StorageGRID® 9.0
Grid Primer

NetApp, Inc.
495 East Java Drive
Sunnyvale, CA 94089 U.S.A.
Telephone: +1 (408) 822-6000
Fax: +1 (408) 822-4501
Support telephone: +1 (888) 463-8277
Web: http://www.netapp.com
Feedback: doccomments@netapp.com

Part number: 215-06989_A0
June 2012

Copyright and trademark information

Copyright information

Trademark information

Contents

1

Introduction to the StorageGRID System

An overview of grid features, components, and configuration

Introduction

This guide is written for users and system administrators who are new to the StorageGRID system. It is also for grid operators whose main function is to monitor the StorageGRID system on a day-to-day basis, in particular operators who log in to the NMS Management Interface (MI) using user-level permissions.

The NetApp StorageGRID System

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

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

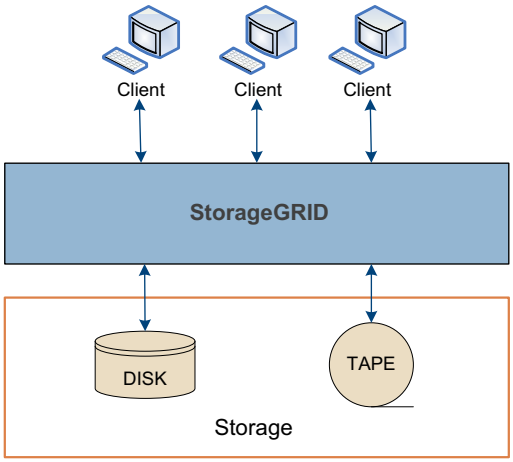

Figure 1: Storage Management System for Massive Volumes of Fixed-Content Data

The StorageGRID system:

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

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

Deployment Topologies

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

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

- Data Center
- Data Center + Disaster Recovery
- Data Center + Disaster Recovery and Satellites
- Federated

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

Data Center

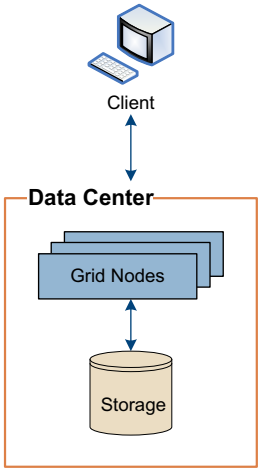

Figure 2: Data Center Deployment

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

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

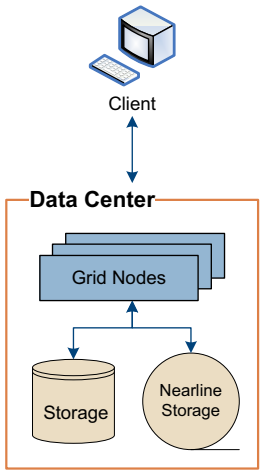

Figure 3: Data Center + Archive Deployment

Data Center + Disaster Recovery

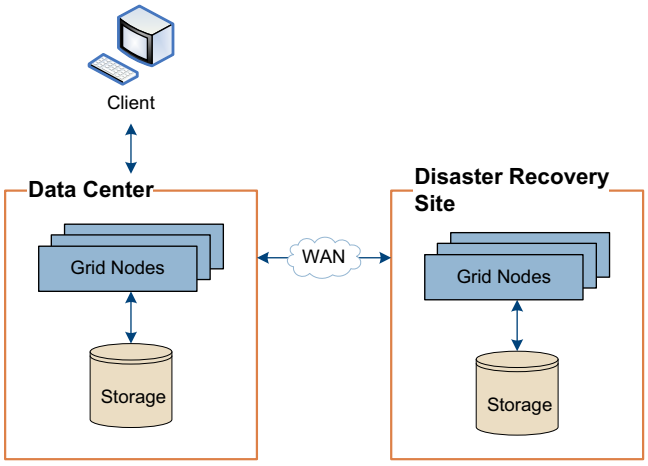

Figure 4: Data Center + Disaster Recovery Deployment

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

Data Center + Disaster Recovery and Satellites

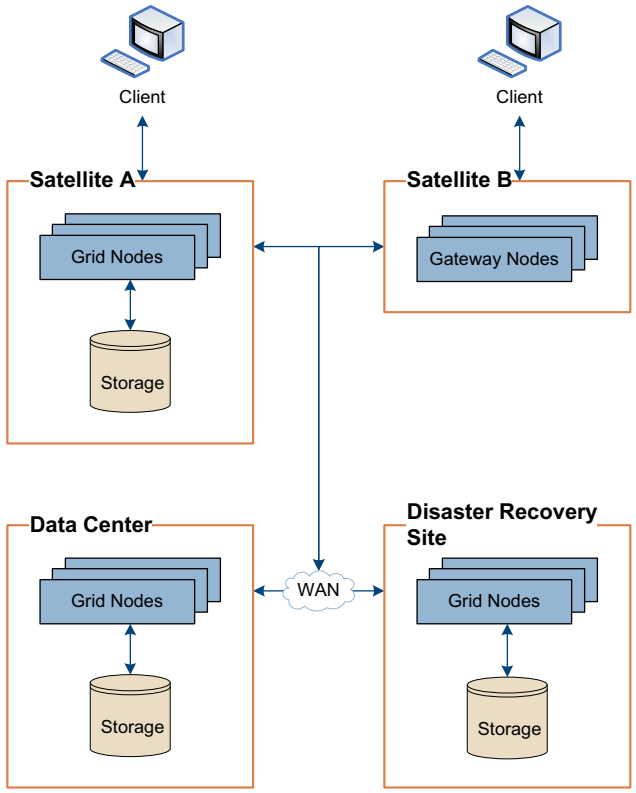

Figure 5: Data Center + Disaster Recovery and Satellites Deployment

A Data Center + Disaster Recovery and Satellites (DC + DR and Satellites) deployment is a type of distributed topology. Individual facilities host a limited amount of local hardware and data is aggregated from all facilities into a centralized DC site. If Satellite sites are designed to support "islanded" operation, data can be ingested and retrieved from the Satellite site while this site is disconnected from the rest of the grid. A DC + DR and Satellites deployment enables fast local access, consolidated operations, and off-site disaster recovery in a network efficient manner.

Federated

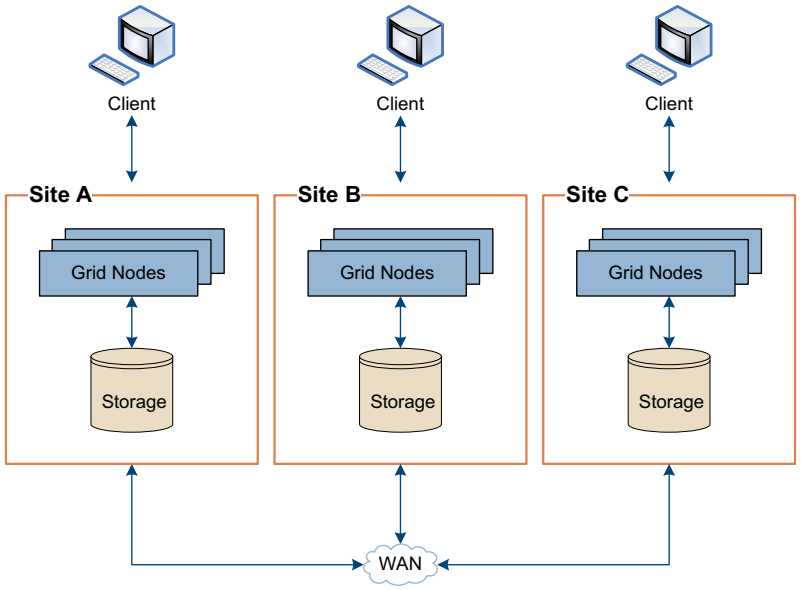

Figure 6: Federated Deployment

A Federated deployment is a "fully distributed" topology, that is, completely decentralized with hardware infrastructure being distributed amongst all facilities. Data sharing and disaster recovery is achieved in a peer-to-peer manner by automatically distributing data to other sites. There is no DC site and one site acts as the Disaster Recovery site for another site.

Grid Nodes

The basic building blocks of a StorageGRID system are grid nodes. A grid node consists of one or more grid services running on a server. A grid service is a software component that performs a specific function. For more information on grid services, see "Grid Services" on page 15.

The basic grid node types are:

- Gateway Node — provides the interface to the grid through which applications, or clients, communicate with the grid
- Storage Node — manages data storage on spinning disks
- Control Node — stores and manages content metadata
- Admin Node — provides grid management services such as grid monitoring, logging, and grid configuration

- Archive Node — for environments where content is also stored on archive media, provides an interface to the middleware that manages the archive media storage device such as a tape library

Figure 7 below displays how grid nodes and services are commonly arranged within the grid.

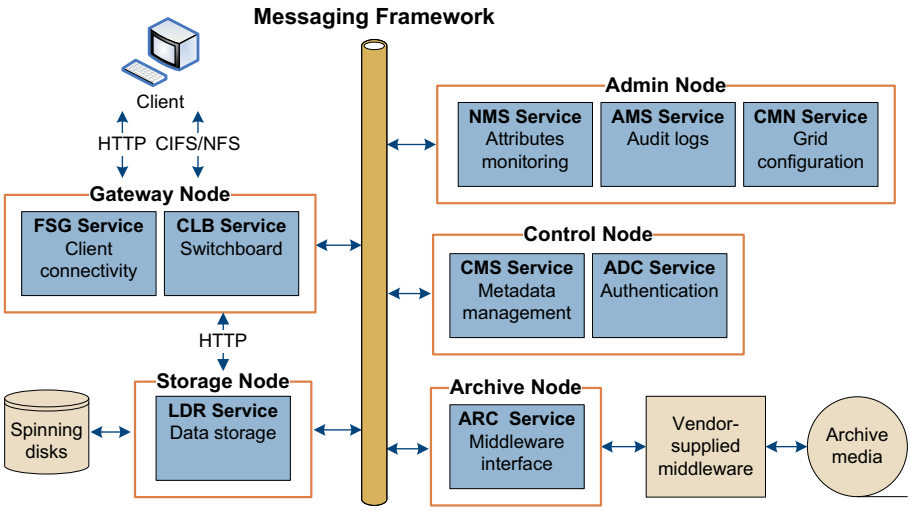

Figure 7: Common StorageGRID System Deployment

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

Table 1: Grid Nodes

Grid Node	Function	Grid Services	Notes
Admin Node	Grid management	NMS CMN AMS SSM	The Admin Node provides tools for grid administration. For purposes of redundancy, grids can have more than one Admin Node. The primary Admin Node hosts the CMN service. (a grid only has one CMN service). The optional High Capacity Admin Cluster (HCAC) allows the NMS service to support a greater number of grid nodes. The HCAC is made up of a reporting Admin Node and a processing Admin Node. The reporting Admin Node hosts the CMN service and is thus the primary Admin Node. For large grids, a separate Audit Node can be included that hosts the AMS service. In this case, the Admin Node does not host the AMS service.

Table 1: Grid Nodes (cont.)

Grid Node	Function	Grid Services	Notes
API Gateway Node	Client connectivity via StorageGRID API applications	CLB SSM	The API Gateway Node provides read-write access to the grid for external StorageGRID API and CDMI applications. API Gateway Nodes do not support NFS and CIFS applications. An API Gateway Node includes only a CLB service; it does not include an FSG service and is therefore not part of a replication group.
Archive Node	Data storage on archive media	ARC SSM	The Archive Node manages storage of data to nearline (neither "online" — instantly available — nor "offline") data storage devices such as tape libraries (via IBM Tivoli® Storage Manager)
Audit Node	Auditing of system events	AMS SSM	The Audit Node logs all audit system events to a text file. Authorization is required to access audit log files. This is an optional node for large grids and is designed to host the AMS service on a separate server.
Control Node	Content management	CMS ADC SSM	The Control Node manages content metadata and content replication. Each grid has at least two Control Nodes for redundancy of metadata storage. Note that best practices suggest that a grid should have at least three Control Nodes.
Gateway Node	Client connectivity	CLB FSG SSM	The Gateway Node provides an interface between client applications and the grid. Each grid must have at least two Gateway Nodes for redundancy. Depending on what client services are required, a Gateway Node may include an FSG service (for CIFS/NFS), a CLB service (for HTTP), or both.
Storage Node	Data storage on spinning disks	LDR SSM	The Storage Node manages data storage on spinning disks. Each grid must have at least two Storage Nodes for redundancy.

Grid Services

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

Table 2: Grid Services

Acronym	Name	Description	Grid Node
ADC 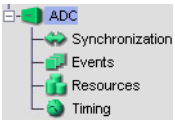	Administrative Domain Controller	Maintains topology information and provides authentication services.	Control Node
AMS 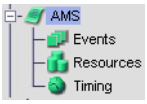	Audit Management System	Keeps logs of grid activity and events.	Admin Node (reporting in HCAC) — or — Audit Node
ARC 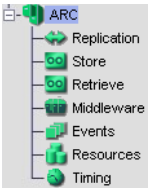	Archive	Communicates with archiving middleware to store and retrieve data to and from archive media.	Archive Node
CLB 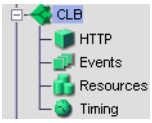	Connection Load Balancer	Acts as switchboard for connecting remote entities to the most efficient LDR. Primary connection point for remote entities using the HTTP protocol.	Gateway Node — or — API Gateway Node
CMN 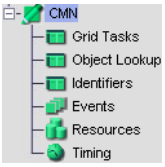	Configuration Management Node	Manages grid-wide configurations: connection profiles, FSG replication groups, grid tasks, and grid options.	Primary Admin Node (reporting in HCAC)

Table 2: Grid Services (cont.)

Acronym	Name	Description	Grid Node
CMS 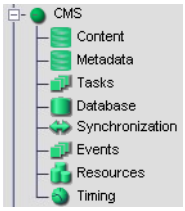	Content Management System	Keeps track of what data is stored on the grid. Stores content metadata and manages content replication based on ILM rules.	Control Node
FSG 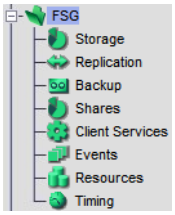	File System Gateway	Allows connections to the grid via a standard file system (NFS or CIFS).	Gateway Node
LDR 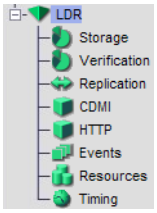	Local Distribution Router	Stores, moves, verifies, and retrieves object data stored on disks.	Storage Node
NMS 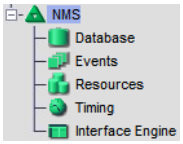	Network Management System	Provides a window into the grid. Used to monitor grid status and to configure the grid. Large grids may be configured with an HCAC to increase the grid's grid services and thus grid node capacity.	Admin Node (both reporting and processing in an HCAC)
SSM 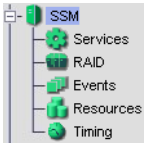	Server Status Monitor	Monitors hardware performance such as key operating system metrics and network metrics. The current operating system installed on the node is listed (SSM ▶ Services).	Present on all nodes

Grid Options

The StorageGRID system may include a number of options that effect how the grid operates. The following table lists the key options that may be deployed with a StorageGRID system. For detailed information regarding options, see the *Administrator Guide*.

Table 3: StorageGRID Options

Option	Description
Audit	Provides user access to the audit logs which contain a complete record of grid activity.
Compression	Compresses objects saved to the grid, reducing file size by roughly 50% for content that is not already in a compressed format.
High Capacity Admin Cluster	Increases the grid's grid services and thus grid node capacity by configuring the grid with a reporting Admin Node and a processing Admin Node.
Deduplication	Deletes unnecessary identical copies of an object from the grid. Deduplication is unavailable on a grid configured with metadata replication. If deduplication is enabled, metadata synchronization must be used. **NOTE** Deduplication and metadata synchronization are deprecated and no longer supported.
Deletion protection	Prevents clients from deleting any content that has been stored to the grid.
Metadata synchronization	Content metadata is synchronized among all CMSs in the grid. After a capacity expansion, metadata is synchronized among all CMSs of the same generation. Required in grids that use deduplication. All other grids use metadata replication. **NOTE** Deduplication and metadata synchronization are deprecated and no longer supported.
Encryption	Enables encrypted storage of all ingested data. Content is encrypted during ingest and objects are stored in an encrypted form so that if a server is compromised no data can be retrieved in any readable form. Note that once enabled, encryption cannot be disabled.
Parallel loading	Enables Gateway Nodes to preload in their cache all the files in a directory upon an initial file request. This enhances performance for systems that store related files in a single directory.

Table 3: StorageGRID Options (cont.)

Option	Description
Secondary preloading	Enables caching on all Gateway Nodes in the Gateway Node replication group as files are ingested or retrieved to speed access via the secondary Gateway Node.
Security partitions	Provides the ability to restrict access to content such that HTTP clients and FGS replication groups only have access to their own data. For more information, see the *StorageGRID API Reference*.

Grid Configuration Information

SAID Package

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

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

Grid Configuration HTML Pages

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

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

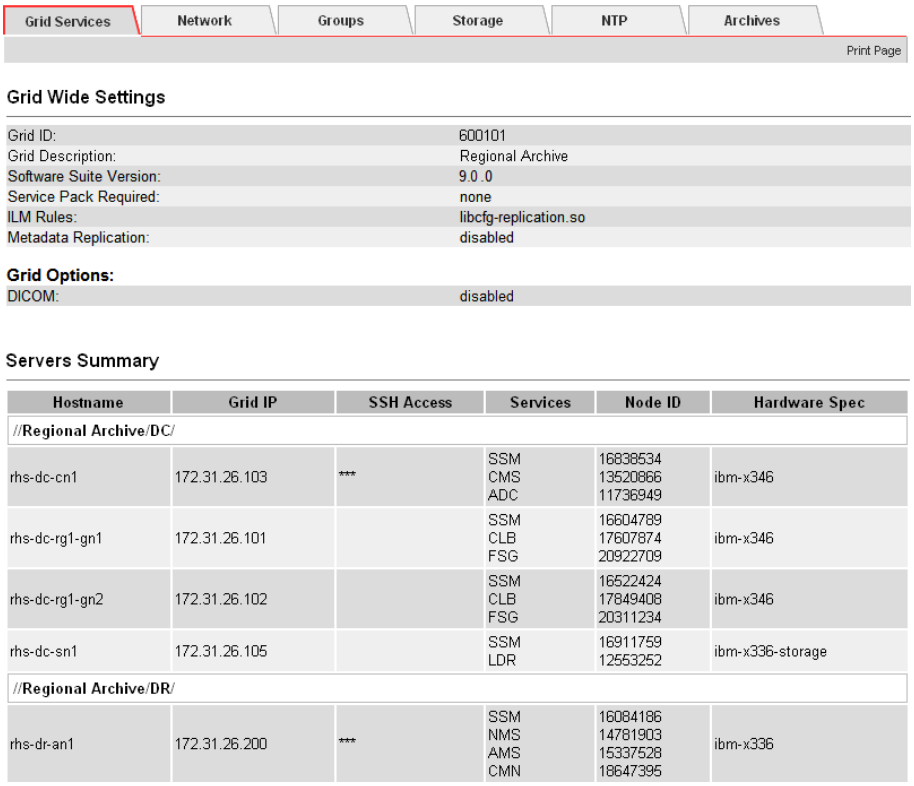

Grid Services | Network | Groups | Storage | NTP | Archives

Print Page

Grid Wide Settings

Grid ID:	600101
Grid Description:	Regional Archive
Software Suite Version:	9.0.0
Service Pack Required:	none
ILM Rules:	libcfg-replication.so
Metadata Replication:	disabled

Grid Options:

DICOM:	disabled

Servers Summary

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

Figure 8: Grid Configuration HTML Pages

Passwords

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

Server Consoles

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

Command Shell Access

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

Log In

- At the server, access a command shell and log in as root using the password listed in the Passwords.txt file.

Log Out

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

Server Manager

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

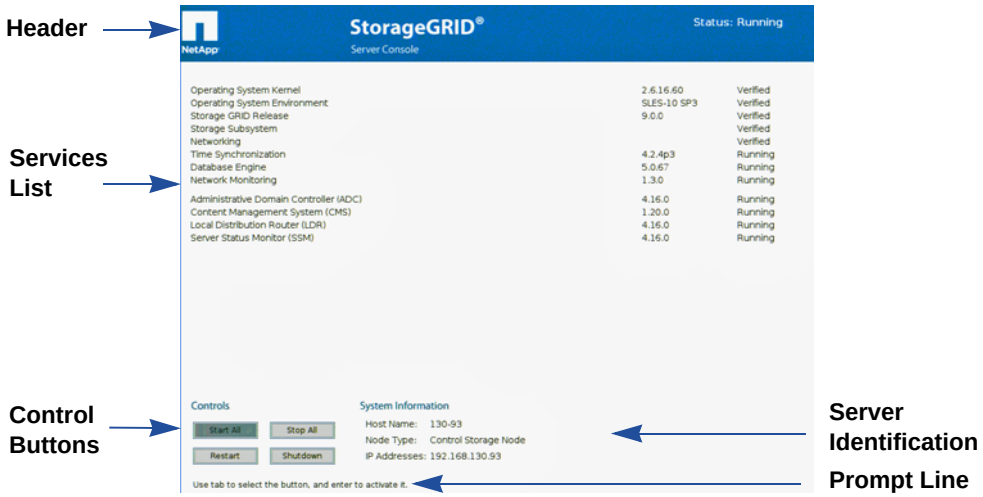

Figure 9: Server Manager

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

For more information on Server Manager, see the *Administrator Guide*.

Related Documentation

Table 4 below lists the documents that are part of the current StorageGRID documentation set.

Table 4: StorageGRID Documentation Set

Document	Notes
Administrator Guide	Contains procedures to customize alarms, set up e-mail notifications, configure grid options, configure the network, set up client shares, run grid tasks, configure ILM policies, and use Server Manager. Also contains background information on grid services.
Audit Message Reference	Describes how to access the AMS log files and interpret common audit messages.
CDMI Reference	Describes StorageGRID specific integration aspects of CDMI used to enable the development of custom applications that integrate with a StorageGRID system.
Documentation Map	Provides guidance as to where to information about the StorageGRID system.
Expansion Guide	Describes how to expand the grid. Includes procedures to add servers, to refresh server hardware, to convert to an HCAC, and to convert a Basic Gateway replication group to a High Availability Gateway replication group.
Grid Designer User Guide	Describes how to use Grid Designer to design and deploy a grid.
Grid Primer	Provides an introduction to the StorageGRID system and describes how to monitor operations using the web based NMS MI.
Installation Guide	Describes the procedures to install a new grid.
Maintenance Guide	Contains procedures to replace failed nodes, fail over Gateway Nodes, verify the integrity of stored data, operate in the absence of the CMN service, recover data, and decommission Storage Nodes.
Release Notes	Provides important details related to late changes and known issues that were not captured in product documentation.
StorageGRID API Reference	Describes the HTTP Application Programming Interface (API) to enable the development of custom applications that integrate with a StorageGRID system.
Troubleshooting Guide	Provides information and troubleshooting procedures to resolve problems that may occur with the StorageGRID system.
Upgrade Guide	Describes how to update to StorageGRID 9.0. Includes a description of new and improved features.

2

NMS Management Interface

How to use the NMS Management Interface: logging in and out, navigating, monitoring alarms, and creating reports

Overview

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

NMS MI Connection Procedure

Access the NMS MI requires a supported web browser with grid access to a known address defined by your system administrator. You must have user name and password to access the system. Each system user is assigned a user name and password when first introduced to the NMS MI.

NOTE Cookies must be enabled

Log In

1. Launch a supported web browser.
2. In the browser's address bar, enter the IP address of the Admin Node. For example: https://199.168.101.108

   The NMS MI Login window opens.

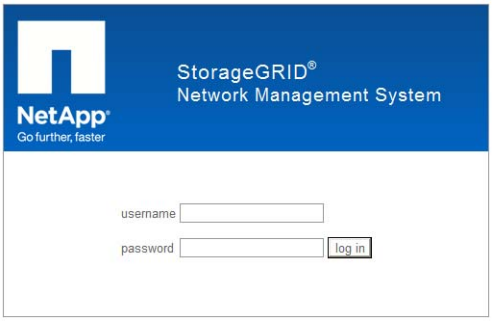

Figure 10: NMS MI Login Window

3. If you are prompted with a security alert, do one of the following:
   - Proceed with this session. The alert will appear again the next time you access this URL.

   — or —

   - View and install the certificate using the browser's installation wizard. The result being that you no longer receive the alert.

   For more information, see "Security Certificate" on page 25.
4. Enter your username and password and click **log in**. Note that both the username and password are case-sensitive.

   The NMS MI appears.

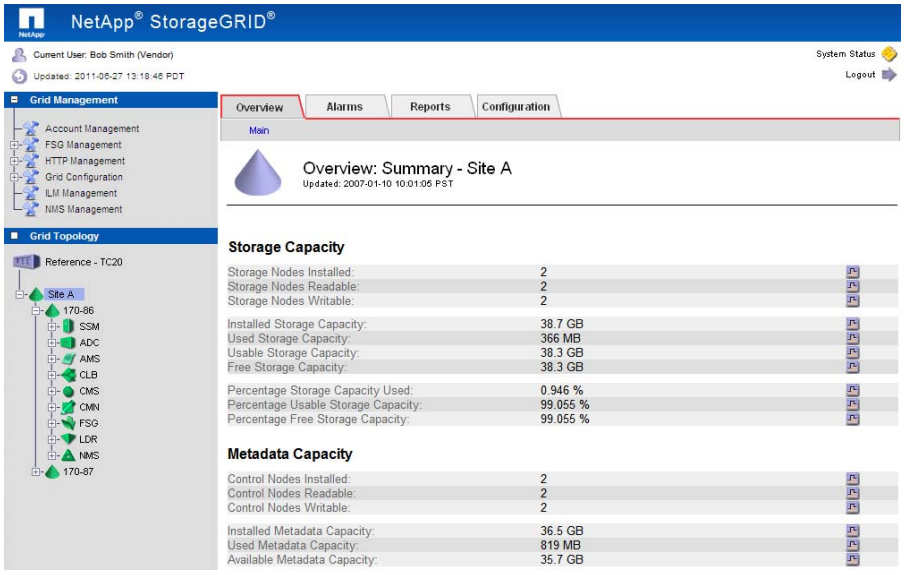

Figure 11: NMS Management Interface

Security Certificate

Depending on your version of Windows and web browser, you may be warned of a problem with the security certificate when you access the NMS MI URL or pop-up window.

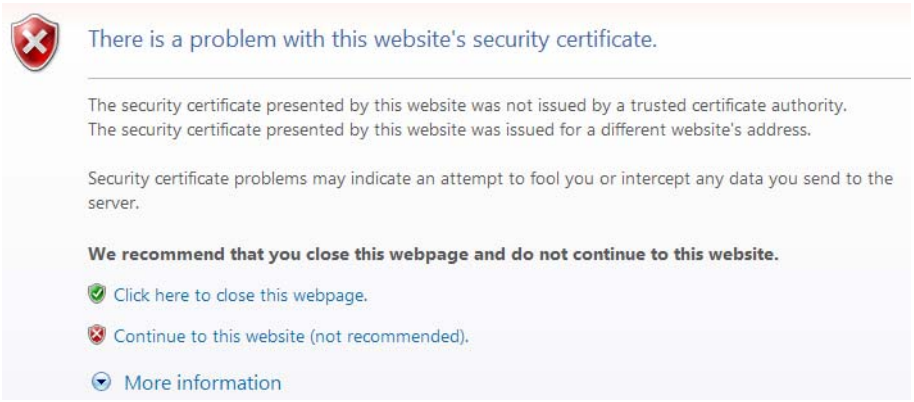

Figure 12: Example of a Security Alert Window

If this appears, you can either:

- Proceed with this session. The alert will appear again the next time you access this URL.
- Install the certificate. Follow the instructions of your browser.

The NMS MI uses a self-signed certificate. For information on importing this certificate into a browser see the browser's documentation. Note that the self-signed certificate used by the NMS MI is based on the grid's IP address. The expected URL to the interface is this IP address and not a domain name. In cases where the domain name is used, browsers may not be able to match the self-signed certificate to the identity of the NMS server. For more information see the browser's documentation.

Enable Pop-ups

To view charts in a pop-up window, view the grid's About window, or configure your account, pops-ups must be enabled for your browser. For more information on enabling pop-up windows, see your browser's documentation.

Log Out

When you have completed your NMS MI session, you must log out of the NMS MI. This keep the system secure.

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 use other applications.

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

NMS Management Interface

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

The main elements of the NMS MI are:

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

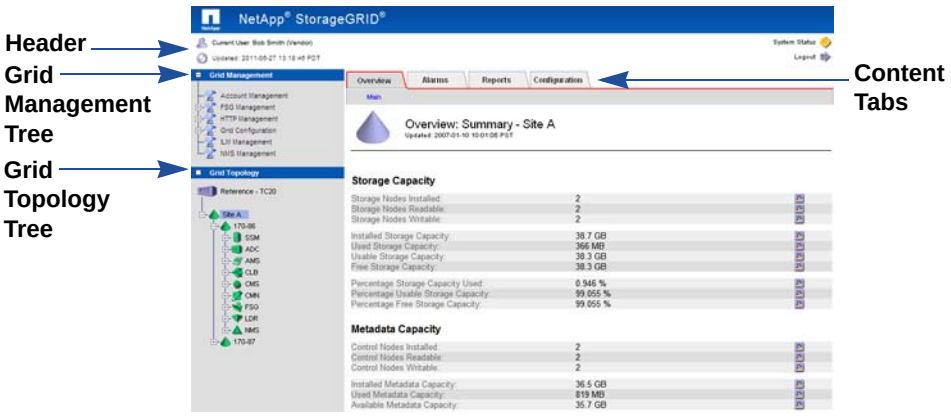

Figure 13: NMS Management Interface Elements

Header

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

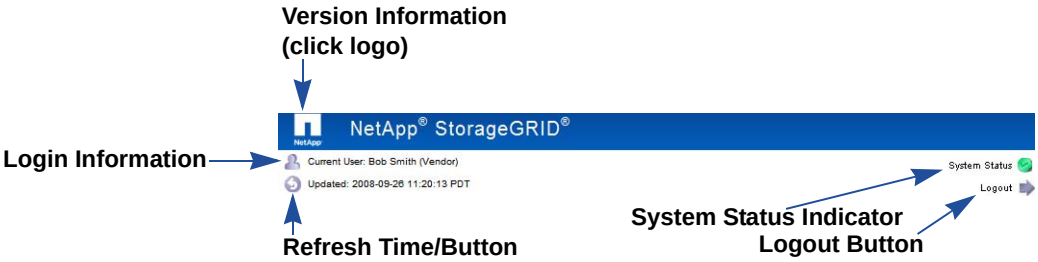

Figure 14: NMS Management Interface — Header

Refreshing the Display

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

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

- Click **Refresh Page** .
- Click the timestamp.
- Click the refresh button for the browser.

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

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

If you have configured the NMS GUI Timeout period to 0 and then minimize your browser for an extended period of time (for example, greater than one week), when you expand the browser, the NMS MI

may not correctly display attribute values and the current time. Refresh the display using the procedure given above.

Displaying NMS MI Version Information

- Click the product name in the header.

  The About window displays the interface version number, the software build number, and copyright information. The copyright information includes a link for easy access to the NetApp web site.

Figure 15: About Window — Version Information

NMS Interface Engine Status

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

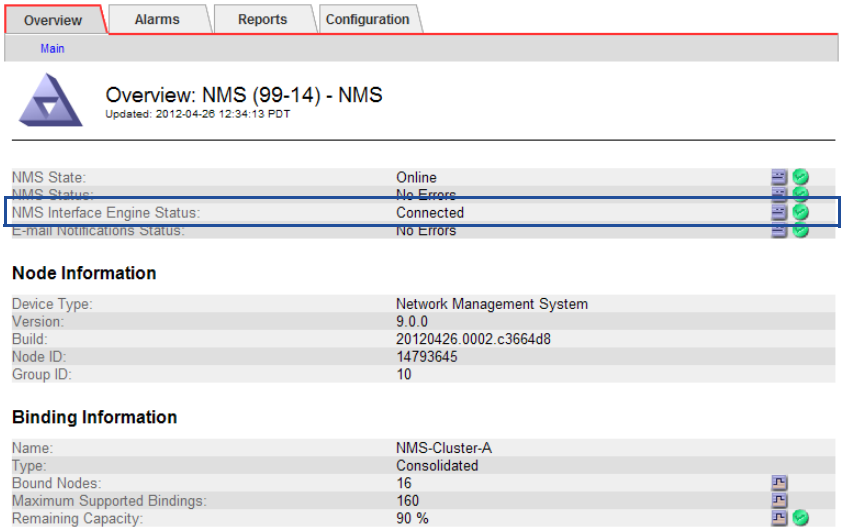

Figure 16: NMS Interface Engine Status

This interface engine status information is useful if your grid is configured with multiple Admin Nodes and thus multiple NMS services.

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

Grid Management Tree

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

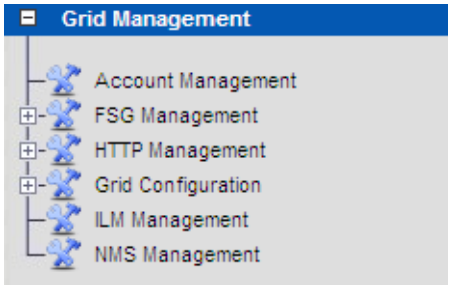

Figure 17: Grid Management Tree

Table 5: Grid Management Tree

Branch	Used to Configure
Account Management	User Accounts. For more information on accounts and user groups, see "User Accounts" on page 34.
FSG Management	FSG settings such as content protection options, cache space, and backups. Configuration is restricted to user accounts that have Maintenance permissions such as the Admin and Vendor accounts. For information on FSG management, see the *Administrator Guide*.
HTTP Management	Grid options such as client permissions, certificates, security partitions, client IP addresses, and HTTP metadata. Configuration is restricted to user accounts that have Maintenance permissions such as the Admin and Vendor accounts. For information on HTTP management, see the *Administrator Guide*.
Grid Configuration	Grid options such as Audit levels, link cost groups, and storage grades. Configuration is restricted to user accounts that have Grid Management permissions such as the Vendor account. For information on grid configuration, see the *Administrator Guide*.
ILM Management	ILM (Information Lifecycle Management) policies. Configuration is restricted to user accounts that have Grid Management permissions such as the Vendor account. For information on ILM, see the *Administrator Guide*.

Table 5: Grid Management Tree

Branch	Used to Configure
NMS Management	NMS overview, custom alarms, alarm notifications, Admin Node name, and GUI timeout period. Configuration is restricted to user accounts that have Maintenance permissions such as the Admin and Vendor accounts. For information on alarm customization and notifications, see the *Administrator Guide*.

Grid Topology Tree

The Grid Topology tree provides quick access to grid elements. These elements are:

- Grid
- Location
- Grid nodes
- Services
- Components

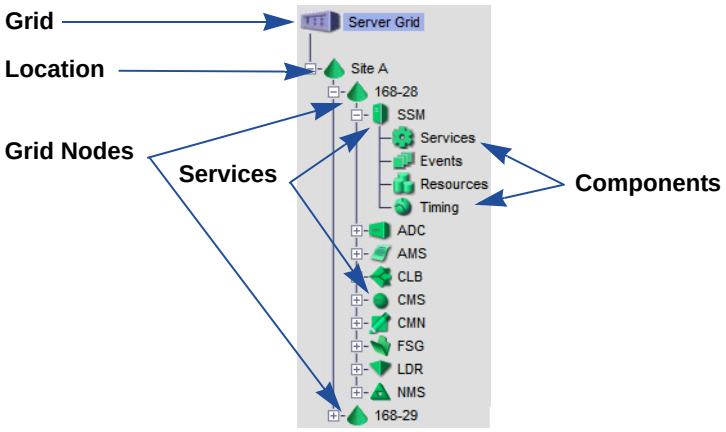

Figure 18: Grid Topology Tree

The grid includes separate locations. Locations can reflect different physical regions such as cities, buildings, floors in a building, or any other grouping. Each location can be expanded to reveal one or more grid nodes, a grid node being a server hosting a collection of one or more grid services. A grid service consists of software components that deliver a particular capability.

Expand or Collapse the Grid Topology Tree

- In the Grid Topology tree, click ⊞ and ⊟.

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

View Grid Elements

- In the Grid Topology tree, click an element's name.

Content Tabs

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

Figure 19: Content Tabs

- **Overview** — The Overview tab is used to monitor grid attributes. Each attribute represents a property, for example the number of managed objects, free storage space, backup size, or service state. These attributes are used to monitor normal grid operation and to detect and troubleshoot abnormal conditions. While there are hundreds of attributes, most of them are used for troubleshooting and only a small number must be monitored on a regular basis to ensure smooth operation. For examples on how to work with attributes, see "Monitoring Operations" on page 103.
- **Alarms** — The Alarms tab is used to view and acknowledge alarms. For more information, see "Alarms and State Indicators" on page 38.
- **Reports** — The Reports tab is used to create charts and text reports. For more information, see "Reports" on page 49.
- **Configuration** — The Configuration tab is used to change configuration settings at the location, grid node, grid service, or component level. Configuration is restricted to user accounts that have Maintenance permissions such as the Admin and Vendor accounts.

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

Figure 20: Page Selector

View an Attribute Description

The NMS MI contains a description of each attribute.

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

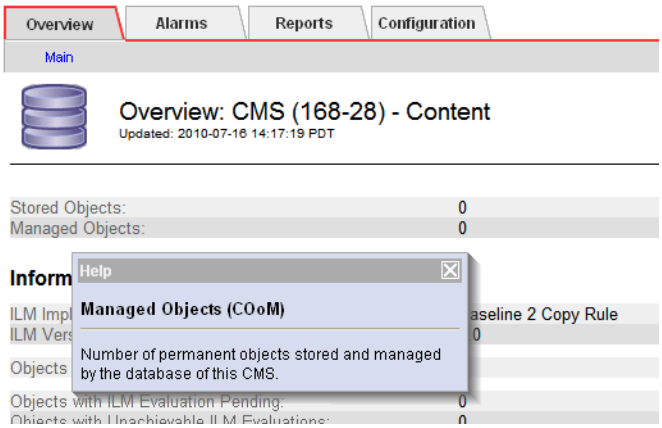

Figure 21: Attribute Description

Interpreting Attribute Values

Attributes are reported by each grid service on a best effort basis. Attribute updates may be lost under some circumstances, such as the crash of a service or the failure and rebuild of a grid node. For example, if a CMS service crashes after it has registered the ingest of a number of objects but before it has committed the updated attribute value to disk, up to two minutes of attribute updates could be lost. The reporting of attributes is also subject to propagation delays within the grid. Updated values for most attributes, except for state attributes, are sent to the NMS MI at fixed intervals. Therefore, it may take a few moments before an update is visible in the NMS MI, and two attributes that change more or less simultaneously may be reported at slightly different times.

Attribute values reported in the NMS MI should therefore be understood as being approximately correct, particularly for "count" attributes such as object count or HTTP connection attempts. Use the

NMS MI for troubleshooting, trend analysis, and to get a general picture of the state of the grid.

Valid Characters

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

Units of Measure

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

When scaled, the value of bytes displayed by the NMS MI on overview pages and in charts uses the "natural" measure of powers of 10. For example 3 MB = 3×10^6 = 3,000,000 bytes. This is not the same as powers of 2 normally used for computing, where 3 MiB = 3×2^{20} = 3,145,728 bytes. To see unscaled values at full accuracy, generate a text report.

NOTE In the NMS MI, scaled byte values are presented in kb, MB or GB (powers of 10), not KiB, MiB, or GiB (powers of 2).

Apply Changes Button

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

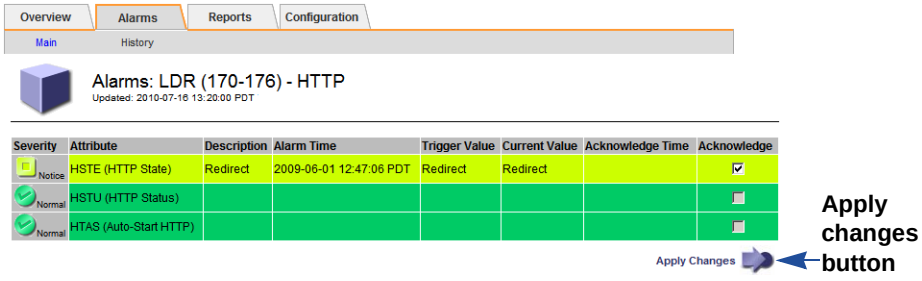

Figure 22: Apply Changes Button

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

Grid Node Path Link

Clicking an underlined grid node path takes you to the grid node's Overview page. For example, in Figure 23 below, clicking **Site A/170-41/FSG** opens the FSG ▶ Overview page for the primary Gateway Node.

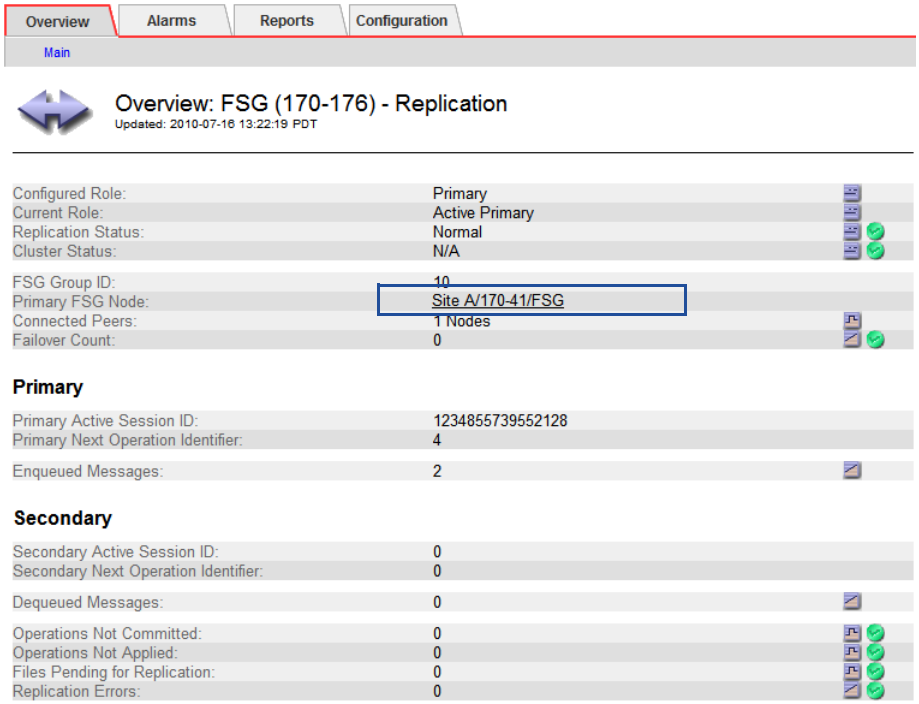

Figure 23: Grid Node Path Links

User Accounts

The StorageGRID system has two built-in user accounts:

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

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

Permissions

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

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

- Grid management
- Maintenance
- Alarm acknowledgement
- Accounts

Table 6 below describes the allowable tasks for each set of permissions.

Table 6: User Groups Permissions

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

Figure 24 below shows the built-in user accounts and group accounts. Depending on the Accounts permission on your account, you may not see all the accounts. See Figure 25 for a comparison.

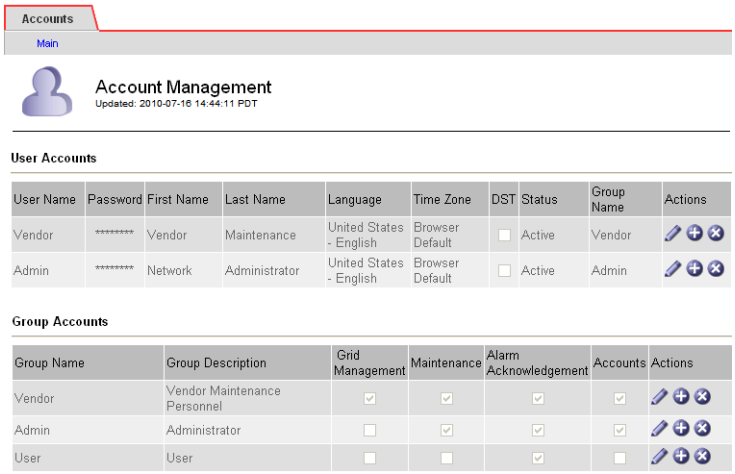

Figure 24: Account Management for Vendor Account

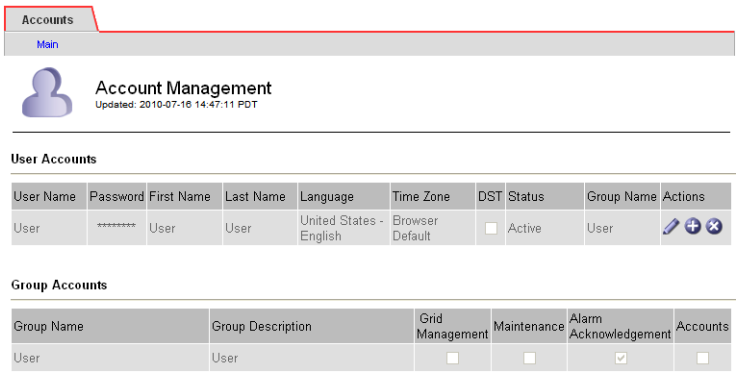

Figure 25: Account Management for Account Without Accounts Permissions

Configuring Your Account

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

NOTE Creating accounts and modifying accounts of other users is restricted to accounts that have "Accounts" permission such as the Admin and Vendor accounts. For more information on creating and modifying accounts, see the *Administrator Guide*.

1. If you intend to change the password, make sure pop-ups are enabled. For more information, see "Enable Pop-ups" on page 25.
2. Go to **Grid Management ▶ Account Management**.
3. Click **Edit** and update entries as needed:

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

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

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

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

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

  Depending on your permissions, Group Name may be read-only.

4. To change your password:
   a. Double-click Password to select the complete field.
   b. Type a new password. Your password must contain between 8 and 32 characters and is case-sensitive.
   c. Press **<Tab>**. A confirmation pop-up window appears.

Figure 26: Password Confirmation Pop-up Window

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

5. To change your name, edit the First Name and Last Name boxes.
6. To change the time zone, select a new time zone from the list.
7. Click **Apply Changes**.

Alarms and State Indicators

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

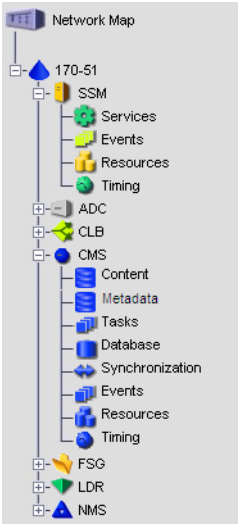

Figure 27: Grid Topology Tree with State and Alarm Colors

If there are no alarms and all services are connected, the icon appears in the normal (green) state. If there is an alarm, the color of the icon reflects the most severe alarm currently active on that branch of the

tree. Locations display the highest alarm level of the grid nodes on that branch. Grid nodes display the color of the most severe state or alarm among their hosted services. Each individual service reflects the highest alarm severity of its components.

Service State Indicators

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

Table 7: Service States

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

Alarm Indicators

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

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

Alarms are generated at the attribute level. There are five alarm severity levels displayed in the NMS MI. Each alarm level has an associated color and icon (see Table 8). Note that a severity level of Normal does not trigger an alarm.

Table 8: Alarm Severity and Indicators

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

Propagation

Alarm Indicators

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

For example, in Figure 28 below, the SSM service has at least two alarms: the Events component has at least one alarm with a severity of Notice and the Resources component has at least one alarm with a severity of Minor Alert. Minor Alert is the more severe of the two alarms and therefore it propagates up the Grid Topology tree so that the SSM service takes on the Minor Alert alarm color, light orange.

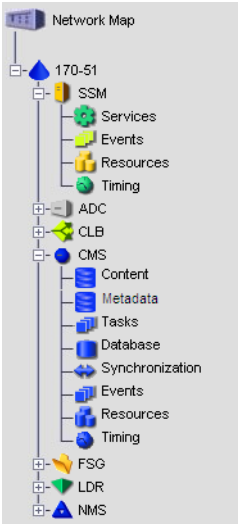

Figure 28: Propagation of Alarm and State Indicators

State Indicators

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

For example, in Figure 28 above, the CMS service has a state of "Unknown." This state overrides any alarms that may have been raised on any of it components. "Unknown" is a more severe state than any alarm and therefore propagates down the Grid Topology tree so that all CMS components take on the "Unknown" state color, blue. The grid node also displays the "Unknown" state color of blue as this is the most critical state of any of its services. The state condition of a node's services propagates up the Grid Topology tree in the same manner that alarms do.

E-mail Notifications

E-mail notifications are automatically sent to designated personnel to alert recipients that an alarm has been triggered or that a service state has changed. Managing e-mail notifications is restricted to user

accounts with Maintenance permissions such as Admin or Vendor. For more information on e-mail notifications, see the *Administrator Guide*.

E-mail Notification Status

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

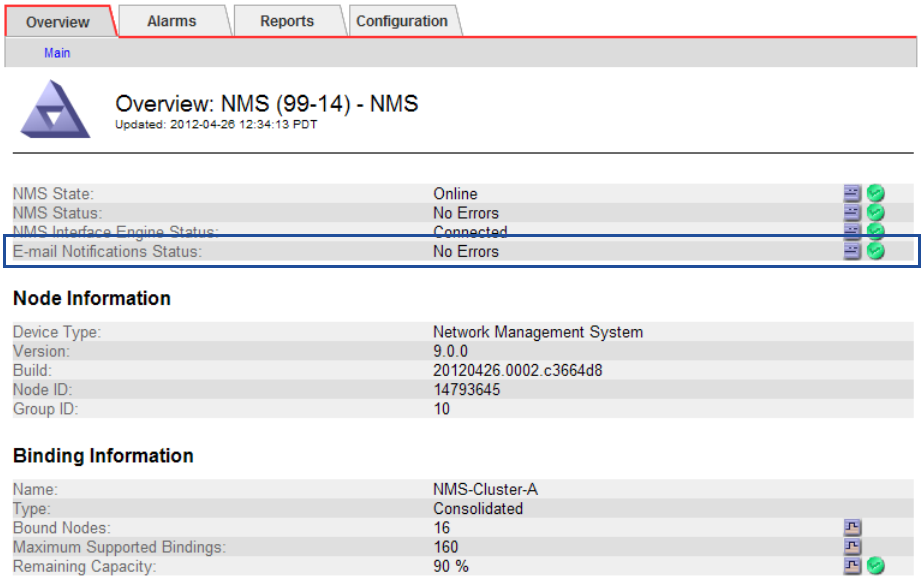

Figure 29: NMS Notification Status

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

Alarm Customization

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

The Configuration ▶ Alarms page of each service or component is used to view configured Default alarms and Global Custom alarms and to create Custom alarms for a service. Access is restricted to user

accounts that have Maintenance permissions such as the Admin and Vendor accounts.

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

For more for information on custom alarms, see the *Administrator Guide*.

Reviewing Alarms

Table 9 below summarizes how to review information on the current status of alarms. This information is located on the:

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

Table 9: Reviewing Alarms

To:	Do this:
Get a list of all current alarms in the grid	• Click **System Status**  in the header to display alarms. Alarms are sorted by severity. 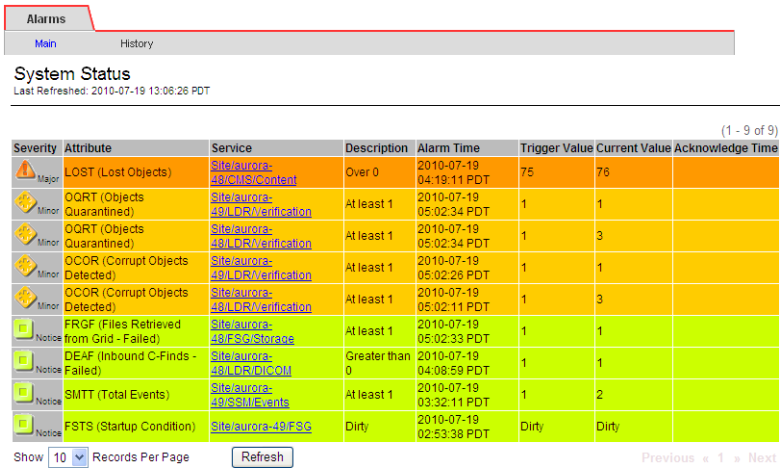 **NOTE** The first thing to look at in the NMS MI is the System Status indicator. It immediately tells you the most serious status (state or alarm) of the grid.
Get a list of all alarms triggered over a period of time	1. Click **System Status** in the header. 2. Click the **History** page. 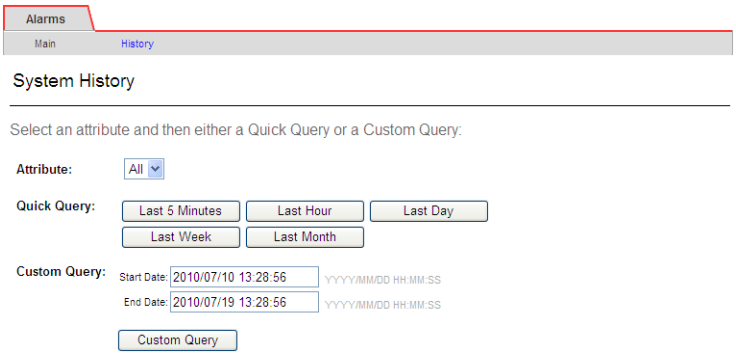 3. Do one of the following: • Click one of the time periods. • Enter a custom range and click **Custom Query**.

Table 9: Reviewing Alarms (cont.)

To:	Do this:
Find out more about an alarm	**1.** In the System Status table, under Service click the path to go to the Alarms tab for the selected alarm (or use the Grid Topology tree). The Alarms tab displays the current status of the attributes for the selected service or component. The colors and icons reflect the severity levels of the alarms. Once an alarm is resolved, the alarm returns to the green "Normal" severity level. 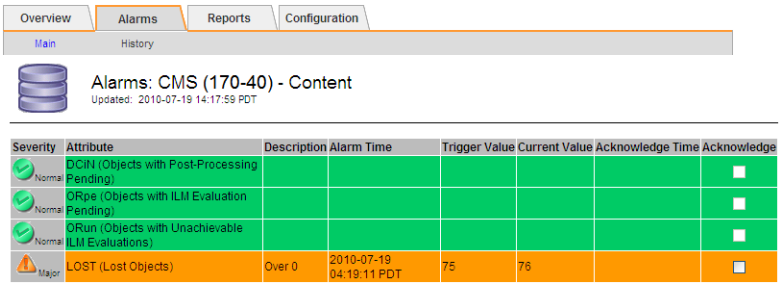 See Table 10 on page 47 for a description of the fields. **2.** Click the alarm to display a description of the attribute. 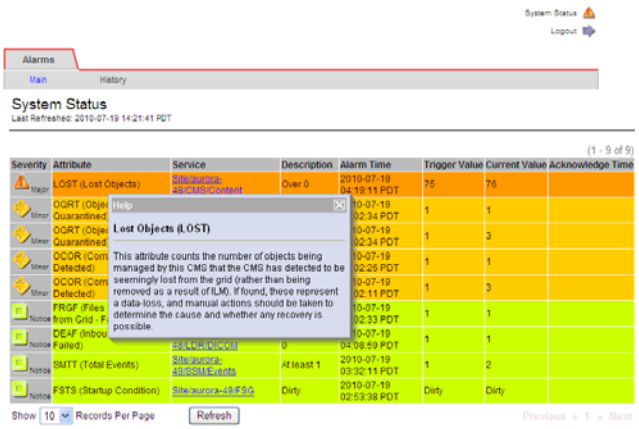 **3.** Look up the four-character attribute in the reference table in the "Alarms Reference" chapter of the *Troubleshooting Guide*. **4.** Click **<*component*> ▶ Overview ▶ Main**. **5.** Locate the alarm, and if it can be charted, click the chart button to view a trend of the attribute over the last hour. Adjust the time period as required. For more information, see "Displaying Charts" on page 52.

Table 9: Reviewing Alarms (cont.)

To:	Do this:
Find out how often alarms have been triggered for a particular attribute	1. Go to the service or component that has the attribute. 2. Click the **Alarms** tab and then the **History** page. 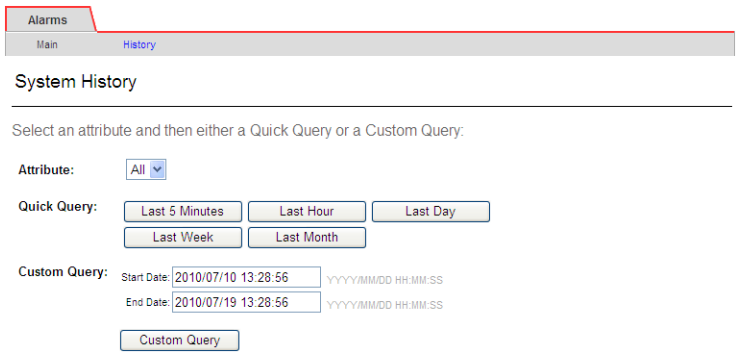 3. Select the attribute from the list. 4. Do one of the following: • Click one of the time periods. • Enter a custom range and click **Custom Query**. 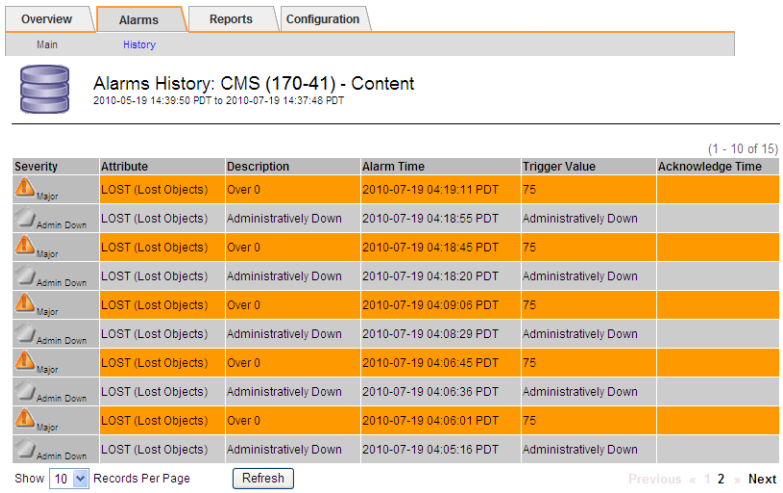 The alarms are listed in reverse chronological order. See Table 10 below for a description of the fields. 5. To return to the alarms history request form, click **History**.
Determine whether an alarm has been disabled globally	If a Default alarm is disabled globally, an asterisk appears beside the alarm on the Configuration ▶ Alarms tab. 

Table 10: Alarms Table Fields

Field	Description
Severity	Color icon indicating the alarm severity level.
Attribute	Code that identifies the attribute and issue being monitored. See the "Alarms Reference" chapter of the *Troubleshooting Guide* for an alphabetical listing of the attributes and reference information on each alarm.
Description	Brief details about the cause of the alarm.
Alarm Time	The date and time in your local time zone at which the trigger value was reported. Blank if the alarm has been acknowledged.
Trigger Value	Value that triggered the alarm.
Current Value	Value of the attribute as last reported by the service or component.
Acknowledge Time	Date and time the alarm was acknowledged.
Acknowledge	Selecting **Acknowledge** acknowledges the alarm. See the procedure below.

Acknowledging Alarms

Depending on the situation, you may choose to acknowledge alarms while you are trying to resolve the underlying issue. Acknowledging alarms is restricted to user accounts that have Alarm Acknowledgement permissions such as the Admin and Vendor accounts.

An acknowledged alarm continues to display as an alarm at the component level on the System Status page. However, once an alarm has been acknowledged, it no longer propagates up the Grid Topology tree. The Grid Topology tree is displayed as "Normal" (green) or the color of the next most severe unacknowledged alarm or more severe service state (see Figure 30 below).

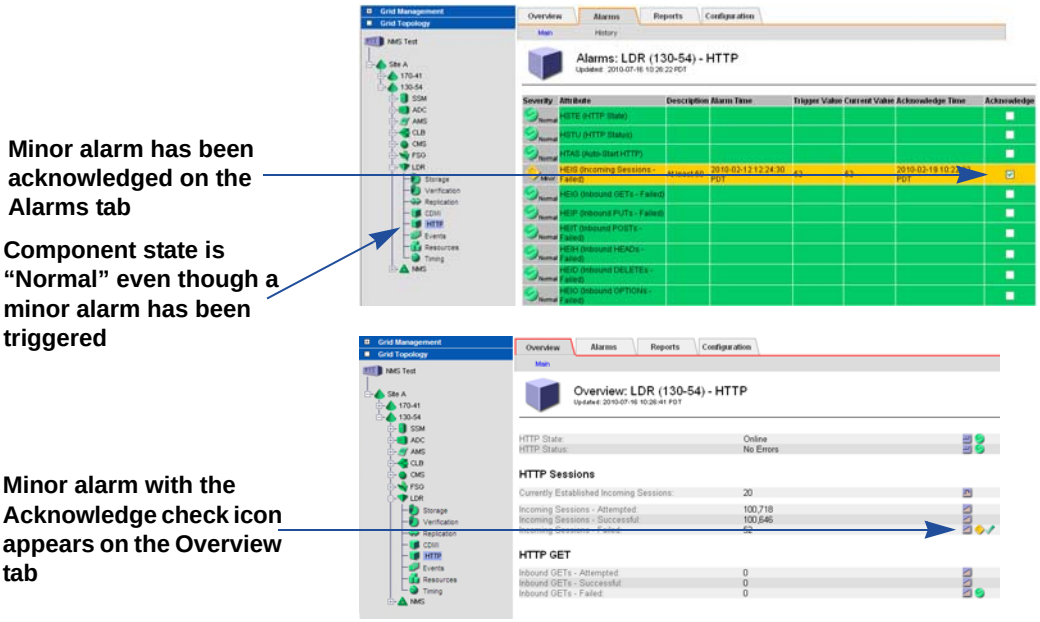

Figure 30: Acknowledging Alarms

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

When a service is in a state of "Administratively Down," all attribute alarms for that service — including acknowledged alarms — are removed. The current attribute values at the time the service restarts will be used to determine if any new alarms are triggered.

Acknowledge an Alarm

1. Go to ***<service or component>*** ▶ **Alarms** ▶ **Main**.
2. Select **Acknowledge** next to the alarm.

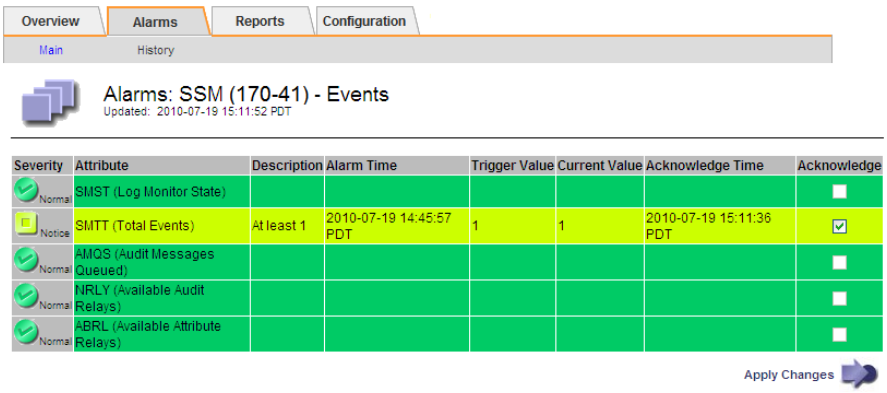

Figure 31: Sample Acknowledged Alarm

3. Click **Apply Changes**.

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

Unacknowledge an Alarm

1. Go to ***<service or component>*** ▶ **Alarms** ▶ **Main**.
2. Clear **Acknowledge** next to the alarm.
3. Click **Apply Changes**.

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

Reports

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

Charts

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

Chart Types

There are three types of charts:

- line graph

- area graph
- state graph

Line Graph

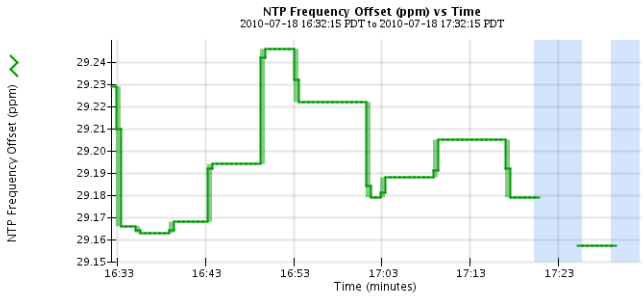

Figure 32: Line Graph

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

Area Graph

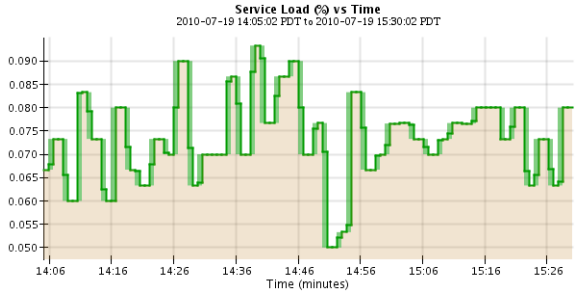

Figure 33: Area Graph

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

State Graph

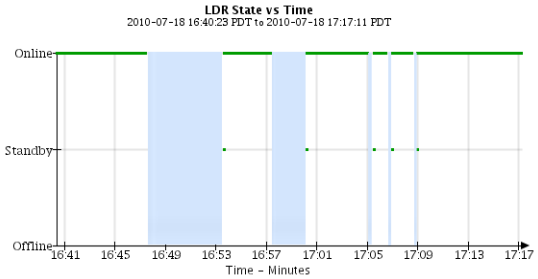

Figure 34: State Graph

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

Interpreting Chart Colors

The chart colors have a specific meaning. Table 11 below describes how to interpret the various colors and line types.

Table 11: Chart Colors and Shading

Sample	Meaning
	Reported attribute values are plotted using dark green lines.
	Light green shading around dark green lines indicates that the actual values in that time range vary and have been “binned” for faster plotting. The dark line represents the weighted average. The range in light green indicates the maximum and minimum values within the bin. Light brown shading is used for area graphs to indicate volumetric data.
	Blank areas (no data plotted) indicate that the attribute values were unavailable. The background may be blue, gray, or a mixture of gray and blue depending on the state of the service reporting the attribute.
	Light blue shading indicates that some or all of the attribute values at that time were indeterminate; the attribute was not reporting because the service was in an unknown state.
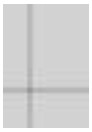	Gray shading indicates that some or all of the attribute values at that time were not known because the service reporting the attributes was administratively down.

Table 11: Chart Colors and Shading (cont.)

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

Displaying Charts

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

You can also manually create reports from the Report tab.

NOTE It is not possible to create charts for all attributes; for example, text attributes such as Node ID, version number, and build number.

Create an Immediate Chart Report

1. Go to ***<component or service>*** ▶ **Overview** ▶ **Main**.

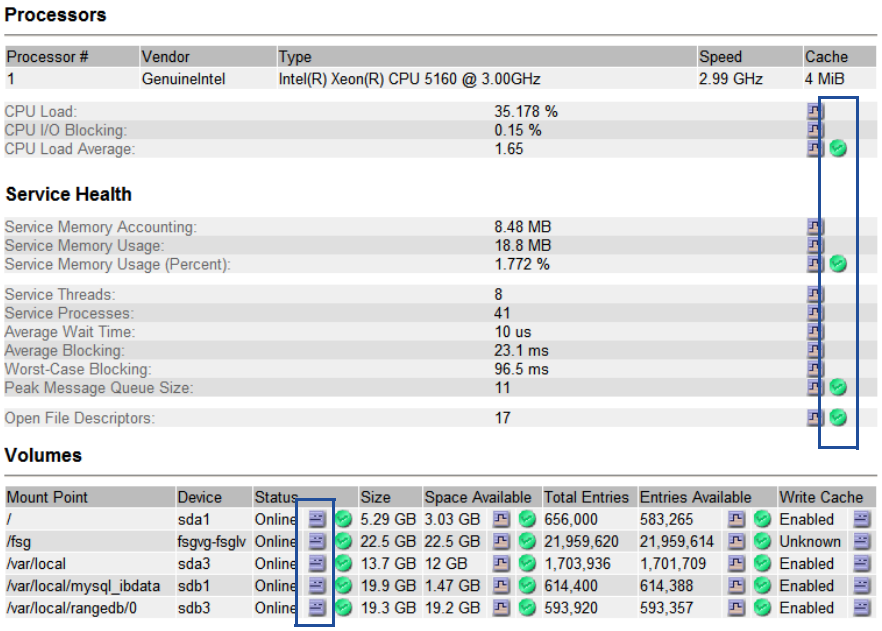

Figure 35: Immediate Chart Buttons

2. Click **Chart** next to the attribute.

   The display automatically changes to the Reports ▶ Charts page. The chart displays the attribute's data over the past hour. To view other time ranges, follow the procedure described below.

Manually Create a Chart Report

1. Go to ***<component or service>*** ▶ **Reports** ▶ **Charts**.
2. Select **Attribute** ▶ ***<attribute>***.
3. To force the Y-axis to start at zero, clear **Vertical Scaling**.
4. To show values at full precision, select **Raw Values**. To round values to a maximum of three decimal places (for example, for attributes reported as percentages), clear **Raw Data**.
5. Select **Quick Query** ▶ ***<time_period>***.

   The chart appears after a few moments. Allow several minutes for long time ranges.
6. To display a chart for a custom time period:
   a. Select **Quick Query** ▶ **Custom Query**.
   b. Enter a **Start Date** and **End Date**.

      Use the format YYYY/MM/DD HH:MM:SS in local time. Leading zeros are required to match the format. For example, 2010/4/8 7:30:00 fails validation; the correct format is 2010/04/08 07:30:00.
   c. Click **Update**.

Displaying Charts in a New Window

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

- Click  to display the current view in new window.

  Click  to close the chart windows.

Text Reports

A text report displays a textual representation of attribute data values that have been processed by the NMS service. For attribute data that is expected to be continuously changing, this attribute data is sampled by the NMS service (at the source) at regular intervals. For attribute

data that changes infrequently (for example, data based on events such as state or status changes) an attribute value is sent to the NMS service when the value changes.

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

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

Raw Text Report

A raw text report displays the following information:

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

Text Results for Services: Load - System Logging

2010-07-18 15:58:39 PDT To 2010-07-19 15:58:39 PDT

Time Received	Sample Time	Value
2010-07-19 15:58:09	2010-07-19 15:58:09	0.016 %
2010-07-19 15:56:06	2010-07-19 15:56:06	0.024 %
2010-07-19 15:54:02	2010-07-19 15:54:02	0.033 %
2010-07-19 15:52:00	2010-07-19 15:52:00	0.016 %
2010-07-19 15:49:57	2010-07-19 15:49:57	0.008 %
2010-07-19 15:47:54	2010-07-19 15:47:54	0.024 %
2010-07-19 15:45:50	2010-07-19 15:45:50	0.016 %
2010-07-19 15:43:47	2010-07-19 15:43:47	0.024 %
2010-07-19 15:41:43	2010-07-19 15:41:43	0.032 %
2010-07-19 15:39:40	2010-07-19 15:39:40	0.024 %
2010-07-19 15:37:37	2010-07-19 15:37:37	0.008 %
2010-07-19 15:35:34	2010-07-19 15:35:34	0.016 %
2010-07-19 15:33:31	2010-07-19 15:33:31	0.024 %
2010-07-19 15:31:27	2010-07-19 15:31:27	0.032 %
2010-07-19 15:29:24	2010-07-19 15:29:24	0.032 %
2010-07-19 15:27:21	2010-07-19 15:27:21	0.049 %
2010-07-19 15:25:18	2010-07-19 15:25:18	0.024 %
2010-07-19 15:21:12	2010-07-19 15:21:12	0.016 %
2010-07-19 15:19:09	2010-07-19 15:19:09	0.008 %
2010-07-19 15:17:07	2010-07-19 15:17:07	0.016 %

Figure 36: Raw Text Report

Aggregate Text Report

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

Each entry displays the following information:

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

Text Results for Attribute Send to Relay Rate

2010-07-11 16:02:46 PDT To 2010-07-19 16:02:46 PDT

Aggregate Time	Average Value	Minimum Value	Maximum Value
2010-07-19 15:59:52	0.271072196 Messages/s	0.266649743 Messages/s	0.274983464 Messages/s
2010-07-19 15:53:52	0.275585378 Messages/s	0.266562352 Messages/s	0.283302736 Messages/s
2010-07-19 15:49:52	0.279315709 Messages/s	0.233318712 Messages/s	0.333313579 Messages/s
2010-07-19 15:43:52	0.28181323 Messages/s	0.241651024 Messages/s	0.374976601 Messages/s
2010-07-19 15:39:52	0.284233141 Messages/s	0.249982001 Messages/s	0.324971987 Messages/s
2010-07-19 15:33:52	0.325752083 Messages/s	0.266641993 Messages/s	0.358306197 Messages/s
2010-07-19 15:29:52	0.278531507 Messages/s	0.274984766 Messages/s	0.283320999 Messages/s
2010-07-19 15:23:52	0.281437642 Messages/s	0.274981961 Messages/s	0.291577735 Messages/s
2010-07-19 15:17:52	0.261563307 Messages/s	0.258318006 Messages/s	0.266655787 Messages/s
2010-07-19 15:13:52	0.265159147 Messages/s	0.258318557 Messages/s	0.26663986 Messages/s

Figure 37: Aggregate Text Report

Create a Text Report

1. Go to ***<component or service>*** ▶ **Reports** ▶ **Text**.
2. Select **Attribute** ▶ ***<attribute>***.
3. Select **Results per Page** ▶ ***<results_per_page>***.
4. To round values to a maximum of three decimal places (for example, for attributes reported as percentages), clear **Raw Data**.
5. Select **Quick Query** ▶ ***<time_period>***.

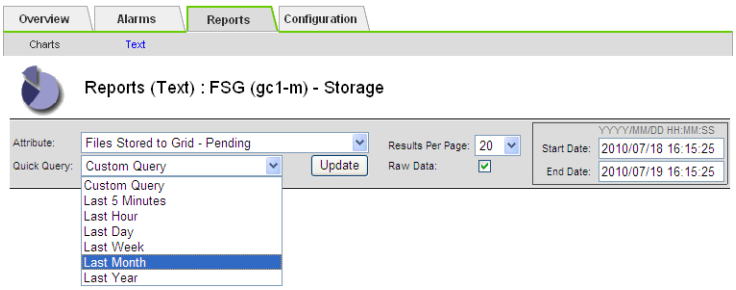

Figure 38: Report Request Form — Quick Text Reports

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

6. To display a report for a custom time period:

   a. Select **Quick Query ▶ Custom Query**.

   b. Enter a **Start Date** and **End Date**.

      Use the format YYYY/MM/DD HH:MM:SS in local time. Leading zeros are required to match the format. For example, 2010/4/8 7:30:00 fails validation. The correct format is: 2010/04/08 07:30:00.

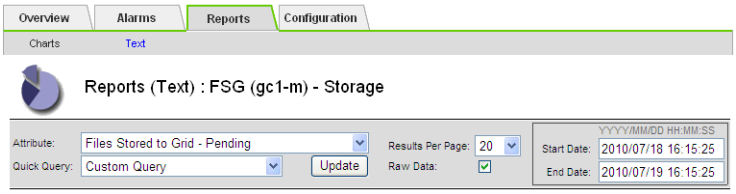

Figure 39: Report Request Form — Custom Query Report

   c. Click **Update**.

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

Export a Text Report

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

1. Create a text report. For more information, see "Create a Text Report" on page 55.

2. Click **Export** .

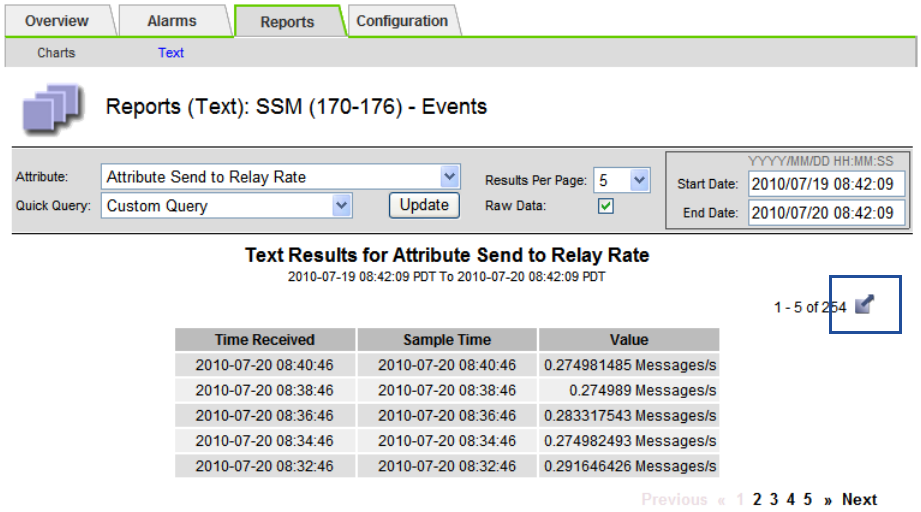

Figure 40: Export Text Report

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

Grid ID: 400019
OID: 2.16.124.113590.2.1.400019.1.1.1.1.16996732.200
Node Path: Site/170-176/SSM/Events
Attribute: Attribute Send to Relay Rate (ABSR)
Query Start Date: 2010-07-19 08:42:09 PDT
Query End Date: 2010-07-20 08:42:09 PDT
Time Received,Time Received (Epoch),Sample Time,Sample Time (Epoch),Value,Type
2010-07-20 08:40:46,1279640446559000,2010-07-20 08:40:46,1279640446537209,0.274981485 Messages/s,U
2010-07-20 08:38:46,1279640326561000,2010-07-20 08:38:46,1279640326529124,0.274989 Messages/s,U
2010-07-20 08:36:46,1279640206556000,2010-07-20 08:36:46,1279640206524330,0.283317543 Messages/s,U
2010-07-20 08:34:46,1279640086540000,2010-07-20 08:34:46,1279640086517645,0.274982493 Messages/s,U
2010-07-20 08:32:46,1279639966543000,2010-07-20 08:32:46,1279639966510022,0.291646426 Messages/s,U
2010-07-20 08:30:46,1279639846561000,2010-07-20 08:30:46,1279639846501672,0.308315369 Messages/s,U
2010-07-20 08:28:46,1279639726527000,2010-07-20 08:28:46,1279639726494673,0.291657509 Messages/s,U
2010-07-20 08:26:46,1279639606526000,2010-07-20 08:26:46,1279639606490890,0.266627739 Messages/s,U
2010-07-20 08:24:46,1279639486495000,2010-07-20 08:24:46,1279639486473368,0.258318523 Messages/s,U
2010-07-20 08:22:46,1279639366480000,2010-07-20 08:22:46,1279639366466497,0.274985902 Messages/s,U
2010-07-20 08:20:46,1279639246469000,2010-07-20 08:20:46,1279639246460346,0.283253871 Messages/s,U
2010-07-20 08:18:46,1279639126469000,2010-07-20 08:18:46,1279639126426669,0.274982804 Messages/s,U
2010-07-20 08:16:46,1279639006437000,2010-07-20 08:16:46,1279639006419168,0.283315503 Messages/s,U

Figure 41: Export Text Report Window

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

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

Printing Reports

1. Create a chart or text report. For more information, see "Reports" on page 49.

2. Right-click the chart or report and select **Print** (for text reports) or **Print Picture** (for chart reports).

   The Print dialog box opens.

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

   The report is printed.

3

Data Flow

The flow of data as objects are ingested, replicated, retrieved, and purged

Key Concepts

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

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

Only a very brief overview is provided here. For more information, see the *Administrator Guide*.

Client Shares and FSG Managed File System

In order to communicate with the FSG service, client applications map a network drive to the FSG file share, for example, /fsg/myDirectory. The FSG service supports CIFS and NFS file share protocols. For more information on client integration, see the *Administrator Guide*.

Applications interface with the grid via the FSG's "managed file system". There is one managed file system per FSG and it is mounted at /fsg. The managed file system contains file pointers that point to the file's location in the grid. The FSG managed file system shown in Figure 42 contains a file pointer for the object /fsg/myDirectory/image.jpg. The FSG uses the object's unique identifier to communicate with the CMS service (via the LDR service) which tracks where the object is stored on the Storage Nodes (LDR service) and on the Archive Nodes (ARC service).

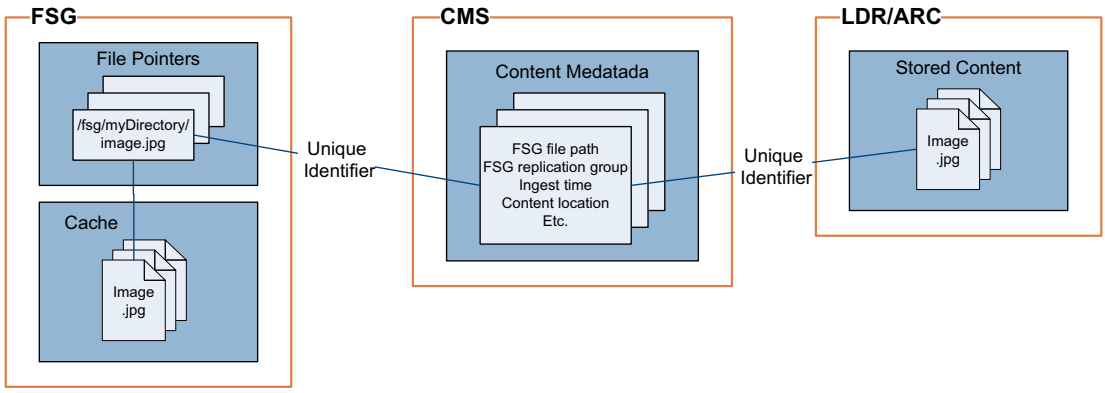

Figure 42: FSG File Pointers

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

For a detailed discussion on FSG cache, see the *Administrator Guide*.

FSG Replication Groups

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

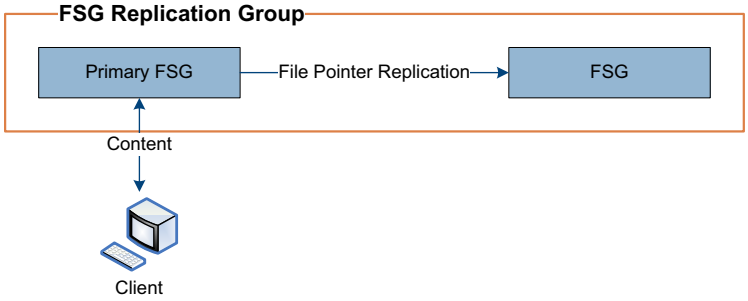

Figure 43: FSG Replication Group

The primary FSG replicates its file pointers (for instance Figure 44 shows the pointer for /fsg/myDirectory/image.jpg) to the secondary FSG.

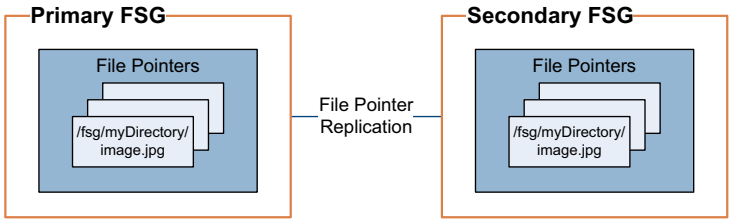

Figure 44: FSG File Pointer Replication

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

The following FSG replication groups are available:

- Basic Gateway replication group
- High Availability Gateway replication group

For a detailed review of each type of FSG replication group, see the *Administrator Guide*.

Topology Queries

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

HTTP Protocol Commands

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

CMS Operation

The CMS service performs two important functions within the grid:

- stores and manages object metadata.

- manages content replication to ensure the Information Lifecycle Management (ILM) policy for the grid is satisfied.

Object metadata is information related to or describing an object stored in the grid, for instance ingest file path, FSG replication group ID, file modification time, storage location, and so on. Metadata is stored in databases maintained by the CMSs. To ensure redundancy, the grid stores multiple copies of the object metadata in different databases.

In addition to managing metadata, the CMS services manage content replication to ensure that the grid's ILM policy is satisfied. The owner CMS carries out the ILM's instructions for storing objects in the grid over time (where to store the objects, how many copies to store, and for how long).

ILM Policy

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

- Geography — the permanent location of an object
- Storage grade — the type of storage used to store an object
- Replication — the number of copies to make of an object
- Retention — the changes over time to the location, storage grade, and number of copies of an object.

Location, storage grade, and the number of copies can vary over time. In the ILM policy example shown in Figure 45, a file is ingested into the grid via an FSG. At ingest, two copies of the object are stored in the Data Center site on SAS disks and one copy is stored in the Disaster Recovery site on SATA disks. One year after ingest, one copy at the DC site is deleted and one copy is created on archive media at the DR site.

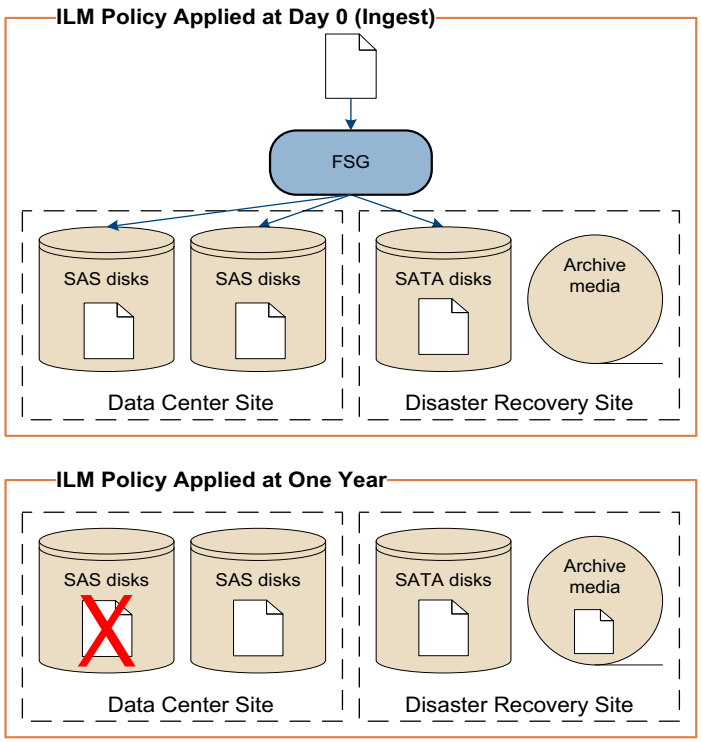

Figure 45: Information Lifecycle Management Policy Example

Owner CMS and Reproduction of Metadata

The first CMS service to get the object metadata when an object is ingested into the grid is referred to as the owner CMS. The owner CMS then copies the metadata to other CMS services.

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

NOTE Metadata synchronization is deprecated and no longer supported.

For a detailed review of the two kinds of CMS and their operation, see the *Administrator Guide*.

Object Content Handle

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

Object Lifecycle

Figure 46 below follows an object as it is ingested, retrieved, becomes inactive, and is finally deleted:

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

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

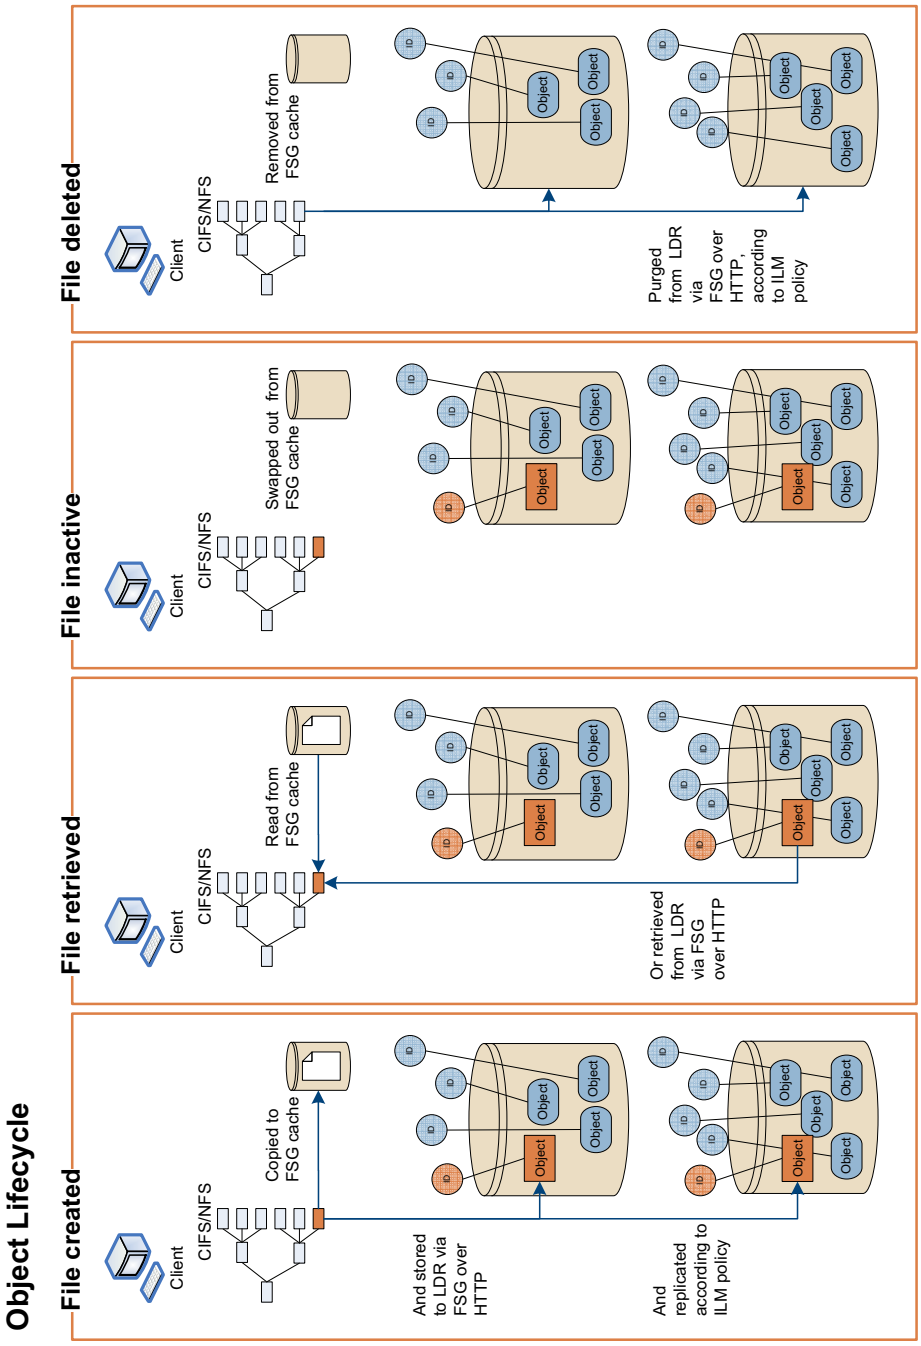

Figure 46: Object Lifecycle

Ingest

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

Data Flow

See Figure 47 below for a simplified step-by-step description of what happens when objects are ingested into the grid.

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

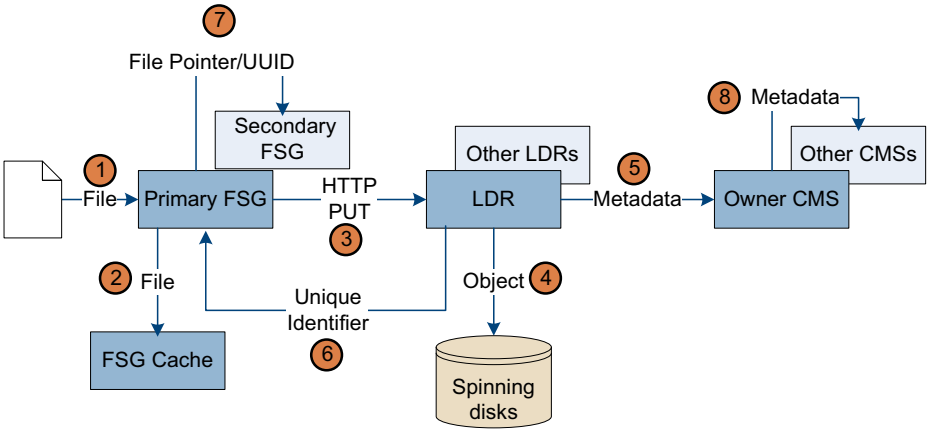

Figure 47: Ingest Data Flow

1. The client application saves a file to the primary FSG managed file system via its mapped network share, for example /fsg/myDirectory. This triggers an FSG "create operation" and as a result the FSG creates a file pointer for this object. The primary FSG replicates the file pointer to the other FSGs in its replication group.

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

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

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

NOTE Metadata synchronization is deprecated and no longer supported.

Related Attributes

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

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

Table 12: Object Ingest Attributes

Component	Attribute Changes
Primary FSG ▶ Storage 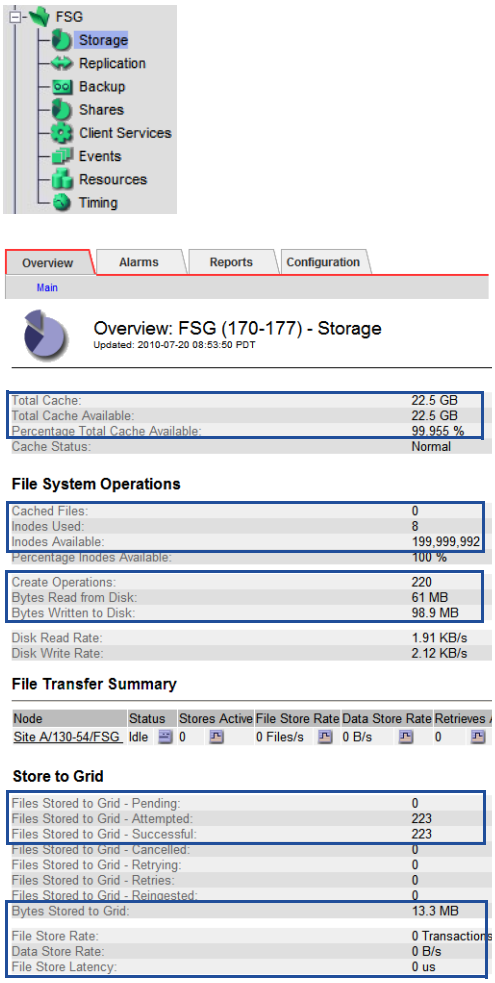	Create Operations (FCRO): The number of new files or folders created in the file system increases by 1. 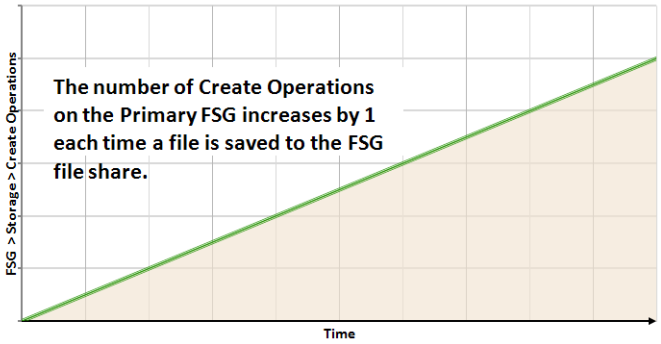 Inodes Used (FSIU): The number of inodes (files and directories) used on the FSG file system increases by 1. Inodes Used and Create Operations can differ if objects are deleted and their inodes are later re-used for new objects. Inodes Available (FSIA): The number of inodes (files and directories) available on the FSG file system decreases by 1. Files Stored to Grid – Successful (FSGC): The number of files stored persistently on the grid increases by 1. The value increases after the FSG receives an acknowledgment from the LDR that the file has been successfully stored into the grid (which is at the same time as it receives the unique identifier from the LDR). 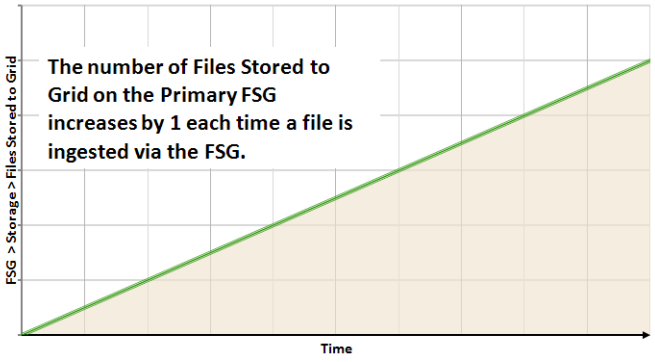 Files Stored to Grid – Attempted (FSGA): The number of initiated file transfers to the grid increases by 1.

Table 12: Object Ingest Attributes (cont.)

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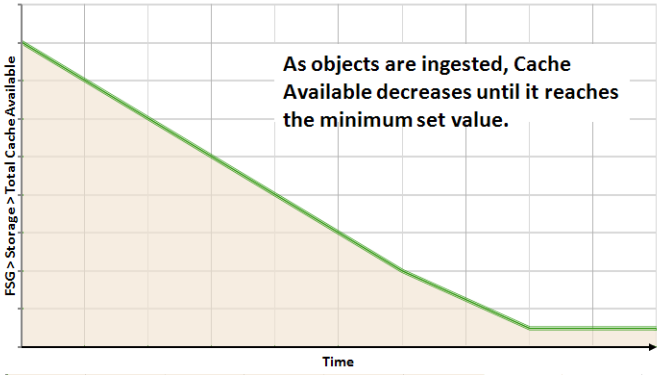

Table 12: Object Ingest Attributes (cont.)

Component	Attribute Changes
	Cached Files (FSCF): The number of cached files on the primary FSG increases by 1 unless the cache is full. The number of cached files will generally stabilize once the Total Cache Available reaches its minimum, but may fluctuate as files are added and removed from the cache depending on file size. 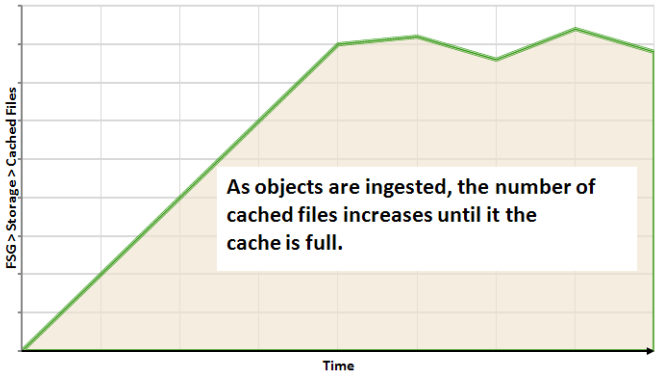
FSG ▶Shares 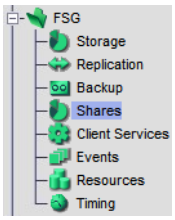 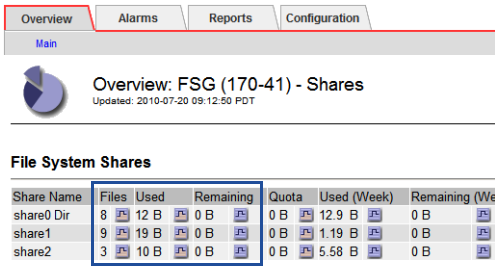 	Files (FSSF): The number of files ingested on this share will increase by 1 the next time a backup takes place. 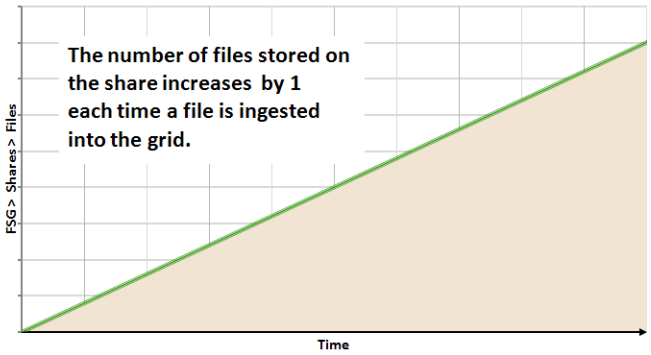 Used (FSSB) and Used (Week) (FSSR): The total size of all files referenced by this share will increase by an amount equivalent to the file size the next time a backup takes place. Remaining (FSSA) and Remaining (Week) (FSSL): If a quota has been configured for the share, the total amount of space remaining on the share will decrease by the size of the file.

Table 12: Object Ingest Attributes (cont.)

Component	Attribute Changes
LDR ▶ HTTP 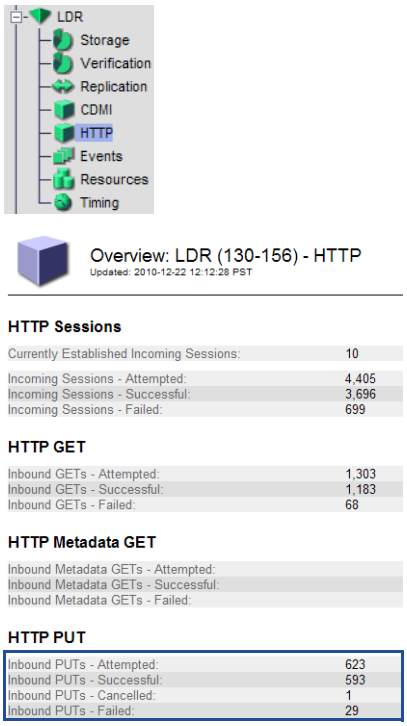 	Inbound PUTs – Successful (HIPC): The total number of HTTP PUT (“content store”) requests that have been completed successfully by the LDR increases by 1. 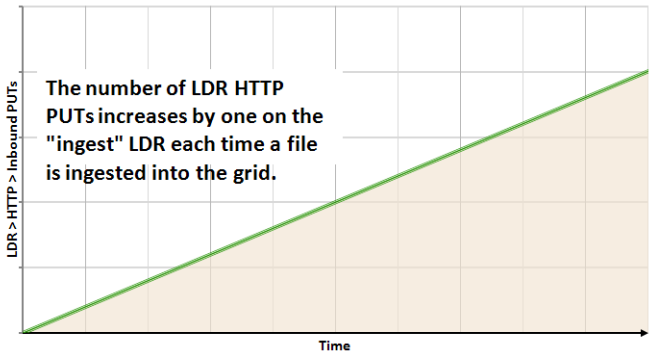 Inbound PUTs – Attempted (HAIP): The total number of HTTP PUT (content store) requests that have been received by the LDR also increases by 1. **NOTE** Transactions that occur via CDMI do not result in an update to these attributes.

Table 12: Object Ingest Attributes (cont.)

Component	Attribute Changes
LDR ▶ Storage 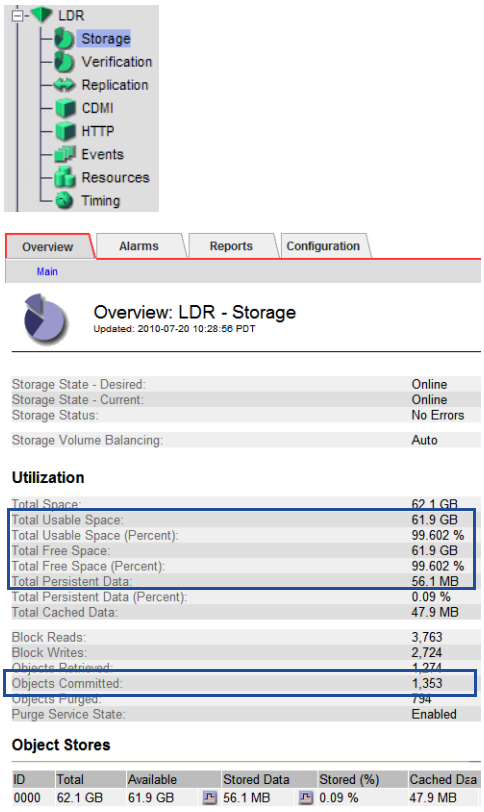 	Total Usable Space (STAS): The total amount of object storage space that is currently available to be used to store objects decreases by an amount roughly equivalent to the file size. 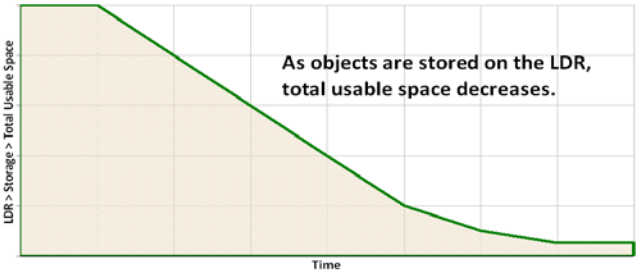 Total Usable Space (Percent) (SAVP): The total amount of object storage space (displayed as a percentage) that is currently available to be used to store objects decreases by an amount roughly equivalent to the file size. Total Free Space (SUSA): The total amount of all free space — available to be used or not — on all object stores decreases by an amount roughly equivalent to the file size. Total Free Space (Percent) (SUSP): The total amount (displayed as a percentage) of all free space — available to be used or not — on all object stores decreases by an amount roughly equivalent to the file size. Total Persistent Data (SPSD): The estimate of the size of the persistently stored data increases by an amount roughly equivalent to the file size. Objects Committed (OCOM): The number of object store operations that have been processed by the LDR increases by 1. 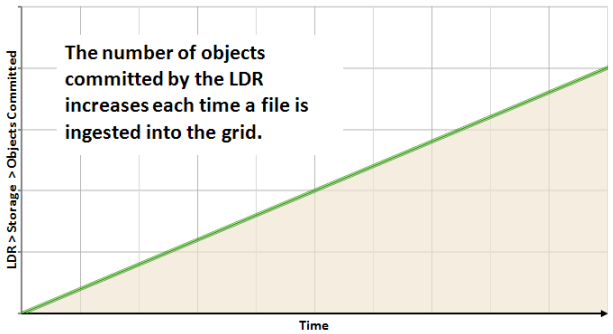

Table 12: Object Ingest Attributes (cont.)

Component	Attribute Changes
CMS ▶ Metadata 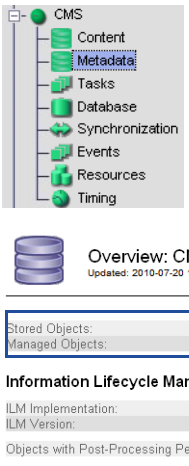 **NOTE** For grids that use synchronized metadata, the attributes Stored Objects and Managed Objects are shown under the Content component. For grids that use metadata replication, these two attributes appear under the Metadata component. Metadata synchronization is deprecated and no longer supported.	Stored Objects (COoT): The number of objects in the CMS metadata database increases by 1 on all CMSs in the same CMS replication group in a grid that uses metadata replication, and on all read-write CMSs in a grid that uses metadata synchronization. 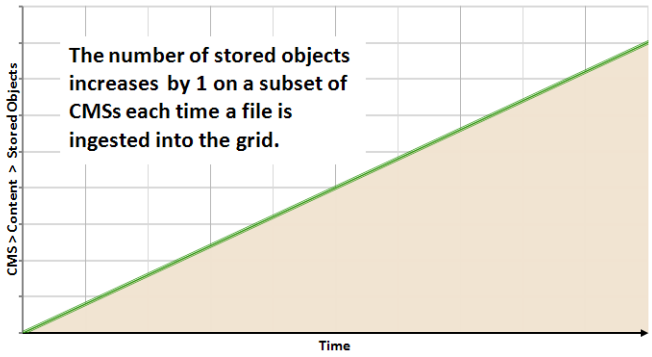 Managed Objects (COoM): The number of objects owned by the owner CMS increases by 1. 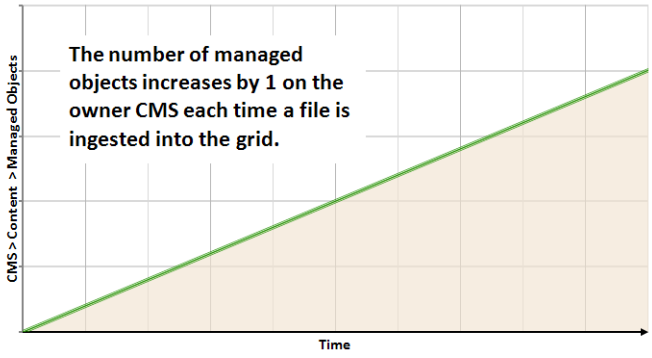

Table 12: Object Ingest Attributes (cont.)

Component	Attribute Changes
CMS ▶ Database 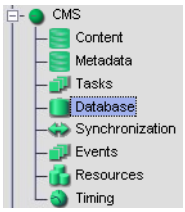 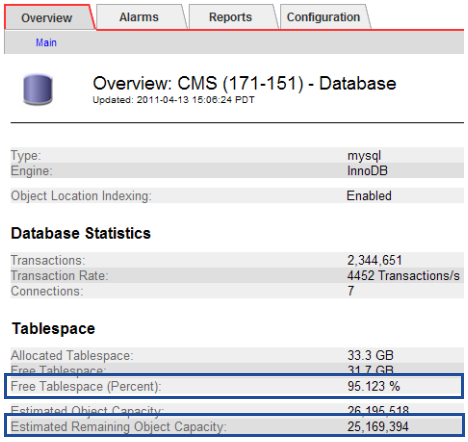 	Free Tablespace Percent (DBSP): The amount of space remaining in the metadata database decreases. MySQL manages free tablespace in chunks. Ingesting (or deleting) a single object may not change the reported free table space. Ingesting (or deleting) many objects will eventually change the free table space. 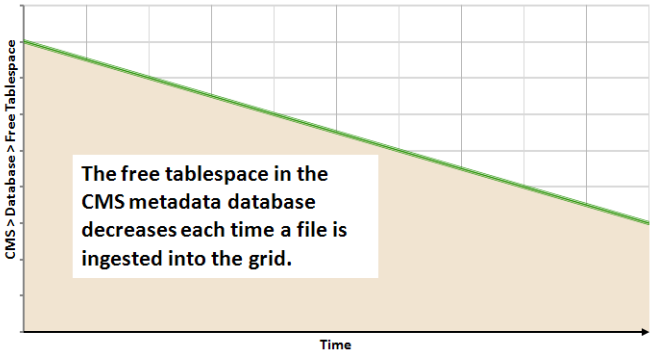 Estimated Remaining Object Capacity (CORS): The estimate of how many more objects can be tracked in the CMS metadata database decreases by 1. 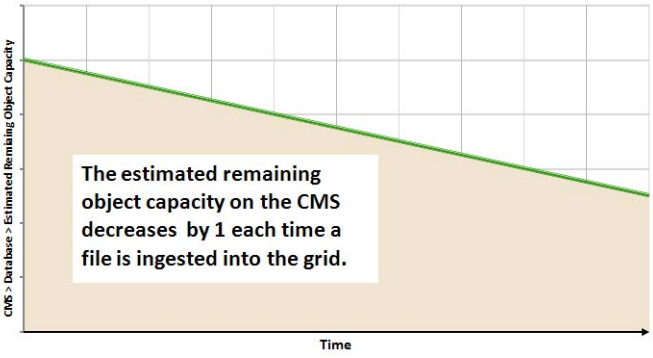

Content Replication

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

Data Flow

The owner CMS, which is the first CMS to receive the object metadata from the LDR, controls the replication, that is, ensures that the correct number of copies are stored in the correct locations for the duration specified by the ILM policy. See Figure 48 below for a simplified step-by-step description of what happens as objects are replicated.

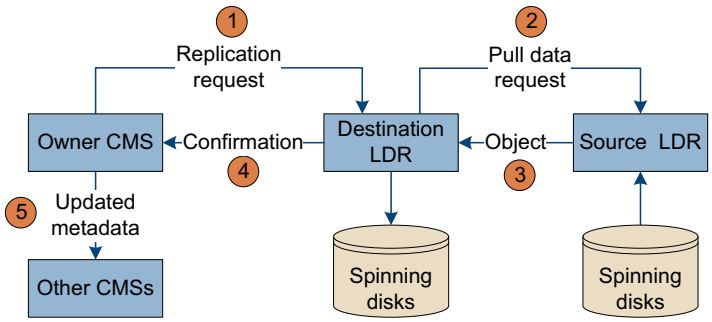

Figure 48: Replication Data Flow

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

Related Attributes

Table 13 below lists some of the NMS MI attributes used to track what happens when a single object is replicated.

Table 13: Object Replication Attributes

Component	Attribute Changes
CMS ▶ Content 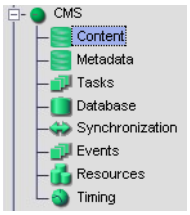 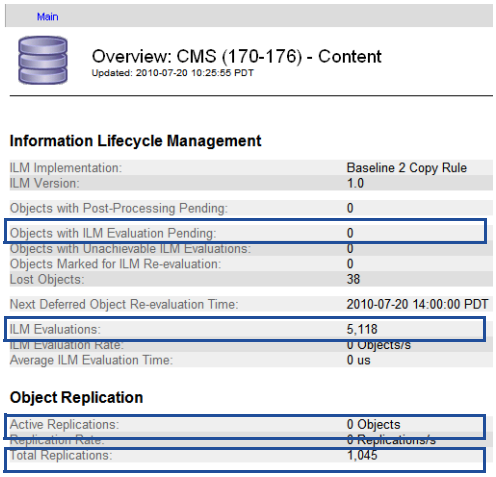 	Objects with ILM Evaluation Pending (ORpe): The number of objects waiting to be processed through the business rules for replication increases by 1 when the object is ingested into the grid and decreases by 1 once replication is complete. Active Replications (DCdA): The number of objects in the process of being replicated increases by 1 when replication starts and decreases by 1 when replication is complete. 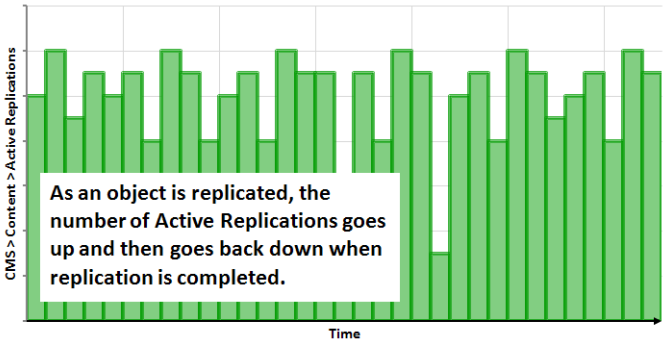 ILM Evaluations (ILev): The total number of ILM evaluations that have been performed to date increases when the object is evaluated after ingest and again when the object has been replicated. **NOTE** The number of evaluations is dependent on the ILM rules for the grid. Total Replications (DCdT): The total number of object replications performed by the owner CMS since grid start-up increases by 1 each time a copy is made. 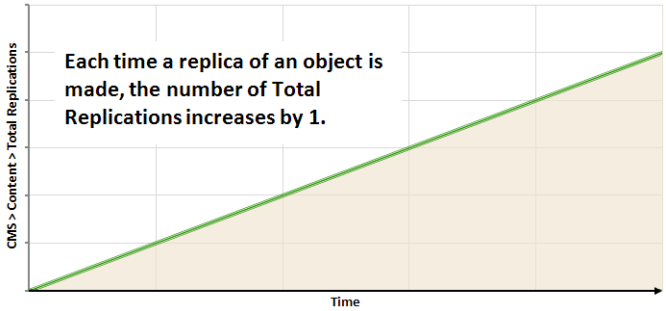

Table 13: Object Replication Attributes (cont.)

Component	Attribute Changes
LDR ▶ Storage 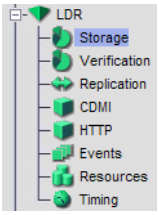 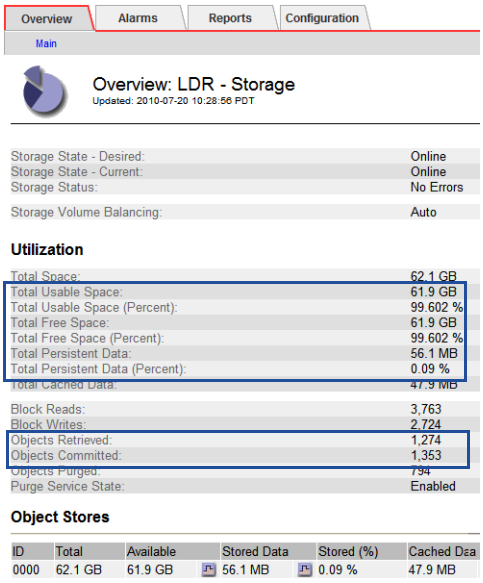 	Total Usable Space (STAS): The total amount of object storage space that is currently available to be used to store objects decreases by an amount roughly equivalent to the file size. Total Usable Space (Percent) (SAVP): The total amount of object storage space (displayed as a percentage) that is currently available to be used to store objects decreases by an amount roughly equivalent to the file size. Total Free Space (SUSA): The total amount of all free space — available to be used or not — on all object stores decreases by an amount roughly equivalent to the file size. Total Free Space (Percent) (SUSP): The total amount (displayed as a percentage) of all free space — available to be used or not — on all object stores decreases by an amount roughly equivalent to the file size. Total Persistent Data (Percent) (SPSD): The estimate of the size of the persistently stored data increases by an amount roughly equivalent to the file size. Total Persistent Data (Percent) (SPDP): The percentage of the total storage space used by persistent data on each destination LDR increases by an amount roughly equivalent to the size of the replicated object. Objects Retrieved (ORET): The number of persistent objects retrieved from the source LDR increases by 1. 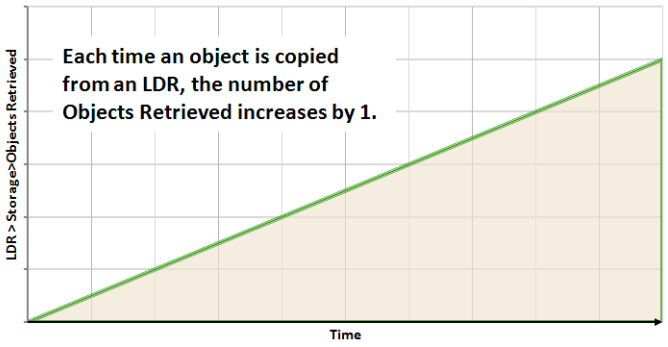 Objects Committed (OCOM): The number of persistent objects stored on each destination LDR increases by 1.

Table 13: Object Replication Attributes (cont.)

Component	Attribute Changes
<*destination*> LDR ▶ Replication 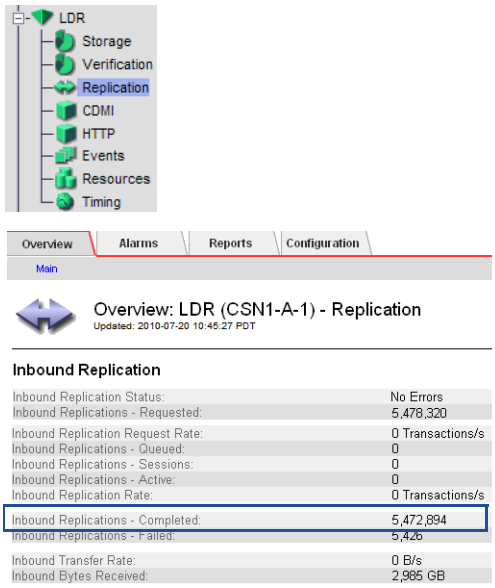 	Inbound Replications Completed (RIRC): The total number of objects replicated to the destination LDR increases by 1. 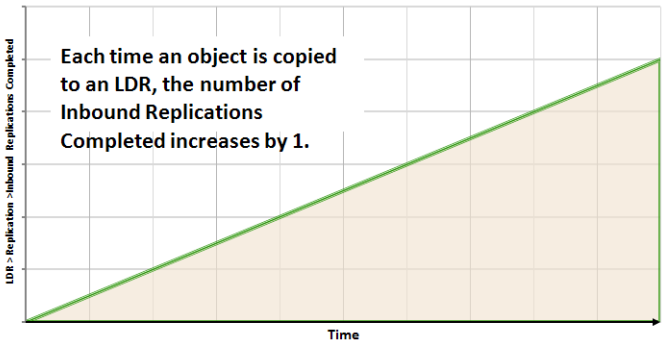
<*source*> LDR ▶ Replication 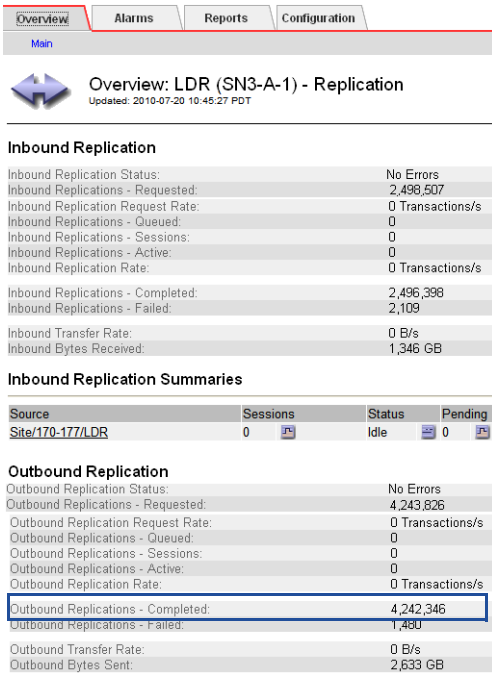 	Outbound Replications Completed (RORC): The total number of objects replicated from the source LDR increases by 1. 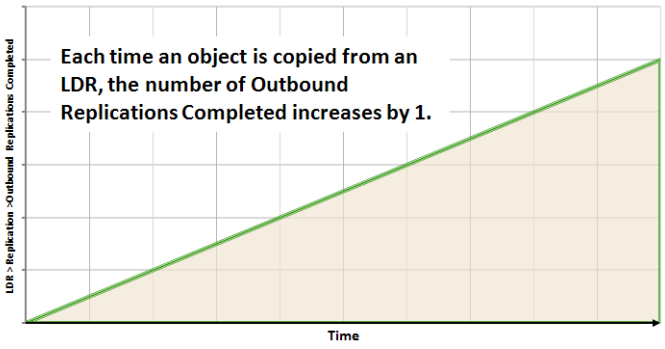

Content Replication to Archive Media

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

Data Flow

If the ILM policy requires an object to be stored on archive media, the CMS sends a request to the Archive Node which in turn sends the object to the middleware (see Figure 49).

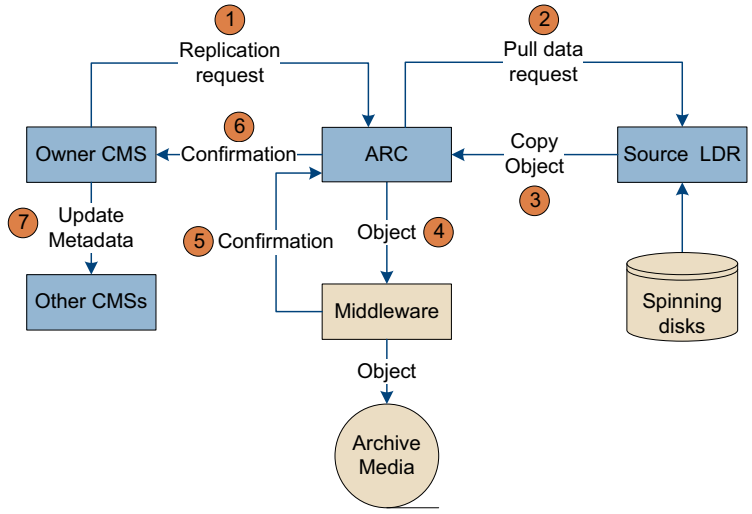

Figure 49: Archiving Data Flow

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

6. The ARC notifies the CMS that the object has been stored.

7. The owner CMS updates the location information and distributes that information to the other CMSs that store metadata for this object.

Related Attributes

Table 14 below lists some of the NMS attributes used to track object replication to archive media.

Table 14: Object Replication to Archive Node Attributes

Component	Attribute Changes
ARC ▶ Store 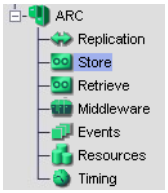 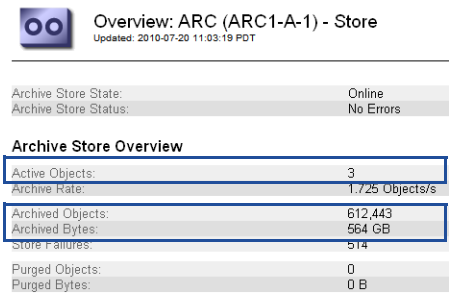 **NOTE** The grid does not know how much installed and available storage is on the archival media device attached to the Archive Node.	Active Objects (AROP): The number of objects in the process of being written to archiving media increases by 1 as the object is being archived and decreases by 1 once the object is archived. 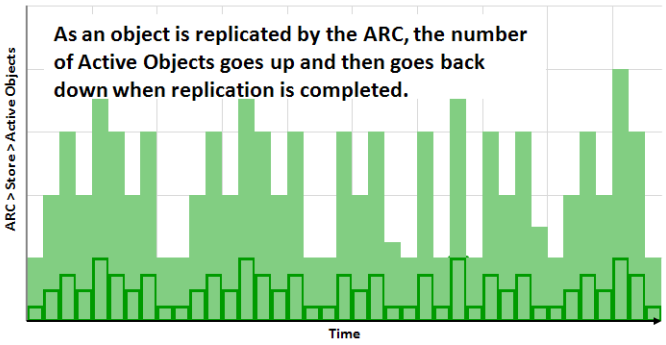 Archived Objects (AROA): The total number of objects written to archive media by this ARC increases by 1. 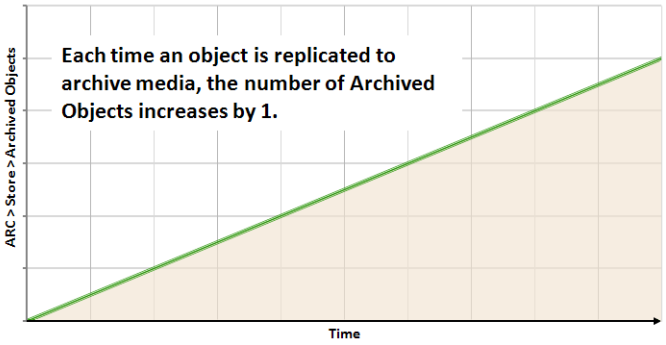 Archived Bytes (ARBA): The total amount of content written to archive media increases by an amount equivalent to the file size.

Table 14: Object Replication to Archive Node Attributes (cont.)

Component	Attribute Changes
ARC ▶ Replication 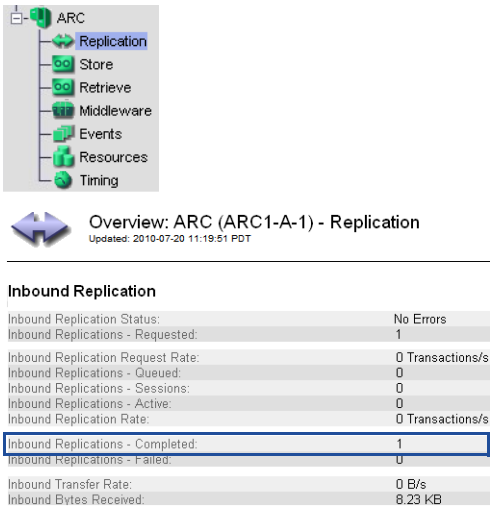	Inbound Replications Completed (RIRC): The total number of objects replicated to the destination ARC increases by 1.
LDR ▶ Storage 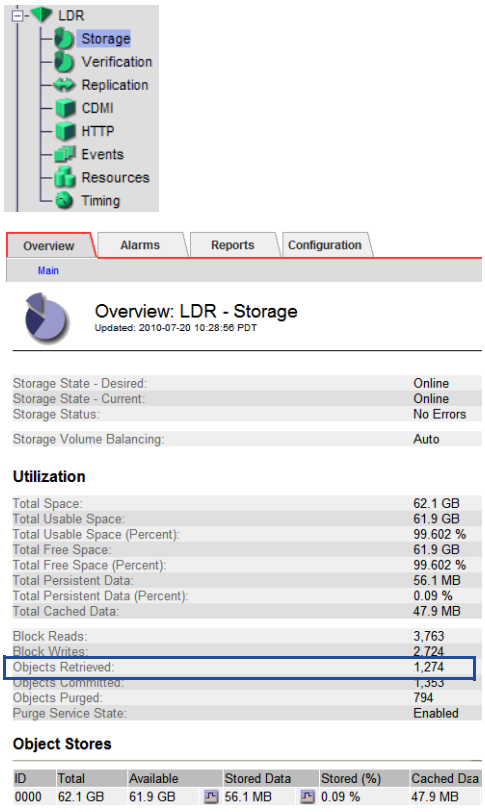	Objects Retrieved (ORET): The number of persistent objects retrieved from the storage system of the source LDR increases by 1 each time the object is replicated from this source LDR to the Archive Node. 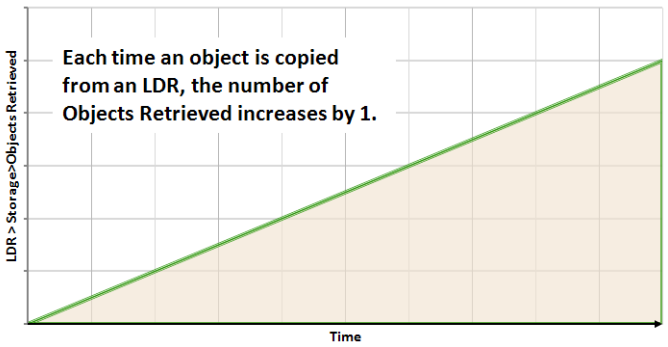

Table 14: Object Replication to Archive Node Attributes (cont.)

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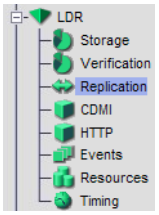 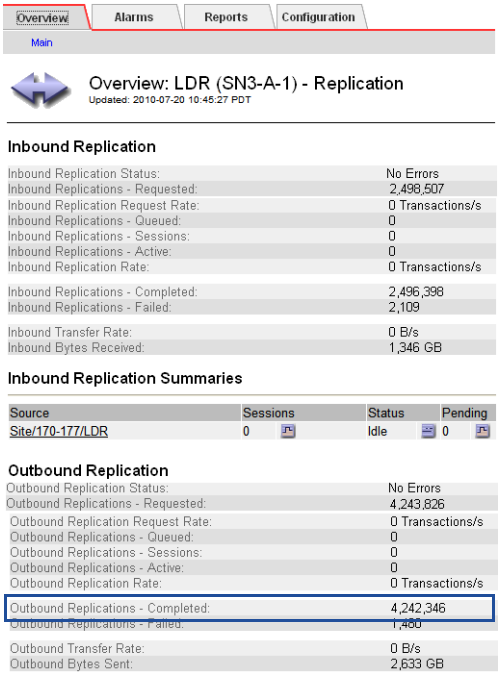	Outbound Replications - Completed (RORC): The total number of objects replicated from the source LDR increases by 1 each time the object is replicated from this source LDR to the Archive Node. 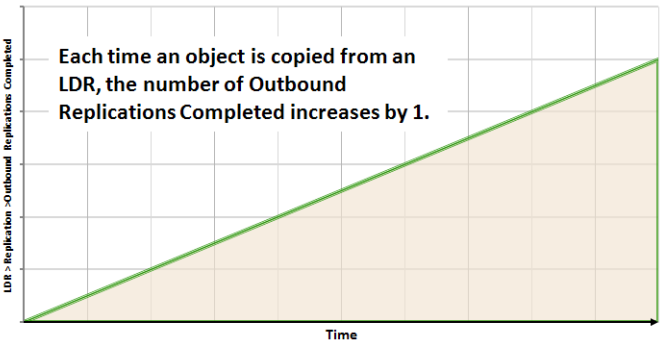

Retrieval

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

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

Data Flow

File in FSG Cache

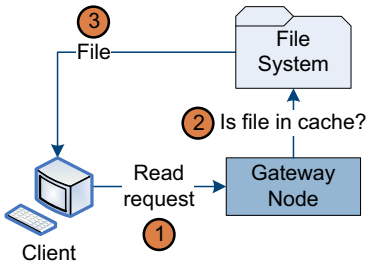

Figure 50: Retrieval Data Flow – File in Cache

1. The client requests the file.
2. The file system (CIFS/NFS) finds the file in the FSG cache.
3. The client reads the file.

File Not in FSG Cache

If the file is not in the FSG cache, the FSG sends a request to an LDR. The LDR returns the file if it has it. Otherwise, the LDR retrieves it from another LDR or an ARC after getting the file location from the CMS (see Figure 51). Retrieval preferentially goes to spinning media under normal performance load balancing. When there is no higher grade copy accessible, the retrieval request is directed to the Archive Node.

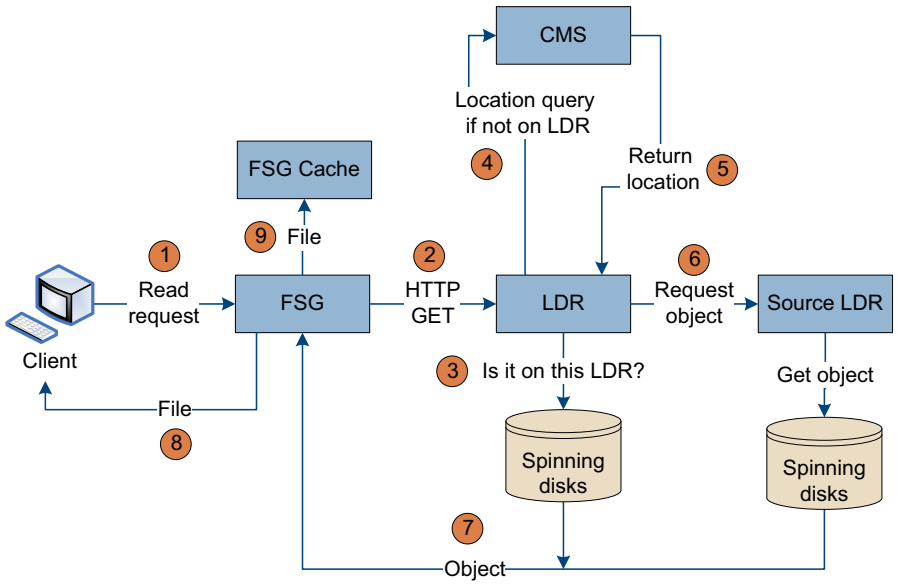

Figure 51: Retrieval Data Flow – File Not in Cache

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

Related Attributes

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

Table 15: Object Retrieval Attributes

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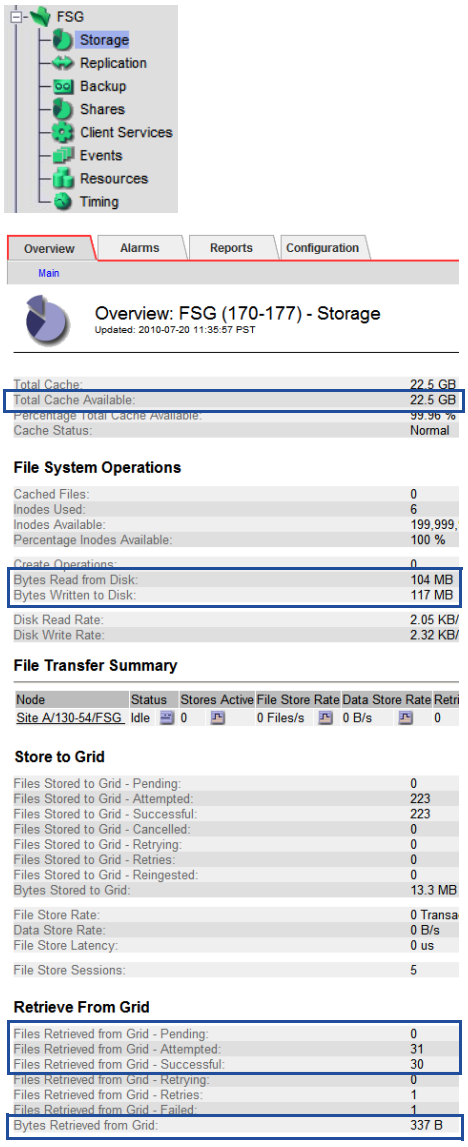 	Total Cache Available (FSTA): The total local cache space that is still available for use does not change if the file was already in the cache. Otherwise, total cache available decreases unless the minimum value has been reached. Bytes Read from Disk (FSRB): If the file is not in the FSG cache, the number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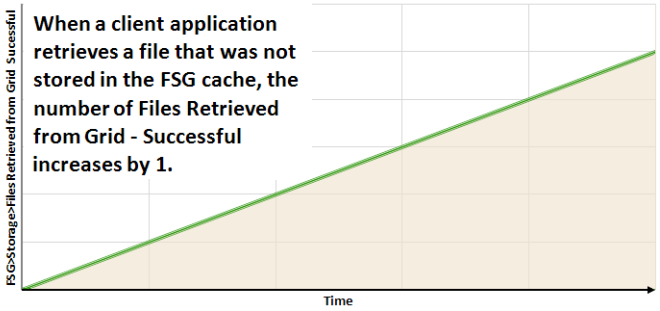 Bytes Retrieved from Grid (FRGB): If the file is not in the FSG cache, the number of bytes retrieved successfully from the grid increases by an amount equivalent to the file size.

Table 15: Object Retrieval Attributes (cont.)

Component	Attribute Changes
LDR ▶ HTTP 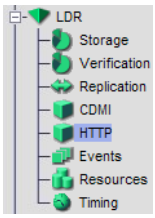  **HTTP Sessions** Currently Established Incoming Sessions: Incoming Sessions - Attempted: Incoming Sessions - Successful: Incoming Sessions - Failed: **HTTP GET** Inbound GETs - Attempted: Inbound GETs - Successful: Inbound GETs - Failed: **HTTP Metadata GET** Inbound Metadata GETs - Attempted: Inbound Metadata GETs - Successful: Inbound Metadata GETs - Failed:	Inbound GETs – Successful (HIGC): If the file is not in the FSG cache, the total number of HTTP GET ("content retrieve") requests that have completed successfully increases by 1. 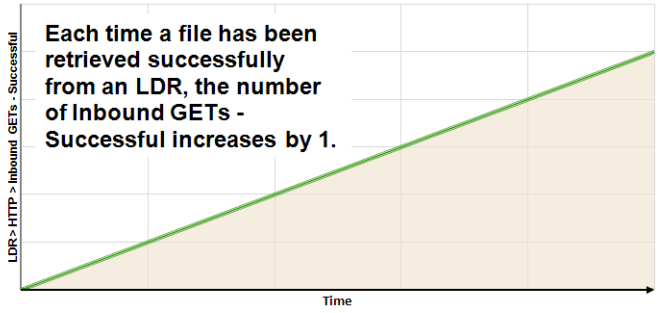 Inbound GETs – Attempted (HAIG): The total number of HTTP GET (content retrieve) requests that have been received by the LDR also increases by 1. **NOTE** Transactions that occur via CDMI do not result in an update to these attributes.

Table 15: Object Retrieval Attributes (cont.)

Component	Attribute Changes
LDR ▶ Storage 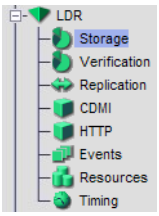 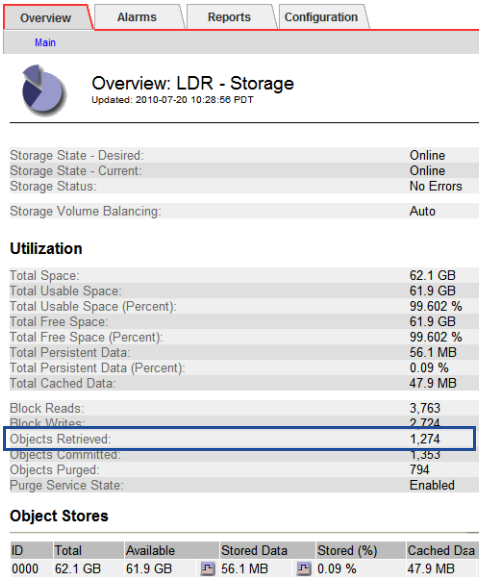	Objects Retrieved (ORET): If the file is not in the FSG cache, the number of persistent objects retrieved from the source LDR increases by 1.

Table 15: Object Retrieval Attributes (cont.)

Component	Attribute Changes
ARC ▶ Retrieve 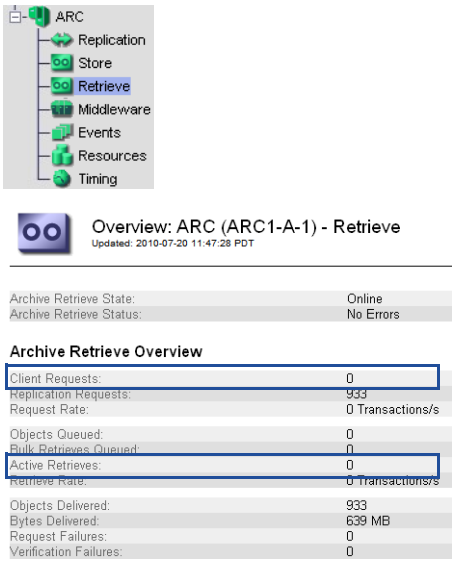	Client Requests (ARCR): The total number of requests received from clients for objects stored on the ARC increases by 1 each time the grid attempts to retrieve an object from this Archive Node. This happens only if the file is not in the FSG cache or on an LDR. 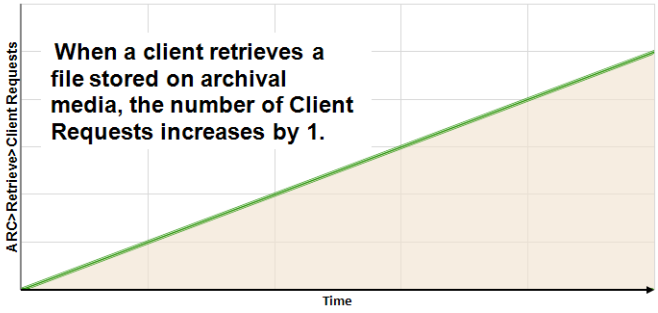 Active Retrieves (ARAR): The number of object retrievals in progress increases by 1 while retrieval is taking place on this Archive Node and then decreases by 1 once retrieval is complete. 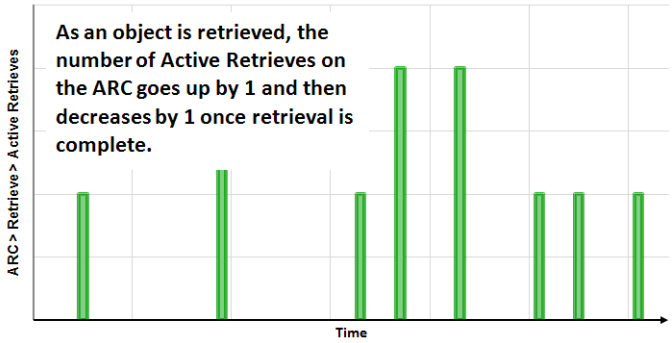

Purging

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

Data Flow

See Figure 52 below for a simplified step-by-step description of what happens when the FSG receives a client request to delete a file.

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

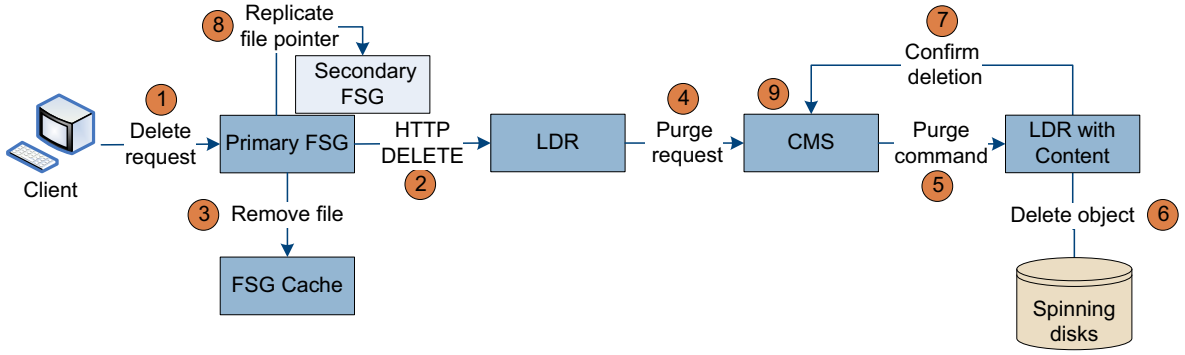

Figure 52: Purging Data Flow

1. The client sends a file delete request to the FSG which removes the file pointer from the FSG file system.
2. The FSG sends an HTTP DELETE command to an LDR. (This releases the content handle.) The LDR that is chosen depends on the result of a topology query.
3. The file is removed from the FSG cache (that is, the file pointer and data are deleted).
4. The LDR notifies a CMS that the content handle has been released.
5. A purge command is sent to all LDRs that have a copy of the object.
6. The LDRs receives the purge command and delete the content.
7. The LDRs notify the CMS that the content has been purged.
8. The primary FSG replicates the file pointer information (deleted) to the other FSGs in its replication group.
9. In a grid that uses metadata replication, after all content locations are purged the metadata is purged. Metadata is never purged in a grid that uses metadata synchronization.

NOTE Metadata synchronization is deprecated and no longer supported.

Related Attributes

Table 16 below lists some of the attributes used to track what happens when a client deletes an object stored in the grid.

Table 16: Object Purging Attributes

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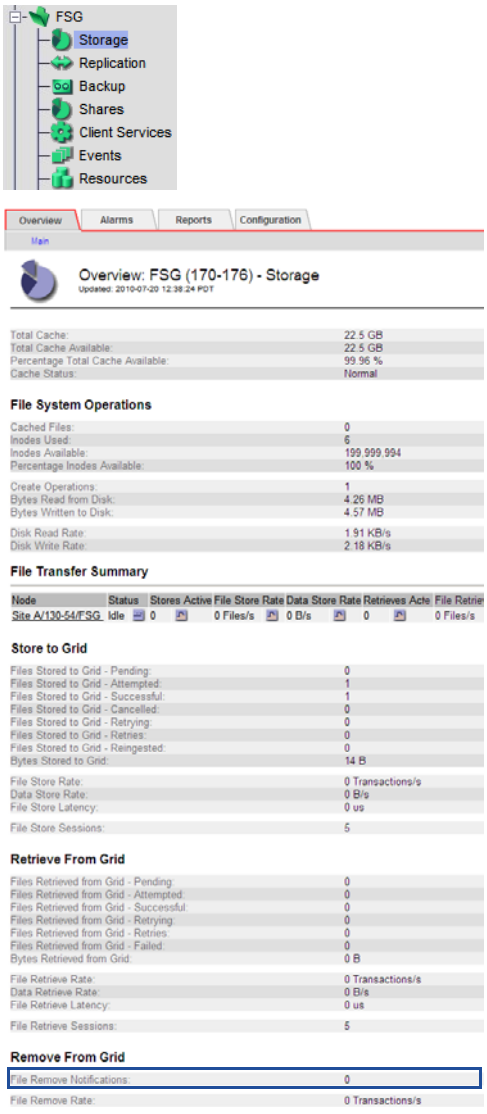	File Remove Notifications (FRGN): The number of content handle release notifications sent by the FSG to LDRs increases by 1 when the FSG receives a request to delete an object. 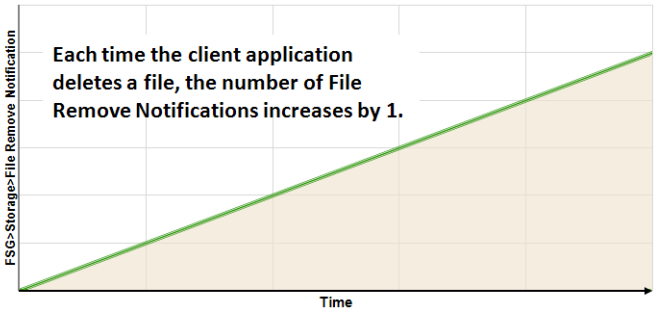

Table 16: Object Purging Attributes (cont.)

Component	Attribute Changes
FSG ▶ Shares 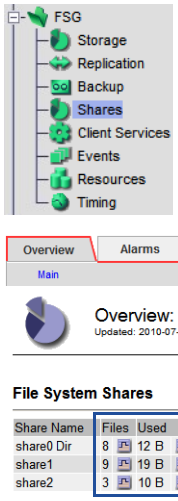 	Files (FSSF): The number of files for the share where the object was ingested decreases by 1. Note that this does not occur until the next FSG backup. Used (FSSB) and Used (Week) (FSSR): The total size of all files referenced by the share decreases by an amount equivalent to the file size. Remaining (FSSA) and Remaining (Week) (FSSL): If a quota has been configured for the share, the total amount of space remaining on the share increases by the size of the file.

Table 16: Object Purging Attributes (cont.)

Component	Attribute Changes
LDR ▶ HTTP 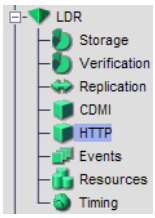 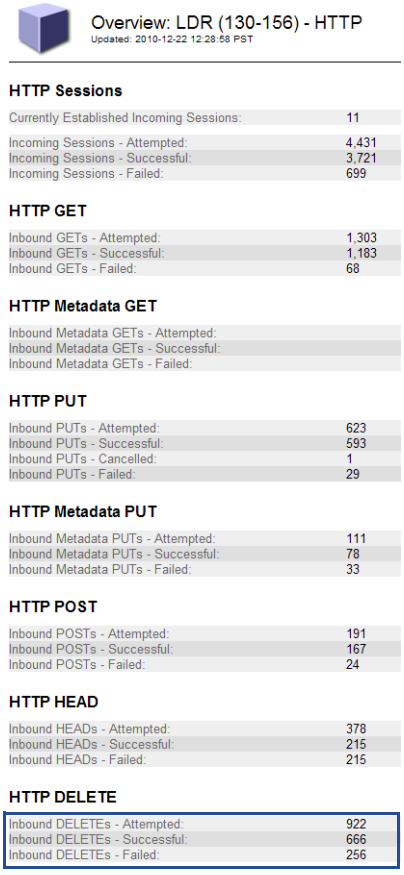 	Inbound DELETEs – Successful (HIDC): The total number of objects for which the content handle has been released increases by 1 after the HTTP DELETE command has been completed successfully. 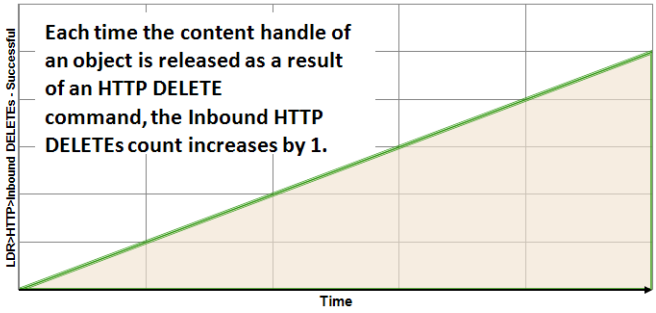 Inbound DELETEs – Attempted (HAID): The total number of HTTP DELETE (content handle release) requests that have been received by the LDR also increases by 1. **NOTE** Transactions that occur via CDMI do not result in an update to these attributes.

Table 16: Object Purging Attributes (cont.)

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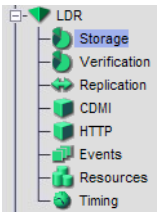 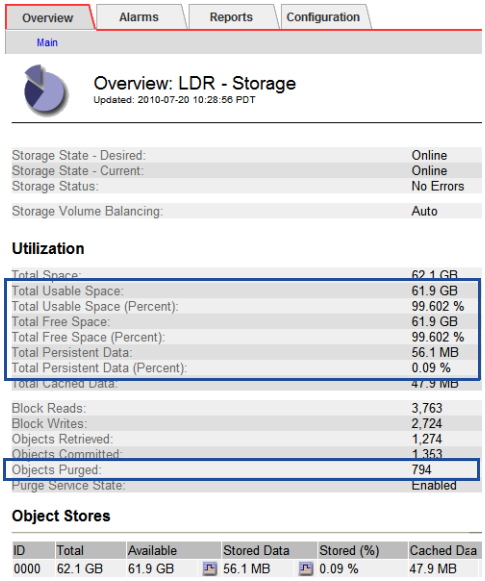 	Objects Purged (OPUR): The number of persistent objects purged from this LDR increases by 1 for each copy of the object purged. 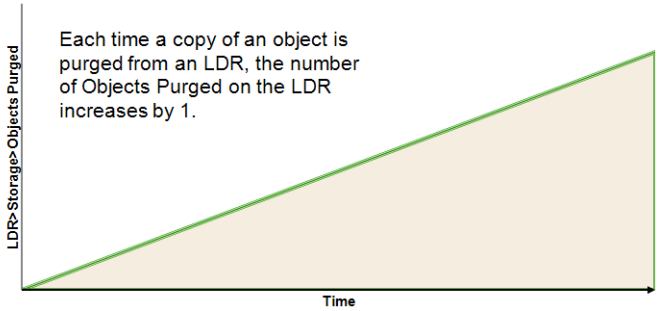 Total Usable Space (Percent) (SAVP): The percentage of object storage space available for use increases by an amount roughly equivalent to the size of the purged object. Total Persistent Data (Percent) (SPSD): The percentage of the total storage space used by persistent data decreases by an amount roughly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

Table 16: Object Purging Attributes (cont.)

Component	Attribute Changes
CMS ▶ Content 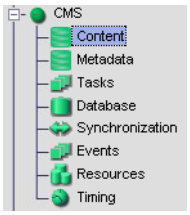 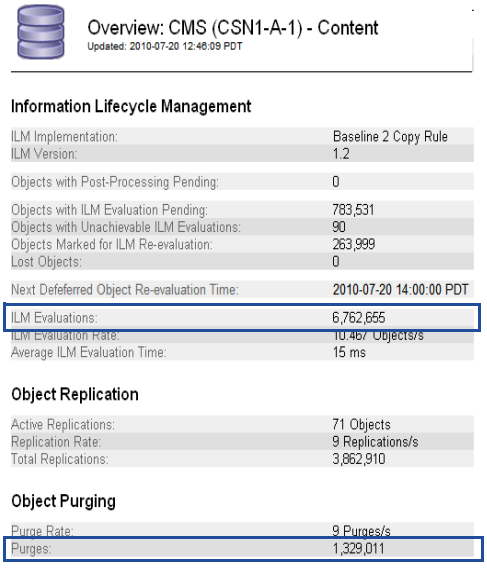 	Purges (DCpT): The number of object copies deleted increases by 1 on the owner CMS for each copy of the object purged from the LDRs and ARCs. 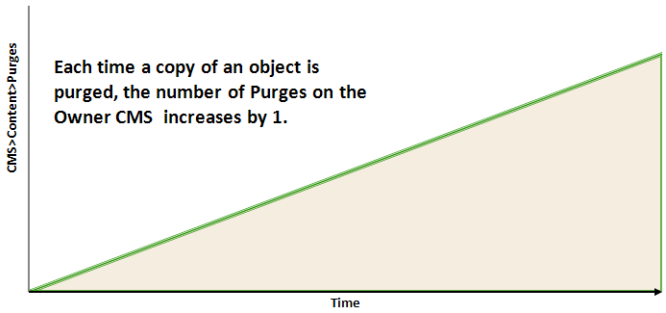 CMS ▶ Database ▶ Free Tablespace Percent (DBSP) and the number of CMS ▶ Metadata ▶ Stored Objects (COoT) decrease in a grid that uses metadata replication. (Note that the owner CMS retains a copy of the metadata of a purged object until it can confirm that all object copies have been deleted, which can lead to inconsistencies in the value of CMS ▶ Metadata ▶ Stored Objects between CMSs. In a grid with a single CMS replication group, these inconsistencies should resolve within a day or two — when the last copy of metadata is purged.) Free Tablespace Percent (DBSP) and the number of Managed Objects (COoM) and Stored Objects (COoT) do not change in a grid that uses metadata synchronization. (With metadata synchronization, COoM is listed under the Content component, not the Metadata component.) **NOTE** Metadata synchronization is deprecated and no longer supported. ILM Evaluations (ILev): The total number of ILM evaluations that have been performed to date increases when an object is purged because an ILM evaluation is triggered each time a content handle is released.

Table 16: Object Purging Attributes (cont.)

Component	Attribute Changes
ARC ▶ Store 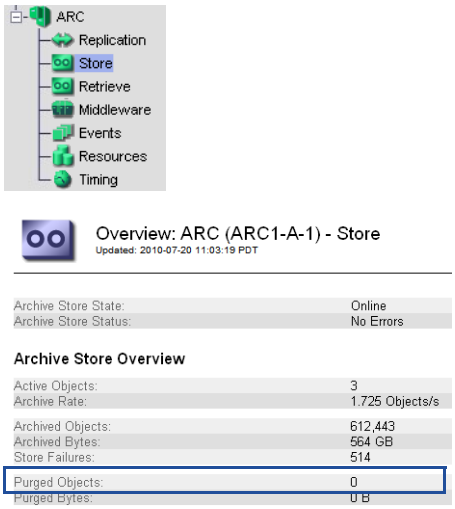	Purged Objects (ADOP): The total number of objects purged from the middleware server by the ARC increases by 1. Whether objects purged from the ARC are actually deleted from archive media depends on the retention settings defined in the middleware server. 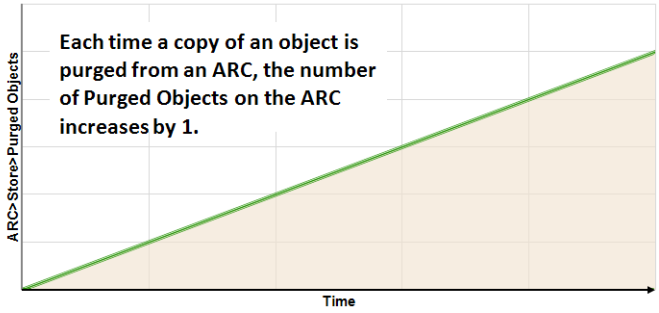

Content Protection

Different levels of deletion protection can be enabled to protect files from being altered or removed after they have been ingested into the grid. For more information, see the *Administrator Guide*.

File Recovery

NOTE File recovery is deprecated and no longer supported.

If file recovery is enabled for an FSG profile, when a file is deleted, it is transferred to a file recovery directory rather than being purged from the grid. The “deleted” file can be accessed and recovered via this file recovery directory and is not purged from the grid until deleted from the file recovery directory. For files transferred to a file recovery directory, a file recovery protection period can also be applied. For more information, see the *Administrator Guide*.

Purging Initiated by ILM Policy

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

File Modification

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

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

Data Flow

See Figure 53 below for a simplified step-by-step description of what happens when the client modifies a file stored in the grid.

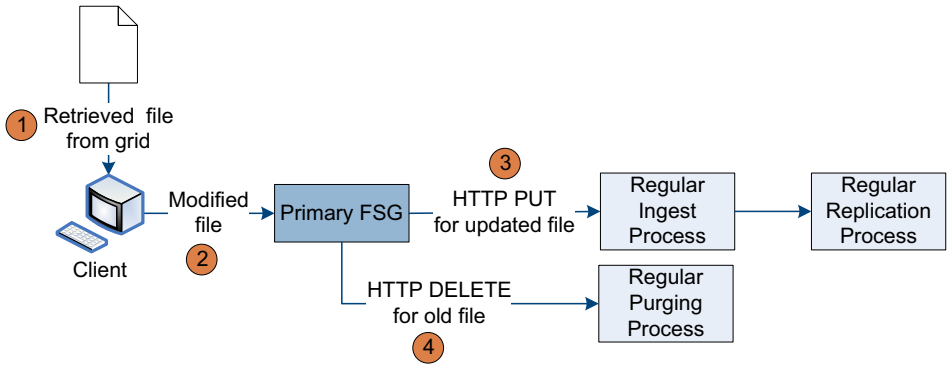

Figure 53: Object Modification Data Flow

1 The client application retrieves the file from the grid via the FSG file share. For more information, see"Retrieval" on page 82.

2. The client application saves the modified file to the grid via the FSG file share.

3. The primary FSG saves a copy of the object in its cache and sends an LDR an HTTP PUT command. The data flow is identical to that for ingest. The LDR allocates a new unique identifier to the object and saves the object to spinning disk. The LDR notifies a CMS that a new piece of content has been ingested and sends the object metadata to the CMS. This CMS becomes the owner CMS. The object is then replicated according to the ILM policy. For more information, see "Ingest" on page 66 and "Content Replication" on page 75.

4. The primary FSG also sends an HTTP DELETE command to an LDR to delete the old version of the object. What happens next is identical to what happens when an object is purged. For more information, see"Purging" on page 88.

File Recovery

NOTE File recovery is deprecated and no longer supported.

If file recovery is enabled for an FSG profile, when a file is modified, the original file is transferred to a file recovery directory rather than being purged from the grid. The original file can be accessed and recovered via this file recovery directory. For more information, see the *Administrator Guide*.

Related Attributes

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

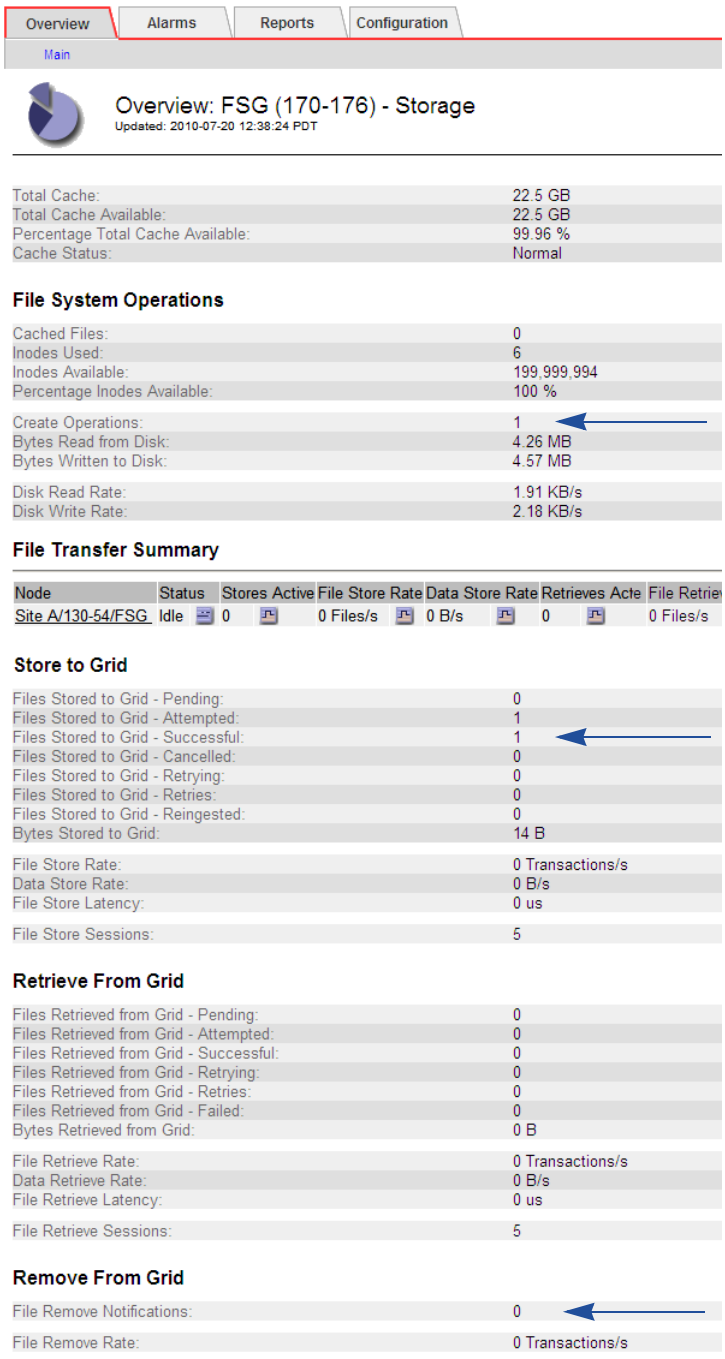

Create Operations (FCRO) does not change when a file is modified

Files Stored to Grid - Successful (FSGC) increases when a file is modified

File Remove Notifications (FRGN) increases when a file is modified

Figure 54: FSG Storage Attributes For Modified Object

FSG Replication

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

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

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

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

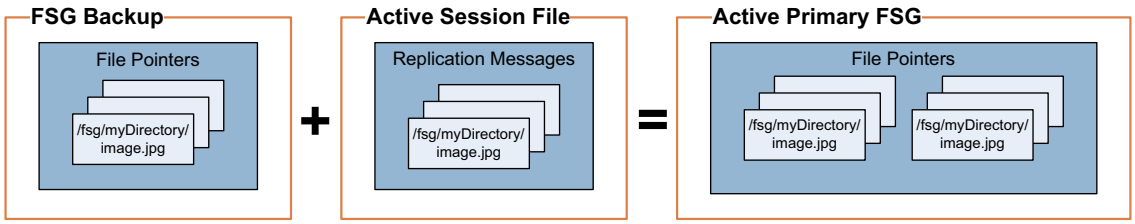

Figure 55: FSG Backups and Active Session File

See Figure 56 below for a simplified description of the FSG replication message flow when a file is ingested, deleted, or modified.

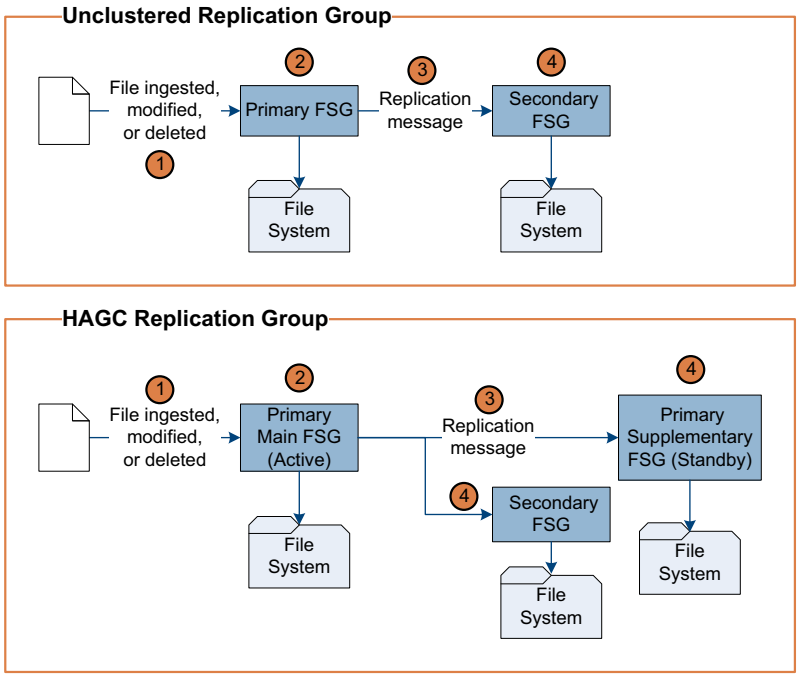

Figure 56: FSG Replication Data Flow

1. The client creates, modifies, or deletes a file via the FSG file share.
2. The primary FSG either creates a file pointer, modifies the file pointer, or deletes the file pointer from its file system.
3. The primary FSG sends a replication messages to the other FSGs in its replication group.
4. The other FSGs in the replication group process the replication messages in real time and update their file system.

Related Attributes

Table 17 below lists some of the NMS MI attributes used to track FSG replication.

Table 17: FSG Replication Attributes

Component	Attribute Changes
Primary FSG ▶ Replication 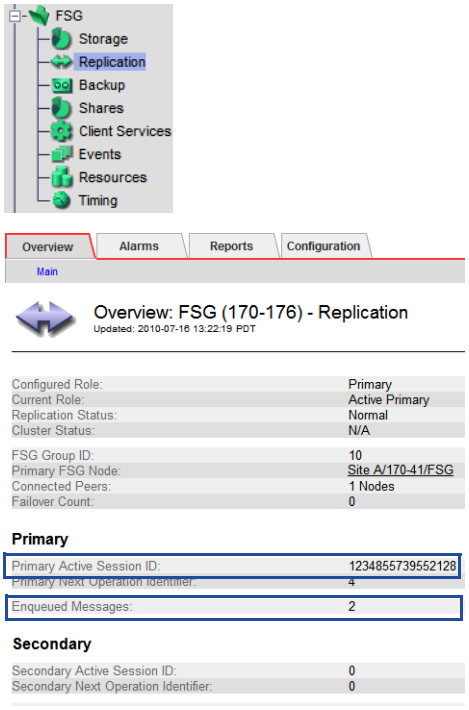	Primary Active Session ID (PAID): A new unique identifier for the current replication session is assigned each time a new session is started by the active primary FSG. A new session is started when an FSG failover occurs or when the size or age of the current session file exceeds an internal threshold. This number is the same as the Secondary Active Session ID (SAID) on the secondary FSG (see below). Enqueued Messages (PEOP): The total count of replication messages generated increases each time a file is ingested, modified or purged. The count increases by more than one for some operations. For example, ingest generates two replication messages (one for the initial file creation and a second message to associate the UUID after the file has been ingested). Modify also generates two replication messages (one to release the old UUID and one to assign the new UUID after ingest). Attribute events on the file (change permissions, etc.) may also generate additional replication messages. The number of enqueued messages matches the number of dequeued messages at the secondary FSG (see below). Because this FSG is the primary FSG, the fields in the "Secondary" section do not apply.

Table 17: FSG Replication Attributes (cont.)

Component	Attribute Changes
Secondary FSG ▶ Replication 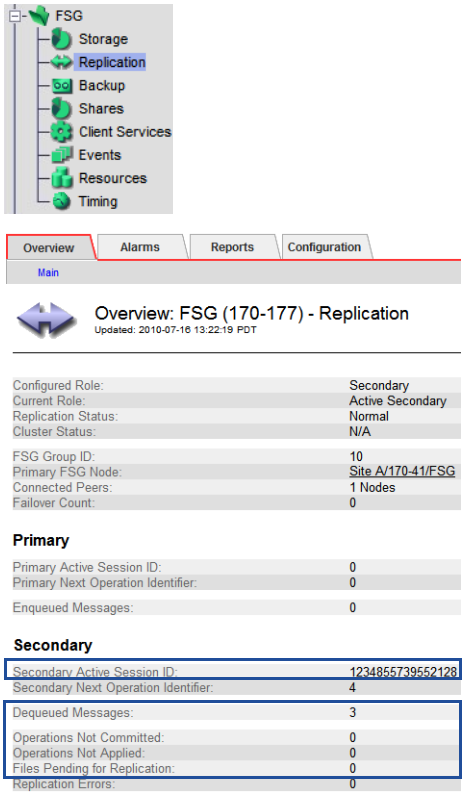	Secondary Active Session ID (SAID): The unique identifier for the current replication session from which messages are being processed is the same as the Primary Active Session ID (PAID) on the primary FSG (see above). Dequeued Messages (SDOP): The total count of dequeued replication messages increases each time a message has been processed by the secondary FSG. The number of dequeued messages matches the number of enqueued messages at the primary FSG (see above). Operations Not Committed (SUOP): The number of replication messages from the primary FSG that have not yet been written to the secondary FSG increases temporarily during periods of high grid activity. Operations Not Applied (SPOP): The number of replication messages to be processed on the secondary FSG in order to catch up to the primary FSG may temporarily increase during periods of high grid activity or backups. (In older systems where “offline backups” were used, SPOP increased during backups. As of Release 8.1, offline backups are deprecated.) If this FSG is the standby primary FSG in an HAGC, the fields in the “Primary” section do not apply. The same is true for a secondary FSG. Files Pending for Replication (FRPP): Number of new files that have not been fully replicated because they are awaiting a replication message from the primary FSG indicating that the file has been successfully stored to the grid. Similar to Operations Not Applied (SPOP), the value for this attribute may temporarily increase during periods of high grid activity.

4

Monitoring Operations

Day-to-day monitoring activities that help ensure smooth grid operation

Monitoring Key Attributes

The NMS MI displays hundreds of attributes; however, most of these attributes are required only for troubleshooting. The list of attributes to monitor routinely, shown in Table 18 below, is much shorter. Tips on how to analyze these attributes are described in the remainder of this chapter.

Table 18: Key Attributes to Monitor

Category	Component	Code	Description	See
LDR content storage capacity	Grid Overview 	PSCU	Percentage Storage Capacity Used: The percentage of installed storage capacity that has been used up for the entire grid.	page 107
	LDR ▶ Storage 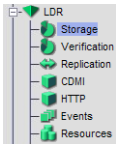	STAS	Total Usable Space: Object storage capacity that is currently available for storage on the LDR.	
		SUSA	Total Free Space: Total amount of all free space — available to be used or not — on all object stores.	
CMS metadata storage capacity	Grid Overview 	PMCA	Percentage Metadata Capacity Available: An estimate of how much metadata capacity remains in the CMS databases in the grid.	
	CMS ▶ Database 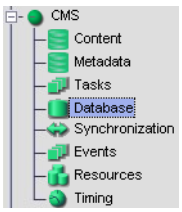	CORS	Estimated Remaining Object Capacity: The number of additional objects that can be managed by this CMS.	page 109
		DBSP	Free Tablespace (Percent): The metadata storage capacity that is still available for use on this CMS.	

Table 18: Key Attributes to Monitor (cont.)

Category	Component	Code	Description	See
Ingest load	FSG ▶ Storage 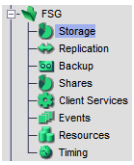	FSGP	Files Stored to Grid - Pending: The number of new ingested files cached locally that are waiting for transfer to the grid for persistent storage.	page 116
Retrieve load	FSG ▶ Storage 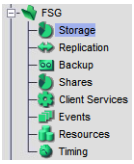	FRTM	File Retrieve Latency (FRTM): The average amount of time required to retrieve a file from the grid. Look at this attribute along with the related attributes Bytes Retrieved from Grid, File Retrieve Rate, and Data Retrieve Rate.	page 118
FSG capacity	FSG ▶ Backup 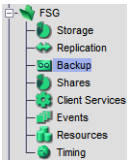	PBNF	Number of Files: The number of files included in the last FSG backup.	page 112
		PBDS	Backup Data Size: The total size of all files referenced by the last FSG backup.	
	FSG ▶ Shares 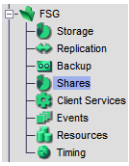	FSSF	Files: The number of files saved to the share. Calculated from the last FSG backup.	
		FSSB	Used: The amount of storage space used by the share. Calculated from the last FSG backup.	
		FSSA	Remaining: If a quota has been configured for the share, the amount of space remaining on the share. Calculated from the last FSG backup.	

Table 18: Key Attributes to Monitor (cont.)

Category	Component	Code	Description	See
FSG capacity	FSG ▶ Storage 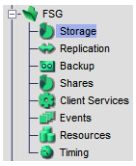	FSIU	Inodes Used: Number of inodes (files and directories) used on the FSG file system.	page 112
		FSIA	Inodes Available: Number of inodes (files and directories) available on the FSG file system.	
		FSIP	Percentage Inodes Available: The percentage of inodes (files and directories) available on the FSG file system.	
Object replication	CMS ▶ Content 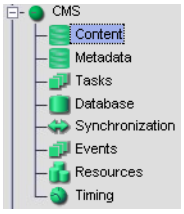	ORun	Objects with Unachievable ILM Evaluations: The number of objects whose ILM business rules cannot be met because the topology or operational state of the grid prevents the rules from being satisfied.	page 120
		ORpe	Objects with ILM Evaluation Pending: The number of objects waiting to be processed through the business rules for replication.	
Metadata synchronization	CMS ▶ Synchronization 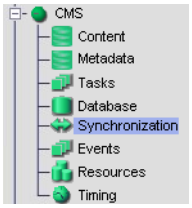	CsQT	Queue Size: The number of outgoing metadata synchronization messages queued to be sent to other CMSs. **NOTE** Metadata synchronization is deprecated and no longer supported.	page 121

Table 18: Key Attributes to Monitor (cont.)

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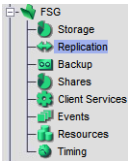	FRPP	Files Pending for Replication: Number of new files that have not been fully replicated because they are awaiting a replication message from the primary FSG indicating that the file has been successfully stored to the grid.	page 122
		SUOP	Operations Not Committed: The number of replication messages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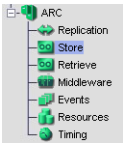	ARBA	Archived Bytes: The total amount of content written to archive media by this ARC.	page 129

Regular Tasks

Table 19 below lists the tasks to be performed on a regular basis.

Table 19: Regular Tasks

Task	Frequency	See
Monitor System Status. Note what has changed from previous day.	Daily	page 44
Monitor system status lights on hardware.	Daily	
Monitor the rate at which LDR storage capacity is being used up.	Weekly	page 107
Monitor the rate at which content metadata storage capacity on the CMS is being used up.	Weekly	page 109
Monitor FSG capacity.	Weekly	page 112

Table 19: Regular Tasks

Task	Frequency	See
Monitor FSG file share usage.	Weekly	page 114
Check available space on the archive media.	Weekly	page 129

Monitor the key attributes regularly to become familiar with grid operations and spot trends before they turn into problems. The important attributes to monitor relate to:

- Content storage capacity on LDRs
- Metadata storage capacity on CMSs
- FSG capacity
- Attribute storage capacity on NMSs
- Ingest load on FSGs
- Retrieve load on FSGs
- Metadata synchronization/replication on CMSs

> **NOTE** Metadata synchronization is deprecated and no longer supported.

- ILM replication on CMSs

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

Content Storage Capacity

The LDRs on Storage Nodes are responsible for storing objects in the grid. You need to monitor the total usable space available on Storage Nodes to make sure the grid does not run out of space to store content. This information is available at the grid level (see Figure 57), at the site level, and at the node level (see Figure 58).

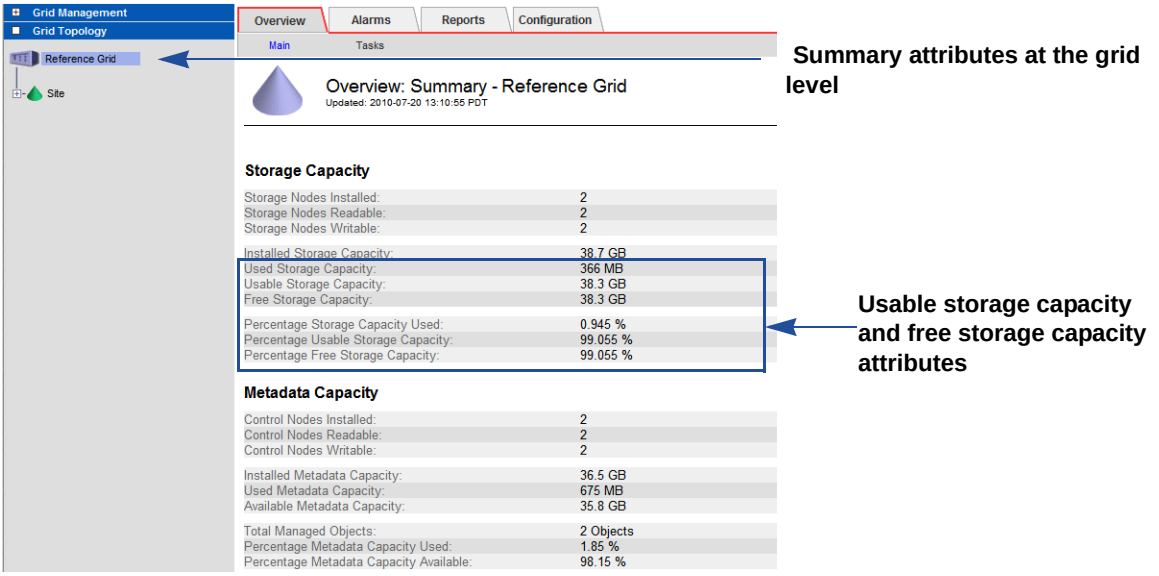

Figure 57: Overall Storage Node Capacity

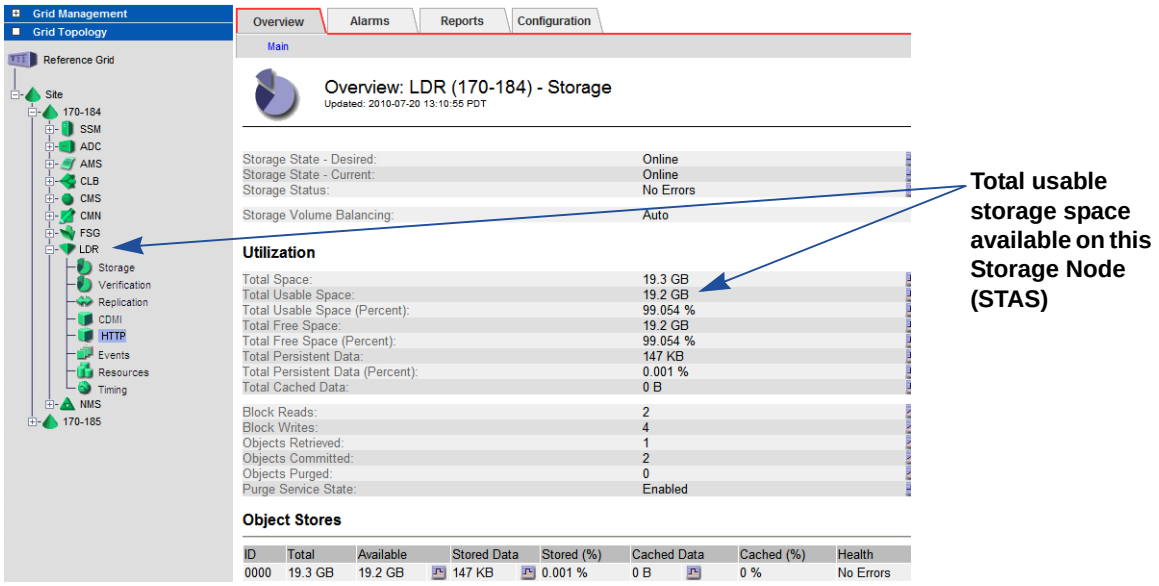

Figure 58: Individual LDR Storage Capacity

On the LDR ▶ Storage ▶ Overview page, monitor the Total Usable Space (STAS) attribute over a period of time to estimate the rate at which usable object storage space is being consumed. Usable space is the actual real amount of storage space available to store objects. For more information on storage space, see the *Administrator Guide*.

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

In the example shown below (Figure 59), usable content storage space is being consumed at a rate of approximately 4% per month, which means that there are 8 months left before this LDR runs out of storage space.

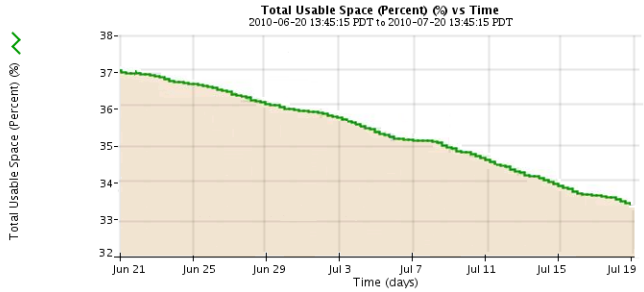

Figure 59: Usable Contents Storage Space

Metadata Storage Capacity

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

Metadata storage capacity information is available at the grid level (see Figure 60 for an example), site level, and node level (see Figure 61 for an example).

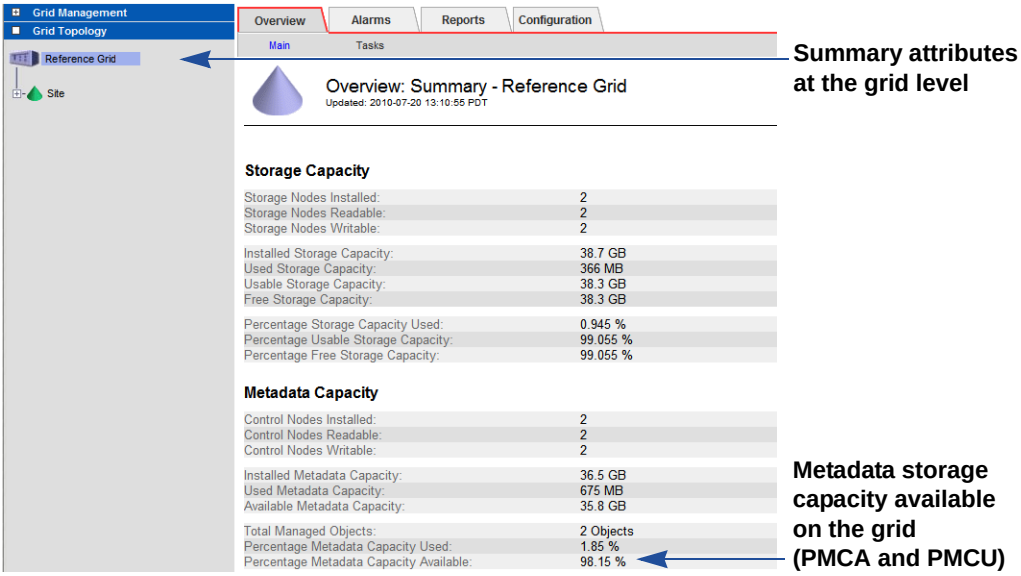

Figure 60: Summary CMS Metadata Storage Capacity

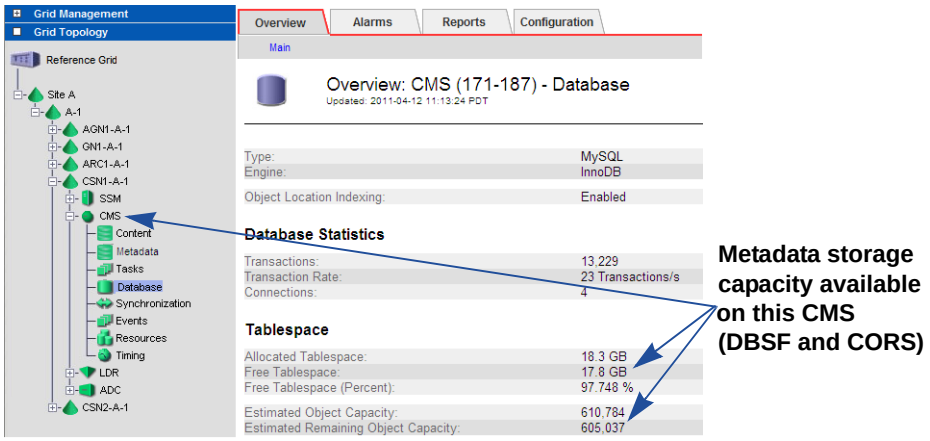

Figure 61: Individual CMS Metadata Storage Capacity

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

NOTE Metadata synchronization is deprecated and no longer supported.

The CMS databases go into “read-only” mode when Free Tablespace drops below 10%. To maintain normal grid operations, you have to

add Control Nodes before the metadata database fills up or if the database size is less than the maximum size of 800GiB perform a hardware refresh to expand the database.

Figure 62 is for the same grid as the one used in Figure 59: “Usable Contents Storage Space”. The migration effect is not noticeable in Figure 59 because the object size is small: the metadata of a large quantity of small objects use more space proportionally than the actual content.

For the example shown in Figure 62, metadata storage capacity is being consumed at a rate of approximately 20 GB per month. However, note how the utilization rate increases towards the end of the period. This could be due, for example, to a data migration that is happening in parallel with regular grid ingest.

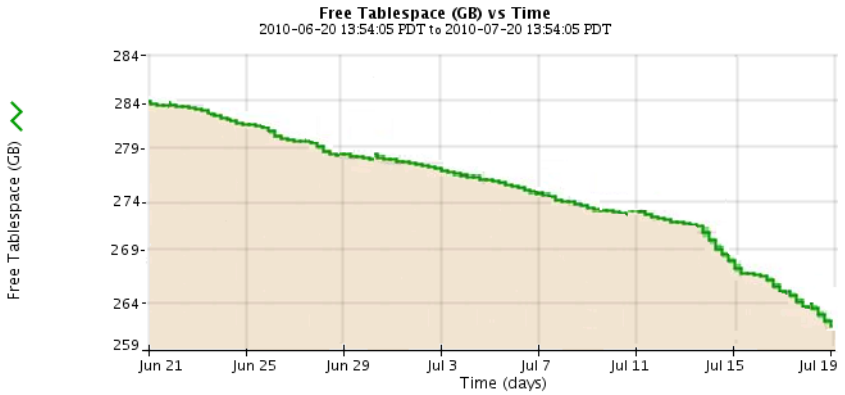

Figure 62: Available CMS Content Storage Metadata Capacity

Gateway Node Capacity and Load

Summary Attributes

Gateway Node information is also available at the grid and site level on the Overview tab (see Figure 63 below).

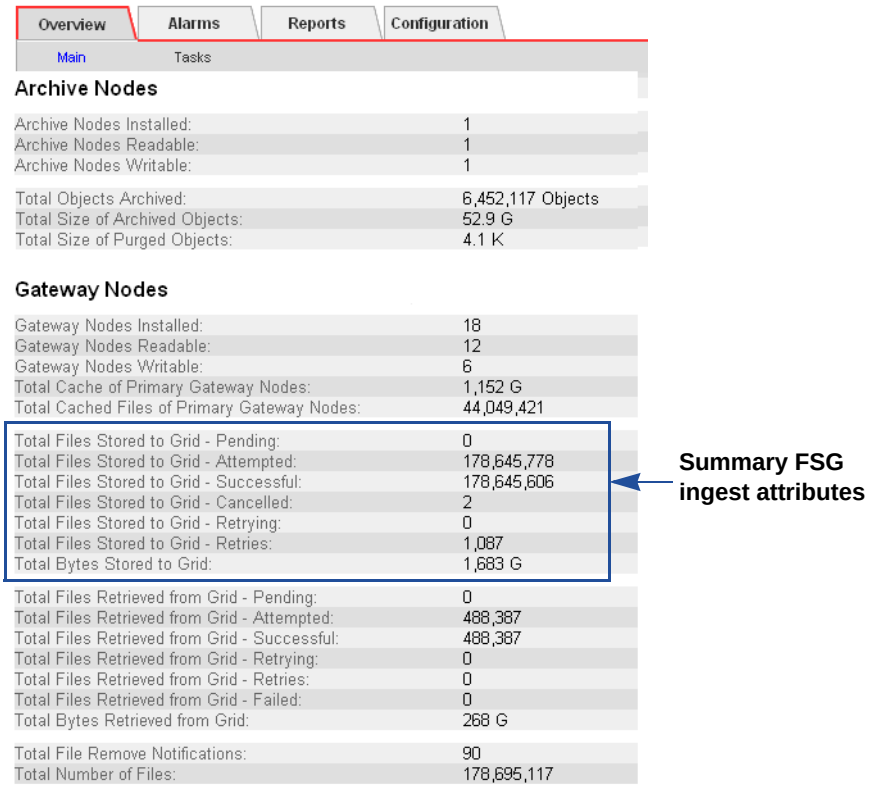

Figure 63: Gateway Node Summary Attributes

FSG Capacity

You can get an estimate of how much content is managed by each FSG replication group by looking at the FSG backup information: Number of Objects (PBNO), Number of Files (PBNF) and Backup Data Size (PBDS) (see Figure 64). The Backup Number of Objects should increase steadily.

The FSG ▶ Backup component displays information on a per FSG basis while the FSG ▶ Shares component displays FSG usage on a per share basis. For more information on the FSG ▶ Shares component, see "File Share Usage" on page 114. To determine the overall amount of capacity remaining on the FSG, use the FSG ▶ Backup component. The backup values reflect the FSG managed file system as of the most recent backup. This includes files that are pending for ingest and files for which ingest into the grid is disabled through FSG profiles.

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

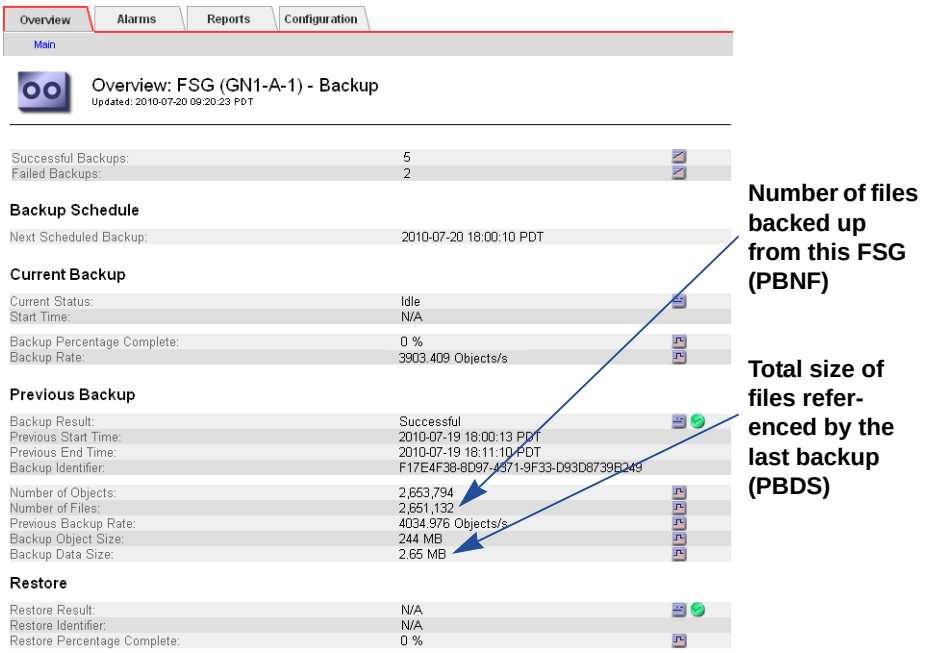

Figure 64: FSG Backup Size

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

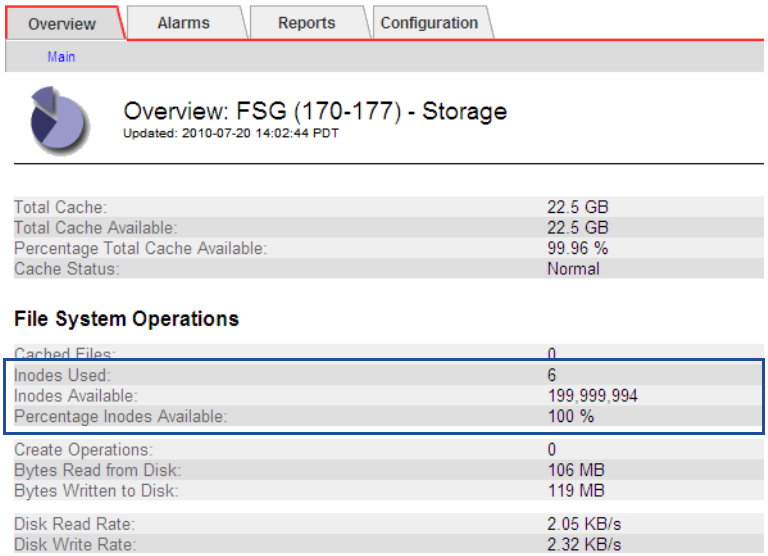

Figure 65: Inodes Used and Availability

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

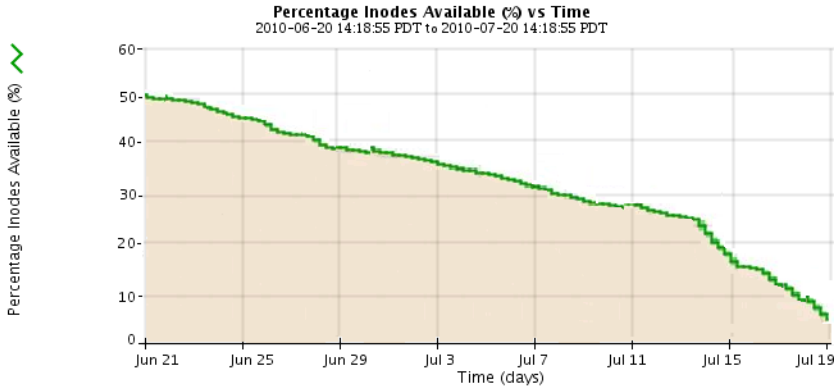

Figure 66: Percentage Inodes Available

File Share Usage

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

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

Important values to monitor are:

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

value exceeds the configured Quota (Weekly) value, an alarm is raised. The administrator can then either increase the quota for the share or remove files from the share.

- **Remaining (Week) (FSSL)** — The amount of space remaining in the share's weekly storage quota.

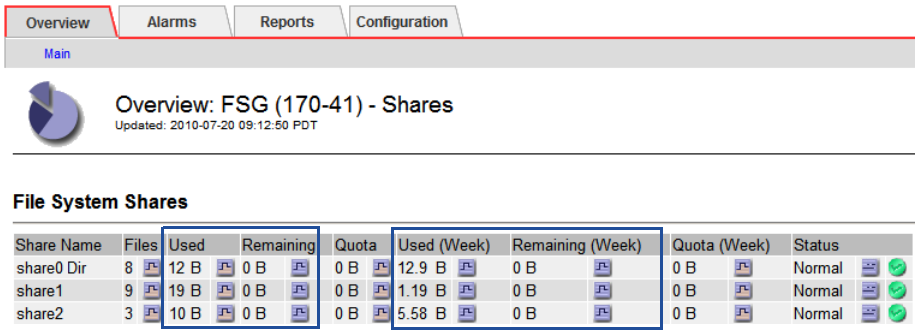

Figure 67: File Share Usage

Ingest Load

To monitor the ingest load on the grid, analyze the trends of these four attributes over time (see Figure 68 below):

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

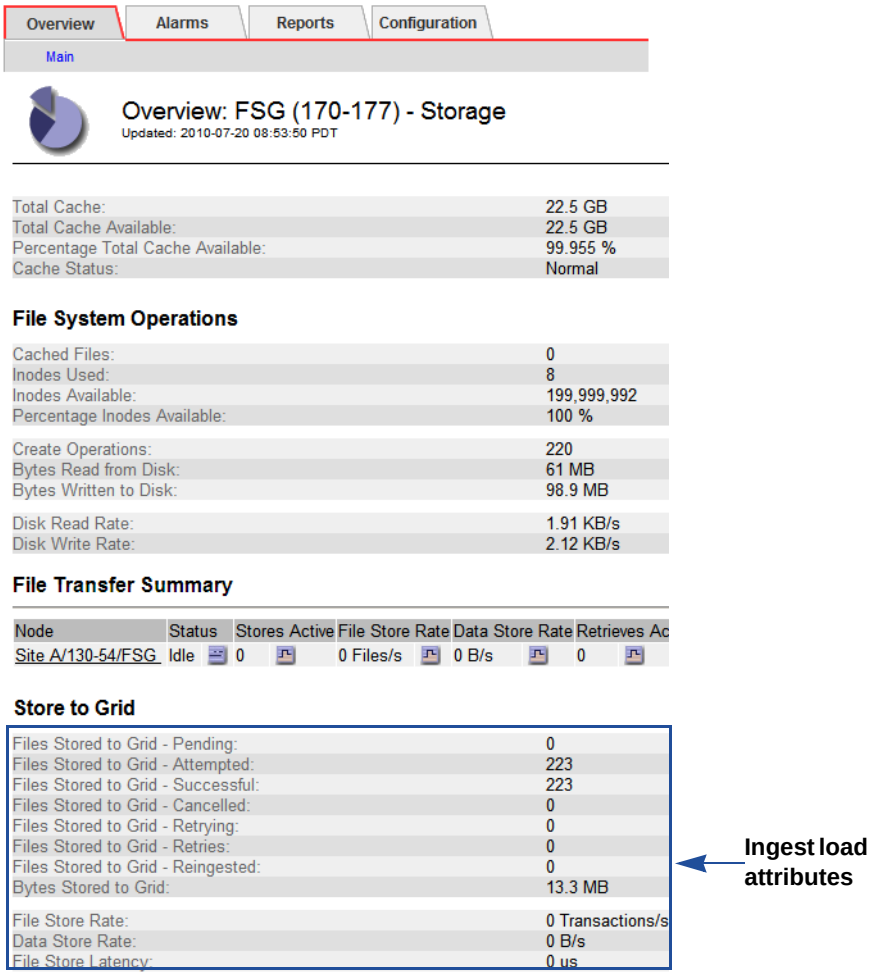

Figure 68: Ingest Load Attributes

Increasing Ingest Load

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

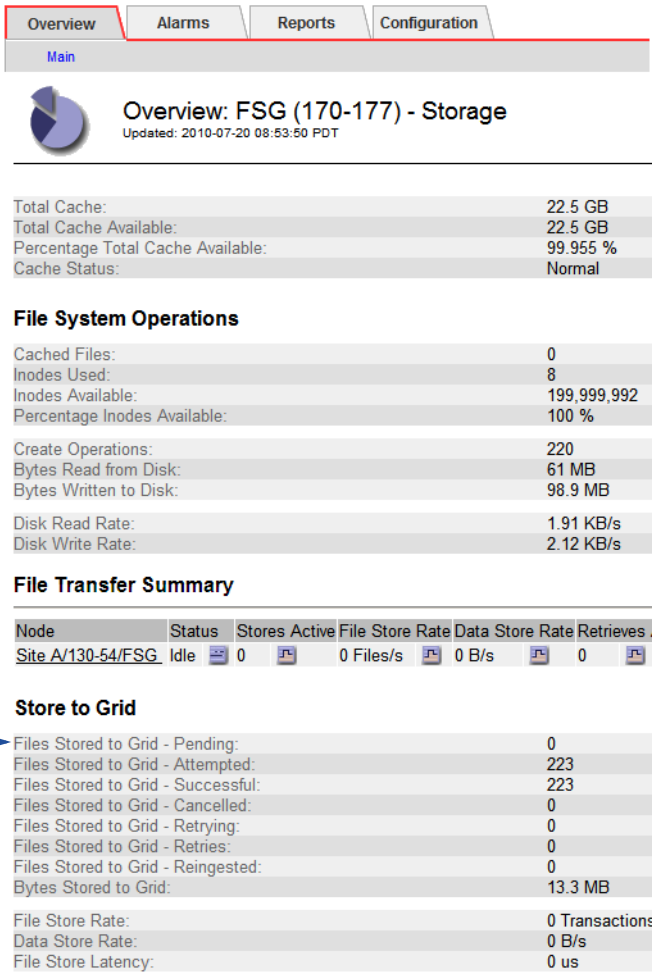

Number of files saved to the FSG by client application and waiting to be stored to the grid (FSGP)

Figure 69: Files Stored to Grid Pending

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

In the example shown in Figure 70 below, the number of files waiting to be stored goes up and down, but remains fairly low. Such a trend could indicate that there was a short term overload due to network throughput, disk I/O performance, grid services availability, and so on.

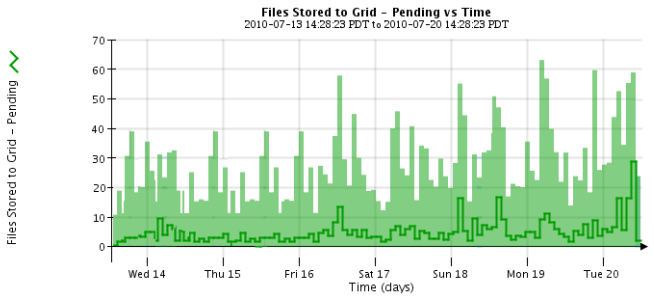

Figure 70: Files Stored to Grid Pending

In contrast, the trend shown in Figure 71 is not sustainable. If the number of files waiting to be stored to the grid starts to increase, make sure that all CMS and LDR services are operating normally. It is also possible that the ingest rate is exceeding the throughput of the grid and that a grid expansion is required.

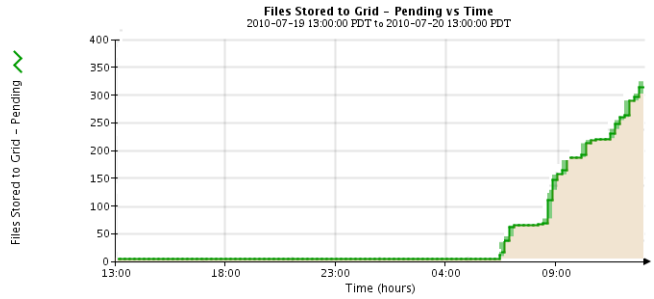

Figure 71: Non Sustainable Ingest Load

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

Retrieve Load

To monitor the retrieve load on the grid, analyze the trends of these four attributes over time (see Figure 72 below).

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

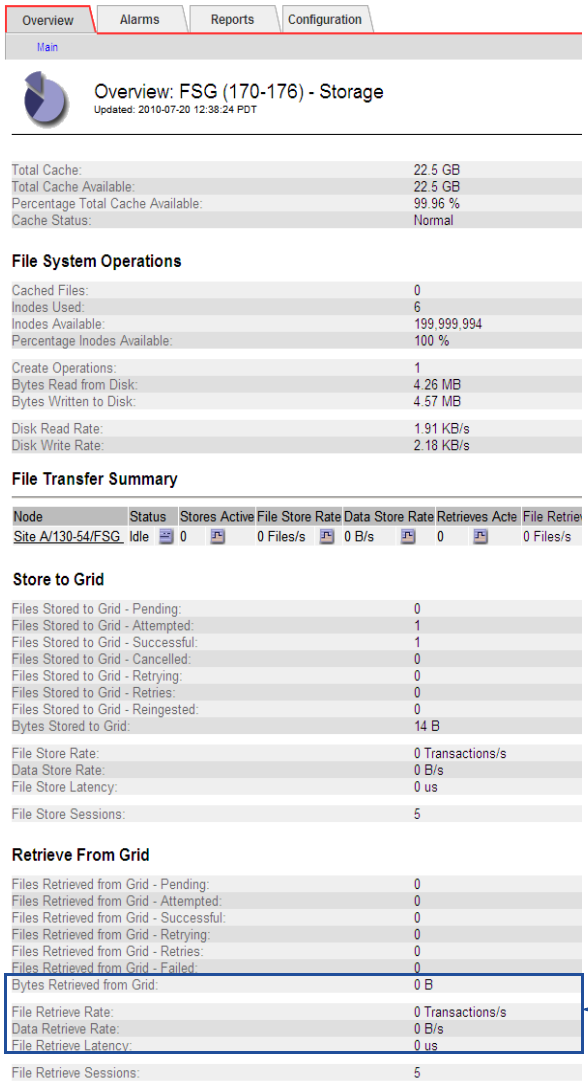

Figure 72: Retrieve Load Attributes

Content Replication and Retention

A grid's CMS services manage the replication of content (copies stored to the grid) by applying the grid's ILM policy. An ILM policy determines how many copies of an object are made, the appropriate placement, and the length of time each copy is retained. You can track what is happening with the ILM policy via the CMS ▶ Content ▶ Overview page by looking at the number of objects in each of the following categories:

Pending — Objects with ILM Evaluation Pending (ORpe)
The number of objects waiting to be processed.

Unachievable — Objects with Unachievable ILM Evaluations (ORun)
The number of objects that require additional copies to be made, but the storage resources specified by the ILM policy are unavailable to store the additional copies.

Future — Objects Marked for ILM Re-evaluation (ORde)
The number of objects that currently satisfy the ILM policy, but are due to be re-evaluated at a scheduled point in the future, for instance because of a rule that says "store a copy to archive media two years after ingest".

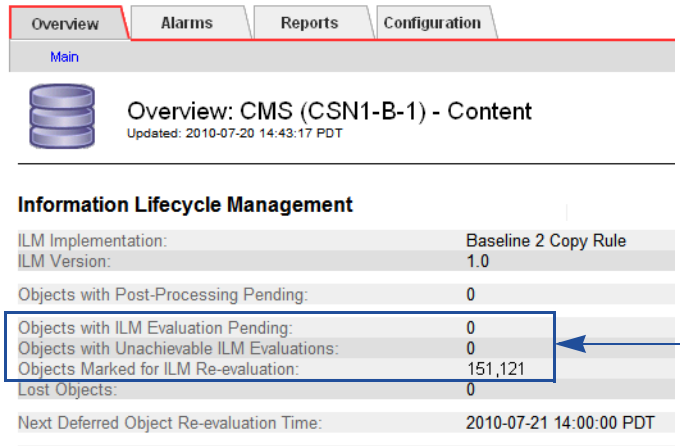

Figure 73: ILM Evaluation

Metadata Replication and Synchronization

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

NOTE Metadata synchronization is deprecated and no longer supported.

Metadata Replication by CMSs

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

Information about metadata replication is shown in the Metadata component of the CMS service (see Figure 74 below). The number of inter-CMS messages that are needed to manage metadata is drastically reduced when a grid uses metadata replication as compared to metadata synchronization (see below).

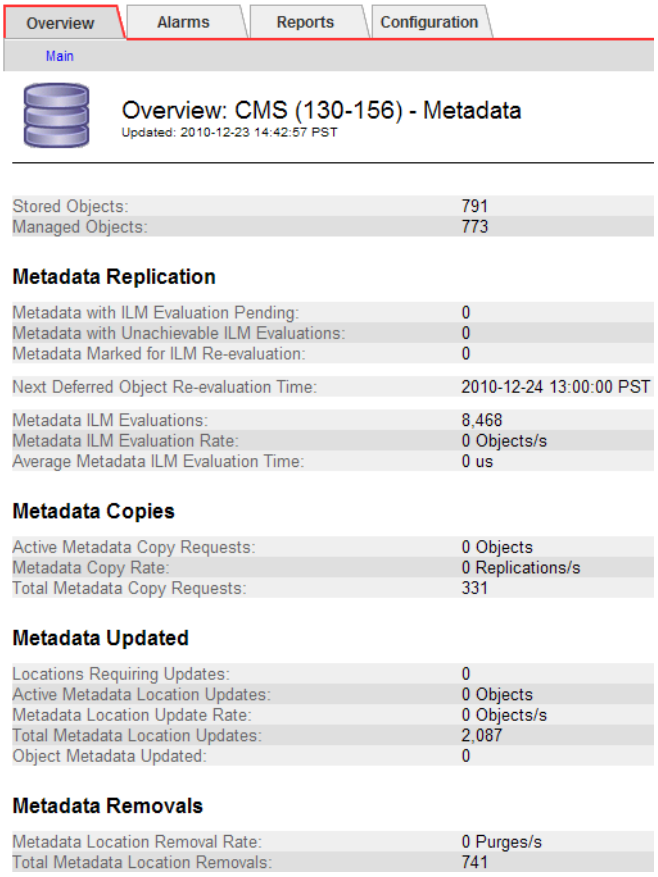

Figure 74: CMS Metadata Component for Metadata Replication

Metadata Synchronization by CMSs

NOTE Metadata synchronization is deprecated and no longer supported.

When an owner CMS service receives new metadata, it stores the metadata in its local database and sends synchronization messages to the other CMSs it must replicate metadata to. The attribute Queue Size tracks the number of messages to be sent to another CMS service. The corresponding attribute Incoming Messages tracks the number of synchronization messages coming from another CMS service and waiting to be processed (Figure 75).

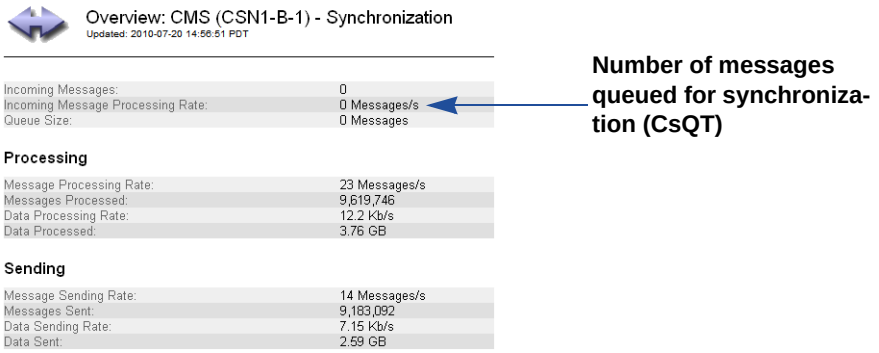

Figure 75: CMS Metadata Synchronization

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

FSG Replication

When an FSG ingests a file, it creates a file pointer to reference the object and it replicates the file pointer to the other FSGs in its replication group. To verify that FSG replication is proceeding normally, look at the attributes Operations Not Committed (SUOP), Operations Not Applied (SPOP), Files Pending for Replication (FRPP), and Replication Errors (RPER) in the FSG Replication component (see Figure 76 and Figure 77).

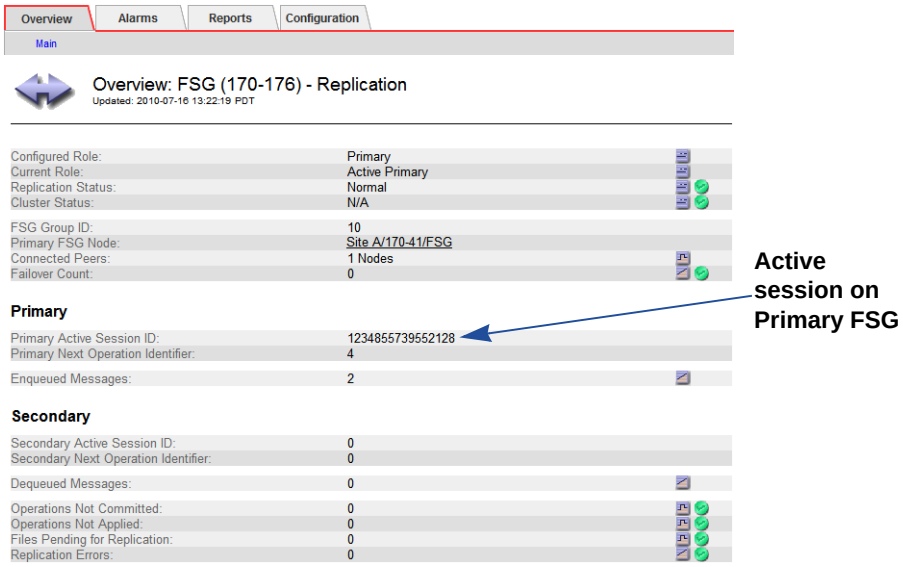

Figure 76: Primary FSG Replication Messages

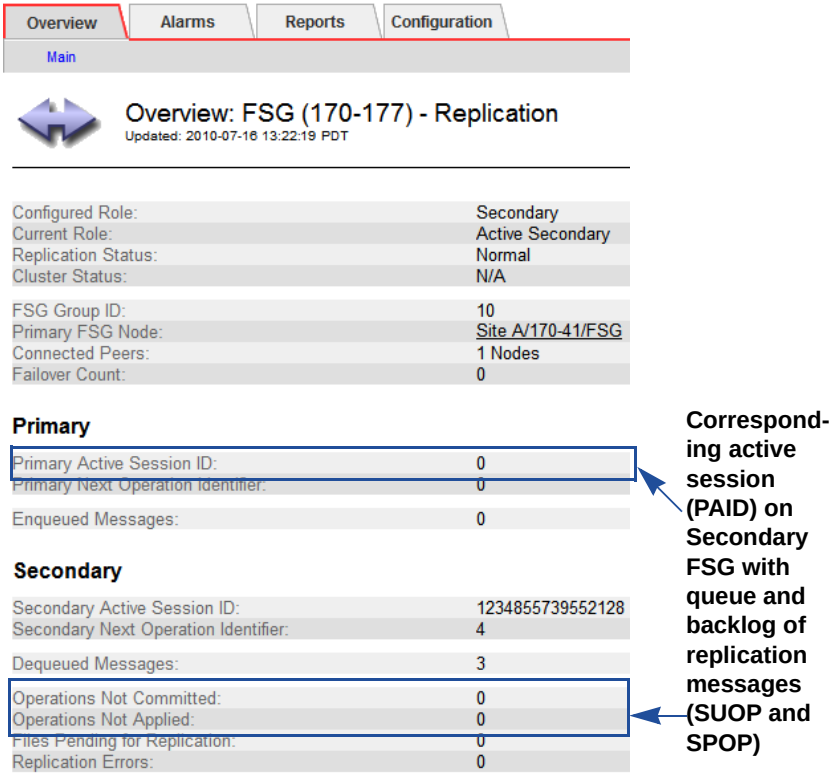

Figure 77: Secondary FSG Replication Messages

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

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

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

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

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

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

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

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

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

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

Gateway Node Failovers

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

Basic Replication Group

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

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

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

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

Following a failure, a manual recovery procedure must be performed to replicate files at risk, even if the failed FSG recovers automatically. For more details on recovery procedures, see the *Maintenance Guide*.

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

High Availability Gateway Replication Group

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

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

Investigate the cause of the failure as soon as possible as another FSG failure in the replication group will render FSG services unavailable. To restore the grid to full redundancy after the failover, fail back to the main primary manually and identify any files at risk of being lost. For more information on failovers and recovery procedures, see the *Maintenance Guide*.

FSG Backups

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

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

You can view information about the FSG backups (for example, backup schedule, backup duration, backup size, backup status) on the FSG Backup component of the Gateway Node (see Figure 78 below).

The value for Number of Files is the total number of files on the FSG managed file system as of the last backup. This includes files that are pending for ingest and files for which ingest into the grid is disabled through FSG profiles. For content information on a per share basis, see

FSG ▶ Shares. Per share data is updated at the time of the FSG backup. For more information, see "File Share Usage" on page 114.

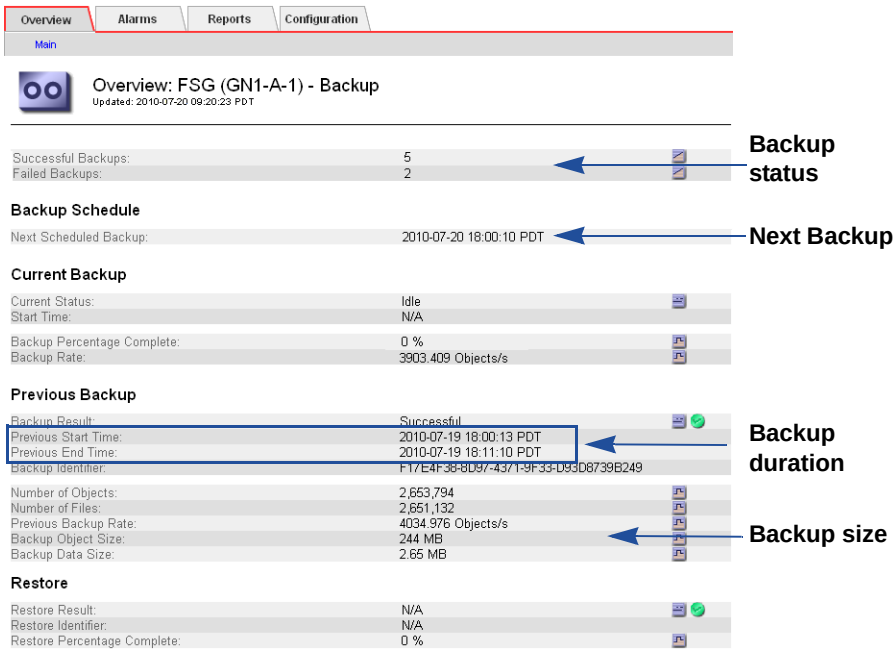

Figure 78: FSG Backup Attributes

For more information on backups, see the *Administrator Guide*.

LDR Verification

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

You can view information about LDR verification, for example progress and number of corrupt objects found, on the Verification component of the LDR (see Figure 79 below). Any corrupt object should be investigated.

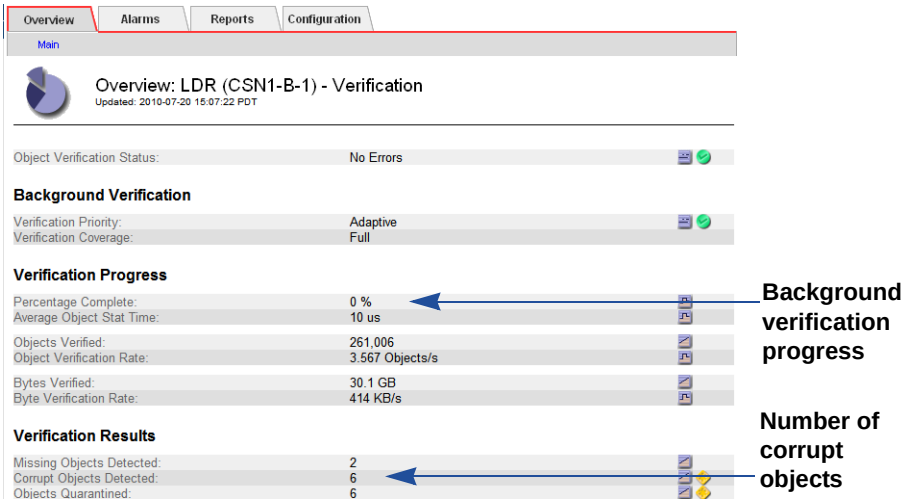

Figure 79: LDR Verification Attributes

There is another type of LDR verification called foreground verification. Foreground verification detects whether objects are missing. The LDR foreground verification procedure is initiated by a grid task and is used mostly during maintenance. See Table 20 for a comparison of the two types of LDR verification.

Table 20: LDR Verification

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

All attributes on Figure 79 above refer to LDR background verification except for Missing Objects Detected which is updated by foreground verification.

For more information on LDR verification, see the *Administrator Guide*.

Archive Node Capacity

Each Archive Node can interface with archival media managed by Tivoli Storage Manager (TSM).

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

NMS Database Usage Rates

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

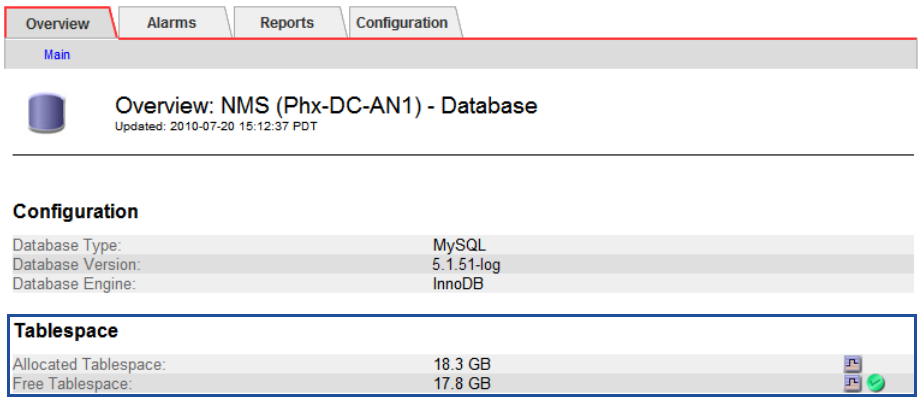

Figure 80: NMS Database

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

In the following example, database rate usage is a steady 0.3 GBs per day.

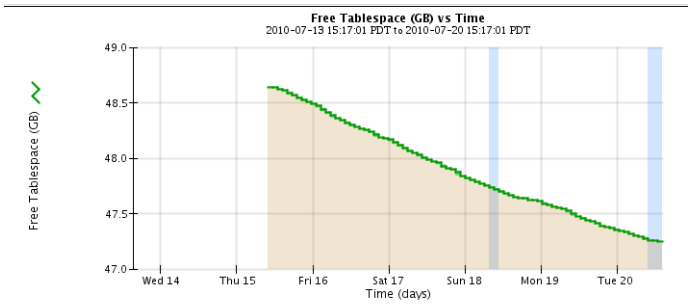

Figure 81: Free Tablespace Usage Rate

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

Grid Tasks

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

You can follow the progress of a grid task from the CMN Grid Tasks Overview tab (see Figure 82 and Table 21). Running grid tasks is restricted to accounts with Maintenance permissions such as the Admin and Vendor accounts.

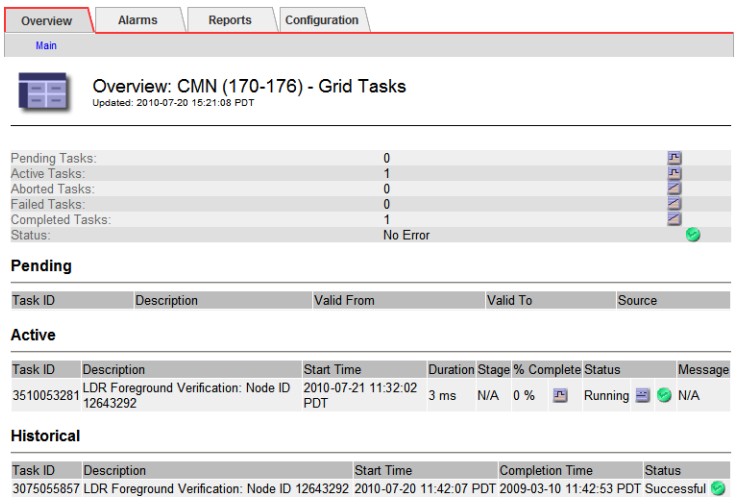

Figure 82: Grid Tasks Overview

Grid tasks go through three distinct phases:

Pending	The grid task has been submitted, but not yet started.
Active	The grid task has been started. It can be either actively running or temporarily paused.
Historical	An 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 21: 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.
Duration	Estimated amount of time since the grid task was started.
Stage	Description of the current stage of the active task.
% Complete	Progress indicator for active tasks.

Table 21: Grid Tasks Overview Fields (cont.)

Field	Description
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 • Rollback Failed • Expired • Aborted • Cancelled • Duplicate • Unauthorized • Invalid
Message	Information about the last stage of the active task.
Completion time	The date and time on which the grid task completed (or failed or expired or was aborted).

For more information on grid tasks, see the *Administrator Guide*.

Common Alarms

Table 22 below lists common alarms that are usually no cause for concerns as long as trends do not develop.

Table 22: Common Alarms

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

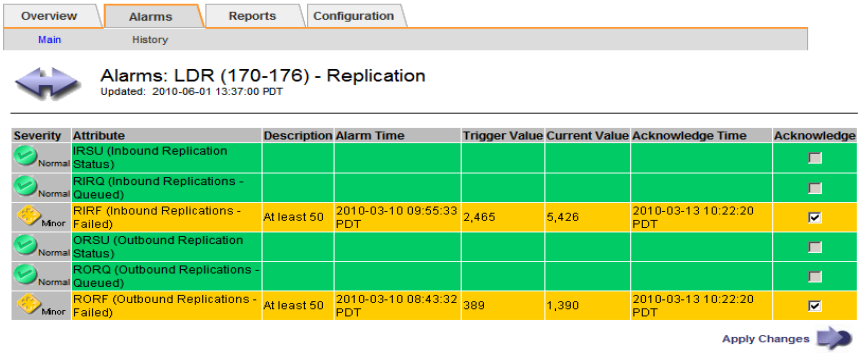

Figure 83: Common LDR Replication Alarms: RIRF & RORF

Table 22: Common Alarms (cont.)

Category	Code	Service	Notes

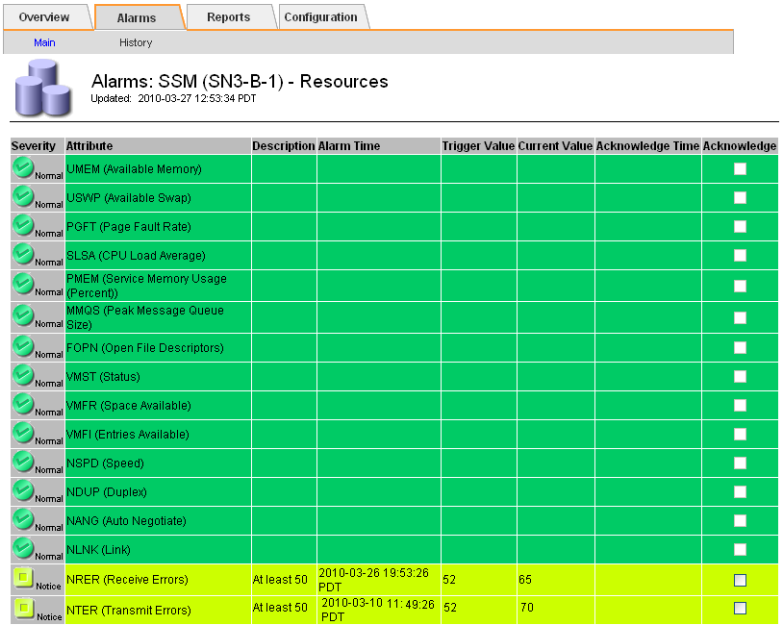

Figure 84: Common Network Alarms: NRER, NTER

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

Table 22: Common Alarms (cont.)

Category	Code	Service	Notes

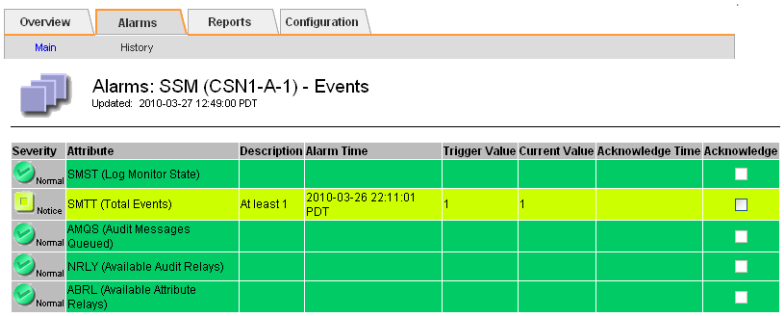

Figure 85: Common Events Alarm: Total Events

For more information on these and other alarms, see the *Troubleshooting Guide*.

Glossary

ACL Access control list—Specifies what users or groups of users are allowed to access an object and what operations are permitted, for example read, write, and execute.

active primary FSG In an HAGC, the FSG that is currently providing read-write service to clients. See also "FSG replication group".

ADC Administrative Domain Controller—A software component of the StorageGRID system. The ADC service maintains topology information, provides authentication services, and responds to queries from the LDR, CMS, CMN, 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 NetApp StorageGRID Software.

Admin Node A building block of the StorageGRID system. The Admin Node provides services for the web interface, grid configuration, and audit logs. See also "reporting Admin Node", "processing Admin Node", "primary Admin Node", "Audit Node", and "HCAC".

AMS Audit Management System—A software component of the StorageGRID system. The AMS service monitors and logs all audited system events and transactions to a text log file. The AMS service is found on the Admin Node—reporting Admin Node in a High Capacity Admin Cluster (HCAC) and the Audit 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.

API Gateway Node Application Programming Interface Gateway Node provides read-write access for HTTP clients (via StorageGRID API or CDMI). API Gateway Nodes are configured to include a "CLB" service, but not an "FSG" service. As a result, API Gateway Nodes do not support NFS/CIFS file systems and are not configured as part of a replication group.

ARC Archive—A software component of the StorageGRID system. The ARC service manages interactions with archiving middleware that controls nearline archival media devices such as tape libraries. The ARC service is found on the Archive Node.

Archive Node A building block of the StorageGRID system. The Archive Node manages storage of data to nearline data storage devices such as such as tape libraries (via IBM Tivoli® Storage Manager).

Audit Node A building block of the StorageGRID system. The Audit Node logs all audit system events. It is an optional grid node that is generally reserved for larger grid deployment.

audit message Information about an event occurring in the StorageGRID 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 StorageGRID API).

AutoYaST An automated version of the Linux installation and configuration tool YaST ("Yet another Setup Tool"), which is included as part of the SUSE Linux distribution.

BASE64 A standardized data encoding algorithm that enables 8-bit data to be converted into a format that uses a smaller character set, enabling it to safely pass through legacy systems that can only process basic (low order) ASCII text excluding control characters. See RFC 2045 for more details.

Basic Gateway replication group A Basic Gateway replication group contains a primary FSG and one or more secondary FSGs.

Binding The persistent assignment of a grid service (for example, an FSG or SSM) to the consolidated NMS service or processing NMS service. This assignment is based on grid topology (consolidated Admin Node or HCAC). See also "Admin Node".

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

business continuity failover A business continuity failover within a Gateway Node replication group is one where a secondary Gateway Node is manually configured to act as a primary after the primary Gateway Node fails. Clients can continue to read and write to the grid after they are manually redirected to the acting primary. This is a temporary measure to maintain service while the primary Gateway Node is repaired.

CBID Content Block Identifier—A unique internal identifier of a piece of content within the StorageGRID system.

CDMI Cloud Data Management Interface—An industry standard defined by SNIA that includes a RESTful interface for object storage. For more information, see http://www.snia.org/cdmi.

CIDR Classless Inter-Domain Routing—A notation used to compactly describe a subnet mask used to define a range of IP addresses. In CIDR notation, the subnet mask is expressed as an IP address in dotted decimal notation, followed by a slash and the number of bits in the subnet. For example, 192.0.2.0/24.

CIFS Common Internet File System—A file system protocol based on SMB (Server Message Block, developed by Microsoft) which coexists with protocols such as HTTP, FTP, and NFS.

CLB Connection Load Balancer—A software component of the StorageGRID system. The CLB service provides a gateway into the grid for clients connecting via the HTTP protocol. The CLB service is part of the Gateway Node.

Cloud Data Management Interface See "CDMI" on page 139.

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

CMS Content Management System—A software component of the StorageGRID system. The CMS service manages content metadata and content replication according to the rules specified by the ILM policy. The CMS service is found on the Control Node.

command In HTTP, an instruction in the request header such as GET, HEAD, DELETE, OPTIONS, POST, or PUT. Also known as an HTTP method.

container A container is a data structure used by the internals of grid software. In the StorageGRID API, an XML representation of a container is used to define queries or audit messages submitted using the POST command. Containers are used for information that has hierarchical relationships between components. The lowest-level component of a container is an atom. Containers may contain 0 to N atoms, and 0 to N other containers.

content block ID See "CBID".

content handle See "UUID".

consolidated Admin Node Admin Node hosting the consolidated NMS service. Can be the primary Admin Node.

consolidated NMS Hosted by the consolidated Admin Node. It is the equivalent of a combined reporting NMS and processing NMS service. See also "NMS".

Control Node A building block of the StorageGRID system. The Control Node provides services for managing content metadata and content replication.

CSTR Null-terminated, variable length string.

DC Data Center site.

deduplication If enabled, when the grid identifies two files as being identical, it "deduplicates" them by redirecting all content handles to point to a single stored instance of the file. The end result is that only the number of copies required by the ILM policy are stored in the grid. The feature was designed for use with applications that save two identical copies of a file to the grid via different Gateway Nodes.

NOTE Deduplication is deprecated and no longer supported.

distributed CMS A CMS that uses metadata replication. See also "metadata replication".

DR Disaster Recovery site.

EMR Electronic Medical Records—A computerized system for managing medical data that may be interfaced to the grid.

Enablement Layer The Enablement Layer for StorageGRID Software CD is used during installation to customize the Linux operating system installed on each grid server. Only the packages needed to support the services hosted on the server are retained, which minimizes the overall footprint occupied by the operating system and maximize the security of each grid node.

FCS Fixed Content Storage—a class of stored data where the data, once captured, is rarely changed and must be retained for long periods of time in its original form. Typically this includes images, documents,

and other data where alterations would reduce the value of the stored information.

Federated A "fully distributed" grid deployment topology that completely decentralizes site deployments. There is no DC or DR site. Data sharing and disaster recovery is achieved in a peer-to-peer manner by automatically distributing data to other Federated sites.

FSG File System Gateway—A software component of the StorageGRID system. The FSG service enables standard network file systems to interface with the grid. The FSG service is found on the Gateway Node.

FSG replication group A replication group is a group of FSGs that provide grid access to a specified set of clients. Within each replication group, there is a primary FSG (or a primary FSG cluster) and one or more secondary FSGs. The primary FSG allows clients read and write access to the grid, while storing file system information (file pointers) for all files saved to the grid. The secondary FSG "replicates" file system information, and backs up this information to the grid on a regular schedule.

Gateway Node A building block of the StorageGRID system. The Gateway Node provides connectivity services for NFS/CIFS file systems and the HTTP protocol.

Gateway Node replication group See "FSG replication group".

GDU Grid Deployment Utility—A StorageGRID software utility used to facilitate the installation and update of software on all grid nodes. GDU is installed and available on the primary Admin Node.

GPT Grid Provisioning Tool—a software tool included with StorageGRID software that permits you to provision a grid for installation, upgrade, maintenance, or expansion. GPT creates and maintains an encrypted "repository" of information about the grid that is required to maintain the grid and recover failed grid nodes.

Grid ID signed text block A BASE64 encoded block of cryptographically signed data that contains the grid ID which must match the grid ID (gid) element in the grid specification file. See also "provisioning".

grid node The name of the StorageGRID system building blocks, for example Admin Node or Control Node. Each type of grid node consists of a set of services running on a server.

Grid Specification File An XML file that provides a complete technical description of a specific grid deployment. It describes the grid topology, and specifies the hardware, grid options, server names, network settings, time synchronization, and gateway clusters included in the grid deployment. The Deployment Grid Specification file is used to generate the files needed to install the grid.

Grid Task A managed sequence of actions that are coordinated across a grid to perform a specific function (such as adding new node certificates). Grid Tasks are typically long-term operations that span many entities within the grid. See also "Task Signed Text Block".

HAGC High Availability Gateway Cluster—An HAGC is a primary gateway cluster that consists of a main FSG and a supplementary FSG. A high availability gateway replication group optionally includes one or more secondary FSGs.

HCAC High Capacity Admin Cluster—An HCAC is the clustering of a reporting Admin Node and processing Admin Node. The result is an increase to a grid's capacity for grid services and thus grid nodes. See also "reporting Admin Node", "processing Admin Node", and "Admin Node".

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, a 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.

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 rack.

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

LUN See "object store".

main primary FSG In an HAGC, the FSG that is configured to be the active primary FSG by default.

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

metadata replication In a grid that uses metadata replication, a CMS makes copies of metadata on the subset of CMSs that are in its CMS replication group, and then applies the grid's ILM policy to content metadata. In the NMS MI, CMSs that use metadata replication display the Metadata component. Called "distributed CMS" in a previous release.

metadata synchronization In a grid that uses metadata synchronization, a CMS synchronizes metadata with all other read-write CMSs in the grid. Called "synchronized CMS" in a previous release.

NOTE Metadata synchronization is deprecated.

MI Management Interface—The web-based interface for managing and monitoring the StorageGRID system provided by the NMS software component. See also "NMS".

namespace A set whose elements are unique names. There is no guarantee that a name in one namespace is not repeated in a different namespace.

nearline A term describing data storage that is neither "online" (implying that it is instantly available like spinning disk) nor "offline" (which could include offsite storage media). An example of a nearline data storage location is a tape that is loaded in a tape library, but is not necessarily mounted.

NFS Network File System—A protocol (developed by SUN Microsystems) that enables access to network files as if they were on local disks.

NMS Network Management System—A software component of the StorageGRID system. The NMS service provides a web-based interface for managing and monitoring the StorageGRID system. The NMS service is found on the Admin Node (both the reporting and processing Admin Nodes in an HCAC). There are three types of NMS service: consolidated, reporting, and processing. See also "MI" and "Admin Node".

node ID An identification number assigned to a grid service within the StorageGRID 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.

object segmentation A StorageGRID process that splits a large object into a collection of small objects (segments) and creates a segment container to track the collection. The segment container contains the UUID for the collection of small objects as well as the header information for each small object in the collection. All of the small objects in the collection are the same size. See also "segment container".

OID Object Identifier—The unique identifier of an object.

primary Admin Node Admin Node that hosts the CMN service. There is one per grid. In an HCAC, the CMN service is hosted by the primary reporting Admin Node. See also "Admin Node" and "HCAC".

primary FSG In an FSG replication group, the FSG that provides read-write services to clients. See also "FSG replication group".

processing Admin Node Performs attribute and configuration processing that is passed on to the reporting Admin Node as part of a High Capacity Admin Cluster. See also "reporting Admin Node" and "HCAC".

processing NMS Hosted by the processing Admin Node. Provides attribute and data processing functionality. Only operates in conjunction with a reporting Admin Node and the reporting NMS. See also "NMS".

provisioning The process of editing the Grid Specification File (if required) and generating a new or updated SAID package and GPT repository. This is done on the primary Admin Node using the provision command. The new or updated SAID package is saved to the Provisioning Media. See also "Grid Specification File" and "SAID".

quorum A simple majority: 50% + 1 of the total number in the grid. In StorageGRID software, some functionality may require a quorum of the total number of some types of service to be available.

reporting Admin Node Reports attribute and configuration information to web clients as part of a High Capacity Admin Cluster. See also "processing Admin Node" and "HCAC".

reporting NMS Hosted by the reporting Admin Node. Reports status information about the grid and provides a browser-based interface. Only operates in conjunction with a processing Admin Node and the processing NMS. See also "NMS".

SAID Software Activation and Integration Data—Generated during provisioning, the SAID package contains site-specific files and software needed to install a grid.

Samba A free suite of programs which implement the Server Message Block (SMB) protocol. Allows files and printers on the host operating system to be shared with other clients. For example, instead of using telnet to log in to 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.

SATA Serial Advanced Technology Attachment—A connection technology used to connect servers and storage devices.

SCSI Small Computer System Interface—A connection technology used to connect servers and peripheral devices such as storage systems.

secondary FSG A read-only FSG that may also perform backups of the FSG replication group. See also "FSG replication group".

security partition If enabled, access to content ingested into the grid is restricted to the application, HTTP client, or FSG replication group that ingested the object.

segment container An object created by StorageGRID during the segmentation process. Object segmentation splits a large object into a collection of small objects (segments) and creates a segment container to track the collection. A segment container contains the UUID for the collection of segmented objects as well as the header information for each segment in the collection. When assembled, the collection of segments creates the original object. See also "object segmentation".

server Used when referring specifically to hardware.

Server Manager Application that runs on all grid servers, supervises the starting and stopping of grid services, and monitors all grid services on the server.

service A unit of the StorageGRID software such as the ADC, CMS or SSM.

SGAPI StorageGRID Application Programming Interface—A set of commands and functions, and their related syntax, that provides HTTP clients with the ability to connect directly to the StorageGRID system (to store and retrieve objects) without the need for a Gateway Node.

SLES SUSE Linux Enterprise Server—A commercial distribution of the SUSE Linux operating system, used with the StorageGRID system.

SQL Structured Query Language— An industry standard interface language for managing relational databases. An SQL database is one that supports the SQL interface.

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 StorageGRID software that monitors hardware conditions and reports to the NMS. Every server in the grid runs an instance of the SSM. The SSMS service is present on all grid nodes.

SSL Secure Socket Layer—The original cryptographic protocol used to enable secure communications over the internet. See also "TLS".

standby primary FSG In an HAGC, the FSG that is available to take over and provide read-write services to clients in event of the failure of the active primary FSG.

Storage Node A building block of the StorageGRID system. The Storage Node provides storage capacity and services to store, move, verify, and retrieve objects stored on disks.

StorageGRID® A registered trademark of NetApp Inc. for their fixed-content storage grid architecture and software system.

StorageGRID API See "SGAPI".

storage volume See "object store".

supplementary primary FSG In an HAGC, the FSG that is configured to be the standby primary FSG by default.

SUSE See "SLES"—SUSE Linux Enterprise Server.

synchronized CMS See "metadata synchronization".

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 StorageGRID. 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.

VM Virtual Machine—A software platform that enables the installation of an operating system and software, substituting for a physical server and permitting the sharing of physical server resources amongst several virtual "servers".

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".

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

Q

R

S

T

U

V

W

X