HP Medical Archive solution

Software version: 8.0

user guide

Document release date: June 2009
Software release date: June 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

Licensing

Copyright notices

Trademark notices

Contents

About this document

The User Guide provides n introduction to:

- grid customization
- grid services
- ILM polices
- grid tasks
- Server Manager

The guide also contains troubleshooting information for the NMS alarms.

For a general introduction to the HP MAS system, see the Grid Primer.

Intended audience

This guide is intended for:

- storage grid 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 grid primer*
- *HP Medical Archive audit message reference*
- *HP Medical Archive DICOM conformance statement*
- *HP Medical Archive IHE integration statement*

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 Grid Configuration

This chapter describes the following grid configuration settings:

- deduplication
- storage compression
- encryption
- security partitioning
- 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.

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 deployment. The deduplication feature was designed for a 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 not supported on grids configured to use distributed metadata replication.
- Deduplication is not supported for all Tape Node configurations.
- 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 anytime. 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 Once enabled, deduplication cannot be disabled.

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

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 and can impact ingest performance. If an application compresses objects before saving them to the grid, disable storage compression.

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

NOTE By default encryption is enabled.

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.

Security Partitions

Security partitions provide a mechanism to isolate content ingested in each Gateway Node replication group. With security partitioning enabled, data ingested via Gateway Nodes can only be retrieved via Gateway Nodes in the same replication group and not via applications using the HTTP API.

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

Security partitioning is disabled by default. When security partitioning is enabled, each Gateway Node replication group is automatically associated with a security partition. All Gateway Nodes in a replication group have ingest and retrieval access to their security partition. Any content ingested into the grid before security partitions are enabled is accessible to all grid clients.

To determine if security partitioning is enabled, log in to the NMS, go to **Grid Management > Grid Configuration > Overview**, and look up the value of Security Partitions Flag.

NOTE By default security partitioning is disabled.

Enabling security partitions is restricted to the Vendor account.

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, 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.

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 1, 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).

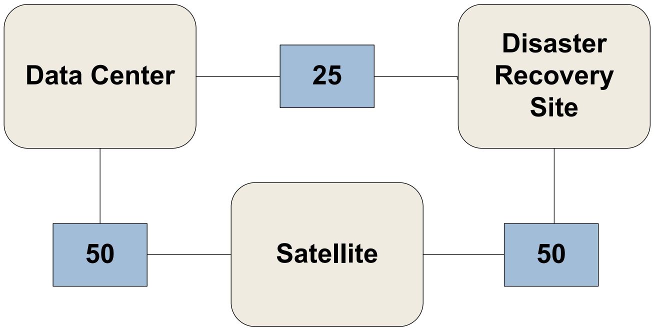

Figure 1 Link Costs Between Nodes

The default values for new installations are shown in <Hyperlink>Table 1.

Table 1 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	25	DC and DR sites usually have an ADC.
Between one group that has an ADC and one that does not	50	A satellite site does not usually have an ADC.
Between groups that do not contain an ADC	100	This is the maximum cost you can assign to a link.

To view the link cost information, log in to the NMS and go to **Grid Management > Grid Configuration > Link Costs**. Link cost 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 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 and stored in the form of text log files.

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

Table 2 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 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 NMS Grid Topology Tree, that is, the items in the navigation 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 NMS navigation tree and elsewhere, in the language selected for the user account
- the sequence that entities appear in the navigation 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 navigation 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

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

Connection Profiles

In most cases, applications external to the grid save and retrieve data via a file share configured on a Gateway Node.

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

- IP Ranges
- HTTP
- HTTP Advanced
- DICOM
- DICOM Advanced

NOTE Gateway Nodes save and retrieve data via the grid's HTTP API and are treated as client applications to the grid. Gateway Nodes are granted access to the grid via the predefined configuration profile FSG_HTTP.

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.

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

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.

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
- IP Ranges

HTTP Profile configuration is restricted to the Vendor account.

HTTP Advanced

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 3 describes each namespace and its use:

Table 3 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		
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 that can be assigned by a third-party application, as long as the standard for generating unique identifiers is followed. The Delete option can be overridden by configuring options on the **Grid Management > Grid Configuration > Configuration** tab. For more information, see HTTP No-Delete Flag (page 15).		
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 and allocate a UUID handle.
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.
HTTP /DICOM Namespace		

Table 3 Grid Management > Grid Configuration > > HTTP Advanced Settings *(continued)*

Prompt	Type	Description
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. • AE Title 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 Grid Management > Grid Configuration > HTTP Advanced > HTTP /DICOM Namespace table. This field is used to assign an AE title to all content submitted by an entity assigned this profile. • Coerce Tag Profile Name 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. This field can be left blank if no partitioning is applied to entities with the profile.		
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.
Coerce Tag Profile Name	Pull-down menu	A Coerce Tag Profile label defined in the DICOM Advanced component. The profile is used to manage logical partitioning of data.
PUT	Check box	Enables use of the HTTP PUT command by clients assigned this profile.
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. Enabling POST in the grid namespace also enables a custom application to store custom audit messages supplied by the application.		
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.
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 through the HTTP API.		

HTTP

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 4 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.

IP Ranges

The IP Ranges component grants devices with the listed IP address access to the grid for interactions that use the permitted protocols (HTTP and/or DICOM).

To open a TCP/IP connection to a grid service, the caller’s IP address must be on the list of addresses configured in **Grid Management > Grid Configuration > IP Ranges**. If the caller is not on the list, the request is ignored.

If an IP address fits more than one range, the device is matched to the first occurrence from the top of the table downward.

Table 5 Grid Management > Grid Configuration > IP Range 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 an IP Range that can access the grid. IP Ranges are specified using one of the following: • a single IP address • a hyphenated list of IP addresses inclusive of the listed values (for example, 174.182.91.00-174.182.91.64) • a range of IP addresses specified using CIDR notation (for example, 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	Provided a connection is permitted, 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. At the time the grid deployment is configured, the services are allocated to logical groups; typically these are numbered: 1, 2, 3, and so on. The main Overview page for each service includes the Group ID attribute, the group to which the service belongs. Communication between groups has a cost assigned when the configuration plan is created for your enterprise. Entities (IP addresses) within the same group are assumed to have a zero cost. You cannot create new groups using this configuration table. 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.
DICOM	Check box	Enable the use of DICOM protocol.
HTTP	Check box	Enable the use of HTTP. If neither HTTP nor DICOM is checked, the caller can make a TCP/IP connection and get a response but cannot access grid content. You could deselect both the HTTP and DICOM protocols for an IP range to temporarily disconnect remote entities because of security concerns, for example, then easily restore connectivity later when the issues are resolved.

Multiple IP Range Matches

It is possible that a caller's IP address may match more than one of the ranges specified in the 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.

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
- IP Ranges

DICOM Advanced

NOTE DICOM is an option, and may not appear in your grid.

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.

Table 6 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).

Table 6 Grid Management > Grid Configuration > DICOM Advanced Component Configuration Settings *(continued)*

Prompt	Type	Description
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.
Behavioral Profile Name	Pull-down menu	Preconfigured profile that directs HP MAS 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 6 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**".

Table 6 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.)
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.)
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.)

DICOM

NOTE DICOM is an option.

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 7 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.
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.

IP Ranges

The IP Ranges component grants devices with the listed IP address access to the grid for interactions that use the permitted protocols (HTTP and/or DICOM). The process is the same for HTTP and DICOM devices. For more information, see IP Ranges (page 22).

2 Information Lifecycle Management

This chapter contains background information on Information Lifecycle Management (ILM).

NOTE 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.

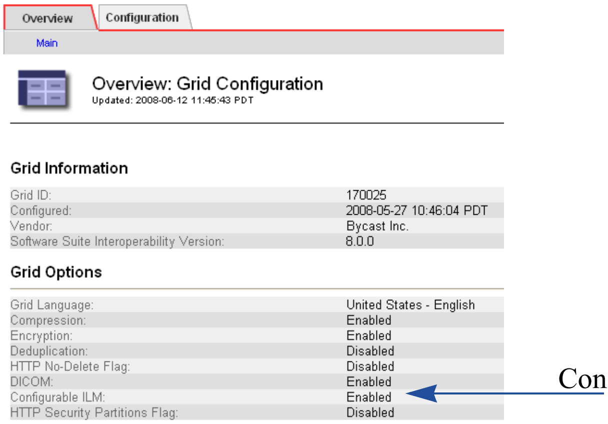

Overview: Grid Configuration
Updated: 2008-06-12 11:45:43 PDT

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 2 Configurable ILM Enabled

The NMS ILM Management menu is used to view and configure ILM policies. This menu is only used if the system has a configurable ILM.

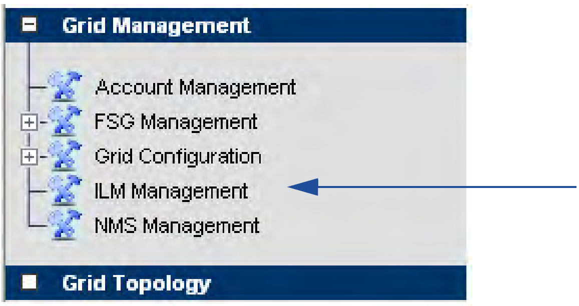

ILM Management

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 3, an object is ingested in the grid via an FSG. At ingest, two copies of the object are stored in the Data Center (DC) on SATA disks and one copy is stored in the Disaster Recovery (DR) site on SATA disks. One year after ingest, one copy at the DC is deleted and one copy is created on archival media at the DR site.

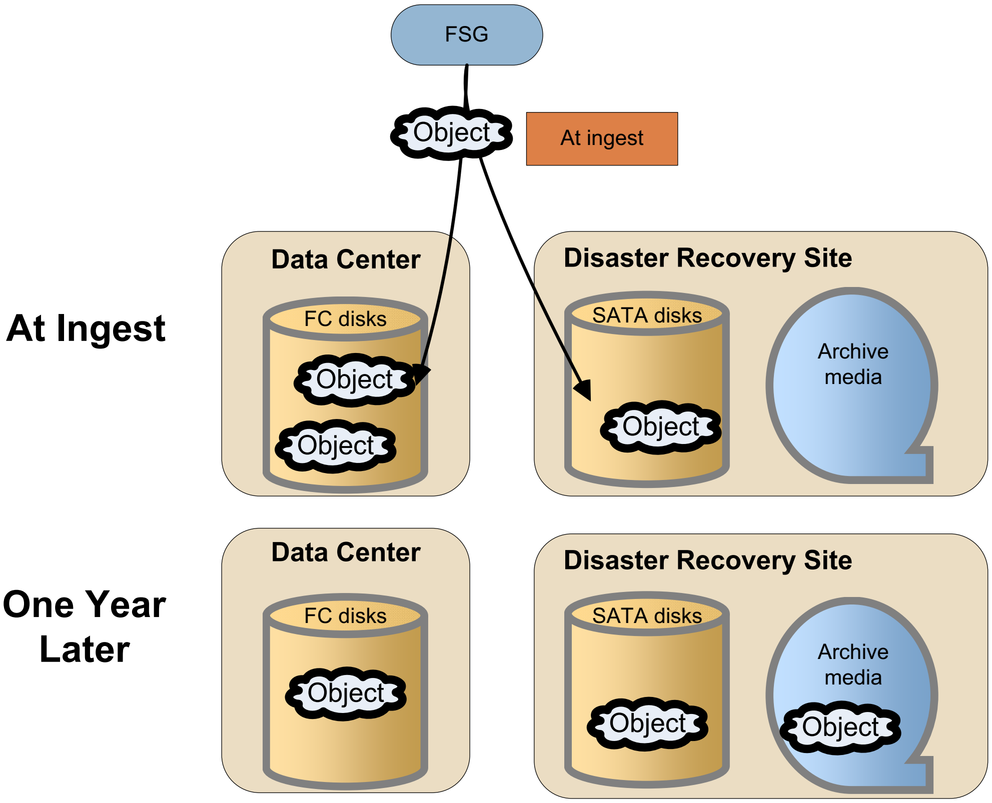

Figure 3 Information Lifecycle Management

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 4) 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.

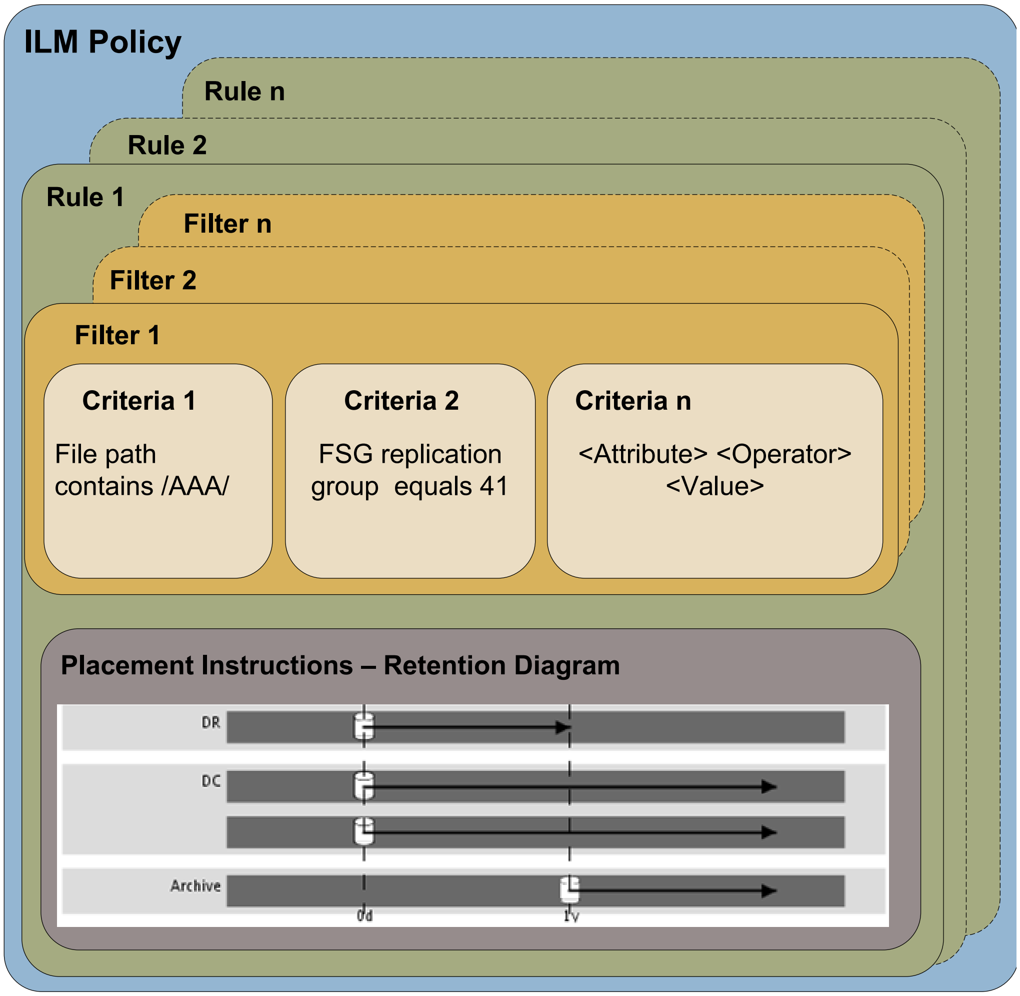

Figure 4 ILM Policy Elements

Retention Diagrams

Retention diagrams represent what happens to an object over time. Figure 5 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, two copies are stored at the DC and one copy at the DR site. After one year, the DR copy is deleted and one copy is made on Archive.

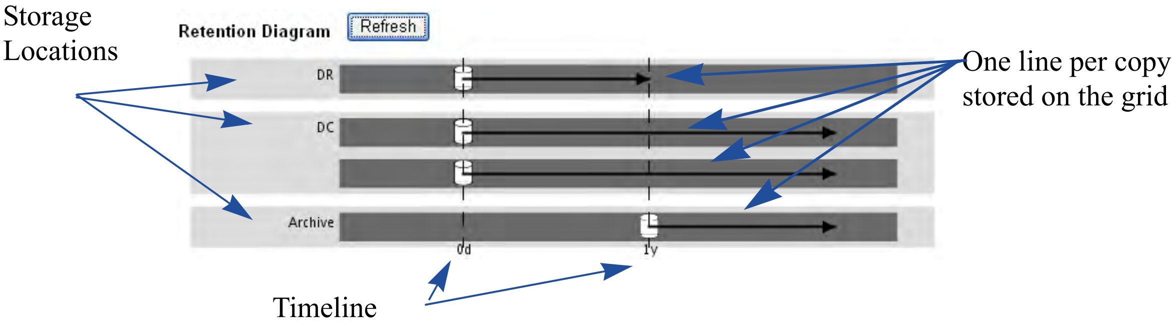

Figure 5 Sample Retention Diagram

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 6, a HP MAS system is deployed across two sites or groups: DC and DR site. Files can be ingested via FSGs or the HTTP API. The DC 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

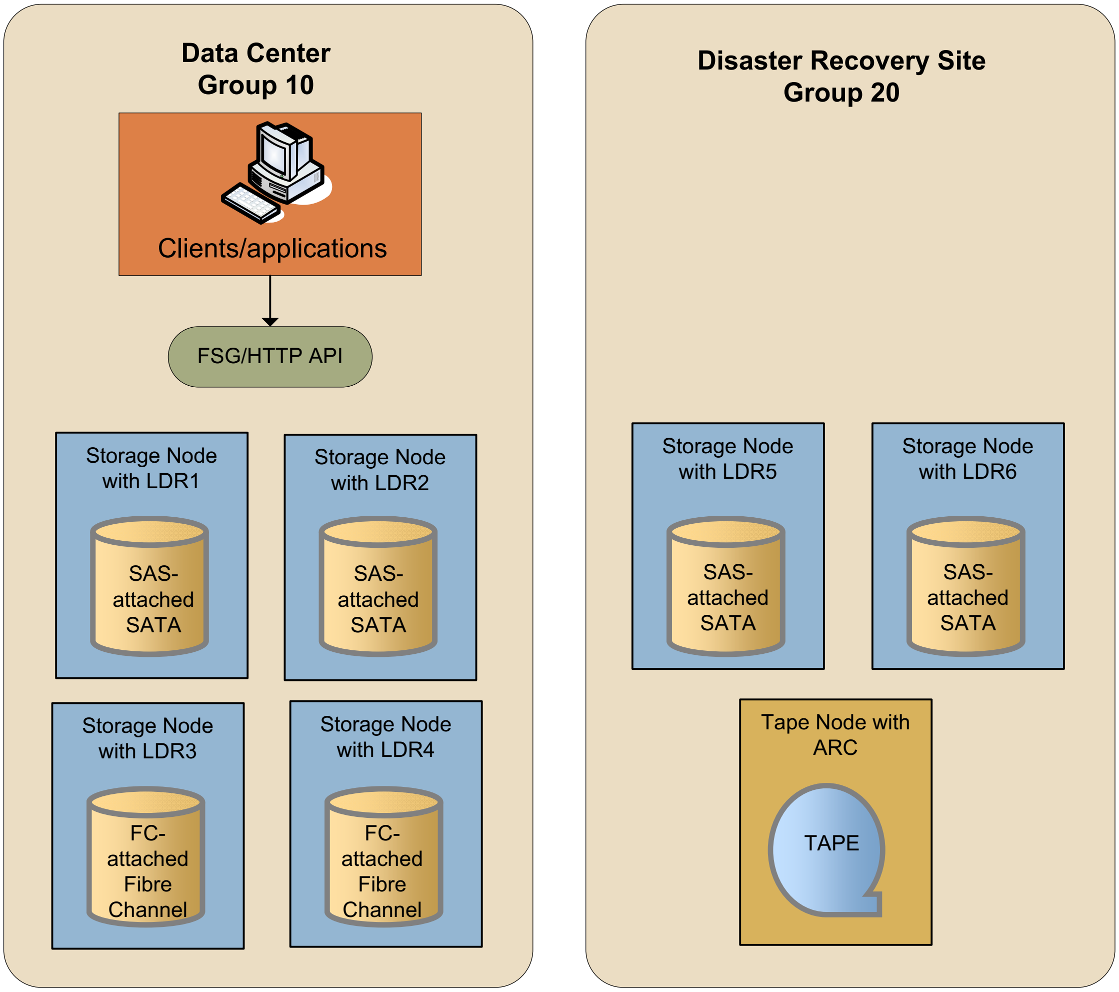

Figure 6 . DC + DR Site with Archive Configuration

The DC has four Storage Nodes. A Storage Node is a server that includes an LDR (Local Distribution Router) 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 ARC (Archive) 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 7 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

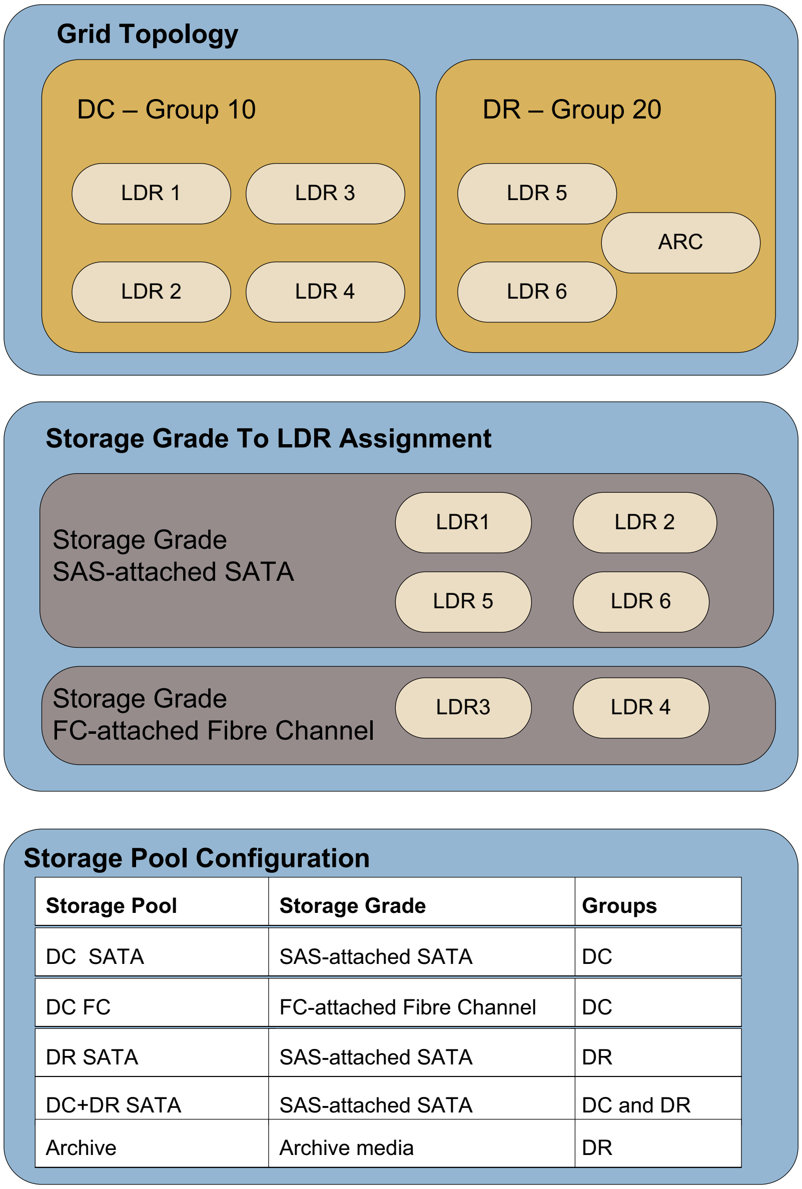

Storage Pool	Storage Grade	Groups
DC SATA	SAS-attached SATA	DC
DC FC	FC-attached Fibre Channel	DC
DR SATA	SAS-attached SATA	DR
DC+DR SATA	SAS-attached SATA	DC and DR
Archive	Archive media	DR

Figure 7 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 8 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 8 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.

Table 8 Application Metadata *(continued)*

Metadata	Code	Definition	Notes on ILM Use
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**).
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

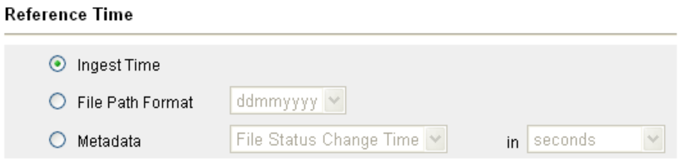

Reference Time

● Ingest Time

○ File Path Format ddmmyyyy

○ Metadata File Status Change Time in seconds

Figure 8 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
- yyyyymmdd

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.

Two 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.

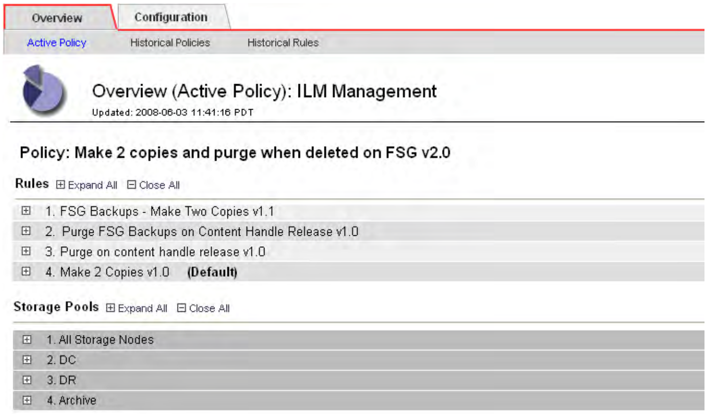

Overview (Active Policy): ILM Management

Updated: 2008-06-03 11:41:16 PDT

Policy: Make 2 copies and purge when deleted on FSG v2.0

Rules ⊞ Expand All ⊟ Close All

⊞ 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 ⊞ Expand All ⊟ Close All

⊞ 1. All Storage Nodes
⊞ 2. DC
⊞ 3. DR
⊞ 4. Archive

Figure 9 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.

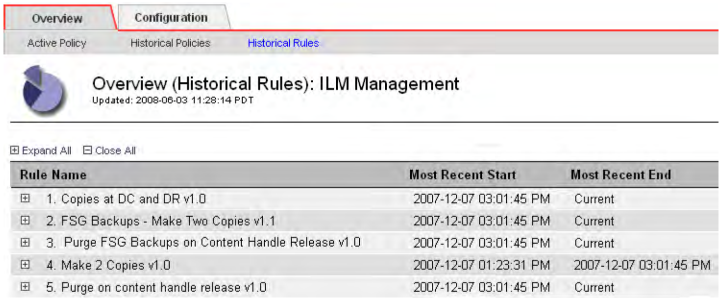

Overview (Historical Rules): ILM Management

Updated: 2008-06-03 11:28:14 PDT

⊞ Expand All ⊟ Close All

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 10 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.

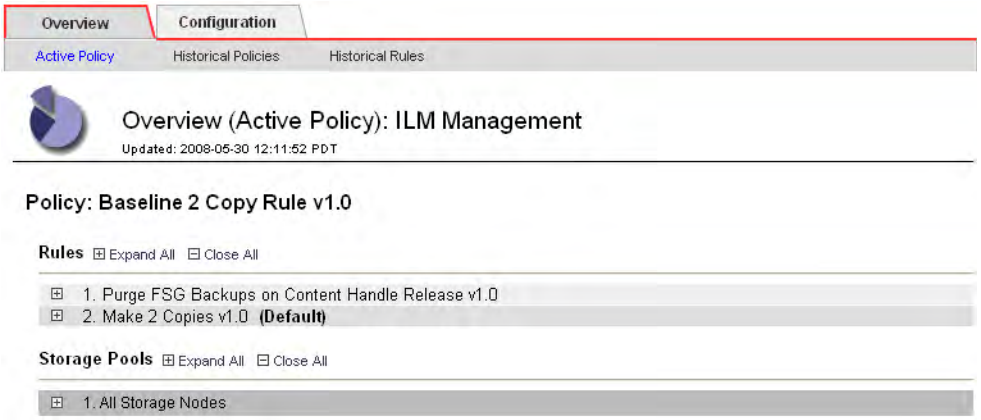

Overview (Active Policy): ILM Management

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 11 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 Versioning

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).

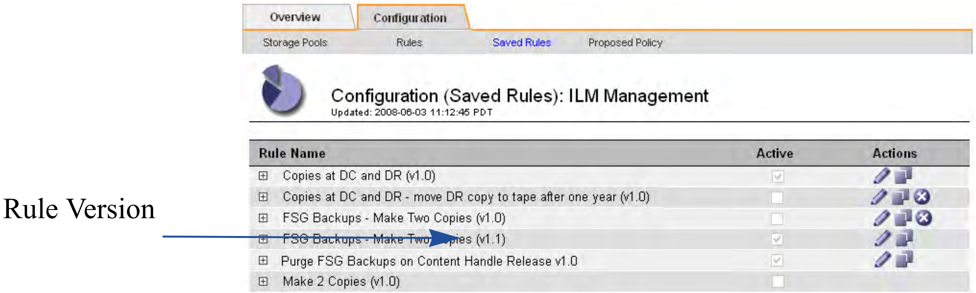

Configuration (Saved Rules): ILM Management

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 12 Rule Versions

Unavailable Resources

In addition to the preferred storage location, a rule should specify a temporary location for the object in case the preferred location is unavailable.

△ CAUTION Use of temporary copies is strongly recommended. Failure to specify a 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 the temporary location if specified.
2. When the preferred location becomes available, the object is placed in the preferred location.
3. After the object has been stored in the preferred location, the copy in the temporary location is purged.

For example, if a rule specifies the following placement for all content at ingest:

- Store two copies in DC (temporary location is DR)
- 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. Once the DC pool becomes available, two copies are made in the DC storage pool and the DR copies are deleted.

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.

Versioning

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).

3 Services and Components

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

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

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

Services

Grid services are programs that run on the physical servers. Each delivers particular services to the grid as shown in Table 9. The NMS interface is used to access the critical information gathered by these services so that you can monitor the condition and performance of your grid.

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

Table 9 HP MAS Services

	Icon	Tag	Service Name	Function	See
Data Plane	▼	LDR	Local Distribution Router	Storing and routing data through the grid.	page 73
		CLB	Connection Load Balancer	Enabling storage and retrieval through the DICOM protocol option.	page 55
	📁	FSG	File System Gateway	Enabling storage and retrieval through standard network file systems.	page 63
		ARC	Archive	Managing long-term nearline storage of data	page 53

Table 9 HP MAS Services *(continued)*

	Icon	Tag	Service Name	Function	See
Control Plane	●	CMS	Content Management System	Storing and managing metadata for the content.	page 57
	▭	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 50
Management Plane	▯	SSM	Server Status Monitor	Monitoring server activity and hardware.	page 77
	▲	NMS	Network Management System	Notifying administrators of alarm conditions and providing the user interface for grid management.	page 76
	✎	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 56
	🗎	AMS	Audit Management System	Logging grid transactions for audit and reporting.	page 52

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

Table 10 lists all of the available service components, in alphabetic order by name. The same symbol may be used by more than one component.

Table 10 Service Components

Icon	Name	Parent Service	Data Provided
	Audit	CMN	Configuration of the audit messages to be logged.
	Backup	FSG	FSG backup settings and status.
	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, as well as attribute transport and the attribute repository.
		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.
	Metadata	CMS	Information on metadata replications. Only used on grids where distributed CMS is enabled.
	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.

Table 10 Service Components *(continued)*

Icon	Name	Parent Service	Data Provided
	Identifiers	CMN	Total number of unique object identifiers installed and available.
	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.
	Retrieve	ARC	Information about object requests and cached objects.
	Services	SSM	Information on the state and resource usage of services being monitored.
	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.
	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.

About NMS Attributes

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

Events Component Attributes

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

Resources Component Attributes

All services have a Resources component which use a standard set of attributes that 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 services have a Timing component which use a standard set of attributes that 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 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 Nodes activity for that part of the grid.

The values of summary attributes are based on estimates.

Top Attributes

The NMS interface displays hundreds of attributes. However, most of these attributes are required only for troubleshooting. The list of attributes to monitor routinely, shown in Figure 11, is much shorter. For tips on how to analyze these attributes, see the *Grid Primer*.

Table 11 Key Attributes to Monitor

Category	Component	Code	Description
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.
	LDR Storage	SAVP	Total Space Available (Percent): Object storage capacity that is still available for use on the LDR.
CMS metadata storage capacity	CMS Database	CORS	Estimated Remaining Object Capacity: The number of additional objects that can be managed by this CMS.
		DBSP	Free Tablespace (Percent): The metadata storage capacity that is still available for use on the CMSs.
FSG capacity	FSG Backup	PBNF	Number of Files: The number of files included in the last FSG backup.
		PBDS	Backup Data Size: The total size of all files referenced by the last FSG backup.
		PBSF	Number of Files: The number of files included in the last backup for each share.
		PBSB	Total File Size: The total size of all files referenced by the last backup for each share.
	FSG Storage	FSIU	Inodes Used: Number of inodes (files and directories) used on the FSG filesystem.
Ingest load	FSG Storage	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.

Table 11 Key Attributes to Monitor ***(continued)***

Category	Component	Code	Description
Retrieve load	FSG Storage	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.
Object replication	CMS Content	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.
		ORpe	Objects with ILM Evaluation Pending: The number of objects waiting to be processed through the business rules for replication.
Metadata synchronization **(for synchronized CMSs only)**	CMS Synchronization	CsQT	Queue Size: The number of outgoing metadata synchronization messages queued to be sent to other CMSs.
FSG replication	Secondary FSG Replication	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 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	ARBA	Archived Bytes: The total amount of content written to archive media by this ARC.

ADC—Administrative Domain Controller

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

Each ADC synchronizes certificates, configuration bundles, and information about services and topology with the other ADCs in the grid to facilitate distributed and islanded operation.

In general, all system nodes maintain a connection to at least one ADC. 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 is unavailable. New nodes can only establish connections via an ADC.

The connection of each node lets the ADC 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 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 supports the following components:

- Synchronization
- Events
- Resources
- Timing

Synchronization

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

Events

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

The ADC acts as both an attribute relay and an audit relay. This means that the ADC 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 ADC Resources component uses the standard set of resources attributes.

Timing

The ADC Timing component uses the standard set of timing attributes.

AMS — Audit Management System

The Audit Management System (AMS) 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, 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 supports the following components:

- Events
- Resources
- Timing

Events

The AMS Events component uses the standard set of events attributes.

Resources

The AMS Resources component uses the standard set of resources attributes.

Timing

The AMS 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 pushes a copy from a Storage Node to the Tape Node. The ARC 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. The ARC requests the object from the middleware, which retrieves the object from the archival media device and sends it to the ARC. 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

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 that uses a Tivoli Storage Manager (TSM) to manage archival media displays Tivoli Storage Manager as its Middleware Type on the **ARC > Middleware** page. The ARC 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.

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 supports the following components:

- Replication
- Store
- Retrieve
- Middleware
- Events

- Resources
- Timing

Replication

The business rules for the grid dictate how many copies of each piece of data are kept, and where these copies are made. This component provides information about the node to node content replication performed to satisfy these business rules.

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 either as a result of a request from a client for an object, or as a result of a request from a CMS 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 ARC Events component uses the standard set of events attributes.

Resources

The ARC Resources component uses the standard set of resources attributes.

Timing

The ARC Timing component uses the standard set of timing attributes.

CLB — Connection Load Balancer

The Connection Load Balancer (CLB) 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 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 supports the following components:

- DICOM
- HTTP
- Events
- Resources
- Timing

DICOM

NOTE 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 CLB Events component uses the standard set of events attributes.

Resources

The CLB Resources component uses the standard set of resources attributes.

Timing

The CLB Timing component uses the standard set of timing attributes.

CMN — Configuration Management Node

The CMN service manages grid-wide configurations of connectivity and protocol features needed by all services in the grid. As well, the CMN 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 supports the following components:
- Grid Tasks
- Object Lookup
- Identifiers
- Events
- Resources
- Timing

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. 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 Grid Tasks (page 117).

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 CMN Events component uses the standard set of events attributes.

Resources

The CMN Resources component uses the standard set of resources attributes.

Timing

The CMN Timing component uses the standard set of timing attributes.

CMS — Content Management System

The Content Management System (CMS) manages content metadata and Information Lifecycle Management (ILM). While the LDRs manage the content, the CMS services provide the logic.

Metadata Management

To ensure redundancy, the grid stores multiple copies of the content metadata in different databases. 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.

The first CMS 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. There are two possible approaches for metadata replication:

- synchronized
- distributed

To determine whether the grid has synchronized CMSs or distributed CMSs, look at the components under CMS. A grid with distributed CMSs has the component Metadata. This component does not appear on grids using synchronized CMSs.

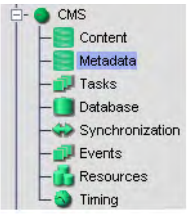

Figure 13 CMS Metadata Component For Distributed CMSs

All grids installed prior to Release 8.0 have synchronized databases even if the software has been updated to 8.0.

Synchronized CMSs

In a synchronized environment, the owner CMS replicates its content to all CMSs. The CMS databases are fully synchronized: they all store all content metadata.

Grids that have synchronized CMSs 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 synchronized with the existing CMSs at the time the new servers were added to the grid. In this case, all Control Nodes remain equivalent.
- If the Control Nodes were added to add CMS capacity, the Control Nodes are no longer all equivalent. Control Nodes are typically added in groups. The CMS added at the same time are equivalent to each other but they are not equivalent to the other CMSs in the grid.

Synchronized CMSs must be used in the following situations:

- Grids using non-configurable ILM
- Grid where DICOM is enabled
- Grids where de-duplication is enabled

Distributed CMSs

In a distributed environment, the owner CMS replicates its content metadata to the CMSs that are in the same CMS replication group as the owner CMS and to the CMS that are in the same groups as the LDR storing the content.

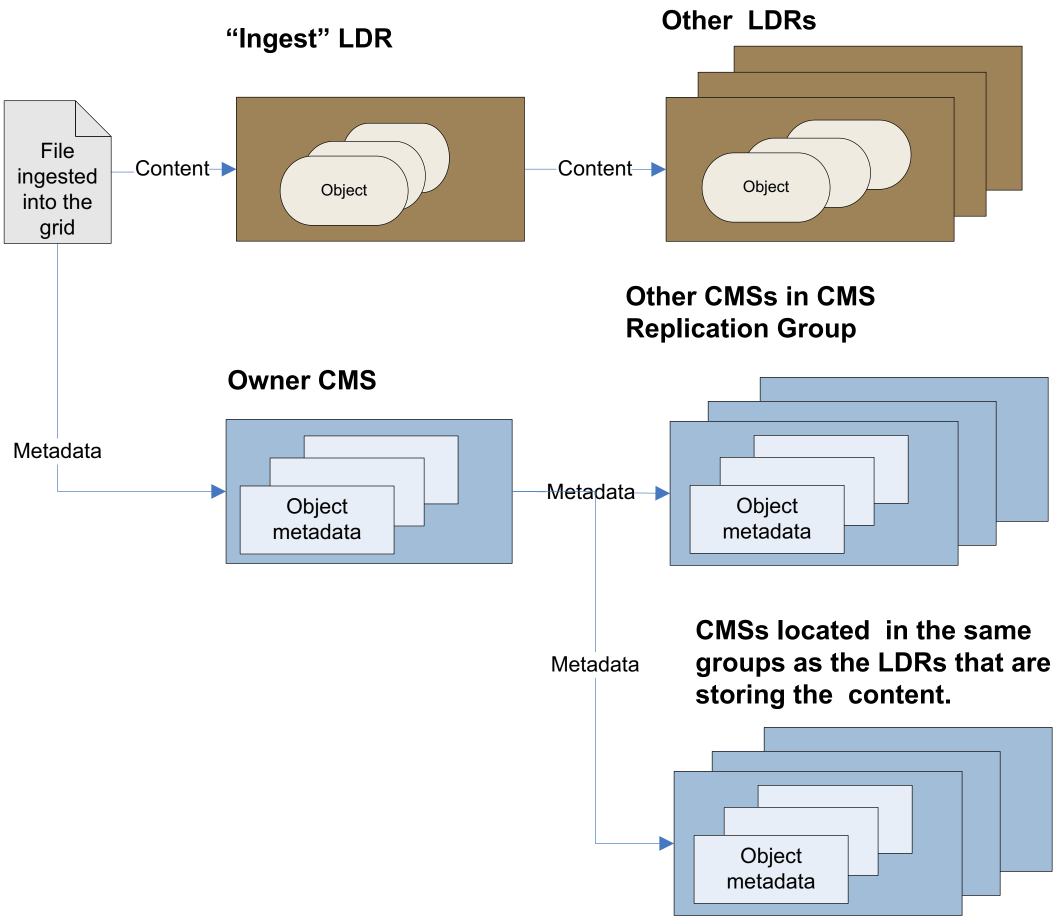
"Ingest" LDR
Other LDRs
File ingested into the grid
Content
Object
Content
Object
Other CMSs in CMS Replication Group
Owner CMS
Metadata
Object metadata
Metadata
Object metadata
CMSs located in the same groups as the LDRs that are storing the content.
Metadata
Object metadata

Figure 14 Distributed CMS Operation

If synchronous metadata commit is enabled, the metadata is sent synchronously at ingest to one other member of the CMS replication group as a safety precaution before the ingest is acknowledged to the LDR.

Information Lifecycle Management

In addition to managing metadata, the CMSs 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). In a grid that uses distributed CMSs the same content replication ILM policy is applied to the content metadata.

Metadata Storage Capacity

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 the "Operations" chapter in the *Grid Primer*.

CMSs become read-only once their 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.

In a distributed CMS environment, if one CMS in a CMS replication group becomes read-only, all of the CMSs in the replication group turn off their "content ingest service". The read-write CMSs are still available for metadata replication, but they do not accept new content.

If all CMS databases in the grid are full, the grid cannot ingest files unless the grid's CMS capacity is increased. For detailed information on Control Node expansion procedures, see the *Expansion Guide*.

CMS Components

The CMS has the following components in the NMS:

- Content
- Metadata
- Tasks
- Database
- DICOM
- Synchronization
- Events
- Resources
- Timing

Content

The Content component reports statistics on the metadata storage and replication activity.

If this is a grid that uses synchronized CMSs, 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. 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. To get the total count of the objects managed by the grid, use the summary attributes at the grid level.

Metadata

NOTE The Metadata component is displayed only if the grid uses distributed CMSs.

The Metadata component reports statistics on the metadata storage and replication activity for distributed CMSs.

If this is a grid that uses distributed CMSs, 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. In a grid that uses distributed CMSs, 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 is processing against its database, and their current status.

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 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 73).)

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 tasks, see Monitoring Grid Tasks (page 118).

Background Tasks

Displays special tasks such as ILM verification tasks and migration tasks.

Recovery Sources

The Recovery Sources table lists what CMS can be used to restore the CMS database in the event of a failure in a grid that uses distributed CMSs.

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

NOTE 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 in a fully synchronized CMS environment. Distributed CMS operation also uses synchronization messages but not to the same extent.

Events

The CMS Events component uses the standard set of events attributes

Resources

The CMS Resources component uses the standard set of resources attributes.

Timing

The CMS Timing component uses the standard set of timing attributes.

FSG — File System Gateway

The File System Gateway (FSG) 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 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 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.

FSG services are most commonly hosted on Gateway Nodes or combined Admin/Gateway Nodes.

FSG Replication Groups

A HP MAS grid 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.

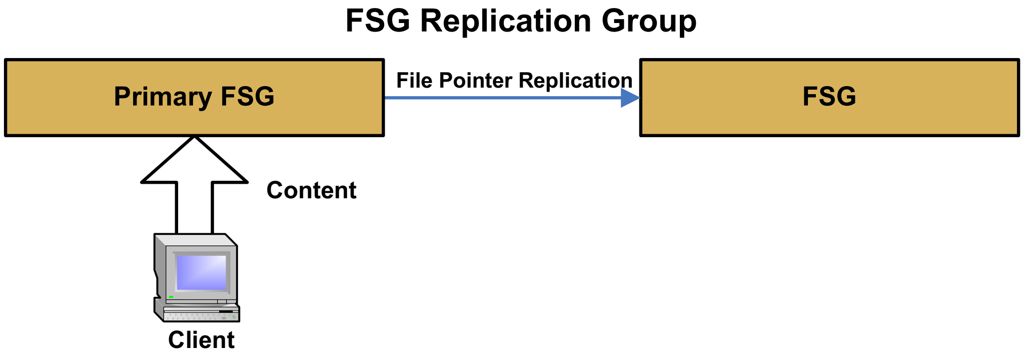

Figure 15 FSG Replication Group—General Case

There are several possible types of replication group:

- unclustered replication groups
- high availability gateway cluster replication groups

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 HAGC Replication Groups (page 66).)

Primary FSGs

Every replication group includes at least one read-write FSG, known as the Primary FSG. A replication group may contain either one stand-alone Primary FSG (in an unclustered replication group) or a set of servers that form a clustered Primary FSG (in a high availability gateway cluster 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 cluster, a stand-alone Secondary is optional.)

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.

All newly installed replication groups now perform online backups. During an online backup, an FSG continues to process replication messages. Therefore, you can retrieve content from the Gateway Node server that hosts the FSG while a backup is in progress, even if the content is newly ingested.

Some older FSGs that were originally installed prior to this release may perform offline backups. While an FSG is processing an offline backup, replication session messages are not processed which means that recently ingested files cannot be retrieved from the backup FSG.

RECOMMENDATION An “offline” backup FSG should not be used for data retrieval.

A Secondary FSG may be designated as the backup FSG in a replication group.

Unclustered Replication Groups

An unclustered 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.

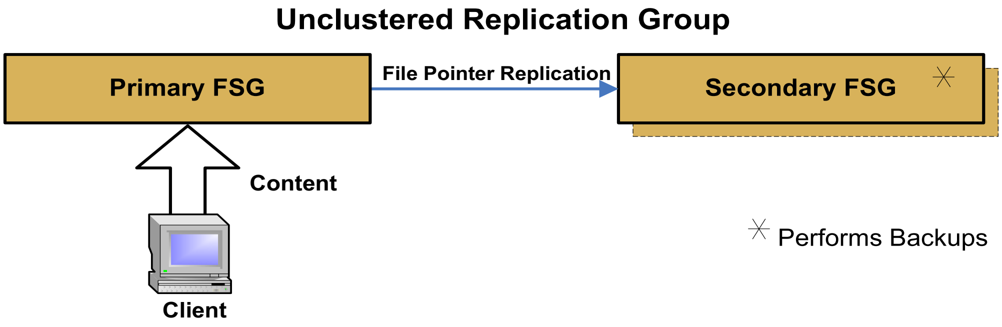

Figure 16 Unclustered Replication Group

Primary FSG

In a unclustered 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 must also include one or more Secondary FSGs. One Secondary FSG acts as the Backup FSG; if there is an additional Secondary in the replication group, its mirrored file system can be used to provide clients with read-only access to the grid. A replication group could include more than two Secondary FSGs if read-only access to the file system 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 an Unclustered 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 reduced while the Secondary FSG is acting as a temporary Primary FSG. If the Secondary is also the backup FSG for the replication group, backups are not performed while it acts as the Primary.

In grids that do not support business continuity failover (many grids do *not*), clients can continue to access files via the read-only file system on the Secondary FSG while the Primary is repaired, but cannot write data to the grid.

The backup file from the grid and the replication session messages on the Secondary FSG are used to restore the file system on the Primary FSG. After the Primary is restored to service, the roles of the FSGs are reversed, clients must be manually redirected back to the Primary, and the Secondary resumes its original role. Full redundancy and functionality are returned to the replication group.

HAGC 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 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.

An HAGC replication group can optionally include one or more stand-alone Secondary FSGs.

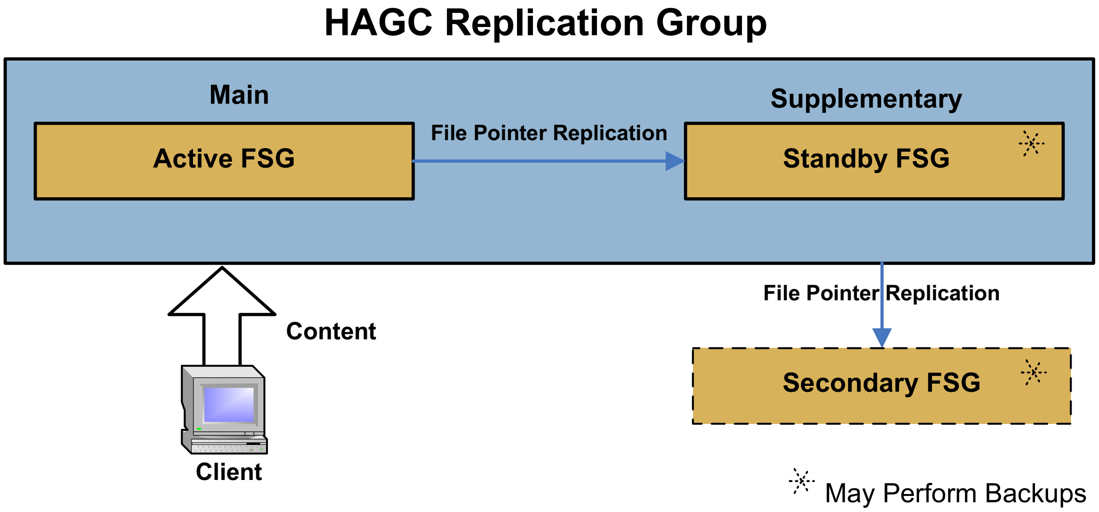

Figure 17 High Availability Gateway Replication Group with an Optional Secondary FSG

Main FSG

Although the two FSGs in the HAGC are functionally equivalent, one FSG is designated as the Main FSG. The Main FSG is the FSG that is Active by default (on system startup, and whenever the cluster is operating normally). The Main FSG is therefore the one that usually provides NFS and CIFS services to clients.

Supplementary FSG

The second FSG in the cluster is the Supplementary. By default, the Supplementary FSG is in Standby (on system startup, and whenever the cluster is operating normally). When it is in the Standby state, the Supplementary FSG receives replication session messages and provides a “mirror” of the file system on the Main FSG. This provides it with the ability to take over from the Main FSG in the event of a failure and permits it to be used to restore the Main FSG to full functionality, if needed. The Supplementary 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 and the Supplementary FSGs) can stand alone as an FSG replication group.

Other configurations are possible: you may also include a Secondary FSG as the third member of an HAGC 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 FSG cannot perform backups while it is acting as the Active Primary. Adding a stand-alone Secondary 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.)

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 Replication Group

A High Availability Gateway Cluster (HAGC) includes two FSGs. When it is operating normally, one of the FSGs in the cluster is Active, and provides the file system interface to clients. The second FSG is on Standby, ready to provide service in case the Active FSG fails. The status of both FSG services in the cluster is monitored by a heartbeat service. If any of the following conditions occur, the Standby FSG becomes Active:

- A hardware failure (such as CPU, memory, etc.) that causes a HP MAS software service or the heartbeat service to stop responding.
- A disk failure (such as disk I/O errors).
- The failure of the heartbeat service, or any of the file sharing services (CIFS, NFS) on the Active FSG.
- The Active FSG becomes disconnected from the customer network.
- A HP MAS software failure.

While the second FSG is Active, it provides full read-write functionality to clients, minimizing disruption to their operations.

When a failover occurs, any client operations that are in progress fail, as do client operations initiated while the Standby FSG makes the transition to Active. Once the CIFS service starts on the newly Active node, Windows CIFS clients should be able to process new operations without remapping their connections to the Primary FSG cluster. NFS clients must remount shares before they can continue to store and retrieve data to the Gateway cluster. Full grid functionality and full grid access is maintained while the second FSG is Active.

To restore the grid to full redundancy, the failed FSG in the cluster must be restored. Once it is fully functional, manually failing back to the restored FSG makes it the Active FSG and also restores the second FSG to Standby status.

The Supplementary FSG in an HAGC can also act as the backup FSG for a replication group, and can be used as a source of the replication messages needed to restore the Main FSG to full functionality after a failure. A stand-alone Secondary FSG is not required as a member of the HAGC replication group.

However, if the HAGC replication group also includes a Secondary FSG, the Secondary FSG can be manually configured to take over from the HAGC and act as the Primary as it would in a unclustered replication group. See Failover in an Unclustered Replication Group (page 66) for details.

Details of HAGC Failover Behavior

In general, an automatic failover within a High Availability Gateway cluster should be seamless for clients of the cluster. However, you should be aware of the following detailed behavior that can affect data access in some cases:

- If the Active FSG becomes disconnected from the rest of the grid by network problems but remains connected to the customer network, the Standby FSG does not take over. Clients can continue to write to the Active FSG as long as space remains in its cache, and they can continue to retrieve files that are cached. Requests to retrieve files that are not cached block until either connectivity is restored, or until the NFS or CIFS client times out. When network connectivity is restored, any files saved to the Active FSG but not yet ingested to the grid are ingested.
- Data that has been saved to an Active FSG, but has not yet been ingested into the grid at the time of a failover, is ingested after that FSG is restored to service.
    - If the file pointers for the new data had been replicated to at least one other FSG in the replication group at the time of the failure, they will be replicated to all FSGs after the Active FSG is restored to service.
    - If the pointers had *not* been replicated, they must be manually replicated to the other FSGs to enable retrieval from these FSGs, and to ensure correct behavior in case of a later failure. Contact Support for assistance.
- If the Active Main FSG in an HAGC fails, the Standby Supplementary FSG automatically takes over and becomes Active. After the Main FSG is restored to service, it moves into a Standby state while it finishes processing replication messages. This means that the Main FSG should be capable of taking over if the Active FSG in the cluster fails. But if the Active FSG fails before the Main FSG finishes processing replication messages, the Main FSG does not automatically become Active. It remains in the Standby state. To restore service to the cluster you must wait until the Main FSG finishes processing replication sessions, and then manually restart the FSG. Contact Support for assistance.
- If the HAGC does not fail over to the Standby FSG when the Active FSG fails, restart the heartbeat service on the Standby FSG by restarting client services. (In the NMS, go to the **Standby FSG > Client Services > Configuration > Main** and select **Stopped**. Then click **Apply Changes**. Return to the **Client Services > Overview** page to confirm that heartbeat is stopped, and then return to the Configuration page to restart Client Services.)

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 12 Summary of Types of FSG Replication Group

Replication Group Type	Primary FSG(s)	Other Members	Backup FSG	Secondary Used for Retrieval?
Online Backups				
High Availability Gateway Cluster	Main and Supplementary	Secondary optional	Supplementary[a] or Secondary	Yes
Unclustered Primary	Single Primary	One or more Secondaries	Secondary	Yes
Offline Backups[b]				
High Availability Gateway Cluster	Main and Supplementary	One or more Secondaries	Secondary	Yes—if grid contains two Secondaries
Unclustered Primary	Primary	One or more Secondaries	Secondary	Yes—if grid contains two Secondaries

[a] If the HAGC replication group does not include a Secondary, 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 the setting **FSG Management > Group X > Backup** 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.

[b] Offline backups were standard in releases prior to 7.5. Some older replication groups may continue to use offline backups after upgrading to 8.0.

Failover

The following table summarizes the failover behavior of each type of Gateway Node replication group.

Table 13 Summary of Failover Behavior

Replication Group Type	Automatic Failover of Primary?	Write Access Maintained After Failover Complete?	Resilient Against More than One Node Failure?
High Availability Gateway Cluster	Yes	Yes—NFS shares must be remounted.	No
Unclustered Primary	No - manual failover required	Yes— in grids that support business continuity failover	No

FSG Components

The FSG supports the following components:

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

Storage

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.

The storage component provides data on the service's storage space, and the objects stored and retrieved.

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.

Each FSG service is configured to assume a certain role by default. For example, all grids support mirrored FSG services to provide failover support should a primary FSG service (or primary FSG cluster) become unavailable. Services that are not clustered are configured simply 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.

This component also reports the status of the high availability 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:

- "Normal" if both FSGs in the high availability cluster are healthy: the Main Primary is active, and the Supplementary Primary is in standby.

- “Failover” if a failover has occurred and the Supplementary Primary is Active and the Main Primary is in Standby (but is available and on the grid).
- “Vulnerable” if only one server in the cluster is available and active.
- “Transitional” if the cluster is not providing FSG service but the clustering mechanism is expected to automatically restore service.
- “Failed” if no FSG in a cluster can provide an active FSG service and the failure cannot be recovered from automatically.

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.

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 FSG Events component uses the standard set of events attributes.

Resources

The FSG Resources component uses the standard set of resources attributes.

Timing

The FSG Timing component uses the standard set of timing attributes.

LDR—Local Distribution Router

The Local Distribution Router (LDR) handles content transport on the grid. Content transport encompasses many tasks including data storage, routing, and request handling. The LDR does the majority of the grid's hard work by handling heavy data transfer loads and extensive data traffic functions.

The LDR handles the following tasks:

- Content storage

The LDR uses the standard overview service attributes with some additions for the protocol interfaces and object verification and replication.

LDR Components

The LDR supports the following components:

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

Storage

This 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 LDR verification component gives you information about 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 Monitoring Grid Tasks (page 118)).

  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.

- Behind the scenes, foreground verification is performed 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 Admin Node
- 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

This component provides additional information about node to node content replication.

DICOM

NOTE 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 LDR Events component uses the standard set of events attributes.

Resources

The LDR Resources component uses the standard set of resources attributes.

Timing

The LDR Timing component uses the standard set of timing attributes.

NMS — Network Management System

The NMS performs two primary 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

The NMS service powers the monitoring, reporting, and configuration options provided by the NMS management interface. The NMS service can itself be monitored via the NMS interface.

In a grid deployment, a single NMS service collects and stores attributes from all services and components in the grid, and provides a management interface to display grid information. A grid deployment that is geographically distributed (for instance, a grid that includes a data center and a disaster recovery site), may contain two or more equivalent NMS services. Each NMS collects and stores information from all grid services, and provides a separate but equivalent management interface to display information about the grid.

NMS Components

The NMS supports the following components:

- Database
- Events
- Resources
- Timing

Database

The Database component provides information about the type of database used by the NMS for tracking attributes.

Events

The NMS Events component uses the standard set of events attributes. As the NMS also acts as an attribute repository, for the NMS this component also reports on the status of the repository.

Resources

The NMS Resources component uses the standard set of resources attributes.

Timing

The NMS Timing component uses the standard set of timing attributes.

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 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 interface.

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 supports the following components:

- Services
- RAID (when a RAID that supports this monitoring is installed)
- Events
- Resources
- Timing

Services

The Services component tracks the services and support modules running on the node. It reports their status, the number of threads (CPU tasks) running and the amount of RAM being used. The grid services are listed, as are support modules (such as time synchronization).

The section on Packages reports overall service suite installations and version numbers to facilitate software update procedures.

RAID

The RAID component only exists on SSMs that are directly monitoring at least one RAID unit. The RAID component displays extensive data on each attached RAID unit

The NMS 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 > Main**.
2. Click the **Reset** boxes for the specific event counters to be reset.
3. Click **Apply Changes** to reset the counters.

Resources

The SSM uses the standard set of resources attributes that report on the service health and all computational, storage resources, and network resources. In addition, the SSM 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.

4 Troubleshooting

Troubleshooting consists primarily of how to respond to various alarms. This chapter includes a list of common alarms followed by a detailed table with the alarm codes and the suggested response. If a serious issue arises to which no answer is provided in this chapter, contact HP Support.

The chapter ends with examples and tips for basic troubleshooting:

- LDR storage troubleshooting
- grid ingest testing

Common Alarms

Table 14 lists common alarms that are usually no cause for concern as long as trends do not develop.

Table 14 Common Alarms

Category	Code	Service	Notes
HTTP	HEIS HEIP HEIG HEID	LDR	HTTP protocol alarms such as Incoming Sessions Failed (HEIS), Inbound PUTs Failed (HEIP), Inbound GETs Failed (HEIG), and Inbound DELETEs - Failed (HEID) are not usually a cause for concern unless the error count escalates. In general, these alarms occur during periods of high load or due to temporary network disruptions. The operations that have failed are retried automatically and usually succeed. NOTE A failed GET will usually be visible to a client as a delayed, failed, or timed out retrieve. A failed GET may indicate that an object has been lost from the grid although this typically triggers other alarms on the FSG or elsewhere.
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.
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 14 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. If available memory falls below 50 MB, contact HP Support.
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. NOTE This alarm is safe to ignore only if the events that triggered the alarm have been investigated.

Alarm Reference Tables

The NMS comes with a series of pre-configured alarms when the system is implemented. Each alarm can be configured to suit your grid management needs. 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.

Alarms are listed alphabetically by alarm code in this chapter.

A

Code	Name	Recommended Action
ABRL	Available Attribute Relays test	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, on an LDR, both the HTTP and DICOM components go into the "Redirect" state. Check and restore connections to a CMS services 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 Support.
ADCE	ADC Device State	If the device goes into "Standby", continue monitoring and if the problem persists, contact HP Support. If the device goes into "Offline" and there are no known server hardware issues (server unplugged) or a scheduled downtime for the device, contact HP Support.
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.

Code	Name	Recommended Action
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	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. Alarms are based on the queue being over a specified threshold. During heavy loads the queue can go over 100,000, but at that level it should be closely monitored. If the alarm persists, view a chart of the queue size. If the number continues increasing without ever decreasing, contact HP Support.
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 (as described in the *Administrator Guide*).
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", check to see if 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.
ARIS	Inbound Relay Sessions	If the number of inbound relay sessions grows too large it may indicate that the grid has become unbalanced, or that it requires expansion. Contact HP Support.

Code	Name	Recommended Action
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. 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. After 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.
ARRV	Verification Failures	Contact HP Support to diagnose and correct this problem. After 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. After the issue has been identified and corrected, reset the counter of failures on the configuration tab of ARC > Store.

Code	Name	Recommended Action
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. If the device is “Unavailable” and there are no known server hardware issues (server unplugged) or scheduled downtime for the device, 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.
BTOF	Offset	You are notified if the service time is too different (hours) 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
CISF	Incoming HTTP Sessions 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 > HTTP component to determine if the failures are client side or node side. Client side faults indicate a problem with the remote entity. To reset the counter: on the **CLB > HTTP > Configuration** page, select **Reset HTTP Counts**. If consistent faults appear on the node side, contact HP Support.
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 “Unavailable” 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. 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. If the device goes into “Unavailable” and there are no known server hardware issues (server unplugged) or scheduled downtime for the device, contact HP Support.
CMST	CMS Status	If you receive an error on CMST, contact HP Support.

Code	Name	Recommended Action
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.
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	More than 50,000 is considered a minor alarm to investigate further. More than 1,000,000 is considered a critical condition. 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 <= 15% Major <= 8% 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.

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 after the cache has been sufficiently cleared in order to allow the FSG to start caching files again.
FCSA	FSG Client Service Status	Restart the service. If the problem persists, contact HP Support.
FCSS	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	FSG Cluster Status	If the cluster is in the Vulnerable state, only one of the two FSGs in the cluster is online and available to provide services. Another failure will render FSG services unavailable: take action to 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. Neither FSG service is available and the cluster cannot recover from the fault automatically. Contact HP Support.
FDPP	Not used by the HP MAS.	N/A
FGWA	FSG Device Status	FSG devices need to be connected to at least one LDR to process transactions. If the service indicates “Waiting for HTTP Services”, the FSG either did not find at least one LDR that was providing HTTP services or did not find an ADC that permitted it to query for such LDRs. Ensure that at least one ADC is online and available, and that at least one LDR is online, available, and providing HTTP services (both ingest and retrieve). “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 Device 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. If the device is “Unavailable” and there are no known server hardware issues (server unplugged) or scheduled downtime for the device, contact HP Support.
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 fifty 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: on the **FSG > Storage > Configuration** page, **select Reset Retrieve from Grid Failure Coun**t.
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.

Code	Name	Recommended Action
FRGR	Files Retrieved from the Grid (Retrying)	Values of five or more are a minor concern. Monitor the situation. If the values do not decrease, contact HP Support.
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 a minor concern, 10,000 a major concern, and 100,000 a critical concern. When FSGP exceeds the critical level, file creation is forbidden on the FSG. 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)	Values of five or more are a minor concern. Monitor the situation. If the values do not decrease, contact HP Support.
FSRI	Files Stored to Grid Reingested	If there is no location in the grid for an object when an object is swapped out of FSG cache, this alarm is raised. FSRI should resolve on its own as the object is stored to a location in the grid. To reset the counter: on the **FSG > Storage > Configuration** page, select **Reset Store to Grid Reingested Count**. If the problem persists, contact HP Support.
FSTA	FSG Total Cache Available	Cache space is managed by the service. Space is freed as objects in the cache are no longer needed. Values less than 20 GB trigger alarms. Ensure all CMS and LDR services are operating normally. Check the following attributes in the FSG > Storage component: “Files Stored to Grid - Pending”, “File Store Rate”, and “File Retrieved Rate”. If these measures of cache activity are high, attempt to reduce demand until the alarm clears. If the problem persists, contact HP Support.
FSTS	FSG Startup Condition.	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. "Restored" may also be reported after initial install if a scheduled or forced backup has completed on another FSG prior to initial install. If there is an existing backup of the filesystem in the grid when the FSG is installed, the FSG will automatically initiate a filesystem restore.

H

Code	Name	Recommended Action
HEID	Inbound DELETE Failed	Values lower than 50 are considered a minor issue. These alarms indicate a problem with executing specific activities in the HTTP protocol. In general these alarms can occur during periods of high load or due to temporary network disruptions. Check the log at the remote entity to determine if the errors are originating with the remote workstation (entity) or the grid. If the entity shows the errors, then adjust the configuration of the entity. If the errors are on the grid side of the connection, contact HP Support. NOTE Deleting or renaming files recently stored to the FSG can produce this alarm if the FSG was actively relaying the file to the grid. Reset the counter. To reset the counters: on the **LDR > Configuration >** page, select **Reset HTTP Counts**.
HEIG	Inbound GET Failed	
HEIH	Inbound HEAD Failed	
HEIO	Inbound OPTION Failed	
HEIP	Inbound PUT Failed	
HEIS	Incoming Sessions Failed	
HEIT	Inbound POST Failed	
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, Restarting the Server (page 136). 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.

LM

Code	Name	Recommended Action
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. If the device is “Unavailable” and there are no known server hardware issues (server unplugged) or scheduled downtime for the device, 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.
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 support normal grid operation. Client requests may time out when nodes are in this condition. Contact HP Support.
MRun	Metadata with Unachievable ILM Evaluations	Applies to distributed CMSs. 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, Restarting the Server (page 136). 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.
NDBC	NMS Connectivity Status	If the service indicates “Authentication Failed, Retrying”, restart the service. If the problem persists, contact HP Support.
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 don’t 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.

Code	Name	Recommended Action
NSCE	NMS Device 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.
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.
NTBR	Free Tablespace	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, you should contact HP Support to arrange for the database allocation to be enlarged.
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.
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.

Code	Name	Recommended Action
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.
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	
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.
NTSU	NTP Status	If the NTP daemon is not running, contact HP Support.

Code	Name	Recommended Action
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. This value is persistent: it is not updated after 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.
ORun	Objects with Unachievable ILM Evaluations	Applies to synchronized CMSs. 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. After it is brought back up, the monitoring module 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 (%)	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.
PSN*n*	Array Status	If the Array Status is not "No Errors" (0), check the reason. If it is either rebuilding or in recovery, wait for a period of time for the array to restore to normal state; if the problem persists contact HP Support. If a drive has failed, replace the drive and allow time for it to rebuild. If the array is unknown or failed, check for error indicators. If the problem persists, contact HP Support.

Code	Name	Recommended Action
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.
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.

Code	Name	Recommended Action
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.
RPER	Replication Errors	Investigate and clear any replication backlogs. To reset the count: on the **FSG > Replication > Configuration** page, select **Reset Replication Errors Count**. Contact HP Support.
RPTE	Object Replication State	If the Replication State is offline, check the status of the replication component and correct any problems.
RPTU	Object Replication Status	If the status is “Waiting for Storage”, check the storage subsystem for errors. If the status is “Storage Unavailable”, the storage may have filled, and be in a read only state.
RSTT	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	FSG 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. If the status is “Error”, contact HP Support.

ST

Code	Name	Recommended Action
SAVP	Total Space Available	Adjusting the alarm threshold allows you to proactively manage when additional storage needs to be purchased. If the available 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.
SCAS	Grid Task Status (Configuration)	See SOAS.
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. Typically above 10 should be investigated as it indicates a fairly high load on the system. Over 25 is considered a critical 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.
SMST	Log Monitor	If the Log Monitor has non-zero values persisting for a period of time, 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.
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.

Code	Name	Recommended Action
SOAS	Grid Task Status (Overview)	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"). After 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.
SOHR	Historical Grid Task Status	If the status of the historical grid task is aborted, investigate the reason and run the task again if required.
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, particularly if the grid uses offline backups, which prevent replication messages from being processed on the secondary while the backup is in progress. This is mostly an issue with offline backups. With online backups, the FSG continues 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 this 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 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. If the device is “Unavailable” and there are no known server hardware issues (server unplugged) or scheduled downtime for the device, contact HP Support.

Code	Name	Recommended Action
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.
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 decline, contact HP Support.
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.
TPOP	Pending operations	A backlog of messages may indicate that the ADC is overloaded. Too few ADCs may be connected to this ADC (see the Tree Neighbors attribute). 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	If the Volume Status indicates “Offline”, verify that the write cache is enabled in the controller card. If it is, contact HP Support.

Code	Name	Recommended Action
VMWC	Write Cache	If the Write Cache status indicates “Unavailable”, contact HP Support.
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 filesystem 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	XCVP (% Completion)	Notice alarm when Foreground Verification is completed.

LDR Storage Troubleshooting

This section describes how to troubleshoot typical problems related to LDR storage:

- Object stores failures
- Offline LDRs
- Unachievable ILM
- Available storage space

Object Store Failures

The underlying data storage on a Storage Node is divided into “object stores” or storage volumes, which are partitions that act as mount points for the storage. You can view the list of object stores for each LDR using the NMS interface (go to **LDR > Storage > Overview**).

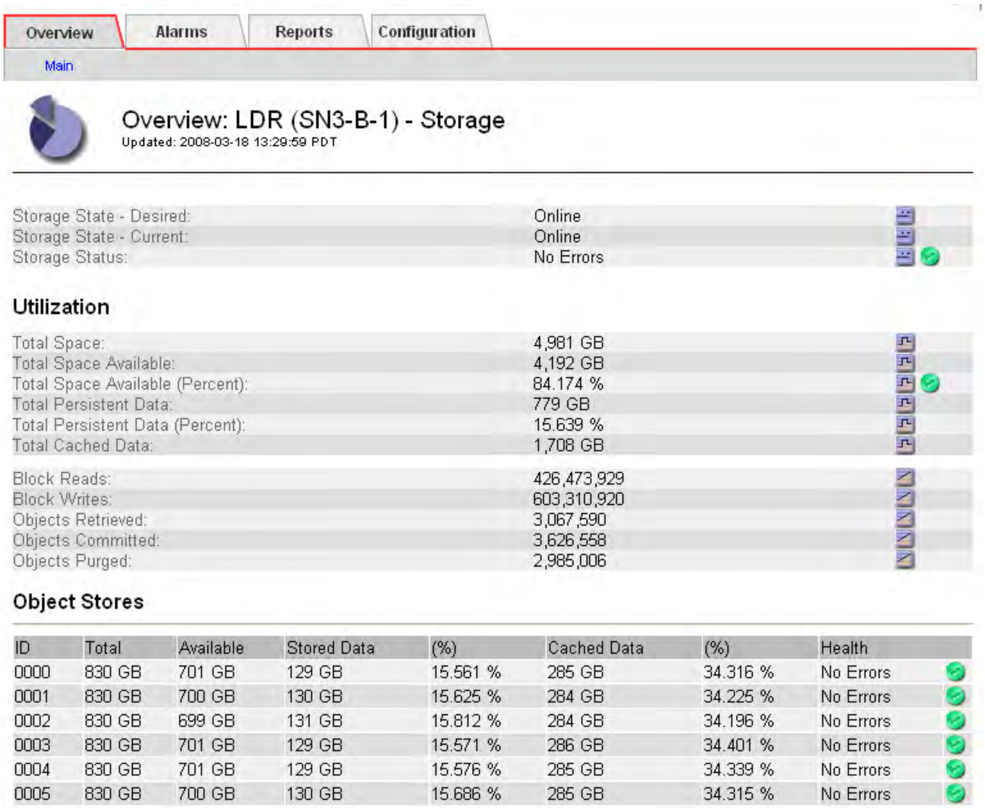

Overview: LDR (SN3-B-1) - Storage
Updated: 2008-03-18 13:29:59 PDT

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

Utilization

Total Space:	4,981 GB
Total Space Available:	4,192 GB
Total Space Available (Percent):	84.174 % ●
Total Persistent Data:	779 GB
Total Persistent Data (Percent):	15.639 %
Total Cached Data:	1,708 GB
Block Reads:	426,473,929
Block Writes:	603,310,920
Objects Retrieved:	3,067,590
Objects Committed:	3,626,558
Objects Purged:	2,985,006

Object Stores

ID	Total	Available	Stored Data	(%)	Cached Data	(%)	Health
0000	830 GB	701 GB	129 GB	15.561 %	285 GB	34.316 %	No Errors ●
0001	830 GB	700 GB	130 GB	15.625 %	284 GB	34.225 %	No Errors ●
0002	830 GB	699 GB	131 GB	15.812 %	284 GB	34.196 %	No Errors ●
0003	830 GB	701 GB	129 GB	15.571 %	286 GB	34.401 %	No Errors ●
0004	830 GB	701 GB	129 GB	15.576 %	285 GB	34.339 %	No Errors ●
0005	830 GB	700 GB	130 GB	15.686 %	285 GB	34.315 %	No Errors ●

Figure 18 Object Stores on LDR

Depending on the nature of the failure, faults with a storage volume may be reflected in an alarm on the storage status or on the health of an object store. If a storage volume fails, you should repair the failed storage volume to restore the Storage Node to full functionality as soon as possible. If necessary, you can place the Storage Node in a read-only state so that the grid can use it for data retrieval while you prepare for a full restoration of the server. Contact HP Support for assistance.

Objects with Unachievable ILM (ORun)

In the following example, the grid topology is DC+DR (see Figure 19). There is one Storage Node at the DC, two Storage Nodes at the DR site, and one CMS at each site. Files are ingested at the DC and the ILM rules require that one copy be kept at the DC and one at the DR site.

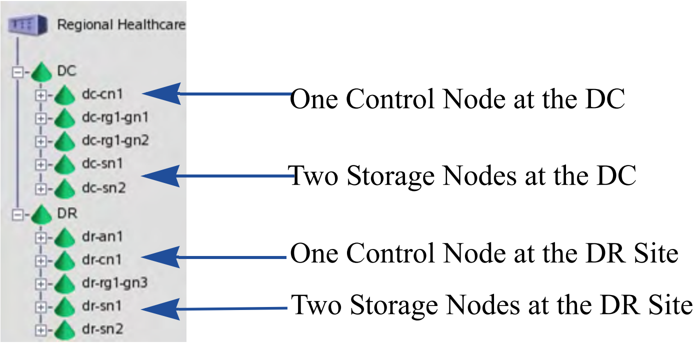

One Control Node at the DC

Two Storage Nodes at the DC

One Control Node at the DR Site

Two Storage Nodes at the DR Site

Figure 19 Grid Topology ORun Error

The ORun alarm has been triggered. To troubleshoot this particular problem, you would follow these steps:

1 Click **System Status** to display the alarms. From the colors of the grid topology tree in Figure 20 you can tell this grid has many alarms. The alarm of interest for this example is ORun.

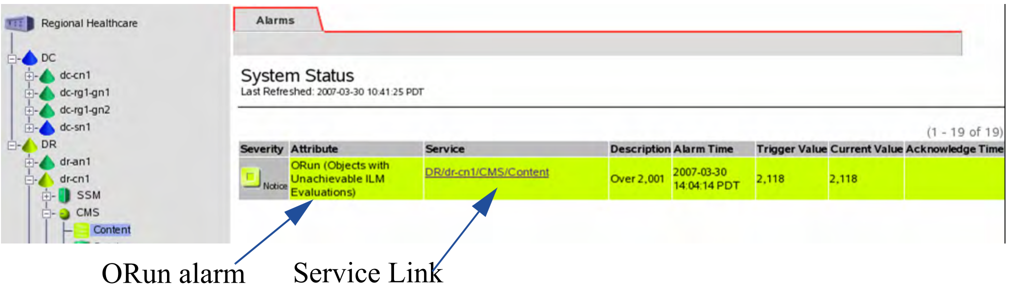

System Status
Last Refreshed: 2007-03-30 10:41:25 PDT

(1 - 19 of 19)

Severity	Attribute	Service	Description	Alarm Time	Trigger Value	Current Value	Acknowledge Time
Notice	ORun (Objects with Unachievable ILM Evaluations)	DR/dr-cn1/CMS/Content	Over 2,001	2007-03-30 14:04:14 PDT	2,118	2,118	

ORun alarm   Service Link

Figure 20 ORun Alarm

2 Click the attribute name to display more information (Figure 21). This alarms means that there is a high number of objects for which ILM business rules cannot be satisfied.

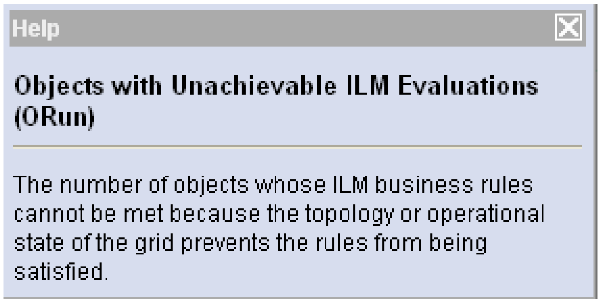

Help

Objects with Unachievable ILM Evaluations (ORun)

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.

Figure 21 ORun Online Help

3 Look up the troubleshooting information for ORun. The recommended action is:

— 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.

— Ensure that you understand the ILM rules for the grid so that you can understand where the grid must make copies.

4 Click the **Service** link in the System Status table to go to **CMS > Content > Alarms** and then click the **Overview** tab (or you could go there using the Grid Topology tree).

5 Notice the alarm indicator on the CMS Content Overview tab (Figure 22). There are almost 4,000 objects for which ILM rules are not being satisfied.

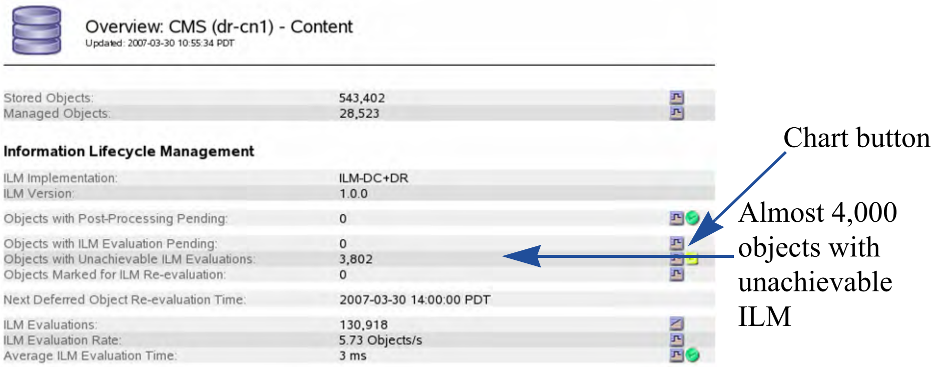

Overview: CMS (dr-cn1) - Content
Updated: 2007-03-30 10:55:34 PDT

Stored Objects:	543,402
Managed Objects:	28,523

Information Lifecycle Management

ILM Implementation:	ILM-DC+DR
ILM Version:	1.0.0
Objects with Post-Processing Pending:	0
Objects with ILM Evaluation Pending:	0
Objects with Unachievable ILM Evaluations:	3,802
Objects Marked for ILM Re-evaluation:	0
Next Deferred Object Re-evaluation Time:	2007-03-30 14:00:00 PDT
ILM Evaluations:	130,918
ILM Evaluation Rate:	5.73 Objects/s
Average ILM Evaluation Time:	3 ms

Figure 22 CMS Content Overview Tab with ORun Alarm

6 Click ⎍ to chart Objects with Unachievable ILM over the last hour. Adjust the time period as required. Make a note regarding when the problem started.

See Figure 23 (page 109).

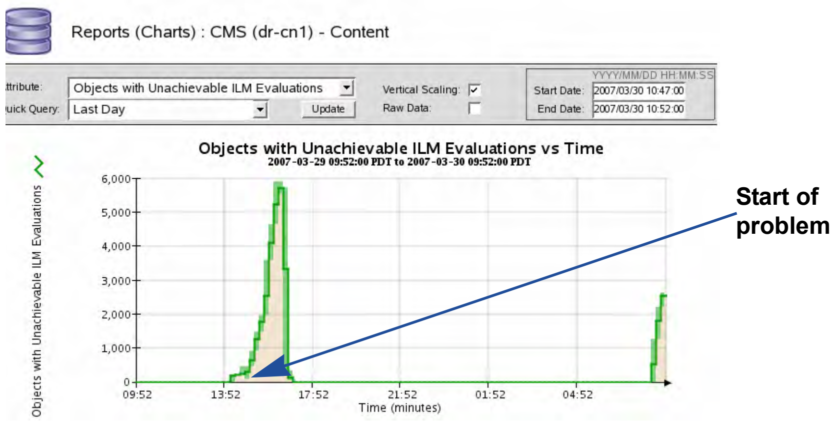

Reports (Charts) : CMS (dr-cn1) - Content

ttribute: Objects with Unachievable ILM Evaluations

uick Query: Last Day

Vertical Scaling: ☑

Raw Data: □

Start Date: 2007/03/30 10:47:00

End Date: 2007/03/30 10:52:00

Update
YYYY/MM/DD HH:MM:SS
Objects with Unachievable ILM Evaluations vs Time
2007-03-29 09:52:00 PDT to 2007-03-30 09:52:00 PDT
Objects with Unachievable ILM Evaluations
6,000
5,000
4,000
3,000
2,000
1,000
0
09:52
13:52
17:52
21:52
01:52
04:52
Time (minutes)

Start of problem

Figure 23 Trend for ORun Over Last Day

7 Check storage state as recommended in step 3. See in Figure 19 how the state of dc-sn1 is Unknown (the node color is blue). Therefore, storage is not available.

8 Chart the state of the LDR storage on that node to determine when storage became unavailable (go to **LDR > Storage > Overview** and click next to Storage State - Current to display a chart; adjust time as required).

9 Compare the charts for Storage State and Objects with Unachievable ILM (Figure 24) and notice the correlation between the two charts. The ORun alarm occurred because storage became unavailable.

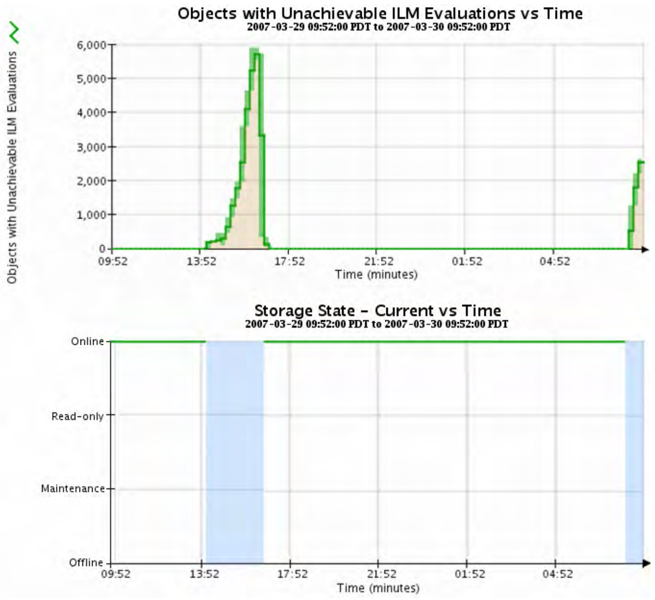

Figure 24 Trends for Objects with Unachievable ILM and Storage State

After the underlying issue is resolved, and the LDR storage is brought back online, the ORun alarm clears. The value of Objects with ILM Unachievable Evaluations is now 0 and the value of Objects with ILM Evaluation Pending is over 4,000.

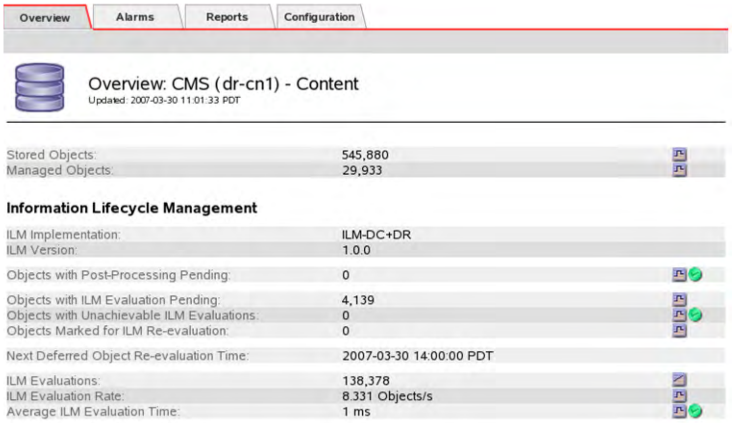

Overview: CMS (dr-cn1) - Content

Updated: 2007-03-30 11:01:33 PDT

Stored Objects:	545,880
Managed Objects:	29,933

Information Lifecycle Management

ILM Implementation:	ILM-DC+DR
ILM Version:	1.0.0
Objects with Post-Processing Pending:	0
Objects with ILM Evaluation Pending:	4,139
Objects with Unachievable ILM Evaluations:	0
Objects Marked for ILM Re-evaluation:	0
Next Deferred Object Re-evaluation Time:	2007-03-30 14:00:00 PDT
ILM Evaluations:	138,378
ILM Evaluation Rate:	8.331 Objects/s
Average ILM Evaluation Time:	1 ms

Figure 25 CMS Content After ORun Alarm Clears

Figure 26, which shows all three attributes (Objects with Unachievable ILM, Objects with ILM Evaluation Pending and Storage State) demonstrates the correlation between ILM evaluation and LDR storage state.

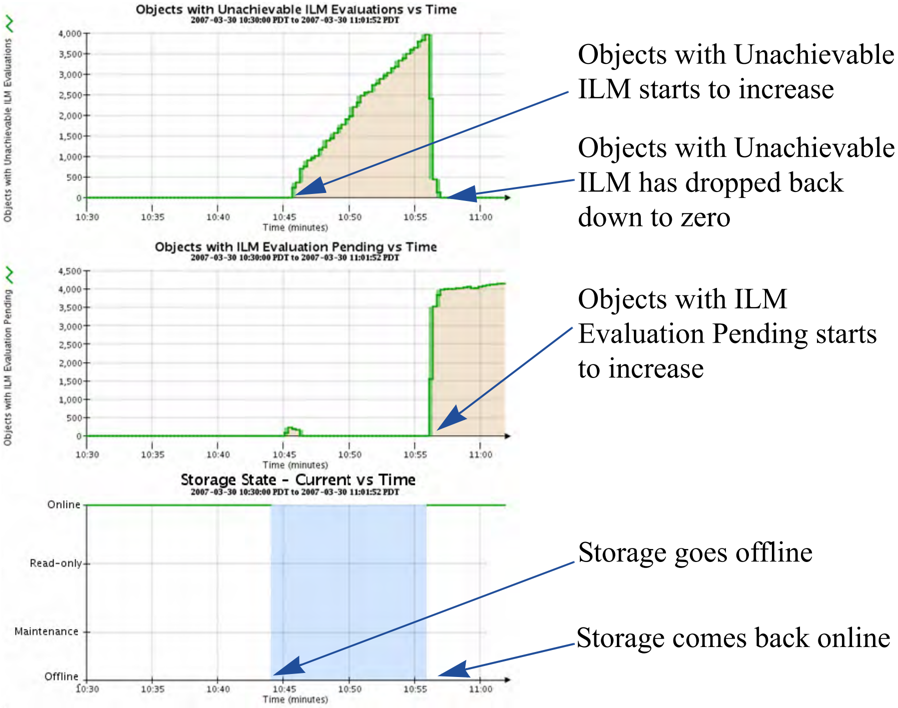

Objects with Unachievable ILM starts to increase

Objects with Unachievable ILM has dropped back down to zero

Objects with ILM Evaluation Pending starts to increase

Storage goes offline

Storage comes back online

Figure 26 Correlation Between Unachievable ILM and LDR Storage State

Total Space Available (SAVP)

In the following example, a Notice SAVP alarm has been triggered.

1 Click **System Status** to display the alarms (Figure 27). Notice the SAVP alarm.

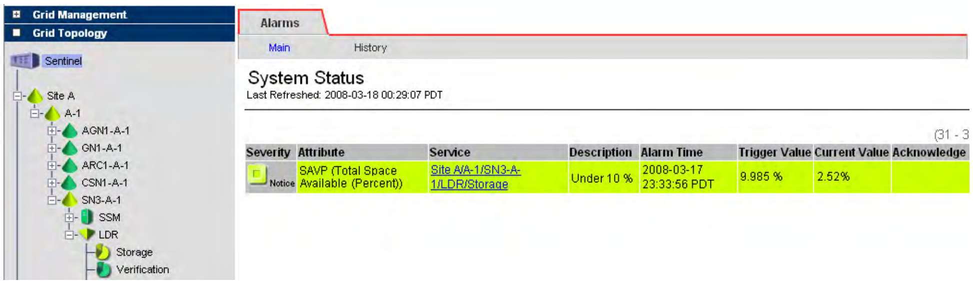

System Status
Last Refreshed: 2008-03-18 00:29:07 PDT

(31 - 3

Severity	Attribute	Service	Description	Alarm Time	Trigger Value	Current Value	Acknowledge
Notice	SAVP (Total Space Available (Percent))	Site A/A-1/SN3-A-1/LDR/Storage	Under 10 %	2008-03-17 23:33:56 PDT	9.985 %	2.52%	

Figure 27 SAVP Alarm

2 Click the attribute name to display more information (Figure 28).

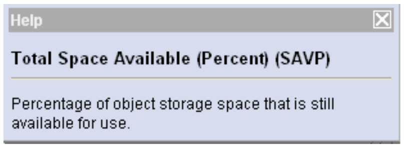

Total Space Available (Percent) (SAVP)

Percentage of object storage space that is still available for use.

Figure 28 SAVP Online Help

3 Look up the troubleshooting information for SAVP. The recommended action is to add storage or migrate data to archive media.

4 To determine how long you have until this LDR runs out of storage, go to the LDR Storage component that triggered the alarm. Click the **Service** link in the System Status table to go to **LDR > Storage > Alarms** and then click the **Overview** tab (or you could go there using the Grid Topology tree). Notice the alarm for Total Space Available (Percent).

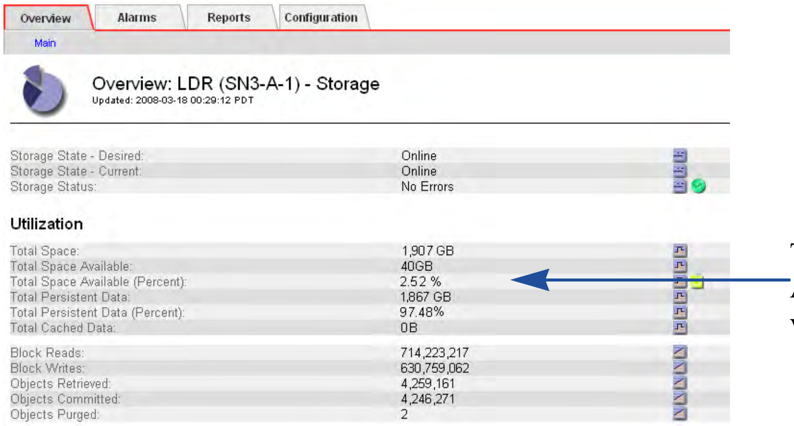

Overview: LDR (SN3-A-1) - Storage
Updated: 2008-03-18 00:29:12 PDT

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

Utilization

Total Space:	1,907 GB
Total Space Available:	40GB
Total Space Available (Percent):	2.52 %
Total Persistent Data:	1,867 GB
Total Persistent Data (Percent):	97.48%
Total Cached Data:	0B
Block Reads:	714,223,217
Block Writes:	630,759,062
Objects Retrieved:	4,259,161
Objects Committed:	4,246,271
Objects Purged:	2

Total Space Available is very low

Figure 29 LDR Storage Overview Tab with SAVP Alarm

5 Click ⎍ to chart space available over the last hour (Figure 30). Total space available has been steadily decreasing.

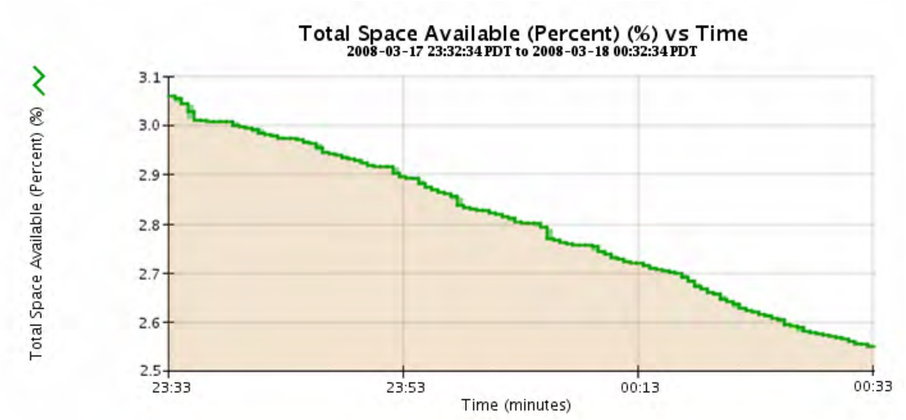

Figure 30 Trend for LDR Total Space Available - Last Hour

6 Look at the trend over the last day to estimate how fast storage capacity is decreasing (Figure 31).

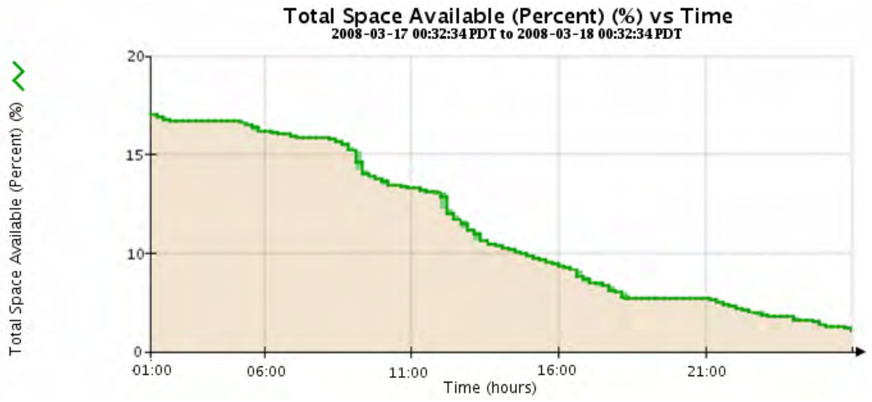

Figure 31 Trend for LDR Total Space Available - Last Day

Total available space on this LDR decreased about 15% in the last day.

NOTE If you have already added more storage capacity, an SAVP alarm on an old Storage Node is not normally an issue.

If there is no available storage on the other Storage Nodes in the grid, you need to add more storage space, for example, by adding more Storage Nodes, adding storage volume to an existing Storage Node, or migrating content to the Tape Node depending on the situation. For more information on how to add storage, contact HP Support.

Offline LDRs

A single root cause can lead to multiple alarms. Consider the situation shown in Figure 32.

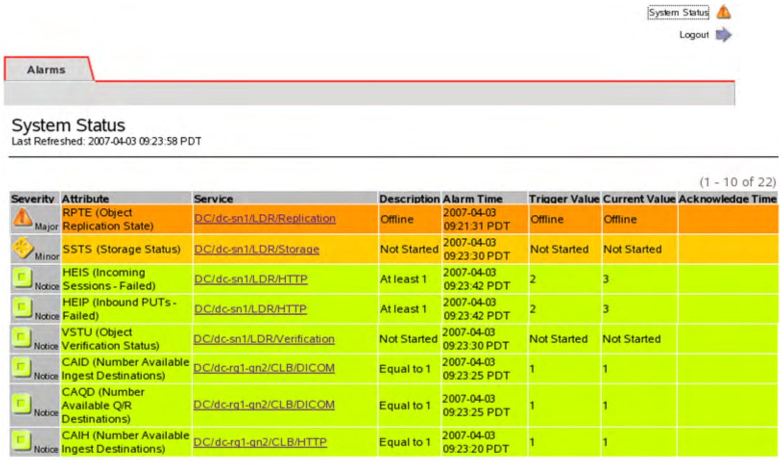

System Status
Last Refreshed: 2007-04-03 09:23:58 PDT

(1 - 10 of 22)

Severity	Attribute	Service	Description	Alarm Time	Trigger Value	Current Value	Acknowledge Time
Major	RPTE (Object Replication State)	DC/dc-sn1/LDR/Replication	Offline	2007-04-03 09:21:31 PDT	Offline	Offline	
Minor	SSTS (Storage Status)	DC/dc-sn1/LDR/Storage	Not Started	2007-04-03 09:23:30 PDT	Not Started	Not Started	
Notice	HEIS (Incoming Sessions - Failed)	DC/dc-sn1/LDR/HTTP	At least 1	2007-04-03 09:23:42 PDT	2	3	
Notice	HEIP (Inbound PUTs - Failed)	DC/dc-sn1/LDR/HTTP	At least 1	2007-04-03 09:23:42 PDT	2	3	
Notice	VSTU (Object Verification Status)	DC/dc-sn1/LDR/Verification	Not Started	2007-04-03 09:23:30 PDT	Not Started	Not Started	
Notice	CAID (Number Available Ingest Destinations)	DC/dc-rg1-gn2/CLB/DICOM	Equal to 1	2007-04-03 09:23:25 PDT	1	1	
Notice	CAQD (Number Available Q/R Destinations)	DC/dc-rg1-gn2/CLB/DICOM	Equal to 1	2007-04-03 09:23:25 PDT	1	1	
Notice	CAIH (Number Available Ingest Destinations)	DC/dc-rg1-gn2/CLB/HTTP	Equal to 1	2007-04-03 09:23:20 PDT	1	1	

Figure 32 Multiple Alarms on System Status Page

To troubleshoot this particular problem, follow these steps:

1 Click **System Status** to display all alarms (Figure 32). The alarms are automatically sorted by severity and then time.

2 Review all alarms. Click the attribute name to display more information (see Figure 33).

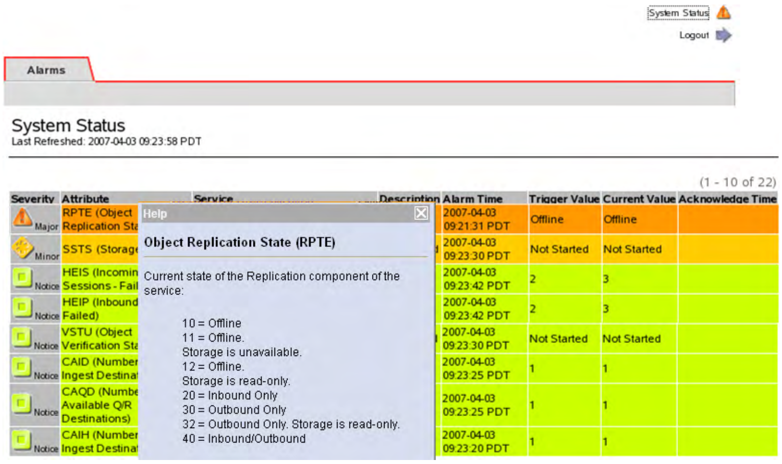

System Status
Last Refreshed: 2007-04-03 09:23:58 PDT

(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 (Storag			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 Destina			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 Destina			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

Figure 33 RPTE Attribute Description

3 See if the alarms have something in common. In this case, a possibility is the LDR Storage component.

4 Given that LDR Storage seems a potential root cause, go to **LDR > Storage > Overview** (see Figure 34).

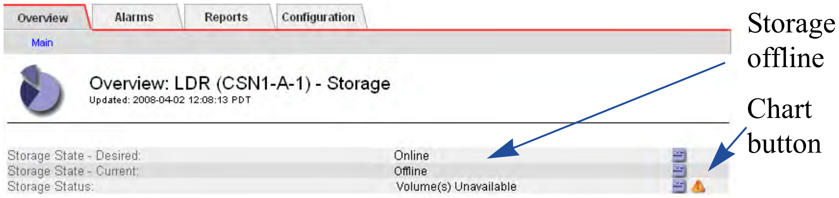

Overview: LDR (CSN1-A-1) - Storage
Updated: 2008-04-02 12:08:13 PDT

Storage State - Desired: Online
Storage State - Current: Offline
Storage Status: Volume(s) Unavailable

Figure 34 LDR Offline

5 Click ▤ to chart storage state over the last hour (see Figure 35). Adjust the time period as required. Study the chart. For example, look at when the LDR storage went offline.

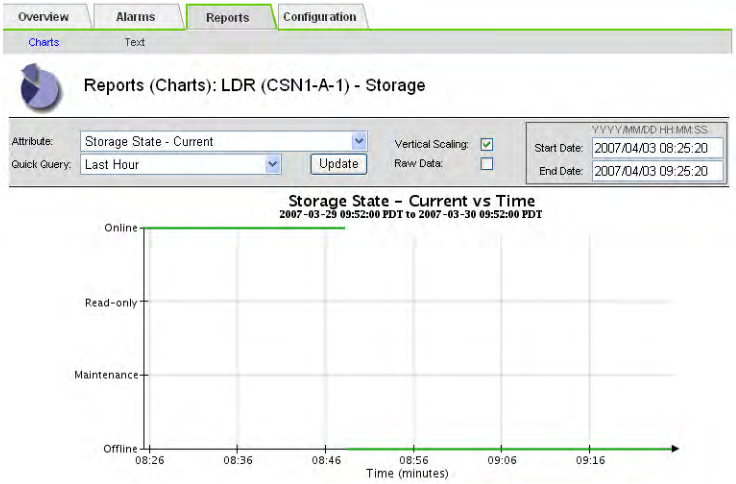

Reports (Charts): LDR (CSN1-A-1) - Storage

Attribute: Storage State - Current
Quick Query: Last Hour
Vertical Scaling: ☑
Raw Data: □
YYYY/MM/DD HH:MM:SS
Start Date: 2007/04/03 08:25:20
End Date: 2007/04/03 09:25:20

Figure 35 Current Storage State

6 Determine whether the issue is hardware error or user error and fix the problem.

7 After the problem is fixed, put the LDR storage back online.

8 Check that the alarms clear. Alarms can take a few minutes to clear as the changes propagate through the grid. In this case, a single change made to the grid resolved at least eight alarms.

Ingest/Retrieval Troubleshooting

Useful Attributes

Table 15 lists a number of attributes that can help you troubleshoot grid ingest and retrieval issues. Chart these attributes over time to gain some insight into the behavior of the grid. For more information on these attributes, see the *Grid Primer*.

Table 15 Ingest and Retrieve Attributes

Component	Name	Description
FSG Storage	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 a file into the grid.
	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 a file from the grid.
FSG Backup	Number of Files (PBSF)	Number of files included in the last backup for this share.
	Total File Size (PBSB)	Total size of all files referenced by the last backup for this share.

5 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.

Monitoring Grid Tasks

You can follow the progress of a grid task from the CMN Grid Tasks Overview tab (see Figure 36 and Table 16). Running grid tasks is restricted to accounts with Maintenance permissions such as the Admin and Vendor accounts.

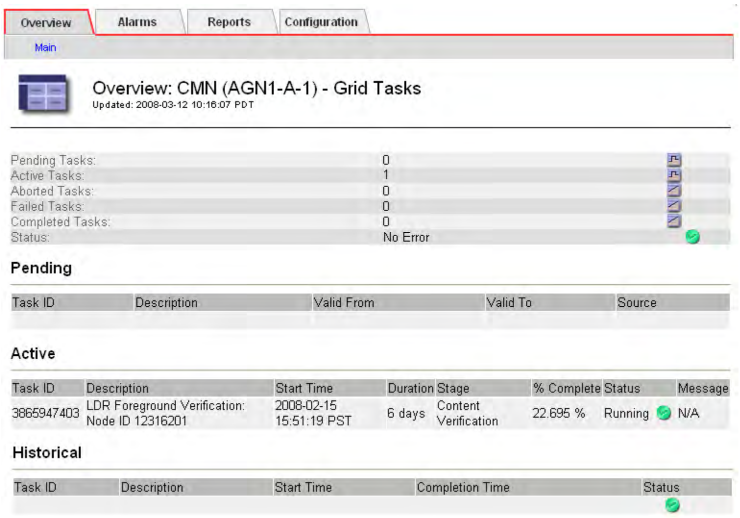

Overview: CMN (AGN1-A-1) - Grid Tasks
Updated: 2008-03-12 10:16:07 PDT

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

Pending

Task ID	Description	Valid From	Valid To	Source

Active

Task ID	Description	Start Time	Duration	Stage	% Complete	Status	Message
3865947403	LDR Foreground Verification: Node ID 12316201	2008-02-15 15:51:19 PST	6 days	Content Verification	22.695 %	Running ●	N/A

Historical

Task ID	Description	Start Time	Completion Time	Status
				●

Figure 36 Grid Tasks Overview

Grid tasks go through several 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 16 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.

Table 16 Grid Tasks Overview Fields *(continued)*

Field	Description
% Complete	Progress indicator for active tasks.
Duration	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 cancelled or expired or was aborted).

Executing Grid Tasks

Grid Tasks That Cannot Run Concurrently

In general, it is possible to run more than one grid task at a time. However, some grid tasks cannot run concurrently, as they require exclusive access to the same grid resources. The following table summarizes grid tasks that cannot run at the same time.

Table 17 Grid Tasks that Cannot Run Concurrently

Grid Task Name	Description	Grid Tasks that Cannot Run at the Same Time
Storage Node decommissioning	Moves all content off the specified Storage Node and then permanently removes it from the grid.	Grid Expansion: Initial Grid Expansion: Add Server Software Upgrade You can only run one grid task that affects a given Storage Node at a time. If any of the following grid tasks is active, you cannot run another task that affects the same Storage Node/LDR: • Storage Node decommissioning • Storage Node hardware refresh (affects both source and destination Storage Nodes) • LDR Foreground Verification • LDR Content Rebalancing It *is* possible to run these grid tasks concurrently for different Storage Nodes. For example, you can run LDR Foreground Verification on one Storage Node while you decommission a second Storage Node.
Storage Node hardware refresh	Moves all content from the source Storage Node to the destination Storage Node, and then permanently removes the source Storage Node from the grid. Used when replacing an existing Storage Node with a new server.	
LDR Foreground Verification: *NodeID*	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.	Grid Expansion: Initial Grid Expansion: Add Server You can only run one grid task that affects a given Storage Node at a time. If any of the following grid tasks is active, you cannot run another task that affects the same Storage Node/LDR: • Storage Node decommissioning • Storage Node hardware refresh (affects both source and destination Storage Nodes) • LDR Foreground Verification • LDR Content Rebalancing It *is* possible to run these grid tasks concurrently for different Storage Nodes. For example, you can run LDR Foreground Verification on one Storage Node while you decommission a second Storage Node.

Table 17 Grid Tasks that Cannot Run Concurrently *(continued)*

Grid Task Name	Description	Grid Tasks that Cannot Run at the Same Time
LDR Content Rebalancing: *NodeID*	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 interface.	You can only run one grid task that affects a given Storage Node at a time. If any of the following grid tasks is active, you cannot run LDR content rebalancing: • Storage Node decommissioning • Storage Node hardware refresh (affects both source and destination Storage Nodes) • LDR Foreground Verification • LDR Content Rebalancing It *is* possible to run these grid tasks concurrently for different Storage Nodes. For example, you can run LDR Foreground Verification on one Storage Node while you rebalance content on a second Storage Node.
Grid Expansion: Initial Grid Expansion: Add Server	Used when adding an additional grid node to an existing deployment.	Software Upgrade Enable DICOM Storage Node decommissioning Storage Node hardware refresh LDR Foreground Verification Bundle Commit Bundle Import
Software Upgrade: *Version*	Used when upgrading grid software to a new version.	Grid Expansion: Initial Grid Expansion: Add Server Enable DICOM Storage Node decommissioning Storage Node hardware refresh LDR Foreground Verification Bundle Commit Bundle Import

Table 17 Grid Tasks that Cannot Run Concurrently *(continued)*

Grid Task Name	Description	Grid Tasks that Cannot Run at the Same Time
Enable DICOM	Used to enable DICOM for a grid.	Grid Expansion: Initial Grid Expansion: Add Server Software Upgrade Bundle Commit Bundle Import
Bundle Import	Used when converting an unclustered Gateway Node replication group to a High Availability Gateway Node Cluster.	Grid Expansion: Initial Grid Expansion: Add Server Storage Node Decommissioning Storage Node Hardware Refresh Software Upgrade Enable DICOM Bundle Commit
Bundle Commit	May be used as part of a software upgrade procedure.	Grid Expansion: Initial Grid Expansion: Add Server Storage Node Decommissioning Storage Node Hardware Refresh Software Upgrade Enable DICOM Bundle Import

Running a Grid Task

NOTE New for Release 8.0

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 execute 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.

NOTE 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 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 Pausing Grid Tasks (page 123) and Aborting Grid Tasks (page 124)).

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.

Pausing 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.

Resuming 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.

Cancelling 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 NMS page to auto-refresh before the change is visible. Do not submit the change again.

Aborting Grid Tasks

A grid task can be aborted from the Active list so that it is no longer available for execution.

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**.
4. When the page refreshes, the Status of the grid task has changed to Paused.
5. Select **Abort** from the Actions menu.
6. Click **Apply Changes**.

The grid task moves to the Historical list with the description Aborted in the Status field. You must wait for the NMS 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 an aborted grid task again but first you have to remove it 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:

- expired—the “task valid before” time has already passed
- invalid—the task was not valid
- unauthorized—the task signature did not pass verification
- duplicate—this task has been previously loaded into the CMN
- aborted—the task was aborted

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 Support.

Grid Task in Error State

If a task enters an error state, Grid Task Status alarms (SOAS and SCAS) are triggered. Look up the task message on the CMN > Grid Tasks > Configuration 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 Support.

Submitting 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 use to access the NMS.
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. Type **<Ctrl>+c** to copy the selected text to the clipboard.
5. Log in to the NMS using the Vendor account.
6. Go to **CMN > Grid Tasks > Configuration > Main**.

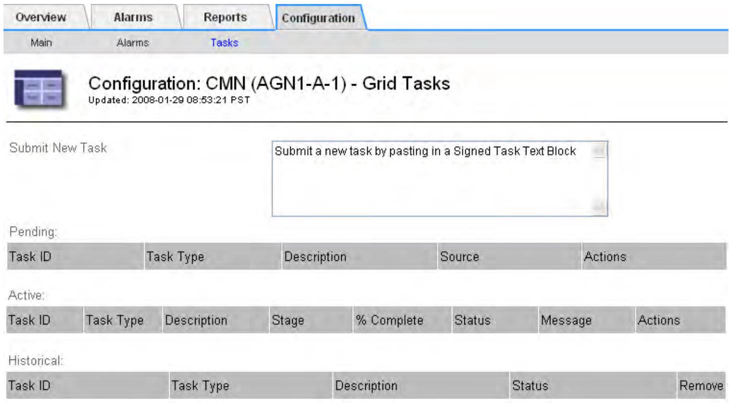

Configuration: CMN (AGN1-A-1) - Grid Tasks
Updated: 2008-01-29 08:53:21 PST

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 37 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. Type **<Ctrl>+v** to paste the Task Signed Text Block from the clipboard 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, you can 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.

6 Introduction to Server Manager

Each server used in a HP MAS deployment runs the Server Manager application. The Server Manager program is used to supervise 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 report faults.

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.

Server Manager provides a graphical interface on the local console of the server, enabling coarse (whole server) control and monitoring of the device. 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. All operations preserve graceful service disconnections and restarts; grid operations are smoothly completed before the services are removed from the grid. Services and support applications are stopped in a sequence that avoids breaking any interdependencies.

Server Manager continues to run at all times, monitoring the services and automatically restarting any that go offline unexpectedly. Should a service go offline unexpectedly or a server suffer an unexpected power cycle or reset, Server Manager ensures all services are brought back online without the need of manual intervention.

Services

The first part of this manual (dealing with the NMS) uses specific terminology for locations, servers, services, and so on. This discussion of the Server Manager refers to “services” in a broader sense, including operating system processes such as networking and file systems and third-party modules such as the MySQL database application. Where reference to only the grid services is needed, the word “grid” is always used.

Capabilities

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
- 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 it must be manually restarted. This enables you to safely power down a server for hardware maintenance.
- Automatically start services if the server is power cycled or reset, and to recover from unintentional restarts.
- Monitor services on an ongoing basis and restart them as needed.

Architecture Overview

Server Manager is implemented in three parts: a back-end daemon (runit) that provides the core functionality of Server Manager, a Graphical User Interface Backend, and a GUI. The GUI Backend is intended to run tightly coupled with the back-end daemon so that it is always available to manage the server. The GUI is run separately, offering administrators a graphical local console for monitoring the server.

At installation, Server Manager is configured to automatically start when the operating system does. Server Manager is always running and available, and is intended to be left unattended.

The core functionality provided by runit provides the logic to:

- Manage dependencies between services, sequentially verifying and starting services to bring the grid server up and gracefully connect to the grid.
- Sequentially stop services to gracefully disconnect from the grid.
- Monitor grid services to restart services that go offline unexpectedly.
- Automatically restart services if the server is restarted or if the power is cycled.
- Respond to commands from the GUI interface.

Automatic Startup and Shutdown

During system startup, Server Manager is automatically started by the operating system. Server Manager 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.

The GUI interface to Server Manager is one of the services started through this sequence of scripts. It is placed early in the sequence so that it can report the status of other services as they start up or shut down.

Display Components

The Server Manager presents a Graphical User Interface (GUI) on the local console of the server. The display uses five panels to present information and controls.

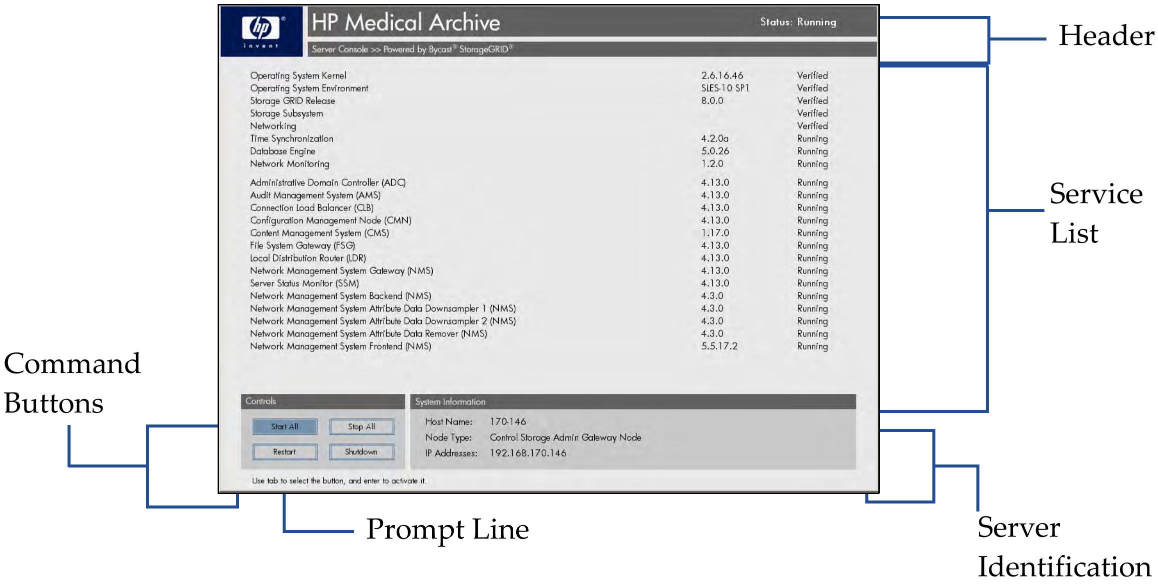

HP Medical Archive   Status: Running

Server Console >> Powered by Bycast® StorageGRID®

```
Operating System Kernel                                        2.6.16.46    Verified
Operating System Environment                                   SLES-10 SP1  Verified
Storage GRID Release                                           8.0.0        Verified
Storage Subsystem                                                           Verified
Networking                                                                  Verified
Time Synchronization                                           4.2.0a       Running
Database Engine                                                5.0.26       Running
Network Monitoring                                             1.2.0        Running

Administrative Domain Controller (ADC)                         4.13.0       Running
Audit Management System (AMS)                                  4.13.0       Running
Connection Load Balancer (CLB)                                 4.13.0       Running
Configuration Management Node (CMN)                            4.13.0       Running
Content Management System (CMS)                                1.17.0       Running
File System Gateway (FSG)                                      4.13.0       Running
Local Distribution Router (LDR)                                4.13.0       Running
Network Management System Gateway (NMS)                        4.13.0       Running
Server Status Monitor (SSM)                                    4.13.0       Running
Network Management System Backend (NMS)                        4.3.0        Running
Network Management System Attribute Data Downsampler 1 (NMS)   4.3.0        Running
Network Management System Attribute Data Downsampler 2 (NMS)   4.3.0        Running
Network Management System Attribute Data Remover (NMS)         4.3.0        Running
Network Management System Frontend (NMS)                       5.5.17.2     Running
```

Host Name: 170-146
Node Type: Control Storage Admin Gateway Node
IP Addresses: 192.168.170.146

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

Figure 38 Server Manager Interface Sample

Header

The top of the screen presents a header line, identifying the Server Manager application and its overall status.

When running normally, the header appears with a dark grey background. If *any* of the services is reporting an error, the background changes to red to draw immediate attention.

Status

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. However, if *any* of the services are in an error state, the Status reports “Error”, and the header background turns red.

Service List

The body of the display is a list of the services being monitored by the 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. 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.

If a grid service reports an error state, the status reported here indicates “Error” and the Server Manager header changes to a red background. Server Manager continuously monitors and reports the state of the service, attempting to restart the service after its dependencies are met. If a service is in an error state, it is not automatically restarted.

Command Buttons

Located in the bottom left corner are four command buttons used to stop and start services and reboot or halt the server. Use of these command buttons ensures the services enter and leave the grid gracefully.

The currently selected button is highlighted.

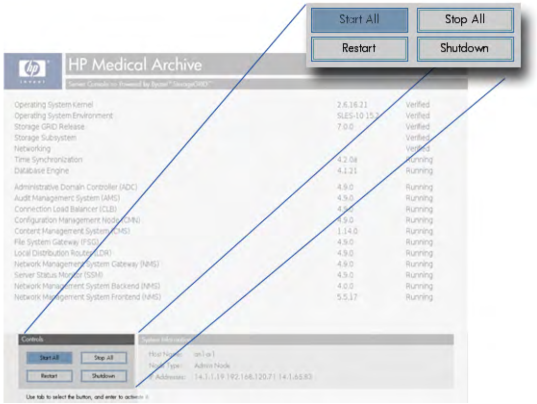

Operating System Kernel	2.6.16.21	Verified
Operating System Environment	SLES-10 15.2	Verified
Storage GRID Release	7.0.0	Verified
Storage Subsystem		Verified
Networking		Verified
Time Synchronization	4.2.0a	Running
Database Engine	4.1.21	Running
Administrative Domain Controller (ADC)	4.9.0	Running
Audit Management System (AMS)	4.9.0	Running
Connection Load Balancer (CLB)	4.9.0	Running
Configuration Management Node (CMN)	4.9.0	Running
Content Management System (CMS)	1.14.0	Running
File System Gateway (FSG)	4.9.0	Running
Local Distribution Router (LDR)	4.9.0	Running
Network Management System Gateway (NMS)	4.9.0	Running
Server Status Monitor (SSM)	4.9.0	Running
Network Management System Backend (NMS)	4.0.0	Running
Network Management System Frontend (NMS)	5.5.17	Running

Host Name: an1-a1
Node Type: Admin Node
IP Addresses: 14.1.1.19 192.168.120.71 14.1.65.83

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

Figure 39 Sample Selected Button

Use of these buttons is described in detail in Command Operations (page 135).

Server Identification

The bottom right portion 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 as specified in the original installation configuration file.

At the time a deployment is configured, a file is prepared that identifies each server and the services it hosts. This file is used to generate the installation disks. The server name used in that file is reported by Server Manager as the Host Name.

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

At the very bottom of the display is a prompt line, providing guidance on what actions you can take. This typically states: “Use tab to select the button, and enter to activate it.” See Command Operations (page 135) for more on using the buttons.

Using the Interface

Display Updates

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.

Command Input

The interface only supports use of the command buttons shown in the bottom left of the display.

The command buttons can be selected using the mouse or the <Tab> and <Enter> keys.

Initiating an Action

To initiate an action:

1. Use the mouse to click the desired command button.

   A confirmation dialog opens allowing you to confirm or cancel the action.

2. Type the Server Manager password (as found in the Passwords.txt file).

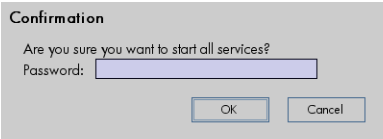

Confirmation

Are you sure you want to start all services?

Password:

Figure 40 Password Confirmation

3. Select **OK** to confirm the action (or Cancel to abort).

   The dialog closes and Server Manager returns to normal operation, and executes the action.

Command Operations

Monitoring Services

While running normally, the Server Manager is constantly monitoring the services. If a service fails, Server Manager attempts to restart it. If there are three failed attempts to start services within five minutes, the service goes down, a line is added to the `servermanager.log` file, and the service fails to start. Server Manager does not attempt another restart.

Starting Services

This command button is used to start services that have been stopped for upgrade or other configuration changes.

Upon server startup, or when commanded by using **Start All**, the 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.

The operating system kernel, environment, and hardware services for storage and networking are verified first. Following the operating system services, time synchronization is verified. This process can take several minutes while the system characterizes and builds trust in the time signal source. Since ntp does not depend on mysql, mysql and ntp start in parallel.

When all services are started, the server joins the grid. Any services already running are not disrupted.

Stopping Services

The command button **Stop All** can be used to gracefully stop services hosted on this 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.

When maintenance tasks are finished, you can either:

- **Start All** services to gracefully start the grid services and rejoin the grid, see Starting Services (page 135).

  —or—

- **Restart** the server; bootstrapping the operating system and then starting grid services, Restarting the Server (page 136).

Restarting the Server

The **Restart** command gracefully stops grid services, and also brings down the operating system to perform an automatic reboot. This is useful for resetting a server that has failed, or starting a new configuration.

The Server Manager performs the same sequence as the Stop All command (Stopping Services (page 135)), 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.

When a server is restarted, the Server Manager application is automatically restarted and resumes normal operation, restarting all services.

Shutting Down the Server

If the server must be powered down for hardware maintenance, upgrades, or reconfiguration, the **Shutdown** command should be used. Similar to Restart, this 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 (Stopping Services (page 135)), 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 its shutdown.

A Connectivity

The appendix describes connection requirements for the NMS interface and the server command shell.

Service Laptop Requirements

The requirements for the service laptop used to connect to the NMS interface are:

- Microsoft Windows operating system.
- Network Interface Card (NIC) adapter for connection to the customer's network. You should also have a cable in case one is not readily available at the customer's site.
- CD drive (to read the Software Activation backup CD)
- Microsoft Internet Explorer v6.0 SP2 or Microsoft Internet Explorer v7.0

  JavaScript and cookies *must* be enabled.

Browser Settings

Internet Options Settings

You need to verify that the Internet Explorer settings for temporary internet files, security and privacy are 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.

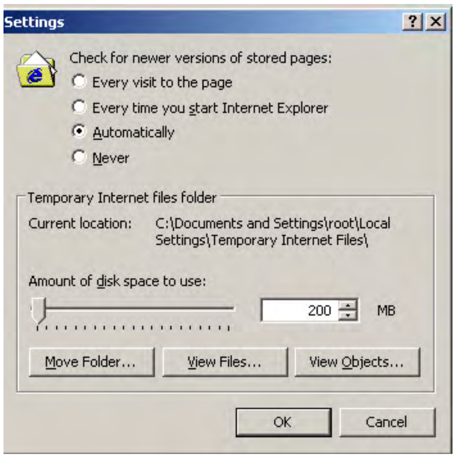

Check for newer versions of stored pages:

○ 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 41 Temporary Files Setting

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

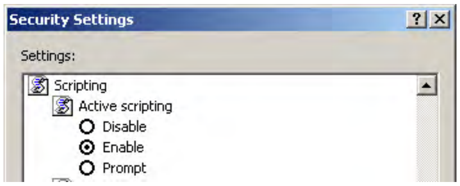

Settings:

Scripting

Active scripting

○ Disable

● Enable

○ Prompt

Figure 42 Active Scripting Setting

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

NMS Connection Procedure

Connecting to the NMS interface at the customer site requires access to the customer’s network in close proximity to the HP MAS cabinets.

To connect to the NMS interface:

1 Work with the customer system administrator to establish the physical network connection to the service laptop. Using the customer’s network rather than a direct connection within the cabinet verifies that the interface is accessible using the same infrastructure the customer uses.

2 Insert the Software Activation backup CD, and open:

— `Configuration.txt`

— `Passwords.txt`

3 From the `Configuration.txt` file, note the IP address of the Admin Node on the customer network. This is needed to access the NMS interface.

4 From the `Passwords.txt` file, note the NMS password for the Vendor account or the Admin account.

5 Launch the web browser.

6 Open the address `https://<IP_address>`.

where `<IP_address>` is the address of the Admin Node hosting the CMN service on the customer network specified in the `Configuration.txt` file.

Security Certificate

Depending on your version of Windows and web browser, you may be prompted with a Security Alert window when you access the NMS URL.

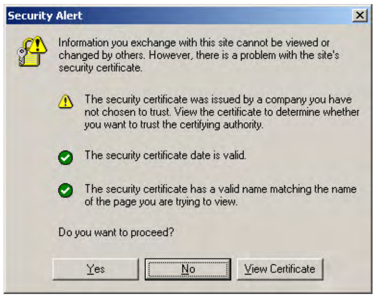

Security Alert

Information you exchange with this site cannot be viewed or changed by others. However, there is a problem with the site's security certificate.

The security certificate was issued by a company you have not chosen to trust. View the certificate to determine whether you want to trust the certifying authority.

The security certificate date is valid.

The security certificate has a valid name matching the name of the page you are trying to view.

Do you want to proceed?

Figure 43 Security Alert Window

If this appears, you can either:

- Click **Yes** to proceed with this session. The alert will appear again the next time you access this URL.

  —or—

- Click **View Certificate** and install the certificate using the installation wizard so that you no longer receive the alert.

Log In

You are presented with the NMS login window.

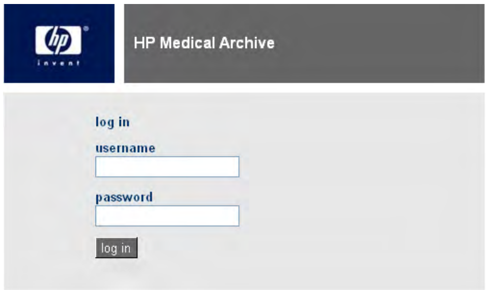

Figure 44 NMS Login Window

To log in:

1 Enter the username Vendor for full access to the NMS. If you are not making grid-wide configuration changes, you can also use the Admin account. For more information, see the *Administrator Guide*.

2 Enter the password for the NMS specified in the `Passwords.txt` file.

Enable Pop-ups

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

To enable pop-ups:

1 Select **Tools > Pop-up Blocker > Turn off Pop-up Blocker** from the Internet Explorer main menu to open the Pop-up Blocker Settings dialog.

NOTE The menu option is a toggle. If the blocker was already disabled, the menu option is to Turn on the Pop-up Blocker.

You may now use the NMS Account Management feature to change access passwords.

Log Out

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

To log out:

1 Click **Logout** Logout ➡ located at the top right corner of the screen. The logging out message appears.

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

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

Command Shell Access Procedures

Log In

To log in at the console of a server:

1. Press **<Alt>+<F1>** to access a command shell.
2. Enter the account name `root`.
3. Enter the password for the server listed in the `Passwords.txt` file.

Accessing a Server Remotely

There are two ways to access a server remotely using SSH:

- from the primary Admin Node using the SSH key password
- from any server using the remote server password

By default, the Admin Node that hosts the CMN service acts as an SSH access point for the grid. When passwordless access is enabled, you can SSH to other servers from the Admin Node without entering their password. The procedure to enable passwordless access to grid nodes is described in the Network Configuration chapter of the Administrator Guide.

To access a remote server using the SSH key password:

1. Log in to the Admin Node hosting the CMN service.
2. Enter: `ssh <IP_address>`

   where `<IP_address>` is the IP address of the remote server.

   —or—

   Enter: `ssh <hostname>`

   where `<hostname>` is the name of the server.

   If the remote server was added as a result of a grid expansion, use the server IP address instead of the hostname.
3. When prompted, enter the SSH private key password listed in the SSH Access Password section of the `Passwords.txt` file.

You may now execute commands on the remote server as if you were logged in to that server directly.

To access a remote server using the remote server password:

1. Log in to any local server.
2. Enter: `ssh <IP_address>`

   where `<IP_address>` is the IP address of the remote server.

   —or—

   Enter: `ssh <hostname>`

where `<hostname>` is the name of the server.

If the remote server was added as a result of a grid expansion, use the server IP address instead of the hostname.

3. When prompted, enter the password for the remote server listed in the `Passwords.txt` file.

You may now execute commands on the remote server as if you were logged into that server directly.

Log Out

To log out of a command shell session:

1. Enter exit to close the current command shell session.
2. Press **<Alt>+<F7>** to return to the Server Manager GUI.

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.

ADC

Administrative Domain Controller—A software component of the HP MAS product. 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 MAS product.

Admin Node

A building block of the HP MAS product. The Admin Node provides services for the web interface, grid configuration, and audit logs.

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 product. 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.

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 product. 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.

bundle
A structured collection of configuration information used internally by various components of the grid. Bundles are structured in container format.

Bycast Enablement Layer
The Bycast Enablement Layer 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.

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, 192.168.110.0/24.

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 product. The CLB service provides a gateway into the grid for clients connecting via HTTP protocol. The CLB service is part of the Gateway Node.

CMN
Configuration Management Node— A software component of the HP MAS product. The CMN service manages system-wide configuration and grid tasks. The CMN service is found on the Admin Node.

CMS
Content Management System—A software component of the HP MAS product. 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.

Control Node

A building block of the HP MAS product. 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.

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.

DR

Disaster Recovery.

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.

FSG

File System Gateway—A software component of the HP MAS product. 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, one FSG is a primary and all others are secondaries. The primary FSG allows clients read and write access to the grid, while storing file system information (stubs) 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 product. The Gateway Node provides connectivity services for NFS/CIFS file systems and the HTTP protocol.

grid node

The name of the HP MAS product 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.

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 product. The LDR service manages the storage and transfer of content within the grid. The LDR service is found on the Storage Node.

metadata

Information related to or describing an object stored in the grid, for example file ingest path or ingest time.

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 product. 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.

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.

provisioning

The process of editing the Grid Specification File (if required) and generating or updating the SAID package. This is done on the Admin Node using the provision command. The new or updated SAID package is saved to the Provisioning USB flash drive. See Grid Specification File and SAID.

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”.

server
Used when referring specifically to hardware.

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 software.

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 TLS.

Storage Node
A building block of the HP MAS product. 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.

SUSE
See SLES—SUSE Linux Enterprise Server.

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