HP Medical Archive solution

Software version: 8.0.5

software release notes

Document release date: October 2009
Software release date: October 2009

Legal notices

Warranty

The only warranties for HP products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. HP shall not be liable for technical or editorial errors or omissions contained herein.

The information contained herein is subject to change without notice.

Restricted rights legend

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

Licensing

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

Copyright notices

Trademark notices

Adobe® and Acrobat® are trademarks of Adobe Systems Incorporated.

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

UNIX® is a registered trademark of The Open Group.

Contents

1 Software release notes

Release: 8.0.5

Introduction

These release notes are intended to provide important details related to late changes and known issues that were not captured in product documentation.

Please read these notes carefully and keep on hand when installing, updating, or maintaining the HP Medical Archive solution software.

This release introduces the HP Medical Archive solution version 4.0 with the following interconnected components or nodes:

- AP733A HP MAS V4 Admin/Gateway Node
- AP734A HP MAS V4 GN2000 Gateway Node
- AP735A HP MAS V4 Control Node
- AP736A HP MAS V4 Base Storage Node, 0TB
- AP737A HP MAS V4 SAS MDL 4 TB Capacity Kit
- AP738A HP MAS V4 SAS ENT 2.25 TB Capacity Kit
- AP739A HP MAS V4 Base Control Storage Node, 0TB (#0D1)
- AP739A HP MAS V4 Ctrl Storage 4TB SAS MDL Node (#101)
- AP739A HP MAS V4 Ctrl Storage 8TB SAS MDL Node (#102)
- AP739A HP MAS V4 Ctrl Stor 3.6TB SAS ENT Node (#201)
- AP 740A HP MAS V4 Foundation Node, 3TB
- AP742A HP MAS V4 Console Bundle
- AP743A HP MAS V4 Network Bundle
- AP744A HP MAS V4 Virtual Node, 2 DR GN VMs[1]
- AP745A HP MAS V4 Virtual Expansion Kit, 2 DR GN VMs[1]

1. Limited functionality, see Virtualized Hardware (page 7).

Release Information

The content of this edition is for the HP MAS Software Release 8.0.5 (as reported on the Software CD, the Server Manager interface, and the **SSM > Services** component) which comprises the HP Medical Archive solution version 4.0.

Fundamental features supported in HP Medical Archive solution 4.0:

- Denser storage capacity per node to enable petabyte scalability.
- Lower cost entry point usability as the base platform for both small and large deployment.

Software Upgrade Path

Updates from Releases 7.5.x are supported.

Customer Sites

Customer sites running previous releases can be updated to this release using the procedure described in the *Software Update Guide*.

SLES 10 Service Pack 2 Support

The HP Proliant DL180G6 server grid node requires the installation of SLES 10 SP2. This revision is supported in the DL180G6 server platform only.

New Features for Release 8.0

Distributed CMS

In earlier releases of HP MAS, all content metadata is replicated to all CMSs: the CMS databases are synchronized. With this release, it is possible to replicate content metadata to a subset of CMSs instead of to all CMSs. Distributed CMS operation improves scalability and performance and allows islanded operations with metadata privacy protection at the satellite sites.

Distributed CMS can only be enabled on new installations. Upgrading a grid to Release 8.0 does not change the type of CMS used in the grid. Conversion to distributed CMS operation is not currently supported for existing grids.

Distributed CMS and DICOM

DICOM and distributed CMS cannot be used together in Release 8.0.

Distributed CMS and Deduplication

Deduplication and distributed CMS cannot be used together in Release 8.0.

HTTPS Protocol and Security Partitioning

As of Release 8.0, all grids only support the HTTPS protocol to connect the Gateway Nodes to the grid. The HTTPS protocol improves security and is required for new functionality such as security partitioning. If enabled, security partitioning isolates content ingested in each Gateway Node replication group such that data ingested via a Gateway Node can only be retrieved via a Gateway Node in the same replication group and not via custom applications that use the HTTP API.

Simplified Grid Specification

In previous releases, customers had to make many detailed decisions about their grid configuration at the time their grid was ordered. Many of these configuration choices can now be made and set at installation time. After the grid is provisioned and installed, the most current revision of the grid specification file can be viewed from the NMS interface.

Virtualized Hardware

As of Release 8.0, you can install up to four secondary Gateway Nodes on a single physical server that is installed with VMware ESX Server 3.5. For each secondary Gateway Node, a virtual machine is created that runs as a virtual hardware server. When added to a grid, each secondary Gateway Node running on a virtual machine is recognized by the grid as a separate instance of a secondary Gateway Node. Refer to the *Managing Virtualized Hardware Guide* for hardware requirements.

NOTE Virtual Secondary Gateway Nodes (SKU # AP744A and AP745A) offer FSG backup functionality only. They do not provide read-only client access to shares.

iSCSI Support

Storage Nodes now support the iSCSI protocol for connecting to NetApp machines and HP StorageWorks Enterprise Virtual Arrays (EVA). For more information, see the *Installation Guide*.

Improved Features for Release 8.0

This release includes the following improvements to existing features.

New Account Permissions

The set of account permissions used with Group Accounts have changed with Release 8.0. Table 1 describes the allowable tasks for the four sets of permissions.

Table 1 User Groups Permissions

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

New Group Account

A new Group Account called “User” has been added. The User Group Account has the “Acknowledge Alarms” permission.

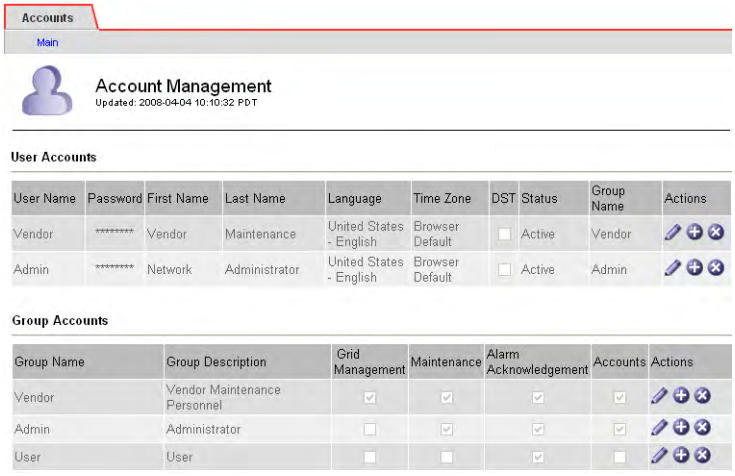

Figure 1 Account Management

ILM Configuration Usability Improvements

A number of changes have been made to the ILM (Information Lifecycle Management) user interface in Release 8.0 to improve usability.

Grid Tasks

As of Release 8.0, grid tasks required for expansion, maintenance or update procedures appear automatically in the Pending list as part of the provisioning process. They are no longer provided in the form of a Task Signed Text Block that must be submitted manually.

Deduplication and Configurable ILM

As of Release 8.0, deduplication can be used with configurable ILM.

FSG Start up when HTTP Services Unavailable

As of Release 8.0, FSGs can start automatically when HTTP services are unavailable. In earlier releases, the FSG does not come online until at least one HTTP Query/Retrieve and one HTTP Ingest service provider is discovered. Since the FSG can provide limited services when no LDRs are available (retrieval of cached content and staging of new content up to watermark levels), this restriction was removed.

ILM Evaluation Pending Calculations

As of Release 8.0, the value of the attribute ILM Evaluation Pending in **CMS > Content** includes the number of active replications (as reported by the attribute **CMS > Content > Active Replications**).

Server Restore Procedure for Custom Nodes

The procedures to restore failed servers that host more than one grid node, for example Control/Storage Nodes or Gateway/Control/Storage Nodes have been improved.

E-Mail Notifications

As of Release 8.0, the symbol + is supported in email addresses.

Documentation Reorganization

The *Grid Primer* is a new manual that provides an introduction to the HP MAS product. The guide is written for grid operators whose main function is to monitor the grid on a day-to-day basis, in particular users who log in to the NMS web interface using User level permissions.

The *NMS User Guide* has been renamed the *Administrator Guide*. The information on how to use the NMS web interface information has been moved to the *Grid Primer*. Grid-wide procedures previously covered in the *Installation Guide* and the *Maintenance Guide* (for example, configuring CIFS file shares, configuring DNS, IP addresses and NTP time servers, and configuring grid options such as storage compression and deduplication) are now consolidated in the *Administrator Guide*.

Operation Notes

Networking

Restarting Networking Takes Time with SLES 10 SP2

Restarting networking takes significantly longer with SLES 10 SP2 than with SLES 10 SP1.

Hardware Monitoring

Raid Controller Model Number is Misreported

The RAID Controller Model Number is not reported correctly for DL180G6 servers on the **SSM > RAID > Overview** page of the NMS Management Interface. P212 is reported as being “544” and P410 is reported as “580”.

Tape Nodes

Multiple Copies Made to Tape

A temporary loss of communications can cause multiple copies of an object to be made by the Tape Node. If the Tape Node becomes isolated from all Control Nodes while writing a volume to tape, it can not inform the Control Node that it has successfully completed its task. As a result, when the communication link is restored the Control Node instructs the Tape Node to make the copy again, and the object is stored to tape twice.

Audit Messages

Audit File Naming Convention

Once a day, the active audit log is closed and saved to an archived log file named YYYY-MM-DD.txt where the date stamp in the file name indicates when the file was archived. The date stamp does not always match the date of the majority of messages logged. In addition, audit log files may contain up to 45 minutes of audit messages for the following day.

ARCE Audit Message for Bulk Retrieve Does Not Include VLID

The VLID (Volume ID) field in ARCE (ARC Retrieve End) audit messages generated from a retrieve initiated through the bulk retrieve service is set to zero instead of the correct volume ID (VOLI) from the CMS metadata.

Control Nodes

Stopping Server Manager May Not Stop mysql

Occasionally, stopping Server Manager does not stop mysql. To stop mysql, enter:

```
mysqladmin --defaults-file=/etc/my.default.cnf shutdown
```

False Report of Lost Objects

Under certain conditions, objects may be reported as lost if they exist on a Storage Node but there is no corresponding metadata on Control Nodes. This could occur for instance if an ingest was aborted after the object was created on the Storage Node but before it was committed to a Control Node. Because the ingest was aborted, the object is not expected to be in the grid.

To determine whether the object is actually lost, use the cqt tool.

No lost object

If the number of locations is 0 and there is no metadata, there is no lost object. You can ignore this false report.

For example:

```
cn2-b-1:~ # /usr/local/cms/bin/cqt --cbid 6FCF94AA6AAB73EE
13020115:6FCF94AA6AAB73EE:owner:13020115
13020115:6FCF94AA6AAB73EE:locations:0
```

Lost object

If the number of locations is 0 and there *is* metadata, the object is lost. Contact Support.

For example:

```
cn2-b-1:~ # /usr/local/cms/bin/cqt --cbid 6FCF94AA6AAB73EE
13020115:6FCF94AA6AAB73EE:owner:13020115
13020115:6FCF94AA6AAB73EE:locations:0
13020115:6FCF94AA6AAB73EE:metadata:APPM:*ctp=application/
octet-stream
13020115:6FCF94AA6AAB73EE:metadata:APPM:CTIM=1158151129
13020115:6FCF94AA6AAB73EE:metadata:APPM:FGID=65534
```

Content May Be Lost In A Two-CMS Grid After One CMS Becomes Read-Only

Content may be lost in a grid with only two CMSs if a CMS becomes read-only before the other CMS becomes read-only, and the grid has not yet been expanded to add more Control Nodes. If the read-only CMS with more content ever fails in the future, the content that it knew about will get lost as cloning the database will not restore that content.

CMS Turns Gray Automatically

CMS turns gray automatically (without the intervention of a system administrator) when instructed by Server Manager to shut down due to a mysql problem. If this happens, a minor SSM alarm is triggered to report that the CMS service is not running on that server.

Gateway Nodes

Error Retrieving Content from Secondary FSG

If a file is overwritten on a Primary Gateway Node, the file's content handle is released and a "release" replication message is sent to the Secondary Gateway Nodes in the same replication group. If a client tries to retrieve this file on the secondary before the "release" replication message has been processed, the client receives an error message.

High Availability Gateway Node Failover

A cluster may not fail over to the Standby Gateway Node if the system drive fails on the Active Gateway Node, as the heartbeat service may not detect the failure.

Failover from the Active FSG to the Standby FSG in a HAGC cluster may fail if the Standby FSG is processing a large backlog of replication messages. The Standby FSG will remain in the Standby state and Cluster Status will change to Failed until the Standby FSG has processed the backlog of replication messages, at which point the FSG will transition to the Active state and the Cluster status will change to Vulnerable.

Files May Be Swapped Out in Different Order

Under some circumstances, files may be swapped out of the Gateway Node cache in order of “least recently ingested/restored to cache” rather than “least recently accessed”. This occurs because the XFS file system used on Gateway Nodes sometimes does not update the access time of files.

File Create Operation Can Cause Assertion

The FSG service asserts when it processes a file create replication message if the file path in question already exists as a directory. This situation was observed to occur while recovering from a server failure on a Primary Gateway Node that was integrated with McKesson Horizon Medical Imaging (which uses a custom sequence of file system operations).

Health Check Timeout

Swapout of large, heavily fragmented files may cause an FSG health check timeout. In a clustered FSG replication group, this may trigger automatic failover to the standby FSG. In an unclustered replication group, client services may be temporarily disabled until the swapout operation completes and the FSG automatically recovers.

FSG Groups Not Updated in NMS Tree After Expansion

After an expansion, the new FSG replication groups are not updated in the tree. To display the groups properly:

Refresh the tree region/page.

—or—

Log out of the NMS and log in again.

FSG Groups Do Not Appear in NMS

If bundles are imported (for example, through a grid task or the CMN console) while the NMS management interface is starting up, the NMS may not subscribe to the FSG Management bundles for each FSG group. If this happens, some FSG groups will not show up under **Grid Management > FSG Management**, and changes to any of the FSG Management configuration options (Settings, Options, Profiles) are not reflected in the NMS after you click **Apply Changes**. If this is the case, contact HP Support.

FSG Management Configuration Settings Are Not Saved

See FSG Groups Do Not Appear in NMS (page 14).

Failure of an Unclustered FSG May Leave No Primary

If the Primary FSG in an unclustered replication group fails due to an issue with the heartbeat service, the replication group stops providing read/write service to clients and all Gateway Nodes display the “no Primary” error on Replication > Replication Status. Fail over to the Secondary Gateway Node to continue, or restart the Primary Gateway Node to resolve the issue.

ILM Management

No filter Name in Built-in Rule Purge FSG Backups on Content Handle Release

The built-in rule Purge FSG Backups on Content Handle Release does not include a filter name. As of Release 8.0, filter names are mandatory. If you need to edit this rule, add a filter name before you click **Apply Changes**.

Cannot Save Changes If Deleting Metadata Items from Filter in a Saved Rule

If you delete metadata items from a filter in a previously saved rule, the change will not be saved. Changes related to adding metadata items or modifying existing metadata items are saved as well as changes to the placement instructions. If you need to delete metadata from a filter, you must create a new rule.

Rules Missing from Active Policy if Policy is not Saved Prior to Activate

When you add rules to a saved proposed policy, both **Activate** and **Save** are enabled. If you click **Activate** without first saving the proposed policy, the rules that you have added will not be in included in the active policy.

NMS

MI Fails to Restart Following Power Failure

If the NMS MI fails to restart following a power failure, restart Server Manager. Enter:

```
/etc/init.d/servermanager restart
```

500: Internal Server Error

If you make a change in the NMS MI and press **<Enter>** instead of clicking **Apply Changes**, the 500: Internal Server Error is generated. To resolve the issue, go back to the page, enter the value again, and wait for the Apply Changes button to appear. Then, click **Apply Changes**.

Grid Task % Completed Value Reported Incorrectly

The % complete value for active grid tasks is sometimes reported incorrectly.

Time Display

Unless a user account is configured with a time zone and daylight savings time (DST) is selected in the account description, the Network Management System interface uses data from the host Windows® computer to determine the local time zone of the user. Some anomalies have been identified that can prevent the time zone from being correctly determined from Windows.

In Newfoundland Canada (GMT-03:30) and Iran (GMT+03:30), the time zone may not be correctly determined if the Windows system does not have the "Automatically adjust clock for daylight savings changes" option selected. Should the information from the browser be invalid, the NMS uses GMT.

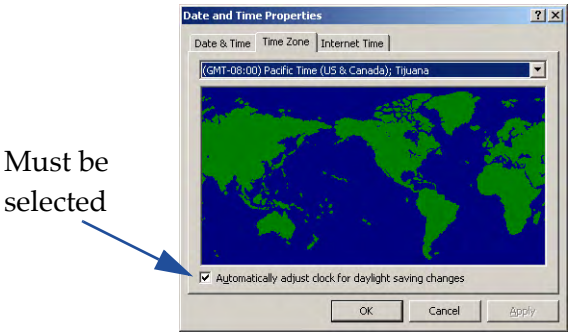

The option is selected by default for new Windows installations.

Also note that if a change to the local time zone selection is made in Windows, the browser must be closed and restarted to correctly pass the time zone change to the NMS interface. It is not sufficient to log out of the NMS interface and then log back in.

Attributes Not Persistent for the CMS

The following attributes (shown on **SSM > Resources**) are persistent across restarts on all services except the CMS.

- Service CPU Seconds
- Transmitted Bytes
- Received Bytes

SSM Updates Write Cache Status Only on Start

The Write Cache status shown in the Volumes table of the **SSM > Resources > Overview** page is only updated when the SSM is started or restarted; it is not dynamically updated when the status of the Write Cache changes.

Summary Attributes May Not Go Indeterminate

There are two sets of summary attributes: the Storage Capacity attributes and the Metadata Capacity attributes. The Storage Capacity attributes show as indeterminate (blue) in the NMS if any of the LDR services in the relevant group are indeterminate. The Metadata Capacity attributes show as indeterminate if any of the CMS services in the relevant group are indeterminate. However, if the Storage Capacity attributes already display as indeterminate in the NMS, the Metadata Capacity attributes will not switch to indeterminate if the relevant CMS services are affected (and vice versa).

Summary attributes Remain Indeterminate When All Services Are Online

Some of the summary attributes may appear blue in the NMS even though all of the services are online without alarms. Summary attributes are updated whenever there is a change for any of the attributes that are used for the summary calculation. If the connection to the NMS is lost, the summary attributes are set to indeterminate (blue). As the values of the related attributes value change, the summary attribute are recalculated and their color will change to black. Attributes whose value does not change frequently may remain in the indeterminate state for a long time.

Drop Downs Rendered On Top Of Help

In Internet Explorer 6 (IE), if you open a drop down menu to display its values while the help text for an attribute is displayed, the drop down menu is rendered on top of the help if they happen to overlap on the screen. This issue does not occur with IE7.

Rows May Not Be Saved in Correct Order

Always check the final result after you change the position of rows in a table, particularly when you change the order of rows back and forth before selecting Apply Changes. In some cases, what is saved may not be the final configuration of rows.

Items in Grid Management Not Highlighted

When you select elements of the Grid Topology (such as a site or server), they are highlighted in the navigation tree. However, when you select items in the Grid Management menu (such as Account Management or FSG Management), they are not highlighted. The highlight remains on the most-recently-selected item in the Grid Topology.

Browser Back Button Displays Wrong Content

If you click **Back** while editing a page under the Grid Management menu (such as Account Management), the wrong content may be displayed. Use the controls in the NMS interface to navigate.

NMS Does Not Validate DICOM AE Titles

The NMS accepts DICOM AE Titles that do not meet the DICOM standard (which states that AE titles must have 16 or fewer ASCII characters, excluding "\".)

Regular Expressions are Not Validated

When configuring an FSG Profile (in **FSG Management > Group X > Configuration > Profiles**), you must enter a regular expression that matches the files to which you want the profile to apply. The NMS does not perform any error checking on the validity of these regular expressions.

Paginated Pages Do Not Auto-Refresh

NMS Pages that are paginated, for example the System Status page, do not auto-refresh. You must refresh the browser manually.

QLogic Events in SSM Events

QLogic events do not show up in **SSM > Events**. However, the QLogic events are captured in the log files.

Storage Nodes

Understanding Storage Capacity Attributes

The attribute **LDR > Storage > Total Space Available** is designed to give you an estimate of how much additional data can be stored to a Storage Node. Understanding how a Storage Node stores data and how this attribute value is calculated can give you a better understanding of when this attribute may appear inaccurate.

As each object is ingested into the grid, the Control Node assigns it a unique identifier called a CBID. Storage volumes are configured such that each object store (storage volume) is associated with a range of CBID values, with larger storage volumes being associated with larger CBID ranges so that storage volumes need not be equal in size. When an object is stored to an LDR it is saved to a particular storage volume based on its CBID value. Because CBIDs are randomly generated, objects are assigned relatively evenly across object stores.

After the first storage volume fills, the Storage Node as a whole cannot accept more data because it can no longer store objects whose CBIDs fall in the full volume's assigned range. The attribute **LDR > Storage > Total Space Available** is therefore calculated as (Storage on volume with least remaining space) X (Number of volumes). This estimate is most accurate when object size is small compared to the storage volume size and is relatively consistent over the lifetime of the Storage Node. There are cases where Total Space Available does not give an accurate assessment. For example, early in the lifetime of a Storage Node, if object size is relatively large compared to storage volume size and by chance many objects in a row are stored to the same storage volume, the estimate

represented by Total Space Available may be much less than the amount of storage that is actually available and usable. Total Space Available is a conservative estimate. In this case, adding up the available space on each object store (using the information in the **LDR > Storage > Object Stores** table) would give a more reliable estimate.

Customizing Alarms for Full Read-Only LDRs

A number of alarms are triggered when a Storage Node becomes read-only because it is full. To remove these alarms once the situation has been investigated, follow this procedure.

1. Log in to the NMS using the Vendor account.
2. Change the desired state of the LDR to read-only.
   a. Go to **<Storage Node> > LDR > Storage > Configuration**.
   b. Select **Read-Only** in the **Storage State - Desired list**.
   c. Click **Apply Changes**.

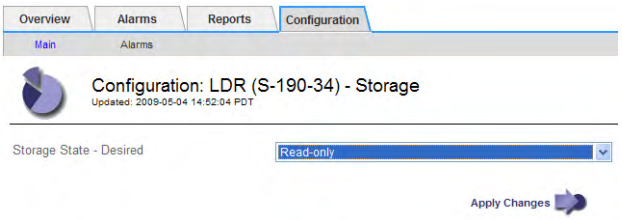

Figure 2 Change the Desired Storage State to Read-only

3. Disable the following Notice alarms using the procedure documented in the *Administration Guide*:

Path	Notice Alarm
LDR > Storage	SAVP (Total Space Available (Percent))
LDR > Replication	RPTE (Object Replication State) “Outbound Only”
LDR > Replication	RPTU (Object Replication Status) “Waiting for Storage...”
LDR > HTTP	HSTE (HTTP State) “Online (Read-Only)”

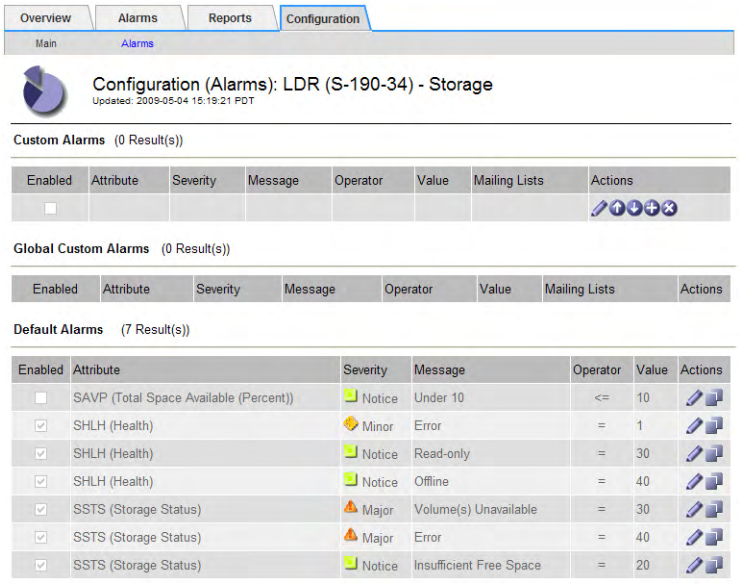

Figure 3 SAVP Alarm Disabled

4 Acknowledge the following Notice alarms using the procedure documented in the *Administration Guide*:

Path	Notice Alarm
LDR > HTTP	HSTE (HTTP State)
LDR > DICOM	MSTE (DICOM State)
LDR > Replication	RPTE (Object Replication State) “Outbound Only”
LDR > Replication	RPTU (Object Replication Status) Waiting for Storage...”

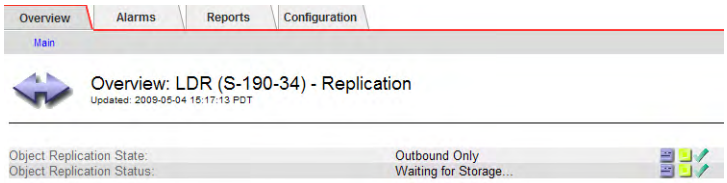

Figure 4 RPTE and RPTU LDR Replication Alarms Acknowledged

5 Reset the HTTP and DICOM failure counts.

a Go to **<Storage Node> > LDR > HTTP > Configuration**.

b Select **Reset HTTP Counts**.

c Click **Apply Changes**.

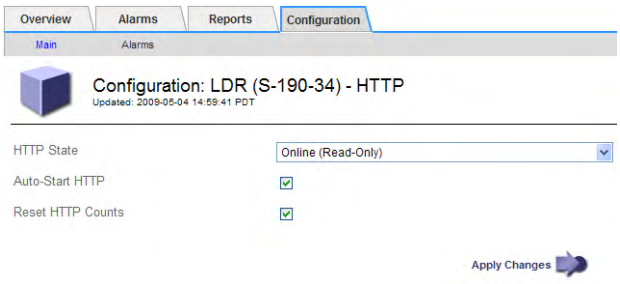

Figure 5 Reset the HTTP Failure Count

d Go to **<Storage Node> > LDR > DICOM > Configuration**.

e Select **Reset DICOM Counts**.

f Click **Apply Changes**.

Storage Expansion

Limitations on Storage Capacity Expansion

Rebalancing Content across storage volumes (LUNs) is not a feature currently available in HP Medical Archive solution 8.0.5 and prior releases. It is expected to be available in a subsequent HP Medical Archive solution 8.0.5 Service Pack. As a result, the Storage Capacity Expansion feature is also unavailable in this release. Adding additional Storage Nodes to an existing grid is a supported feature.

Additional Known Limitations

The following issues exist in the 8.0 software release. Please review this list before reporting a new issue.

Maximum Number of Filters in ILM Rules Limited to 6

Creating an Information Lifecycle Management (ILM) rule with more than six filters causes an IE script error. In addition, the browser may become sluggish.

Upgrade Fails to remove previous version of Custom ILM CD

The update script does not remove the 7.5 custom Information Lifecycle Management (ILM) CD which is no longer required because the ILM is on the 8.0 Software CD.

Workaround: Remove the 7.5 Custom ILM CD from /var/local/cis/iso, and then run:

```
cis register-iso-images
```

Summary Attributes not Updated after Storage Node Decommissioning

After decommissioning a Storage Node, the values of summary attributes such as Installed Storage Capacity are not updated in the NMS MI.

Workaround: Restart the Admin Node to update the values of summary attributes.

Switching to Configurable ILM can Trigger ILM Replication Backlogs

On switching over to configurable ILM, all items that were in the old replication queues, including those deferred for future re-evaluation, are moved to the new replication queues used for configurable ILM. ILM re-evaluation will continue until all items in the new replication queues are evaluated against the configurable ILM rules.

Storage Node Hardware Refresh grid task not generated

If a Storage Node has been refreshed in a previous version, and not all Storage Nodes were refreshed, the Storage Node Hardware Refresh (LMIG) grid task is not generated for the later version.

Workaround: Remove from the GPT the “replacing” attribute from the server stanza of the Storage Nodes that were refreshed in previous revisions so that this attribute is only present for the servers being refreshed in the current revision.

GRID Services not correctly handling CONNECTION_CLOSED Messages

GRID services not correctly handling CONNECTION_CLOSED messages for a connection that has not yet been supplied a SessionID block. This can prevent the LDR from shutting down.

High Availability Gateway Node Failover

Failover from the Active File System Gateway (FSG) to the Standby FSG in a HAGC cluster may fail if the Standby FSG is processing a large backlog of replication messages. The Standby FSG remains in the Standby state and the Cluster Status changes to Failed until the Standby FSG has processed the backlog of replication messages, at which point the FSG transitions to the Active state and the Cluster status changes to Vulnerable.

Secondary Gateway Nodes may reference deleted content

When a file is deleted on a primary gateway, the object may be removed from the grid before the reference to the object is removed from the secondary gateways.

High Availability Gateway Node Failover

A cluster may not fail over to the Standby Gateway Node if the system drive fails on the Active Gateway Node, as the heartbeat service may not detect the failure.

Provisioning fails if Decommissioned Node has unused Hardware Profile

Provisioning tries to load hardware profiles for inactive servers.

Workaround: Supply hardware profiles of missing servers to bypass error state.

CMS Does Not Purge Locations from LDR with least Available Space

Purging of an object from multiple storage nodes does not purge from the node with the least available space. This may result in aggravating grids with pre-existing difference in available space. These differences do not tend to resolve themselves over time.

LUN Rebalancing Grid Task should Restart the LDR upon Completion

Upon completion of the LUN Rebalancing Grid Task, the operator must manually restart the LDR to finish the activity.

Allow the FSG HTTP Client Request Timeout to be Configurable

Currently the File System Gateway (FSG) has a set 60 second timeout for HTTP requests. If there is any latency within this timeout period, the request fails and is retried by the FSG. Subsequent retries may continue to timeout indefinitely.

LDR may Crash During Remote Retain Transfer if Network Connection is Disrupted

The remote retain module is used during ingest with dual commit enabled, during Storage Node hardware refresh and during Storage Node recovery when the only other copy exists on an Archive Node. If a network disruption occurs when these processes are transferring data, it is possible for the LDR to fail on an assertion.

Delay ILM Re-evaluation of Deferred Items if there is Already ILM Re-evaluations Pending

The nightly queuing of batches of deferred items for ILM re-evaluation can trigger large ILM backlogs, potentially causing alarms. The greater problem is that new ingests coming in while these batches are being processed are at risk for an extended period. If the queues are not cleared daily, each nightly batch can add more into an ever-growing queue, leaving an increasing amount of content at an increasing level of risk.

Control Nodes may take a Long Time to Shutdown

Content Management System (CMS) internal scans and/or background tasks can prevent shutdown sequence until they are completed. Operators run the risk of database inconsistencies if forcing a shutdown during this time.

Storage Node Health Check may fail during LUN Rebalancing Grid Task

During LUN rebalancing, a destination volume can become full however the node is in Maintenance mode and should not trigger a health-check failure.

CMS may trigger MySQL Memory Leak Resulting in Segmentation Fault

Content Management System (CMS) activities may trigger MySQL memory leak. This will result in increasing memory consumption and finally a segmentation fault.

Metadata Capacity Displayed Incorrectly

Due to a known MySQL bug: http://bugs.mysql.com/bug.php?id= 20862, upgrading a grid to version 8.0 (MySQL 5.0) results in the incorrect ibdata size being used by mysql. This leads to the attribute values for the CMS and NMS databases being incorrectly reported to, displayed, and used in the NMS. Customers are not able to monitor and plan for metadata capacity as no alarm will be triggered.

Remote Retain Process may get stuck if Network Disruption Occurs During Transfer

The remote retain module is used during ingest with dual commit enabled, during Storage Node hardware refresh and during Storage Node recovery when the only other copy exists on an Archive Node. If a network disruption occurs when these processes are transferring data it is possible for the process to get stuck.

Storage Node Hardware Refresh Results in Phantom Alarms in the NMS

Network Management System (NMS) does not clean up alarms associated to decommissioned nodes. Alarms appear in the NMS without a parent node.

Procedure to Add LUNS to Existing Storage Node Takes a Long Time

The procedure to add additional LUNS to an existing Storage Node can take an excessive amount of time to create the new LUNS if the existing LUNS contain a large amount of content. The script to create the new rangedb structure is missing a max depth parameter that limits operations to just the directory structure.

Remote Retain Failure May Halt Ingest

In the event of a remote retain session failure, ingests may be halted. Remote retain sessions are used for dual-commit. This failure may occur during a race condition of a blocked remote queue being destroyed before it is unblocked.

LDR May Fail to Shutdown

The remote retain module is used during ingest with dual commit enabled, during Storage Node hardware refresh and during Storage Node recovery when the only other copy exists on an Archive Node. If a remote retain session fails, the LDR may fail to shutdown properly. Sessions may fail if network connections are unreliable during these operations.

Ingesting Multiple Duplicates May Result in UUID Mapped to Non-existent Content Block

In tests where multiple duplicate objects are ingested at the same time sometimes a UUIDs may be mapped to non-existent content blocks.

Limitations of Siemens syngo Imaging PACS

Certain storage status features of Siemens syngo Imaging PACS system VB30 and later may report inaccurate information; this does not effect PACS or grid operations.


Java™
COMPATIBLE