total Usable Space:	1,835 GB
Total Usable Space (Percent):	34.809 %
Total Free Space:	4,624 GB
Total Free Space (Percent):	87.686 %
Total Persistent Data:	648 GB
Total Persistent Data (Percent):	12.296 %
Total Cached Data:	1,718 GB
Block Reads:	1,229,920,019
Block Writes:	573,601,962
Objects Retrieved:	8,228,629
Objects Committed:	2,809,815
Objects Purged:	1,771,044
Purge Service State:	Enabled

Object Stores

ID	Total	Available	Stored Data	Stored (%)	Cached Data
0000	292 GB	262 GB	30.1 GB	10.293 %	114 GB

Each time a copy of an object is purged from an LDR, the number of Objects Purged on the LDR increases by 1.

Table 15 Object Purging Attributes *(continued)*

Component	Attribute Changes
CMS > Content 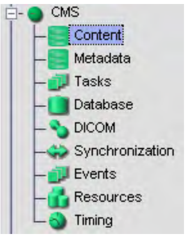 CMS Content Metadata Tasks Database DICOM Synchronization Events Resources Timing 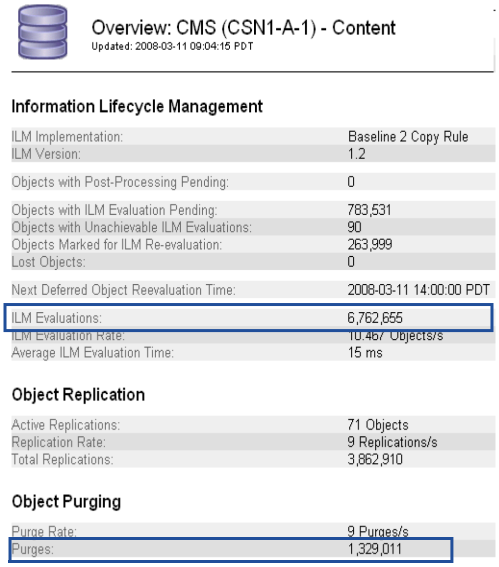 Overview: CMS (CSN1-A-1) - Content Updated: 2008-03-11 09:04:15 PDT **Information Lifecycle Management** ILM Implementation: Baseline 2 Copy Rule ILM Version: 1.2 Objects with Post-Processing Pending: 0 Objects with ILM Evaluation Pending: 783,531 Objects with Unachievable ILM Evaluations: 90 Objects Marked for ILM Re-evaluation: 263,999 Lost Objects: 0 Next Deferred Object Reevaluation Time: 2008-03-11 14:00:00 PDT ILM Evaluations: 6,762,655 ILM Evaluation Rate: 10.467 Objects/s Average ILM Evaluation Time: 15 ms **Object Replication** Active Replications: 71 Objects Replication Rate: 9 Replications/s Total Replications: 3,862,910 **Object Purging** Purge Rate: 9 Purges/s Purges: 1,329,011	Purges (DCpT): The number of object copies deleted increases by 1 on the owner CMS for each copy of the object purged from the LDRs and ARCs. 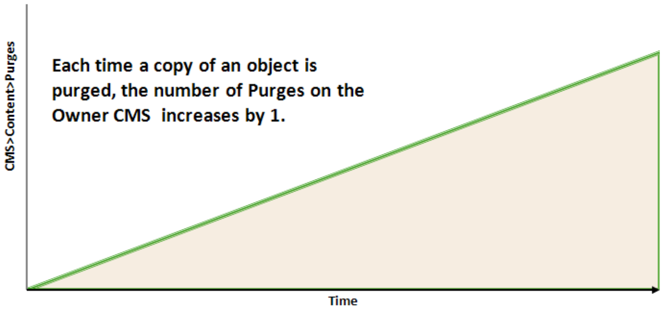 Each time a copy of an object is purged, the number of Purges on the Owner CMS increases by 1. CMS > Database > Free Tablespace Percent (DBSP) and the number of CMS > Metadata > Stored Objects (COoT) decrease in a grid that uses metadata replication. (Note that the owner CMS retains a copy of the metadata of a purged object until it can confirm that all object copies have been deleted, which can lead to inconsistencies in the value of CMS > Metadata > Stored Objects between CMSs. In a grid with a single CMS replication group, these inconsistencies should resolve within a day or two—when the last copy of metadata is purged.) Free Tablespace Percent (DBSP) and the number of Managed Objects (COoM) and Stored Objects (COoT) do not change in a grid that uses metadata synchronization. (With metadata synchronization, COoM is listed under the Content component, not the Metadata component.) ILM Evaluations (ILev): The total number of ILM evaluations that have been performed to date increases when an object is purged because an ILM evaluation is triggered each time a content handle is released.
ARC > Store 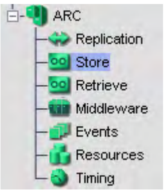 ARC Replication Store Retrieve Middleware Events Resources Timing 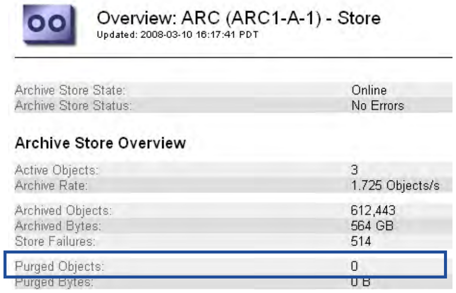 Overview: ARC (ARC1-A-1) - Store Updated: 2008-03-10 16:17:41 PDT Archive Store State: Online Archive Store Status: No Errors **Archive Store Overview** Active Objects: 3 Archive Rate: 1.725 Objects/s Archived Objects: 612,443 Archived Bytes: 564 GB Store Failures: 514 Purged Objects: 0 Purged Bytes: 0 B	Purged Objects (ADOP): The total number of objects purged from the middleware server by the ARC increases by 1. Whether objects purged from the ARC are actually deleted from archive media depends on the retention settings defined in the middleware server. 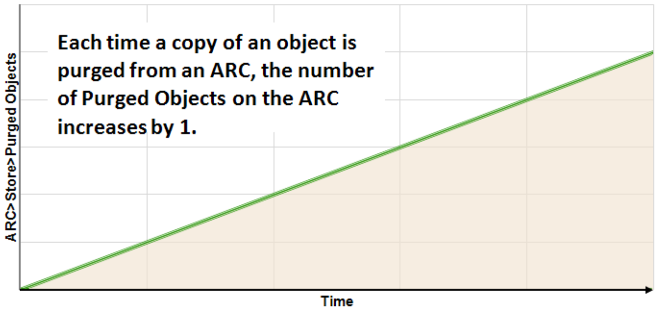 Each time a copy of an object is purged from an ARC, the number of Purged Objects on the ARC increases by 1.

Deletion Protection

Different levels of deletion protection can be enabled to protect files from being altered or removed after they have been ingested into the grid.

Purging Initiated by ILM Policy

Purging can also happen without a client request. For instance, an ILM policy could mandate that all content be automatically deleted two years after ingest. However, if objects are deleted from the grid without being deleted from the FSG first, links from the application to the objects will be broken and attempts to retrieve the objects from the FSG will fail.

File Modification

If the grid configuration allows it, client applications may modify content that has already been ingested into the grid.

When a file is modified, the content handle of the original file is released and the object is purged from the grid according to the ILM policy. The modified file is assigned a different unique identifier and the object is treated like a new file ingest. The grid does not track the multiple versions of the object.

Data Flow

See Figure 52 for a simplified step-by-step description of what happens when the client modifies a file stored in the grid.

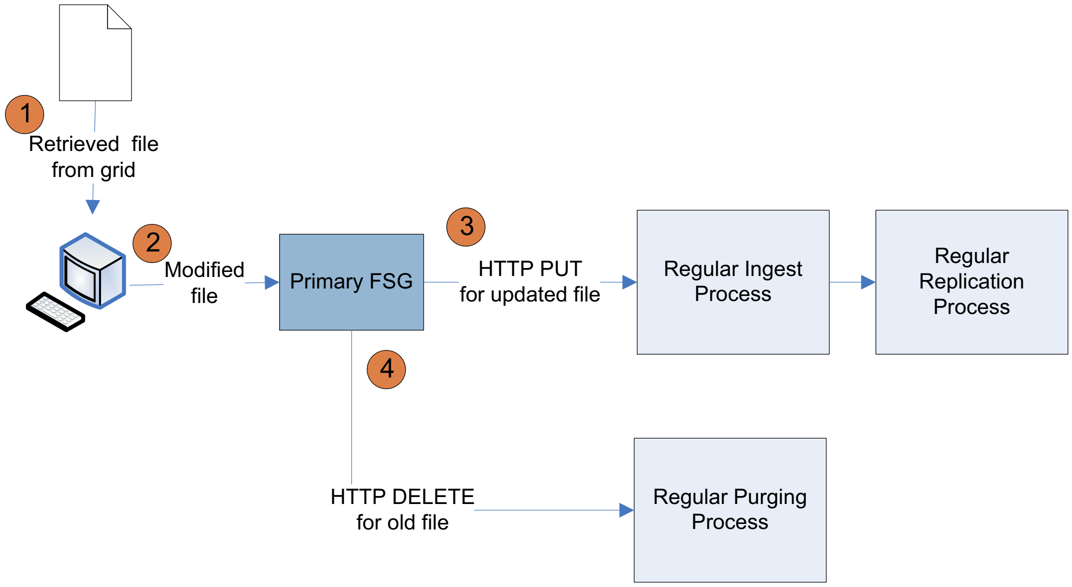

Figure 52 Object Modification Data Flow

1	The client application retrieves the file from the grid via the FSG file share. For more information, see Retrieval (page 85).
2	The client application saves the modified file to the grid via the FSG file share.
3	The primary FSG saves a copy of the object in its cache and sends an LDR a HTTP PUT command. The data flow is identical to that for ingest. The LDR saves the object to spinning disk and allocates a new unique identifier to the object. The LDR notifies a CMS that a new piece of content has been ingested and sends the object metadata to the CMS. This CMS becomes the owner CMS. The object is then replicated according to the ILM policy. For more information, see Ingest (page 67) and Content Replication (page 76).
4	The primary FSG also sends a HTTP DELETE command to an LDR to delete the old version of the object. What happens next is identical to what happens when an object is purged. For more information, see Purging (page 90).

Related Attributes

The key attributes that change when a file is modified are essentially the same as when a file is ingested and replicated, and then purged. One notable exception is Create Operations. You can tell when a file has been modified if File Remove Notifications and Files Stored To Grid have both increased by 1 but Create Operations has not changed.

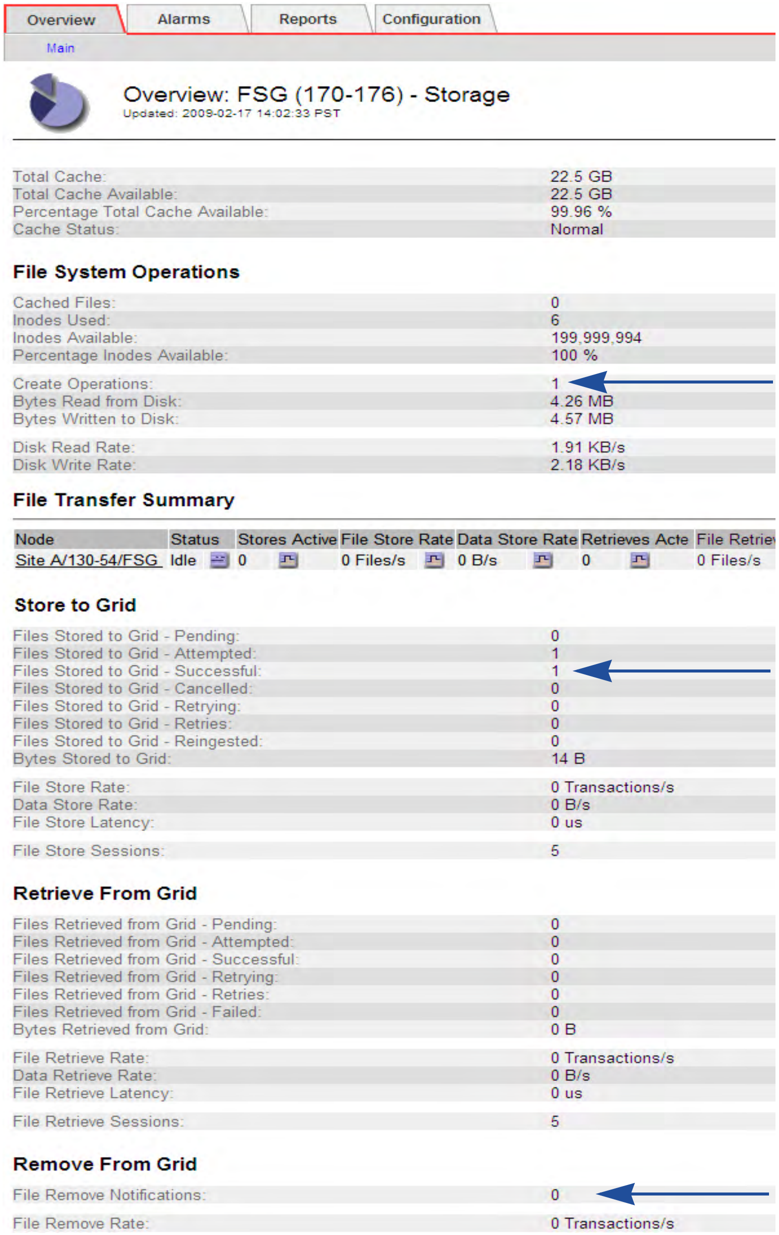

Overview: FSG (170-176) - Storage
Updated: 2009-02-17 14:02:33 PST

Total Cache:	22.5 GB
Total Cache Available:	22.5 GB
Percentage Total Cache Available:	99.96 %
Cache Status:	Normal

File System Operations

Cached Files:	0
Inodes Used:	6
Inodes Available:	199,999,994
Percentage Inodes Available:	100 %
Create Operations:	1
Bytes Read from Disk:	4.26 MB
Bytes Written to Disk:	4.57 MB
Disk Read Rate:	1.91 KB/s
Disk Write Rate:	2.18 KB/s

File Transfer Summary

Node	Status	Stores Active	File Store Rate	Data Store Rate	Retrieves Acte	File Retriev
Site A/130-54/FSG	Idle	0	0 Files/s	0 B/s	0	0 Files/s

Store to Grid

Files Stored to Grid - Pending:	0
Files Stored to Grid - Attempted:	1
Files Stored to Grid - Successful:	1
Files Stored to Grid - Cancelled:	0
Files Stored to Grid - Retrying:	0
Files Stored to Grid - Retries:	0
Files Stored to Grid - Reingested:	0
Bytes Stored to Grid:	14 B
File Store Rate:	0 Transactions/s
Data Store Rate:	0 B/s
File Store Latency:	0 us
File Store Sessions:	5

Retrieve From Grid

Files Retrieved from Grid - Pending:	0
Files Retrieved from Grid - Attempted:	0
Files Retrieved from Grid - Successful:	0
Files Retrieved from Grid - Retrying:	0
Files Retrieved from Grid - Retries:	0
Files Retrieved from Grid - Failed:	0
Bytes Retrieved from Grid:	0 B
File Retrieve Rate:	0 Transactions/s
Data Retrieve Rate:	0 B/s
File Retrieve Latency:	0 us
File Retrieve Sessions:	5

Remove From Grid

File Remove Notifications:	0
File Remove Rate:	0 Transactions/s

Create Operations does not change when a file is modified

Files Stored to Grid increases when a file is modified

File Remove Notifications increases when a file is modified

Figure 53 FSG Storage Attributes For Modified Object

FSG Replication

The primary FSG maintains a system of file pointers to the objects stored in the grid. As seen in Figure 53, the file pointer system is modified each time a file is ingested, changed, or deleted.

The primary FSG must replicate its file pointer system to the secondary FSG to ensure redundancy in case the primary FSG becomes unavailable.

During normal operation, the file pointers are replicated in real time, as files are ingested, modified, or deleted.

The backup together with the active session file can be used to restore the managed file system should it become corrupted. The active session file is a log of the FSG activity.

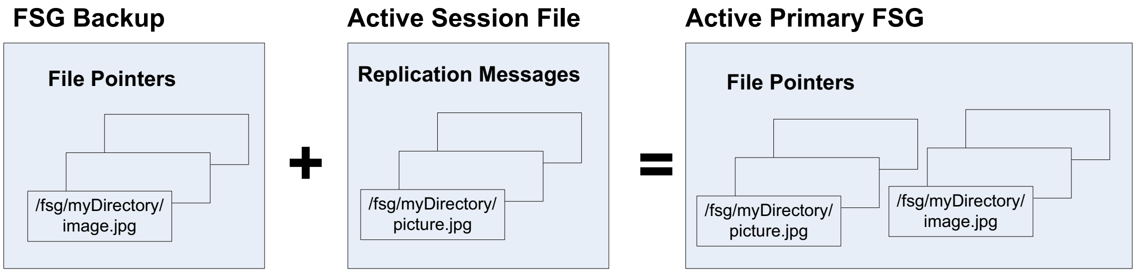
FSG Backup
File Pointers
/fsg/myDirectory/ image.jpg
+
Active Session File
Replication Messages
/fsg/myDirectory/ picture.jpg
=
Active Primary FSG
File Pointers
/fsg/myDirectory/ picture.jpg
/fsg/myDirectory/ image.jpg

Figure 54 FSG Backups and Active Session File

See Figure 52 for a simplified description of the FSG replication message flow when a file is ingested, deleted, or modified.

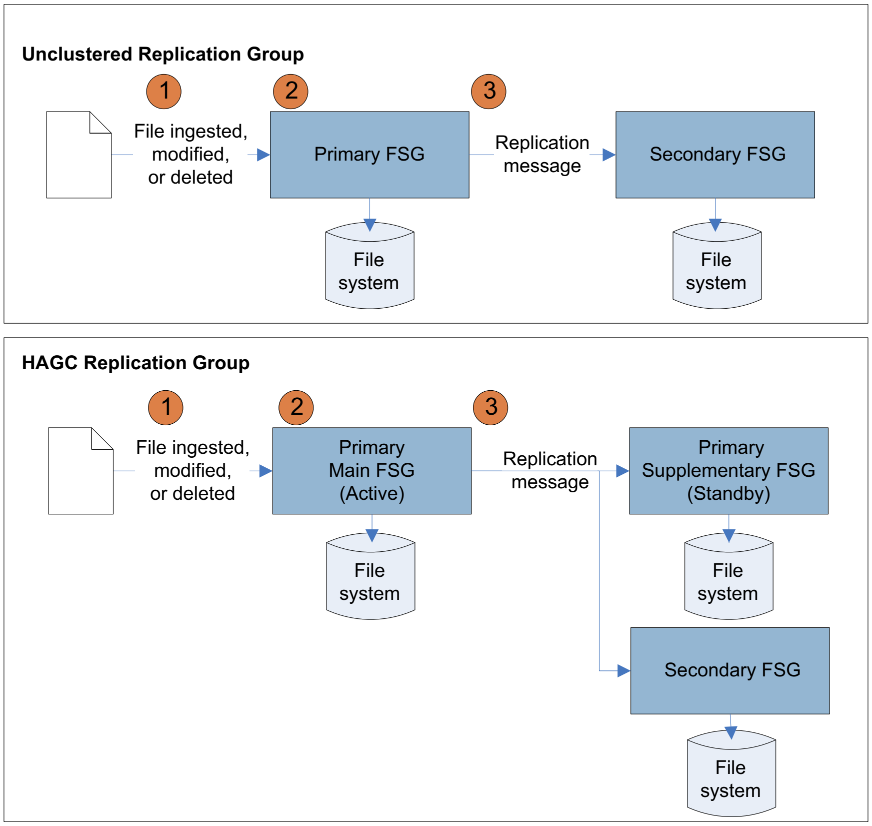
Unclustered Replication Group
1
2
3
File ingested, modified, or deleted
Primary FSG
Replication message
Secondary FSG
File system
File system
HAGC Replication Group
1
2
3
File ingested, modified, or deleted
Primary Main FSG (Active)
Replication message
Primary Supplementary FSG (Standby)
File system
File system
Secondary FSG
File system

Figure 55 FSG Replication Data Flow

1 The client creates, modifies, or deletes a file via the FSG file share.

2 The primary FSG either creates a file pointer, modifies the file pointer, or deletes the file pointer from its file system.

3 The primary FSG sends a replication messages to the other FSGs in its replication group.

4 The other FSGs in the replication group process the replication messages in real time and update their file system.

Related Attributes

Table 16 lists some of the NMS MI attributes used to track FSG replication.

Table 16 FSG Replication Attributes

Component	Attribute Changes
Primary FSG > Replication 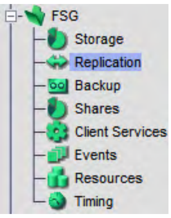 FSG Storage Replication Backup Shares Client Services Events Resources Timing 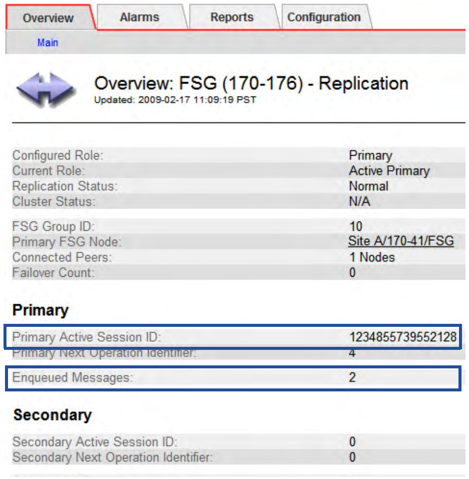 Overview: FSG (170-176) - Replication Updated: 2009-02-17 11:09:19 PST Configured Role: Primary Current Role: Active Primary Replication Status: Normal Cluster Status: N/A FSG Group ID: 10 Primary FSG Node: Site A/170-41/FSG Connected Peers: 1 Nodes Failover Count: 0 **Primary** Primary Active Session ID: 1234855739552128 Primary Next Operation Identifier: 4 Enqueued Messages: 2 **Secondary** Secondary Active Session ID: 0 Secondary Next Operation Identifier: 0	Primary Active Session ID (PAID): A new unique identifier for the current replication session is assigned each time a new session is started by the active primary FSG. A new session is started when a FSG failover occurs or when the size or age of the current session file exceeds an internal threshold. This number is the same as the Secondary Active Session ID (SAID) on the secondary FSG (see Secondary FSGs (page 197)). Enqueued Messages (PEOP): The total count of replication messages generated increases each time a file is ingested, modified or purged. The count increases by more than one for some operations. For example, ingest generates two replication messages (one for the initial file creation and a second message to associate the UUID after the file has been ingested). Modify also generates two replication messages (one to release the old UUID and one to assign the new UUID after ingest). Attribute events on the file (change permissions, etc.) may also generate additional replication messages. The number of enqueued messages matches the number of dequeued messages at the secondary FSG (see page 101). Because this FSG is the primary FSG, the fields in the "Secondary" section do not apply.

Table 16 FSG Replication Attributes *(continued)*

Component	Attribute Changes
Secondary FSG > Replication 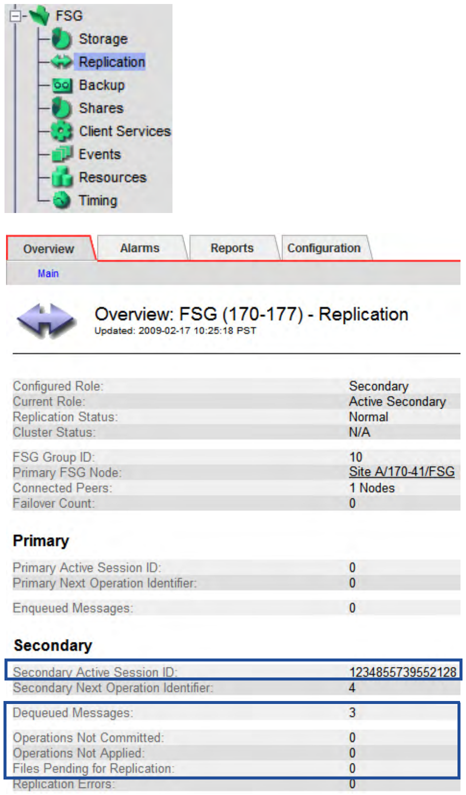 FSG Storage Replication Backup Shares Client Services Events Resources Timing Overview Alarms Reports Configuration Main Overview: FSG (170-177) - Replication Updated: 2009-02-17 10:25:18 PST Configured Role: Secondary Current Role: Active Secondary Replication Status: Normal Cluster Status: N/A FSG Group ID: 10 Primary FSG Node: Site A/170-41/FSG Connected Peers: 1 Nodes Failover Count: 0 **Primary** Primary Active Session ID: 0 Primary Next Operation Identifier: 0 Enqueued Messages: 0 **Secondary** Secondary Active Session ID: 1234855739552128 Secondary Next Operation Identifier: 4 Dequeued Messages: 3 Operations Not Committed: 0 Operations Not Applied: 0 Files Pending for Replication: 0 Replication Errors: 0	Secondary Active Session ID (SAID): The unique identifier for the current replication session from which messages are being processed is the same as the Primary Active Session ID (PAID) on the primary FSG (see page 100). Dequeued Messages (SDOP): The total count of dequeued replication messages increases each time a message has been processed by the secondary FSG. The number of dequeued messages matches the number of enqueued messages at the primary FSG (see page 100). Operations Not Committed (SUOP): The number of replication messages from the primary FSG that have not yet been written to the secondary FSG increases temporarily during periods of high grid activity. Operations Not Applied (SPOP): The number of replication messages to be processed on the secondary FSG in order to catch up to the primary FSG may temporarily increase during periods of high grid activity or backups. (In older systems where “offline backups” were used, SPOP increased during backups. As of Release 8.1, offline backups are deprecated.) If this FSG is the standby primary FSG in an HAGC, the fields in the “Primary” section do not apply. The same is true for a secondary FSG. Files Pending for Replication (FRPP): Number of new files that have not been fully replicated because they are awaiting a replication message from the primary FSG indicating that the file has been successfully stored to the grid. Similar to Operations Not Applied (SPOP), the value for this attribute may temporarily increase during periods of high grid activity.

4 Operations

This chapter describes common routine tasks that you perform as a grid operator:

- Monitor trends
- Monitor FSG backups
- Monitor LDR verification
- Monitor the Tape Node
- Monitor grid tasks

Top Attributes

The NMS MI displays hundreds of attributes; however, most of these attributes are required only for troubleshooting. The list of attributes to monitor routinely, shown in Table 17, is much shorter. Tips on how to analyze these attributes are described in the remainder of this chapter.

Table 17 Key Attributes to Monitor

Category	Component	Code	Description	See
LDR content storage capacity	Grid Overview	PSCU	Percentage Storage Capacity Used: The percentage of installed storage capacity that has been used up for the entire grid.	page 108
	LDR Storage 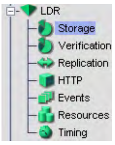 LDR Storage Verification Replication HTTP Events Resources Timing	STAS	Total Usable Space: Object storage capacity that is currently available for storage on the LDR.	
		SUSA	Total Free Space: Total amount of all free space—available to be used or not—on all object stores.	

Table 17 Key Attributes to Monitor *(continued)*

Category	Component	Code	Description	See
CMS metadata storage capacity	Grid Overview	PMCA	Percentage Metadata Capacity Available: An estimate of how much metadata capacity remains in the CMS databases in the grid.	
	CMS Database 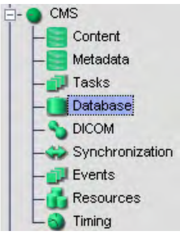 CMS Content Metadata Tasks Database DICOM Synchronization Events Resources Timing	CORS	Estimated Remaining Object Capacity: The number of additional objects that can be managed by this CMS.	page 110
		DBSP	Free Tablespace (Percent): The metadata storage capacity that is still available for use on this CMS.	
Ingest load	FSG Storage 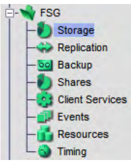 FSG Storage Replication Backup Shares Client Services Events Resources Timing	FSGP	Files Stored to Grid - Pending: The number of new ingested files cached locally that are waiting for transfer to the grid for persistent storage.	page 116
Retrieve load	FSG Storage 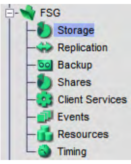 FSG Storage Replication Backup Shares Client Services Events Resources Timing	FRTM	File Retrieve Latency (FRTM): The average amount of time required to retrieve a file from the grid. Look at this attribute along with the related attributes Bytes Retrieved from Grid, File Retrieve Rate, and Data Retrieve Rate.	page 118

Table 17 Key Attributes to Monitor *(continued)*

Category	Component	Code	Description	See
FSG capacity	FSG Backup 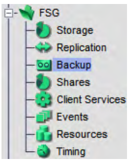 FSG Storage Replication Backup Shares Client Services Events Resources Timing	PBNF	Number of Files: The number of files included in the last FSG backup.	page 112
		PBDS	Backup Data Size: The total size of all files referenced by the last FSG backup.	
	FSG Shares 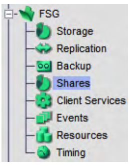 FSG Storage Replication Backup Shares Client Services Events Resources Timing	FSSF	Files: The number of files saved to the share. Calculated from the last FSG backup.	
		FSSB	Used: The amount of storage space used by the share. Calculated from the last FSG backup.	
		FSSA	Remaining: If a quota has been configured for the share, the amount of space remaining on the share. Calculated from the last FSG backup.	
	FSG Storage 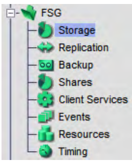 FSG Storage Replication Backup Shares Client Services Events Resources Timing	FSIU	Inodes Used: Number of inodes (files and directories) used on the FSG filesystem.	
		FSIA	Inodes Available: Number of inodes (files and directories) available on the FSG filesystem.	
		FSIP	Percentage Inodes Available: The percentage of inodes (files and directories) available on the FSG filesystem.	
Object replication	CMS Content 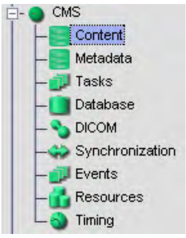 CMS Content Metadata Tasks Database DICOM Synchronization Events Resources Timing	ORun	Objects with Unachievable ILM Evaluations: The number of objects whose ILM business rules cannot be met because the topology or operational state of the grid prevents the rules from being satisfied.	page 120
		ORpe	Objects with ILM Evaluation Pending: The number of objects waiting to be processed through the business rules for replication.	
Metadata synchronization	CMS Synchronization 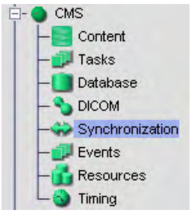 CMS Content Tasks Database DICOM Synchronization Events Resources Timing	CsQT	Queue Size: The number of outgoing metadata synchronization messages queued to be sent to other CMSs.	page 121

Table 17 Key Attributes to Monitor *(continued)*

Category	Component	Code	Description	See
FSG replication	Secondary FSG Replication 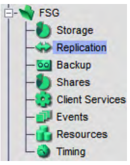 FSG Storage Replication Backup Shares Client Services Events Resources Timing	FRPP	Files Pending for Replication: Number of new files that have not been fully replicated because they are awaiting a replication message from the primary FSG indicating that the file has been successfully stored to the grid.	page 122
		SUOP	Operations Not Committed: The number of replication messages from the primary FSG that have not been written to the secondary FSG yet.	
		SPOP	Operations Not Applied: The number of replication messages to be processed by the secondary FSG in order to catch up to the primary FSG.	
Amount of content written to archive media	ARC Store 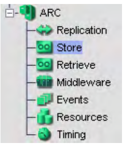 ARC Replication Store Retrieve Middleware Events Resources Timing	ARBA	Archived Bytes: The total amount of content written to archive media by this ARC.	page 129

Regular Tasks

Table 18 lists the tasks to be performed on a regular basis.

Table 18 Regular Tasks

Task	Frequency	See
Monitor System Status. Note what has changed from previous day.	Daily	page 45
Monitor system status lights on hardware.	Daily	
Monitor the rate at which LDR storage capacity is being used up.	Weekly	page 108
Monitor the rate at which content metadata storage capacity on the CMS is being used up.	Weekly	page 110
Monitor FSG capacity.	Weekly	page 112
Monitor FSG file share usage.	Weekly	page 114
Check available space on the archive media.	Weekly	page 129

- Monitor the key attributes regularly to become familiar with grid operations and spot trends before they turn into problems. The important attributes to monitor relate to:
- Content storage capacity on LDRs
- Metadata storage capacity on CMSs
- FSG capacity
- Attribute storage capacity on NMSs
- Ingest load on FSGs
- Retrieve load on FSGs
- Metadata synchronization/replication on CMSs
- ILM replication on CMSs

In the case of the capacity attributes—for example, LDR content storage space—you must not only look at the absolute value, but also at the rate at which capacity is being consumed.

Content Storage Capacity

The LDRs on Storage Nodes are responsible for storing objects in the grid. You need to monitor the total usable space available on Storage Nodes to make sure the grid does not run out of space to store content. This information is available at the grid level (see Figure 56), at the site level, and at the node level (see Figure 57).

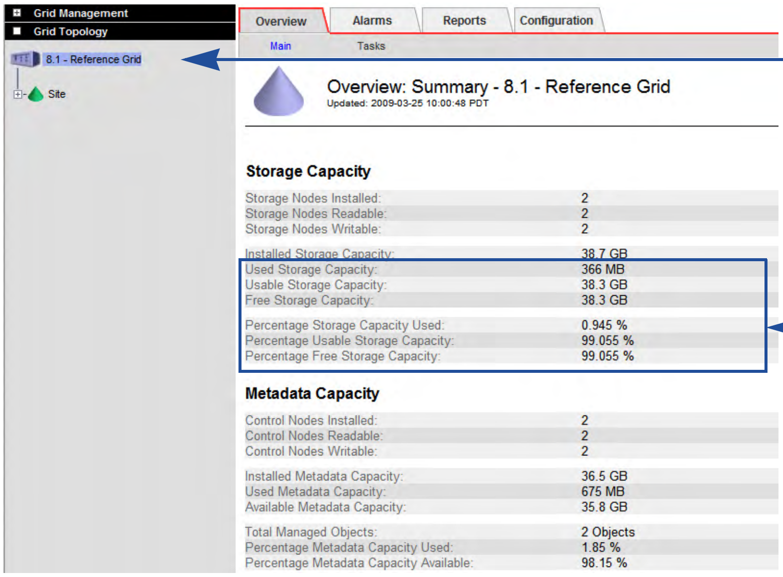

Overview: Summary - 8.1 - Reference Grid
Updated: 2009-03-25 10:00:48 PDT

Storage Capacity

Storage Nodes Installed:	2
Storage Nodes Readable:	2
Storage Nodes Writable:	2
Installed Storage Capacity:	38.7 GB
Used Storage Capacity:	366 MB
Usable Storage Capacity:	38.3 GB
Free Storage Capacity:	38.3 GB
Percentage Storage Capacity Used:	0.945 %
Percentage Usable Storage Capacity:	99.055 %
Percentage Free Storage Capacity:	99.055 %

Metadata Capacity

Control Nodes Installed:	2
Control Nodes Readable:	2
Control Nodes Writable:	2
Installed Metadata Capacity:	36.5 GB
Used Metadata Capacity:	675 MB
Available Metadata Capacity:	35.8 GB
Total Managed Objects:	2 Objects
Percentage Metadata Capacity Used:	1.85 %
Percentage Metadata Capacity Available:	98.15 %

Summary attributes at the grid level

Usable storage capacity and free storage capacity attributes

Figure 56 Overall Storage Node Capacity

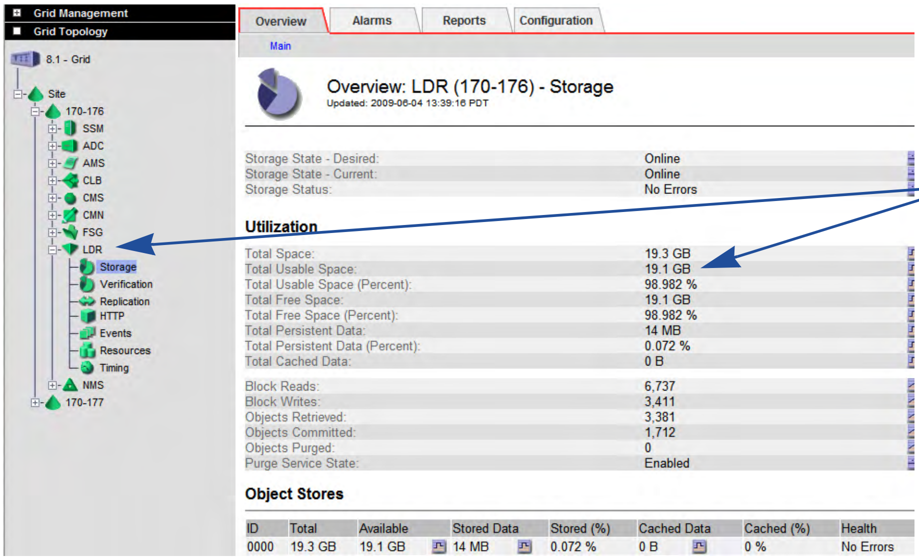

Overview: LDR (170-176) - Storage
Updated: 2009-06-04 13:39:16 PDT

Storage State - Desired:	Online
Storage State - Current:	Online
Storage Status:	No Errors

Utilization

Total Space:	19.3 GB
Total Usable Space:	19.1 GB
Total Usable Space (Percent):	98.982 %
Total Free Space:	19.1 GB
Total Free Space (Percent):	98.982 %
Total Persistent Data:	14 MB
Total Persistent Data (Percent):	0.072 %
Total Cached Data:	0 B
Block Reads:	6,737
Block Writes:	3,411
Objects Retrieved:	3,381
Objects Committed:	1,712
Objects Purged:	0
Purge Service State:	Enabled

Object Stores

ID	Total	Available	Stored Data	Stored (%)	Cached Data	Cached (%)	Health
0000	19.3 GB	19.1 GB	14 MB	0.072 %	0 B	0 %	No Errors

Total usable storage space available on this Storage Node (STAS)

Figure 57 Individual LDR Storage Capacity

On the LDR > Storage > Overview page, monitor the Total Usable Space (STAS) attribute over a period of time to estimate the rate at which usable object storage space is being consumed. Usable space is the actual real amount of storage space available to store objects. Thus, the attribute Total Usable Space (STAS) displays the amount of usable storage space available on the object store with the least amount of usable space remaining (the fullest object store) multiplied by the number of object stores on the Storage Node. This number may differ and be less than the Total Free Space (SUSA).

To maintain normal grid operations, you have to add Storage Nodes, or add storage volumes, or migrate content to archive media before the storage disks' usable space fills up.

In the example shown in Figure 58, usable content storage space is being consumed at a rate of approximately 4% per month, which means that there are 8 months left before this LDR runs out of storage space.

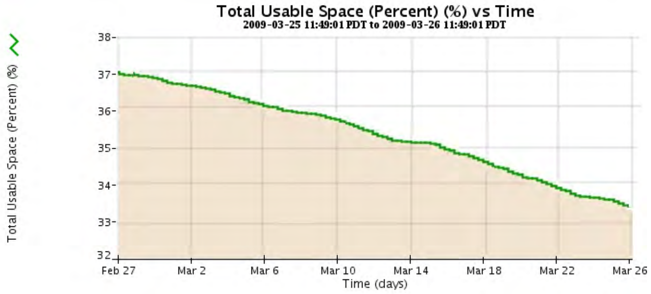

Figure 58 Usable Contents Storage Space

Metadata Storage Capacity

The CMS databases on the Control Nodes are responsible for storing metadata, e.g., the information about the objects ingested into the grid. You need to monitor the total space available on Control Nodes to make sure the grid does not run out of space to store metadata. If all CMS databases in the grid fill up, the grid can no longer ingest files until additional metadata capacity is added to the grid.

Metadata storage capacity information is available at the grid level (see Figure 59 for an example), site level, and node level (see Figure 60 for an example).

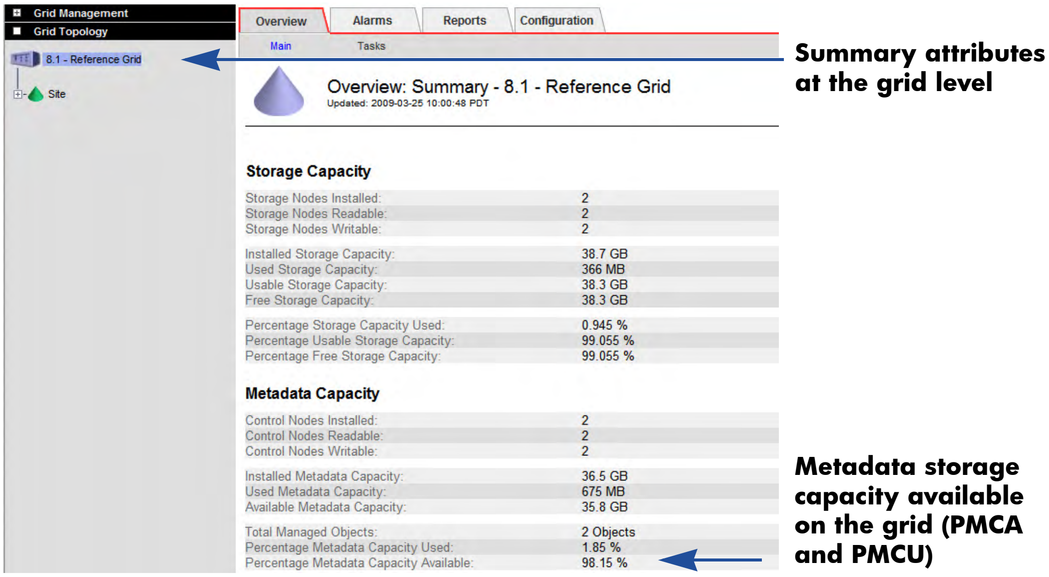

Overview: Summary - 8.1 - Reference Grid
Updated: 2009-03-25 10:00:48 PDT

Storage Capacity

Storage Nodes Installed:	2
Storage Nodes Readable:	2
Storage Nodes Writable:	2
Installed Storage Capacity:	38.7 GB
Used Storage Capacity:	366 MB
Usable Storage Capacity:	38.3 GB
Free Storage Capacity:	38.3 GB
Percentage Storage Capacity Used:	0.945 %
Percentage Usable Storage Capacity:	99.055 %
Percentage Free Storage Capacity:	99.055 %

Metadata Capacity

Control Nodes Installed:	2
Control Nodes Readable:	2
Control Nodes Writable:	2
Installed Metadata Capacity:	36.5 GB
Used Metadata Capacity:	675 MB
Available Metadata Capacity:	35.8 GB
Total Managed Objects:	2 Objects
Percentage Metadata Capacity Used:	1.85 %
Percentage Metadata Capacity Available:	98.15 %

Summary attributes at the grid level

Metadata storage capacity available on the grid (PMCA and PMCU)

Figure 59 Summary CMS Metadata Storage Capacity

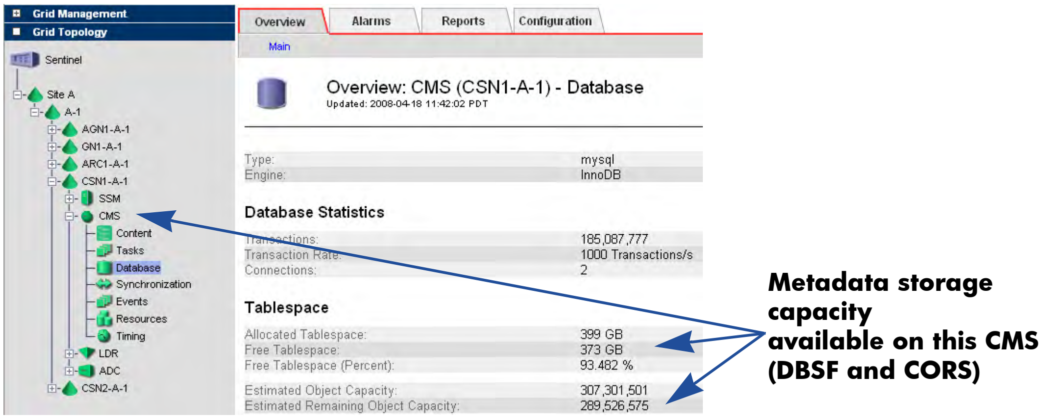

Overview: CMS (CSN1-A-1) - Database
Updated: 2008-04-18 11:42:02 PDT

Type:	mysql
Engine:	InnoDB

Database Statistics

Transactions:	185,087,777
Transaction Rate:	1000 Transactions/s
Connections:	2

Tablespace

Allocated Tablespace:	399 GB
Free Tablespace:	373 GB
Free Tablespace (Percent):	93.482 %
Estimated Object Capacity:	307,301,501
Estimated Remaining Object Capacity:	289,526,575

Metadata storage capacity available on this CMS (DBSF and CORS)

Figure 60 Individual CMS Metadata Storage Capacity

Track CMS Database Free Tablespace and CMS Database Estimated Remaining Object Capacity over a period of time to estimate the rate at which the available database space is being consumed. The databases of CMSs that are in the same CMS replication group fill at approximately the same time. In a grid that uses metadata synchronization, the databases of CMSs in the same generation fill at approximately the same time.

The CMS databases go into “read-only” mode when Free Tablespace drops below 10%. To maintain normal grid operations, you have to add Control Nodes before the metadata database fills up.

Figure 61 is for the same grid as the one used in Figure 58 (page 109). The migration effect is not noticeable in Figure 58 because the object size is small: the metadata of a large quantity of small objects use more space proportionally than the actual content.

For the example shown in Figure 61, metadata storage capacity is being consumed at a rate of approximately 20 GB per month. However, note how the utilization rate increases towards the end of the period. This could be due, for example, to a data migration that is happening in parallel with regular grid ingest.

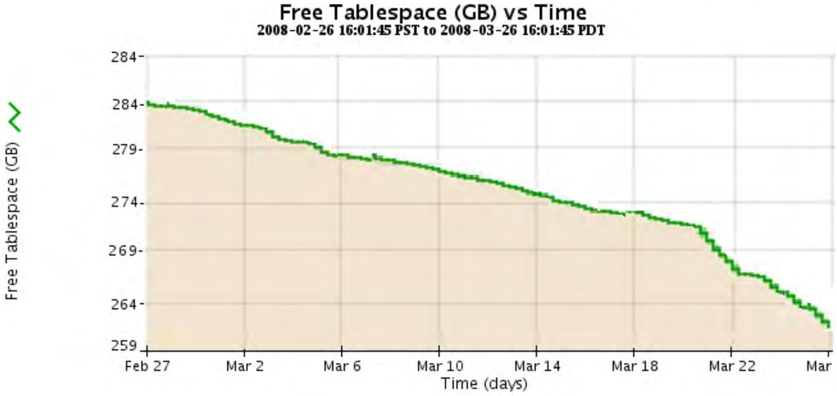

Figure 61 Available CMS Content Storage Metadata Capacity

Gateway Node Capacity and Load

Summary Attributes

Gateway Node information is also available at the grid and site level on the Overview tab (see Figure 62).

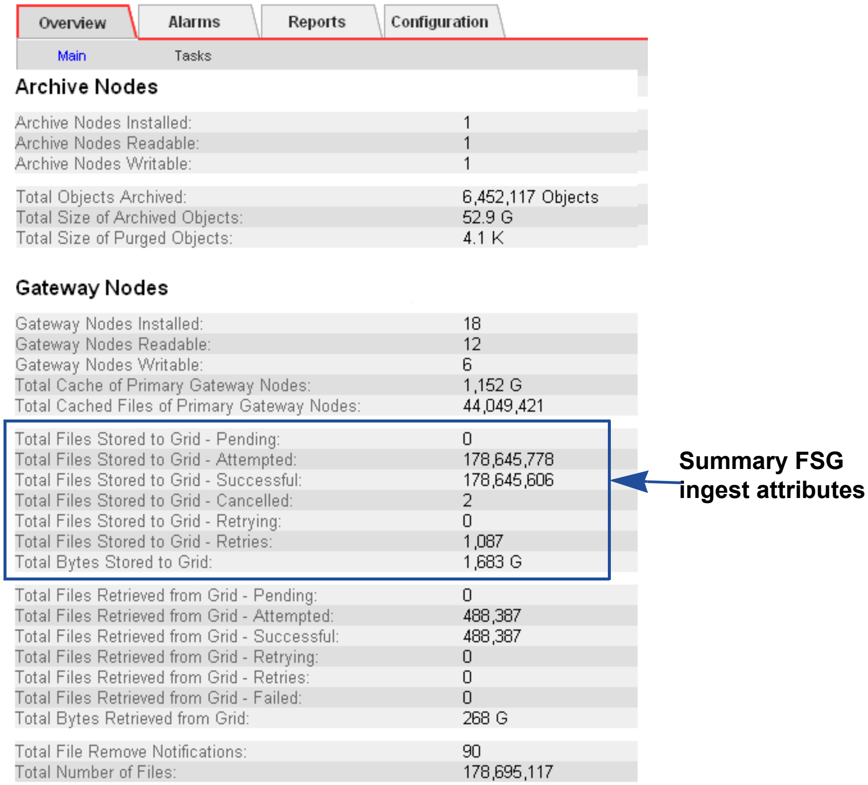

Archive Nodes

Archive Nodes Installed:	1
Archive Nodes Readable:	1
Archive Nodes Writable:	1
Total Objects Archived:	6,452,117 Objects
Total Size of Archived Objects:	52.9 G
Total Size of Purged Objects:	4.1 K

Gateway Nodes

Gateway Nodes Installed:	18
Gateway Nodes Readable:	12
Gateway Nodes Writable:	6
Total Cache of Primary Gateway Nodes:	1,152 G
Total Cached Files of Primary Gateway Nodes:	44,049,421
Total Files Stored to Grid - Pending:	0
Total Files Stored to Grid - Attempted:	178,645,778
Total Files Stored to Grid - Successful:	178,645,606
Total Files Stored to Grid - Cancelled:	2
Total Files Stored to Grid - Retrying:	0
Total Files Stored to Grid - Retries:	1,087
Total Bytes Stored to Grid:	1,683 G
Total Files Retrieved from Grid - Pending:	0
Total Files Retrieved from Grid - Attempted:	488,387
Total Files Retrieved from Grid - Successful:	488,387
Total Files Retrieved from Grid - Retrying:	0
Total Files Retrieved from Grid - Retries:	0
Total Files Retrieved from Grid - Failed:	0
Total Bytes Retrieved from Grid:	268 G
Total File Remove Notifications:	90
Total Number of Files:	178,695,117

Summary FSG ingest attributes

Figure 62 Gateway Node Summary Attributes

FSG Capacity

You can get an estimate of how much content is managed by each FSG replication group by looking at the FSG backup information: FSG Backup Number of Objects (PBNO), Number of Files (PBNF) and FSG Backup Data Size (PBDS) (see Figure 63). The Backup Number of Objects should increase steadily.

The FSG > Backup component displays information on a per FSG basis while the FSG > Shares component displays FSG usage on a per share basis. For more information on the FSG > Shares component, see File Share Usage (page 114). To determine the overall amount of capacity remaining on the FSG, use the FSG > Backup component

The backup values reflect the FSG managed file system as of the most recent backup. This includes files that are pending for ingest and files for which ingest into the grid is disabled through FSG profiles.

Since FSG capacity limits are not actively enforced, proactive monitoring is necessary to identify when an FSG has reached its capacity. At that point, client ingests should be directed to a new FSG replication group to avoid problems that may occur by exceeding the supported capacity.

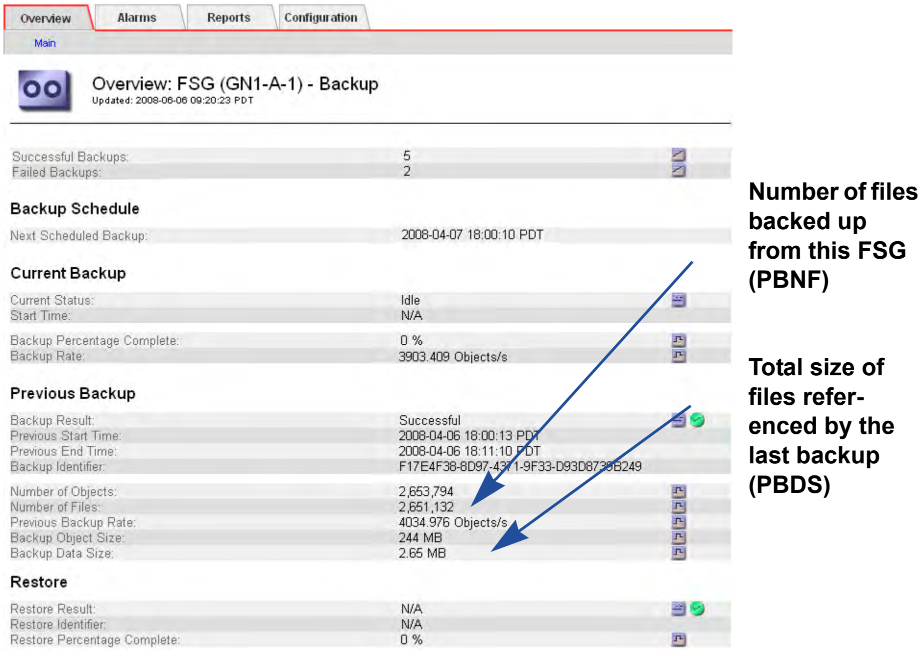

Overview: FSG (GN1-A-1) - Backup
Updated: 2008-06-06 09:20:23 PDT

Successful Backups:	5
Failed Backups:	2

Backup Schedule

Next Scheduled Backup:	2008-04-07 18:00:10 PDT

Current Backup

Current Status:	Idle
Start Time:	N/A
Backup Percentage Complete:	0 %
Backup Rate:	3903.409 Objects/s

Previous Backup

Backup Result:	Successful
Previous Start Time:	2008-04-06 18:00:13 PDT
Previous End Time:	2008-04-06 18:11:10 PDT
Backup Identifier:	F17E4F38-8D97-4371-9F33-D93D8738B249
Number of Objects:	2,653,794
Number of Files:	2,651,132
Previous Backup Rate:	4034.976 Objects/s
Backup Object Size:	244 MB
Backup Data Size:	2.65 MB

Restore

Restore Result:	N/A
Restore Identifier:	N/A
Restore Percentage Complete:	0 %

Number of files backed up from this FSG (PBNF)

Total size of files referenced by the last backup (PBDS)

Figure 63 FSG Backup Size

When an FSG replication group reaches capacity, the grid must be expanded by adding a new FSG replication group. To estimate the number of objects per FSG replication group and remaining capacity (and thus estimate when you must expand the grid), monitor the inodes: the number and percentage available.

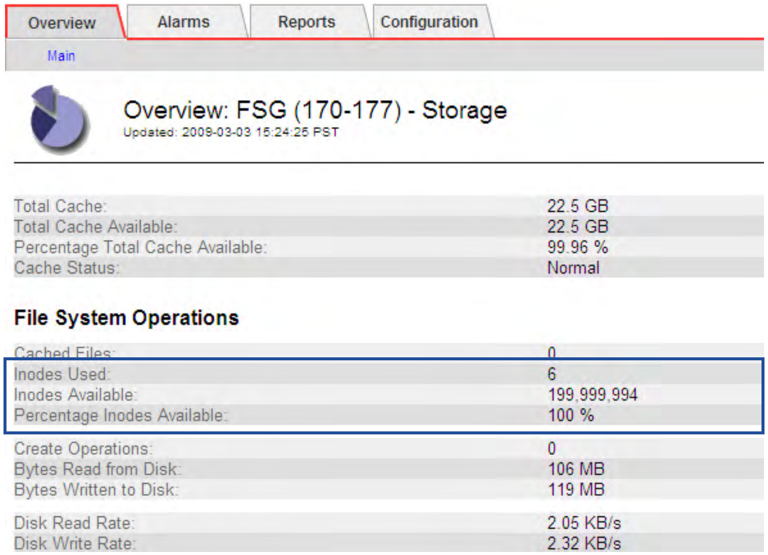

Overview: FSG (170-177) - Storage
Updated: 2008-03-03 15:24:25 PST

Total Cache:	22.5 GB
Total Cache Available:	22.5 GB
Percentage Total Cache Available:	99.96 %
Cache Status:	Normal

File System Operations

Cached Files:	0
Inodes Used:	6
Inodes Available:	199,999,994
Percentage Inodes Available:	100 %
Create Operations:	0
Bytes Read from Disk:	106 MB
Bytes Written to Disk:	119 MB
Disk Read Rate:	2.05 KB/s
Disk Write Rate:	2.32 KB/s

Figure 64 Inodes Used and Availability

When a minor alarm is triggered for the Percentage Inodes Available (FSIP) attribute (by default when remaining capacity is 20%), calculate how long before the remaining 20% is used. This assists in determining when to add a new replication group.

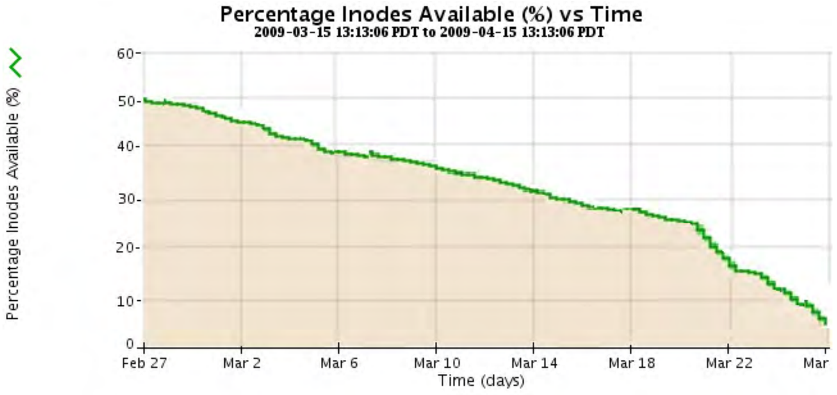

Figure 65 Percentage Inodes Available

File Share Usage

FSG file share usage is available to monitor the amount of data being saved to each FSG share. This information is useful in determining whether a client is saving too much data (or too little) to a share. If multiple clients are mapped to the same share, this information can help determine if clients should be remapped to other shares. Note that share usage values are only updated at the time of the FSG backup. By default FSG backup is once per day.

Quotas can be set to monitor share usage. A client may have a limit to the amount of data it can store to a share. An alarm is raised if the quota is exceeded; however, exceeding the quota does not prevent clients from continuing to save to the share.

Important values to monitor are:

- **Used (FSSB)**—The total amount of storage space used by the share. If this value exceeds the configured Quota value, an alarm is raised. The administrator can then either increase the quota for the share or remove files from the share.
- **Remaining (FSSA)**—The total amount of storage space remaining for the share in its quota.
- **Used (Week) (FSSR)**—The amount of storage space used by the share in the past seven days (as of the last FSG backup). If this value exceeds the configured Quota (Weekly) value, an alarm is raised. The administrator can then either increase the quota for the share or remove files from the share.
- **Remaining (Week) (FSSL)**—The amount of space remaining in the share's weekly storage quota.

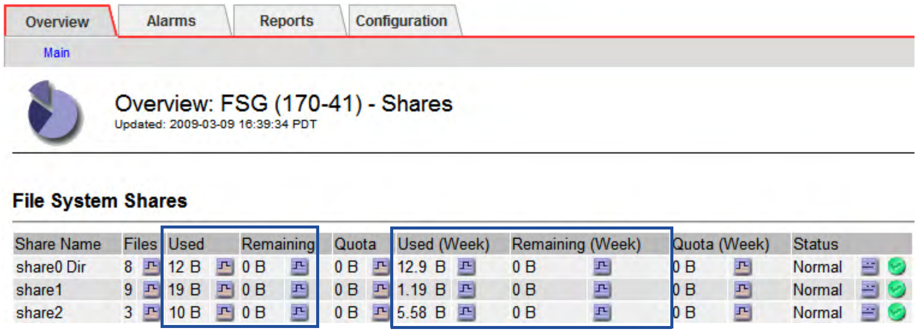

Overview: FSG (170-41) - Shares

Updated: 2009-03-09 16:39:34 PDT

File System Shares

Share Name	Files	Used	Remaining	Quota	Used (Week)	Remaining (Week)	Quota (Week)	Status
share0 Dir	8	12 B	0 B	0 B	12.9 B	0 B	0 B	Normal
share1	9	19 B	0 B	0 B	1.19 B	0 B	0 B	Normal
share2	3	10 B	0 B	0 B	5.58 B	0 B	0 B	Normal

Figure 66 File Share Usage

Ingest Load

To monitor the ingest load on the grid, analyze the trends of these four attributes over time (see Figure 67):

- **Bytes Stored to Grid (FSGB)**—The number of bytes ingested successfully into the grid.
- **File Store Rate (FSRA)**—The rate at which files are successfully stored to the grid (number of transactions per second).
- **Data Store Rate (FSBA)**—The rate at which data is successfully stored to the grid (in bytes per second)
- **File Store Latency (FSTM)**—The average amount of time required to store the entire file into the grid. The average is calculated over the last sampling period.

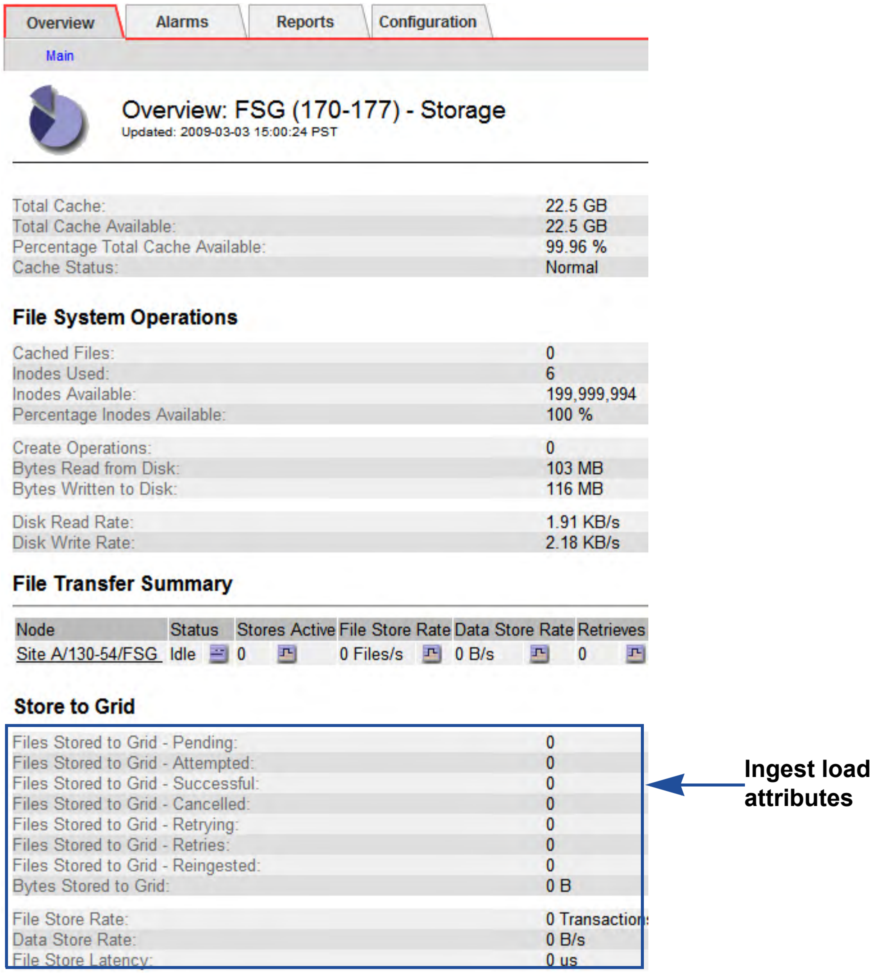

Overview: FSG (170-177) - Storage
Updated: 2009-03-03 15:00:24 PST

Total Cache:	22.5 GB
Total Cache Available:	22.5 GB
Percentage Total Cache Available:	99.96 %
Cache Status:	Normal

File System Operations

Cached Files:	0
Inodes Used:	6
Inodes Available:	199,999,994
Percentage Inodes Available:	100 %
Create Operations:	0
Bytes Read from Disk:	103 MB
Bytes Written to Disk:	116 MB
Disk Read Rate:	1.91 KB/s
Disk Write Rate:	2.18 KB/s

File Transfer Summary

Node	Status	Stores Active	File Store Rate	Data Store Rate	Retrieves
Site A/130-54/FSG	Idle	0	0 Files/s	0 B/s	0

Store to Grid

Files Stored to Grid - Pending:	0
Files Stored to Grid - Attempted:	0
Files Stored to Grid - Successful:	0
Files Stored to Grid - Cancelled:	0
Files Stored to Grid - Retrying:	0
Files Stored to Grid - Retries:	0
Files Stored to Grid - Reingested:	0
Bytes Stored to Grid:	0 B
File Store Rate:	0 Transactio
Data Store Rate:	0 B/s
File Store Latency:	0 us

Ingest load attributes

Figure 67 Ingest Load Attributes

Increasing Ingest Load

During normal operations, it is possible for the ingest load to exceed the rate at which the services on the grid process the objects. When this happens, services may queue operations that can no longer be fulfilled in real time. For instance, the value of File Stored to Grid - Pending (that is, the number of ingested files cached locally that are waiting for transfer to the grid for persistent storage) may increase temporarily.

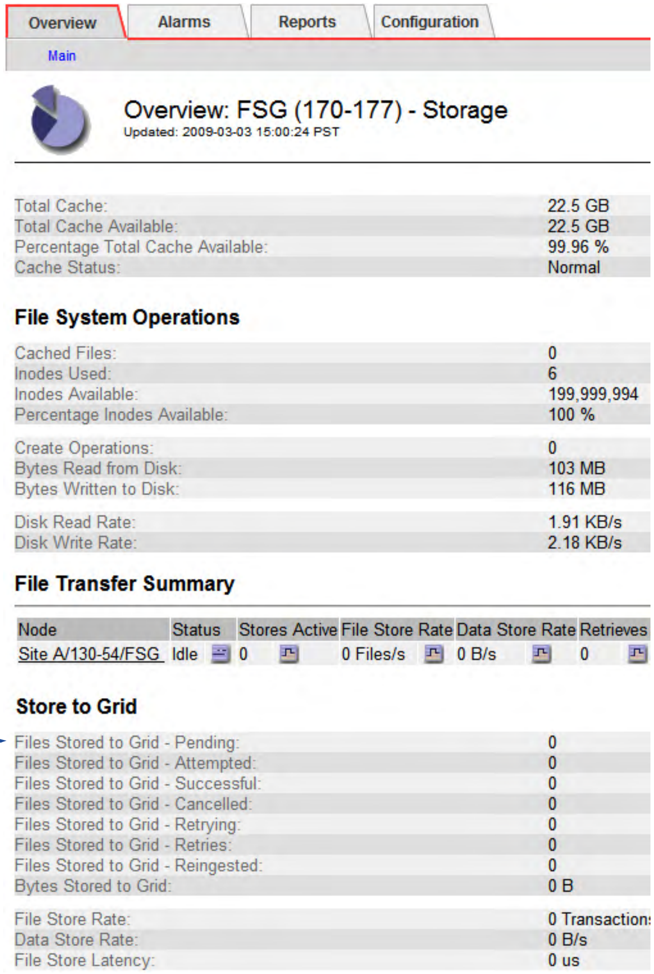

Overview: FSG (170-177) - Storage
Updated: 2009-03-03 15:00:24 PST

Total Cache:	22.5 GB
Total Cache Available:	22.5 GB
Percentage Total Cache Available:	99.96 %
Cache Status:	Normal

File System Operations

Cached Files:	0
Inodes Used:	6
Inodes Available:	199,999,994
Percentage Inodes Available:	100 %
Create Operations:	0
Bytes Read from Disk:	103 MB
Bytes Written to Disk:	116 MB
Disk Read Rate:	1.91 KB/s
Disk Write Rate:	2.18 KB/s

File Transfer Summary

Node	Status	Stores Active	File Store Rate	Data Store Rate	Retrieves
Site A/130-54/FSG	Idle	0	0 Files/s	0 B/s	0

Store to Grid

Files Stored to Grid - Pending:	0
Files Stored to Grid - Attempted:	0
Files Stored to Grid - Successful:	0
Files Stored to Grid - Cancelled:	0
Files Stored to Grid - Retrying:	0
Files Stored to Grid - Retries:	0
Files Stored to Grid - Reingested:	0
Bytes Stored to Grid:	0 B
File Store Rate:	0 Transaction
Data Store Rate:	0 B/s
File Store Latency:	0 us

Number of files saved to the FSG by client application and waiting to be stored to the grid (FSGP)

Figure 68 Files Stored to Grid Pending

Files Stored to Grid—Pending includes files actively being written but for which the ingest delay period has not expired. Therefore, Files Stored to Grid - Pending will always be non-zero when ingests are active.

In the example shown in Figure 69, the number of files waiting to be stored goes up and down, but remains fairly low. Such a trend could indicate that there was a short term overload due to network throughput, disk I/O performance, grid services availability, and so on.

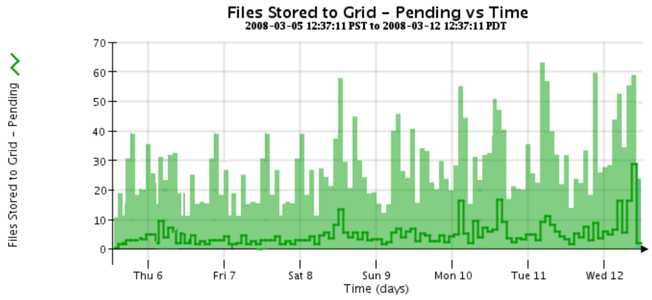

Figure 69 Files Stored to Grid Pending

In contrast, the trend shown in Figure 70 is not sustainable. If the number of files waiting to be stored to the grid starts to increase, make sure that all CMS and LDR services are operating normally. It is also possible that the ingest rate is exceeding the throughput of the grid and that a grid expansion is required.

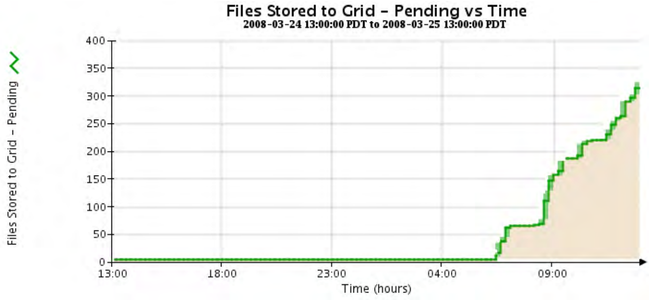

Figure 70 Non Sustainable Ingest Load

Total Cache Available may also be monitored: if new files are ingested faster than existing cached files can be swapped out, the amount of cache available may dip below the "Swapout No Create Watermark" (defined in FSG Management). When this happens, the creation of new files is temporarily disallowed until enough space is freed. This is uncommon, but may occur; for example, when the grid has ingested many small files followed by many large files. In this case, the FSG may not be able to swap out the small files fast enough to make room for the large files. If this happens, the client must throttle its ingest rate.

Retrieve Load

To monitor the retrieve load on the grid, analyze the trends of these attributes over time (see Figure 71).

- **Bytes Retrieved from Grid (FRGB)**—The number of bytes retrieved successfully from the grid.
- **File Retrieve Rate (FRRA)**—The rate at which files are successfully retrieved from the grid (number of transactions per second).
- **Data Retrieve Rate (FRBA)**—The rate at which data is successfully retrieved from the grid (in bytes per second).
- **File Retrieve Latency (FRTM)**—The average amount of time required to retrieve the entire file from the grid. The average is calculated over the last sampling period.

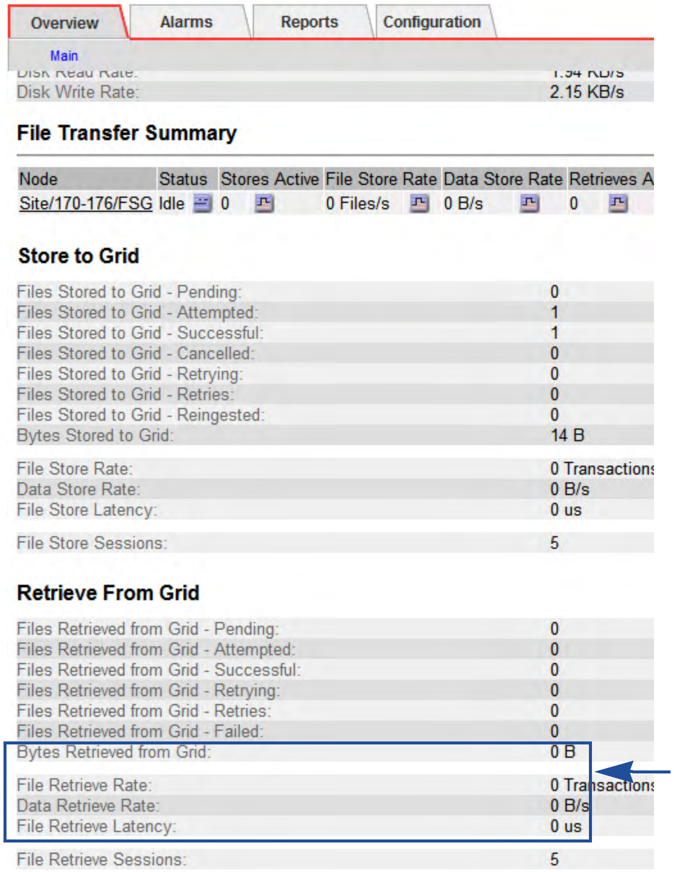

Disk Read Rate:	1.94 KB/s
Disk Write Rate:	2.15 KB/s

File Transfer Summary

Node	Status	Stores Active	File Store Rate	Data Store Rate	Retrieves A
Site/170-176/FSG	Idle	0	0 Files/s	0 B/s	0

Store to Grid

Files Stored to Grid - Pending:	0
Files Stored to Grid - Attempted:	1
Files Stored to Grid - Successful:	1
Files Stored to Grid - Cancelled:	0
Files Stored to Grid - Retrying:	0
Files Stored to Grid - Retries:	0
Files Stored to Grid - Reingested:	0
Bytes Stored to Grid:	14 B
File Store Rate:	0 Transactions
Data Store Rate:	0 B/s
File Store Latency:	0 us
File Store Sessions:	5

Retrieve From Grid

Files Retrieved from Grid - Pending:	0
Files Retrieved from Grid - Attempted:	0
Files Retrieved from Grid - Successful:	0
Files Retrieved from Grid - Retrying:	0
Files Retrieved from Grid - Retries:	0
Files Retrieved from Grid - Failed:	0
Bytes Retrieved from Grid:	0 B
File Retrieve Rate:	0 Transactions
Data Retrieve Rate:	0 B/s
File Retrieve Latency:	0 us
File Retrieve Sessions:	5

Retrieve load attributes (FRGB, FRRA, FRBA, FRTM)

Figure 71 Retrieve Load Attributes

ILM Replication

A grid’s CMS services manage ILM replication. ILM replication refers to the process of making copies of the object and keeping the copies in the appropriate storage locations for a pre-determined length of time. You can track what is happening with ILM replication by looking at the number of objects in each of these categories:

Pending	Objects with ILM Evaluation Pending (ORpe) The number of objects waiting to be processed.
Unachievable	Objects with Unachievable ILM Evaluations (ORun) The number of objects that require additional copies to be made but the storage resources specified by the ILM policy are unavailable to store the additional copies.
Future	Objects Marked for ILM Re-evaluation (ORde) The number of objects that currently satisfy the ILM policy but are due to be re-evaluated at a scheduled point in the future, for instance because of a rule that says “store a copy to archive media two years after ingest”.

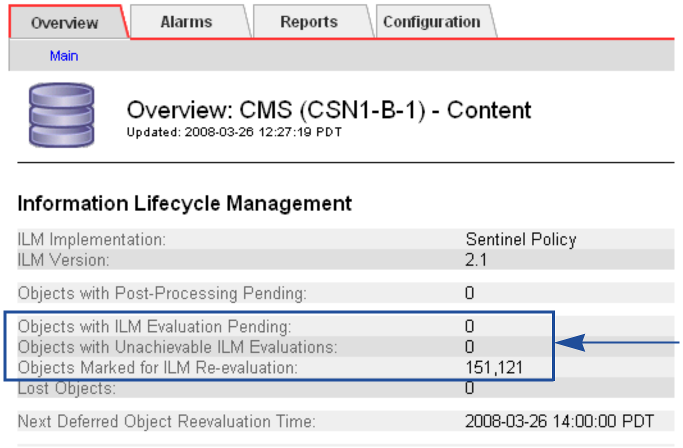

Overview: CMS (CSN1-B-1) - Content

Updated: 2008-03-26 12:27:19 PDT

Information Lifecycle Management

ILM Implementation:	Sentinel Policy
ILM Version:	2.1
Objects with Post-Processing Pending:	0
Objects with ILM Evaluation Pending:	0
Objects with Unachievable ILM Evaluations:	0
Objects Marked for ILM Re-evaluation:	151,121
Lost Objects:	0
Next Deferred Object Reevaluation Time:	2008-03-26 14:00:00 PDT

Number of objects in each category: pending, unachievable, and future (ORun, ORpe, ORde)

Figure 72 ILM Evaluation

Metadata Replication and Synchronization

Metadata management depends on whether the CMS service uses metadata replication or metadata synchronization.

Metadata Replication by CMSs

The Metadata component is only displayed for CMSs that use metadata replication.

Information about metadata replication is shown in the Metadata component of the CMS service (see Figure 73). The number of inter-CMS messages that are needed to manage metadata is drastically reduced when a grid uses metadata replication as compared to metadata synchronization (see Metadata Synchronization by CMSs (page 122)).

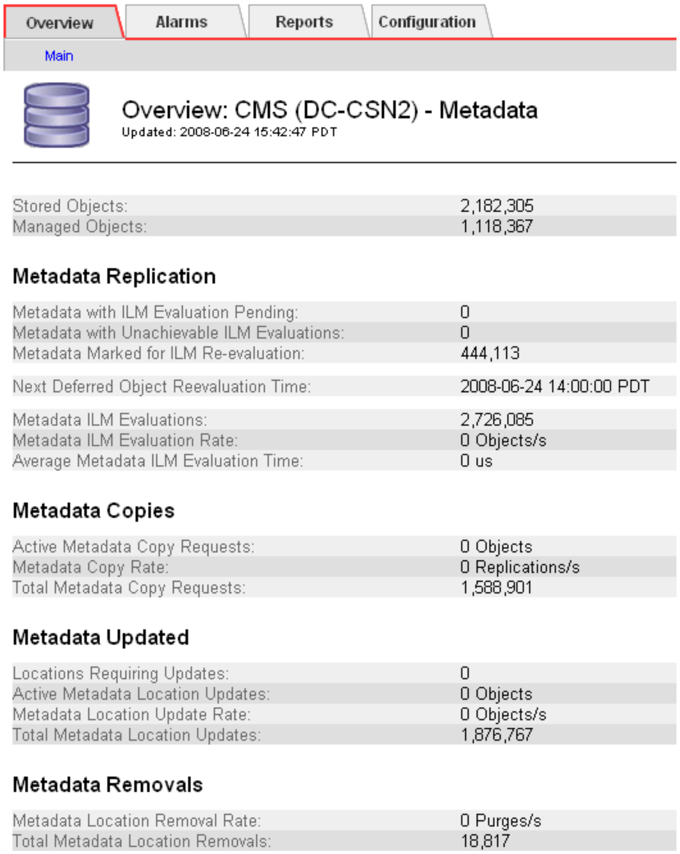

Overview: CMS (DC-CSN2) - Metadata

Updated: 2008-06-24 15:42:47 PDT

Stored Objects:	2,182,305
Managed Objects:	1,118,367

Metadata Replication

Metadata with ILM Evaluation Pending:	0
Metadata with Unachievable ILM Evaluations:	0
Metadata Marked for ILM Re-evaluation:	444,113
Next Deferred Object Reevaluation Time:	2008-06-24 14:00:00 PDT
Metadata ILM Evaluations:	2,726,085
Metadata ILM Evaluation Rate:	0 Objects/s
Average Metadata ILM Evaluation Time:	0 us

Metadata Copies

Active Metadata Copy Requests:	0 Objects
Metadata Copy Rate:	0 Replications/s
Total Metadata Copy Requests:	1,588,901

Metadata Updated

Locations Requiring Updates:	0
Active Metadata Location Updates:	0 Objects
Metadata Location Update Rate:	0 Objects/s
Total Metadata Location Updates:	1,876,767

Metadata Removals

Metadata Location Removal Rate:	0 Purges/s
Total Metadata Location Removals:	18,817

Figure 73 CMS Metadata Component for Metadata Replication

Metadata Synchronization by CMSs

When an owner CMS service receives new metadata, it stores the metadata in its local database and sends synchronization messages to the other CMSs it must replicate metadata to. The attribute Queue Size tracks the number of messages to be sent to another CMS service. The corresponding attribute Incoming Messages tracks the number of synchronization messages coming from another CMS service and waiting to be processed (Figure 74).

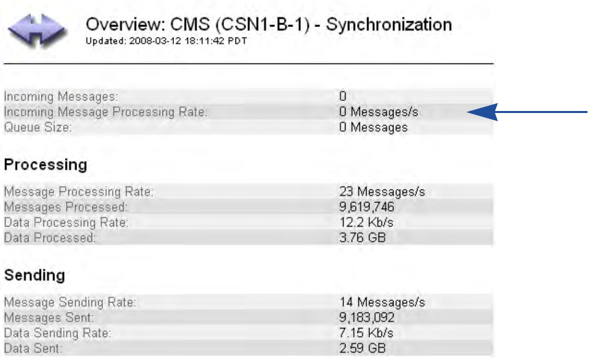

Overview: CMS (CSN1-B-1) - Synchronization
Updated: 2008-03-12 18:11:42 PDT

Incoming Messages:	0
Incoming Message Processing Rate:	0 Messages/s
Queue Size:	0 Messages

Processing

Message Processing Rate:	23 Messages/s
Messages Processed:	9,619,746
Data Processing Rate:	12.2 Kb/s
Data Processed:	3.76 GB

Sending

Message Sending Rate:	14 Messages/s
Messages Sent:	9,183,092
Data Sending Rate:	7.15 Kb/s
Data Sent:	2.59 GB

Number of messages queued for synchronization (CsQT)

Figure 74 CMS Metadata Synchronization

During normal operations, it is possible for the ingest load to exceed the rate at which the CMS services can synchronize metadata. This temporary overload solution will resolve itself over time. However, if synchronization messages start to accumulate, ensure that the other CMS services are running normally and if the trend continues, escalate the issue as it could be that the ingest rate is exceeding the throughput of the grid.

FSG Replication

When an FSG ingests a file, it creates a file pointer to reference the object and it replicates the file pointer to the other FSGs in its replication group. To verify that FSG replication is proceeding normally, look at the attributes Operations Not Committed (SUOP), Operations Not Applied (SPOP), Files Pending for Replication (FRPP), and Replication Errors (RPER) in the FSG Replication component (see Figure 75 and Figure 76).

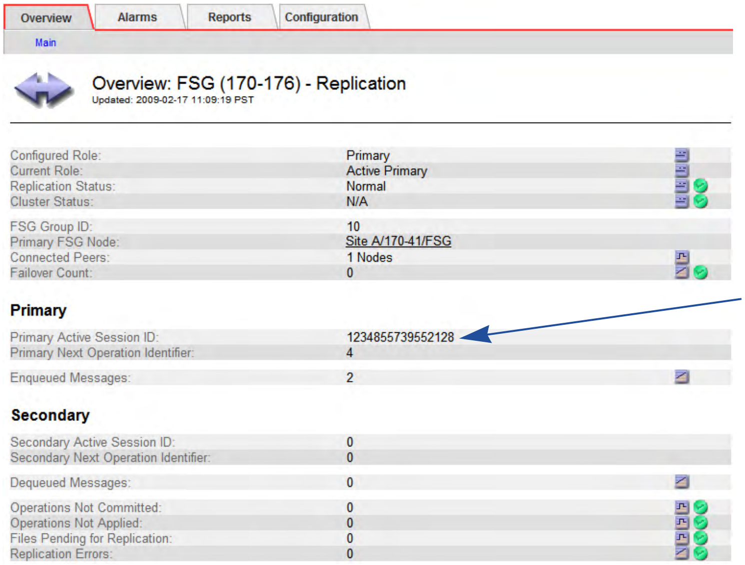

Overview: FSG (170-176) - Replication
Updated: 2009-02-17 11:09:19 PST

Configured Role:	Primary
Current Role:	Active Primary
Replication Status:	Normal
Cluster Status:	N/A
FSG Group ID:	10
Primary FSG Node:	Site A/170-41/FSG
Connected Peers:	1 Nodes
Failover Count:	0

Primary

Primary Active Session ID:	1234855739552128
Primary Next Operation Identifier:	4
Enqueued Messages:	2

Secondary

Secondary Active Session ID:	0
Secondary Next Operation Identifier:	0
Dequeued Messages:	0
Operations Not Committed:	0
Operations Not Applied:	0
Files Pending for Replication:	0
Replication Errors:	0

Active Session on Primary FSG

Figure 75 Primary FSG Replication Messages

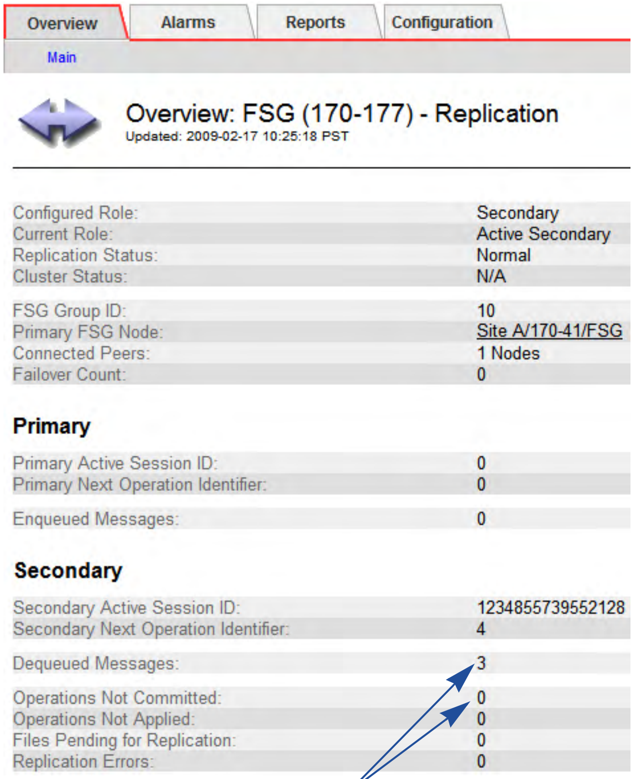

Overview: FSG (170-177) - Replication
Updated: 2009-02-17 10:25:18 PST

Configured Role:	Secondary
Current Role:	Active Secondary
Replication Status:	Normal
Cluster Status:	N/A
FSG Group ID:	10
Primary FSG Node:	Site A/170-41/FSG
Connected Peers:	1 Nodes
Failover Count:	0

Primary

Primary Active Session ID:	0
Primary Next Operation Identifier:	0
Enqueued Messages:	0

Secondary

Secondary Active Session ID:	1234855739552128
Secondary Next Operation Identifier:	4
Dequeued Messages:	3
Operations Not Committed:	0
Operations Not Applied:	0
Files Pending for Replication:	0
Replication Errors:	0

Corresponding active session on Secondary FSG with queue and backlog of replication messages (SUOP and SPOP)

Figure 76 Secondary FSG Replication Messages

- **Operations Not Committed (SUOP)**—The number of replication messages from the primary FSG session that have not been written to the secondary FSG yet.

An upwards trend indicates that the grid is ingesting files faster than the secondary FSG can process transactions. The number will go back down during periods of reduced grid activity. If it does not go down, escalate the issue as it could be that the ingest rate is exceeding the throughput of the grid.

- **Operations Not Applied (SPOP)**—The number of replication messages to be processed on the secondary FSG in order to catch up to the primary FSG.

  An upwards trend indicates that the primary FSG is ingesting files faster than the secondary FSG can process transactions. The number of messages to be processed on the secondary FSG in order to catch up to the primary FSG may temporarily increase during periods of high grid activity or backups. (In older systems where "offline backups" were used, SPOP increased during backups. As of Release 8.1, offline backups are deprecated.)

  The number will go back down during periods of reduced grid activity. If it does not go down, escalate the issue. This could be an indication that there is an FSG "synchronization" problem or that the ingest rate is exceeding the throughput of the grid.

  Note that a temporary increase will also occur on an FSG that is being restored (as it processes the replication backlog).

  Contact HP Support, who can check if FSGs have been converted to online backups.

- **Files Pending for Replication (FRPP)**—The number of new files that have not been fully replicated because they are awaiting a replication message from the primary FSG indicating that the file has been successfully stored to the grid.

  If an upwards trend persists or value does not decrease and all FSG services are operating correctly, and the Operations Not Applied attribute on the secondary FSG service and Files Stored to Grid - Pending attribute on the primary FSG service are zero, replication inconsistencies may exist. If the value does not go down escalate the issue.

- **Replication Errors (RPER)**—The number of replication messages that the secondary FSG cannot apply to its file system because of a conflict. For example, if the secondary FSG is asked to create a reference to a file in a directory that it has no record of, this is recorded as a replication error.

  When the FSG encounters a replication error, it increments RPER and issues a log message with additional information about the replication error. This log message is picked up by the SSM events monitor to provide some visibility in the NMS MI of files that have been affected by replication errors. If the value does not go down or persists escalate the issue.

Gateway Node Failovers

If the primary Gateway Node fails, the RSTU FSG Replication Status alarm is triggered, displaying a status of No Primary or No Session. What happens next depends on the type of replication group.

Basic Replication Group

If the primary FSG in a Basic Gateway replication group fails, immediately notify a grid administrator who has access to the Admin or Vendor account.

A manual failover procedure can be performed if the grid supports business continuity failover. The secondary FSG can be manually configured to act as a primary. After clients are manually redirected to the acting primary, they can continue to read and write to the grid. This is a temporary measure to maintain service while the primary FSG is repaired: grid access is interrupted until manual failover is completed, and the redundancy of file system information in the grid is reduced while the secondary FSG is an acting primary FSG.

In grids that do not support business continuity failover, clients can continue to access files via the read-only file system on the secondary FSG while the primary is repaired, but they cannot write to the grid.

Situations may arise where the replication sessions on the FSG services within a replication group may become inconsistent with each other. Monitor FSGs to determine if replication sessions within a replication group become inconsistent with each other. One indication of this is if the active secondary Gateway Node's Secondary Active Session ID is zero or non-changing.

Following a failure, a manual recovery procedure must be performed to replicate files at risk, even if the failed FSG recovers automatically.

Note that an alarm is not raised if replication sessions become inconsistent with each other; however, if a situation triggers a No Session alarm on RSTU Replication Status, it may also affect replication sessions.

High Availability Gateway Replication Group

If the active primary FSG in a High Availability Gateway Cluster (HAGC) fails, the standby primary FSG becomes the active primary FSG without any manual intervention. The cluster status changes to Vulnerable and the alarm FCST Cluster Status is triggered.

When a failover occurs, any client operations that are in progress fail, as do client operations initiated while the standby primary FSG makes the transition to active. After the CIFS service starts on the new active primary FSG, Windows CIFS clients should be able to process new operations without remapping their connections to the primary FSG cluster. NFS clients may have to remount shares before they can continue to store and retrieve data to the Gateway cluster (depending on the NFS client). Full grid functionality and full grid access is maintained while the second FSG is active.

Investigate the cause of the failure as soon as possible as another FSG failure in the replication group will render FSG services unavailable. To restore the grid to full redundancy after the failover, fail back to the main primary manually and identify any files at risk of being lost. Contact HP Support for assistance with failover and recovery procedures.

FSG Backups

Within an FSG replication group, one FSG is designated as the backup FSG. This FSG backs up the replication group managed file system (the file pointer references) daily. The backup files are ingested into the grid and by default are automatically deleted after 14 days.

The greater the number of objects, the longer it takes to back up the FSG. The backup must complete each day.

You can view information about the FSG backups (for example, backup schedule, backup duration, backup size, backup status) on the FSG Backup component of the Gateway Node (see Figure 77).

The value for Number of Files is the total number of files on the FSG managed file system as of the last backup. This includes files that are pending for ingest and files for which ingest into the grid is disabled through FSG profiles. For content information on a per share basis, see FSG > Shares. Per share data is updated at the time of the FSG backup. For more information, see File Share Usage (page 114).

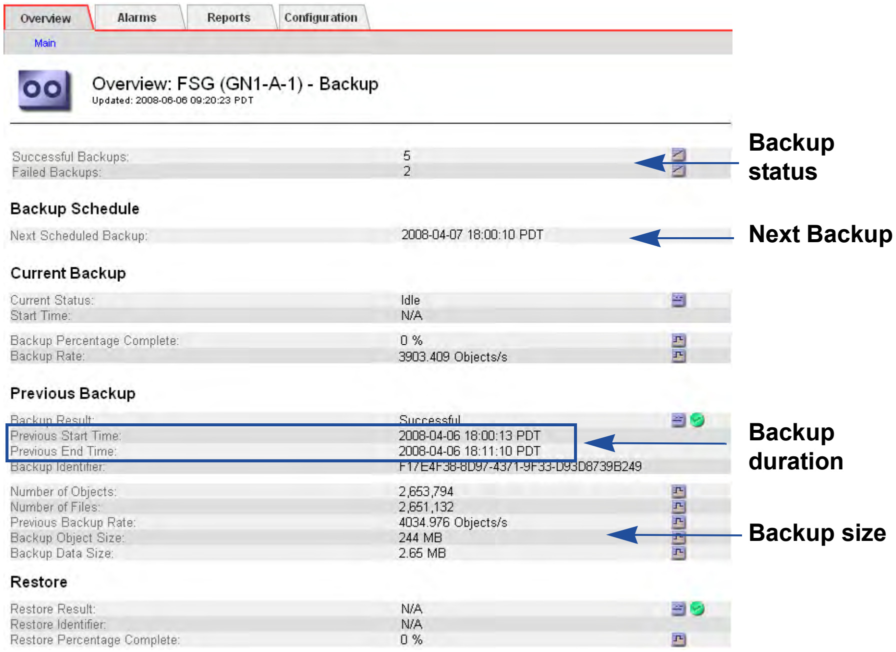

Overview: FSG (GN1-A-1) - Backup
Updated: 2008-06-06 09:20:23 PDT

Successful Backups:	5
Failed Backups:	2

Backup Schedule

Next Scheduled Backup:	2008-04-07 18:00:10 PDT

Current Backup

Current Status:	Idle
Start Time:	N/A
Backup Percentage Complete:	0 %
Backup Rate:	3903.409 Objects/s

Previous Backup

Backup Result:	Successful
Previous Start Time:	2008-04-06 18:00:13 PDT
Previous End Time:	2008-04-06 18:11:10 PDT
Backup Identifier:	F17E4F38-8D97-4371-9F33-D93D8739B249
Number of Objects:	2,653,794
Number of Files:	2,651,132
Previous Backup Rate:	4034.976 Objects/s
Backup Object Size:	244 MB
Backup Data Size:	2.65 MB

Restore

Restore Result:	N/A
Restore Identifier:	N/A
Restore Percentage Complete:	0 %

Figure 77 FSG Backup Attributes

LDR Verification

The grid is said to be self-healing. This means that the grid checks the integrity of the ingested objects via a process called background verification. If a corrupt object is found on an LDR, the object is quarantined and replaced with a copy of an uncorrupted object stored elsewhere on the grid. The existence of corrupt objects can indicate disk corruption or data tampering. Background verification operates at an adaptive priority to avoid interfering with normal grid operations.

You can view information about LDR verification, for example progress and number of corrupt objects found, on the Verification component of the LDR (see Figure 78). Any corrupt object should be investigated.

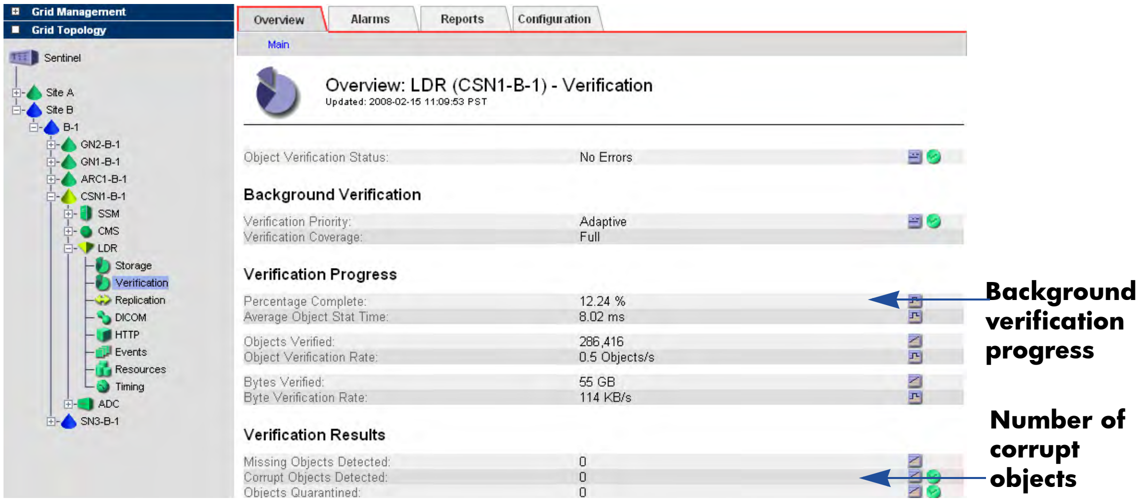

Overview: LDR (CSN1-B-1) - Verification
Updated: 2008-02-15 11:09:53 PST

Object Verification Status: No Errors

Background Verification

Verification Priority: Adaptive
Verification Coverage: Full

Verification Progress

Percentage Complete: 12.24 %
Average Object Stat Time: 8.02 ms

Objects Verified: 286,416
Object Verification Rate: 0.5 Objects/s

Bytes Verified: 55 GB
Byte Verification Rate: 114 KB/s

Verification Results

Missing Objects Detected: 0
Corrupt Objects Detected: 0
Objects Quarantined: 0

Background verification progress

Number of corrupt objects

Figure 78 LDR Verification Attributes

There is another type of LDR verification called foreground verification. Foreground verification detects whether objects are missing. The LDR foreground verification procedure is initiated by a grid task and is used mostly during maintenance. See Table 19 for a comparison of the two types of LDR verification.

Table 19 LDR Verification

Background	Foreground
Runs continuously at a low level	Used in maintenance procedures
Automatic	Initiated by a grid task
Identifies corrupt objects	Identifies missing objects
Performed by the Storage Nodes	Performed by Control Nodes
Slower process	Faster process
Adaptive, lower priority	Higher priority

All attributes on Figure 78 refer to LDR background verification except for Missing Objects which is updated by foreground verification.

Tape Node Capacity

Each Tape Node can interface with the supported type of archival storage: archival media managed by Tivoli Storage Manager (TSM).

In a grid that includes a TSM Tape Node, the TSM middleware has no way to inform the Tape Node when the TSM database or the archive media managed by the TSM is near capacity. The Tape Node will continue to accept objects for archiving after the TSM stops accepting new content and the Store Failures (ARVF) alarm is triggered.

NMS Database Usage Rates

Attribute values are saved to the NMS database. As attribute data is saved to the NMS database the size of the NMS database grows as the amount of free tablespace decreases.

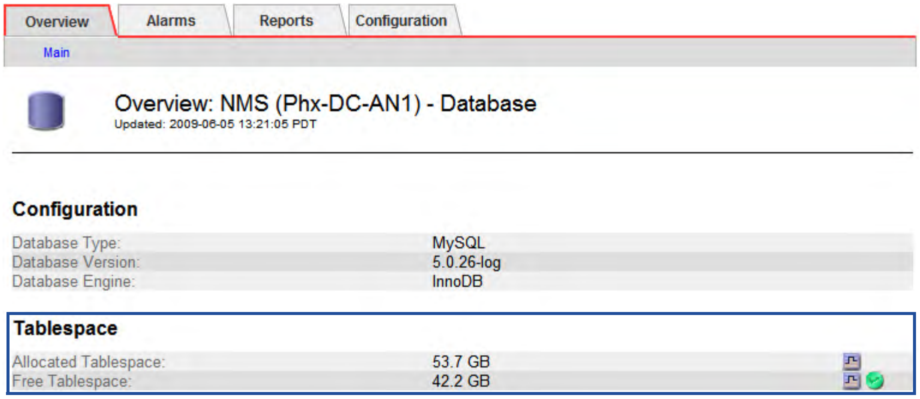

Overview: NMS (Phx-DC-AN1) - Database

Updated: 2009-06-05 13:21:05 PDT

Configuration

Database Type:	MySQL
Database Version:	5.0.26-log
Database Engine:	InnoDB

Tablespace

Allocated Tablespace:	53.7 GB
Free Tablespace:	42.2 GB

Figure 79 NMS Database

Monitor database usage rates to determine when the amount of free tablespace remaining will reach a critical level. When the NMS database begins to run out of free tablespace the Admin Node must be refreshed. (The NTBR alarm is triggered.)

In the following example, database rate usage is a steady .3 GBs per day.

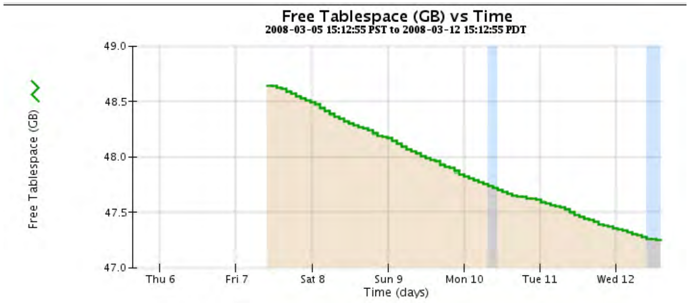

Figure 80 Free Tablespace Usage Rate

You may notice spikes in the Free Tablespace chart. At regular intervals, raw data and downsampled data is purged from the NMS database. This will reclaim some tablespace. Note that not all attribute data is purged from the NMS database.

Grid Tasks

A grid task is a program that performs grid procedures that involve several grid services automatically. For instance, LDR foreground verification is performed via a grid task. Most maintenance and expansion procedures involve running grid tasks.

You can follow the progress of a grid task from the CMN Grid Tasks Overview tab (see Figure 81 and Table 20). Running grid tasks is restricted to accounts with Maintenance permissions such as the Admin and Vendor accounts.

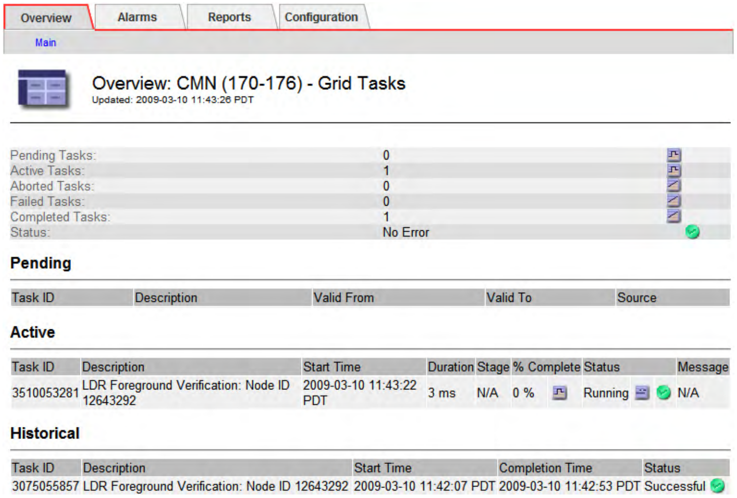

Overview: CMN (170-176) - Grid Tasks
Updated: 2009-03-10 11:43:26 PDT

Pending Tasks:	0
Active Tasks:	1
Aborted Tasks:	0
Failed Tasks:	0
Completed Tasks:	1
Status:	No Error

Pending

Task ID	Description	Valid From	Valid To	Source

Active

Task ID	Description	Start Time	Duration	Stage	% Complete	Status	Message
3510053281	LDR Foreground Verification: Node ID 12643292	2009-03-10 11:43:22 PDT	3 ms	N/A	0 %	Running	N/A

Historical

Task ID	Description	Start Time	Completion Time	Status
3075055857	LDR Foreground Verification: Node ID 12643292	2009-03-10 11:42:07 PDT	2009-03-10 11:42:53 PDT	Successful

Figure 81 Grid Tasks Overview

Grid tasks go through three distinct phases:

Pending	The grid task has been submitted, but not yet started.
Active	The grid task has been started. It can be either actively running or temporarily paused.
Historical	A historical grid task is a task that has been submitted but is no longer active. This includes grid tasks that completed successfully, grid tasks that were rejected (for example because the valid time period had expired), grid tasks that were cancelled or aborted, and grid tasks that terminated in error

Table 20 Grid Tasks Overview Fields

Field	Description
Task ID	Unique identifier assigned when the task is created.
Description	Brief description of the purpose of the task.
Valid From	Date from which the task is valid. The grid task will be rejected if it is submitted before this date.
Valid To	Date until which the task is valid. The grid task will be rejected if it is submitted after this date.
Source	The author of the grid task.
Start Time	Date and time on which the grid task was started.
Stage	Description of the current stage of the active task.
% Complete	Progress indicator for active tasks.
Duration	Estimated amount of time since the grid task was started.
Status	Current status of the active or historical task. For active tasks, one of: • Starting • Running • Pausing • Paused • Error: An error has been encountered. User action is required. • Aborting • Abort Paused: Task failed to be aborted and is paused in error. For historical tasks, one of: • Successful • Expired • Aborted • Cancelled • Duplicate • Invalid
Message	Information about the last stage of the active task.
Completion time	The date and time on which the grid task completed (or failed or expired or was aborted).

Common Alarms

Table 21 lists common alarms that are usually no cause for concerns as long as trends do not develop.

Table 21 Common Alarms

Category	Code	Service	Notes
Content replication	RIRF RORF	LDR ARC	Replication alarms (Inbound Replications - Failed RIRF and Outbound Replications - Failed RORF) occur in general during periods of high load or due to temporary network disruptions. After grid activity goes back down, these alarms should clear. If the count of failed replications continues to increase, look for network problems and verify that the source and destination LDRs and the ARCs are online and available.

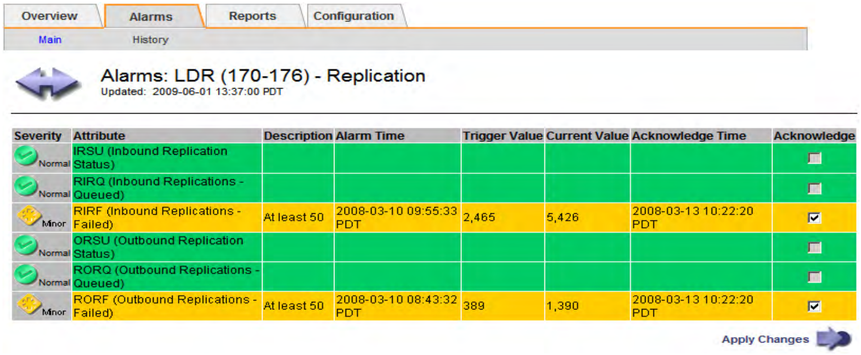

Alarms: LDR (170-176) - Replication

Updated: 2009-06-01 13:37:00 PDT

Severity	Attribute	Description	Alarm Time	Trigger Value	Current Value	Acknowledge Time	Acknowledge
Normal	IRSU (Inbound Replication Status)						□
Normal	RIRQ (Inbound Replications - Queued)						□
Minor	RIRF (Inbound Replications - Failed)	At least 50	2008-03-10 09:55:33 PDT	2,465	5,426	2008-03-13 10:22:20 PDT	☑
Normal	ORSU (Outbound Replication Status)						□
Normal	RORQ (Outbound Replications - Queued)						□
Minor	RORF (Outbound Replications - Failed)	At least 50	2008-03-10 08:43:32 PDT	389	1,390	2008-03-13 10:22:20 PDT	☑

Common LDR Replication Alarms: RIRF & RORF

Network	NRER NTER	SSM	Network interface errors (Receive Errors NRER and Transmit Errors NTER) are fairly common with some network interface adapters. These errors may clear without being manually reset. If they do not clear, check the network hardware.

Table 21 Common Alarms *(continued)*

Category	Code	Service	Notes

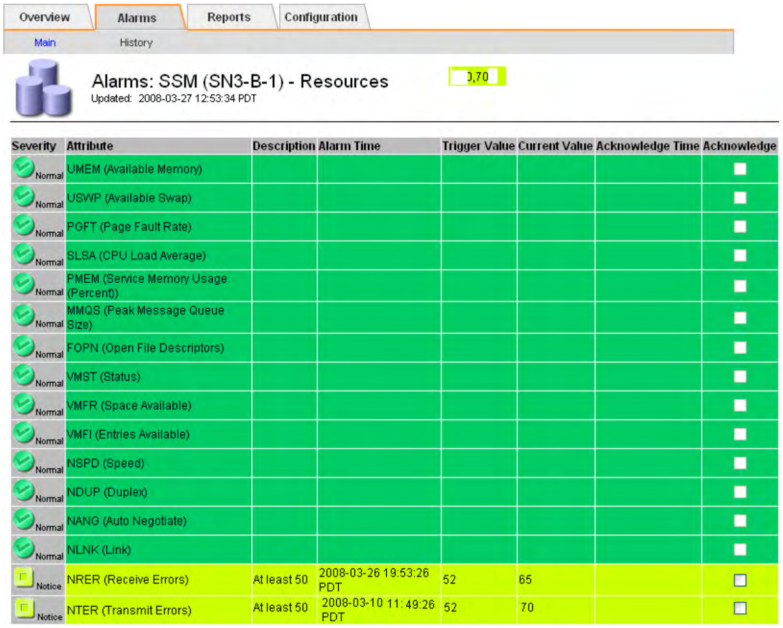

Alarms: SSM (SN3-B-1) - Resources

Updated: 2008-03-27 12:53:34 PDT

Severity	Attribute	Description	Alarm Time	Trigger Value	Current Value	Acknowledge Time	Acknowledge
✓ Normal	UMEM (Available Memory)						□
✓ Normal	USWP (Available Swap)						□
✓ Normal	PGFT (Page Fault Rate)						□
✓ Normal	SLSA (CPU Load Average)						□
✓ Normal	PMEM (Service Memory Usage (Percent))						□
✓ Normal	MMQS (Peak Message Queue Size)						□
✓ Normal	FOPN (Open File Descriptors)						□
✓ Normal	VMST (Status)						□
✓ Normal	VMFR (Space Available)						□
✓ Normal	VMFI (Entries Available)						□
✓ Normal	NSPD (Speed)						□
✓ Normal	NDUP (Duplex)						□
✓ Normal	NANG (Auto Negotiate)						□
✓ Normal	NLNK (Link)						□
□ Notice	NRER (Receive Errors)	At least 50	2008-03-26 19:53:26 PDT	52	65		□
□ Notice	NTER (Transmit Errors)	At least 50	2008-03-10 11:49:26 PDT	52	70		□

Common Network Alarms: NRER, NTER

Table 21 Common Alarms *(continued)*

Category	Code	Service	Notes
Resource utilization	UMEM	SSM	Minor alarms for Available Memory (the amount of system RAM available for system operations) set by default at 100 MB are not a cause for concern unless available memory continues to decrease. This could indicate a serious problem.
Total events	SMTT	SSM	The total number of logged error or fault events (Total Events SMTT) includes errors such as network errors and FSG replication errors. Unless these errors have been cleared (that is, the count has been reset to 0), total events alarms may be triggered. This alarm is safe to ignore only if the events that triggered the alarm have been investigated.

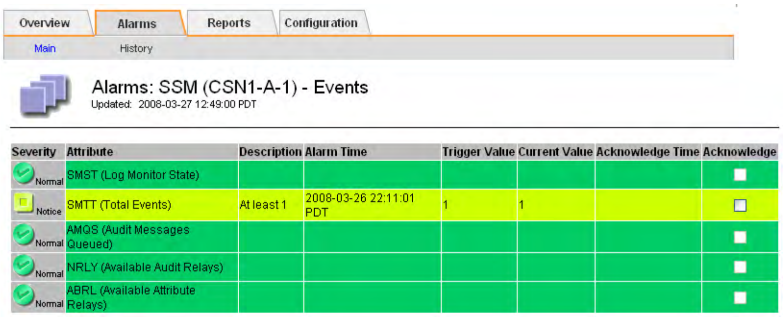

Alarms: SSM (CSN1-A-1) - Events

Updated: 2008-03-27 12:49:00 PDT

Severity	Attribute	Description	Alarm Time	Trigger Value	Current Value	Acknowledge Time	Acknowledge
✓ Normal	SMST (Log Monitor State)						□
□ Notice	SMTT (Total Events)	At least 1	2008-03-26 22:11:01 PDT	1	1		□
✓ Normal	AMQS (Audit Messages Queued)						□
✓ Normal	NRLY (Available Audit Relays)						□
✓ Normal	ABRL (Available Attribute Relays)						□

Common Events Alarm: Total Events

5 Grid Configuration

This chapter describes the following grid configuration settings:

- deduplication
- storage compression
- encryption
- security partioning
- HTTP No-Delete flag
- link costs
- audit levels
- NMS entities
- connection profiles

Deduplication

When deduplication is enabled, the grid detects duplicate objects ingested at multiple locations and stores only the number of copies required by the ILM policy.

NOTE Deduplication is unavailable if the grid is configured for metadata replication. If deduplication is enabled, metadata replication cannot be enabled.

The main purpose of the deduplication feature is to allow the grid to be used as an image archive in a multi-site environment with client applications that perform their own cross-site data replication to ensure data redundancy. This feature is specifically intended for GE Enterprise Archive Server.

When a file is ingested into the grid, it is assigned a unique content handle and a file type identifier that allows the grid to identify the file as eligible for deduplication. When the grid identifies two files as being identical, it “deduplicates” them by redirecting all content handles to point to a single stored instance of the file. Files are then stored and replicated normally by the grid’s ILM policy. As a result, only the number of copies specified by the ILM policy are permanently stored in the grid. Without deduplication, each of the identical copies of the file would be stored and replicated by the grid, consuming at least twice the necessary amount of storage space.

A request to retrieve any of the identical files stored to the grid returns the same deduplicated object.

Deduplication and GE Optimized Store

A GE Enterprise Archive may be configured to save two identical copies of each file to its storage device from parallel archive servers. In normal operation, when an HP MAS system is used as a storage device, it automatically makes multiple copies of each ingested file according to the ILM policy defined for the grid. However, when GE Optimized Store is enabled, the grid is prevented from creating unnecessary copies of files. The optimized store feature inspects ingested objects to identify identical pairs saved from a GE Enterprise Archive Server (version 3.0 or later), and "deduplicates" the identical objects.

The GE Optimized Store feature is only effective when identical objects can be identified successfully. If the parallel Enterprise Archive servers run different versions or patch levels of Enterprise Archive software, the grid can not identify files saved from these servers as being identical and unnecessary copies of each file are created in the grid. When integrating HP MAS software with a GE Enterprise Archive Server, ensure that both Enterprise Archive Servers use identical version and patch levels of software. When updating the software on GE Enterprise Archive servers, either take care to update both servers promptly, or interrupt data storage until the update is complete. If you do not, grid storage is unnecessarily consumed by multiple unnecessary copies of files.

GE Enterprise Archive Server version 3.0 or later saves consolidated objects to the grid. Consolidated files are related files that have been compressed into one .tar file. Along with the .tar file, GE Enterprise Archive Server may also save related .dcm files that have not been saved in a .tar file. Deduplication is more complex when saving consolidated (.tar) files to the grid. Because each archive server creates its own .tar file, the .tar files have differing time stamps and are not binary identical. Therefore, consolidated files saved to the grid are first "sliced" into two components: a container that includes the TAR file headers, and a content file. If the grid identifies two or more content files as being binary identical, the grid points the object identifiers of these files to a single instance. The container files and the content file are then stored and replicated according to the grid's ILM policy.

When a "sliced and deduplicated" object is requested from the grid, the grid finds the appropriate container object and a copy of the content file. It then re-assembles the requested object, and returns it to the requester. The content returned to the requester is binary identical to the one saved to the grid.

Deduplication Details

The behavior of the deduplication feature depends on the grid deploment. The deduplication feature was designed for a Data Center + Disaster Recovery (DC + DR) configuration used as a storage archive for GE Enterprise Archive Server version 3.0.

- Object metadata related to the ingest location and time is not guaranteed to be correct for deduplicated objects.
- In a grid configured not to purge content on content handle release, deduplication will not result in any storage space savings.

- Deduplication may not work as intended on legacy grids that have multiple generations of Control Nodes.
- Deduplication is unavailable on grids configured to use metadata replication (that is, the CMS uses metadata replication).
- When deduplication is used in a grid with a Tape Node, the ILM must be chosen appropriately.
- If the ILM policy makes a copy of the object on tape at ingest, enabling deduplication may provide no benefit because the TSM/SSAM may not be able to reclaim space when the object copy is purged.
- Deduplication is performed on a best-effort basis. The grid ensures data safety before performing deduplication. When the grid is under heavy load or when there are failed nodes or network connectivity problems in the grid, objects may not be deduplicated.
- The computational load of performing deduplication is high, such that the additional cost of enabling it in non-standard configurations can be higher than that of storing extra copies.

Enabling Deduplication

Deduplication is disabled by default. Deduplication may be enabled at any time. If deduplication is enabled on a running grid that has pre-existing content, only content ingested after deduplication is enabled is eligible for deduplication.

△ CAUTION After it is enabled, deduplication cannot be disabled.

To determine if deduplication is enabled, log in to the NMS MI, go to **Grid Management > Grid Configuration> Overview,** and look up the value of Deduplication.

NOTE Deduplication is unavailable if the grid is configured for metadata replication. If deduplication is enabled, metadata replication cannot be enabled.

Enabling configuration is restricted to the Vendor account.

Storage Compression

If enabled, the lossless compression of objects stored to the grid reduces the size of the objects and can result in the ability to save up to twice as much data. By default storage compression is enabled.

Applications saving an object to the grid may or may not compress the object before saving it. If an application compresses an object before saving it to the grid, enabling compression does not reduce the object's size.

To determine if storage compression is enabled, log in to the NMS MI, go to **Grid Management > Grid Configuration > Overview**, and look up the value of Compression.

NOTE By default storage compression is enabled.

Disabling storage compression is restricted to the Vendor account.

Encryption

The encryption option enables encrypted storage of all stored data so that if a server is compromised no data can be retrieved in any readable form.

To determine if encryption is enabled, log in to the NMS MI and go to **Grid Management > Grid Configuration > Overview** and look at the value of Encryption.

Disabling encryption is restricted to the Vendor account.

When encryption is disabled, new content ingested into the grid is not stored in an encrypted form. However, existing content that was previously encrypted remains encrypted.

HTTP No-Delete Flag

The HTTP No-Delete flag overrides the delete permissions defined in the HTTP Advanced profiles. If enabled, all HTTP DELETE operations are denied.

To determine if HTTP No-Delete is enabled, log in to the NMS MI, go to **Grid Management > Grid Configuration > Overview**, and look at the value of HTTP No-Delete Flag.

NOTE By default HTTP No-Delete Flag is disabled.

Enabling HTTP No-Delete is restricted to the Vendor account.

Security Partitions

Security partitioning is a mechanism that restricts access to content ingested into the grid—content ingested either directly through the HTTP API, or via a Gateway Node. Security partitioning prevents the query or retrieval of an object from the grid by applications other than the application that ingested the object.

Security partitions are used when it is imperative for security and privacy reasons that applications be prevented from retrieving data ingested at another location or by another HTTP API client.

Enabling and configuring security partitions is restricted to the `Vendor` account.

To determine if security partitioning is enabled:

1. Log in to the NMS MI.
2. Go to **Grid Management > Grid Configuration > Overview**.
3. Look up the value of Security Partitions.

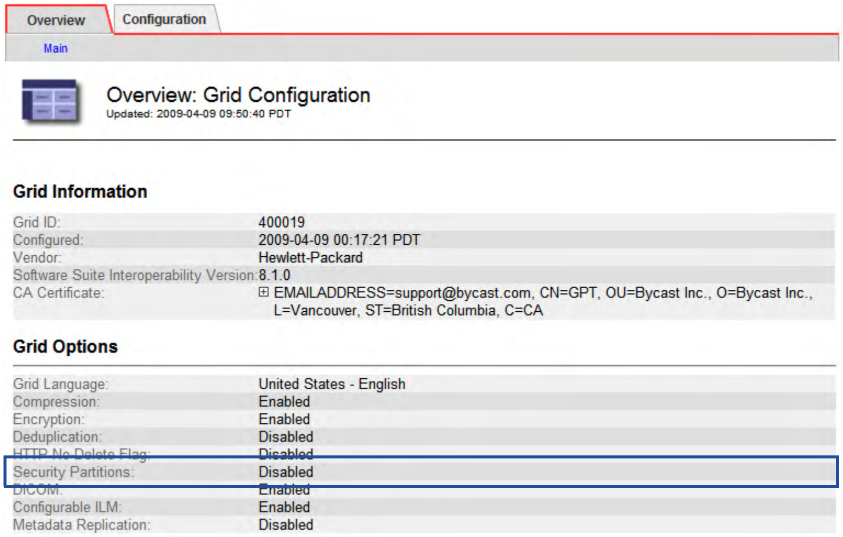

Overview: Grid Configuration
Updated: 2009-04-09 09:50:40 PDT

Grid Information

Grid ID:	400019
Configured:	2009-04-09 00:17:21 PDT
Vendor:	Hewlett-Packard
Software Suite Interoperability Version:	8.1.0
CA Certificate:	⊞ EMAILADDRESS=support@bycast.com, CN=GPT, OU=Bycast Inc., O=Bycast Inc., L=Vancouver, ST=British Columbia, C=CA

Grid Options

Grid Language:	United States - English
Compression:	Enabled
Encryption:	Enabled
Deduplication:	Disabled
HTTP No-Delete Flag:	Disabled
Security Partitions:	Disabled
DICOM:	Enabled
Configurable ILM:	Enabled
Metadata Replication:	Disabled

Figure 82 Security Partitions Enabled or Disabled

Link Costs

Link costs refer to the relative costs of communicating between different groups of services within the grid. Link costs are used to determine which server in the grid should provide a requested service. For example, link cost information is used to determine which Storage Nodes are used to retrieve objects. All else being equal, the server with the lowest link cost is preferred.

In the example shown in Figure 83, if a user at the Disaster Recovery (DR) site retrieves an object that is stored both at the Data Center (DC) and the Satellite site, the Storage Node at the DC is responsible for sending the object because the link cost from the DC to the DR site is 25, which is lower than the link cost from the Satellite site to the DR Site (50).

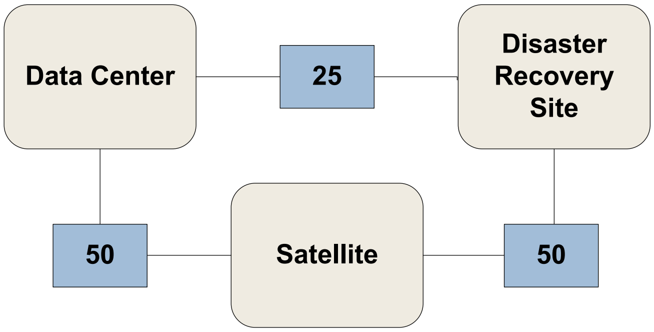

Figure 83 Link Costs Between Nodes

The default values for new installations are shown in Table 22.

Table 22 Default Link Costs

Link	Link Cost	Notes
Between servers within a group	0	Usually a high speed link exists between all servers in the same group.
Between groups that both have an ADC service	25	DC and DR sites usually have an ADC service.
Between one group that has an ADC service and one that does not	50	A satellite site does not usually have an ADC service.
Between groups that do not contain an ADC service	100	This is the maximum cost you can assign to a link.

To view the link cost information, log in to the NMS MI and go to **Grid Management > Grid Configuration > Link Costs**. Link cost configuration is restricted to the Vendor account.

Storage

Storage displays the storage watermarks and the ports for ingest, query, and retrieve.

To view the storage watermarks information, log in to the NMS MI and go to **Grid Management > Grid Configuration > Storage**. Storage configuration is restricted to the Vendor account.

Storage Grades

Storage Grades are used to configure the storage pools used in configurable ILMs. To view the storage grade information, log in to the NMS MI and go to **Grid Management > Grid Configuration > Storage Grades**. Storage Grade configuration is restricted to the Vendor account.

Audit Levels

Audit messages are generated as grid services perform various activities. These messages are processed by the AMS service and stored in the form of text log files.

The audit level determines the amount of details contained in the audit messages.

Table 23 Audit Levels

Audit Level	Description
Off	No audit messages are logged.
Error	Only error messages are logged.
Debug	Trace messages are logged. This is for troubleshooting only.
Normal	Standard transactional messages are logged

To view the audit level information, log in to the NMS MI and go to **Grid Management > Grid Configuration > Audit**. Audit level configuration is restricted to the Vendor account.

NMS Entities

NMS entities refer to the entities shown in the grid topology tree, that is, the items in the grid topology tree that appear above the component level (the names of the grid, locations, cabinets, nodes, and services). NMS entity settings determine:

the name that appears for the entity in the grid topology tree and elsewhere, in the language selected for the user account

the sequence that entities appear in the grid topology tree

Names are allocated to each entity by using a system of Object IDs (OIDs) that are unique to each entity while being hierarchically organized. Each row in the NMS Entities table allocates a name in a specified language to an entity OID. The combination of OID hierarchy and position in the table determines the sequence of appearance of names in the grid topology tree.

To view the NMS Entities information, log in to the NMS and go to **Grid Management > Grid Configuration >** `NMS Entities`. NMS Entities configuration is restricted to the Vendor account.

DICOM Indexes

DICOM is optional, and may not appear in your grid.

The DICOM Indexes page lists the DICOM tags used by the grid for content query and management, if DICOM is enabled for the grid. This information is read-only and cannot be altered through the NMS MI.

To view the DICOM Indexes information, log in to the NMS MI and go to **Grid Management > Grid Configuration > DICOM Indexes**.

Connection Profiles

DICOM is optional, and may not appear in your grid.

Access to the grid for external applications that use the grid's HTTP API or DICOM is controlled by configuring a connection profile in the NMS MI. The following components of the Grid Management > Grid Configuration menu show the connection profiles used by client applications or DICOM entities:

- Link Cost Groups
- HTTP
- HTTP Advanced
- HTTP Security Partitions (optional)
- DICOM
- DICOM Advanced

In general, a connection profile is a named definition of capabilities or behaviors that can be assigned to one or more external applications or DICOM entities. By defining a profile with a given set of permissions, that set of permissions can then be granted to an application or applications by referencing the profile name.

The grid applies security checks and enforces access permissions to prevent external applications from making unauthorized access. To understand the configuration options for connection profiles, you need to understand the layers of connectivity and a little about the protocols used to communicate with the grid.

At the top level, connections to the grid are made using TCP/IP. Before a TCP/IP connection is opened, the grid checks to see if the external application is on its list of "friendly" IP addresses. If the connection request does not originate from a listed IP address, the grid ignores the connection request; the application is not aware that there was any device at the called address.

Assuming that the grid permits the TCP/IP connection to be made, the grid further restricts the caller to permit it to interact with the grid only via the permitted connection protocol: HTTP or DICOM (or both). Within a permitted protocol, the grid grants or declines permission for an external application to carry out specific operations. These permissions are defined in "profiles" that can be allocated to individual applications or groups of applications, based on their IP address.

Security partitioning, while not directly a configuration element of connection profiles, is nevertheless related. If security partitioning is enabled, data ingested through an HTTP API client can only be retrieved by the same client (multiple clients can have read-only access to the same partition), while data ingested via a Gateway Node can only be retrieved by applications with permission to mount to that particular Gateway Node's FSG service (or HTTP API clients with read-only access to that replication group's content). For more information, see Security Partitions (page 139).

For DICOM entities (AEs), the grid can also create logical partitions of the data. This means that selected entities can have their access restricted to a partitioned subset of the grid data.

Some DICOM modalities are limited in their ability to handle particular actions or data transfer formats (syntax). Configurations can be set to restrict activities of these entities to a subset of operations or transfer formats. Setting these types of restrictions, or enabling associations from private or unsupported SOP classes, requires you to be familiar with the DICOM standard.

A configuration profile definition may take more than one line in a table. All lines using the same (case sensitive) profile name are part of the profile. The sequence in the table is significant; the table is processed from the top downward. Therefore it is possible to specify specific exceptions to a general rule by listing the exceptions closer to the top of the table, and then specifying the general rule that applies to all remaining entities.

HTTP Configuration Profile

To understand the configuration profile of an external application that accesses the grid via HTTP, you need to view the following **Grid Management > Grid Configuration** components:

- HTTP Advanced
- HTTP
- Link Cost Groups

HTTP Profile configuration is restricted to the Vendor account.

HTTP Advanced Component

Each external application that interacts with the grid via HTTP must be assigned a profile that defines which activities the application can perform on the grid.

Table 24 describes each namespace.

Table 24 Grid Management > Grid Configuration > HTTP Advanced Settings

Prompt	Type	Description
HTTP /CBID Namespace		
The Content Block ID (CBID) namespace is owned by the grid and is used *within* the HP MAS system software. The content referenced by a CBID may be deleted by the grid (according to grid business rules) at any time. External use of the CBID namespace is deprecated in favor of the use of the /UUID namespace.		
Profile Name	Text	User-defined profile name. This name is referenced in the HTTP component.
GET	Check box	Enables use of the HTTP GET command by clients assigned this profile.
POST	Check box	Enables use of the HTTP POST (query) command by clients assigned this profile.
HTTP /UUID Namespace		

Table 24 Grid Management > Grid Configuration > HTTP Advanced Settings *(continued)*

Prompt	Type	Description
The Universal Unique ID (UUID) namespace is the preferred namespace used to store, access, and delete (hide) data using a unique identifier. This UUID provides an abstracted primary key (handle) for stored content.		
Profile Name	Text	User-defined profile name. This name is referenced in the HTTP component.
PUT	Check box	Enables use of the HTTP PUT command to store files.
GET	Check box	Enables use of the HTTP GET command by clients assigned this profile.
POST	Check box	Enables use of the HTTP POST (query) command by clients assigned this profile.
DELETE	Check box	Enables use of the HTTP DELETE command to release a UUID handle. The DELETE option can be overridden by enabling HTTP No-Delete Flag on Grid Management > Grid Configuration. For more information, see HTTP No-Delete Flag (page 139).
HTTP /DICOM Namespace		
The Digital Imaging and Communications in Medicine (DICOM) namespace is used to store data in grids where the DICOM option has been purchased and configured.		
Profile Name	Text	User-defined profile name. This name is referenced in the HTTP component.
AE Title	Text	The AE Title allocated to the remote entity. Used to assign an AE title to all content submitted by an entity assigned this profile. DICOM objects must be saved to the grid with the AE title of the sending device. However, when using the HTTP interface, you do not establish a DICOM association between the grid and the device sending the object. Therefore the sender cannot provide an AE title when it submits an object. Instead, you must specify the AE title of the sending device in the HTTP /DICOM Namespace table.
Coerce Tag Profile Name	Pull-down menu	A Coerce Tag Profile label defined in the DICOM Advanced component. This is a reference to another configuration profile that is used for DICOM transactions. When an entity assigned this HTTP profile submits a DICOM file, the behavior associated with the Coerce Tag Profile named here is also applied. The profile is used to manage logical partitioning of data. This field can be left blank if no partitioning is applied to entities with the profile. For more information, see the *DICOM Integration Guide.*
PUT	Check box	Enables use of the HTTP PUT command by clients assigned this profile.

Table 24 Grid Management > Grid Configuration > HTTP Advanced Settings *(continued)*

Prompt	Type	Description
HTTP /GRID Namespace		
The GRID namespace is used to query the grid for information about nodes and the services they provide. The GRID namespace is used by custom applications developed to store or retrieve grid content via the UUID or DICOM namespaces.		
Profile Name	Text	User-defined profile name. This name is referenced in the HTTP component.
POST	Check box	Enables use of the HTTP POST command by clients assigned this profile. Enabling POST in the grid namespace also enables a custom application to store custom audit messages supplied by the application.
HTTP Metadata Indexing		
HTTP Metadata Indexing is not a namespace, and is not configurable via the NMS. This table lists the custom metadata defined for use with this grid via the HTTP API.		

HTTP Component

In the HTTP component, the configuration profiles listed in the HTTP Advanced component are assigned to IP address ranges that include the IP address of the external application that uses the profile.

External applications that do not have an assigned HTTP profile cannot access grid content. The IP addresses of all entities intended to use HTTP to access the grid appear somewhere on this list.

Table 25 Grid Management > Grid Configuration > HTTP Settings

Prompt	Type	Description
Description	Text	User-defined description for the set of entities in this IP range / profile group.
IP Range	Dotted decimal or CIDR	Sets the IP address or range of remote devices to which the grid assigns the profile. A hyphen or slash indicates a range, inclusive of the values entered. For example: • 192.168.120.0/24 (CIDR format) • 192.168.142.20-192.168.142.28 (dotted decimal) An abbreviated format for masks in eight-bit steps is available. For example: `192.168.142.0` is equivalent to the CIDR notation `192.168.142.0/24`. This can be extended: *n*.*n*.0.0 is equivalent to *n*.*n*.0.0/16.
Profile Name	Pull-down menu	Case sensitive reference to a profile name defined in the HTTP Advanced component. The profile governs the permitted activities for connections from this IP range. The IP Ranges specified here can be subsets of the addresses from the IP Ranges component. This enables a workgroup to be broken down into particular activity profiles on the grid.

Link Cost Groups (Client Group IP Ranges Table)

The table of Client Group IP Ranges in the Link Cost Groups component is used to assign groups to the clients. The clients access the grid for queries via these groups.

A caller's IP address may match more than one of the ranges specified in the Client Group IP Range table. When the grid is validating the caller, it searches the table from the top downward. The first match found is used to assign the protocol permissions to the caller.

Table 26 Grid Management > Grid Configuration > Link Cost Groups Component: Client Group IP Ranges Settings

Prompt	Type	Description
IP Range Name	Text	User-defined label to identify the IP range, usually a location or workgroup.
IP Range	Dotted decimal or CIDR	Specifies the IP range to assign to the grid group. IP ranges are specified using one of the following: • a single IP address • a hyphenated list of IP addresses inclusive of the listed values (e.g. 174.182.91.00-174.182.91.64) • a range of IP addresses specified using CIDR notation (e.g. 192.168.120.0/27) • an IP address in the form A.B.C.D, where at least one of A, B, C, or D is zero (dotted decimal). • If D is 0, then IPs between A.B.C.0 and A.B.C.255 will be in range. If C and D are both 0, then IPs between A.B.0.0 and A.B.255.255 will be in range. If B, C, and D are all 0, then IPs between A.0.0.0 and A.255.255.255 will be in range. If A, B, C, and D are all 0, then all IPs will be in range. For example: 192.168.142.0 is equivalent to the CIDR notation 192.168.142.0/24. This can be extended: n.n.0.0 is equivalent to n.n.0.0/16.
Group ID	Text	The grid can improve performance by knowing which grid server group the external application is associated with. This helps the grid route data more efficiently by using services that operate at a lower "cost". To achieve this, the configuration includes a group ID for each external application. If no group is assigned, or the group ID specified is not known to the grid, the connection cost to that location is assumed to be zero.For more on groups, see Link Costs (page 141).

DICOM Configuration Profile

To understand the configuration profile of a DICOM device, you need to view the following Grid Management > Grid Configuration components:

- DICOM Advanced
- DICOM
- Link Cost Groups

DICOM Advanced

Each DICOM device that uses the grid is allocated a configuration profile that defines the capabilities the grid offers to that device. The DICOM Advanced component defines the profiles that can be assigned to the DICOM device.

The grid can be partitioned in a way that restricts visibility of content to particular entities. To achieve this, the grid can apply, or coerce, a tag in the DICOM file to a particular value. The tag can be selectively applied at ingest time to allocate data to a partition, and used during queries from entities assigned to a particular profile to restrict access.

The Advanced Config Profile table is used during the association handshake to control the types of actions permitted and the format of files in the transfer syntax. Some entities present problems when a particular action (SOP class) or transfer syntax is used. If an Advanced Config Profile is applied, the grid enforces a complex rule set on the association. For more information, see the *DICOM Integration Guide*. Before making entries to the DICOM Profiles table that use either Coerce Tag or Advanced Config profiles, you must first define the target profile name in the associated table below. Input validation permits you to enter only profiles that already exist in the Profile Name fields.

Table 27 Grid Management > Grid Configuration > DICOM Advanced Component Configuration Settings

Prompt	Type	Description
DICOM Profiles		
Profile Name	Text	User-defined name for this profile. This name is referenced in the DICOM component.
S	Check box	Send to GRID—enables a client with this profile to send DICOM files to the grid (ingest content).
R	Check box	Receive from GRID—enables a client with this profile to receive DICOM files from the grid (retrieve content).
F	Check box	Find on GRID—enables a client with this profile to issue C-Find commands for DICOM content.
M	Check box	Move—enables a client with this profile to request the grid to open an association to relay an object (C-Move).
C	Check box	Commit Storage—enables a client with this profile to request acknowledgement of object storage in the grid.
Coerce Tag Profile Name	Pull-down menu	A Coerce Tag Profile label used to manage logical partitioning of data.
Advanced Config Profile Name	Pull-down menu	Case sensitive reference to a profile name defined in the Advanced Config Profiles table of the DICOM Advanced component. This optional profile may be used to permit Private SOP classes, and/or to restrict actions and dictate preferences for transfer syntax.

Table 27 Grid Management > Grid Configuration > DICOM Advanced Component Configuration Settings *(continued)*

Prompt	Type	Description
Behavioral Profile Name	Pull-down menu	Preconfigured profile that directs the HP MAS system to adjust its DICOM behavior for compatibility with specific DICOM clients. Current values are: • Xcelera A (Required for Philips Xcelera, version 1.2 L4 SP1) • syngo Dynamics A (Required for Siemens syngo Dynamics/ KinetDX PACS, version 5.0)

Table 27 Grid Management > Grid Configuration > DICOM Advanced Component Configuration Settings *(continued)*

Prompt	Type	Description
Coerce Tag Profiles		
Coerce Tag Profile Name	Text	User-defined name for this coerce tag profile. This name is referenced in the HTTP Advanced component for the /DICOM namespace and by the DICOM Profiles.
Tag	DICOM private tag name	A DICOM private tag name (as per DICOM specifications) that is stored with the CMS metadata, not with the payload object.
Tag Value	Text	Value assigned to the tag indicating the data partition or forced value.
Query	Check box	Restricts entities using this profile to only have visibility to the named data partition.
Ingest	Check box	Allocates data submitted to the grid from an entity with this profile to a specific partition.
Description	Text	User-defined description of the data partition.
Advanced Config Profiles		
Advanced Configuration Profile Name	Text	User-defined name for this advanced configuration profile. This name is referenced by the DICOM Profiles. If an Advanced Configuration Profile is specified, it must be given a name. (This name need not be unique, as you can use multiple table rows to build a list of allowed or disallowed classes.)
Behavior	Pull-down menu	Can be set to either: • Allow • Disallow The behavior is assigned to the specified SOP Class. By default, all classes listed in the grid's *DICOM Conformance Statement* are allowed. Use multiple table rows with the same profile name to build a list of allowed or disallowed classes. If the profile is to control transfer syntax only, or if the profile is to be used to grant access to a Private SOP class (one that does not have a prefix of 1.2.840.10008), set the Behavior to **Allow**. For more information, see the *DICOM Integration Guide*.

Table 27 Grid Management > Grid Configuration > DICOM Advanced Component Configuration Settings *(continued)*

Prompt	Type	Description
SOP Class	DICOM standard string	Class (action) to which the Behavior applies. Can be one of: • a valid OID that begins with a prefix of 1.2.840.10008 and has fewer than 64 characters, representing a supported DICOM SOP class, as per the grid's *DICOM Conformance Statement*. • a valid OID with fewer than 64 characters that represents a private or unsupported SOP class. • an asterisk "*" indicating "all supported classes" (as listed in the grid's *DICOM Conformance Statement*). Use multiple table rows with the same profile name to build a list of allowed or disallowed classes. If the profile is to control transfer syntax only, set this to "*" to apply the profile to all supported SOP classes as listed in the grid's *DICOM Conformance Statement*. (Profiles that use "*" are not applied to private or unsupported SOP classes, even if profiles exist that permit associations using these classes.) For more information, see the *DICOM Integration Guide*.
Preferred Transfer Syntax	DICOM standard string	Preferred format for images, if supported. This is an optional parameter that can be left blank. If specified, it must be a list of valid OIDs, separated by commas. (Valid OIDs begin with a prefix of 1.2.840.10008 and have fewer than 64 characters.) For more information, see the *DICOM Integration Guide*.
Required Transfer Syntax	DICOM standard string	Format required in a request. This is an optional parameter that can be left blank. If specified, it must be a list of valid OIDs, separated by commas. (Valid OIDs begin with a prefix of 1.2.840.10008 and have fewer than 64 characters.) For more information, see the *DICOM Integration Guide*.

DICOM

In the DICOM component, the configuration profiles listed in the DICOM Advanced component are assigned to a DICOM Application Entity (AE) address ranges that include the IP address of the external application that uses the profile. An entity is uniquely defined by the combination of AE Title, IP address and port, and the Grid AE Title to which it connects. For example, entities can share an AE Title and port number while still being distinguished by unique IP addresses within the defined range.

Table 28 Grid Management > Grid Configuration > DICOM Component Configuration Settings

Prompt	Type	Description
Description	Text	User-defined description of the entity or group of entities in this profile group.
AE Title	Text	The AE Title of the remote entity.
IP Range	Dotted decimal or CIDR	IP address (or inclusive range) to which the grid assigns the profile.
Port	Number	The well-known port used for DICOM by the remote device.
Via LDR	Check box	Indicates whether the LDR can make direct connections with remote entities. If *not* selected, connections are routed through a CLB service.The HP MAS must use the CLB (check box is *not* selected) if the grid uses a private network for communications between grid servers.
GRID AE	Text	AE Title of the grid to which the entity can associate. For more information on AE Title, see the *DICOM Integration Guide.*
Profile Name	Pull-down menu	Case sensitive reference to a profile name defined in the DICOM Advanced component. The profile governs the permitted activities for entities in this IP range.

Link Cost Groups (Client Group IP Ranges Table)

The table of Client Group IP Ranges lists what groups the DICOM AE use to access the grid).

Support for IPv6

As of Release 8.1, the HP MAS system offers limited Internet Protocol version 6 (IPv6) support. IPv6 is the successor to IPv4, the currently dominant Internet Protocol version.

IPv6 support is only available for new installations. Existing grids updated to Release 8.1 cannot use IPv6.

IPv6 is currently supported for external interfaces such as the NMS MI, the HTTP API, and the external NTP sources. An IPv6 grid must use IPv6 for all external interfaces. For instance, you cannot use IPv6 for the IP address of the NMS MI and IPv4 for the IP address of the NTP sources.

Grids that use FSGs or DICOM must use IPv4. Only grids that use HTTP API clients exclusively can use IPv6.

IPv6 is only supported for SLES 10 SP2. It is not compatible with SLES 10 SP1.

Table 29 summarizes IPv6 support for the HP MAS software.

Table 29 IPv6 Grid

IP Address	IPv6	IPv4	Notes
External NTP time sources	✓		Related procedures include: • Specify the IP address of the NTP servers in Grid Designer • Run the script setup_ntp.rb to add or remove external NTP time servers • Edit the tags ntp > sources > ip in the grid specification file
NMS MI	✓		Related procedures include: • Specify the IP address of the NMS web interface in Grid Designer • Run the script apply-pending-changes to change the IP address • Edit the tags network > ip and grid network > ip in the grid specification file IPv6 support is restricted to Internet Explorer 7 on Vista.
Client IP Range configuration	✓		This is specified in the NMS MI: configure IP Range in Grid Management > Grid Configuration > HTTP and in Grid Management > Grid Configuration > Link Cost Groups
E-mail server	✓		This is specified in the NMS MI: configure the IP address of the SMTP mail server in Grid Management > NMS Management > E-mail > Server
DNS server	✓		Add or remove DNS server using `setup_resolv.rb`
SNMP Trap/ Manager	✓		
Internal IP address of grid servers		✓	IPv6 support does not include internal grid communications. The following procedures must use IPv4: • Remote access via ssh, for example., ssh <***Gateway_Node_IP_address***> • Entries in the `etc/hosts` file • Maintenance scripts such as the `fsg-reconcile` tool or the CMS database cloning scripts.
NetApp machine	✗	✓	IPv6 support does not include internal grid communications such as NetApp and LDR storage.
syngo IP	✗	✓	syngo integration is done via FSGs and IPv6 support does not include FSGs.
TSM middleware server used by the ARC	✗	✓	

Enter IP addresses carefully. No validation is performed to check the IP address format.

6 Grid Services and Components

This section gives an overview of the complete set of grid services that can be included in a grid and of the components that are included in these grid services.

Not all grids use every grid service and component listed here. Each grid uses only those grid services and components suited to the options purchased and the functionality required.

Information is organized by grid service and component, with grid services listed alphabetically.

Grid Services

Grid services are programs that run on the physical servers. Each delivers particular services to the grid as shown in Table 30. The NMS MI is used to access the critical information gathered by these grid services so that you can monitor the condition and performance of your grid.

By clicking on the grid service name in the NMS MI, you can view detailed data about the service and each of its component functions. The NMS MI provides basic functional statistics, alarm status, reporting functionality, and configuration options for each service and node.

Table 30 HP MAS Services

	Icon	Tag	Service Name	Function	See
Data Plane		LDR	Local Distribution Router	Storing and routing data through the grid.	page 204
		CLB	Connection Load Balancer	Enabling storage and retrieval through the DICOM protocol option.	page 168
		FSG	File System Gateway	Enabling storage and retrieval through standard network file systems.	page 195
		ARC	Archive	Managing long-term nearline storage of data	page 166

Table 30 HP MAS Services *(continued)*

	Icon	Tag	Service Name	Function	See
Control Plane	●	CMS	Content Management System	Storing and managing metadata for the content.	page 171
	■	ADC	Administrative Domain Controller	Authenticating and managing IP address ranges, and reporting network topology. Acts as an attribute relay and as an audit relay.	page 163
Management Plane	▮	SSM	Server Status Monitor	Monitoring server activity and hardware.	page 217
	▲	NMS	Network Management System	Notifying administrators of alarm conditions and providing binding information. Provides the user interface for grid management.	page 208
	✎	CMN	Configuration Management Node	Managing system configuration through the NMS interface. Access to changing settings in this service is restricted to the Vendor account.	page 169
	▤	AMS	Audit Management System	Logging grid transactions for audit and reporting.	page 165

Within a grid, the same grid service can be installed and used on more than one server. The settings made to a service on one server do *not* affect the settings on the same service installed on a different server. System-wide configurations can be made by a service technician under the Vendor account using components of Grid Management > Grid Configuration.

Service Components

Each grid service is composed of one or more components that manage one piece of the service's functionality. In the NMS MI, you can select a service's component to view the attributes that the component manages. The NMS MI also provides alarm status, reporting functionality, and configuration options for each component.

Table 31 lists all of the available service components, in alphabetic order by name. Note that the same symbol may be used by more than one component.

Table 31 Service Components

Icon	Name	Parent Grid Service	Data Provided
▦	Audit	CMN	Configuration of the audit messages to be logged.
oo	Backup	FSG	FSG backup settings and status.
▦	Connect to Support	CMN	Creates an ssh connection to Support's troubleshooting server allowing remote NMS and ssh access from Support to the grid.
⚙	Client Services	FSG	Status of file sharing and heartbeat services.
🛢	Content	CMS	Statistics on grid content ingest, replication, and purging.
🛢	Database	CMS, NMS	Statistics on database type and transactions.
●	DICOM	CLB, CMS	Statistics on DICOM connections and transactions, if DICOM is enabled for the grid.
		LDR	Connectivity settings and statistics for the DICOM interface, if DICOM is enabled for the grid.
❐	Events	ADC	Statistics on the audit messages received and committed, as well as attribute transport and attribute relays.
		AMS, ARC, CLB, CMN, CMS, LDR, FSG	Statistics on the audit messages received and committed, as well as information on attribute transport.
		NMS	Statistics on the audit messages received and committed, attribute transport and the attribute repository, and data connections.
		SSM	Server hardware and driver logs, as well as statistics on audit messages and attribute transport.
▦	Grid Tasks	CMN	Management of grid-wide programmed tasks.
■	HTTP	CLB, LDR	Connectivity settings and statistics for the HTTP interface.
▦	Identifiers	CMN	Total number of unique object identifiers installed and available.

Table 31 Service Components *(continued)*

Icon	Name	Parent Grid Service	Data Provided
▤	Interface Engine	NMS	Information on the NMS MI. This includes notifications status, queue information, and the connection pool.
⛁	Metadata	CMS	Information on metadata replications. Displayed on grids where the CMS uses metadata replication rather than metadata synchronization.
▥	Middleware	ARC	Information about the Middleware that manages the archival media device attached to the ARC.
▤	Object Lookup	CMN	Permits advanced users to look up object metadata, given an object's unique identifier, for advanced troubleshooting.
▥	RAID	SSM	Status of attached RAID hardware and drives.
⇔	Replication	ARC, FSG, LDR	Replication configuration settings and activity.
⛁	Resources	ADC, ARC, AMS, CLB, CMN, CMS, FSG, LDR, NMS, SSM	Information on the hardware and system resources available to the node.
oo	Retrieve	ARC	Information about object requests and cached objects.
⚙	Services	SSM	Information on the state and resource usage of services being monitored. The current operating system installed on the server is listed.
◔	Shares	FSG	Information on data amounts saved to an FSG share.
oo	Store	ARC	Information about objects being written to removable storage.
◔	Storage	LDR, FSG	Statistics on storage used and space available.
⇔	Synchronization	CMS	Statistics on message processing and information synchronization between CMS services.
❐	Tasks	CMS	Information on the status of tasks being performed by the CMS service.

Table 31 Service Components *(continued)*

Icon	Name	Parent Grid Service	Data Provided
🕐	Timing	ADC, ARC, AMS, CLB, CMN, CMS, FSG, LDR, NMS, SSM	Information about local service time and its synchronization with other grid services.
		SSM	Also includes information about NTP status and configuration for the server.
◔	Verification	LDR	Object verification settings and activity.

Accessing Attribute Information

The NMS MI presents information about the wide variety of attributes that are reported for each location, group of servers, grid node, service, and service component in the grid. These attributes and their values form the basis for NMS alarm notifications and reporting, and are used both to monitor normal grid operation and to detect and troubleshoot abnormal conditions.

The exact organization of information within the NMS MI depends upon the design of a particular grid deployment. However, all grids contain summary information about the grid as a whole and about server groupings (such as the group of servers hosted at a single physical location), as well as detailed information about each service and service component hosted on individual grid servers.

Service Overview Attributes

All grid services use a standard set of overview attributes for state, status, and server information. These attributes appear on the Overview tab for the grid service.

Events Component Attributes

All grid services have an Events component which uses a standard set of attributes to report information on the audit messages and attribute values generated by the node, and their progress to relay grid services which forward these messages or values to their final destination.

Resources Component Attributes

All grid services have a Resources component which uses a standard set of attributes to provide information about the computational, storage, and network resources available to the node, as well as giving low-level information about the operation of the service.

Timing Component Attributes

All grid services have a Timing component which uses a standard set of attributes to report on the state of the node's time and the time recorded by neighboring nodes.

Summary Attributes

Summary attributes provide information about:

- the grid as a whole
- physical locations within the grid
- groups of servers at a location

Summary attributes are calculated from the values of attributes for individual nodes and grid services, and provide a convenient synopsis of information about the grid, location, or server group.

The Summary attributes on the Overview > Main tab for the grid, location, or server group give an overview of the storage capacity, metadata capacity, Tape Node storage, and Gateway Node activity for that part of the grid.

The values of summary attributes are based on estimates.

ADC — Administrative Domain Controller

The Administrative Domain Controller (ADC) service authenticates the grid nodes and their connections with each other. For two nodes to connect, the ADC service must have certificates for both nodes. The ADC service also maintains information about grid topology and the location and availability of each grid service. When a node requires information from another node or an action to be performed by another node type, it contacts an ADC service to find the best node to process its request. In addition, the ADC service retains a copy of the grid's configuration bundles, allowing any node to retrieve current configuration information.

To facilitate distributed and islanded operation, each ADC service synchronizes certificates, configuration bundles, and information about grid services and topology with the other ADC services in the grid.

In general, all system nodes maintain a connection to at least one ADC service. This ensures that the nodes are always accessing the latest certificates, bundles, and information. When nodes connect, they cache other nodes' certificates, enabling systems to continue functioning with known nodes even when an ADC service is unavailable. New nodes can only establish connections via an ADC service.

The connection of each node lets the ADC service gather topology information. This node information includes the CPU load, the amount of available disk space (if it has storage), the supported services, and the node's group ID (location). The grid's LDRs, CMSs, and CLBs ask the ADC for topology information through topology queries. The ADC service responds to each query with the latest information received from the grid.

Group IDs are in the form of 10X00Y, where X and Y are pre-assigned based on the site and cabinet number. For example, cabinet B-3 would have a Group ID of 102003. These are used in placing replications under business rules, to permit the grid to disperse data in the most robust and efficient manner possible within the available topology.

ADC Components

The ADC service includes components listed and described in the following sections.

Synchronization

The Synchronization component is used to monitor attributes related to the discovery and monitoring of services and configuration in the grid by the ADC service.

Events

The Events component uses the standard set of events attributes plus other information.

The ADC service acts as both an attribute relay and an audit relay. This means that the ADC service displays information on the transport of attributes from other services to attribute repositories. It also displays information on the transport of audit messages (including audit messages generated by the ADC) to audit repositories.

Resources

The Resources component uses the standard set of resources attributes.

Timing

The Timing component uses the standard set of timing attributes.

AMS — Audit Management System

The Audit Management System (AMS) service logs all audited system events to a text file on the server. The grid uses positive acknowledgement to prevent loss of audit messages. A message remains queued at a service until the AMS service, or an intermediate audit relay service, has acknowledged control of it.

Access to the audit log files is limited to authorized technical support staff unless the customer has purchased the audit option. For those with the audit option, a separate document detailing access and log content is provided.

AMS Components

The AMS service includes components listed and described in the following sections.

Events

The Events component uses the standard set of events attributes.

Resources

The Resources component uses the standard set of resources attributes.

Timing

The Timing component uses the standard set of timing attributes.

ARC—Archive Service

The Archive (ARC) service manages storage and retrieval operations for content stored on devices that use nearline archival media, such as tape.

When the ILM rules for the grid specify that a copy of an object be saved to archival media, the CMS service pushes a copy from a Storage Node to the Tape Node. The ARC service sends these objects to middleware that operates the physical storage device. The storage device then writes the objects to its archival media.

When a client requests an object that must be retrieved from archival media, a Storage Node requests the object from the ARC service. The ARC service requests the object from the middleware, which retrieves the object from the archival media device and sends it to the ARC service. The ARC verifies the object and forwards it to the Storage Node.

The ARC service supports all the standard overview service attributes with some additions for the state and status of the middleware and the archive components.

Supported Archival Media Devices

service Tivoli Storage Manager

A Tivoli Storage Manager (TSM) server provides a logical interface for storing and retrieving data to random or sequential access storage devices, including tape libraries.

An ARC service that uses a TSM to manage archival media displays Tivoli Storage Manager as its Middleware Type on the ARC > Middleware page. The ARC service acts as a client to the TSM server, using the TSM as middleware for communicating with the archival media device. The Middleware Account section of the page displays information about the ARC's account on the TSM server.

Note that the TSM interface has no way to inform the Tape Node when the TSM database or the archival media managed by the TSM is near capacity. The Tape Node continues to accept objects for archiving after the TSM stops accepting new content. This content cannot be written to the TSM managed media. An alarm will be triggered if this happens. Avoid this situation through proactive administration of the TSM.

ARC Components

The ARC service includes components listed and described in the following sections.

Replication

This Replication component provides information about the node to node content replication performed to satisfy grid business rules. The business rules for the grid dictate how many copies of each piece of data are kept, and where these copies are made.

Store

The Store component reports on the process of writing objects to the archival storage device.

Retrieve

The Retrieve component tracks the status of objects requested from the Tape Node. Objects can be requested from an ARC service either as a result of a request from a client for an object, or as a result of a request from a CMS service to replicate objects to another location in the grid to satisfy business rules.

Requests for archived objects are managed to increase the efficiency of retrievals. For example, requests may be ordered such that objects stored in sequential order on tape are requested in that same sequential order. Requests are then queued for submission to the storage device. Depending upon the middleware and archival device, multiple requests for objects on different volumes may be processed simultaneously.

Middleware

The Middleware component provides information on the middleware used to write data to archival media. If the Middleware Type is Tivoli Storage Manager, the component displays information about the middleware and the ARC service's account on the middleware server.

Events

The Events component uses the standard set of events attributes.

Resources

The Resources component uses the standard set of resources attributes.

Timing

The Timing component uses the standard set of timing attributes.

CLB — Connection Load Balancer

The Connection Load Balancer (CLB) service directs incoming content to the optimal storage service (LDR), making its decision using factors such as availability and system load. When the optimal storage service has been chosen, the CLB service establishes an outgoing connection and forwards the traffic to the chosen node.

HTTP and DICOM connections from the grid to an external device use the CLB to act as a proxy, unless the grid is configured to make the connection “Via LDR”. The CLB serves as a connection pipeline between the remote entity and an LDR for DICOM and HTTP.

CLB Components

The CLB service includes components listed and described in the following sections.

DICOM

DICOM is optional.

The DICOM component handles forwarding of DICOM traffic to optimal services and tracks TCP/IP connectivity for DICOM connections. The number of available destinations for query and retrieval
(Q/R) and for ingest via DICOM are reported, as are statistics on connections.

HTTP

The HTTP component handles the forwarding of HTTP session traffic to optimal services and tracks TCP/IP connectivity for HTTP connections. The number of available destinations for query and retrieval (Q/R) and for ingest via HTTP are reported, as are statistics on connections.

Events

The Events component uses the standard set of events attributes.

Resources

The Resources component uses the standard set of resources attributes.

Timing

The Timing component uses the standard set of timing attributes.

CMN — Configuration Management Node

The Configuration Management Node (CMN) service manages grid-wide configurations of connectivity and protocol features needed by all services in the grid. As well, the CMN service is used to run and monitor grid tasks.

The attributes of this service and its components include:

- State and history of Grid Tasks.
- Grid ID number
- Other information related to servicing the grid.

CMN Components

The CMN service includes components listed and described in the following sections.

Grid Tasks

The Grid Tasks component monitors the status of grid-wide programmed maintenance tasks, such as foreground verification of the content on an LDR service. When exceptional maintenance is required—adding new node certificates during grid expansion for example—a special Grid Task program is entered and run. Preparing and running these special Grid Tasks is usually an HP technical support activity. However, as an administrator, you may be asked to initiate grid tasks under the direction of Support. In addition, you may need to initiate grid tasks for Storage Foreground Verification.

The Overview tab for this component displays information about all Grid Tasks, whether user-entered or system-generated. For more information on grid tasks, see Chapter 8, Grid Tasks.

Object Lookup

The Object Lookup component is used to permit maintenance users with Configuration access to request object metadata about any object using its unique identifier (as assigned to the object by the grid). This functionality is included for advanced troubleshooting procedures.

Identifiers

Every object that is saved to the grid is given a unique object identifier by the CMS. Each grid is allocated a range of globally unique object identifiers at the time that the grid is configured. The CMN grants blocks of these object identifiers as needed to each CMS, and tracks how many remain.

Events

The Events component uses the standard set of events attributes.

Resources

The Resources component uses the standard set of resources attributes.

Timing

The Timing component uses the standard set of timing attributes.

CMS — Content Management System

The Content Management System (CMS) service manages object metadata and Information Lifecycle Management (ILM). While the LDRs manage the content, the CMS services provide the logic.

CMS Operation

The CMS service performs two important functions within the grid:

- stores and manages object metadata.
- manages content replication to ensure the Information Lifecycle Management policy for the grid is satisfied.

Object metadata is information related to or describing an object stored in the grid, for instance ingest file path, FSG replication group ID, file modification time, storage location, and so on. Metadata is stored in SQL databases maintained by the CMSs. To ensure redundancy, the grid stores multiple copies of the object metadata in different databases.

The first CMS service to get the object metadata when an object is ingested into the grid is referred to as the owner CMS. The owner CMS manages the object's metadata and replicates it to other CMSs.

In addition to managing metadata, the CMS services manage content replication to ensure that the grid's ILM policy is satisfied. The owner CMS carries out the ILM's instructions for storing objects in the grid over time (where to store the objects, how many copies to store, and for how long). For detailed information on how to configure an ILM policy, see Chapter 7, Information lifecycle mangement.

The CMS service performs metadata management and content management in parallel.

A CMS service becomes read-only when its database is 90% full. A read-only CMS can respond to queries and update information about objects that it tracks, but it cannot accept information about new objects.

You need to monitor the total space available on Control Nodes to ensure that the grid does not run out of storage space for metadata. For information on how to track metadata storage capacity, see Capacity and Expansion (page 174).

Types of CMS

Within HP MAS software, there are two different approaches used to ensure metadata redundancy:

- metadata replication
- metadata synchronization

All CMS services in a grid use one of these two approaches. In a grid where the CMS services use metadata replication, the owner CMS replicates metadata to a subset of CMS services, and applies ILM policies to the data and metadata to make additional copies of metadata as required. In a grid with synchronized CMS databases, the owner CMS synchronizes metadata to every other read-write CMS in the grid.

To determine whether the grid uses metadata replication or metadata synchronization, look at the components under CMS in the NMS MI. A grid that uses metadata replication has the component Metadata. This component does not appear on grids that use metadata synchronization.

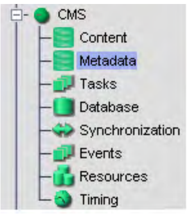

Figure 84 CMS Metdata Component for Metadata Replication

Metadata Replication

In a grid that uses metadata replication, the grid includes one or more CMS replication groups. Each CMS replication group includes two or more CMS services: all CMS services are assigned to a CMS replication group.

Overview of Operation

When an object is ingested into the grid, the owner CMS replicates the object metadata to every CMS service in its CMS replication group. In parallel, it applies the ILM policy and makes any additional required copies of metadata in other groups (that is, in other link cost groups, which usually correspond to locations). When the grid processes a request for an object, it must query at most one CMS service from each CMS replication group to ensure that it finds a record of the object in a CMS database.

In a grid that uses metadata replication, metadata replication within the CMS replication group ensures that there are multiple copies of object metadata and that each CMS database in a replication group has an equivalent database that can be used to restore a failed CMS service. Applying the ILM policy to object metadata ensures that “metadata follows content”, so that a local copy of object metadata is available in the same group (location) as the objects to facilitate data retrieval.

Only the owner CMS replicates metadata to other CMS services within its replication group. Therefore, the CMS services within a CMS replication group may not have an identical set of object metadata. For example, when the owner CMS evaluates the ILM and makes a metadata copy in a different location, the CMS service at that location may be in a different replication group than the owner. In this case, the non-owner CMS service in the second replication group does not replicate metadata to its CMS replication group. Only one additional copy of metadata is made as a result of the operation of the ILM. For an example, see Example: DC+DR+Satellite Site (page 179).

CMS replication groups cannot be used as a method of partitioning metadata. A CMS service at the site where an object is ingested usually becomes the owner CMS, but this is not guaranteed. The grid always selects the “best CMS” to store the initial copy of the metadata based on a number of factors such as availability, location (link cost), and how busy the CMS service is. If the CMS service at the ingest site is very busy or temporarily unavailable, the grid picks a “next best” CMS to be the owner, and this CMS service may be at a different site and in a different CMS replication group than the preferred CMS service.

It is not a supported configuration to place a single CMS in a CMS replication group, or to include a CMS service in the grid design that is not a member of a CMS replication group. Such a CMS immediately transfers ownership to another CMS service elsewhere in the grid, and is not recoverable via the standard CMS recovery procedure in case of a failure.

ILM Evaluation for Metadata

The owner CMS replicates object metadata to the other CMSs in its replication group and evaluates the ILM policy for the grid. The ILM policy consists of one or more rules that describe which objects should be stored in which storage pools at which times (as described in Chapter 7, Information lifecycle mangement).

A storage pool consists of a user-defined storage grade and the group of the storage. (In this context, the “group” of the storage is the link cost group of the LDR, which usually corresponds to its physical location). When placing an object in a storage pool according to the ILM policy, the CMS service applies the same rules to the object metadata as it does to the object. The ILM ensures that there is one copy of metadata on a CMS in each group (location) that contains an object copy. That is, the ILM ensures that metadata follows content.

When deciding whether to make copies of metadata at a location, the ILM first counts the number of metadata copies that might already be there due to metadata replication. For example, if the ILM rule says “Make one copy in the DC and one copy in the DR site” but there is already one copy of metadata at each site because of metadata replication within a CMS replication group, the ILM does not make any additional metadata copies. See the example on Example: DC+DR Grid (page 176) for an illustration.

Note that an ILM may require more than one object copy in a single location, which means that the location must also have more than one copy of the metadata. Because each metadata copy must be made on a different CMS, this places restrictions on the grid design. For example, a second ILM might say “Make one copy in the DR site on SATA storage, and one copy in the DR site on SCSI storage”. Satisfying this ILM policy requires having two copies of the object and two copies of the metadata at the DR site. That is, there must be at least two CMSs at the DR site.

In addition to making object and metadata copies at ingest, the CMS applies the grid’s ILM policy throughout the object’s lifetime. If the ILM policy purges an object copy after a specified time, the corresponding copy of the object metadata is purged and the owner CMS replicates the change to all other CMSs that have a copy of the metadata. If all copies of an object are removed from the grid, all copies of object metadata are also removed.

Note that data and metadata are not purged simultaneously. The owner CMS must retain a copy of the object's metadata until it can confirm that all object copies have been purged. After the owner CMS confirms that all object copies are purged, the object's metadata is queued for purging at the next CMS ILM Evaluation Deferred Time. Because of this behavior, in a grid where objects are being purged the value of CMS > Metadata > Stored Objects may be different on each CMS in a CMS replication group. (In a simple grid configuration that has a single CMS metadata replication group, the numbers of Stored Objects on all CMSs should synchronize a day or two after object purging stops. In a grid with more than one replication group, the effect of purging may be harder to see because the ILM may make copies of metadata for other objects on a single CMS in this replication group.)

Capacity and Expansion

CMS replication groups fill independently of one another; therefore, in general the overall remaining free object capacity of the grid is the sum of the remaining object capacities of the CMS replication groups within the grid[1]. To ensure that the grid does not run out of storage space for metadata, you must monitor the grid's metadata storage capacity as described in Metadata Storage Capacity (page 129).

An individual CMS becomes read-only when its database reaches 90 % of its total capacity. This leaves a margin of database space to handle updates to object metadata for objects owned by that CMS, as well as space to replicate object metadata from other CMSs in the same CMS replication group (in the case that they have accepted ownership of objects but have not yet replicated their metadata within the group). When a CMS becomes read-only, all of the CMSs in its replication group turn off their "content ingest service". The read-write CMSs are still available to make copies of metadata as required to meet the grid's ILM rules (or to replicate metadata within the CMS replication group). But read-only CMSs and CMSs that have turned off their content ingest service do not accept ownership of new objects.

If you ignored capacity alarms in the NMS MI and permitted a CMS database in each CMS replication group to fill until they became read-only, the grid would no longer be able to ingest files. To prevent an interruption of service, add additional metadata capacity before the first CMS fills up. The metadata capacity of a grid is increased by adding additional CMS replication groups.

In some cases, metadata capacity may be freed by purging objects from the grid. When an object is purged, its metadata is also purged. Therefore, if the ILM policy systematically purges objects from the grid (for example, when objects are removed after a specified time) the CMS can reclaim the metadata capacity and make it available for the ingest of new objects. However, when only a few objects are purged (for example, if the only object that is purged is the FSG backup), the CMS cannot reclaim the metadata capacity for reuse due to database fragmentation.

[1] This may not be true if the ILM requires additional copies of metadata. For example, if a grid has two CMS replication groups and the ILM requires metadata copies in each group, the overall object capacity of the grid is reduced.

After a CMS becomes read-only, it remains read-only unless it is able to reclaim database space due to metadata purging. If it is able to free a total of 20 % of its database space, the CMS becomes read-write once more. The CMS replication group remains read-only until this change is replicated to the other CMSs in the same group and all CMSs can turn on their content ingest service.

Edge Metadata Cache (Satellite Sites)

A grid can be designed with one or more “edge metadata caches”, or “satellite sites” designed to provide local access to a subset of the information that has been ingested at a site (usually the most recently ingested data), and to provide islanded operation in case of network failure. Data from all sites is aggregated at a central Data Center and/or Disaster Recovery site, and older objects and their metadata are purged from the satellite sites after a fixed amount of time has passed. Satellite sites are designed to provide only temporary data and metadata storage at the satellite site.

When an object is ingested at the satellite site, a copy of the object is generally made on a local LDR and a copy of the metadata is made on a local CMS, which becomes the owner. The owner CMS replicates the object metadata to the other CMSs in its CMS replication group and evaluates the ILM policy to ensure that object copies and metadata copies are made as required to meet the ILM policy.

If the grid design includes two CMSs and two LDRs at the satellite site, the grid has redundant copies of objects ingested at that site (and their metadata) even when it is isolated from the rest of the grid. If the satellite site includes a single CMS and LDR, data ingested at the satellite site while it is islanded is at risk of loss if either the CMS or LDR fails while the site is islanded.

During “islanded operation” clients can store and retrieve data at the satellite site while it is disconnected from the rest of the grid. When the site is reconnected, the grid automatically synchronizes data and metadata, making additional copies as required to meet ILM rules and to satisfy the requirements of metadata replication.

In a grid deployment where clients are most likely to retrieve recently ingested data, a satellite site offers better ingest and retrieval performance for clients at the satellite site by removing WAN latency for most grid operations, in addition to the added operational redundancy of being able to operate while disconnected from the rest of the grid.

To create a satellite site that acts as an edge metadata cache, a grid is designed with:

- CMSs that use metadata replication.
- local storage for temporary copies of data at the satellite site, and a gateway to support file ingest.
- one or two CMSs at the satellite site. These CMSs store temporary local copies of metadata.
- a CMS at the DC or DR site that is a member of the same CMS replication group as the CMS(s) at the satellite site. This CMS provides metadata redundancy, and can be used as a recovery source if the CMS at the satellite site should fail.

- a link cost group that includes the satellite CMSs and satellite LDRs (so that a storage pool can be created for use by the ILM policy).
- a link cost group that includes only the "satellite" CMS at the DC or DR site. (That is, the CMS at the DC or DR site does not have the same group ID as other grid nodes at the DC or DR site, or the same group ID as the CMSs at the satellite site.) This is required to avoid issues with metadata replication in various edge cases, as discussed in Example: DC+DR+Satellite Site (page 179).
- The link costs to this group can be set to any value, as this group is not referenced in the ILM policy.
- an ILM policy that keeps temporary copies of data and metadata at the satellite site, purging them after a fixed amount of time.

For a detailed example that discusses how metadata is replicated in a grid with a satellite site, see Example: DC+DR+Satellite Site (page 179).

If the grid includes more than one satellite site, there will be one "satellite" CMS service at the DC or DR site for each remote site. It is possible to place all such "satellite" CMSs at the DC or DR sites in the same group, instead of putting each one in its own group. For more information, see the Example: DC + Two Satellite Sites (page 181).

CMS Recovery

In a grid that uses metadata replication, when a Control Node fails you recover it by reinstalling the server and cloning its database from an equivalent CMS in the same CMS replication group. You can find a list of potential recovery sources in the NMS MI under CMS > Tasks.

Recovery for CMSs at a satellite site follows the same procedure as for DC or DR CMSs.

Metadata Replication in Action: Examples

The following examples illustrate how metadata replication works in a variety of grid scenarios.

Example: DC+DR Grid

The first example is a DC+DR grid that uses metadata replication and has one CMS replication group. The ILM policy states that for each object, one copy is stored at the DC site and one copy is stored at the DR site.

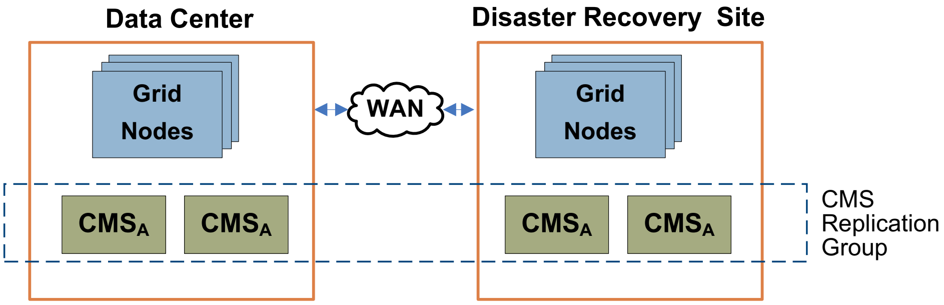

Figure 85 DC + DR Grid With 1 CMS Replication Group

When an object is stored to the DC site, a CMS service at that site becomes the owner. The owner CMS evaluates the ILM policy and makes object copies and metadata copies at the DC and the DR. The owner CMS also replicates the object metadata to every other CMS service in its replication group. In this case, the owner makes two additional copies of the object metadata for a total of four copies, one on each CMS in the grid.

Example: Grid with 2 CMS Replication Groups

The DC+DR grid in this example also uses metadata replication and has the same ILM policy as the previous example: for each object, create one copy at the DC site and one copy at the DR site. However, the customer anticipates ingesting many small objects over the first year that the grid is in place, so it was deployed with two CMS replication groups to increase the grid's object capacity.

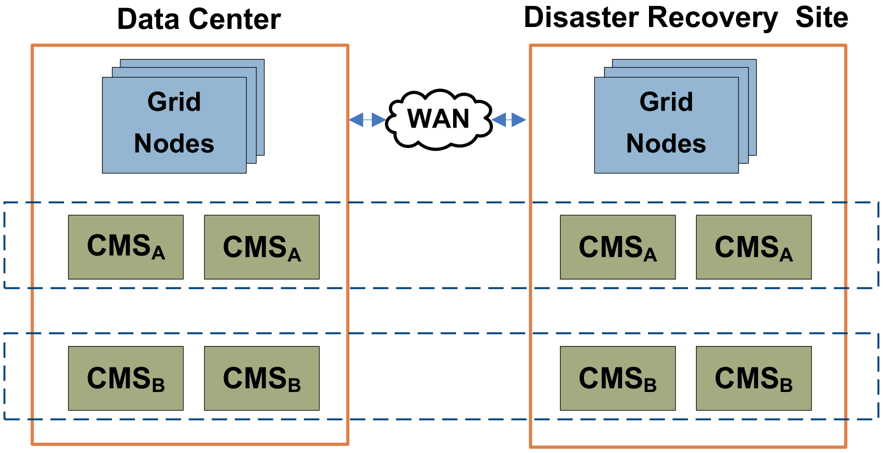

Figure 86 DC + DR Grid With 2 CMS Replication Groups

When an object is stored to the DC site, a CMS service at that site (for example, a CMS service in CMS replication group A) becomes the owner. The owner CMS replicates the object metadata to every other CMS in CMS replication group A, making three additional copies of the object metadata (for a total of four copies). No metadata copies are made in CMS replication group B. When the owner CMS evaluates the ILM policy, it makes the required object copies at the DC and DR sites. When it evaluates the ILM for the metadata, it finds that there is already at least one copy of the metadata at each required location, so no additional copies of the object metadata are made. The CMS replication groups fill independently of one another.

Example: DC+DR+ARC

In this example, a DC+DR grid is configured to use metadata replication and has a Tape Node at the DR site. The ILM policy states that at ingest, one copy is stored on disk at the DC site, one copy is stored on disk at the DR site, and one copy is stored on tape at the DR. After one year, the copy on disk at the DR site is purged.

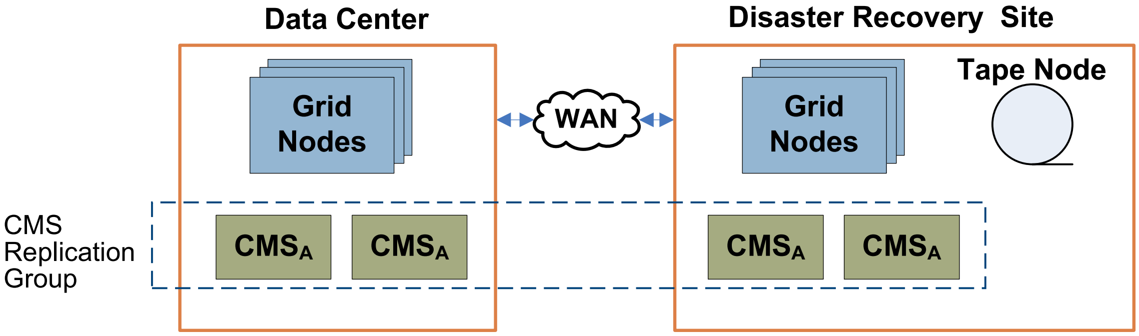

Figure 87 DC+DR+Tape Node Grid

When an object is ingested at the DC site, a CMS service at the DC site becomes the owner and evaluates the ILM. In this grid, the disk storage at the DC forms one storage pool, the disk storage at the DR forms a second storage pool, and the Tape Node at the DR site forms a third storage pool. At ingest, the ILM makes one object copy in each storage pool, and one metadata copy in the same group (location) as each storage pool. This means that one metadata copy is made at the DC site and two metadata copies are made at the DR site. When the owner applies metadata replication, it replicates the object metadata to every other CMS in its CMS replication group, which means that one additional copy of object metadata is made at the DC site. (No additional metadata copies are made at the DR site, as both DR CMSs already have a copy of the metadata due to the application of the ILM policy.) The end result is four copies of metadata: two at the DC site and two at the DR site.

After a year, the ILM policy is re-evaluated. One disk copy of the object at the DR site is purged, while the copy of the object at the DC and the tape copy at the DR are retained. The ILM dictates that a copy of the metadata at the DC site be retained, while a copy of the metadata at the DR site can be purged. However, because all CMSs at the DR site are in the same CMS replication group as the DC CMSs, the owner CMS preserves both copies of metadata at the DR site due to metadata replication.

If the ILM then went on to say “after five years, retain one copy on tape”, after five years the ILM would be re-evaluated again. If the owner CMS is at the DC site, the DC CMS purges the disk copy of the object at the DC site and realizes that its copy of metadata can also be purged. Therefore, it transfers ownership to a

CMS at the DR site. However, because all CMSs in the grid are in the same replication group, the new owner CMS retains all four copies of metadata in the grid.

Example: DC+DR+Satellite Site

In this example, a DC+DR+Satellite Site grid is configured to use metadata replication. The ILM policy states that for each object, one copy is stored at the DC site and one copy is stored at the DR site. In addition, for each object ingested at the satellite site, a copy is saved at that site for 30 days, and then is purged.

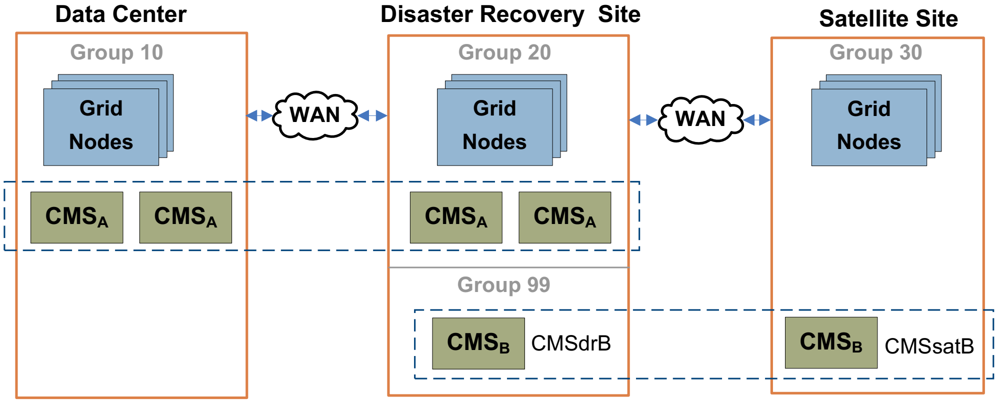

Figure 88 DC+DR+Satellite Grid

As described in Edge Metadata Cache (Satellite Sites) (page 175), this grid requires the following configuration to function correctly:

- CMS services that use metadata replication.
- local storage at each satellite site for temporary copies of data.
- one or two CMS services at each satellite site. These CMSs store temporary local copies of metadata and provide local access to object metadata.
- for each satellite site, a CMS service at the DC site that is a member of the same CMS replication group as the CMS service at the satellite site. This CMS service provides metadata redundancy, and can be used as a recovery source if the CMS service at the satellite site should fail.
- a (link cost) group that includes the satellite CMS services and satellite LDRs (so that a storage pool can be created for use by the ILM policy).
- an ILM policy that keeps temporary copies of data and metadata at the satellite site, purging them after a fixed amount of time.
- a (link cost) group that includes only the “satellite” CMSs that are located at the DC site.

When an object is ingested at the satellite site, the CMS at the satellite site becomes the owner. It replicates the object’s metadata to the other CMS service in CMS replication group B (CMSdrB, at the DR site) and evaluates the ILM policy. The ILM policy states that one copy must exist at each of the DC and DR sites, with a copy at the satellite site if it is less than 30 days since ingest. The owner CMS makes object copies at the DC, DR, and at the satellite site, and looks at the metadata locations.

A metadata copy already exists at the Satellite site. So, the CMS makes a metadata copy on a CMS at the DC site and a metadata copy at the DR site on a CMS in replication group A. CMSdrB has a copy of the metadata already, but while CMSdrB is physically located at the DR site, it is logically in its own “group” (link cost group) that is not referenced in the ILM. Thus, CMSdrB is not a valid metadata location for the DR copy. The end result is a total of four copies of metadata: two copies in CMS replication group B, one copy at the DC site in CMS replication group A, and one copy at the DR site in CMS replication group A.

After 30 days, the ILM policy is re-evaluated. The owner CMS at the satellite site realizes that it should purge the object copy at the satellite site and its own copy of the object metadata, so it transfers ownership to a CMS in replication group A that already has a copy of the metadata. The new owner CMS makes copies on all of the remaining CMSs in replication group A at the DC and DR sites. It also evaluates the ILM, and purges the metadata copy on the CMS at the satellite site and its replication group copy on CMSdrB. In total, there are four copies of the object metadata, one on each CMS in replication group A.

What happens if the CMS at the satellite site is temporarily unavailable when an object is ingested? The grid sends the object metadata to another CMS service in the grid which becomes the owner. There are two possibilities as to what happens next:

- The grid selects a CMS in CMS replication group A to be the owner; for example, a CMS at the DC site. This CMS makes metadata copies on every other CMS in CMS replication group A via metadata replication and evaluates the ILM policy for the grid, which requires a metadata copy on the CMS at the satellite site. This metadata copy is made when the CMS at the satellite site becomes available. The end result is five copies of object metadata for the first 30 days, and four copies afterwards (when the satellite site copy is purged).
- The grid selects the CMS in replication group B at the DR site to be the owner. CMSdrB evaluates the ILM and makes a copy of the metadata at the DC site and a copy of metadata on a CMS in replication group A at the DR site. When the satellite site’s CMS service becomes available, the owner makes a copy on the CMS at the satellite site. Because CMSdrB is not a valid location according to the ILM policy, CMSdrB transfers ownership to another CMS as soon as another metadata location becomes available. In practice, this means that CMSdrB will usually transfer ownership to a CMS in replication group A at the DC or DR.

  After 30 days, the owner CMS transfers ownership of the object to a CMS in replication group A (if the owner is not in replication group A already), and the metadata copies in replication group B are purged, leaving 4 copies total of metadata in the grid.

Why must the CMSdrB be in its own (link cost) group, and not be placed in either the same group as the other CMSs at the DR site, or in the same group as the other satellite CMS? The short answer is that metadata replication will not function as expected in some edge cases. The long answer is that unless CMSdrB is in its own group you may encounter one of these two problems:

- **Satellite Purge Problem**—If CMSdrB is placed in the same group as the other CMSs at the DR site, the satellite purge problem may occur. When objects are ingested at the satellite site, the satellite CMS (CMSsatB) becomes the owner and immediately makes a copy on CMSdrB. When CMSsatB evaluates the ILM policy, it makes a copy of the metadata on one CMS at the DC site. However, there is already a copy of the metadata at the DR site (on CMSdrB), so no further metadata copies are made at the DR site. There are three copies of metadata in the grid.

  After 30 days, the CMSsatB notes that its copy of metadata should be purged and transfers ownership to a CMS at the DC or DR site. If it transfers ownership to a CMS at the DC site, that CMS purges the copy of metadata at the satellite site and makes additional copies on the remaining CMSs in replication group A. However, if CMSsatB transfers ownership to CMSdrB, metadata replication requires that the CMS at the satellite site retain a copy of the metadata. Therefore, the copy of the metadata at the satellite site is not purged and space is not freed on this CMS for new objects ingested at the satellite site. Eventually, both CMSsatB and CMSdrB will fill up.

- **Satellite Copy Problem**—If CMSdrB is physically placed at the DR site but is logically placed in the same group at the CMS at the satellite site, you do not encounter the satellite purge problem. However, you may encounter the satellite copy problem. If the CMS at the satellite site is unavailable when an object is ingested, another CMS assumes ownership. For example, a CMS at the DC site could become the owner and make copies on all other CMSs at the DC and DR sites that are in CMS replication group A. When the owner CMS evaluates the ILM policy, it realizes that it needs to make a copy on a CMS at the satellite site. However, both CMSs in CMS replication group B are logically “at the satellite site” because they are in the same group. If the owner CMS makes a copy on CMSdrB, the ILM is satisfied and no copy of the metadata is made at the remote satellite location. Therefore, the satellite site cannot operate effectively while it is islanded because it does not contain a copy of the metadata for all objects ingested in the last 30 days.

Placing CMSdrB in its own (link cost) group prevents both the satellite purge and the satellite copy problems.

Example: DC + Two Satellite Sites

In this example, the grid uses metadata replication and we have a DC site and two satellite sites. The ILM policy states that for each object, two copies are stored at the DC site. For each object ingested at a satellite site, a copy is saved at that satellite site for 30 days and then is purged.

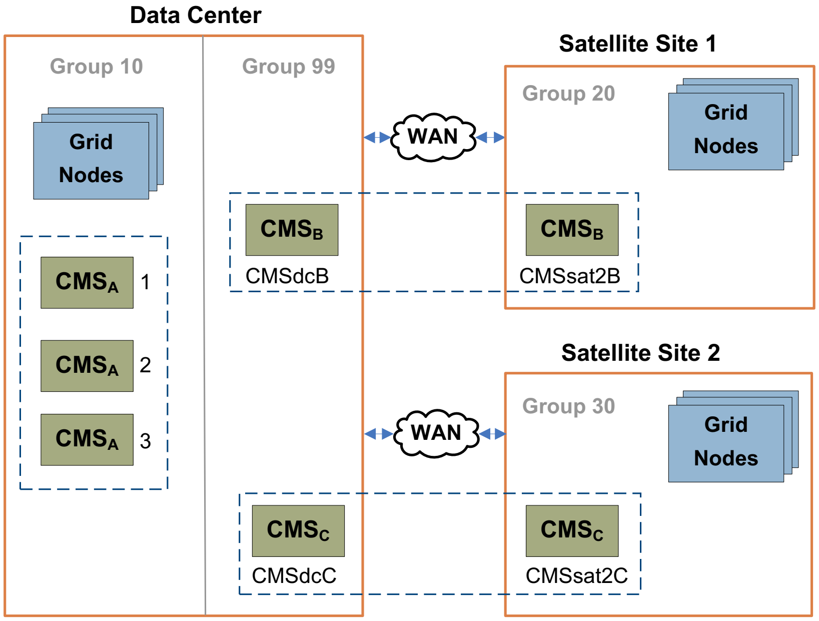

Figure 89 Data Center with Two Satellite Sites

As with the DC+DR+Satellite grid described in the previous example, this grid requires the following configuration to function correctly:

- CMS services that use metadata replication.
- local storage at each satellite site for temporary copies of data.
- one or two CMS services at each satellite site. These CMSs store temporary local copies of metadata and provide local access to object metadata.
- for each satellite site, a CMS service at the DC site that is a member of the same CMS replication group as the CMS service at the satellite site. This CMS service provides metadata redundancy, and can be used as a recovery source if the CMS service at the satellite site should fail.
- a (link cost) group that includes the satellite CMS services and satellite LDRs (so that a storage pool can be created for use by the ILM policy).
- an ILM policy that keeps temporary copies of data and metadata at the satellite site, purging them after a fixed amount of time.
- a (link cost) group that includes only the “satellite” CMSs that are located at the DC site.

That is, the group IDs and the CMS replication groups for each CMS are configured as shown in Table 32.

Table 32 Example Configuration for a DC+2 Satellite Site Grid

Name	Location	Group ID	CMS Replication Group
CMSdcA1	Data Center	10	A
CMSdcA2	Data Center	10	A
CMSdcA3	Data Center	10	A
Other grid nodes	Data Center	10	Not applicable
CMSdcB	Data Center	99	B
CMSsat1B	Satellite Site 1	20	B
Other grid nodes	Satellite Site 1	20	Not applicable
CMSdcC	Data Center	99	C
CMSsat2C	Satellite Site 2	30	C
Other grid nodes	Satellite Site 2	30	Not applicable

It is equally valid to configure each “satellite” CMS at the DC to have its own unique Group ID. (For example, CMSdcB could be in group 88, while CMSdcC could be in group 99.) As long as the group or groups for the “satellite” CMSs at the DC are not referenced in the ILM policy, these two configurations are equivalent.

How objects and object metadata are replicated within this grid is very similar to how they are replicated within the DC+DR+Satellite Site grid.

In a grid with two satellite sites, each satellite site operates independently. That is, according to the ILM policy for the grid, for each object ingested at satellite site 2, two copies of the object and its metadata are permanently stored at the DC site and a temporary copy of the object and its metadata is kept at satellite site 2 for 30 days. For each object ingested at satellite site 1, two copies of the object and its metadata are permanently stored at the DC site, and a temporary copy of the object and its metadata is kept at satellite site 1 for 30 days.

However, you cannot exclude the possibility that metadata from objects ingested at one satellite site might be stored temporarily on the CMS at another satellite site. If CMSsat2C is temporarily unavailable when an object is ingested at Satellite Site 2, it is possible that CMSsat1B might assume ownership of the object. In this case, when CMSsat1B evaluates the ILM policy it will realize that it is not a valid location for the metadata (only the DC site and satellite site 2 are valid locations). Therefore CMSsat2C will transfer ownership to another CMS that has a metadata copy (likely a CMS in CMS replication group A at the DC site). The new owner purges the metadata copy on CMSsat1B.

You can reduce the likelihood of one satellite CMS assuming ownership of an object ingested at another satellite CMS by setting a high link cost between the two satellite sites (see Link Costs (page 161) for details).

Example: Grid Expansion DC+DR

As in the first example, at installation our grid is a DC+DR with one CMS replication group, and the CMSs use metadata replication. The ILM policy states that for each object, one copy is stored at the DC site and one copy is stored at the DR site.

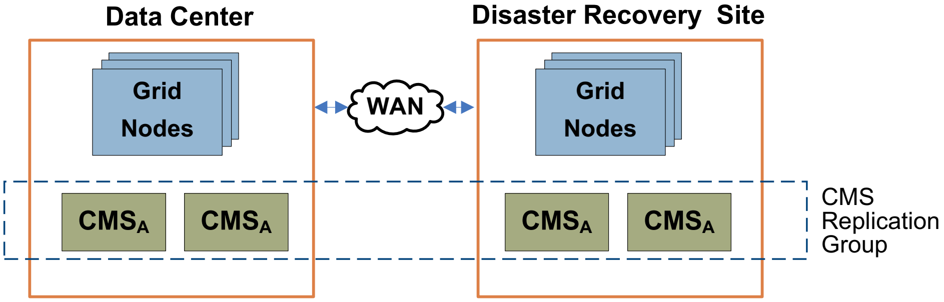

Figure 90 Pre-expansion Grid: DC+DR, 1 CMS Replication Group

In this example, when the metadata capacity of the grid reaches 80 % full the grid administrator plans a capacity expansion. Because the ingested objects are small, sufficient object capacity remains for some period of time at present data rates. So the grid administrator expands the grid to add only metadata capacity by adding an additional CMS replication group, but does not add additional Storage Nodes.

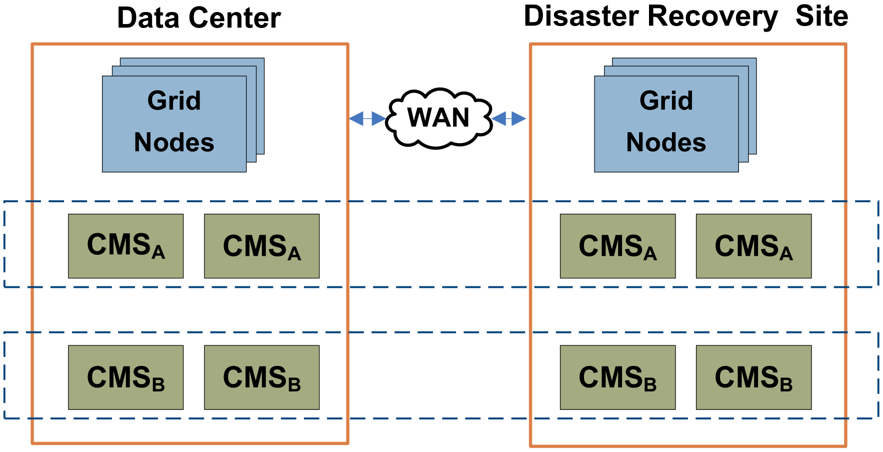

Figure 91 DC+DR Grid After Expansion: 2 CMS Replication Groups

After the expansion is complete, the grid includes two CMS replication groups, and works just as it would if the grid had been originally installed that way. (For more information, see Example: Grid with 2 CMS Replication Groups (page 177).) Metadata copies are made in CMS replication group A or CMS replication group B, depending on which CMS initially assumes ownership of a newly ingested object. Object copies are made on the existing grid storage.

If CMS replication group A was installed first, when the first CMS in CMS replication group A fills up, that replication group turns off its content ingest service. After this time, all content ingest is handled by CMSs in CMS Replication group B.

Example: Expansion DC+DR+Satellite Site

As in the example shown in Example: DC+DR+Satellite Site (page 179), at installation we have a DC+DR+Satellite site grid that uses metadata replication and has two CMS replication groups. CMS Replication group A has four CMSs: two at the DC and two at the DR. CMS replication group B supports the operation of an edge metadata cache, and has one CMS at the satellite site and one CMS at the DR site. CMSdrB is in its own group. The ILM policy states that at ingest at the satellite site, one copy is stored in each location (DC, DR, and Satellite). After 30 days, the copy at the satellite site is purged.

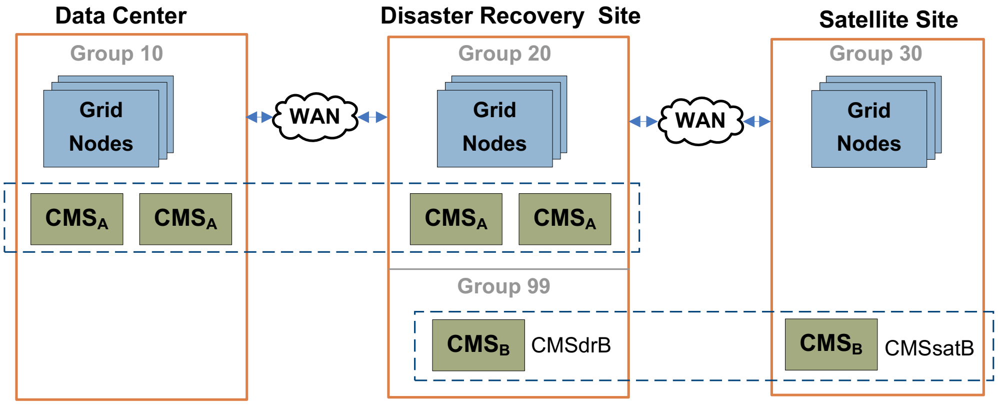

Figure 92 Pre-expansion: DC+DR+Satellite Site

In this example the grid begins to fill up after it has been in service for two years, so the grid administrator plans a capacity expansion that adds both additional storage and additional metadata capacity (an additional CMS replication group at the DC+DR sites, as well as additional Storage Nodes that are not shown in Figure 93).

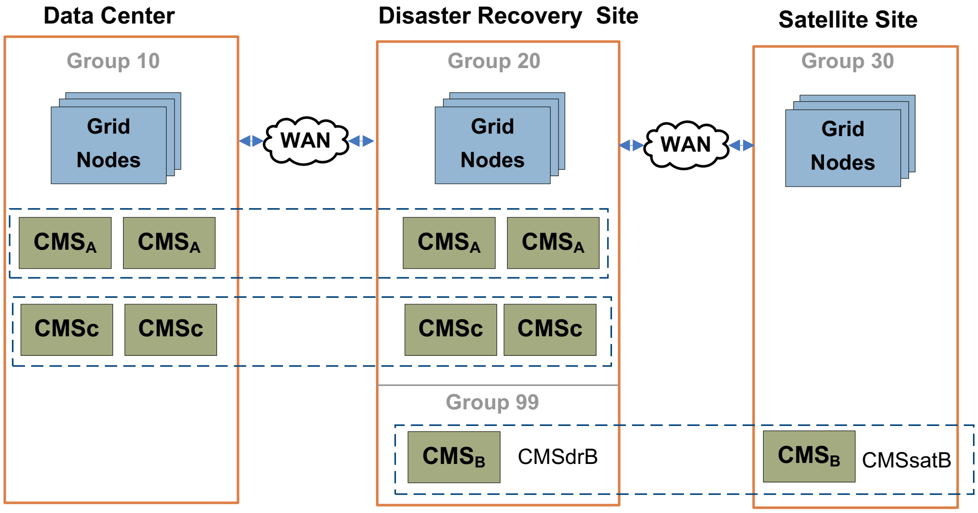

Figure 93 After Expansion: DC+DR+Satellite Site

After the expansion is complete, the grid includes two CMS replication groups at the DC+DR sites: CMS replication group A and C. No additional CMSs are needed to support the satellite site. (This is in contrast to the situation for CMSs that use metadata synchronization as described in Example: DC+DR+Satellite Grid (page 192).) The operation of the grid remains the same as it did before expansion, except that metadata copies for the DC and DR site may be made in either CMS replication group A or CMS replication group C. Object copies are made either on the new or the existing grid storage while the existing grid storage still has free storage space.

When the first CMS in CMS replication group A fills, that replication group turns off its content ingest service. After this time, permanent copies of metadata at the DC and DR site are created on CMSs in CMS Replication group C. Temporary copies of metadata for objects ingested at the satellite site continue to be made in CMS replication group B and are purged after 30 days.

Example: Read-only CMS Becomes Read-Write

In this example, we have a grid that uses metadata replication and has a single Data Center site. It has one CMS replication group, and an ILM that says “Keep three copies of all objects for three years. After three years, purge content that was ingested to file paths that begin with `/fsg/Legacy/Temp`”.

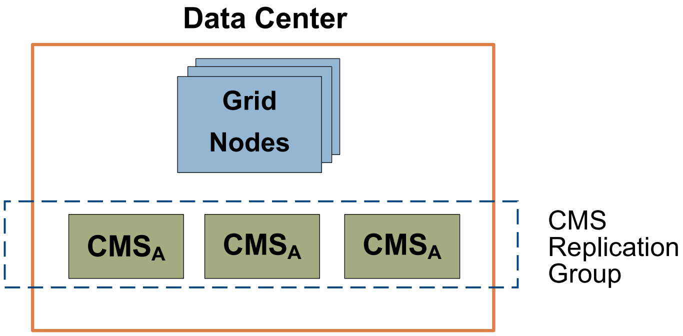

Figure 94 Data Center Grid

During the first year that the grid is in place, a large quantity of legacy data is migrated into the grid, mostly to subdirectories of the file path /fsg/Legacy/Temp. Because of the large amount of legacy data, the CMSs in replication group A begin to fill up shortly after the grid’s first birthday. Before the first CMS goes read-only, the grid is expanded to add a second CMS replication group, as shown in Figure 95.

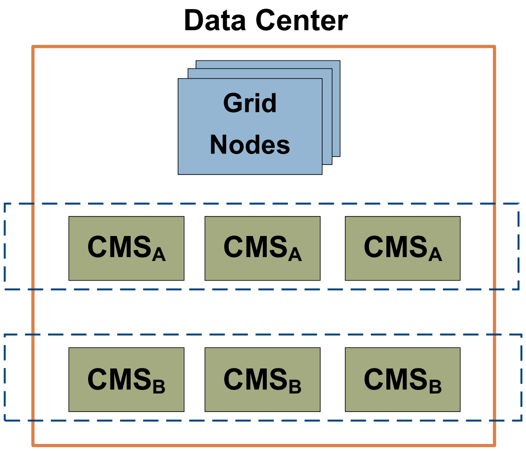

Figure 95 Data Center Grid with 2 CMS Replication Groups

After the first CMS in replication group A goes read-only, content ingest is handled by CMSs in CMS replication group B, which now assume ownership of all new content.

Three years after the grid was first installed, the ILM begins purging legacy content (and its associated metadata) that was ingested to sub-directories of /fsg/Legacy/Temp. All of the metadata for this purged content is owned by CMSs in CMS replication group A, as the legacy data migration completed before the grid was expanded to add CMS replication group B. Because large amounts of data and metadata are systematically being purged from the grid, blocks of database space are freed. Some time after purging begins, the CMSs in CMS replication group A free 20 % of their total database space, and become read-write again.

After the CMSs in CMS replication group A become read-write, any CMS in the grid (in either replication group A or replication group B) may assume ownership of new content. As new content is ingested, the two CMS replication groups proceed to fill independently of one another.

Metadata Synchronization

In a grid that uses metadata synchronization, the owner CMS synchronizes its metadata to all other read-write CMSs in the grid.

The following grid types use CMSs that synchronize metadata:

- Any grid that uses a non-configurable ILM.

  —or—

- Any grid where deduplication is enabled.

  —or—

- Any grid that was originally installed with metadata synchronization, that has not yet been converted to use metadata replication. These grids include:
  - Grids originally installed with a HP MAS software release that did not support metadata replication (Release 7.5.x or earlier)

— Grids originally installed with Release 8.0software using metadata synchronization.

Simply updating a grid to Release 8.1 software does not automatically change the type of metadata management that it uses. You must perform a conversion to metadata replication.

Overview of Operation

When an object is ingested into the grid, the owner CMS synchronizes the object metadata to every other CMS in the grid. The synchronization of metadata ensures that there are multiple copies of object metadata in the grid, and that more than one CMS database in the grid has an equivalent database that can be used to restore a failed CMS.

Grids that use metadata synchronization do not necessarily have identical CMS databases. It depends on whether the grid was expanded to add more Control Nodes, and on the reason for the expansion:

- If Control Nodes were added to increase redundancy, their CMS databases were cloned from the existing CMSs at the time the new grid nodes were added. In this case, all CMSs remain equivalent.
- If the Control Nodes were added to increase object capacity, they were added with empty CMS databases. In this case, all Control Nodes are not equivalent. The set of CMSs added at the same time form a single “generation” and are equivalent to each other, as are the original set of CMSs in the grid.

When an object is retrieved, to ensure that it finds an object location the grid must query a CMS from each CMS generation.

Because each CMS must synchronize metadata to every other CMS in the grid, during normal operations it is possible for the ingest load to exceed the rate at which CMSs can synchronize metadata. (For more information, see Chapter 4, Operations.)

Capacity and Expansion

In a grid that uses metadata synchronization, the overall remaining free metadata capacity for the grid is the same as the remaining capacity of the read-write CMSs. To ensure that the grid does not run out of storage space for metadata, you must monitor the grid’s remaining metadata storage capacity as described in Chapter 4, Operations.

When an object is purged from a grid that uses metadata synchronization, the object’s metadata is not purged. Therefore, purging objects does not free database space, and does not increase the grid’s capacity to ingest new objects. After a CMS becomes read-only, it remains read-only.

An individual CMS becomes read-only when its database reaches 90 % of its total capacity. This leaves a margin of database space to handle updates to object metadata for objects owned by that CMS. Because all read-write CMSs in a grid synchronize metadata, all CMSs fill at approximately the same time. So when one CMS becomes read-only, the remaining read-write CMSs in the grid become read-only shortly afterwards.

The metadata capacity of a grid is increased by adding an additional “generation” of CMSs with empty databases. If all CMS databases in the grid are full, the grid cannot ingest files. To prevent an interruption of grid service you *must* expand the grid before the first CMS database in the first generation fills up. However, because the new CMSs synchronize with all old and new read-write CMSs, “extra” copies of metadata are made while both generations of CMS remain read-write. To prevent wasting database space, the new CMSs should be added only a short time before the previous generation fills up.

⚠ WARNING! Do not ignore alarms on metadata storage capacity and let CMSs that use metadata synchronization fill up before expanding the grid. If you do, some metadata may be permanently lost if a single CMS fails. Associated objects then become unretrievable.

However, you *must* add the new generation of CMSs before the *first* CMS in the first generation fills up to prevent metadata loss in event of a failure. This potential loss of metadata occurs because the CMS databases do not fill at exactly the same time. After the first CMS in the grid becomes read-only, the remaining CMSs in the same generation continue to accept ownership of newly ingested objects until they fill up. And while the read-write CMSs continue to send synchronization messages to the read-only CMSs, the read-only CMSs cannot add the information to their databases. Therefore, if a new generation of CMSs is *not* present, the last CMS to go read-only is the only CMS to have information about the most recently-ingested objects. If this CMS subsequently fails and is recovered via cloning from another CMS database, all information about the last objects to be ingested is lost forever.

This situation does not occur if a new generation of CMSs is added to the grid before the first CMS fills up. As CMSs from the first generation go read-only, the entire generation of new CMSs remains read-write. After the last CMS of the first generation goes read-only, the new generation of CMSs accept ownership of all new content and make copies of the metadata for all ingested objects. Throughout the overlap period, there are multiple copies of metadata made for every object ingested.

This prevents the loss of metadata in event of a failure of any CMS in the grid. For example, if the last CMS to go read-only from the first generation fails, during recovery its database is cloned from another CMS in the first generation. By definition, this CMS has less metadata than the lost CMS so that after recovery the recovered CMS contains less metadata than it had before the failure. But nothing is lost from the grid: all new CMSs have copies of the metadata for these objects, which were ingested during the overlap period. So while the grid has lost an “extra” copy of metadata for some objects, the grid continues to include multiple copies of metadata for every ingested object.

CMS at Satellite Sites

It is possible to include a satellite site in a grid that uses metadata synchronization, and for that satellite site to provide islanded operation. However, in a grid with metadata synchronization, the CMSs at satellite sites are equivalent to CMSs at the Data Center and/or Disaster Recovery sites. Therefore, the satellite site holds permanent copies of object metadata for all objects in the entire grid, and the satellite site cannot be configured to function as an “edge metadata cache” that flushes older data, as it can when the grid uses metadata replication (as described in Edge Metadata Cache (Satellite Sites) (page 175)).

When the first generation of CMSs in the grid fill up, additional CMSs must also be added at the satellite site to enable access to recently ingested data. Therefore in a grid that uses metadata synchronization, the CMSs at satellite sites must be sized the same as DC or DR CMSs, and over time the number of CMSs at each satellite site will grow making this a more hardware-intensive solution. For an example, see Example: DC+DR+Satellite Grid (page 192).

CMS Recovery

In a grid that uses metadata synchronization, when a Control Node fails you recover it by reinstalling the server and cloning its database from an equivalent CMS of the same generation as the failed Control Node. However, the grid does not record which CMSs are equivalent. You must track the history of a grid and its expansions and use this history to find an equivalent CMS database for the cloning operation.

Metadata Synchronization in Action: Examples

The following examples illustrate how metadata synchronization works in a variety of grid scenarios.

Example: DC+DR Grid

In the first example, we have a DC+DR grid that uses metadata synchronization. The ILM policy states that for each object, one copy is stored at the DC site and one copy is stored at the DR site.

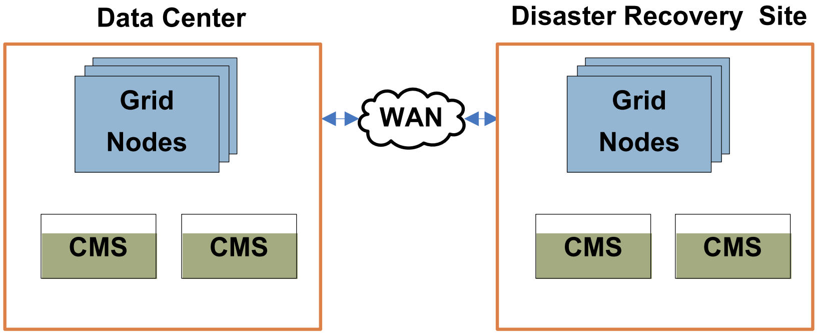

Figure 96 DC+DR Grid with Metadata Synchronization

When an object is stored to the DC site, a CMS service at that site becomes the owner. The owner CMS synchronizes object metadata to every other CMS service in the grid, making three additional copies of the object metadata for a total of four copies.

In the diagram, coloring of the CMS boxes represents the amount of metadata stored in each CMS. As you can see in Figure 96, the databases on all CMSs are partially full, and all CMSs have an equivalent set of metadata (same color).

Example: Grid Expansion

As in the first example, the grid is a DC+DR that uses metadata synchronization. It has an ILM policy that states that for each object, one copy is stored at the DC site and one copy is stored at the DR site.

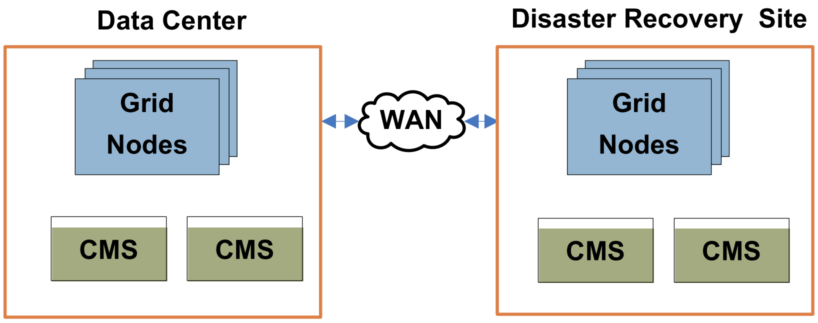

Figure 97 Metadata Synchronization: Before Expansion

After the grid has been in service for some time, the available metadata capacity of the grid begins to run low (as shown in Figure 97), so the grid administrator plans a capacity expansion. Because the ingested objects are small, sufficient object capacity remains for some time at present data rates. So the grid administrator expands the grid to add metadata capacity by adding additional CMSs with empty databases, but does not add additional Storage Nodes.

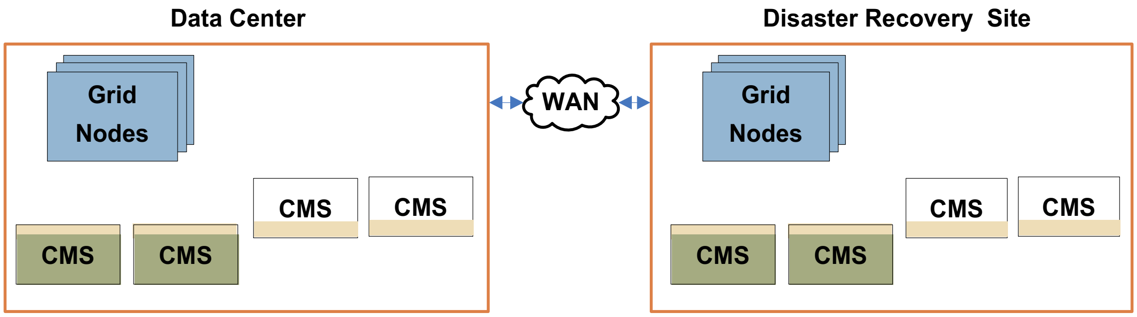

Figure 98 Metadata Synchronization: Grid After Expansion

After the expansion is complete, the grid includes an additional four CMSs. The new CMSs were added with empty databases before the original CMSs were full. While the original set of CMSs are still read-write, both the original CMSs and the expansion CMSs synchronize all metadata to one another, creating an overlap of content between the original and expansion CMSs. After the original set of CMSs become read-only, the expansion CMSs remain read-write and continue to synchronize object metadata.

This is displayed in Figure 98. The original CMSs contain the same set of metadata (shown in green). Before they fill completely, expansion CMSs are added to the grid. While the original CMSs remain read-write, all CMSs synchronize metadata (shown in gold), so there is overlap between the content of the original and expansion CMSs. After the original CMSs fill, the expansion CMSs continue to accumulate (gold) metadata.

Object copies are made on the existing grid storage both before and after CMSs are added to the grid.

Example: DC+DR+Satellite Grid

In this example, a grid that uses metadata synchronization has three sites: a Data Center, a Disaster Recovery Site, and a Satellite site, as shown in Figure 99. When the grid is first installed, there are a total of five CMSs in the grid: two at the DC, two at the DR, and one at the satellite site (as shown in green in Figure 99). All of the CMSs synchronize metadata for all objects ingested into the grid, so all CMSs in the grid are equivalent.

When the grid nears its metadata capacity, additional CMSs are added: one for each existing CMS in the grid, including an additional CMS at the satellite site (as shown in Figure 99). An expansion CMS must be added at the satellite site to preserve the ability for this site to have local access to metadata for recently ingested objects, and for this site to continue to have the ability to operate while islanded from the rest of the grid.

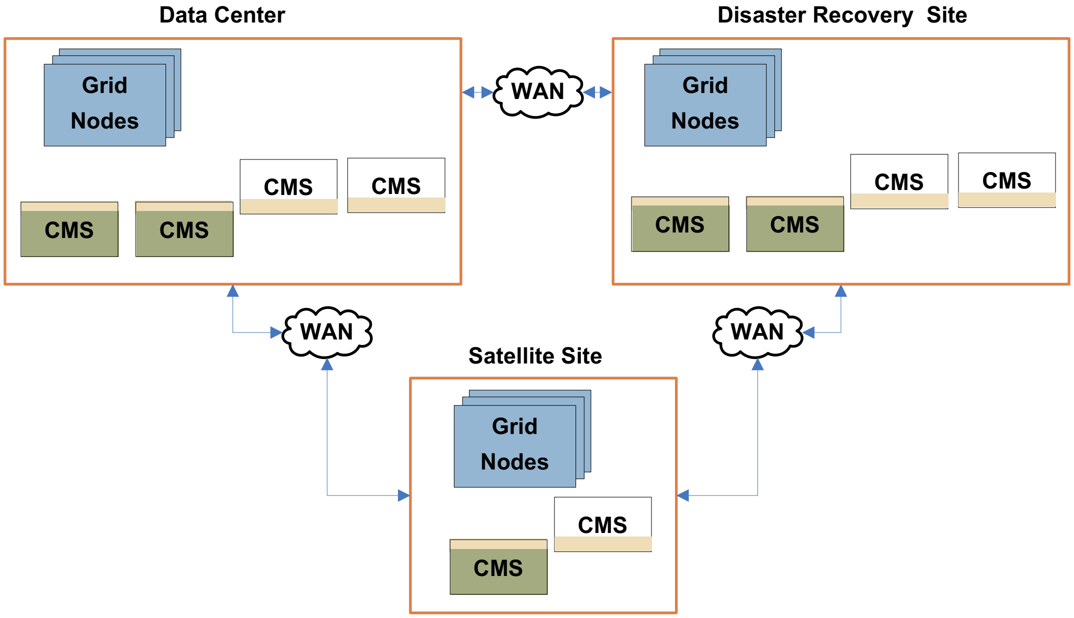

Figure 99 Metadata Synchronization: DC+DR+Satellite site

CMS Components

The CMS service includes components listed and described in the following sections.

Content

The Content component reports statistics on the metadata storage and replication activity.

If the CMS services in this grid use metadata synchronization, the page displays the attributes for Stored Objects and Managed Objects. The number of stored objects is the number of objects in the database of each CMS service. This number includes objects that have been purged from the grid. The number of managed objects is the number of objects owned by the CMS service. To get the total count of the objects managed by the grid, use the summary attributes at the grid level.

Metadata

The Metadata component is displayed only if the CMSs use metadata replication.

If the CMSs in this grid use metadata replication, the Metadata component reports statistics on the metadata storage and replication activity.

If present, this page displays the attributes for Stored Objects and Managed Objects. The number of stored objects is the number of objects in the database of each CMS service. For CMS services that use metadata replication, the number of stored objects does not include the number of purged objects. The number of managed objects is the number of objects owned by the CMS. To get the total count of the objects managed by the grid, use the summary attributes at the grid level.

Tasks

The Tasks component provides information about the tasks that the CMS service is processing against its database, and their current status.

Table 33 Elements in the CMS Tasks Component

Section	Description
Foreground Verification	The Foreground Verification table records the status of foreground verification tasks being performed by this CMS for each LDR currently undergoing foreground verification. (When you initiate foreground verification of an LDR, it directs each CMS service to actively check its database for objects recorded as being on that LDR, and verify their presence. For more information on foreground verification, see LDR Components Verification (page 206).) CMS processing of LDR foreground verification tasks is persistent across restarts. If a CMS is restarted or its operation is otherwise interrupted, foreground verification resumes where it left off after the CMS rejoins the grid. If an LDR restarts or goes offline, the CMS pauses foreground verification for that LDR until it is available, automatically resuming verification when the LDR rejoins the grid. Multiple CMSs process the foreground verification request for each LDR: to view the overall status of verification for an LDR across all CMSs, monitor the grid task. For more information, see Monitor Grid Tasks (page 270).
Background Tasks	Displays special tasks such as ILM verification tasks and migration tasks.
Recovery Sources	For CMSs that use metadata replication, the Recovery Sources table lists which CMSs can be used to restore this CMS database in the event of a failure.

Database

The Database component provides information about the type of database used by the CMS for metadata tracking, and data transaction statistics with that database. Each addition and each query to the database are considered transactions.

General database statistics include the number of transactions to date, the data transaction rate, and the number of connections to the database.

DICOM

DICOM is optional.

The DICOM component reports the number of studies and instances managed by the CMS (and *not* the number managed by the grid as a whole).

Synchronization

The Synchronization component provides information on message processing and data transmission. The Synchronization component is useful for viewing detailed statistics about the rate that data is being processed and sent from the CMS, as well as determining whether the CMS is operating efficiently and data is being processed in a timely manner. Synchronization messages are mainly used by CMSs that use metadata synchronization. CMSs that use metadata replication also use synchronization messages, but not to the same extent.

Events

The Events component uses the standard set of events attributes

Resources

The Resources component uses the standard set of resources attributes.

Timing

The Timing component uses the standard set of timing attributes.

FSG — File System Gateway

The File System Gateway (FSG) service provides a virtual file system interface using CIFS (Windows) or NFS (UNIX/Linux). This allows any type of fixed content, in any file format, to be stored to the grid without a need for proprietary interfaces. Applications can seamlessly store and retrieve images on the grid as if they were using ordinary disk storage.

FSG Operation

When an application saves a file to the File System Gateway, the FSG service stores a local copy, creates a file system reference for the file, and streams the file to permanent storage on multiple LDRs, whose locations are selected according to the business rules defined for the grid. A transient copy of the file is retained in the cache of the FSG where the file was ingested, as is a permanent pointer to the file and its file system reference.

Cached copies of files are used to speed retrieval of files requested via the FSG file system. The FSG service generally manages its cache by deleting the least-recently accessed copies of files (or files with a lower caching priority) as needed to make space. The FSG retains a permanent copy of the virtual file system and file pointers, regardless of the state of its cache, enabling the FSG to store and retrieve very large volumes of data.

Many details of FSG behavior can be customized to meet particular customer needs. For more information, see FSG Management (page 31).

FSG services are most commonly hosted on Gateway Nodes or combined Admin/ Gateway Nodes.

FSG Replication Groups

A HP MAS deployment always has two or more FSGs, arranged in a replication group. The primary purpose of a replication group is to facilitate data recovery and to provide service in the event that an FSG fails.

Within the replication group, at least one FSG provides read and write access to the grid for a set of clients. As files are saved to this FSG, file pointers and file system references are replicated to at least one other FSG in the replication group. This second FSG "mirrors" the file system of the first FSG to provide redundancy.

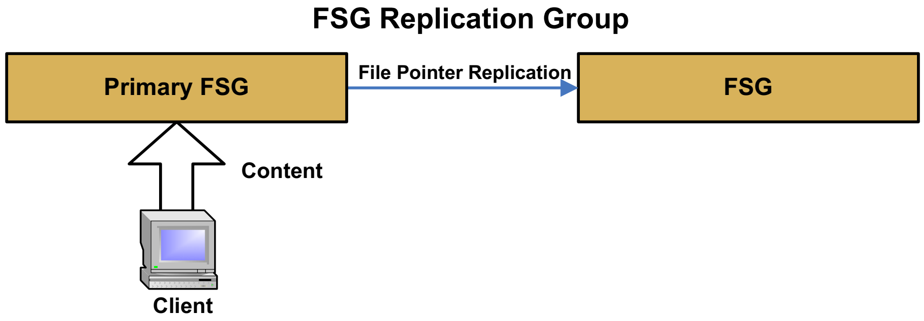

Figure 100 FSG Replication Group—General Case

These are the possible types of replication group:

- Basic Gateway replication group
- High Availability Gateway replication group

Each of these types of replication group provide different levels of redundancy and data availability, and can be customized to meet particular customer needs.

A HP MAS deployment may contain one or more FSG replication groups. There are several possible reasons for including multiple replication groups in a single grid. For example, each replication group can be used to provide grid access to a separate set of clients. This can be useful when write access is required from separate physical locations. Multiple replication groups may also be required if clients at a single location use incompatible authentication methods. If more than one CIFS authentication method is required (that is, both Windows Workgroup and Windows Active Directory), the grid must include a separate FSG replication group for each set of clients. Multiple replication groups can also be useful when different users have different requirements for data availability, prompting them to specify different types of replication group. For example, users who cannot tolerate the downtime associated with a manual failover require a clustered gateway solution as described in High Availability Gateway Replication Groups (page 199). As well, multiple replication groups can also be useful for scaling throughput when ingest loads are high. This allows for load distribution and more efficient ingest rates.

Primary FSGs

Every replication group includes at least one read-write FSG, known as the primary FSG. A replication group may contain either one primary FSG (in a Basic Gateway replication group) or a set of servers that form a clustered primary FSG (in a High Availability Gateway replication group).

Secondary FSGs

A secondary FSG receives file system replication messages from the primary, and provides a mirror of the primary's file system which can be used to provide read-only file system access to clients.

Read-only access to files saved on the primary FSG from the secondary FSG is not instantaneous. You must wait until the file system references have been replicated to the secondary. If you try to retrieve a file from a secondary FSG before the file system reference to that file has been completely replicated, the retrieval operation freezes. The operation resumes when the object replication completes or when the FSG service is halted and restarted.

Most replication groups include at least one secondary FSG. In a High Availability Gateway replication group, a secondary FSG is optional.

NOTE Secondary Gateway Nodes installed on virtualized hardware can only be used to back up the managed filed system. Secondary Gateway Nodes installed on virtualized hardware cannot be used for file retrieval.

Backup FSG

Within a replication group, one FSG automatically stores a backup of the managed file system into the grid each day. While the backup is in progress, the FSG is read-only. Therefore, the backup cannot be performed by any FSG that is currently acting as a primary FSG.

Activity that takes place after the backup is performed is retained as a "replication session". The file system backup and the interim replication session messages can be used to restore the file system of an FSG in case of failure. Because every FSG in a replication group retains the replication session messages for all data ingested into the group, you can restore any FSG in a replication group if a single member survives.

While it is performing a backup, the FSG continues to process replication session messages from the primary to keep its file system up to date with changes on the primary. Therefore, you can retrieve content from the Gateway Node while a backup is in progress, even if the content is newly ingested.

NOTE By default, preloading is disabled on the FSG designated to perform backups. If preloading is enabled, it is temporarily disabled on the FSG while it performs a backup. This applies to any FSG performing a backup. For more information on preload functionality, see FSG—File System Gateway (page 195).

A secondary FSG may be designated as the backup FSG in a replication group.

As one step of the update to Release 8.1, FSGs are updated to use online backups. In releases prior to 7.5, FSGs performed offline backups. While an FSG performed an offline backup, replication session messages were not processed and recently ingested files could not be retrieved from the backup FSG. Backup FSGs that performed offline backups were not recommended for data retrieval.

Basic Gateway Replication Groups

A Basic Gateway replication group is one that includes a single primary FSG and one or more secondary FSGs.

This is the simplest type of FSG replication group.

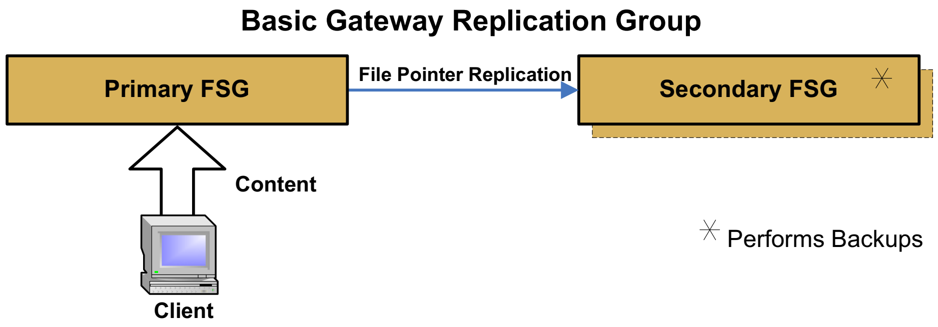

Figure 101 Basic Gateway Replication Group

Primary FSG

In a Basic Gateway replication group, there is a single primary FSG that provides read-write access to clients.

Secondary FSGs

When the primary FSG is not clustered, then the replication group includes one or more secondary FSGs. One secondary FSG acts as the backup FSG, and its mirrored file system can be used to provide clients with read-only access to the grid. A replication group could include more than one secondary FSG if read-only access is required from more than one location.

Any secondary can be used to restore a failed primary FSG from the daily file system backups and the interim replication messages copied to each member of the replication group.

Failover in a Basic Gateway Replication Group

In normal operation, the primary FSG in the replication group provides full read-write access to the grid for NFS and CIFS clients. The secondary FSG may provide read-only access to grid content.

In grids that support business continuity failover, if the primary FSG fails, a secondary FSG can be manually configured to act as a primary. After clients are manually redirected to the acting primary, they can continue to read and write to the grid. This is a temporary measure to maintain service while the primary FSG is repaired: grid access is interrupted until manual failover is completed, and the redundancy of file system information in the grid is temporarily reduced while the secondary FSG is an acting primary. If the secondary is also the backup FSG for the replication group, backups are not performed while it acts as the primary.

High Availability Gateway Replication Groups

A replication group can optionally be configured with a High Availability Gateway cluster (HAGC) that acts as the primary FSG for the group.

Each High Availability Gateway cluster includes two servers: a main FSG and a supplementary FSG. These two servers share a single virtual IP address, and provide a single access point to the grid for clients, whose NFS or CIFS file shares are present on both members of the cluster. Both servers in the HAGC must be co-located.

A High Availability Gateway replication group can optionally include one or more secondary FSGs.

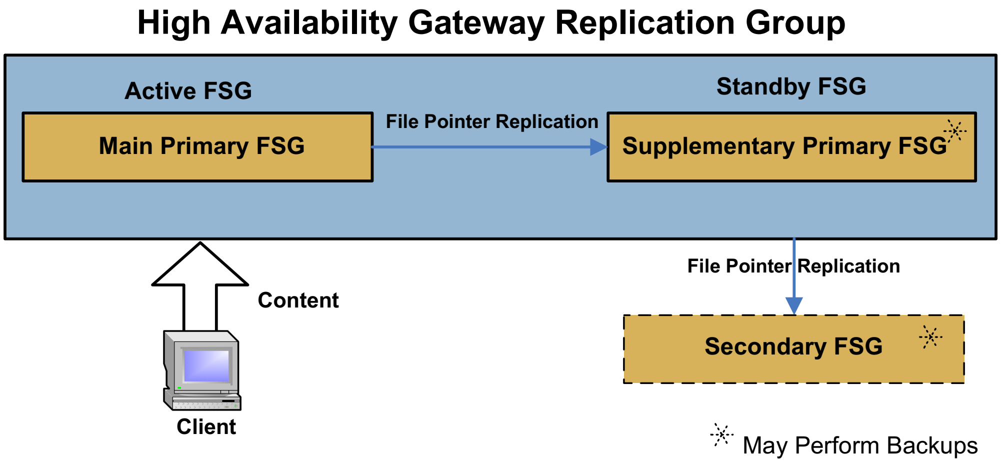

Figure 102 High Availability Gateway Replication Group with an Optional Secondary FSG

Main Primary FSG

Although the two FSGs in the HAGC are functionally equivalent, one FSG is designated as the main primary FSG. The main primary FSG is the FSG that is active by default (on system startup, and whenever the cluster is operating normally). The main primary FSG is therefore the one that usually provides NFS and CIFS services to clients.

Supplementary Primary FSG

The second FSG in the cluster is the supplementary primary. By default, the supplementary primary FSG is in standby (on system startup, and whenever the cluster is operating normally). When it is in the standby state, the supplementary primary FSG receives replication session messages and provides a “mirror” of the file system on the main primary FSG. This provides it with the ability to take over from the main primary FSG in the event of a failure and permits it to be used to restore the main primary FSG to full functionality, if needed. The supplementary primary FSG also usually performs the daily file system backup (unless the replication group includes an optional secondary FSG).

Replication Group Members

Every FSG replication group has a minimum of two members—a read-write FSG that provides client services, and a second FSG that provides redundancy and performs backups. The two members of a High Availability Gateway cluster (the main primary and the supplementary primary FSGs) can stand alone as an FSG replication group.

Other configurations are possible: you may also include one or more secondary FSGs in a High Availability Gateway replication group. A secondary may be desirable:

- at a Disaster Recovery site, to enable FSG services to be made available after the failure of the Data Center site
- if read-only access to files is required from another location (as both members of an HAGC must be co-located).
- to ensure that file system backups complete in a failover situation.

  The supplementary primary FSG cannot perform backups while it is acting as the active primary. Adding a secondary FSG that performs backups ensures that backups complete during failover. This provides better data security in the case that a failover lasts for an extended period of time.

Note that only one FSG in a replication group is configured to perform backups. Backups are usually performed by the supplementary primary (or whichever FSG in the cluster is currently in standby), unless the replication group includes a secondary FSG.

Failover in an HAGC

A High Availability Gateway Cluster (HAGC) includes two FSGs. When it is operating normally, one of the FSGs in the cluster is the active primary, and provides the file system interface to clients. The second FSG is the standby primary, ready to provide service in case the active primary FSG fails. The cluster health is monitored by a heartbeat service. Upon failure detection, the cluster automatically fails over from the active primary to the standby primary.

NOTE There are several situations that will not trigger an automatic fail over. For more information, contact HP Support.

Replication Group Summary

Replication Group Members

The following table summarizes the most common types of FSG replication groups. It also indicates which members of the group are commonly designated as backup FSGs.

Table 34 Summary of Types of FSG Replication Group

Replication Group Type	Primary FSG(s)	Other Members	Backup FSG	Secondary Used for Retrieval?
High Availability Gateway	Main and Supplementary	Secondary is optional	Supplementary[a] or Secondary	Yes
Basic Gateway	Single Primary	One or more Secondaries	Secondary	Yes

a. If the HAGC replication group does not include a secondary FSG, backups are performed by whichever FSG in the cluster is currently in Standby. This can be either the main or the supplementary FSG. Therefore, in the NMS MI, the setting Backup FSG on the FSG Management > <*replication group*> > Configuration > Settings page shows the primary main FSG as both the Primary FSG and the Backup FSG, indicating that backups are performed in the Primary HAGC cluster. However, backups are always performed by the FSG that is currently in Standby. By default this is the supplementary FSG.

In releases prior to 7.5, FSGs performed offline backups and it was recommended that the replication group include two secondary FSGs if a secondary was to be used for data retrieval. During the update to Release 8.1 software, FSGs are updated to perform online backups.

FSG Components

The FSG service includes components listed and described in the following sections.

Storage

The Storage component provides data on the service's storage space, and the objects stored and retrieved.

As files are ingested through the FSG, they are cached locally and forwarded to the grid for persistent storage. The FSG maintains the file system directory tree locally. Payload data is stored in the grid, making the file system appear to have very high capacity, even though the FSG server itself has relatively modest local resources.

Replication

The Replication component reports information about the replication status of the service. It also reports information about the role of the FSG service within its replication group and/or cluster, and the status of the cluster itself.

Role

Each FSG service is configured to assume a certain role by default. For example, all grids support mirrored FSG services to provide failover support if a primary FSG service (or primary FSG cluster) becomes unavailable. In a basic replication group, FSGs are configured as primary or secondary FSGs. In a High Availability Gateway Cluster, the server that is active by default is configured as a main primary while the server that is standby by default is configured as a supplementary primary.

Each FSG may temporarily assume a role other than its configured role when required. That is, a secondary FSG can be configured to act as a primary, and a primary can be configured to act as a secondary. The sections for Primary and Secondary are populated with data based on the current role of this FSG service; if acting as a primary, the Secondary section contains no meaningful data and vice versa.

Replication Status

The value of replication status is one of:

Normal	The replication session is proceeding normally. There are no errors.
No Primary	There is no active primary in the replication group. If this persists, it could indicate that the active primary is shut down, failed, or disconnected from the secondary FSG.
No Session	The secondary FSG cannot determine the next incoming replication session in the chain. If this persists, it could indicate that an FSG replication session inconsistency that must be manually corrected.
Starting	A transitional state indicating that the FSG replication system is starting up, or waiting for configuration bundles, or if the FSG is being restored from backup.
Stopping	A transitional state indicating that the FSG replication system is shutting down.
Proxying Session	The FSG reporting this state became the active primary as a result of a failover and is proxying any remaining replication messages due to pending ingests from the old active primary into the new replication session. This state is reported until the old active primary has closed the previous primary replication session. The session will not be closed until the FSG has recovered and is back online.

Primary Session Interrupted	This is the converse of Proxying Session. It is reported on the old active primary after a failover has occurred. This state is reported until the old primary replication session has been closed.
Error	An error state, possibly due to configuration errors, I/O errors, and so on.
Stopped	The FSG replication system is offline.
Session Inconsistent	The replication session between the primary and secondary FSGs are no longer consistent with each other.

Cluster Status

The replication component also reports the status of the High Availability Gateway Cluster. If an FSG is not part of a cluster, its cluster status is reported as "N/A" (not applicable).

In a High Availability Gateway Cluster, the cluster status is one of:

Normal	Both FSGs in the cluster are healthy: the main primary is active, and the supplementary primary is in standby (or vice versa).
Vulnerable	Only one server in the cluster is available and active.
Transitional	The cluster is not providing FSG service but the clustering mechanism is expected to automatically restore service.
Failed	No FSG in a cluster can provide an active FSG service and the failure cannot be recovered from automatically.

Note that the Failover cluster status is deprecated with release 8.1. If a failover within a replication group occurs, the "Failover Count (RPFO)" is incremented by one on the new active primary FSG. Historical attribute data for FSGs upgraded to 8.1 may include the "Failover" cluster status.

Backup

The Backup component provides information about backups of the managed file system. The FSG is configured at installation to make regular backup copies of the file system (folder and file information) and ingest them into the grid. Should the shared file system become corrupted, a backup can be used to restore the system. By default, backups are kept for only 14 days, and by default are not saved to removable media.

Shares

The Shares component provides information on share usage, allowing you to monitor the amount of data being saved on a per share basis. Quotas can be configured that when exceeded will trigger an alarm and notification.

Client Services

The Client Services component provides information about the support services used to manage the file system shares (NFS, CIFS), client integrations (Siemens RSH Responder), or clustering software (heartbeat). If the relevant support service is not running you do not have access to the grid's managed file system, the client integration, or clustering services.

Events

The Events component uses the standard set of events attributes.

Resources

The Resources component uses the standard set of resources attributes.

Timing

The Timing component uses the standard set of timing attributes.

LDR — Local Distribution Router

The Local Distribution Router (LDR) service handles content transport on the grid. Content transport encompasses many tasks including data storage, routing, and request handling. The LDR service does the majority of the grid's hard work by handling data transfer loads and data traffic functions.

The LDR service handles the following tasks:

- Content storage
- Content caching
- Content transfers from another LDR
- Data storage management
- Protocol interfaces (DICOM and HTTP)

The LDR service uses the standard overview service attributes with some additions for the protocol interfaces and object verification and replication.

Object Stores

The underlying data storage on an LDR is divided into a fixed number of "object stores" or storage volumes, which are partitions that act as mount points for the storage. When files are saved to an LDR, they are saved to an object store based on the file's unique identifier (its CBID), which is assigned to the file by the CMS as the object is ingested. Object stores are configured such that each one is associated with a subset of the total range of CBID values. Because CBIDs are randomly generated, objects are therefore assigned relatively evenly across the object stores. Note that object stores need not be identical in size: larger object stores are associated with a larger range of CBIDs.

After the first storage volume fills, the LDR as a whole cannot accept more data because it can no longer store objects whose CBIDs fall in that volume's assigned range.

Calculating Capacity

The capacity of an LDR can be calculated in different ways. The attribute LDR > Storage > Total Usable Space is calculated as (Storage on the object store with the least remaining space) multiplied by (Number of object stores). This gives an estimate of how much additional data can be saved to the LDR before the first object store fills. This estimate is most accurate when object size is small compared to the size of the object store and is relatively consistent over the lifetime of the LDR.

The attribute LDR > Storage > Total Free Space is calculated by adding up the remaining space on each object store. This may be a more accurate reflection of the capacity in some cases. For example, early in the lifetime of an LDR, if object size is relatively large compared to the size of the object stores, and by chance many objects in a row are stored to the same object store, the object stores may appear rather unbalanced. In this case the estimate represented by Total Usable Space could be much less than the value of Total Free Space.

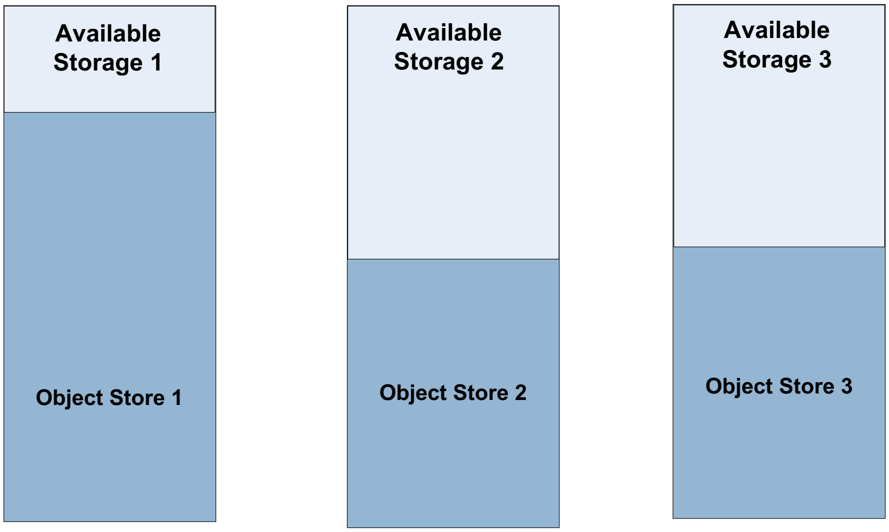

Free Space = Available Storage 1 + Available Storage 2 + Available Storage 3

Usable Space = Available Storage 1 x 3 (the number of Object Stores)

Figure 103 Calculating Storage Capacity of an LDR

LDR Components

The LDR service includes components listed and described in the following sections.

Storage

The Storage component provides information on the object storage space and the objects processed. The Storage component tracks the total amount of object storage space being used and the space available. If the available space falls below the configured amount, a notice alarm state occurs. This allows you to manage the storage proactively and purchase additional capacity only when necessary.

Verification

The Verification component provides information on the current state of the background data verification process, and allows you to modify the type and scope of verification performed on stored data.

Background verification runs automatically by default, proactively verifying the integrity of every object stored on the LDR. If an object is found to be corrupt, the object is quarantined and replaced using an uncorrupted copy from elsewhere in the grid. Background verification operates at an adaptive priority that adjusts its activity to avoid interference with grid operations, throttling its activities to prevent over-stressing the storage hardware.

Background verification is performed by the LDR itself, and the attributes displayed on the Overview > Main page give information on the status, progress, and results of background verification.

Alternatively, you can choose to initiate foreground verification on all or part of an LDR's storage. Foreground verification is driven by Control Nodes, whose CMS databases contain a record of every object in the grid, and verifies that each object recorded as being stored on the LDR is present. If an object is not present, the CMS creates another copy of the missing object. The replacement copy can be made in any location that satisfies the ILM rules for the grid: the copy is not necessarily made on the LDR that is being verified.

Foreground verification checks the existence rather than the integrity of each stored object, and completes much more quickly than background verification. It is appropriate as a way of quickly discovering if a storage device has issues.

Performing foreground verification on a single LDR has grid-wide implications; it initiates actions by every CMS in the grid. Therefore:

- To check the progress of foreground verification for an LDR, refer to a page at the grid level: Grid Name > Overview > Tasks summarizes the progress of verification for an LDR over all Control Nodes (see Monitor Grid Tasks (page 270)).

  NOTE When foreground verification completes for an LDR (that is, all CMSs finish verifying objects on that LDR), by default you are alerted via a Notice Alarm on the % Completion (XCVP) attribute for the LDR.

- Foreground verification is controlled by the grid task framework. To pause, resume, or abort a foreground verification task that is in progress, you must pause, resume, or abort the associated grid task. You can access these controls via:

  — a link on the LDR > Verification > Configuration > Main tab

  - the overview page listing all tasks for the grid as a whole (Grid Name > Configuration > Tasks)
  - the CMN > Grid Tasks component on the primary Admin Node (reporting Admin Node in an HCAC)
- Foreground verification puts load on all Control Nodes in a grid. To prevent undue impact on normal grid operations, no more than 25% of the total number of LDRs in a grid can be placed into foreground verification at the same time.

Replication

The Replication component provides additional information about node to node content replication.

DICOM

DICOM is optional.

The DICOM component tracks connectivity over the DICOM interface and records statistics on DICOM transactions.

HTTP

The HTTP component tracks connectivity over the HTTP interface and records statistics about HTTP transactions. The LDR uses the HTTP interface for transactions with the grid content.

Events

The Events component uses the standard set of events attributes.

Resources

The Resources component uses the standard set of resources attributes.

Timing

The Timing component uses the standard set of timing attributes.

NMS — Network Management System

The Network Management System (NMS) service powers the monitoring, reporting, and configuration options provided by the NMS MI. The NMS service can itself be monitored via the NMS MI.

The NMS service is hosted by the Admin Node. Depending on your grid's requirements, the grid may be deployed with a primary Admin Node or a High Capacity Admin Cluster (HCAC). These deployment types host differing types of the NMS service. For more information, see NMS Service Types.

Primary Admin Node

The primary Admin Node node is an Admin Node that host the CMN service. There is only one CMN service per grid and thus only one primary Admin Node per grid.

NMS Service Types

There are three NMS service types that can be used in a grid:

- consolidated NMS
- reporting NMS
- processing NMS

The reporting NMS and processing NMS service work together in an HCAC configuration to provide the same functionality as a consolidated NMS service.

Consolidated NMS Service

Hosted by a primary Admin Node and—if applicable—a second Admin Node, the consolidated NMS service collects and stores attributes from all services and components in the grid. The consolidated NMS service provides a management interface to display grid information.

NOTE The consolidated NMS service is installed on grids that are not configured with an HCAC.

The consolidated NMS service performs two main functions:

- It is a monitoring system that notifies you of problems when the status of key hardware or software changes.
- It is a browser-based interface making the system easily available to multiple users for:
  - Reporting status information about the grid so you can monitor and resolve grid issues.
  - Creating, viewing, and printing reports on current and historic data about each grid component based on your selection of report criteria.

— Configuring grid components and customizing the notification settings according to your criteria

Reporting NMS Service

The reporting NMS service is hosted by the reporting Admin Node of an HCAC (see High Capacity Admin Cluster (page 211)). The reporting NMS service provides a browser-based interface making the system easily available to multiple users for:

- Reporting status information about the grid so you can monitor and resolve grid issues.
- Creating, viewing, and printing reports on current and historic data about each grid component based on your selection of report criteria.
- Configuring grid components and customizing the notification settings according to your criteria.

NOTE The reporting NMS service must exists with a processing NMS service in the same HCAC.

Processing NMS Service

The processing NMS service is hosted by the processing Admin Node of an HCAC. It is a monitoring system that notifies you of problems when the status of key hardware or software changes. The processing NMS service provides attribute and data processing, and storage.

NOTE The processing NMS service must exists with a reporting NMS service in the same HCAC.

NMS Service Information

Current NMS service information can be viewed via the <*Admin Node*> > NMS > Overview page. The NMS Overview page displays the type of NMS: consolidated, reporting, or processing. Also displayed is grid services capacity information (for more information, see Grid Capacity (page 210)), MI system status (NMS Interface Engine Status), and e-mail notifications status.

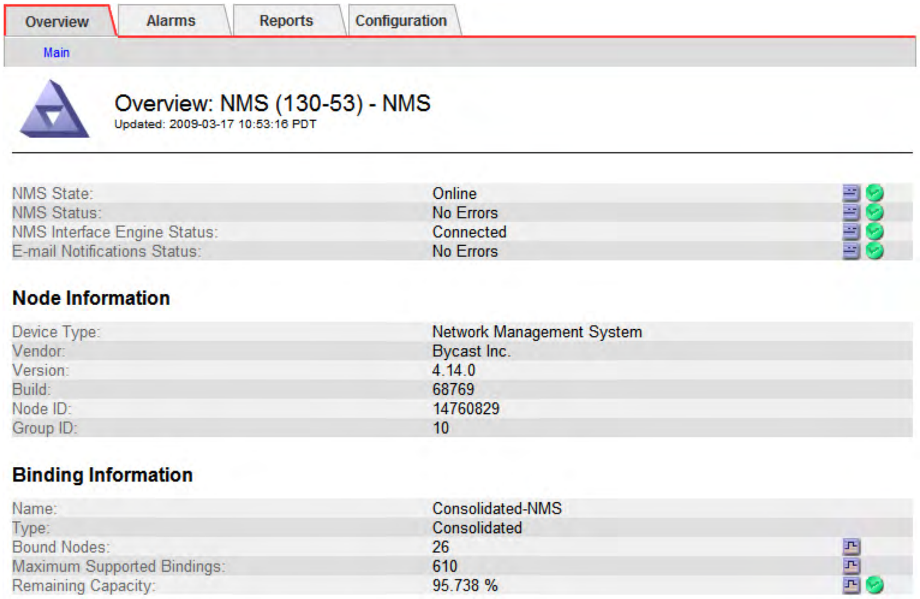

Overview: NMS (130-53) - NMS

Updated: 2009-03-17 10:53:16 PDT

NMS State:	Online
NMS Status:	No Errors
NMS Interface Engine Status:	Connected
E-mail Notifications Status:	No Errors

Node Information

Device Type:	Network Management System
Vendor:	Bycast Inc.
Version:	4.14.0
Build:	68769
Node ID:	14760829
Group ID:	10

Binding Information

Name:	Consolidated-NMS
Type:	Consolidated
Bound Nodes:	26
Maximum Supported Bindings:	610
Remaining Capacity:	95.738 %

Figure 104 Admin Node > NMS > Overview Page

Grid Capacity

Bindings

A binding is the persistent assignment of a grid service (for example, an FSG or SSM) to the consolidated NMS service or processing NMS service. This assignment is based on grid topology (consolidated Admin Node or HCAC). Each grid service in a HP MAS deployment is bound to the NMS service and there is a maximum number of bindings that an NMS service can support. The size of a grid (the number of bindings used) dictates whether the grid is deployed with a consolidated Admin Node or an HCAC.

NOTE Bindings are designated by the provisioning system and do not change during the lifetime of the grid.

If a grid node is decommissioned, its services are reclaimed and can be reused for expansion. The Bound Nodes count (<Admin Node> > NMS > Overview page) decreases by the number of reclaimed bindings.

Current Grid Services Capacity

You can confirm the current NMS service's grid services capacity (the number of bindings supported) via the <Admin Node> > NMS > Overview page. This information can be critical when performing a grid expansion. Before adding grid services, confirm that the current grid configuration can support these new bindings. If these bindings cannot be supported, then the Admin Node must be converted to an HCAC.

NOTE If the HCAC's processing Admin Node cannot support an increase in the number of bindings, contact HP Support.

To view current grid services capacity:

- In the NMS MI, go to **`<Admin Node>` > NMS > Overview**.

  Under Binding Information, Bound Nodes displays the current number of services bound to the Admin Node and Maximum Supported Bindings displays the current services capacity of the grid.

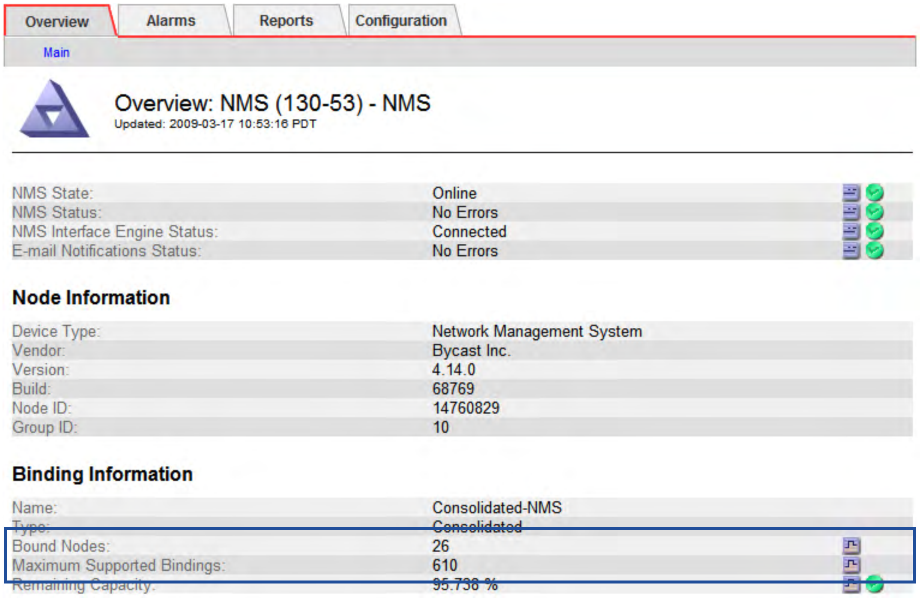

Overview: NMS (130-53) - NMS

Updated: 2009-03-17 10:53:16 PDT

NMS State:	Online	✓
NMS Status:	No Errors	✓
NMS Interface Engine Status:	Connected	✓
E-mail Notifications Status:	No Errors	✓

Node Information

Device Type:	Network Management System
Vendor:	Bycast Inc.
Version:	4.14.0
Build:	68769
Node ID:	14760829
Group ID:	10

Binding Information

Name:	Consolidated-NMS	
Type:	Consolidated	
Bound Nodes:	26	
Maximum Supported Bindings:	610	
Remaining Capacity:	95.738 %	✓

Figure 105 Grid Services Capacity

Conversion to a High Capacity Admin Cluster

To increase a grid's grid services and thus grid node capacity, a grid can be converted to an HCAC configuration at any time. During the conversion process, new servers are added and the old Admin Node is decommissioned.

High Capacity Admin Cluster

There is an upper limit to the number of grid nodes and their services that can be supported by a consolidated Admin Node. With large grids, the attribute management and alarm processing load can exceed the computational resources of the consolidated NMS service. To allow for these large grids and the limitations of a consolidated NMS service, a grid can be configured with a High Capacity Admin Cluster (HCAC).

NOTE Grid node capacity varies with the server hardware used.

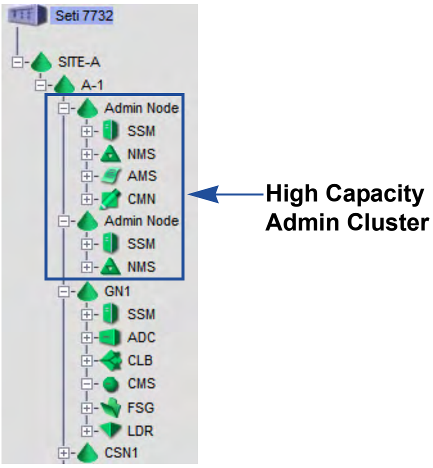

Figure 106 Grid Topology Tree with an HCAC

An HCAC consists of a reporting Admin Node and a processing Admin Node. The result is an increase to the capacity of the processing Admin Node's NMS service and thus an increase to the number of grid services that the grid can support. In an HCAC, all grid services are bound to the processing Admin Node's processing NMS service while the reporting Admin Node only provides the management interface to display grid information.

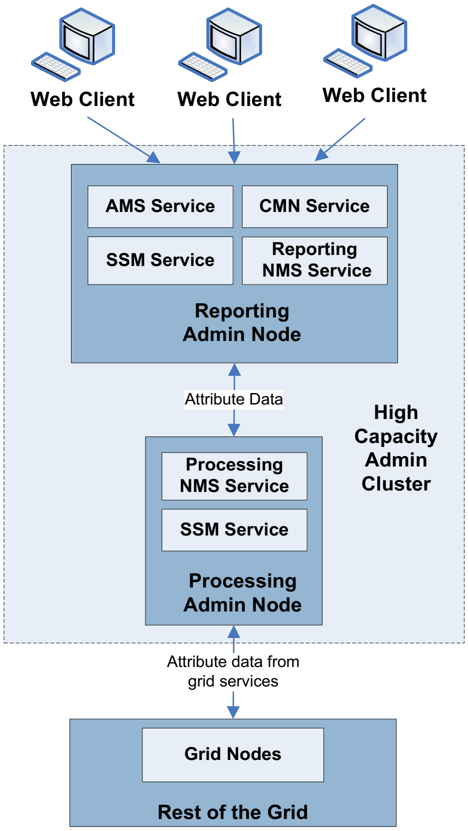

Figure 107 High Capacity Admin Cluster

The reporting Admin Node forwards information from the processing Admin Node to web clients via the NMS MI. Attribute management, reports, and alarm processing including notifications are handled by the processing Admin Node.

HCAC Grid Services

The reporting and processing Admin Nodes include the following grid services:

Table 35 HCAC Admin Node Services

Grid Node	Grid Service
Reporting Admin Node	Reporting NMS AMS SSM CMN (primary Admin Node only)
Processing Admin Node	Processing NMS SSM

Redundancy

To add redundancy, the grid can be configured with two Admin Nodes or HCACs. Redundancy provides a separate—but equivalent—NMS MI to display grid information and to perform configuration task available though the NMS MI. Redundancy does not provide failover protection—except for e-mail notifications. For more information, see E-mail Notifications (page 215).

Bindings to the primary Admin Node or HCAC are duplicated at the second Admin Node or HCAC. Thus, with two HCACs, data is always available to web clients (as long as at least one Admin Node or HCAC is running). If an Admin Node or HCAC becomes unavailable, web clients can reconnect to the available Admin Node or HCAC and continue to view and configure the grid. With redundancy if an Admin Node or HCAC becomes unavailable, attribute processing continues, alarms are still triggered and related notifications sent.

Primary Admin Node Redundancy

With Admin Node redundancy, the grid is configured with two Admin Nodes: a primaryAdmin Node and a second Admin Node that does not host the CMN service. Both Admin Nodes host a consolidated NMS service. Note that alarm acknowledgments made from one Admin Node are not copied to the other Admin Node. For more information, see Alarm Acknowledgments (page 215).

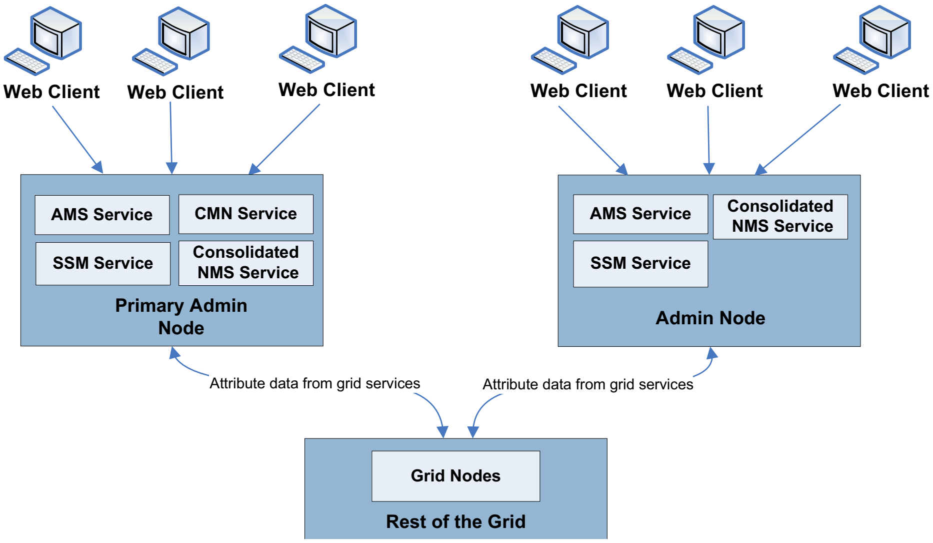

Figure 108 Admin Node Redundancy

High Capacity Admin Cluster Redundancy

With HCAC redundancy, the grid is configured with two HCACs one of which is configured with a primary reporting Admin Node hosting the CMN service. Note that alarm acknowledgments made from one Admin Node are not copied to the other Admin Node. For more information, see Alarm Acknowledgments (page 215).

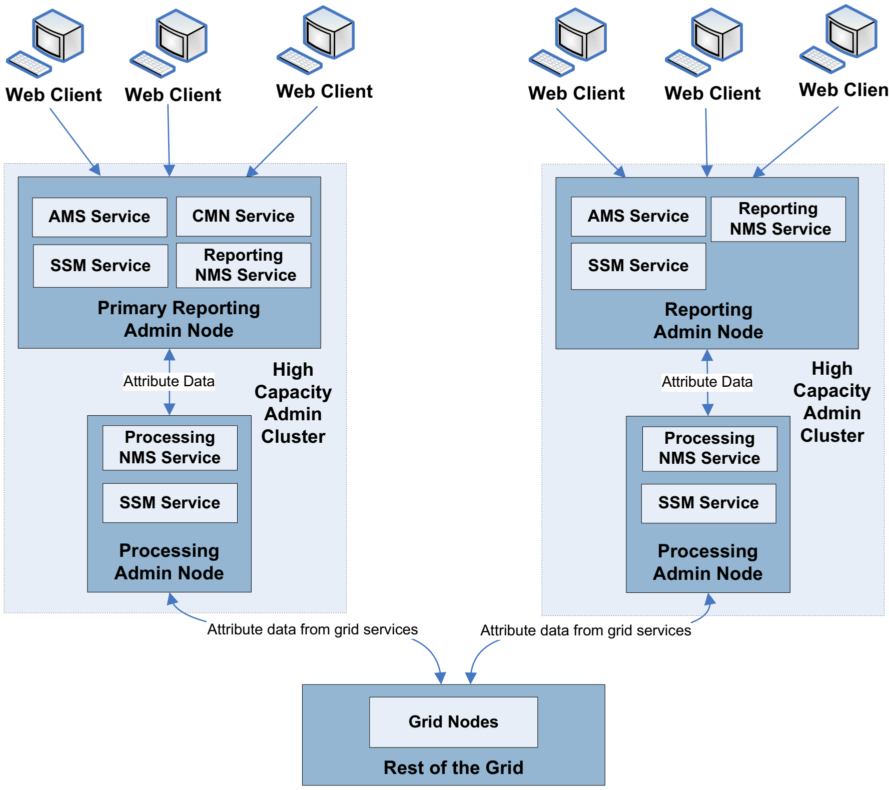

Figure 109 High Capacity Admin Cluster Redundancy

E-mail Notifications

In a grid configured with a second Admin Node or HCAC, one Admin Node or HCAC is configured as the preferred sender of e-mail notifications. This preferred sender can be either Admin Node or HCAC. The Admin Node or HCAC not set as the preferred sender becomes the "standby" sender.

Under normal grid operations, only the preferred sender sends notifications. The standby sender monitors the preferred sender (primary Admin Node or processing Admin Node of an HCAC) and if it detects a problem, the standby sender switches to online status and begins sending notifications. For more information and procedures on setting the preferred sender, contact HP Support.

The Preferred vs. Standby Sender

There are two scenarios in which both the preferred sender and the standby sender may send notifications:

- It is possible that while the grid is running in this "switch-over" scenario, whereby the standby sender begins sending notifications, the preferred sender will maintain the ability to send notifications. If this occurs, duplicate notifications are sent: one from the preferred sender and one from the standby sender. When the Admin Node or HCAC configured as the standby sender no longer detects errors on the preferred sender, it switches to "standby" status and stops sending notifications. Notifications are once again only sent by the preferred sender. For more information on notifications, see E-mail Notifications (page 215).
- If the standby sender cannot detect the preferred sender, the standby sender switches to online and begins to send notifications. In this scenario, the preferred sender and standby senders are "islanded" from each other. Each sender (Admin Node or HCAC) may be operating and monitoring the grid normally, but because the standby sender cannot detect the other Admin Node of the preferred sender both the preferred sender and the standby sender send notifications.

When sending a test e-mail, all NMS services send a test e-mail.

Alarm Acknowledgments

Alarm acknowledgments made from one Admin Node or HCAC are not copied to the other Admin Node or HCAC. Because acknowledgments are not copied to the other Admin Node or HCACs, it is possible that the grid topology tree will not look the same on each Admin Node or HCAC. This difference can be useful when connecting web clients to the NMS MI. Web clients can have different views of the grid based on the client's needs.

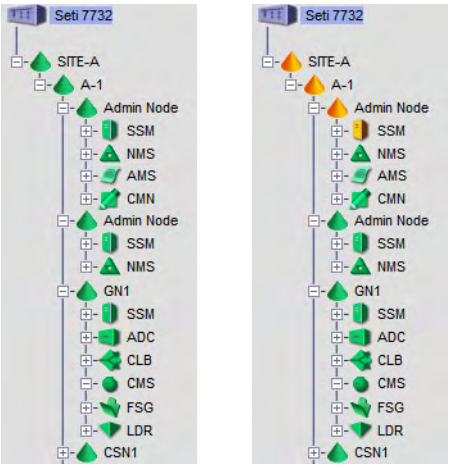

Figure 110 Grid Topology Tree with Differing Alarm Acknowledgments

Note that notifications are sent from the Admin Node or HCAC where the acknowledgment occurs. This may not be the preferred sender.

NMS Components

The NMS service includes components listed and described in the following sections.

Database

The Database component provides information about the type of database used by the NMS service for tracking attributes.

Events

The Events component uses the standard set of events attributes. As the NMS service also acts as an attribute repository, for the NMS service this component reports on the status of the repository. As well, notification status, including queued notifications, is reported.

Resources

The Resources component uses the standard set of resources attributes.

Timing

The Timing component uses the standard set of timing attributes.

Interface Engine

The Interface Engine component provides information about the status of the NMS MI, notification events, and connection pool. The connection pool displays information regarding connections to the database.

SSM—Server Status Monitor

The Server Status Monitor (SSM) is a service present on *all* nodes. Each node on the grid has its own SSM service to monitor that node's status, services, and the system log. It monitors the condition of the node and related hardware, polls the server and hardware drivers for information, and displays the processed data via the NMS MI.

The information monitored includes:

- CPU information (type, mode, speed)
- Memory information (available, used)
- Performance (system load, load average, uptime, restarts)
- Volumes (status, available space)

SSM Components

The SSM includes components listed and described in the following sections.

Services

The Services component tracks the services and support modules running on the node. It reports the service's current version, status, the number of threads (CPU tasks) running, the current load on the CPU, and the amount of RAM being used. The grid services are listed, as are support modules (such as time synchronization). Also listed is the operating system installed on the grid node.

The status of a service is either Running or Not Running. A service is listed with a status of Not Running when its state is Administratively Down. For more information on service state indicators, see Service State Indicators (page 40).

The section on Packages reports overall service suite installations and version numbers to facilitate software update procedures.

RAID

The RAID component only exists on SSM services that are directly monitoring at least one RAID unit. The RAID component displays extensive data on each attached RAID unit

The NMS MI displays as much status and basic hardware information as each physical RAID communicates. This information may include the RAID controller status, the array status, and the status of each of the available drives.

Events

The Events component relays logged events from the hardware drivers. Interpretation of these numbers depends on the hardware and drivers in use on your system. You can treat this data as a general indicator of problems with the server.

The error event counters can be individually reset to zero.

To reset error event counters:

1. Go to **SSM > Events > Configuration**.
2. Select the **Reset** boxes for the specific event counters to be reset.
3. Click **Apply Changes**.

Resources

The SSM uses the standard set of resources attributes that report on the service health and all computational, disk device, and network resources. In addition, the Resources attributes report on memory, hardware and network interfaces.

Timing

The SSM uses the standard set of timing attributes that report on the state of the node's time and the time recorded by neighboring nodes. In addition, the SSM Timing attributes report on NTP synchronization.

7 Information lifecycle mangement

This chapter contains background information on Information Lifecycle Management (ILM).

Configurable ILM may not be enabled on your grid.

An ILM policy defines how and where objects are stored in the grid. HP MAS either uses a custom ILM or a configurable ILM. If the grid has a configurable ILM, the value of the grid option Configurable ILM in Grid Management > Grid Configuration is Enabled.

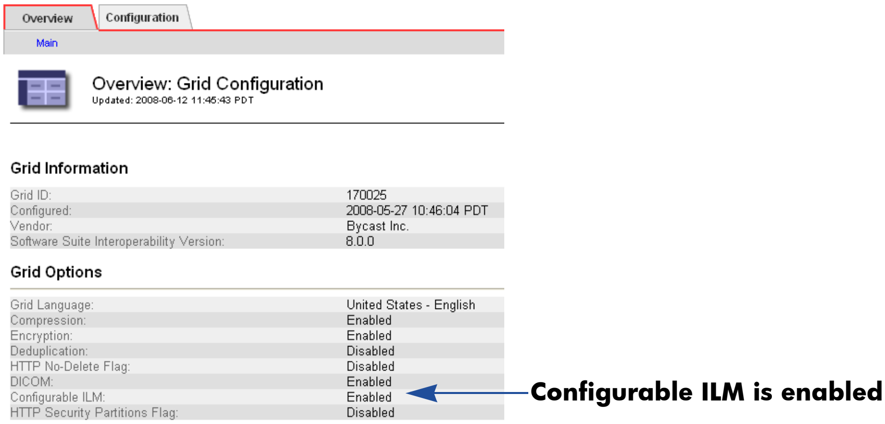

Overview: Grid Configuration
Updated: 2008-06-12 11:45:43 PDT

Grid Information

Grid ID:	170025
Configured:	2008-05-27 10:46:04 PDT
Vendor:	Bycast Inc.
Software Suite Interoperability Version:	8.0.0

Grid Options

Grid Language:	United States - English
Compression:	Enabled
Encryption:	Enabled
Deduplication:	Disabled
HTTP No-Delete Flag:	Disabled
DICOM:	Enabled
Configurable ILM:	Enabled
HTTP Security Partitions Flag:	Disabled

Configurable ILM is enabled

Figure 111 Configurable ILM Enabled

The Grid Management > ILM Management menu is used to view and configure ILM policies. This menu is only used if the system has a configurable ILM.

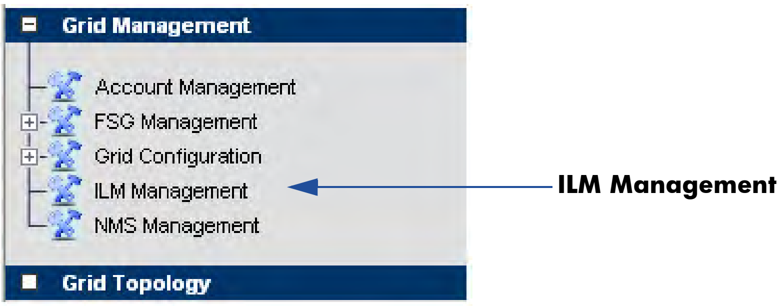

ILM Management

Figure 112 ILM Management Menu

NOTE ILM configuration is restricted to user accounts with Grid Management permissions such as the Vendor account. Other user accounts can view the ILM information but cannot modify it.

What is Information Lifecycle Management?

ILM polices define how and where objects are stored in the grid. At a high level, ILM policies dictate:

- geography—the location of the files
- storage grade—what type of storage to use
- replication—the number of copies to make

Location, storage grade and number of copies can vary over time.

In the ILM example shown in Figure 113, an object is ingested in the grid via an FSG. At ingest, one copy of the object is stored in the Data Center (DC) site on SATA disks, one copy is stored in the Disaster Recovery (DR) site on SATA disks, an done copy is stored on archival media at the DR site. After on year, the copy on SATA disks at the DR site is deleted.

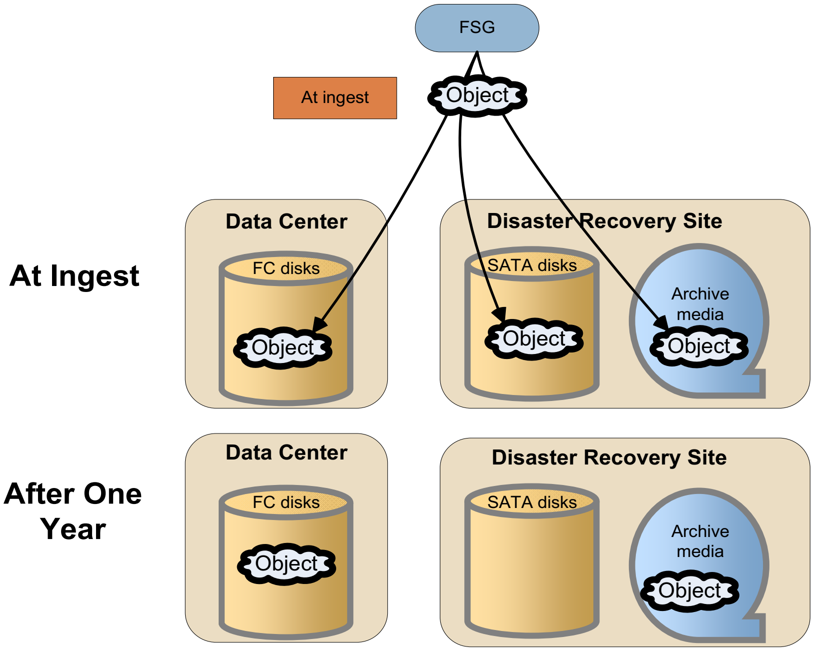

Figure 113 Information Lifecycle Management

ILM Policy Elements

ILM policies manage where objects are placed at ingest and dictate how objects are replicated, moved, and deleted over time. An ILM policy is a set of prioritized rules. A rule (see Figure 114) contains instructions for placing objects in the grid over time. A rule specifies:

- one or more filters used to determine if the rule applies to the object
- the number of copies to store
- the retention period
- the storage location

File path and FSG replication group ID are common criteria to determine whether a rule applies to an object.

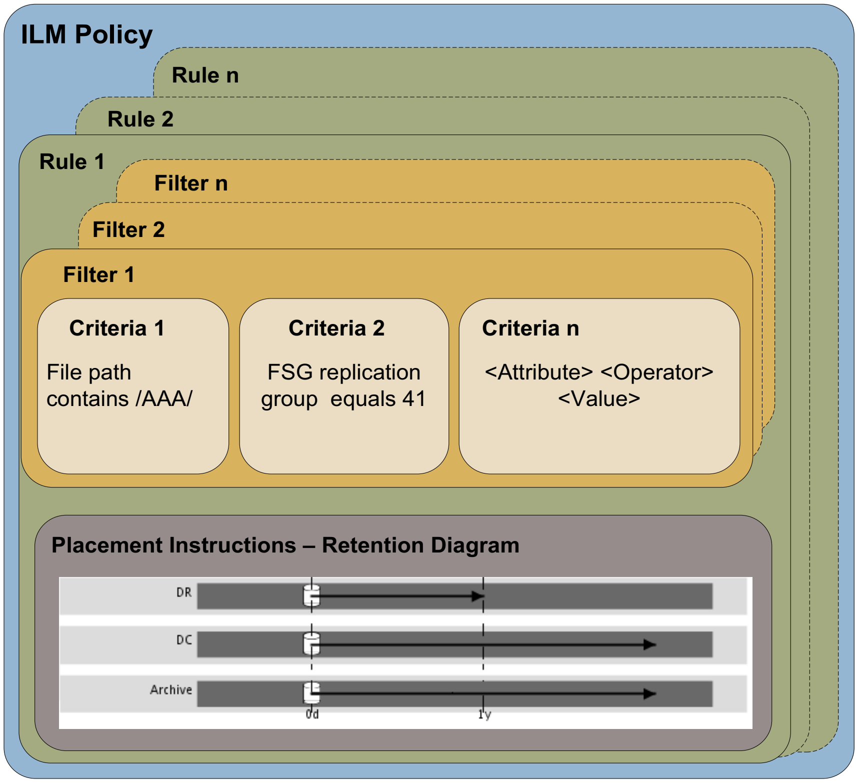

Figure 114 ILM Policy Elements

Retention Diagrams

Retention diagrams represent what happens to an object over time. Figure 115 is an example of a retention diagram.

Storage location is shown at the left. In this example, there are three storage locations: DC, DR, and Archive.

The bottom axis represents time elapsed since the reference time (typically, since ingest). In this example, there are two periods: the first period starts at ingest (Day 0) and the second one starts 1 year after ingest.

The placement instructions also specify the number of copies to keep. In this example, at ingest, one copy is stored at the DC site, one copy at the DR site, and one copy on Archive. After one year, the DR site copy is deleted.

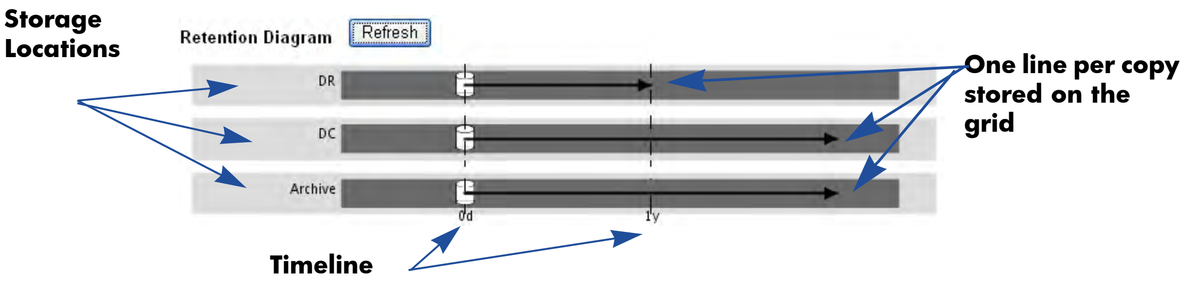

Figure 115 Sample Retention Diagram

Storage Pools

A HP MAS deployment may incorporate multiple spinning and archive media storage technologies, for instance SCSI, SAN, SATA, or Fibre Channel.

In the example shown in Figure 116, a HP MAS system is deployed across two sites or groups: DC and DR sites. Files can be ingested via FSGs or the HTTP API. The DC site has SAS-attached SATA disks and FC-attached Fibre Channel disks, and the DR site has SAS-attached SATA disks and a tape library. Storage grade refers to the type of storage. In this case, there are three storage grades: SAS-attached SATA, FC-attached Fibre Channel, and archive media.

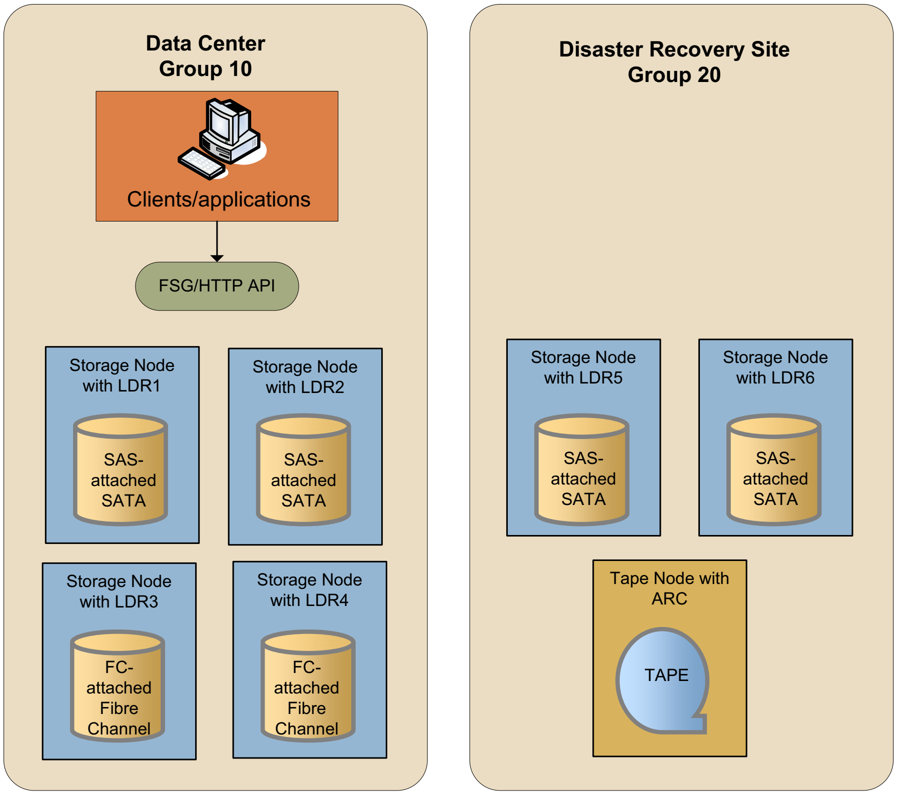

Figure 116 . DC + DR Site with Archive Configuration

The DC site has four Storage Nodes. A Storage Node is a server that includes an Local Distribution Router (LDR) service. The LDR handles content transport on the grid. This encompasses many tasks including data storage, routing, and request handling.

The DR site has two Storage Nodes and one Tape Node. A Tape Node is a server that includes an Archive (ARC) service. The ARC service manages storage and retrieval operations for content stored on devices that use archive media, tape libraries for instance.

For the purpose of ILM content replication, storage is organized into “storage pools”. A storage pool is a logical grouping of Storage Nodes or Tape Nodes. Storage pools allow you to group Storage Nodes or Tape Nodes based on factors such as performance, location, reliability, and cost. A storage pool has two attributes: storage grade and group. Storage grade typically refers to the type of storage, for example, SATA, SAN, Fibre Channel, archive media, and so on. The group is the physical location of the storage, for example DC or DR site.

Figure 117 shows examples of storage pool definitions for a DC+DR + Archive grid.

- DC SATA: All LDRs at the DC with SAS-attached SATA disks
- DC FC: All LDRs at the DC with Fibre Channel-attached Fibre Channel disks
- DR SATA: All LDRs at the DR site with SAS-attached SATA disks
- DC+DR SATA: All LDRs in the grid with SAS-attached SATA disks
- Archive: Archive media at the DR site

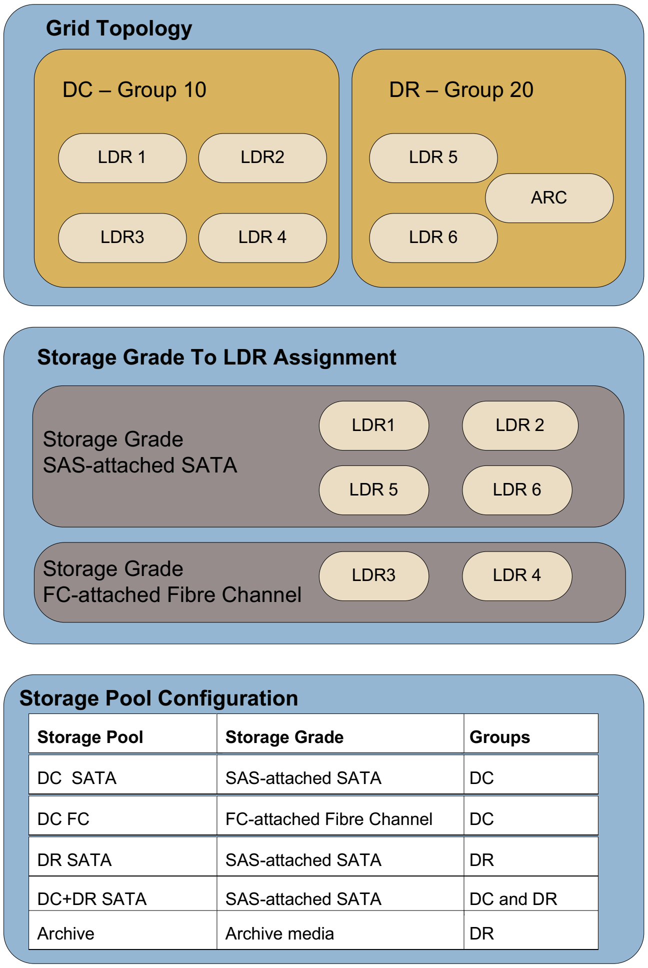

Storage Pool	Storage Grade	Groups
DC SATA	SAS-attached SATA	DC
DC FC	FC-attached Fibre Channel	DC
DR SATA	SAS-attached SATA	DR
DC+DR SATA	SAS-attached SATA	DC and DR
Archive	Archive media	DR

.

Figure 117 Storage Pool Definition

About Rules

This section describes the theory behind rules.

Filters and Criteria Evaluation

A rule contains the instructions for placing objects in the grid over time. A rule contains one or more filters. A filter is a set of criteria used to evaluate whether the rule applies to a specific object. Criteria are metadata such as file path, FSG replication group ID, and backup identifier.

Metadata Tags

Metadata tags can be either strings or integers. Strings are used for metadata such as file path. Strings are case-sensitive. Integers are used for metadata such as group IDs. Table 36 lists the application metadata that can be used to create ILM rules. Custom metadata defined in custom applications developed using the HP MAS HTTP API will also appear in the list of criteria metadata and may be used to define ILM rules.

Table 36 Application Metadata

Metadata	Code	Definition	Notes on ILM Use
Backup Identifier	BUID	If present, indicates that the object is a FSG Backup.	Used in the FSG backup rules.
External Application Identifier	XAID	The external application that created the object.	If the HTTP API is used directly, the application can specify settings for XAID, XTYP, XVER and the settings can be used in the rule filters. If ingest happens through the FSG, XAID = BFSG and XTYP = BKUP or FILE.
External Object Type	XTYP	The type of the object, as defined by the external application.	
External Object Version	XVER	The version of the object, as defined by the external application	
File Modification Time	MTIM	The modification time associated with the file at grid ingest time.	Used to create a criteria where “MTIM exists” if MTIM is used as a reference time.
File Replication Group	FGRP	The replication group of the FSG that the file was stored on.	Very useful when setting up a policy for distributed topologies such as “DC+DR and Satellites” and “Federated”. The File Replication Group allows you to distinguish between different locations. You can get the replication group ID values for your grid from the NMS interface (go to **Grid Management > FSG Management**).

Table 36 Application Metadata *(continued)*

Metadata	Code	Definition	Notes on ILM Use
File Status Change Time	CTIM	Status change time associated with the file at grid ingest time. Changed by writing or setting inode information (that is, owner, group, link count, mode, and so on).	Used to create a criteria where “CTIM exists” if you CTIM is used as a reference time.
FSG File Path	FPTH	The file path at ingest time.	One of the most useful metadata since file path often contains information such as department name, time stamp, etc. Could also be used in a “DC+DR and Satellites” environment to determine where content was ingested (must be used with care since it is more likely to change than the FSG Replication Group).
HTTP Protocol Handler Version	PHTP	The presence of this metadata indicates that the content was ingested into the grid using the HTTP API (this also includes all content ingested via a FSG).	Used in earlier releases of the software to create filters that match all content ingested via HTTP API or FSG (HTTP Protocol Handler Version Greater Than 0).
Source Node ID	INID	The node id of the service which first processed the data and created the original object.	Normally not used for ILM except in very special cases.

Reference Time

The placement instructions specify the point in time from which the replication instructions are valid.

There are three ways to specify the reference time for the ILM rules.

- Ingest time (default)
- Timestamp embedded into the FSG file path
- Time specified by a metadata tag

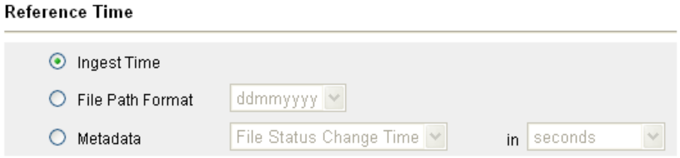

Reference Time

● Ingest Time

○ File Path Format ddmmyyyy

○ Metadata File Status Change Time in seconds

Figure 118 Reference Times for ILM Rules

Ingest Time

When the reference time is Ingest, all placement instructions are specified relative to when the object was ingested into the grid. Time can be specified in days or years. Ingest Time is the default reference time for ILM rules.

FSG File Path

FSG File Path can be used as a reference time in situations where the directory name includes a reference to the time. This scenario is common in migration situations. Three time formats are supported:

- mmddyyyy
- ddmmyyyy
- yyyymmdd

If the reference timestamp cannot be extracted from the file path (for instance, if the format is wrong), ingest time will be used as reference time.

If the reference timestamp lies in the future, the placement instructions of the first time period (at day 0) are applied until that period expires, based on the reference timestamp.

Metadata Tag

The information contained in custom integer metadata tags can also be used as reference time. Custom metadata tags that have been defined in custom applications developed using the HTTP API will appear in the list of metadata and may be used to define a reference time. The time of these custom metadata tags can be expressed in seconds, milliseconds, or microseconds elapsed since January 1, 1970.

If the reference timestamp cannot be extracted from the custom metadata tag (for instance, the tag does not exist), ingest time will be used as reference time.

These HP MAS built-in metadata tags are supported:

- File Status Change Time: the status change time associated with the file at ingest
- File Modification Time: the content modification time associated with the file at ingest

If using File Status Change Time or File Modification Time as reference time in a rule, it is a good practice to include a filter criteria in the rule that refers to this metadata (File Status Change Time Exists or File Modification Time Exists).

Types of Rules

Active Rules

The active rules are the rules currently enforced, that is the rules that make up the active policy that the grid is using to replicate content. You can view the list of active rules in the Overview > Active Policy page of the NMS ILM Management interface.

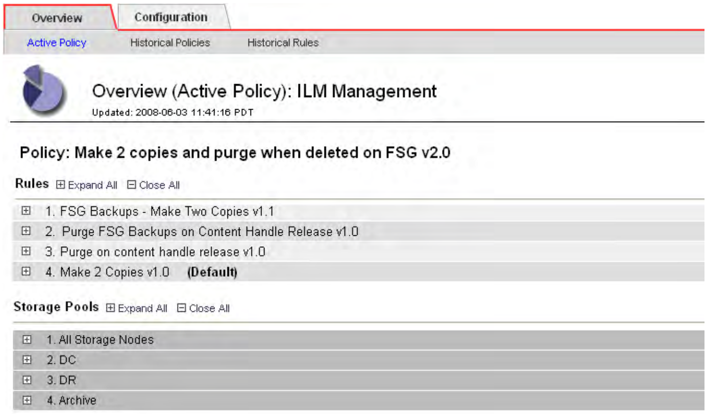

Overview (Active Policy): ILM Management
Updated: 2008-06-03 11:41:16 PDT

Policy: Make 2 copies and purge when deleted on FSG v2.0

Rules

1. FSG Backups - Make Two Copies v1.1
2. Purge FSG Backups on Content Handle Release v1.0
3. Purge on content handle release v1.0
4. Make 2 Copies v1.0 **(Default)**

Storage Pools

1. All Storage Nodes
2. DC
3. DR
4. Archive

Figure 119 Active Rules in Active Policy

Historical Rules

Historical rules are rules that were activated at some point. This means rules that are currently active and rules that are no longer active. You can view the list of historical rules with their most recent start and end times in the Overview > Historical Rules page.

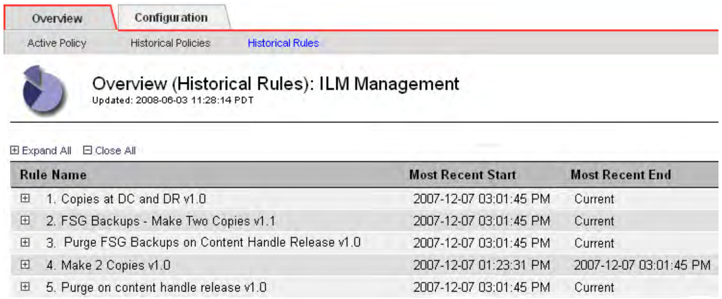

Overview (Historical Rules): ILM Management
Updated: 2008-06-03 11:28:14 PDT

Rule Name	Most Recent Start	Most Recent End
1. Copies at DC and DR v1.0	2007-12-07 03:01:45 PM	Current
2. FSG Backups - Make Two Copies v1.1	2007-12-07 03:01:45 PM	Current
3. Purge FSG Backups on Content Handle Release v1.0	2007-12-07 03:01:45 PM	Current
4. Make 2 Copies v1.0	2007-12-07 01:23:31 PM	2007-12-07 03:01:45 PM
5. Purge on content handle release v1.0	2007-12-07 03:01:45 PM	Current

Figure 120 Historical Rules

Built-in Rules

When configurable ILM is first enabled, the active policy contains two built-in rules: Make 2 Copies V1.0 and Purge FSG Backups on Content Handle Release v1.0.

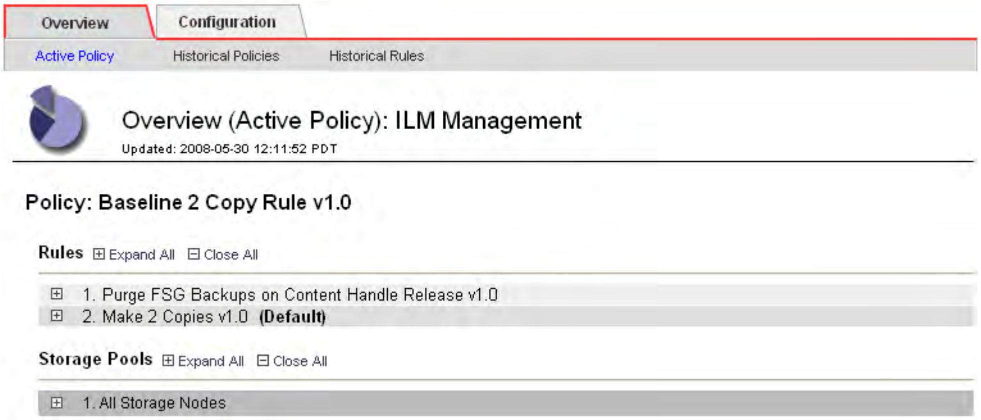

Overview (Active Policy): ILM Management

Updated: 2008-05-30 12:11:52 PDT

Policy: Baseline 2 Copy Rule v1.0

Rules ⊞ Expand All ⊟ Close All

⊞ 1. Purge FSG Backups on Content Handle Release v1.0

⊞ 2. Make 2 Copies v1.0 **(Default)**

Storage Pools ⊞ Expand All ⊟ Close All

⊞ 1. All Storage Nodes

Figure 121 Built-in ILM Policy

The Make 2 Copies rule specifies that two copies of every ingested object be stored indefinitely on any disk in the grid (that is, in the All Storage Nodes storage pool). When configurable ILM is first set up, this built-in rule acts as a default rule, serving as insurance that all ingested objects are preserved. This rule cannot be modified or deleted. However, you may choose to deactivate it if it does not meet your requirements.

The Purge FSG Backups on Content Handle Release automatically deletes copies of the FSG backups. Unlike the Make 2 Copies rule, the Purge FSG Backups on Content Handle Release can be modified.

Rule Version Numbers

Each rule is automatically assigned a version number in the form n.n. The first time a rule is created, its version number is 1.0. The version number is incremented each time the rule is modified (for instance, the version number changes from 1.2 to 1.3 or from 1.9 to 2.0).

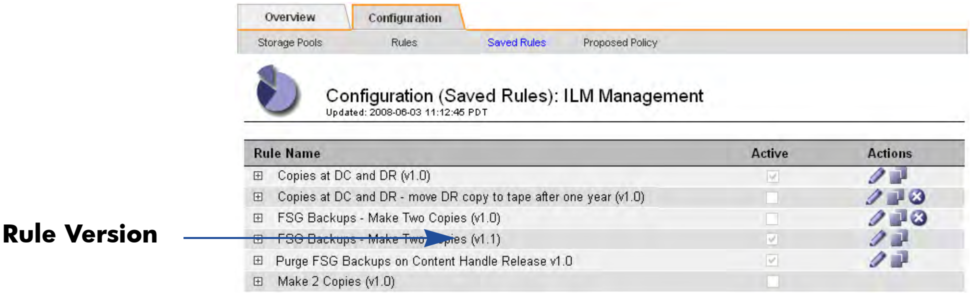

Configuration (Saved Rules): ILM Management

Updated: 2008-06-03 11:12:45 PDT

Rule Name	Active	Actions
⊞ Copies at DC and DR (v1.0)	☑	
⊞ Copies at DC and DR - move DR copy to tape after one year (v1.0)	□	
⊞ FSG Backups - Make Two Copies (v1.0)	□	
⊞ FSG Backups - Make Two Copies (v1.1)	☑	
⊞ Purge FSG Backups on Content Handle Release v1.0	☑	
⊞ Make 2 Copies (v1.0)	□	

Figure 122 Rule Versions

Unavailable Resources

In addition to the preferred storage location, a rule should specify an alternate location for a temporary copy of the object. This location is used in case the preferred location is unavailable.

△ CAUTION Use of temporary copies is strongly recommended. Failure to specify an alternate storage location temporary storage location puts data at risk if the preferred location is unavailable.

If the preferred location is unavailable (for example if the Storage Nodes in the storage pool are full), the following happens:

1. The object is copied to each preferred location if possible.
2. For each preferred location that is unavailable, the object is copied to the alternate location if specified.
3. When a preferred location becomes available, the object is placed in the preferred location.
4. After the object has been stored in all preferred locations, the temporary copies in all alternate locations are purged.

For example, if a rule specifies the following placement for all content at ingest:

- Store two copies in DC site (alternate location for a temporary copy is the DR site)
- Store one copy at local site

If the DC storage pool is unavailable when the object is ingested, two temporary copies are stored in the DR storage pool and one copy is stored at the local site. After the DC pool becomes available, two copies are made in the DC storage pool and the DR copies are deleted.

One special case where an alternate location may not be appropriate is when the preferred location is archive media. In addition, Tape Nodes should not be used to store temporary copies.

About Policies

An ILM policy consists of a set of prioritized rules which specify the instructions for placing objects in the grid over time. A policy lists:

- the rules to be applied
- the order in which the rules are evaluated
- the default rule to be used if the object does not match the filters specified in any of the rules

When the HP MAS system is first installed with configurable ILM, the active policy is a built-in policy called Baseline 2 Copy Rule v1.0 that contains the built-in rules Make 2 copies and Purge FSG Backups on Content Handle Release.

Version Numbers

Each policy is automatically assigned a version number in the form *major.minor*. The version number of the built-in policy is 1.0. The version number is incremented each time the policy is changed.

A change to the policy name triggers a major revision (for instance, the version number changes from 2.2 to 3.0).

Changes to the rule set—removing a rule, adding a rule, changing the default rule, changing a storage pool definition—without changing the policy name trigger a minor revision (for instance, the version number changes from 2.2 to 2.3). The 10th such change without a change to the policy name changes the version number to “10” (for instance, from 1.9 to 1.10).

8 Grid Tasks

A grid task is a predefined set of procedures that are executed on one or more grid services in the background. For instance, LDR foreground verification is performed via a grid task. Most maintenance and expansion procedures involve running grid tasks. For example, to decommission a Storage Node the content on that server must be migrated to other locations on the grid. This can involve thousands of objects being relocated. To execute this type of operation, a grid task can be executed and run in the background to normal grid activity.

Monitor Grid Tasks

You can follow the progress of a grid task from the **CMN > Grid Tasks > Overview** page (see Figure 123 and Table 37). Running grid tasks is restricted to accounts with Maintenance permissions such as the Admin and Vendor accounts.

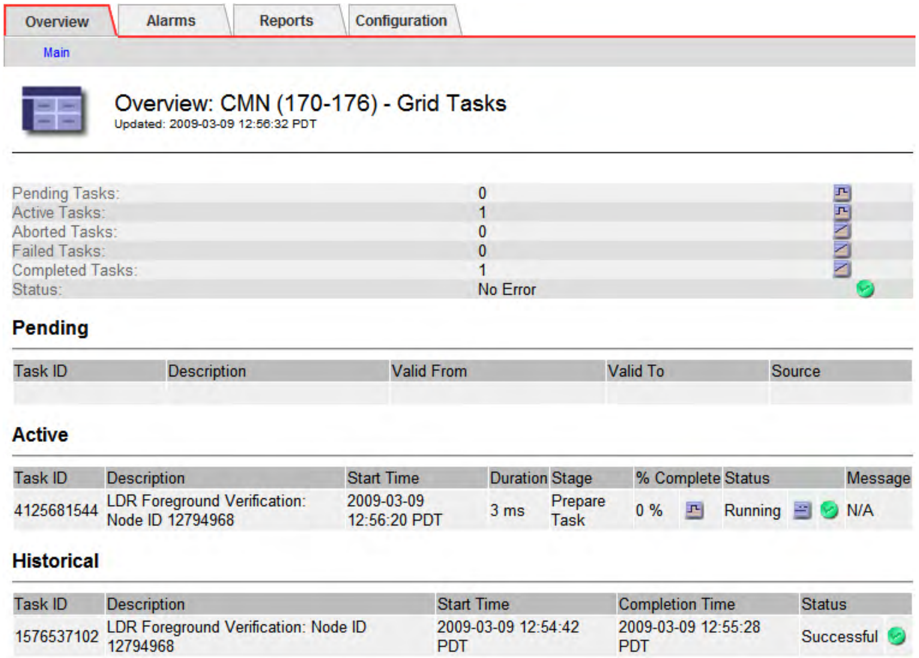

Overview: CMN (170-176) - Grid Tasks
Updated: 2009-03-09 12:56:32 PDT

Pending Tasks:	0
Active Tasks:	1
Aborted Tasks:	0
Failed Tasks:	0
Completed Tasks:	1
Status:	No Error ✓

Pending

Task ID	Description	Valid From	Valid To	Source

Active

Task ID	Description	Start Time	Duration	Stage	% Complete	Status	Message
4125681544	LDR Foreground Verification: Node ID 12794968	2009-03-09 12:56:20 PDT	3 ms	Prepare Task	0 %	Running ✓	N/A

Historical

Task ID	Description	Start Time	Completion Time	Status
1576537102	LDR Foreground Verification: Node ID 12794968	2009-03-09 12:54:42 PDT	2009-03-09 12:55:28 PDT	Successful ✓

Figure 123 Grid Tasks Overview

Grid tasks go through distinct phases:

Pending. The grid task has been submitted, but not started yet.

Active. The grid task has been started. It can be either actively running or temporarily paused.

Historical. A historical grid task is a task that has been submitted, but is no longer active. This includes grid tasks that completed successfully, grid tasks that were rejected (for example because the valid time period had expired), grid tasks that were cancelled or aborted, and grid tasks that terminated in error.

Table 37 Grid Tasks Overview Information

Field	Description
Task ID	Unique identifier assigned when the task is created.
Description	Brief description of the purpose of the task.
Valid From	Date from which the task is valid. The grid task will be rejected if it is submitted before this date.
Valid To	Date until which the task is valid. The grid task will be rejected if it is submitted after this date.
Source	The author of the grid task.
Start Time	Date and time on which the grid task was started.
Stage	Description of the current stage of the active task.
% Complete	Progress indicator for the current stage of the active tasks.
Duration	Amount of time since the grid task was started.

Table 37 Grid Tasks Overview Information *(continued)*

Field	Description
Status	Current status of the active or historical task. For active tasks, one of: • Starting • Running • Pausing • Paused (either paused by the user or automatically paused by the task) • Error: An error has been encountered. User action is required. • Aborting • Abort Paused: Task failed to be aborted and is paused in error. For historical tasks, one of: • Successful • Rollback Failed • Expired • Aborted • Cancelled • Unauthorized • Duplicate • Invalid
Message	Information about the last stage of the active task.
Completion time	The date and time on which the grid task completed (or cancelled or expired or was aborted).

Charting Grid Tasks

While active, the progress of grid tasks—particularly those that take a long time to complete—can be charted via the **CMN > Grid Tasks > Reports > Charts** page. This allows you to visual track a grid task's progress and determine if the task is stalled. For more information on stalled grid tasks, see Grid Task Hangs (page 242).

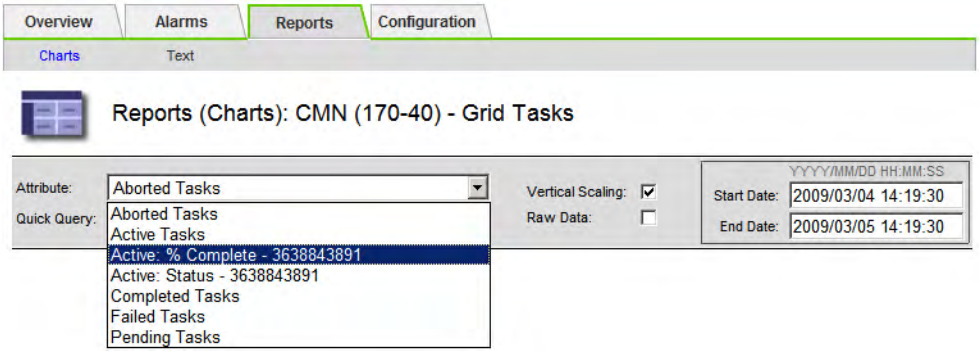

Reports (Charts): CMN (170-40) - Grid Tasks

Attribute: Aborted Tasks

Quick Query:

- Aborted Tasks
- Active Tasks
- Active: % Complete - 3638843891
- Active: Status - 3638843891
- Completed Tasks
- Failed Tasks
- Pending Tasks

Vertical Scaling: ☑

Raw Data: ☐

Start Date: 2009/03/04 14:19:30

End Date: 2009/03/05 14:19:30

Figure 124 Grid Task Chart

A grid task is only available for charting while it is active.

Grid Task Listing

Table 38 lists the most frequently used grid tasks.

Table 38 Common Grid Tasks

Code	Grid Task Name	Description
BCMT	Bundle Commit	May be used as part of a software upgrade procedure.
BDLI	Bundle Import	Used when a grid configuration bundle needs to be updated, usually as one step of a larger grid maintenance procedure. For example, it is used when converting a basic Gateway Node replication group to a High Availability Gateway replication group.
CMSC	Import Bundles for Metadata Replication	Used during the procedure to convert a grid using metadata synchronization to using metadata replication
CSRF	Storage Node hardware refresh —or— Control Node hardware refresh —or— Control/Storage Node hardware refresh	Used during the hardware refresh procedure for any server containing a Storage Node or a Control Node. Moves all content and/or metadata from the source Storage Node or Control Node or Control/Storage Node to the destination server, and then permanently removes the source server from the grid.
DDCM	Disable DICOM	Used to disable DICOM in a grid.
EDCM	Enable DICOM	Used to enable DICOM in a grid.
GDCM	Admin Node decommissioning —or— Gateway Node decommissioning —or— Admin/Gateway Node decommissioning	Used during the hardware refresh procedure for any server that contains an Admin Node or Gateway Node.
GEXP	Grid Expansion: Initial Grid Expansion: Add Server	Used during the expansion procedure to add servers to an existing grid.
ILME	ILM Evaluation	Used to apply a new ILM policy to existing content

Table 38 Common Grid Tasks *(continued)*

Code	Grid Task Name	Description
LBAL	LDR Content Rebalancing	Rebalances content across all storage volumes on a Storage Node. Used when adding additional storage volumes to an existing Storage Node. NOTE The LDR Rebalancing Grid Task is started from the ADE console, not from the NMS MI.
LDCM	Storage Node decommissioning	Moves all content off the specified Storage Node to other Storage Nodes and then permanently removes the Storage Node from the grid.
LFGV	LDR Foreground Verification	Verifies that all content on the specified Storage Node (LDR) exists. NOTE Foreground Verification of an LDR is initiated from the LDR > Verification > Configuration page in the NMS MI.
RNMC	Remove NMS Cluster Bindings	Used prior to the Admin Node decommissioning grid task in the hardware refresh procedure. Removes all attribute repository bindings to the NMS cluster.
SWUP	Software Upgrade: *Version*	Used when upgrading the grid software to a new version.

Mutually Exclusive Grid Tasks

Most grid tasks cannot be run concurrently because they require exclusive access to the same grid resources. See Figure 126 for a summary of which grid tasks can be run at the same time.

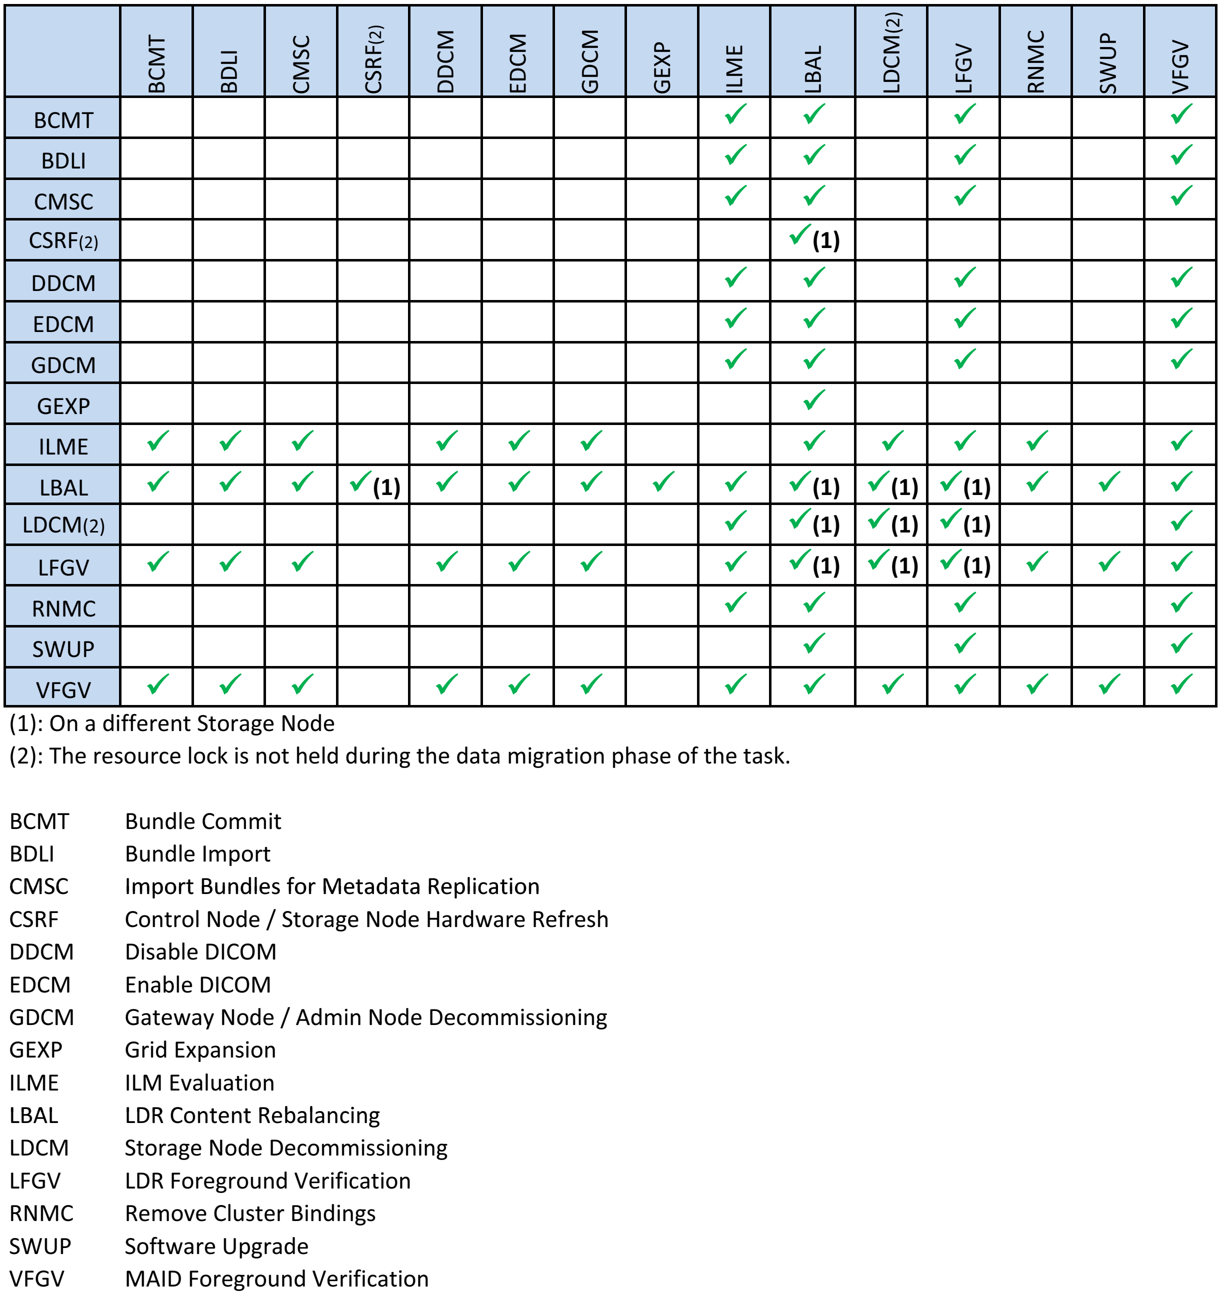

	BCMT	BDLI	CMSC	CSRF(2)	DDCM	EDCM	GDCM	GEXP	ILME	LBAL	LDCM(2)	LFGV	RNMC	SWUP	VFGV
BCMT									✓	✓		✓			✓
BDLI									✓	✓		✓			✓
CMSC									✓	✓		✓			✓
CSRF(2)										✓(1)					
DDCM									✓	✓		✓			✓
EDCM									✓	✓		✓			✓
GDCM									✓	✓		✓			✓
GEXP										✓					
ILME	✓	✓	✓		✓	✓	✓			✓	✓	✓	✓		✓
LBAL	✓	✓	✓	✓(1)	✓	✓	✓	✓	✓	✓(1)	✓(1)	✓(1)	✓	✓	✓
LDCM(2)									✓	✓(1)	✓(1)	✓(1)			✓
LFGV	✓	✓	✓		✓	✓	✓		✓	✓(1)	✓(1)	✓(1)	✓	✓	✓
RNMC									✓	✓		✓			✓
SWUP										✓		✓			✓
VFGV	✓	✓	✓		✓	✓	✓		✓	✓	✓	✓	✓	✓	✓

(1): On a different Storage Node

(2): The resource lock is not held during the data migration phase of the task.

BCMT	Bundle Commit
BDLI	Bundle Import
CMSC	Import Bundles for Metadata Replication
CSRF	Control Node / Storage Node Hardware Refresh
DDCM	Disable DICOM
EDCM	Enable DICOM
GDCM	Gateway Node / Admin Node Decommissioning
GEXP	Grid Expansion
ILME	ILM Evaluation
LBAL	LDR Content Rebalancing
LDCM	Storage Node Decommissioning
LFGV	LDR Foreground Verification
RNMC	Remove Cluster Bindings
SWUP	Software Upgrade
VFGV	MAID Foreground Verification

Figure 125 Grid Tasks That Can Be Run Concurrently

The Storage Node/Control Node hardware refresh task (CSRF) and the LDR decommissioning task (LDCM) release their lock after the actual data migration phase starts. In other words, you can run another task, for example, the grid expansion (GEXP) task during the migration phase of the Storage Node hardware refresh or LDR decommissioning task as long as there are no other resource contentions. Use the NMS MI to determine the phase of the grid task (see Monitor Grid Tasks (page 233)).

Run Grid Tasks

Under normal circumstances, grid tasks required for expansion, maintenance or update procedures appear automatically in the Pending list as part of the provisioning process. They are no longer provided in the form of a Task Signed Text Block.

During the period when a grid task is valid, it can be started from the list of Pending tasks. To run a grid task:

1. Go to **CMN > Grid Tasks > Configuration > Main**.
2. Select **Start** from the Actions menu for the Pending task you want to run.
3. Click **Apply Changes** to execute the task.

Information on the Configuration tab does not update automatically. Go to the Overview tab to monitor the progress of a task and come back to the Configuration tab to make further changes.

The grid task moves from the Pending list to the Active list. You must wait for the NMS MI page to auto-refresh before the change is visible. Do not submit the change again.

The task continues to execute until it completes, is paused or aborted manually (see Pause Grid Tasks).

When the task completes successfully, it moves to the Historical list with the description Successful in the Status field. If the task fails, it moves to the Historical list with a description of the error in the Status field.

Pause Grid Tasks

An active grid task can be paused before execution is complete. This may be necessary if the grid becomes particularly busy and you need to suspend the background grid task for a while.

To pause an active grid task:

1. Go to **CMN > Grid Tasks > Configuration > Main**.
2. Select **Pause** from the Actions menu for the Active task you want to suspend temporarily.
3. Click **Apply Changes** to pause the task.

The grid task remains on the Active list with its status changed to Paused. This can be seen by returning to the Overview tab.

Resume Grid Tasks

A paused grid task can be resumed when conditions permit.

To resume a paused grid task:

1. Go to **CMN > Grid Tasks > Configuration > Main**.

2. Select **Run** from the Actions menu for the Active task you want to resume.
3. Click **Apply Changes** to resume the task.

The grid task remains on the Active list with its status restored to Active. This can be seen by returning to the Overview tab.

Cancel Grid Tasks

A grid task can be canceled from the Pending list so that it is no longer available for execution.

To cancel a grid task:

1. Go to **CMN > Grid Tasks > Configuration**.
2. Select **Cancel** from the Actions menu on the row of the Pending task that you want to cancel.
3. Click **Apply Changes** to cancel the task.

   The grid task moves to the Historical list with the description Cancelled in the Status field. You must wait for the page to auto-refresh before the change is visible. Do not submit the change again.

Abort Grid Tasks

A grid task can be aborted from the Active list so that it is no longer available for execution.

NOTE Not all grid tasks can be aborted. For instance, *never* abort the LDR content rebalancing grid task. To determine whether a grid task can be aborted, follow the guidelines in the procedures where the grid task is discussed.

To abort a grid task:

1. Go to **CMN > Grid Tasks > Configuration > Main**.
2. Select **Pause** from the Actions menu for the Active task you want to abort.
3. Click **Apply Changes**.

   When the page refreshes, the Status of the grid task has changed to Paused.
4. Select **Abort** from the Actions menu.
5. Click **Apply Changes**.

   The grid task moves to the Historical list with the description Aborted in the Status field. You must wait for the page to auto-refresh before the change is visible. Do not submit the change again.

> Aborting an active grid task triggers it to take programmed action to leave the affected entities in a reliable state. This may require time to roll back some actions or set appropriate device states. The task may remain on the Active list with a status of "Aborting" for a period of time. When the programmed abort process is complete, the grid task moves to the Historical list.

You can run a task you aborted again. First remove it from the historical list (see Remove Grid Task from the Historical List) and then submit the grid task again using the task signed text block (see Submit a Task Signed Text Block (page 243)).

Remove Grid Task from the Historical List

To remove a task from the Historical list:

1. Go to **CMN > Grid Tasks > Configuration > Main**.
2. Select the **Remove** box in the row of the task.
3. Click **Apply Changes**.

Do not click Apply Changes more than once. Wait for the page to refresh.

Troubleshooting

Grid Task Fails And Moves to Historical List

Check the explanation in the Status field. Typical reasons for load failure are:

- aborted—the task was aborted
- cancelled—the task was cancelled
- duplicate—the task has been previously loaded into the CMN
- expired—the "task valid before" time has already passed
- invalid—the task was not valid
- unauthorized—the task signature did not pass verification

The most common reason for a grid task to fail is that the grid task has expired before it is started. If a task expires, it can never be run. A new task must be created and run instead.

To try to run the task again, you must first remove it from the Historical list.

Grid Task Hangs

If a grid task halts and you can not identify a reason, try (in order):

- Pause and then restart the task
- Restart the CMN
- Abort the task, remove it from the historical table, then resubmit it
- Contact HP Support.

Grid Task in Error State

If a task enters an error state, the Grid Task Status alarm (SCAS) are triggered. Look up the task message on the CMN > Grid Tasks > Overview page. It will display some information about the error (for example, "check failed on node 12130011"). After you have investigated and corrected the problem, start the task again (go to **CMN > Grid Tasks > Configuration** and select **Run** from the **Actions** menu).

Aborting Grid Tasks

If a task you are aborting enters an Internal Error state, try to abort the task again (go to **CMN > Grid Tasks > Configuration** and select **Abort** from the **Actions** menu). If the task is still unable to complete the abort sequence, contact HP Support.

Submit a Task Signed Text Block

If the grid task you need to run is not in the Pending list, you can load the grid task by submitting the Task Signed Text Block.

To load a grid task:

1. Get the grid task from the Grid_Tasks folder of the SAID package.
2. Copy the Task Signed Text Block file to the same computer that you will use to access the NMS MI.
3. Open the file that contains the grid task (Task Signed Text Block) using WordPad or a programmer's text editor.
4. Copy the Task Signed Text Block to the clipboard:
   a. Select the text including the opening and closing delimiters:.

      ```
      -----BEGIN TASK-----
      AAAOH1RTSUJDT05UAAANB1RCTEtjbmN0AAAM+1RCTEtDT05UAAAA
      EFRWRVJVSTMyAAAAAQAAABBUU0lEVUkzMoEecsEAAAAYVFNSQ0NTV
      ...
      s5zJz1795J3x7TWeqBAInHDVEMKg95O95VJUW5kQij5SRjtoWLAYXC
      -----END TASK-----
      ```

      If the block has a readable description above the opening delimiter, it may be included but is ignored by the HP MAS system.

   b. Copy the selected text.
5. Log in to the NMS MI using the Vendor account.
6. Go to **CMN > Grid Tasks > Configuration > Main**.

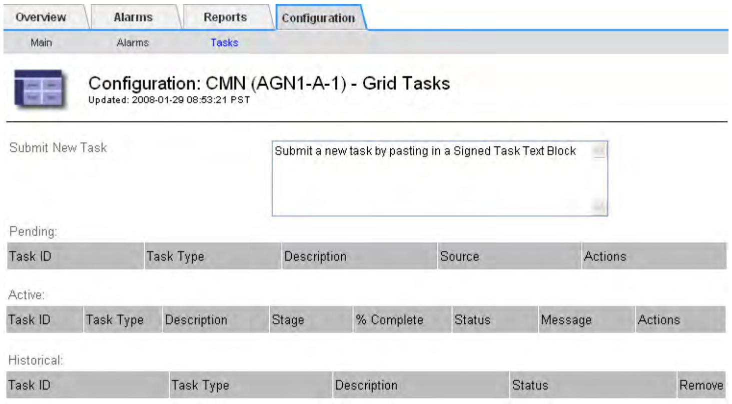

Configuration: CMN (AGN1-A-1) - Grid Tasks
Updated: 2008-01-29 08:53:21 PST

Submit New Task: Submit a new task by pasting in a Signed Task Text Block

Pending:

Task ID	Task Type	Description	Source	Actions

Active:

Task ID	Task Type	Description	Stage	% Complete	Status	Message	Actions

Historical:

Task ID	Task Type	Description	Status	Remove

Figure 126 CMN Grid Task Configuration

7. Select the prompt text in the Submit New Task text box so that you can replace it by the Task Signed Text Block.
8. Paste the Task Signed Text Block into the text field, replacing the prompt.
9. Click **Apply Changes** to load the grid task.

   The grid validates the Task Signed Text Block and either rejects the grid task or adds it to the list of Pending tasks.

Grid Performance

Grid-wide tasks are prioritized lower than normal grid operations (such as processing retrieval requests from clients). However, because running grid-wide tasks puts additional load on the system, they may have an impact on overall grid performance at times of peak load. Moreover, if a number of grid-wide tasks are in progress, they compete for resources, delaying the completion of individual tasks.

If either of these situations occur, use the **CMN > Grid Tasks > Overview** tab to identify which tasks are in progress, and which are closest to completion. This permits you to select which tasks to pause to either reduce the overall load on the grid, or to permit selected tasks to complete more quickly.

9 Server Manager

Server Manager is an application that runs on every server in the grid. Server Manager supervises the starting and stopping of services on the server, ensuring services gracefully join and leave the grid. It also monitors services on the server and attempts to restart any that reports faults, without the need for manual intervention.

Grid services have a variety of dependencies on support packages such as networking, timing synchronization, and third-party programs such as databases. Server Manager ensures that these support services are brought online in the correct sequence to meet dependencies.During system startup, Server Manager is automatically started by the operating system. Server Manager then executes a sequential series of scripts to verify that support services are running, and starts them as needed. The startup and shutdown sequences are reversed, ensuring that dependent services are in place as needed, and are not removed prematurely.

Server Manager provides the following capabilities:

- Stop and start all services on a server for:
  - Restarting grid services that have gone offline
  - Stopping the services in preparation for an upgrade
  - Bringing up the services after a reconfiguration
- Monitor services on an ongoing basis and restart them as needed.
- Automatically start services if the server is power cycled or reset, and to recover from unintentional restarts.
- Detect OS shutdown and gracefully close services.
- Restart a server (bring down everything, including the OS, and reboot the machine from the BIOS up).
- Shut down a server to the point where the server must be manually restarted. This enables you to safely power down a server for hardware maintenance.

About the GUI

Server Manager's graphical interface on the local console of the server enables monitoring and coarse (whole server) control. The state of services is reported, along with server identity information such as IP addresses. Control buttons allow you to stop services to perform upgrades, reboot the server, or shut the server down for hardware maintenance.

The display of service states is updated synchronously with changes in the services. However, in some cases there may be processing delays of up to 15 seconds before the GUI accurately reflects the current state of all services.

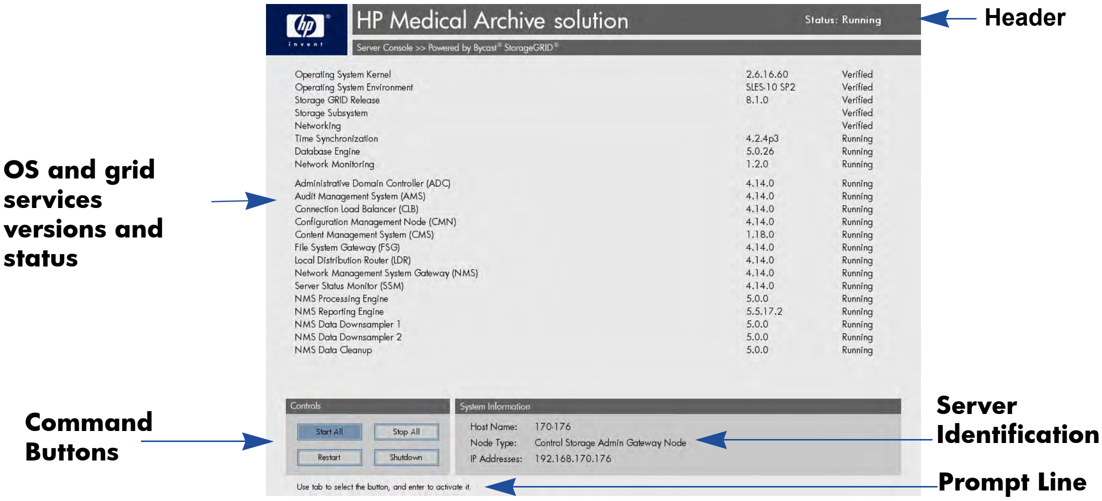

Operating System Kernel	2.6.16.60	Verified
Operating System Environment	SLES-10 SP2	Verified
Storage GRID Release	8.1.0	Verified
Storage Subsystem		Verified
Networking		Verified
Time Synchronization	4.2.4p3	Running
Database Engine	5.0.26	Running
Network Monitoring	1.2.0	Running
Administrative Domain Controller (ADC)	4.14.0	Running
Audit Management System (AMS)	4.14.0	Running
Connection Load Balancer (CLB)	4.14.0	Running
Configuration Management Node (CMN)	4.14.0	Running
Content Management System (CMS)	1.18.0	Running
File System Gateway (FSG)	4.14.0	Running
Local Distribution Router (LDR)	4.14.0	Running
Network Management System Gateway (NMS)	4.14.0	Running
Server Status Monitor (SSM)	4.14.0	Running
NMS Processing Engine	5.0.0	Running
NMS Reporting Engine	5.5.17.2	Running
NMS Data Downsampler 1	5.0.0	Running
NMS Data Downsampler 2	5.0.0	Running
NMS Data Cleanup	5.0.0	Running

Host Name: 170-176
Node Type: Control Storage Admin Gateway Node
IP Addresses: 192.168.170.176

Use tab to select the button, and enter to activate it.

Figure 127 Server Manager Interface

Header

The header line identifies the Server Manager application and its overall status. When running normally, the header has a dark grey background. If any of the services is reporting an error, the background changes to red to draw immediate attention.

The status line is an aggregate representation of the state of the services under its control. If they are all stopped, the status is Stopped; all running, Running, and so on. If any of the services are in an error state, the Status is Error. Similarly, if any of the services are in the Stopped state, the Status is Stopped.

Service List

The body of the display lists the services monitored by Server Manager on this server.

Service Name—The name of each service is shown. Some services may appear that are not identified through the NMS MI. These are services that can run independently and provide support capabilities to the grid services.

Version—Where a version number is available from the service, it is displayed in the list. This provides a quick reference for administrators to verify the currency of services and identify if updates are required.

Status—Nominally these all report “Verified” (operating system services) or “Running”. As the services are being started or stopped, the status may report the transition stage, such as “Stopping...” or “Starting...”. Additionally, “Stopped” indicates a service that has been ordered stopped by the Server Manager and will not restart without a Server Manager command.

Command Buttons

Command buttons are used to stop and start services and reboot or shut down the server. Use of these buttons ensures the services enter and leave the grid gracefully.

Start All

Use the Start All button to start services that have been stopped for upgrade or other configuration changes. Server Manager executes a series of scripts to initiate dependent applications in a sequence that ensures prerequisite services are running before starting the grid services themselves. All existing error states in all service are cleared.

When all services are started, the server joins the grid. Any services already running are not disrupted.

Stop All

Use the Stop All button to gracefully stop services hosted on the server, effectively disconnecting it from the grid. The operating system and Server Manager continue to run on the server, allowing you to perform tasks such as configuration changes, software updates, and similar maintenance that requires the operating system. Grid services and third-party applications are stopped.

Restart

Use the Restart command to stop grid services and bring down the operating system to perform an automatic reboot. This is useful for resetting a server that has failed or when starting a new configuration.

The Server Manager performs the same sequence as the Stop All command, then continues to bring down the operating system. The Server Manager and GUI are closed by the operating system as part of its shutdown. Settings in the operating system are used to trigger a reboot, which in turn restarts the Server Manager application. Server Manager then restarts the GUI and all services.

Shutdown

Use the Shutdown command to power down the server for hardware maintenance, upgrades, or reconfiguration. Similar to Restart, this command gracefully closes services and halts the operating system. Unlike Restart, this command does not trigger a reboot. The system is fully halted, and power is turned off. Power must be turned on to start the system.

The Server Manager performs the same sequence as the Stop All command (see Stop All (page 247)), then continues to bring down the operating system to a halted state. The Server Manager and GUI are closed by the operating system as part of the shutdown.

Initiating an Action

To initiate an action:

1. Use the mouse or the **<Tab>** and **<Enter>** keys to select the desired command button.

   A confirmation dialog opens allowing you to confirm or cancel the action.

2. Type the Server Manager password listed in the `Passwords.txt` file.

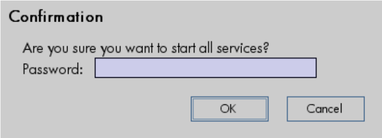

Confirmation

Are you sure you want to start all services?

Password:

Figure 128 Password Confirmation

3. Click **OK** to confirm the action.

Server Identification

The bottom right area of the display provides information about the server itself.

Node Type—This name is allocated at installation to describe the type of grid services hosted on the server, such as “Control Node” or “Storage Node”. If a Control Node and Storage Node are both installed on the same server, they will be reported as “Control Storage Node.”

Host Name—This is the name of the server specified when the server is added to grid at configuration time.

IP Addresses—The server’s operating system networking service is used to report the IP addresses assigned to the network interface(s). Most servers have only one address; however, if multiple adapters/addresses are available, all are reported.

Prompt Line

The prompt line provides guidance on what actions you can take. This typically states: “Use tab to select the button, and enter to activate it.” For more information, see Command Buttons (page 247).

At the server, press **<Alt>+<F1>** to access a command shell and log in as root using the password listed in the `Passwords.txt` file.

At the server, press **<Alt>+<F1>** to access a command shell and log in as root using the password listed in the `Passwords.txt` file.

10 Alarm Reference

Alarm Reference Tables

This section lists the NMS MI pre-configured alarms alphabetically by alarm code. Responses are assigned according to the severity of the alarm. This may vary if you customize the alarm settings to fit your system management approach.

A

Code	Name	Recommended Action
ABRL	Available Attribute Relays	Restore connectivity to a service (an ADC) running an Attribute Relay Service as soon as possible. If there are no connected attribute relays, the device cannot report attribute values to the NMS. The NMS can no longer monitor the status of the service, or update attributes for the service. If this condition persists, contact HP Support.
ACMS	Available Metadata Services	When an LDR or an ARC does not have a connection to any CMS, it cannot process ingest or retrieve transactions. As a result, ingest and retrieve operations to an LDR that cannot connect to any CMS services may be delayed (if the unavailability or CMS services is only a brief transient issue) or fail. The ACMS alarm may be triggered when a Control Node is shut down. It is safe to ignore this alarm when a Control Node is purposely shut down; for example, when performing maintenance procedures. This alarm clears automatically if other Control Nodes are online. Check and restore connections to a CMS service to clear this alarm and return the service to full functionality.
ADCA	ADC Device Status	If there is an Error, check the Overview and Alarm tabs for the device to find the cause of the error and to troubleshoot the problem. If the problem persists, contact HP alarms, 39 to 50 .
ADCE	ADC Device State	If the device goes into Standby, continue monitoring and if the problem persists, contact HP Support.

Code	Name	Recommended Action
AITE	Archive Retrieve State	If the state is Waiting for Middleware, check the middleware server and ensure that it is operating correctly. If the service has just been added to the grid, ensure that the middleware server is correctly configured. If the alarm text is Offline, try bringing Archive Retrieve online using the drop down on the ARC > Retrieve > Configuration tab.
AITU	Archive Retrieve Status	If the status is Middleware Error, go to the Middleware server to check for errors. If the status is Session Lost, check the middleware server to ensure it is online and operating correctly. Check the network connection with the middleware server. If the status is Unknown Error, contact HP Support.
ALIS	Inbound Attribute Sessions	If the number of inbound attribute sessions on an attribute relay grows too large, it may be an indication that the grid has become unbalanced. This can lead to performance issues. Contact HP Support.
ALOS	Outbound Attribute Sessions	The ADC has a high number of attribute sessions, and is becoming overloaded. Contact HP Support.
ALUR	Unreachable Attribute Repositories	Check network connectivity with the NMS to ensure that the service can contact the attribute repository. If network connectivity is good, contact HP Support.
AMQS	Audit Messages Queued	If audit messages cannot be immediately forwarded to an audit relayor repository, the messages are stored in a disk queue. During heavy loads the queue can go over 100,000, but at that level it should be closely monitored. If the alarm is triggered, check the load on the system—if there have been a significant number of transactions this may be normal and will resolve itself over time. If the alarm persists, view a chart of the queue size. If the number continues increasing without ever decreasing, contact HP Support. In rare instances, the disk queue may be large enough to cause a thread deadlock when the AMS starts. Contact HP Support if this occurs.
AOTE	Archive Store State	If the state is Waiting for Middleware, check the middleware server and ensure that it is operating correctly. If the service has just been added to the grid, ensure that the middleware server is correctly configured. If the Store component is offline, check the Archive Store status and correct any problems before moving the Archive Store State back to Online.

Code	Name	Recommended Action
AOTU	Archive Store Status	If the status is Session Lost check that the middleware server is online. If the status is Middleware Error check the middleware server for errors. If the Status is Unknown Error, contact HP Support.
ARCA	ARC Status	If the status is Waiting for CMSs, confirm that at least one read-write CMS is available.
ARCE	ARC State	The ARC service waits in Standby until all ARC subsystems (Replication, Store, Retrieve, Middleware) have started, and then transitions to Online. If the Standby state does not clear, check the status of the ARC subsystems.
AROQ	Objects Queued	There is a Notice alarm sent when Objects Queued reaches 3000, and a Minor alarm when it reaches 6000. This alarm may occur if the removable storage device is running slowly due to problems on the middleware server, or if the middleware server encounters multiple read errors. Check the middleware server for errors, and ensure that it is operating correctly. In some cases, this error may occur as a result of a high rate of data requests. Monitor the number of objects queued as grid activity declines. Note that a single CT exam could have 1000-2000 images.
ARRC	Remaining Capacity	If the remaining capacity of the attribute repository drops too low, the grid may require expansion. Contact HP Support.
ARRF	Request Failures	If an object retrieval from archive media fails, the Tape Node retries the retrieval as the failure may be due to a transient issue. However, if the object is corrupted on tape or the tape has been marked as being permanently unavailable by the TSM administrator, the retrieve action does not fail—the Tape Node continuously retries, and the value of ARRF continues to increase. This alarm may indicate that the storage media holding the requested data has become corrupt. See the middleware server to further diagnose the problem. If it is found that the object is no longer on the archive media, the object will have to be removed from the grid. Contact HP Support for assistance. Once the problem that caused this alarm is addressed, you can reset the count of failures on the configuration tab of ARC > Retrieve.
ARRS	Repository Status	If the attribute repository becomes disconnected from the NMS interface, contact HP Support.

Code	Name	Recommended Action
ARRV	Verification Failures	Contact HP Support to diagnose and correct this problem. Once the problem has been identified and corrected, reset the count of failures on the configuration tab of ARC > Retrieve.
ARVF	Store Failures	This alarm may occur as a result of middleware errors or media errors. Check the middleware server to diagnose the problem. Once the issue has been identified and corrected, reset the counter of failures on the configuration tab of ARC > Store.
ASXP	Audit Share Exported	The alarm is triggered when the value is Unknown. It could indicate a problem with the installation or configuration of the software on the AMS server. Contact HP Support.
AUMA	AMS Status	If the service indicates DB Connectivity Error, restart the server. If this does not clear the problem, contact HP Support.
AUME	AMS State	If the device goes into Standby, continue monitoring and if the problem persists, contact HP Support.
AUXS	Audit Export Status	A minor alarm is triggered if an error occurs. Correct the underlying problem and then restart the AMS service. If errors persist, contact HP Support.

BC

Code	Name	Recommended Action
BASF	Available Object Identifiers	The CMN is allocated a fixed number of object identifiers when the grid is provisioned. In the unlikely event that the grid begins to exhaust its supply and triggers this alarm, contact HP Support to be allocated more identifiers.
BASS	Identifier Block Allocation Status	A major alarm is triggered when object identifiers cannot be allocated because ADC quorum cannot be reached. Identifier block allocation on the CMN requires a quorum (50% + 1) of the ADCs to be online and connected. If quorum is unavailable, the CMN is unable to allocate new identifier blocks until quorum is re-established. If quorum is lost, there is generally no immediate impact on the grid (clients can still ingest and retrieve content), as the CMSs maintain a store (about a month's worth) of cached identifiers; however, if the condition persists eventually it will not be possible to ingest new content. If an alarm is triggered, investigate the reason for the loss of quorum (for example, network or Control Node failure(s)) and take corrective action. If the alarm persists, contact HP Support.
BTOF	Offset	You are notified if the service time (seconds) differs significantly from the operating system time. Under normal conditions, the service should resynchronize its time itself. If the time drifts too far from operating system time, grid operations can be affected.
BTSE	Clock State	You are notified if the service’s time is not synchronized with the time tracked by the operating system. Under normal conditions, the service should resynchronize its time itself. If the time drifts too far from operating system time, grid operations can be affected.
CAID	Available Ingest Destinations	Normally there is more than one in a grid. When there is only one, a notice is set; having none is a major alarm condition. These alarms clear when underlying issues of available LDR services are corrected. Ensure the DICOM (for CAID/CAQD) and HTTP (for CAIH/CAQH) components of LDRs are online and running normally.
CAIH		
CAQD	Available Q/R Destinations	
CAQH		
CIAF	Incoming Associations Failed	More than 50 is considered a minor alarm to investigate further. Failures could be caused by either end of the association. Check the CLB > DICOM component to determine if the failures are client side or node side. To reset the counter: on the **CLB > DICOM > Configuration** page, select **Reset DICOM Counts**. Client side faults indicate a problem with the remote entity. If consistent faults appear on the node side, contact HP Support.

Code	Name	Recommended Action
CLBA	CLB Status	If there is an Error, check the Overview and Alarm tabs for the device to find the cause of the error and to troubleshoot the problem. If the problem persists, contact HP Support.
CLBE	CLB State	If the device goes into Standby, continue monitoring and if the problem persists, contact HP Support. If the device is Offline and there are no known server hardware issues (server unplugged) or scheduled downtime for the device, contact HP Support.
CMNA	CMN Status	If there is an Error, check the Overview and Alarm tabs for the device to find the cause of the error and to troubleshoot the problem. The alarm No Online CMN is triggered during a hardware refresh of the primary Admin Node when the CMNs are switched (the old CMN service is put in standby and the new CMN service is put online). If the problem persists, contact HP Support.
CMNE	CMN State	If the device goes into Standby, continue monitoring and if the problem persists, contact HP Support.
CMSS	CMS State	If you receive an error on CMSS, contact HP Support.
CMST	CMS Status	If you receive an error on CMST, contact HP Support.
COAF	Outgoing Associations Failed	More than 50 is considered a minor alarm to investigate further. Failures could be caused by either end of the association. Check the CLB > DICOM component to determine if the failures are client side or node side. To reset the counter: on the **CLB > DICOM > Configuration** page, select **Reset DICOM Counts**. Client side faults indicate a problem with the remote entity. If consistent faults appear on the node side, contact HP Support.

Code	Name	Recommended Action
CPRC	Remaining Capacity	A minor alarm is triggered if the remaining capacity (number of available connections that can be opened to the database) falls below ten percent. If triggered, contact HP Support.
CsQL	Incoming CMS Synchronization Messages Queue Size	More than 50,000 is considered a minor alarm to investigate further. More than 1,000,000 is considered a critical condition. Check the load on the system—if there have been a significant number of transactions this may be normal and will resolve itself over time. If the alarm persists, view a chart of the queue size. If the number continues increasing without ever decreasing, contact HP Support.
CsQT	Outgoing CMS Synchronization Messages Queue Size	You receive a notification when the number of synchronization messages reaches 1000. More than 30,000 messages is considered a minor alarm. More than 300,000 messages is considered a major alarm. Investigate all alarms promptly. Ensure that other CMS services are online and running normally. Check the load on the system—if there have been a significant number of transactions this may be normal and will resolve itself over time. If the alarm persists, view a chart of the queue size. If the number continues increasing without ever decreasing, contact HP Support.

D

Code	Name	Recommended Action

DBSP	Free Table Space (Percent)	Adjusting the alarm threshold allows you to proactively manage when additional database storage needs to be added to the grid.
		The following levels are set by default:
		Notice = 25%
		Minor = 20%
		Major = 15%
		Critical = 5%
		When you receive a notice alarm, chart the available tablespace (**CMS > Database > Charts**) over a period of time to estimate the rate at which the available database space is being consumed. To maintain normal grid operations, you must add additional CMS capacity before the metadata database fills. Contact HP Support.

DCiN	Objects With Post-Processing Pending	More than 1000 is considered cause for investigation. If the alarm clears by itself, it may indicate that there was a short term overload only. If the number persists for some time after new content has been stored, contact HP Support.

Code	Name	Recommended Action
DEAE	Inbound C-Echo Failed	50 is considered a minor issue. These alarms indicate a problem with executing specific activities in the DICOM protocol. Check the log at the remote entity to determine if the errors are originating with the modality or the grid. If the modality shows the errors, then adjust the configuration of the entity. If the errors are on the grid side of the connection, contact HP Support. To reset the counters: on the **LDR > DICOM Configuration** page, select **Reset DICOM Counts**.
DEAF	Inbound C-Find Failed	
DEAI	Inbound C-Store Failed	
DEAM	Inbound C-Move Failed	
DEAO	Outbound C-Store Failed	
DEEO	Outbound C-Echo Failed	
DEIA	Inbound Associations Failed	
DEOA	Outbound Associations Failed	
DESC	Inbound Storage Commitment Failed	
DRIA	Incoming Associations - Rejected	
DROA	Outgoing Associations - Rejected	
DSN*n*	HD Status	The letter *n* indicates which RAID, as some servers are attached to multiple RAID units. All the drives are assigned to their respective array utilizing the Cid to distinguish each individual drive. Should a drive indicate Missing, Failed, or Unknown, check the physical hardware and resolve any hardware issues. Often the drive needs to be replaced with a spare. See also PSNn (page 270).

F

Code	Name	Recommended Action
FCCH	Flush Cache	A notice alarm is triggered when the Flush Cache option is enabled. The Notice alarm is a reminder that the option should be disabled once the cache has been sufficiently cleared in order to allow the FSG to start caching files again.
FCSA	Client Services Status	Restart the service. If the problem persists, contact HP Support.
FCSS	Status (Individual FSG client service status)	An alarm occurs if an error is detected for the client service. Contact HP Support.

Code	Name	Recommended Action
FCST	Cluster Status	If the cluster is in the Vulnerable state, one of the FSGs in the cluster is not available to provide services. Another failure may render FSG services unavailable. Investigate why the FSG is unavailable and correct the issue with the cluster. If the cluster is in the Failed state, the cluster is not providing FSG services to the grid. In an HAGC, neither FSG service is available. The cluster cannot recover from the fault automatically. Contact HP Support. The Failover state is deprecated as of release 8.1. If a failover within a replication group occurs, Failover Count (RPFO) is incremented by one on the new active primary FSG.
FDPP	Not used by the HP MAS.	N/A
FGWA	FSG Status	Waiting for configuration means that the FSG has not yet received configuration information from an ADC. Ensure that at least one ADC is online and reachable.
FGWE	FSG State	The device may go into Standby if it encounters disk I/O errors. Check for disk errors. After you correct them, restart the FSG. After the FSG has restarted, check to see if there was any data loss (for example, look for any loss of files being written at the time of the problem), and check to see if there are inconsistencies between FSGs in a replication group.
FOPN	Open File Descriptors	800 is considered an unusually high number. This value can become large during peak activity. If it does not diminish during periods of slow activity, contact HP Support.
FRGF	Files Retrieved from the Grid (Failed)	Values of 50 or more are a minor concern. Ensure all CMS and LDR services are operating normally. If the value continues to increase, contact HP Support. When the underlying problem is resolved, reset the counter to clear the alarm: go to **FSG > Storage > Configuration** and select **Reset Retrieve from Grid Failure Count**.
FRGP	Files Retrieved from the Grid (Pending)	Values of 20 or more are a minor concern. Ensure all CMS and LDR services are operating normally. Reduce demand on the grid if possible. If the problem persists, contact HP Support.
FRGR	Files Retrieved from the Grid (Retrying)	By default, a Minor alarm is triggered with a value of at least 15 and a Notice alarm is triggered with a value of at least 5. Monitor the situation. Check the state of LDRs and connectivity between FSGs and LDRs. If the values do not decrease, contact HP Support.

Code	Name	Recommended Action
FRPP	Files Pending for Replication	Objects are arriving faster than they can be replicated or objects cannot be stored to the grid by the primary FSG service. Ensure that the primary FSG service is operating correctly. Check connectivity. Monitor the situation. If the value of files pending does not decrease, contact HP Support. If the value does not decrease and all FSG services are operating correctly, and the Operations Not Applied attribute on the secondary FSG service and Files Stored to Grid - Pending attribute on the primary FSG service are zero, replication inconsistencies may exist.
FRSE	File Retrieve Sessions	A Major alarm is triggered if all established HTTP connections for content retrieve sessions are lost. The FSG is unable to establish HTTP GET sessions to the LDR. Check connectivity and check the state of the LDR. Go to **LDR > Storage** and check the **Storage State and Storage Status** attributes. If the problem persists, contact HP Support.
FSCS	Cache Status	A Major alarm indicates that Total Cache Available has dropped below the Swapout No Create watermark (NCWM) value set under **FSG Management > <*replication group*> > Configuration**. Monitor the situation. Wait for the FSG cache to swap out files and for the FSG's current free space to become higher than Swapout Free Space Watermark (FSWM). If free space does not increase, contact HP Support. NOTE To prevent cycling between NCWM and FSWM, file creation is disabled when free space drops below NCWM until free space goes back above the higher FSWM value.
FSGP	Files Stored to Grid—Pending	Objects are arriving faster than the grid can manage, or all Storage Nodes are unavailable. Values of 3000 or more are brought to your attention via a Notice level alarm, values of 10,000 or more are a minor concern, and values of 25,000 or more are a major concern. Ensure all CMS and LDR services are operating normally, and can be accessed by the FSG. Reduce ingests to the grid if possible. If the problem persists, contact HP Support.
FSGR	Files Stored to Grid—Retrying	By default, a Minor alarm is triggered with a value of at least 15 and a Notice alarm is triggered with a value of at least 5. Monitor the situation. Check the state of LDRs and connectivity between FSGs and LDRs. If the values do not decrease, contact HP Support.

Code	Name	Recommended Action
FSIP	Percentage Inodes Available	A Major alarm indicates that the percentage of available inodes has fallen below the configured value (default is 10%). Expand the grid with a new FSG replication group.
FSRI	Files Stored to Grid—Reingested	A notice alarm is triggered when no copies of a file that has to be swapped out from the cache exist in the grid and the value of Minimum Copies Before Swapout is non-zero. This causes the file to be re-ingested. FSRI should resolve on its own as the object is stored to a location in the grid. However, this alarm may need further investigation since it could be an indication that a failure in the grid caused all locations of the object to be lost. To reset the counter, go to **FSG > Storage > Configuration** and select **Reset Store to Grid Reingested Count**. If the problem persists, contact HP Support.
FSSE	File Store Sessions	A Major alarm is triggered if all established HTTP connections for content ingest sessions are lost. The FSG is unable to establish HTTP PUT sessions to the LDR. Check connectivity and check the state of the LDR. Go to **LDR > Storage** and check the Storage State and Storage Status attributes. If the problem persists, contact HP Support.

Code	Name	Recommended Action
FSST	Status	A notice alarm is triggered if share usage exceeds the configured quota. Monitor share usage and adjust quotas as needed. Clients using the share can also remove files from the share to set the share usage value back under the share quota.
FSTA	Total Cache Available	Values less than 20 GB trigger alarms. • Critical: 10 GB • Major: 15 GB • Minor 20 GB If this alarm is triggered, ensure all CMS and LDR services are operating normally. Check the following attributes in **FSG > Storage**: Files Stored to Grid-Pending, File Store Rate, and File Retrieve Rate. If these measures of grid activity are high, attempt to reduce demand until the alarm clears. If the problem persists, contact HP Support.
FSTS	Startup Condition	A Failed startup condition indicates that the FSG has failed to start up successfully, typically as a result of replication session inconsistencies or hardware faults. The server should be inspected and appropriate recovery actions taken. The value Dirty may be reported if the FSG does not start up cleanly. This may occur, for example, after a system or application crash, or after a power failure. The FSG will recover automatically and does not need to be restarted. Acknowledge the alarm and investigate the cause of the “dirty” startup. If additional unexpected “dirty” startups continue to occur, contact HP Support. If any value is displayed other than Init on initial install or Restored after the FSG is replaced or restored after a failure, the FSG has not been started cleanly. Inspect the system for other conditions. If the problem persists, contact HP Support. If there is an existing backup of the file system in the grid when the FSG is installed, the FSG will automatically initiate a file system restore. Restored may also be reported after initial install if a scheduled or forced backup has completed on another FSG prior to initial install.

HI

Code	Name	Recommended Action
HSTE	HTTP State	If you are using the FSG service, it is critical that the HTTP protocol be online and running without errors. Check the state of the LDR and the related Storage component. Ensure all are online. Check that the HTTP component is configured to autostart when the service is restarted. If the HTTP State of an LDR is Redirect, check the number of Available Metadata Services (on Events > Overview). If the number is zero, check and restore connectivity to at least one CMS.
HSTU	HTTP Status	
HTAS	Auto-Start HTTP	Specifies whether to start HTTP services automatically on startup. This is a user-specified configuration option.
IRSU	Inbound Replication Status	An alarm indicates that inbound replication has been disabled manually: check **LDR > Replication > Configuration**.

LM

Code	Name	Recommended Action
LATA	Average Latency	A Notice level alarm is triggered if the average latency period for NMS data connections is greater than 60 seconds (60000000 us). Check for connectivity issues. Check grid activity to confirm that there is an increase in grid activity. An increase in grid activity will result in an increase to attribute data activity. This increased activity will result in a delay to the processing of attribute data. This may be normal grid activity and will subside. If this alarm continues, contact HP Support. Check for multiple alarms on the grid. An increase in average latency times may indicate an excessive number of triggered alarms. If this continues, contact HP Support.
LATW	Worst-Case Latency	A Notice level alarm is triggered if a latency period for NMS data connections is greater than 120 seconds (120000000 us). Check for connectivity issues. Check grid activity to confirm that there is an increase in grid activity. An increase in grid activity will result in an increase to attribute data activity. This increased activity will result in a delay to the processing of attribute data. This may be normal grid activity and will subside. If this alarm continues, contact HP Support. Check for multiple alarms on the grid. An increase in average latency times may indicate an excessive number of triggered alarms. If this continues, contact HP Support.
LDRA	LDR Device Status	LDR devices need to be connected to at least one CMS to process transactions. If the service indicates Waiting for CMSs, ensure all CMS services are operating normally. If problems persist, contact HP Support.
LDRE	LDR Device State	If the device goes into Standby, continue monitoring and if the problem persists, contact HP Support.
LOST	Lost objects	The number of objects lost from the grid. Lost objects represent a loss of data. If this error occurs, contact HP Support to get help with investigating the cause of the problem and data recovery if possible.
MINQ	E-mail Notifications Queued	A Notice level alarm is triggered if notifications in the mail queue reach a count of 1000 or more. Check the network connections of the servers hosting the NMS service and the external mail server. Also confirm that the NMS e-mail server configuration is correct.
MINS	E-mail Notifications Status	A minor alarm is triggered if the NMS service is unable to connect to the mail server. Check the network connections of the servers hosting the NMS service and the external mail server. Also confirm that the NMS e-mail server configuration is correct.

Code	Name	Recommended Action
MISS	NMS Interface Engine Status	A minor alarm is triggered if the NMS interface engine on the Admin Node that gathers and generates interface content is disconnected from the grid. Check Server Manager to determine if the server individual application is down.
MMQS	Peak Message Queue Size	An alarm indicates that the node is overloaded, and may not be able to process operations at a high enough rate to HP Support normal grid operation. Client requests may time out when nodes are in this condition. Contact HP Support.
MRpe	Metadata with ILM Evaluation Pending	Either objects are arriving into the grid faster than the grid can evaluate the ILM for the object metadata, or a large number of objects that require an ILM re-evaluation are being processed. (For example, if the ILM changes the location of data based on content age, then the ILM for these objects must be re-evaluated and the metadata locations updated.) Values of 1,000 or more are brought to your attention via a Notice level alarm. Values of 10,000 or more are a minor concern and values over 250,000 a major concern. Plot the value of MRpe over the course of a day or week, and check that at times of low grid activity the number of metadata items with ILM evaluation pending drops, and tends towards zero. If the problem persists, contact HP Support.
MRun	Metadata with Unachievable ILM Evaluations	Check that all sites in the grid are connected and available, specifically that all CMSs are available. Find and correct any issues.
MSTE	DICOM State	Check the state of the LDR and the related Storage component. Ensure all are online. Check that the DICOM component is configured to autostart when the service is restarted. If the DICOM State of an LDR is Redirect, check the number of Available Metadata Services (on **Events > Overview**). If the number is zero, check and restore connectivity to at least one CMS.
MSTU	DICOM Status	Restart the service. If the error persists, it may indicate an installation problem. Contact HP Support.

NO

Code	Name	Recommended Action
NANG	Network Auto Negotiate Setting	Check the network adapter configuration. The setting must match preferences of your network routers and switches. An incorrect setting can have a severe impact on grid performance.
NDUP	Network Duplex Setting	Check the network adapter configuration. The setting must match preferences of your network routers and switches. An incorrect setting can have a severe impact on grid performance.
NLNK	Network Link Detect	Check the network cable connections on the port and at the switch. Check the network router, switch, and adapter configurations. Restart the server. If the problem persists, contact HP Support.
NRER	Receive Errors	These errors may clear without being manually reset. If they do not clear, check the network hardware. Check that the adapter hardware and driver are correctly installed and configured to work with your network routers and switches. When the underlying problem is resolved, reset the counter: on the **SSM > Resources > Configuration** page, select Reset Receive Error Count. On servers that use the bnx2 driver for the Broadcom Corporation NetXtreme II BCM5708 Gigabit Ethernet controller, the NMS may report spurious Network Receive Errors or Network Transmit Errors. These errors may clear without being manually reset. Disable the default alarm on these two attributes to eliminate these “nuisance” alarms.
NRLY	Available Audit Relays	If there are no connected Audit Relays (ADCs), the device cannot report audit events and they are queued and unavailable to users until the connection is restored. Restore connectivity to a service running an Audit Relay Service (an ADC) as soon as possible. If this condition persists, contact HP Support.
NSCA	NMS Device Status	If the service indicates DB Connectivity Error, restart the service. If the problem persists, contact HP Support.
NSCE	NMS Device State	If the device goes into Standby, continue monitoring and if the problem persists, contact HP Support.
NSPD	Speed	You receive a notice if the network connection is 10BaseT, and a minor alarm if the negotiated speed is reported as unavailable, unknown, or unsupported. This may be caused by network connectivity or driver compatibility issues. Contact HP Support.

Code	Name	Recommended Action
NTBR	Free Tablespace	If this alarm is triggered, check how fast database usage has been changing. A sudden drop (as opposed to a gradual change over time) indicates an error condition. Contact HP Support. Adjusting the alarm threshold allows you to proactively manage when additional storage needs to be allocated. If the available space is reaching a low threshold, contact HP Support to change the database allocation.
NTER	Transmit Errors	These errors may clear without being manually reset. If they do not clear, check the network hardware. Check that the adapter hardware and driver are correctly installed and configured to work with your network routers and switches. When the underlying problem is resolved, reset the counter: on the **SSM > Resources > Configuration** page, select **Reset Transmit Error Count**. On servers that use the bnx2 driver for the Broadcom Corporation NetXtreme II BCM5708 Gigabit Ethernet controller, the NMS may report spurious Network Receive Errors or Network Transmit Errors. These errors may clear without being manually reset. Disable the default alarm on these two attributes to eliminate these "nuisance" alarms.
NTFQ	NTP Frequency Offset	If the frequency offset exceeds the configured threshold, there is likely a hardware problem with the local clock. Contact HP Support to arrange a replacement.
NTLK	NTP Lock	If the NTP daemon is not locked to an external time source, check network connectivity to the designated external time sources, their availability, and their stability.
NTOF	NTP Time Offset	If the time offset exceeds the configured threshold, there is likely a hardware problem with the oscillator of the local clock. Contact HP Support to arrange a replacement.
NTSA	NTP Sources Available	If this server is configured to act as a primary NTP server for the grid, then this attribute tracks the number of external NTP time sources available. It is normal for this number to fluctuate if there are a large number of external time sources available. If the server is configured to act as a secondary NTP time server or an NTP client, the server uses other grid servers as its NTP time sources. If the number of NTP time sources available falls below the configured minimum, the accuracy and consistency of local time on the server may suffer. If the number of NTP time sources falls to zero, local server time will drift out of synchronization with the time recorded by other services. In extreme cases, this can disrupt grid operations. Correct the issue as quickly as possible.

Code	Name	Recommended Action
NTSD	Chosen Time Source Delay	These values give an indication of the reliability and stability of the time source that NTP on the local server is using as its reference. If an alarm is triggered, it may be an indication that the time source's oscillator is defective, or that there is a problem with the WAN link to the time source.
NTSJ	Chosen Time Source Jitter	
NTSO	Chosen Time Source Offset	
NTSU	NTP Status	If the NTP daemon is not running, contact HP Support.
OCOR	Corrupt Objects Detected	Any corrupt objects are worthy of investigation. More than 10 indicates a major problem. If there are several corrupt objects identified, contact HP Support. Note that this value is persistent: it is not updated once the corrupt objects have been restored. If corrupt objects are detected, change the background verification priority from Adaptive to High to speed up verification and assess the magnitude of the problem (on the **LDR > Verification > Configuration** page, select **High** in the Verification Priority box) After the underlying problem is resolved, reset the counter to clear the alarm: on the **LDR > Verification > Configuration** page, select **Reset Corrupt Objects Count**.
OQRT	Objects Quarantined	After the objects are automatically restored by the grid, the quarantined objects must be manually removed from the quarantine directory. Contact HP Support. After the quarantined objects are removed, the value of OQRT value is updated and the alarm will clear.

Code	Name	Recommended Action
ORpe	Objects with ILM Evaluation Pending	Either objects are arriving into the grid faster than the grid can evaluate the ILM rules for these objects, or a large number of objects that require an ILM re-evaluation are being processed. (For example, if the ILM changes the location of data based on content age, then the ILM for these objects must be re-evaluated and the number or type of object locations updated.) Values of 1,000 or more are brought to your attention via a Notice level alarm. Values of 10,000 or more are a minor concern and values over 250,000 a major concern. Plot the value of ORpe over the course of a day or a week, and check that at times of low grid activity the number of objects with ILM evaluation pending drops, and tends towards zero. This alarm may also be triggered during CMS evaluation at the deferred time.
ORSU	Outbound Replication Status	An alarm indicates that outbound replication is not possible: storage is in a state where objects cannot be retrieved from it. A notice alarm is triggered if outbound replication has been disabled manually: check **LDR > Replication > Configuration**. A major alarm is triggered if the LDR is unavailable for replication: check the storage component of the LDR.
ORun	Objects with Unachievable ILM Evaluations	Check that all sites in the grid are connected and available. Check that the storage at each site is online, and has not moved into a read only state. Find and correct any issues.

PR

Code	Name	Recommended Action
PAT*n*	RAID controller status	Under certain circumstances the monitoring module can become shutdown or disconnected. Once it is brought back up, it initializes and rebuilds the array if necessary, then goes into Online state again. Check the RAID monitoring serial cable at the back of the server to make sure it is connected properly. If this problem persists for a period of time, unplug the serial cable from the back and plug it back in. If the RAID monitoring still does not initialize, contact HP technical support.
PBST	Backup Status	If the backup reports Failed, ensure there are no coincident HTTP alarms on LDR services and that capacity remains on some Storage Nodes. Backups are run automatically each day, and retained for two weeks.
PGFT	Page Fault Rate	If the page fault rate is too high, the software is spending too much time swapping information in and out of physical memory. The server may need more physical memory. Contact HP Support.
PMEM	Service Memory Usage (Percent)	Can have values Over Y% RAM where Y represents the percentage of memory being used by the server. Figures under 80% are normal. Over 90% is considered a major issue. If the memory usage is fairly high for a single device/service, this should be monitored and investigated. Should it reach over 90% RAM and this alarm continues to persist, contact HP Support.
PSA*n*	Array State	If the array is offline, check the physical hardware to make sure it is alright and that it is cabled correctly. If this does not appear to be a hardware problem, contact HP Support.

Code	Name	Recommended Action
PSN*n*	Array Status	If the code is not OK (0), investigate the cause of the alarm. Rebuilding (1): Failed drive(s) have been replaced, and the controller is currently rebuilding redundant parity information. Wait for a period of time for the array to restore to normal state; if the problem persists contact HP Support. Ready to Rebuild (2): Failed drive(s) have been replaced, and the controller is about to begin rebuilding redundant parity data. Wait for a period of time for the array to restore to normal state; if the problem persists contact HP Support. Recovery (3): Although the logical drive can still operate, one or more drives has failed. Replace the failed drives as soon as possible. Drive Fail (4) If a drive has failed, replace the drive and allow time for it to rebuild. Unknown (5): contact HP Support. Failed (6): The array has failed. Run hardware diagnostics and checking for loose or failed cables. If the problem persists, contact HP Support. See also DSNn (page 257).
RDTE	Archive Middleware State	If the state is Offline, check the status of the Archive Middleware component (RDTU), and resolve any problems. Try bringing the component back online using the drop-down on the **ARC > Middleware > Configuration** tab.
RDTU	Archive Middleware Status	If the status is Configuration Error and the service has just been added to the grid, ensure that the middleware server is correctly configured. If the status is Connection Failure, or Connection Failure, Retrying check the network configuration on the middleware server, and the network connection between the server and the grid. If the status is Authentication Failure, or Authentication Failure, Reconnecting the grid can connect to the middleware server, but cannot authenticate the connection. Check that the middleware server is configured with the correct user, password, and permissions, and restart the service. If the status is Session Failure, an established session was lost unexpectedly. Check the network connection between the middleware server and the grid. Check the middleware server for errors. If the status is Unknown Error, contact HP Support.

Code	Name	Recommended Action
RIRF	Inbound Replications - Failed	The threshold for a notice alarm is 10 objects, while greater than 50 objects triggers a minor alarm. Replication alarms (Inbound Replications - Failed RIRF and Outbound Replications - Failed RORF) can occur during periods of high load or due to temporary network disruptions. After grid activity goes back down, these alarms should clear. If the count of failed replications continues to increase, look for network problems and verify that the source and destination LDRs and the ARCs are online and available. To reset the count, go to **ARC** or **LDR > Replication > Configuration**, and select **Reset Inbound Replication Failure Count**.
RIRQ	Inbound Replications - Queued	The threshold for a notice alarm is 5000 objects, and 10,000 objects for a minor alarm. Ensure that the CMS is online and running without errors with synchronization. If the problem persists, contact HP Support.
RMSn	RAID Monitor Status	A status of Disconnected does not mean that the RAID is faulty, only that the SSM is unable to monitor it. If the monitor is disconnected for an extended period, restart the server. If this does not clear the problem, contact HP Support.
RORF	Outbound Replications - Failed	The threshold for a notice alarm is 10 objects, while greater than 50 objects triggers a minor alarm. Replication alarms (Inbound Replications - Failed RIRF and Outbound Replications - Failed RORF) can occur during periods of high load or due to temporary network disruptions. After grid activity goes back down, these alarms should clear. If the count of failed replications continues to increase, look for network problems and verify that the source and destination LDRs and the ARCs are online and available. To reset the count, go to **ARC** or **LDR > Replication > Configuration**, and select **Reset Outbound Replication Failure Count**.
RORQ	Outbound Replications - Queued	The threshold for a notice alarm is 5000 objects, and 10,000 objects for a minor alarm. The outbound replication queue contains both objects being replicated by business rules and objects requested by clients. An alarm may occur as a result of a grid overload. Wait to see if the alarm clears when grid activity declines. If the alarm recurs, you may need to increase grid capacity by adding LDRs.

Code	Name	Recommended Action
RPER	Replication Errors	This alarm is related to the FSG Replication Errors system event found under **SSM > Events**. When the FSG encounters a replication error, it increments RPER and issues a log message with additional information about the replication error. This log message is picked up by the SSM events monitor to provide some visibility in the NMS MI of files that have been affected by replication errors. Investigate and clear any replication backlogs. To reset the counter: go to the **FSG > Replication > Configuration >** page, select **Reset Replication Errors Count**, and click **Apply Changes**.
RPFO	Failover Count	Investigate the cause of failovers. The failover count is incremented on the new active primary FSG within a replication group whenever a failover occurs. This includes manual failovers initiated through the FSG Management component of the NMS MI and automatic failovers within a High Availability Gateway Cluster. To reset the count: on the **FSG > Replication > Configuration** page, select **Reset Failover Count**.
RSST	Restore Result	An alarm is triggered if the result of the previous restore process is Failure. A failure could be due to a transient condition such as a temporary network disruption that requires no action, or a problem such as a failed disk on the FSG that requires action to resolve. If the result is Failure, identify and correct the cause of the restore failure, and restart the restore procedure from the beginning.
RSTU	Replication Status	If the device reports No Primary or No Session, ensure the primary FSG service of the replication group is online and running normally. An alarm is triggered if replication sessions between the primary and secondary FSG service become inconsistent with each other. If the status is Error, contact HP Support.

ST

Code	Name	Recommended Action
SAVP	Total Usable Space (Percent)	If usable space is reaching a low threshold, you should start looking into either purchasing additional storage or migrating some data to archive, depending on your available options. Adjusting the alarm threshold allows you to proactively manage when additional storage needs to be purchased.
SCAS	Grid Task Status	If the status is Error, look up the task message. It should display some information about the error (for example, "check failed on node 12130011"). Once you have investigated and corrected the problem, start the task again (go to the Configuration page and select **Run** from the **Actions** menu). If a task you are aborting enters an Internal Error state, try to abort the task again (go to the Configuration page and select **Abort** from the **Actions** menu). If the task is still unable to complete the abort sequence, contact HP Support as it will be necessary to use an ADE console command to unload the task from the Active table.
SCHR	Historical Grid Task Status	If the status of the historical grid task is Aborted, investigate the reason and run the task again if required.
SHLH	Health of Object Store	Diagnose an error as a disk problem on the server. Check and correct: • problems with the volume being mounted • file system errors
SLSA	CPU Load Average	The higher the value the busier the system. Values above 15 should be investigated as it indicates a fairly high load on the system. Over 60 is considered a major issue. If the CPU Load Average persists at a high value, the number of transactions (seen by graphing network bandwidth usage) in the system should be investigated to determine whether this is due to heavy load at the time. If there does not appear to be a heavy load on the system, and the alarm persists, contact HP Support.
SMTT	Total Events	If the Total Events becomes non-zero, check if there are known events (such as network failures or RAID failures) that could have caused this. Unless these errors have been cleared (that is, the count has been reset to 0), total events alarms may be triggered. When an issue is resolved, you can reset the counter to clear the alarm: go to **SSM > Events > Configuration**. If there have not been any known events in the SSM, or the number increases and the alarm persists, contact HP Support.
SMST	Log Monitor State	If the Log Monitor has non-zero values persisting for a period of time, contact HP Support.

Code	Name	Recommended Action
SNST	Grid Task Framework Status	An alarm indicates that there is a problem storing the grid task bundles to the grid. For both types of alarms (Checkpoint Error and Quorum Not Reached), make sure that a majority of ADCs are connected to the grid (half + 1) and then wait for a few minutes. If the alarm does not clear, contact HP Support.
SPOP	Operations Not Applied	This alarm is triggered when the replication backlog is approximately equal to one day's worth of replications. It generally indicates that the replication backlog is growing continuously. This may occur because of the size of the backup and the amount of time that the backup needs to complete. It is much more likely to occur if the replication group uses offline backups. An FSG that performs offline backups does not process replication messages while it is performing a backup. Offline backups were standard for FSGs in software releases prior to release 7.5, but all FSGs should have been converted to online backups during the update to release 8.1 software. If the SPOP alarm is triggered on a grid that was originally installed prior to release 7.5, check to see if the FSGs have been updated to use online backups. Newer FSGs use online backups and continue to apply replication messages while the backup is running. Since the performance of the FSG is reduced somewhat while it is performing an online backup, SPOP backlog is possible but unlikely. If the SPOP alarm triggers, check how long the backup is taking to complete. (**FSG > Backup > Previous Backup**). If the backup is taking an excessive amount of time to complete (more than a few hours), contact HP Support to get the backup frequency changed to 2 days and the value of the alarm trigger altered.
SSMA	SSM Device Status	If there is an Error, check the Overview and Alarm tabs for the device to find the cause of the error and to troubleshoot the problem. If the problem persists, contact HP Support.
SSME	SSM Device State	If the device goes into Standby, continue monitoring and if the problem persists, contact HP Support.
SSTS	Storage Status	If the Storage Status has Insufficient Free Space, there is no more available storage on the server; data ingests are redirected to other available nodes. Data retrieval requests can continue to be delivered from this node. Additional storage should be added to this server. It is not impacting end user functionality, but the alarm persists until additional storage is added. The Volume(s) Unavailable message indicates that part of the storage is unavailable. Storage and retrieval from these volumes is not possible: check the volume health table for more information. If Storage Status has an error condition, contact HP Support.

Code	Name	Recommended Action
SUOP	Operations Not Committed	If this rises, it indicates that you are storing data faster than the secondary FSG can process transactions. This normally declines during periods of reduced grid activity. Stop storing files into the primary FSG of the group. Wait a few minutes to determine if the value of this attribute declines. When the alarm clears, resume normal activity on the primary FSG. If the value does not decrease, contact HP Support.
SUSP	Total Free Space (Percent):	Adjusting the alarm threshold allows you to proactively manage when additional storage needs to be purchased. If the total free space is reaching a low threshold, you should start looking into either purchasing additional storage or migrating some data to archive, depending on your available options. See also SAVP (page 273).
SVST	Service Status	This alarm clears when other alarms related to a non-running service are resolved. Track the source service alarms to restore operation. A service is listed with a status of Not Running when its state is Administratively Down. The service's status may be listed as Not Running for the following reasons: • The service is manually stopped (**`/etc/init.d/<service> stop)`** • There is an issue with the MySQL database and Server Manager shuts down the MI or CMS service • A server has been added to the grid, but is not started • During installation, a server has not yet connected to the Admin Node If a service is listed as Not Running, restart the service (**`/etc/init.d/<service> restart`**) If the problem continues, contact HP Support.
TPOP	Pending Operations	A backlog of messages may indicate that the ADC is overloaded. Too few ADCs may be connected to this ADC. In a large grid, the ADC may need adding computational resources, or the grid may require additional ADCs.

UV

Code	Name	Recommended Action
UMEM	Available Memory	Anything over 100 MB is considered normal. Below 50 MB is a major issue; below 10 MB is critical. If the available RAM gets low, determine whether this is a hardware or software issue. If it is not a hardware issue, or if available memory falls below 50 MB, contact HP Support *immediately.*
USWP	Available Swap	Anything over 1.2 GB is considered normal. Below 1 GB is a minor issue; below 500 MB is a major issue. If the available swap space gets low, the system may not handle peak loads as effectively. Contact HP Support.
VMFI	Entries Available	Values below 100,000 start to raise a concern. When the value dips below 25,000 the situation is a major issue that should be addressed. Values below 10,000 are considered critical. This is an indication that additional storage is needed on the server. Contact HP Support for assistance in installing and configuring new object storage.
VMFR	Space Available	Concern is not raised until the free space falls below 1 GB. At 100 MB the situation is considered major. Values below 10 MB are a critical concern. If the Free Space gets fairly low, it needs to be investigated as to whether there are log files growing out of proportion, or files taking up too much disk space that need to be reduced or deleted. If this problem persists, contact HP Support.
VMST	Volume Status	A minor alarm is triggered if the Volume Status is Unknown. An Unknown or Offline status may indicate that the volume cannot be mounted or accessed due to a problem with the file system or underlying storage device.
VPRI	Verification Priority	By default, storage verification priority is adaptive. If the priority is set to High, the NMS gives notice that this priority has been set because storage verification may slow normal operations of the service.
VSTU	Object Verification Status	Look for other problems on the Storage component. Check the Verification status. If it is Verify Location Synchronize Failed check that the service is connected to at least one CMS. Also check the operating system for any signs of block-device or file system errors. If the status is Maximum Number of Failures Reached, it usually indicates a low-level file system or hardware problem (I/O error) that prevents the Storage Verification task from accessing stored content. This message may also occur when there is a high number of content errors indicating that data was invalid. If the status is Unknown Error, contact HP Support.

X

Code	Name	Recommended Action
XAMS	Unreachable Audit Repositories	Check network connectivity to the server hosting the AMS; contact HP Support.
XCVP	% Completion	Notice alarm when Foreground Verification is completed.

Glossary

ACL
Access control list—Specifies what users or groups of users are allowed to access an object and what operations are permitted, for example read, write, and execute.

active primary FSG
In an HAGC, the FSG that is currently providing read-write service to clients. See also FSG replication group.

ADC
Administrative Domain Controller—A software component of the HP MAS system. The ADC service maintains topology information, provides authentication services, and responds to queries from the LDR, CMS, and CLB. The ADC service is found on the Control Node.

ADE
Asynchronous Distributed Environment—Proprietary development environment used as a framework for grid services within the HP Medical Archive solution Software.

Admin Node
A building block of the HP MAS system. The Admin Node provides services for the web interface, grid configuration, and audit logs. See also reporting Admin Node, processing Admin Node, primary Admin Node and HCAC.

AE title
Application Entity Title—The identifier of a DICOM node communicating with other DICOM AEs.

AMS
Audit Management System—A software component of the HP MAS system. The AMS service monitors and logs all audited system events and transactions to a text log file. The AMS service is found on the Admin Node—reporting Admin Node in a High Capacity Admin Cluster (HCAC).

API
Application Programming Interface—A set of commands and functions, and their related syntax, that enable software to use the functions provided by another piece of software.

ARC
Archive—A software component of the HP MAS system. The ARC service manages interactions with archiving middleware that controls nearline archival media devices such as tape libraries. The ARC service is found on the Tape Node.

Association
A connection protocol between two DICOM Application Entities (AEs), typically a local and remote AE. The AEs use the Association Establishment to negotiate the type of data to exchange and the format of data encoding.

audit message
Information about an event occurring in the HP MAS system that is captured and logged to a file.

atom

Atoms are the lowest-level component of the container data structure, and generally encode a single piece of information. (Containers are sometimes used when interacting with the grid via the HTTP API).

AutoYaST

An automated version of the Linux installation and configuration tool YaST (“Yet another Setup Tool”), which is included as part of the SUSE Linux distribution.

BASE64

A standardized data encoding algorithm that enables 8-bit data to be converted into a format that uses a smaller character set, enabling it to safely pass through legacy systems that can only process basic (low order) ASCII text excluding control characters. See RFC 2045 for more details.

Basic Gateway replication group

A Basic Gateway replication group contains a primary FSG and one or more secondary FSGs.

bundle

A structured collection of configuration information used internally by various components of the grid. Bundles are structured in container format.

business continuity failover

A business continuity failover within a Gateway Node replication group is one where a secondary Gateway Node is manually configured to act as a primary after the primary Gateway Node fails. Clients can continue to read and write to the grid after they are manually redirected to the acting primary. This is a temporary measure to maintain service while the primary Gateway Node is repaired.

cabinet

Used to house hardware components, a cabinet includes the physical rack and all power and network wiring required for an installation.

cabinet connectivity kit

Used to link cabinets. One Cabinet Connectivity Kit is required at Single Site installations that have more than one cabinet, and one is required at each site in a Single Site + DR installation. The Cabinet Connectivity Kit is installed in the Base Cabinet. See also WAN Connectivity Kit.

CBID

Content Block Identifier—A 64-bit number that uniquely identifies a piece of content within the HP MAS system. CBIDs are represented as a zero-padded, 16-character, hexadecimal number when used to refer to a unique piece of content using the HTTP interface.

CIDR

Classless Inter-Domain Routing—A notation used to compactly describe a subnet mask used to define a range of IP addresses. In CIDR notation, the subnet mask is expressed as an IP address in dotted decimal notation, followed by a slash and the number of bits in the subnet. For example, 192.0.2.0/24.

CIFS

Common Internet File System—A file system protocol based on SMB (Server Message Block, developed by Microsoft) which coexists with protocols such as HTTP, FTP, and NFS.

CLB

Connection Load Balancer—A software component of the HP MAS system. The CLB service provides a gateway into the grid for clients connecting via DICOM and HTTP protocols. The CLB service is part of the Gateway Node.

CMN

Configuration Management Node— A software component of the HP MAS system. The CMN service manages system-wide configuration and grid tasks. The CMN service is found on the primary Admin Node.

CMS

Content Management System—A software component of the HP MAS system. The CMS service manages content metadata and content replication according to the rules specified by the ILM policy. The CMS service is found on the Control Node.

command

In HTTP, an instruction in the request header such as GET, HEAD, DELETE, OPTIONS, POST, or PUT. Also known as an HTTP method.

connectivity kit

See cabinet connectivity kit and WAN Connectivity Kit.

container

A container is a data structure used by the internals of grid software. In the HTTP API, an XML representation of a container is used to define queries or audit messages submitted using the POST command. Containers are used for information that has hierarchical relationships between components. The lowest-level component of a container is an atom. Containers may contain 0 to N atoms, and 0 to N other containers.

content block ID

See CBID.

content handle

See UUID.

consolidated Admin Node

Admin Node hosting the consolidated NMS service. Can be the primary Admin Node.

consolidated NMS

Hosted by the consolidated Admin Node. It is the equivalent of a combined reporting NMS and processing NMS service. See also NMS.

Control Node

A building block of the HP MAS system. The Control Node provides services for managing content metadata and content replication.

C-STORE

A DICOM operation to send data between devices.

CSTR

Null-terminated, variable length string.

DC

Data Center site.

deduplication

If enabled, when the grid identifies two files as being identical, it “deduplicates” them by redirecting all content handles to point to a single stored instance of the file. The end result is that only the number of copies required by the ILM policy are stored in the grid. The feature was designed for use with applications that save two identical copies of a file to the grid via different Gateway Nodes.

DICOM

Digital Imaging and COmmunications in Medicine—A standard developed by ACR-NEMA (an alliance of the American College of Radiology and the National Electrical Manufacturer’s Association) for communications between medical imaging devices.

distributed CMS

A CMS that uses metadata replication. See also metadata replication.

DR

Disaster Recovery site.

Enablement Layer

The Enablement Layer for StorageGRID Software CD is used during installation to customize the Linux operating system installed on each grid server. Only the packages needed to support the services hosted on the server are retained, which minimizes the overall footprint occupied by the operating system and maximize the security of each grid node.

EVA

Enterprise Virtual Array—an HP product that uses virtual arrays to allocate SAN or Fibre Channel storage resources to different uses.

Fibre Channel

A networking technology primarily used for storage. The standard connection type for SAN.

FCS

Fixed Content Storage—a class of stored data where the data, once captured, is rarely changed and must be retained for long periods of time in its original form. Typically this includes images, documents, and other data where alterations would reduce the value of the stored information.

Federated

A “fully distributed” grid deployment topology that completely decentralizes site deployments. There is no DC or DR site. Data sharing and disaster recovery is achieved in a peer-to-peer manner by automatically distributing data to other Federated sites.

FSG

File System Gateway—A software component of the HP MAS system. The FSG service enables standard network file systems to interface with the grid. The FSG service is found on the Gateway Node.

FSG replication group

A replication group is a group of FSGs that provide grid access to a specified set of clients. Within each replication group, there is a primary FSG (or a primary FSG cluster) and one or more secondary FSGs. The primary FSG allows clients read and write access to the grid, while storing file system information (file pointers) for all files saved to the grid. The secondary FSG “replicates” file system information, and backs up this information to the grid on a regular schedule.

Gateway Node

A building block of the HP MAS system. The Gateway Node provides connectivity services for NFS/CIFS file systems and the HTTP and DICOM protocols. A Gateway Node always hosts an FSG.

Gateway Node replication group

See FSG replication group.

GDU

Grid Deployment Utility—A HP MAS software utility used to facilitate the installation and update of software on all grid nodes. GDU is installed and available on the primary Admin Node.

Grid ID signed text block

A BASE64 encoded block of cryptographically signed data that contains the grid ID which must match the grid ID (gid) element in the grid specification file. See also provisioning.

grid node

The name of the HP MAS system building blocks, for example Admin Node or Control Node. Each type of grid node consists of a set of services running on a server.

Grid Specification File

An XML file that provides a complete technical description of a specific grid deployment. It describes the grid topology, and specifies the hardware, grid options, server names, network settings, time synchronization, and gateway clusters included in the grid deployment. The Deployment Grid Specification file is used to generate the files needed to install the grid.

Grid Task

A managed sequence of actions that are coordinated across a grid to perform a specific function (such as adding new node certificates). Grid Tasks are typically long-term operations that span many entities within the grid. See also Task Signed Text Block.

HAGC

High Availability Gateway Cluster—An HAGC is a primary gateway cluster that consists of a main FSG and a supplementary FSG. A high availability gateway replication group optionally includes one or more secondary FSGs.

HCAC

High Capacity Admin Cluster (HCAC) is the clustering of a reporting Admin Node and processing to increase a grid's grid services and thus grid node capacity. See also reporting Admin Node, processing Admin Node, and Admin Node.

HP MAS

HP Medical Archive solution — Fixed-content grid storage system from Hewlett-Packard. The solution is sold under the HP brand and is serviced and supported by the HP services/support organization worldwide. The HP MAS Solution is powered by Bycast® StorageGRID® software.

HTTP

Hyper-Text Transfer Protocol—A simple, text based client/server protocol for requesting hypertext documents from a server. This protocol has evolved into the primary protocol for delivery of information on the World Wide Web.

HTTPS

Hyper-Text Transfer Protocol, Secure—URIs that include HTTPS indicate that the transaction must use HTTP with an additional encryption/authentication layer and often, a different default port number. The encryption layer is usually provided by SSL or TLS. HTTPS is widely used on the internet for secure communications.

ILM

Information Lifecycle Management—A process of managing content storage location and duration based on content value, cost of storage, performance access, regulatory compliance and other such factors.

inode

On UNIX/Linux systems, data structure that contains information about each file, for example, permissions, owner, file size, access time, change time, and modification time. Each inode has a unique inode number.

instance

A DICOM term for an image. One or more instances for a single patient are collected in a “study”. For example, each “slice” of an MRI is an instance; together, the full set of slices is a study.

KVM

Keyboard, Video, Mouse—A hardware device consisting of a keyboard, LCD screen (video monitor), and mouse that permits a user to control all servers in a cabinet.

LAN

Local Area Network—A network of interconnected computers that is restricted to a small area, such as a building or campus. A LAN may be considered a node to the Internet or other wide area network. Contrast with WAN.

latency

Time duration for processing a transaction or transmitting a unit of data from end to end. When evaluating system performance, both throughput and latency need to be considered. See also throughput.

LDR

Local Distribution Router—A software component of the HP MAS system. The LDR service manages the storage and transfer of content within the grid. The LDR service is found on the Storage Node.

main primary FSG

In an HAGC, the FSG that is configured to be the active primary FSG by default.

metadata

Information related to or describing an object stored in the grid, for example file ingest path or ingest time.

metadata replication

In a grid that uses metadata replication, a CMS makes copies of metadata on the subset of CMSs that are in its CMS replication group, and then applies the grid’s ILM policy to content metadata. In the NMS MI, CMSs that use metadata replication display the Metadata component. Called “distributed CMS” in a previous release.

metadata synchronization

In a grid that uses metadata synchronization, a CMS synchronizes metadata with all other read-write CMSs in the grid. Called “synchronized CMS” in a previous release.

MI

Management Interface—The web-based interface for managing and monitoring the HP MAS system provided by the NMS software component. See also NMS.

namespace

A set whose elements are unique names. There is no guarantee that a name in one namespace is not repeated in a different namespace.

nearline

A term describing data storage that is neither “online” (implying that it is instantly available like spinning disk) nor “offline” (which could include offsite storage media). An example of a nearline data storage location is a tape that is loaded in a tape library, but is not necessarily mounted.

NFS

Network File System—A protocol (developed by SUN Microsystems) that enables access to network files as if they were on local disks.

NMS

Network Management System—A software component of the HP MAS system. The NMS service provides a web-based interface for managing and monitoring the HP MAS system. The NMS service is found on the Admin Node (both the reporting and processing Admin Nodes in an HCAC). There are three types of NMS service: consolidated, reporting, and processing. See also MI and Admin Node.

node ID

An identification number assigned to a grid service within the HP MAS system. Each service (such as an CMS or ADC) in a single grid must have a unique node ID. The number is set during system configuration and tied to authentication certificates.

NTP

Network Time Protocol—A protocol used to synchronize distributed clocks over a variable latency network such as the internet.

object store

A configured file system on a disk volume. The configuration includes a specific directory structure and resources initialized at system installation.

PACS

Picture Archiving and Communication System—A computerized system of patient records management responsible for short and long term (archival) storage of images.

presentation context

A combination of a DICOM SOP Class and a transfer syntax; the type and format of a DICOM transaction.

primary Admin Node

Admin Node that hosts the CMN service. There is one per grid. See also Admin Node.

primary FSG

In an FSG replication group, the FSG that provides read-write services to clients. See also FSG replication group.

processing Admin Node

Performs attribute and configuration processing that is passed on to the reporting Admin Node as part of a High Capacity Admin Cluster. See also reporting Admin Node and HCAC.

processing NMS

Hosted by the processing Admin Node. Provides attribute and data processing functionality. Only operates in conjunction with a reporting Admin Node and the reporting NMS. See also NMS.

provisioning

The process of editing the Grid Specification File (if required) and generating or updating the SAID package. This is done on the primary Admin Node using the provision command. The new or updated SAID package is saved to the Provisioning USB flash drive. See also Grid Specification File and SAID.

quorum

A simple majority: 50% + 1 of the total number in the grid. In HP MAS software, some functionality may require a quorum of the total number of some types of service to be available.

reporting Admin Node

Reports attribute and configuration information to web clients as part of a High Capacity Admin Cluster. See also processing Admin Node.

reporting NMS

Hosted by the reporting Admin Node. Reports status information about the grid and provides a browser-based interface. Only operates in conjunction with a processing Admin Node and the processing NMS. See also NMS.

SAID

Software Activation and Integration Data—Generated during provisioning, the SAID package contains site-specific files and software needed to install a grid.

Samba

A free suite of programs which implement the Server Message Block (SMB) protocol. Allows files and printers on the host operating system to be shared with other clients. For example, instead of using telnet to log into a Unix machine to edit a file there, a Windows user might connect a drive in Windows Explorer to a Samba server on the Unix machine and edit the file in a Windows editor. A Unix client called “smbclient”, built from the same source code, allows FTP-like access to SMB resources.

SAN

Storage Area Network—A high-speed network connecting heterogeneous storage devices to servers.

SATA

Serial Advanced Technology Attachment—A connection technology used to connect servers and storage devices.

SCP

Storage Class Provider—A device that provides images (a storage class) to a DICOM compliant system. Contrast with “SCU”.

SCSI

Small Computer System Interface——A connection technology used to connect servers and peripheral devices such as storage systems.

SCU

Storage Class User—A device that receives (uses) images (a storage class) from a DICOM compliant system. Contrast with “SCP”.

secondary FSG

A read-only FSG that may also perform backups of the FSG replication group. See also FSG replication group.

security partitions

If enabled, access to content ingested into the grid via a Gateway Node or configured HTTP API client is restricted to the application or client that ingested the object.

server

Used when referring specifically to hardware.

Server Manager

Application that runs on all grid servers, supervises the starting and stopping of grid services, and monitors all grid services on the server.

service

A unit of the HP MAS software such as the ADC, CMS or SSM.

SLES

SUSE Linux Enterprise Server—A commercial distribution of the SUSE Linux operating system, used with the HP MAS system.

SOP class

Service-Object Pair Class—The combination of an information object description (IOD) and the set of Services that are useful for a given purpose.

SOP instance

Service-Object Pair (SOP) Instance—A specific occurrence of an Information Object.

SQL

Structured Query Language— An industry standard interface language for managing relational databases. An SQL database is one that supports the SQL interface.

SSAM

IBM® System Storage™ Archive Manager.

ssh

Secure Shell— A Unix shell program and supporting protocols used to log in to a remote computer and execute commands over an authenticated and encrypted channel.

SSM

Server Status Monitor—A unit of the HP MAS software that monitors hardware conditions and reports to the NMS. Every server in the grid runs an instance of the SSM. The SSMS service is present on all grid nodes.

SSL

Secure Socket Layer—The original cryptographic protocol used to enable secure communications over the internet. See also TLS.

standby primary FSG

In an HAGC, the FSG that is available to take over and provide read-write services to clients in event of the failure of the active primary FSG.

Storage Node

A building block of the HP MAS system. The Storage Node provides storage capacity and services to store, move, verify, and retrieve objects stored on disks.

StorageGRID®

A registered trademark of Bycast Inc. for their fixed-content storage grid architecture and software system.

study

A DICOM term for a collection of images (instances) related to an individual patient or subject.

supplementary primary FSG

In an HAGC, the FSG that is configured to be the standby primary FSG by default.

SUSE

See SLES—SUSE Linux Enterprise Server.

synchronized CMS

See metadata synchronization.

Tape Node

A building block of the HP MAS product. The Tape Node manages storage of data to nearline data storage devices such as such as tape libraries (via IBM Tivoli® Storage Manager).

Task Signed Text Block

A BASE64 encoded block of cryptographically signed data that provides the set of instructions that define a grid task.

TCP/IP

Transmission Control Protocol / Internet Protocol—A process of encapsulating and transmitting packet data over a network. It includes positive acknowledgement of transmissions.

throughput

The amount of data that can be transmitted or the number of transactions that can be processed by a system or subsystem in a given period of time. See also latency.

TLS

Transport Layer Security—A cryptographic protocol used to enable secure communications over the internet. See RFC 2246 for more details.

transfer syntax

The parameters, such as the byte order and compression method, needed to exchange data between systems.

TSM

Tivoli® Storage Manager—IBM storage middleware product that manages storage and retrieval of data from removable storage resources.

URI

Universal Resource Identifier—A generic set of all names or addresses used to refer to resources that can be served from a computer system. These addresses are represented as short text strings.

UTC

A language-independent international abbreviation, UTC is neither English nor French. It means both “Coordinated Universal Time” and “Temps Universel Coordonné”. UTC refers to the standard time common to every place in the world.

UUID

Universally Unique Identifier—Unique identifier for each piece of content in the HP MAS. UUIDs provide client applications with a content handle that permits them to access grid content in a way that does not interfere with the grid’s management of that same content. A 128-bit number which is guaranteed to be unique. See RFC 4122 for more details.

XFS

A scalable, high performance journaled file system originally developed by Silicon Graphics.

WAN

Wide Area Network—A network of interconnected computers that covers a large geographic area such as a country. Contrast with LAN.

WAN Connectivity Kit

Required for linking the primary and DR sites of an HP MAS deployment. The WAN Connectivity Kit is installed in the Base Cabinet at both locations. See also cabinet connectivity kit.
XML
eXtensible Markup Language—A text format for the extensible representation of structured information; classified by type and managed like a database. XML has the advantages of being verifiable, human readable, and easily interchangeable between different systems.

Index