Hitachi Unified Storage (HUS) File Module NAS Operating System SU 12.2 and SMU 12.2 Release Notes

This document provides late-breaking information about the Hitachi Unified Storage (HUS) File Module NAS operating system, as well as a list of known issues and solutions. These release notes highlight (SU) 12.2.3753.10, and system management unit (SMU) 12.2.3753.11. This is a maintenance release that includes defect fixes only.

NAS operating system, which includes server update SU 12.2.3753.10 and SMU 12.2.3753.11, supports Hitachi NAS Platforms 4040, 4060, 4080, 4100, 3080, G1, 3080 G2, 3090 G1 and 3090 G2 models.

Intended audience

This document is intended for Hitachi Data Systems field and support personnel, customers, and authorized service partners.

Additions or changes

If you believe additions or changes are needed for these Release Notes, please send an HDS internal email to HNAS Doc Input HNAS_Doc_Input@hds.com.

Getting help

If you have any difficulties installing or configuring your NAS Platform, please call HDS Technical Support at one of the numbers listed:

North America:	**1-800-446-0744**
Outside North America:	**1-858-547-4526**

Hitachi Data Systems
2845 Lafayette Street
Santa Clara, California 95050-2627
https://portal.hds.com

Contents

Document history

The following changes have been made to this document.

Revision	Description
RN-92USF022-01	Initial publication for server version 12.2.3753.08 and SMU 12.2.3753.10
RN-92USF022-02	Maintenance release for server version 12.2.3753.10 and SMU 12.2.3753.11

New features at a glance

SU 12.2.3753.10 includes defect fixes only; there are no new features in this release. The following table lists the features included in the previous HNAS version, 12.2.3753.08 and their compatibility with HNAS server platforms. These features are explained in greater detail in this document.

Feature	HNAS 30x0 G1	HNAS 30x0 G2	4040	4060	4080	4100
Improved Dedupe throttling	✓	✓	✓	✓	✓	✓
Optimize deletion of clones	✓	✓	✓	✓	✓	✓

New features in detail

This section covers each of the key features and HNAS enhancements in greater detail. Please refer to the HNAS user guides for details on using these features.

Improved Dedupe throttling

The current Dedupe throttling mechanism was found to be insufficient in minimizing impact to file serving performance. This enhancement improves Dedupe performance by adding additional throttling parameters. These new parameters include:

- Throttling the number of Dedupe operations issued per second, based on the current load of the OBJ chip.

- Linking read-ahead operations to the above throttle to control their impact.
- Implementing a separate Dedupe read queue in DI to allow prioritization of file serving requests.

Optimize deletion of clones

This feature improves some of the latency and speed issues experienced when deleting clones. When deleting a clone, the object is truncated in chunks until its length reaches 0 then deleted. This involves a series of setlength operations that can take a long time to complete when the necessary metadata is not available in OBJ cache. This results in an extended checkpoint dead time and an unacceptable latency to clients.

In order to improve the setlength performance, the software issues read aheads to load the onode branches for the region being truncated. Each read ahead loads up to two direct onodes at once without loading the user data. However this remains suboptimal as it doesn't prevent loading or accessing the undiverged onodes unecessarily during the clone's deletion.

The deletion can also be optimized for completely undiverged clones as well as diverged clones with an undiverged leaf onode pointer.

Features included in the previous release of 12.2

The following features were introduced in the previous release, SU 12.2.3719.07.

Support for Multi-Tenancy (for data access)

The HNAS multi-tenancy feature provides HNAS application service providers (ASPs) with another configuration mode option in addition to the standalone HNAS individual EVS security feature. Both options provide support for multiple file serving Enterprise Virtual Servers (EVSs) on a single HNAS host or multiple hosts. However, the multi-tenancy option extends the functionality of the stand-alone option and provides additional security and configuration enhancements.

Understanding multi-tenancy

Multi-tenant architecture provides companies, such as application service providers (ASPs), the ability to support more than one customers' services on a single server, but still keep them logically separate. In an HNAS server implementation, this architecture is sometimes called real EVS separation.

Note: The ASP has the responsibility of managing the storage, file systems,

shares, and exports to which each tenant has access.

HNAS multi-tenancy configuration mode provides enhancements to the

previous stand-alone mode in the following ways:

- Supports tenant configurations in logically separate serving environments
- on a single physical server or cluster.
- Extends HNAS individual security mode to provide true separation by
- maintaining per-EVS variables and connection states.
- Supports serving environments for tenants with single or multiple EVSs,

- configured separately and possibly sharing file serving interfaces.
- Provides per-EVS IP routing and networking settings to support duplicate
- or overlapping server IP addresses. Includes support for both IPv4 and
- IPv6.
- Helps detect and prevent EVS crosstalk that can occur when duplicate IP
- ranges are used. EVS crosstalk can lead to server unresponsiveness.
- Provides CLI EVS context usability improvements.
- Provides support scripts and tools for migration.

HNAS multi-tenancy benefits

Using HNAS multi-tenancy can help you avoid some of the challenges faced with traditional multi-tenant environments. Commonly, HNAS customers who are ASPs (Internet services providers and managed services providers) sell their services to their customers. Their customers are the tenants in a multi-tenant environment. The ASPs cannot force their tenants into a specific subnet, which means that the ASPs run into issues when some tenants use the same network address scheme. In the past, this situation caused overlapping IP addresses and networks on the HNAS EVSs. The IP routing and networking settings were global on an HNAS server--per-EVS settings were unsupported. The HNAS multi-tenancy feature allows you to set up all the different tenant networks as VLANs and then allocate them to the specific EVSs. These networks may have the same IP subnet but may be different gateways in their VLAN-segregated networks.

Note: Multi-tenancy is licensed using the EVS Security license.

Tree directory delete (improved file and directory delete responsiveness)

The tree delete feature provides a mechanism to immediately remove a directory tree from its position in the file system and to perform the deletion as a background job. A directory tree consists of a specified directory and the hierarchy of subdirectories and files below it. When a directory tree is targeted for deletion, a tree delete job is created and added to the job queue. The targeted directory tree is immediately removed from the file system namespace, moved to the system trash directory, and scheduled for background deletion.

The Tree delete interfaces are provided in the form of management APIs, new CLI commands, and SOAP interfaces.

Tree delete is supported on WFS-2 file systems. No license is required.

The tree delete feature provides the following benefits, compared to deleting a directory tree via a network client:

- The instantaneous removal of a directory tree from the listing of the parent directory, allowing the client to proceed with further actions.
- The server-side delete eliminates the need for a client to recursively delete the directory tree over the network, therefore using less system resources.
- The multi-threaded implementation allows parallel deletion of the contents of the directory tree.

Support for symlinks on CIFS shares over SMB2

This feature adds symlink support for SMB2+, this will allow Windows clients from Vista onwards, as well as Mac OSX clients from 10.9 (Mavericks) onwards, to use symlinks over SMB on HNAS.

Applications such as Firefox and Thunderbird sometimes need to create symlinks in the user's home directory. Previously, when that home directory is hosted on HNAS, this caused problems as symlinks aren't supported over any version of SMB.

Universal Migrator improvements

You can now enable Dedupe on file systems that contain associations. This was not possible in versions of HNAS. Additionally, an issue is fixed, where snapshots at the start of migration were disabled due to a possible deadlock. This has now been fixed and the snapshots are re-enabled.

AES Crypto Support for NFS

The HNAS Kerberos implementation has been updated with the Advanced Encryption Standard (AES), the latest and so far the strongest available cryptosystem.

The Data Encryption Standard (DES) has been deprecated and is not secure. The following AES crypto profiles are supported:

- AES128-CTS-HMAC-SHA1-96
- ES256-CTS-HMAC-SHA1-96

Universal Migrator: non-modifying virtualization mode (for copying CIFS ACLs)

Universal Migrator is capable of migration over NFS only, because it has no knowledge about CIFS metadata; therefore, it is necessary to copy CIFS metadata by external process. A special mode is required in order not to apply CIFS metadata to LNAS over NFS.

A new mode of operation added to Universal Migrator which behaves as follows:

- During Virtualization – There will be no client access to the virtualization target file system with the exception of RoboCopy.
- There is read-write access for RoboCopy process running on a separate machine. RoboCopy process applies all CIFS metadata found on LNAS to HNAS. In contrary to current approach such requests are only performed on HNAS (not both LNAS and HNAS as currently implemented).
- During Migration - Clients are permitted to access HNAS. Metadata changes (both CIFS and non-CIFS specific) are applied only to HNAS (as there is no way to pass them over NFSv3). The only metadata bit applied both to LNAS and HNAS is file length as it affects actual file data.

Allow the SMU's TLS versions and ciphers to be disabled (post JDK 8)

By default, all protocols and cipher suites are enabled. However, occasionally a protocol or cipher suite may be no longer secure and an admin can now use the Security Options page in the SMU Administration menu to prevent a browser from communicating with

the SMU using that protocol or suite. It is necessary to have at least one protocol and cipher suite remain enabled.

The SMU communicates with the browser using HTTPS, and makes available a list of SSL/TLS protocols and cipher suites from which the browser can choose to encrypt communication with the SMU. The feature allows an admin to restrict the use of individual protocols or suites as desired. The GUI supporting the feature is found on the Security Options page under the SMU Administration menu.

Occasionally a customer feels that a protocol or cipher suite is no longer secure and wants to prevent a browser from communicating with the SMU using that protocol or suite. The lack of configurability has proved to be in some cases a "purchase roadblock" for potential customers with corporate security policies prohibiting protocols or suites that we make available.

As in the past, SMU upgrades may add or delete protocols or suites, and as before newly added items are enabled by default, and deleted items are no longer available.

No feature specific installation is required beyond the installation of the SMU application itself.

HCP 7.0 support for IPv6 (Validation)

Support for HCP 7.0 (including IPv6) is now included. There is no specific installation or usage requirements for this enhancement.

Improved background truncation and throttling

Improvements were made to address the performance of the File clone delete feature. There is no specific installation or usage requirements for this enhancement.

Important considerations to read before installation

Please read the following sections before installing and using SU 12.2.3719.07.

Important note on downgrading from 12.x versions

If you need to downgrade from 12.x to an earlier version, the following considerations must be kept in mind.

When version 12.x places a file system in a storage pool's recycle bin, Cod is written in a format that version 11.x (or earlier) cannot read. The result is that the file system will not load. Before downgrading from 12.x, ensure compatibility with earlier releases by recycling all deleted file systems on all storage pools, using:

filesystem-recycle --all-spans --all-filesystems

Note: Perform this step carefully, as it will stop **filesystem-undelete** from working.

If you need to downgrade to 11.x (or earlier) then, before downgrading, you also have to run **span-rewrite-cod** on each span for which **filesystem-recycle** recycled at least one file system.

Performing an emergency downgrade

If you have to perform an emergency downgrade and don't get the opportunity to run **filesystem-recycle**, then, for each span whose file systems will not load, follow these steps after performing the downgrade:

Use **sd-back-up-cod** to take a single-SD Cod backup.

Use the Cod Converter in 11.x to convert it to an integrated Cod backup (ICB).

Use **span-restore-cod** to restore the ICB.

Use **span-rewrite-cod** to downgrade the Cod, so that it loads after the next reboot.

Note on Shellshock CVE-2014-6271 “Bash” Security Vulnerability

A recent security vulnerability known as CVE-2014-6271 has come to our attention. This vulnerability affects UNIX-based Bash (Bourne shell) and has the potential to arbitrarily execute code within UNIX environments. Some native services and applications may allow remote unauthenticated attackers to provide environment variables and exploit this issue. At this time, there are no known HNAS vulnerabilities.

For up-to-date details, customers and partners can log in to the HDS Support Portal or PartnerXchange (select “Support Portal” on the upper-right tab once logged into PartnerXchange) and click the link to the customer letter on the homepage.

Note on supported AMS storage arrays

Please note that AMS storage arrays are supported on all HNAS 30x0 and 4000 systems. Previously, the *Storage Subsystem Administration Guide* stated that AMS arrays were not supported on the 4040 series. All series 4000 HNAS storage servers support the AMS 2100, 2300, and 2500 storage arrays.

Note on data spillage from tier 1 to tier 0

A file system consists of files and directories. Data within the file system (both user data and metadata) is stored on the storage media of a storage subsystem. In HNAS, storage subsystems are classified into “tiers,” which are then used to manage storage resources.

In a tiered file system, metadata is stored on the highest performance tier of storage, and user data is stored on a lower-performance tier. It is possible for tier 1 data (user data) to spill over into tier 0 (metadata). This will only occur if the tier 1 file system is full, and additional data is written to the file system.

Please note that if tier 1 data spills over to tier 0, performance may be degraded, including reduced write performance. In upcoming versions of HNAS, users will be alerted if such spillage occurs, enabling them to better allocate data.

Note on using Hitachi Dynamic Provisioning on HNAS before and after v12.1

HNAS supports Hitachi Dynamic Provisioning (HDP) thin provisioning. This note concerns file system size, and full capacity mode use, in versions, pre and post 12.1.

HNAS supports HDP thin provisioning, and in versions prior to 12.1, full capacity mode must be enabled.

- Before 12.1, Full Capacity Mode = Enable is mandatory
- From 12.1 and up, Full Capacity Mode = Disable is mandatory

Prior to v12.1, if you divided the HDP-POOL in equal sized DP-VOLs, there was a risk that you could not create any file systems in the storage pool because, after diving the HDP-POOL into DP-VOLs, there may be a small leftover chunk size. For example:

- Having a chunk size of 18GB, and the leftover chunk is more than 18GB (e.g., 19GB), will not cause a problem.
- Having a chunk size of 18GB, and the leftover chunk is less than 18GB e.g., (17GB) you may get the following error message: "Failed: Can't create or expand the filesystem: the host span has too few free chunks."

In version 12.1 onward, there is no need for full capacity mode on thick provisioned HDP, as HNAS will now recognize any leftover chunk size.

Possible boot loop on upgrade to 12.x firmware on 4060/4080/4100 systems

Previously, after upgrading to 12.0 or 12.1 versions of firmware, the node might get stuck in a boot loop. It is believed that a change introduced in 12.0 exposed a marginal timing condition on certain FPGAs. If a system has been upgraded to 12.x code, and is running normally, then it is not likely to be susceptible to this issue. This issue did not affect 3080,3090 or 4040 HNAS systems.

Note that this issue is resolved in version 12.1.3163.10 on.

Recovery

If you should need to recover from this issue, roll back the upgrade to 11.x code or replace the node.

Note: Recovering a node in a boot loop is not always possible. It may become necessary to replace the node, but there is a risk the new node could have the same issue.

Important note on SMU code release 12.2.3753.11

This release uses SMU code release 12.2.3753.11, and it requires the CentOS 6.2 operating system. Before you install the 12.2.3753.11 SMU code, you must install the 6.2 CentOS operating system on external SMUs. See the section, Upgrading from 8.x to 10.x, in this document, for more information.

SU 12.2.3753.11 is not compatible with SMU 200. Before updating, check the current SMU version. If it is SMU 200, upgrade to SMU 300 before proceeding. You can use the **smu-info** command to check the current SMU version. The syntax for this command is:

```
[root@hdsml tmp]# smu-info
usage: /usr/local/bin/smu-info <output file>
```

Alternatively, you can use the vSMU, or, in a non-cluster environment, the embedded SMU.

For complete instructions on upgrading the SMU, see "Upgrading the SMU and Server Software", in the *System Installation Guide*. If you need to update the SMU, note that you must also save the current SMU configuration files to the updated version. Save the base configuration to the network element, and then restore the configuration. See "Backing Up Configuration Files," in the *System Installation Guide.* This guide is for internal HDS personnel only.

SMU, server, and cluster compatibility

HNAS 30x0 servers running HNAS OS version 12.2 require the system management unit to be running SMU 12.2 software; however, SMU 12.2 can manage HNAS 3x00 and HNAS 30x0 servers on the latest HNAS OS version 8.1, 8.2. 10.2,11.x and all released versions of 12. The following chart shows SMU, server, and cluster compatibility.

	Supported Server Releases	# of Supported Servers / Cluster			
		SMU 200	SMU 300	SMU 400	VMware
SMU 6.1	HNAS OS 6.1	8	8	8	--
SMU 6.5	HNAS OS 6.1, 6.5	8	8	8	--
SMU 7.0	HNAS OS 7.0	8	8	8	--
SMU 8.0	HNAS OS 7.0, 8.0	5	5	5	2
SMU 8.1	HNAS OS 7.0, 8.1	5	5	5	2
SMU 10.0 (Cent OS 6.0)	HNAS OS 8.1, 10.0	2	5	5	2
SMU 10.1, 10.2 (Cent OS 6.2)	HNAS OS 8.1, 8.2, 10.0, 10.2	2	5	5	2
SMU 11.x (Cent OS 6.2)	HNAS OS 8.1, 8.2, 10.2, 11.x	NOT SUPPORTED	5	5	2
SMU 12.x (Cent OS 6.2)	HNAS OS 8.1, 8.2, 10.2, 11.x, 12.x	NOT SUPPORTED	5	5	2

Notes on installing, upgrading and downgrading

Notes on this release include:

- NAS platforms 4040, 4060, 4080, 4100, 3080, G1, 3080 G2, 3090 G1 and 3090 G2 models, with cluster support up to four nodes.
- When establishing a cluster, locate the node that is assigned the Storage (Terabyte/TB) license key, and configure it as the first node in the cluster, as it is required for proper license activation.,
- The Web Manager for the SMU uses cookies and sessions to remember user selections on various pages. Therefore, you should open only one web browser window, or tab to the SMU from any given workstation.

Upgrade Path Flowchart

When upgrading from previous versions, use the following chart for the recommended upgrade paths:

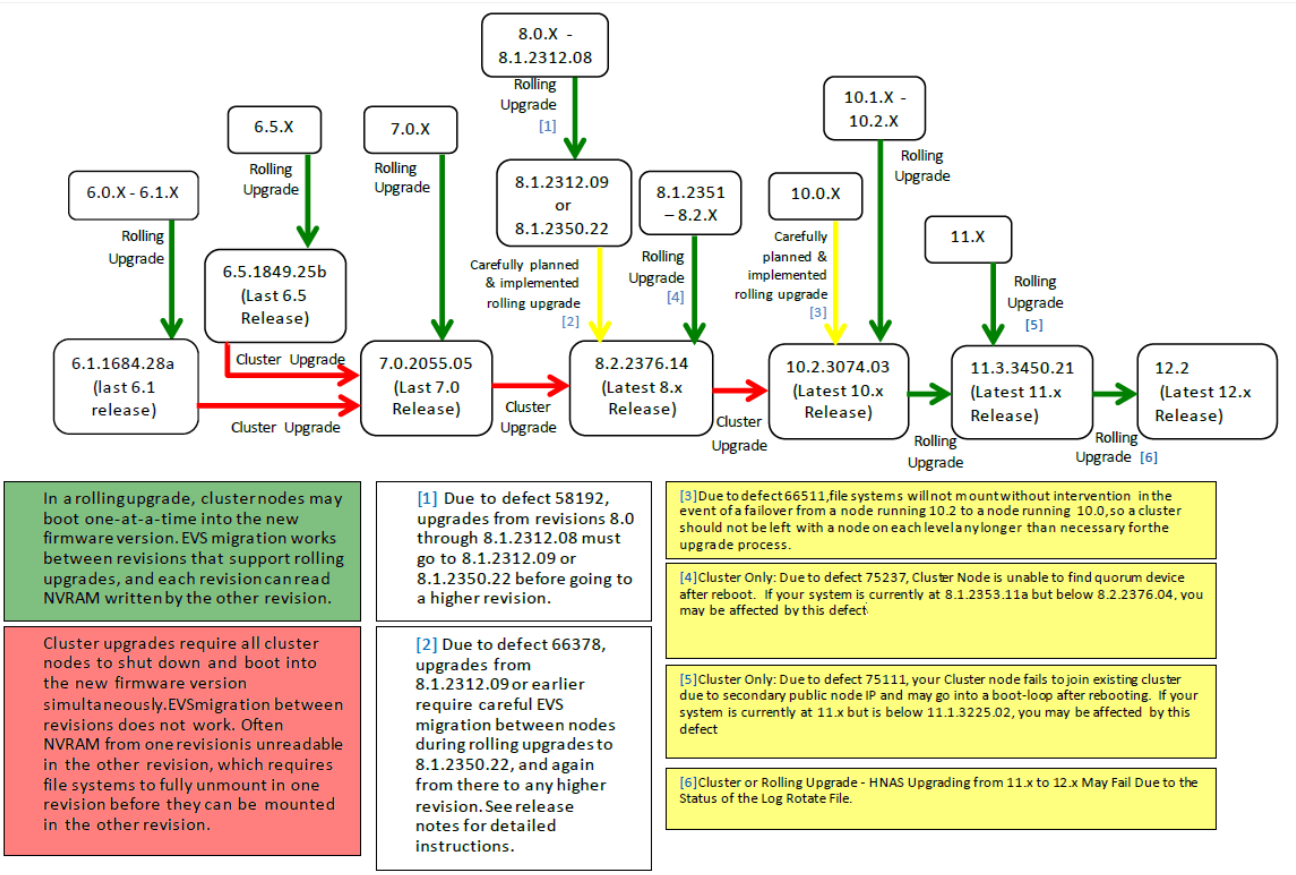

For more details on issue [5] above, see: **Technical Bulletin** **HNAS Cluster Node Fails to Join the Existing Cluster Due to the Presence of a Second pnode ip Address**

For more details on issue [6] above, see: **Technical Bulletin** **HNAS Upgrading from 11.x to 12.x May Fail Due to the Status of the Log Rotate File**

Performing a rolling upgrade on 8.x versions

When performing the rolling upgrade, 8.0.X-8.1.2312.09 → 8.1.2350.22 → 8.2.2374.14, follow the steps in the section, Performing a Rolling Upgrade to 8.2.2374.12, in this document. Note that you must upgrade to 8.1.2312.09 before upgrading to a later version of 8.x or 10.x. If you are already running 8.2, you can upgrade directly to 10.x.

Warning: File systems will not mount without intervention in the event of a failover from a node running 10.2 to a node running 10.0, so a cluster should not be left with a node on each level any longer than necessary for the upgrade process (defect 66511). When

performing a rolling upgrade from 10.0 to 10.2, the following steps are highly recommended:

1. Reboot node 1 from 10.0 to 10.2.
2. Confirm that the node booted into the new code.
3. Migrate all EVSs to node 1.
4. Reboot node 2.
5. Do steps 3 and 4 as soon as possible, after step 1.

Upgrading from NAS OS 7.0

For this major release, upgrade instructions and requirements can vary based on the storage system configuration, see the *Server Cluster and Administration Guide* for detailed instructions for upgrading from NAS OS 7.0 to 8.x.

Notes:

- New systems can be shipped with 7.0.2052.02 server code from our manufacturer. As part of your installation or upgrade plans, assess the customer implementation requirements. If the system uses CNS serving CIFS clients, you can downgrade to server code version 7.0.2051.07, or upgrade to 7.0.2053 or a later version (as available).
- If the system is running a NAS OS 6.x release, or earlier, you must first upgrade to 7.0 before installing 8.x.
- Rolling upgrades are not supported between major releases at this time. Ensure you schedule system downtime when performing this upgrade.

Upgrading from 8.x to 10.x

The SMU release 10.2.x uses the CentOS 6.2 operating system. You must upgrade the SMU operating system to CentOS 6.2 before you upgrade the SMU software.

Important: For the first installation of the SMU software on a 10.1 system, you must use the CLI. The instructions are included in this section.

Caution: Before upgrading or making configurations changes to the system, it is highly recommended that you back up the configurations and save them in a safe, external location. See the *System Installation Guide* for more information.

Note: It is a recommended best practice that system diagnostics are captured both before and after any upgrade or change to the system configuration.

Install CentOS 6.2 on the SMU

1. Back up the SMU configuration to an external source (the installation procedure destroys the configuration so you need to reinstall it after installing the SMU software).
2. Insert the CentOS 6.2 DVD and reboot the SMU.
3. At the prompt, enter **clean-kvm** or **clean-serial** depending on your connection.
4. After the system loads the CentOS software, replace the CentOS DVD with the SMU (Uplands) DVD and reboot the SMU (this is the default option on the monitor).

Install the SMU software using the CLI

1. Log in as **root** (using serial or KVM) using the default password **nasadmin**.
2. If you are installing from an ISO image, skip to step 4.
3. Insert the SMU software DVD and enter the following commands:

```
a) mount /media/cdrom <-- Some SMUs may have /media/cdrecorder>
b) /media/cdrom/autorun <-- Some SMUs may have /media/cdrecorder/autorun>
```

   The system reboots when the installation is complete.

4. If you are installing from an ISO image (*SMUsetup_uplands_xx.iso*), rather than a DVD, copy *SMUsetup_uplands_xx.iso* to /tmp on the SMU (use **scp**) and use the following commands:

```
a) su - root
b) cd /tmp
c) mount -o loop /tmp/SMUsetup_uplands_xx.iso /media/cdrom
d) /media/cdrom/autorun
```

5. Log in as **root** using the default password (**nasadmin**) and run the **smu-config** script.
6. Restore the SMU configuration from your backup configuration using the CLI or Web Manager.

Note: If an external or embedded SMU is running an SMU software 7.x release, or an earlier release, it must be upgraded. Use the procedure in the *System Installation Guide* to upgrade the SMU software.

- Upgrading from an external SMU release before 8.x to an 8.x release, an external SMU will also require upgrading to CentOS 4.8.1 as a first step.
- An internal SMU does not require upgrade to CentOS.

Install the SU software using the GUI

Caution: Before upgrading or making configurations changes to the system, it is highly recommended that you back up the configurations and save them in a safe, external location. See the System Installation Guide for more information.

Note: It is a recommended best practice that system diagnostics are captured both before and after any upgrade or change to the system configuration.

Note: 10.0, and later releases, do not support the 2xxx , 3100 and 3200 platforms.

To upgrade servers running the NAS OS:

1. Open a supported browser and enter the server SMU IP address to launch Web Manager.
2. Log in as **admin** (default password **nasadmin**).
3. Click **Home** > **Server Settings** > **Firmware Package Management**.
4. Ensure there are fewer than three existing packages (**excluding** any "patch" *.tar.gz* files). If there are more than three packages, remove the oldest files by selecting the checkbox next to its name, and clicking **delete** in the Package View section.
5. Select **upload package** in the Node View section.
6. Select a managed server, and then click **OK**.
7. Click the **Browse** button, and select the new server software package; ensure the **Set as default package** and **Restart file serving**, **and Reboot the server if necessary** options are enabled, click **Apply**, and then click **OK** when the warning is displayed to start the install.

The Package Upload Progress page is displayed. At the end of the process, the file server restarts. After the page refreshes, "Upgrade completed successfully" is displayed.

The status indicator might appear red, and the system displays the following message: "Server is not responding to management protocol. Connection refused." If this happens, refresh the page to resolve the issue.

Refer to the *System Installation Guide* for additional installation and upgrade information, and the documentation CD for platform documentation.

Performing a rolling upgrade to 8.2.2374.12

These instructions are written with the assumption that all cluster nodes are currently running 8.1.2312.09. The instructions transition you through 8.1.2350.22 to get to a later version of 8.1 or 8.2.

Note: This special procedure is not necessary when a cluster is running 8.1.2351.00 or later. In such cases, a regular rolling upgrade can be performed, in other words: install the new version on every node, and then reboot each node in turn.

Important: Where 8.2.2374.12 appears in these instructions, you may substitute any version of 8.1 or 8.2 later than 2352.08. This procedure is necessary to cross a barrier between 8.1.2350 and 8.1.2352 and it applies to all later 8.1 and 8.2 builds. A simultaneous cluster-wide upgrade need not use 8.1.2350.22 in this manner.

The following steps make the assumption that all nodes in the cluster have the following builds loaded or installed:

- 8.1.2312.09
- 8.1.2350.22
- 8.2.2374.12

1. To prevent EVSs from migrating back to their preferred node during the procedure, disable "auto fail-back". To disable auto fail-back run the following command:

   ```
   evsmap autofb off
   ```

2. Set the default firmware package on all nodes to 8.1.2350.22, this firmware version is a prerequisite before moving to 8.2.2374.12.

   To set the package on Hitachi 3080/3090 series servers:

   ```
   cn all package-set-default mercury-8.1.2350.22.tar
   ```

   To set the package on HNAS 2000/3000 series platform servers:

   ```
   cn all package-set-default titan-8.1.2350.22.pkg
   ```

3. All EVSs should be migrated to a single node in the cluster e.g. node 1.

   ```
   evs migrate --all <from node ID> -n <to node ID>
   ```

   Examples:

   ```
   evs migrate --all 2 -n 1

   evs migrate --all 4 -n 1
   ```

4. After all EVSs have been migrated to an individual node (e.g. node 1), reboot all other nodes in the cluster, waiting for each node to complete the reboot before moving on to the next node. On Hitachi 3080/3090 series, restarting "file services" (Bali) is adequate. On the HNAS 2000/3000 series servers, a full reboot is required.

   Hitachi 3080/3090 series:

   ```
   cn <node ID> reboot --app
   ```

   Example:

   ```
   cn 2 reboot --app
   ```

   HNAS 2000/3000 series:

   ```
   cn <node ID> reboot
   ```

   Example:

   ```
   cn 2 reboot
   ```

After the reboot:

- Ensure `cluster-show` reports the health of the cluster as `robust`
- Some pauses may be seen, but I/O should not terminate or hang indefinitely
- Mirroring will be automatically enabled in one direction (e.g. node 1 to node 2), which can be observed with `cn all nvmirroring`

5. Migrate all EVSs to an individual node running 8.1.2350.22, for example node 2.

```
evs migrate --all <from node ID> -n <to node ID>
```

Examples:

```
evs migrate --all 1 -n 2

evs migrate --all 4 -n 2
```

6. After all EVSs have been migrated to a node running 8.1.2350.22 release (for example, node 2), reboot the node running 8.1.2312.09 (for example, node 1). On Hitachi 3080/3090 series servers, restarting "file services" (Bali) is adequate. On the HNAS 2000/3000 series servers, a full reboot is required.

Hitachi 3080/3090 series:

```
cn <node ID> reboot --app
```

Example:

```
cn 1 reboot --app
```

HNAS 2000/3000 series servers:

```
cn <node ID> reboot
```

Example:

```
cn 1 reboot
```

After the reboot:

- Ensure `cluster-show` reports the health of the cluster as `robust`
- Some pauses may be seen, but I/O should not terminate or hang indefinitely
- Mirroring will be automatically enabled in one direction (for example; node 2 to node 1), which can be observed with `cn all nvmirroring`

7. Set the default package on all nodes of the cluster to 8.2.2374.12

Hitachi 3080/3090 series:

```
cn all package-set-default mercury-8.2.2374.12.tar
```

HNAS 2000/3000 series servers:

```
cn all package-set-default titan-8.2.2374.12.pkg
```

8. Reboot all nodes **not** currently hosting EVSs, waiting for each node to complete the reboot before moving on to the next node. On Hitachi 3080/3090 series, restarting "file services" (Bali) is adequate. HNAS 2000/3000 series servers, a full reboot is required.

   Hitachi 3080/3090 series:

   ```
   cn <node ID> reboot --app
   ```

   Example:

   ```
   cn 1 reboot --app
   ```

   HNAS 2000/3000 series servers:

   ```
   cn <node ID> reboot
   ```

   Example:

   ```
   cn 1 reboot
   ```

   After the reboot:

   - Ensure `cluster-show` reports the health of the cluster as `robust`
   - Some pauses may be seen, but I/O should not terminate or hang indefinitely
   - Mirroring will be automatically enabled in one direction (e.g. node 2 to node 1), which can be observed with `cn all nvmirroring`

9. Migrate all EVSs to an individual node running 8.2.2374.12, for example node 1.

   ```
   evs migrate --all <from node ID> -n <to node ID>
   ```

   Examples:

   ```
   evs migrate --all 2 -n 1

   evs migrate --all 4 -n 1
   ```

10. After all EVSs have been migrated to an individual node running 8.2.2374.12 (e.g. node 1), reboot all other nodes in the cluster, waiting for each node to complete the reboot cycle before moving on to the next node. On Hitachi 3080/3090 series servers, restarting "file services" (Bali) is adequate. On the HNAS 2000/3000 series servers, a full reboot is required.

    Hitachi 3080/3090 series:

    ```
    cn <node ID> reboot --app
    ```

    Example:

    ```
    cn 2 reboot --app
    ```

HNAS 2000/3000 series servers:

```
cn <node ID> reboot
```

Example:

```
cn 2 reboot
```

After the reboot:

- Ensure `cluster-show` reports the health of the cluster as `robust`
- Some pauses may be seen, but I/O should not terminate or hang indefinitely
- Mirroring will be automatically enabled in one direction (e.g. node 1 to node 2), which can be observed with `cn all nvmirroring`

11. Distribute the EVSs, as desired, between the cluster nodes

```
evs migrate -e <EVS ID> -n <to node ID>
```

Example:

```
evs migrate  -e 3 -n 3
```

Some pauses may be seen, but I/O should not terminate or hang indefinitely

Mirroring will be automatically enabled in both directions, which can be observed with `cn all nvmirroring`

12. Re-enable EVS auto failback within the cluster, if it was enabled before the procedure or is otherwise desired.

```
evsmap autofb on
```

Upgrading from 10.x to 11.3

The SMU 11.3 release uses the CentOS 6.2 operating system. If you are running SU 10.0 you must upgrade to the CentOS 6.2 operating system. Follow the instructions in the section, Install CentOS 6.2 on the SMU, in this document. If you are running SU 10.x or greater, you are already running the CentOS 6.2 operating system.

Important: When upgrading from 10.x to 11.3, the SU version your system is running must be within the following SU ranges:

Server Releases		
10.0.3036.00	to	10.0.3099.99
10.1.3070.00	to	10.1.3099.99
10.1.3100.00	to	10.1.3199.99
10.2.3071.00	to	10.2.3099.99
10.3.3100.00	to	10.3.3199.99

Caution: Before upgrading or making configuration changes to the system, it is highly recommended that you back up the configurations and save them in a safe, external location. See the *System Installation Guide* for more information.

Note: It is a recommended best practice that system diagnostics are captured both before and after any upgrade or change to the system configuration.

Use the following chart to find the correct CentOS version to use with SMU releases.

SMU Release	OS
SMU 12.x	CentOS 6.2
SMU 11.2	CentOS 6.2
SMU 11.1	CentOS 6.2
SMU 11.0	CentOS 6.2
SMU 10.2	CentOS 6.2
SMU 10.1	CentOS 6.2
SMU 10.0	CentOS 6.0
SMU 8.x	CentOS 4.8.1
SMU 7.x	CentOS 4.4
Internal SMU	No CentOS

Install the SMU software using the GUI

From the SMU:

1. Navigate to **SMU Administration** > **SMU Upgrade**. The system displays the SMU Upgrade page.
2. Click **Browse**. The system displays a search dialog you use to search for the .iso upgrade file.
3. Click **Apply**. The system displays a message warning that the upgrading the SMU will restart the SMU web application and may take in excess of ten minutes.
4. After the SMU restarts, install the SU software.

Install the SU software

Use the instructions in the Install the SU software using the GUI section to install the SU OS software on your 11.2 system.

Downgrading from 11.2

From the SMU:

1. Select **Server Settings** > **Upgrade Firmware** dialog. The system displays the Server Upgrade Selection dialog.
2. Select **Managed Server** and then a server from the drop-down list.
3. Click **OK**. The system displays the Upgrade Firmware dialog.
4. Click **Browse** and select the upgrade file you want to load on the nodes.
5. Click **apply**.

After downgrading the nodes using the SMU, the nodes reboot. If a sufficient number of nodes boot to form a cluster without the need for the quorum device (three nodes, in the case of a four node cluster), the nodes form a cluster. This cluster is in a **Critical** state, and the user must manually reconfigure the quorum device by running the following commands on a cluster node:

```
Cluster-1$ quorumd remove
Cluster-1$ quorumd add <SMU-NAME>
```

Important considerations if downgrading from version 11.2

HNAS version 11.2 introduced support for more than 1023 chunks per file system. If you have set up non-standard parameters using `span-tune-allocator`, restore the defaults before downgrading from 11.2. Allocator settings are stored in storage pool Cod, and releases earlier than 11.2 can't parse Cod with non-standard settings.

If a file system has more than 1,024 chunks, that file system, and all others on the same storage pool, will be unusable after a downgrade from 11.2. You can find the number of chunks on a storage pool's file systems by running `file system-scalability` on the storage pool.

If any storage pool has more than 16,384 chunks, that storage pool will be unusable in releases earlier than 11.2. In 11.2, you can find the number of chunks in a span by running, `span-list-chunks –terse span-instance-name | tail`

Note: Chunks are numbered beginning from 0, not from 1.

File-based replication between different HNAS firmware levels

The ability to replicate between systems is determined by the version of the firmware that is running on those systems. The model number of the server is not a factor for interoperability for replication purposes. If both the destination and target servers are running the same major software version (for example, 6.x), replication as 'managed servers' is fully supported. If the destination and target servers are running different major software versions (for example, 7.x to 8.x), one of the servers is configured as an 'unmanaged' server. Replication continues to be fully supported within the constraints of replication between managed and unmanaged servers.

Object-based replication between different HNAS firmware levels

Object replication was first introduced in HNAS OS v8.0 and has been enhanced with each release. For example in v10.1, an enhancement was made so that during incremental replication, objects maintained their sparseness. v11.1 has the ability to preserve file clone states during replication. To ensure interoperability, feature flags are

negotiated when object replication occurs between servers running at different version levels.

Object replication between servers is supported up to one major version away. For example, object replication between servers running v8.x and v10.x and v10.x and v11.x are supported.

Note: Object replication between servers that are more than one major releases apart may work (for example, between v8.x and v11.x) – but this is not supported.

Important update: Mandatory new SyncDR Release v2.0.5 for HNAS OS 11.2

Release 11.2 introduced a revised license key name for the SyncDR license.

To support the new license key name, a new version of SyncDR is required. Existing license keys do not have to be upgraded, but still requires the new SyncDR version.

SU 11.2 requires SyncDR v2.0.5. Older versions of SyncDR will *not* work with release 11.2.

SyncDR v 2.0.5 supports all previous HNAS versions, from version 8.0, on.

Availability and upgrade.

The new version is a drop-in replacement for the existing version, and can be installed as usual. Please consult with your HDS field representative for details on installation.

Changes:

The following changes are implemented by **SyncDR v2.0.5**:

- Support for legacy and new SyncDR license key name
- Improved failure detection
- The monitor will now check the HNAS system more often. Together with other adjustments, SyncDR will now trigger a failover between 30-50 seconds faster than before
- Improved SSC resiliency

  In earlier versions, unstable SSC connectivity caused SyncDR to abort failover and monitoring immediately (due to blocked SSC sessions). Version 2.0.5 of SyncDR is more resilient and will retry failed commands.

- Cleaner dblog

  SyncDR monitor had the tendency to fillup the HNAS dblog with health check messages this is no longer the case. Monitor messages are not logged to the dblog anymore.

- Span names including "." Are now supported.

  SyncDR now recognizes spans which include a dot (".”). However the use of special characters in a storage pool or file system labels is not recommended.

- Smaller bug fixes and adjusted log messages

Licensing

This section provides important information on license types and clustering.

A ModelType license type will be implemented, that can be set to "none", 4040, or 4080:

- A 4060/4080 node that has no ModelType license will report its model number as 4060.
- A 4060/4080 node that has the 4080 license applied will report its model number 4080.
- If clustered, all 4060 nodes in the cluster will be affected by this license (and will thus become 4080 nodes).
- 4100 is unaffected by the ModelType license, and will always report its model number 4100.

Only nodes of the same model can be clustered together:

- Adding a 4060 to a 4080 cluster would result in the 4060 picking up the 4080 license, and would thus become a 4080.

File size, maximum per entity, and throughput are directly controlled by the model number, that is, these values are fixed in the software, based on the model number that results from the detected hardware type, and any relevant ModelType license.

Cluster nodes and usable capacity maximum per Entity are already controlled by the cluster and TB license types respectively. The values will therefore be the maximum allowable limits:

- Adding a cluster license of four to a 4060 node would still only allow two nodes in a cluster (as it is beyond the max limit for a 4060).
- Adding a cluster license of three to a 4080 node would allow three nodes (as it is less than the maximum limit for a 4080).
- A server would, by default, include a license that matches the maximum limit for that model type.
- The same logic applies for the TB license.

Note: New license keys are typically firmware-version specific. Upon upgrading firmware to this release, all previous licenses present on the system will continue in force.

To request upgrade keys

- New features with a sale price will be purchased by the customer per normal HDS channel policies and procedures.
- Non-sale feature requests will be routed based on server branding until such time as the relicensing process has been fully integrated.
- HDS Server Request Routing

- The emailed request shall include the following information:
  - Customer Name
  - MAC-ID of the HNAS Unit (the MAC-ID format is XX-XX-XX-XX-XX-XX), the serial # is not needed or acceptable to issue new keys.
  - If you have not followed normal upgrade procedures, please indicate details of your current situation and indicate if a new full set of keys are required. Also, if your server is part of a cluster, please indicate if the MAC Address is a "Primary" server of the cluster and how many units are in the cluster.
- All permanent upgrade key requests will be handled by way of email sent to TBKeys@HDS.com. Turnaround time on all requests is targeted within 24 hours. Standard working hours for this distribution list (dlist) are 8am to 5pm Pacific Standard Time. See below for emergency situations.
- Should your need for upgrade keys be an emergency, please contact the HDS Support Centers, where Temporary Keys for these features can be provided.
- An email to TBKeys@hds.com should also be sent to receive your permanent keys.

Fixes and known issues

Fixes in 12.2.3753.10

Issue ID	Severity	Summary	Explanation
91069	D	Record in the per-fs-throttle man page that it only applies to file systems accessed by supported protocols	This issue is resolved. **per-fs-throttle** man page now makes it clear that it applies only to NFSv2 and NFSv3.
106470	B	fatal assert SI/T2_SI_REG/perx_stats_fatal_error from SiM2	Fixed issue where running a PIR while a tape back-up was in progress would cause DI to crash with fatal assert SI/T2_SI_REG/perx_stats_fatal_error.
100500	C	The SMB2 server calls into the file system without a file system context: smb::v2::FileSystemContext::begin(smb::v2::FileSystemContextResources const&	A cause of server resets when accessing files via SMB has been corrected.
103409	C	Assert failure panic - main/fsb/wfs/WFSVolume.cpp:314: in function static WFSVolume& WFSVolume::getCurrentFromFLS():	This issue is resolved. A cause of server resets when accessing files via SMB has been corrected, and will not result in: Assert failure panic - main/fsb/wfs/WFSVolume.cpp:314: in function static WFSVolume& WFSVolume::getCurrentFromFLS():
104531	C	SIGSEGV (Segmentation fault) ResolveLinkName<unicodechar<="" td="">	Fixed so customers will not experience an erroneous node reset signature.
104532	C	SIGSEGV (Segmentation fault) ResolveLinkName<unicodechar<="" td="">	This issue is resolved. Fixed so customers will not experience this an erroneous node reset signature.
94052	C	start of "assertion failure" panic in 'fsb/dedupe/main/lib/DedupeQueue.cpp'	Disabling and enabling the dedupe service caused incorrect flags in the dedupe config, causing assertion failure panic. This issue is resolved.
100994	C	"invalid netmask" panic in function void net::ipv4::ensureContiguousNetmask triggered by wildcard in cifs-dc filter	This issue is resolved. Fixed regression that changed choice of source ip addresses for NDMP data connections.
106539	C	NDMP data connections are made from an inappropriate source address	This issue is resolved. Fixed regression that changed choice of source IP addresses for NDMP data connections.
103629	C	CHD-1 became unresponsive and had to be reset (RST button)	This issue is resolved. Code modified so that deadlock seen at customer sites can no longer occur.

104725	C	Node2 became unresponsive and had to be reset (RST button)	This issue is resolved. Code modified so that deadlock seen at customer sites can no longer occur.
104710	D	Migration path list in SMU times out after 1 minute but server may take longer to respond.	This issue is resolved. When HNAS had many external data migration paths on a single EVS, the SMU may have timed out when retrieving those paths from HNAS. As such, the SMU was unable to display all paths. The timeout has been increased from 1 minutes to 5 minutes per EVS. This affects the page **Home** > **Storage Management** > **Data Migration Paths.**
107791	C	High numbers of smb::v2::AsyncConnectionCloseWork and connection exhaustion.	This issue is resolved. The handling of named stream lookups in the SMB2 QueryDirectory has been modified to report the correct error.
97028	C	Cannot mount FS after large file length change & panic due to shutdown timeout	This issue is resolved. Mounting a file system after a large file length change will not cause aVLSI hung panic.
104912	D	scsi-queue-limits-X displays VSP G600 along with VSP G400/G600	This issue is resolved. **scsi-queue-limits** commands will no longer show VSP G600 storage devices.
106194	C	Record logins in the management log.	This issue is resolved. Added server logging as part of a feature that optionally monitors all console commands and the users who ran them.
95188	C	Reset: FSBStatus ResolveLinkNameLeaf<unicodechar, FSBStatus ResolveLinkName<unicodechar >	This issue is resolved. An issue which could result in a panic when handling SMB1 read request to a file under the C$ (the unified root share) has been addressed.
101037	C	Effective Access tool returns "Code 0x80070057 the parameter is incorrect" from Windows 2012 Server and Windows 10 clients when accessing a UNC path.	This issue is resolved. A change to HNAS has been made to enable Explorer on Windows Server 2012 and Windows 8 clients to perform Effective Access checks against HNAS shares.
103916	C	"Half open session destructor queue" stuck thread due to deadlock between SMB2 close request triggering delete-on-close on file and change notify on directory, and SMB2 lock request against apparently the same file.	This issue is resolved. A coding problem which could result in a deadlock when handling SMB2 byte range locks has been fixed.
99309	B	start of "MFB vlsi fatals OBJ/T2_OBJ_MISC_ASSERTS/fdp_ack_is_rsp_error" panic	This issue is resolved to prevent an internal VLSI response fifo overflow.
105448	C	VLSI hung panic due to what appears to be another large setlength to zero.	This issue is resolved to prevent a VLSI hung panic due to what appears to be another large setlength to zero.

Fixes in 12.2.3753.08

Issue ID	Severity	Summary	Explanation
71946	C	Node reset due to watchdog reset (believed to be leap second insertion issue).	This issue is resolved. A new Linux kernel containing a fix for "the leap second" problem is included in this build. Additionally, to address this issue in previous releases, upgrade to a fixed version (12.2.3753.08 and 12.1.3613.16) of the HNAS OS, prior to the next scheduled leap second insertion.
96722	C	Obj_return_blocks slow after upgrading to 12.1.	This issue is resolved. Fixed a problem where customers upgrading to affected builds may see warnings about OBJ_RETURN_BLOCKS taking a long time.
102287	C	Broken logrotate resulted into 82GB of kern.log and syslog files making /var full.	This issue is resolved. This fix supplies code to check and correct corrupted logrotate status files. If the status file is corrupted, a system file system will eventually fill up, causing loss of service.
95205	C	Dedupe throttling: SD queuing should have separate queue for dedupe to avoid VDBench sequential read testing showing ~40% drop in performance with dedupe running.	This issue is resolved. New disk queue added for Dedupe generated reads to prevent dedupe from excessively holding up other disk read accesses.
102203	C	Retrieval of host netgroups via LDAP implicated in node hangs and prolonged outage after upgrading to 12.1.	This issue is resolved. Fixed to reduce the number of LDAP queries used by default to check a host's netgroup on mount.
104160	C	Need to know when DI is given duff onodes from disk.	This issue is resolved. DI now chips the meta data signature when it reads from its data Cache, and sets a warning assert if it is incorrect.
98000	C	Performance degradation when testing credentials for uid and gid membership for file security.	Fixed to resolve security performance regression.
103763	C	SIGSEGV:Mappings::MappingsContainer< mappings::mappingtraits >::LookupBySID(NTSecurity::SIDHeader</ mappings::mappingtraits	Fixes a crash in the user/group mappings code that could occur during access to the user/group mappings container after an evs going offline.
103867	C	Watch-kitten and loss of clustering due to soap request for local groups.	This issue is resolved. To protect the HNAS server, the API that populates the SMU's local group page will clip the returned results when the total of groups and members reaches 3000.
103924	C	ntsecurity trouble reporter takes unreasonably long.	Resolves issues with the trouble reporter taking too long; it now only looks up names of users and groups when absolutely necessary.
84628	B	File System corrupted with "CFSSnapshotBlocksListPointerIsOfProhib	This issue is resolved. A defect in the **fs-convert-to-support-dedupe** command meant that

		itedType" during converting file system to 8 bitmap.	conversion to support Dedupe was sometimes allowed when it should not have been, because the file system still possessed object-based snapshots that were not fully deleted. The result was the potential corruption of the file system upon mounting it, post-conversion. The defect has been fixed, but note that this does not fix any corruption previously introduced by the defect. Affected file systems require **fixfs** to be run.
100983	B	Recovered fatal on MMB: "referencing after deletion started" panic.	This issue is resolved. Addresses a possible cause of a reset in the SMB2.1 lease (oplock) handling, and adds diagnostics to help with any further problems.

Fixes in SU 12.2.3719.07

Issue ID	Severity	Summary	Explanation
93545	C	warning assert WLOG/M1_INQRSP_COMBSPLIT/fdp_stalled_5s from WlogM1 (MFB1).	This issue is resolved. When NI is continuously very busy (85%+), some RX blocks did not have high enough priority to access the CMEM which in turn stall the TX blocks that are dependent on them.
69417	C	warn - (now) (now) FDP/T2_PROT_BUFF_MGR/fsm_ni_rx_out_of_buffers.	This issue is resolved. The outstanding inquiry alerts would output the volume name in normal scenarios but when the system crashes the alerts would output the device id instead.
90911	C	Unable to backup/restore using DM with XVL unless DAR is disabled with overwrite enabled.	This issue is resolved. Make DAR recovery from tape backups work with externally migrated files.
94680	C	repeated full replication using recreate_links for XVLs and test-before-recreate will delete existing XVLs on the target and their migrated file.	This issue is resolved. It is now harder to misconfigure NDMP file replication when the source and destination both have external migration paths configured.
95910	C	There is insufficient heap: dropping additional connection request from 10.18.222.148:59159->10.18.128.201:445	This issue is resolved. We now drain the fsi cache to bounds before taking such drastic action as denying a connection.
94526	C	DRB contains integer overflow causing out of balance condition.	This issue is resolved. Fixed an issue that could prevent 4KiB formatted file systems larger than 16TiB, or, 32KiB formatted file system larger than 128 TiB from being correctly read balanced.
92519	C	A shared file system lock has been used for an excessive time by "bossock task" (@0x00007F72E526FEA0, id 74, state SW [LOCK:cloudfoundry(shared:707274ms!)acquired]): 11.8 min: this event happened	This issue is resolved. A long standing deadlock has been fixed.

		once since reset on the MMB1.	
93176	C	Universal Migrator: Virtualizing NetApp and attempting to copy SMB ACL's via RoboCopy to HNAS fails with Access Denied.	This issue is resolved. Copying SMB ACL's via RoboCopy to HNAS will not fail.
86992	C	unpooled allocations: long-term: typedef std::map FSAccumulatorsMap.	This issue is resolved. Moved per-filesystem performance accumulators to a heap pool to reduce heap fragmentation.
94041	C	Checkpoint inquiries are being issued to OBJ with an unknown activity source.	Resolved. Creating a V2I backup, which in turn caused deletion of the older snapshots, will not affect performance in the VM infrastructure.
101569	C	HNAS sent DHnQ create context responses with an incorrect length.	When an HNAS EVS was configured to use SMB 2.1 or SMB 3, Microsoft Office for Mac 2011 applications including Word and Excel running on Apple OS X 10.9.5 and 10.10 clients were unable to save files to a share mounted from the HNAS EVS. This issue is resolved. DHnQ create context responses are now the correct length.
100242	C	Namespace links not resolving correctly to cifs-shares when migrated from one node to another during fail-over.	This issue is resolved. An issue with cifs-share deletions not always being correctly propagated round a cluster has been addressed.
100996	C	Segmentation fault reset / Address not mapped to object in NetworkPerfMonitorStore::showLogs()	This issue is resolved. Fixed to avoid unlikely possibility of a reset when gathering a PIR.
96600	C	"MFB vlsi fatal SI/T2_SI_REG/scmem_sdram_ecc_multiple_bit_error" panic occurs when trying to join the cluster.	Previously, when bringing up two nodes in a cluster, node 2 failed to power up; node 1 failed to come up due to the "MFB vlsi fatal SI/T2_SI_REG/scmem_sdram_ecc_multiple_bit_error" panic:
99868	C	ndmputil's SIGPIPE handler can fail with SIGPIPE, monopolizing CPU throughput.	This issue is resolved. ndmputil's SIGPIPE handler will not fail with SIGPIPE.
101105	C	Long term logging of network protocol performance monitor stats all zeros.	This issue is resolved. Previously, network-protocol-performance-monitor-report was only reporting zeros. The correct values are now reported.

Fixes in SU 12.1.3613.06

Issue ID	Severity	Summary	Explanation
89762	C	Heap fragmentation due to extant NFSv4::LockOwnerState objects.	This issue is resolved. NFSv4::LockOwnerState objects are now treated as long-term allocations, making them less likely to cause heap fragmentation.

92535	C	Packets not routed correctly when sent via Eth0 or Eth1.	This issue is resolved. A problem related to the use of Linux networking has been fixed. Linux now bases its routing decisions in the same manner that the server does, for file-serving interfaces.
96333	C	CIFSSecurityDescriptorSerialiser should not modify historic Windows SIDs.	This issue is resolved. For migrated domain accounts, HNAS used to return SIDs from new domain, whereas historic SIDs (from old domain) were expected. This is now fixed.
90685	C	Motherboard memory errors are not being reported correctly.	This issue is resolved. The code that parses a file will not parse an empty line, as a duplicate of the previous line. The empty line is always treated as representing an uncorrectable error.
73900	C	An NDMP fiber(ndmp_subtask_fiber) has timed out on exit waiting for 1 references to be removed: assertion failed (m_action.get() == 0x00007FF8F53E44A0) asserted to be == (0 == 0x0).	This issue is resolved. A fix has been made to prevent an "assertion failure" panic caused by NDMP replication.
79977	C	"assertion failure" panic in function bool FiberOwner::BeforeLock()	This issue is resolved. Previously, restarting an NDMP file replication could cause the system to reset. Restarting an NDMP file replication will not cause an error.
92915	C	Heap fragmented by short-lived >= 32768 B allocations from Max Data Count from SMB1.	This issue is resolved. A potential and subtle source of heap fragmentation for SMB1 customers has been addressed.
91459	C	Warning asserts following snapshot delete after object replication.	This issue is resolved. The server was raising severe event warnings when a snapshot was deleted, when outstanding readahead had been issued on the snapshot. The warning has been downgraded to an info assert.
92067	C	VLSI hung panic with READ_AHEAD target set greater than half the file system size.	This issue is resolved. To prevent this issue, HNAS now restricts the read ahead target value to a maximum of one quarter of the file system size.
92619	C	Specifying a name longer than 30 chars with the NDMP_BLUEARC_USE_SNAPSHOT_RULE env variable causes an "assertion failure" panic	This issue is resolved. Specifying a rule name longer than 30 characters for the NDMP environment variable NDMP_BLUEARC_USE_SNAPSHOT_RULE will no longer cause a panic.

Fixes in SU 12.0.3528.04

Issue ID	Severity	Summary	Explanation
78876	C	Breaking an oplock causes Exception: 0006 Invalid opcode __cxa_guard_acquire in OplockManager::GetWorkTeam	NFSv4 operations which break SMB oplocks could trigger a reset, if two or more such operations happened at the same time.
92312	B	Renaming with SMB2 fails when an	When the variable symlink_caching_support is set

		uncached symlink is involved.	-1 to disable symlink cache, renaming symlinked files can fail over SMB2 for shares with mixed security. This fix fixes the issue.
86807	C	Deadlock NetFileEnum and NetFileClose calls on the srvsvc pipe	Fixes an issue where simultaneous requests to the Srvsvc pipe can deadlock the HNAS.
91422	C	Indirection object reuse count may be non-zero, if formatted by an old enough build; software utilities don't cope well with this.	Software utilities are able to work on file systems formatted in stone.
89140	C	OBJ leaf onode cache should probably somehow scale with FSA cache.	The OBJ FPGA caches leaf onodes for the free space object. The size of this cache now scales with the size of the cache in the FSA FPGA.
92351	B	Reset: Segmentation fault fsb::net::vlsi::FsmEventUdp::Process	Fixed a panic that could occur under some circumstances when receiving a VLAN-tagged UDP packet.
88923	C	NTPSoftware file filtering fails to send FileRename and DirectoryRename notifications.	A problem has been fixed which prevented some FileRename and some DirectoryRename notifications being sent to NTPSoftware QFS for file filtering.
89975	C	ICC software drops the oldest packet when the queue fills old, which probably hurts bluestone performance worse than dropping the newest packet.	The management request transport over the ICC links was too weak to sustain an abnormally parallel load, rendering the management Ethernet a single point of failure in situations where such a load presents.
86084	C	HNAS local user feature not working for entire GE PACS solution. Works with Windows or Jcifs, not both.	The local-password facility has been enhanced to optionally ignore the domain supplied by the client. This is off by default but documented in "man local-password-set".
86654	C	Exception level: 1 "assertion failure" panic SMBFile::AddToTDP()	Fixes an issue where un-setting delete on close on a file can cause the server to crash when SMB1 autoinquiry is enabled.
86596	C	unpooled allocations: long-term: ShareMap::ShareMapType	The server's ability to sustain configuration changes while under file serving load has been marginally improved.
86597	C	unpooled allocations: long-term: Share	The server's ability to sustain configuration changes while under file serving load has been marginally improved.
80646	C	"FC watchdog" panic caused by delay in READDIR response becoming valid	Long system stalls caused by heavy load could cause software to be notified to fetch readdir data before it was valid. Exceptionally long stalls could cause the software to time out, resulting in a crash. It is now impossible for the notification to arrive before the readdir data is valid, so stalls will no longer result in this crash.
88428	C	Filesystem went offline because an OBJ_WRITE from the processor failed	Prevent the filesystem from going offline if an object is deleted, a checkpoint taken, then a

		with ERR_OBJ_HIDDEN_OBJ_TYPE after having modified the filesystem	modifying inquiry (a write, for example) comes in for the deleted object.
90835	C	localnetgroup causes peak in cpu	Fix performance degradation due to localnetgroup update fault.
81491	C	create-group-table-from-active-directory.rb script was barfing on the "--no-warnings" parameter passed to it by this crontab entry	The shipping version of this script now accepts the arguments that are supplied by the shipping crontab.
92575	C	Object Replication needs to extend the indirection object in a fashion that doesn't negatively impact other threads (Backport D80040)	An object replication operation on an internal file system object at the start of the replication session on replication target file systems now shares server resources more fairly with other file systems on the same node.
91997	C	Aruba fails to recover backup stream from Netvault (Bakbone).	Make Aruba accept "" as a valid alternative to "/" in the original_path field of the nlist in the NDMP_START_RECOVER request.

New, modified, and deleted CLI commands

feature-list-resolutions

During a rolling upgrade or downgrade, any features supported only by the newer release must not be used until all servers have been upgraded. If this rule is disregarded, a down-rev server (one that is running an earlier software release) may go into a boot loop until the features in question are no longer in use. The event log on the running server that denies a down-rev machine permission to boot will record the name of the feature causing the problem. You can feed a keyword from this feature name to a new **feature-list-resolutions** command, which will tell you how to stop using the feature for long enough to upgrade the down-rev machine and get the whole cluster booted. The Resolution text of the event logged by the running machine reminds you to run this command.

route-host-add

route-host-delete

route-net-add

route-net-delete

route-gateway-add

route-gateway-delete

route-mtu-add

route-mtu-delete

route-flush

cluster-node-route-host-add

cluster-node-route-host-delete

cluster-node-route-net-add

cluster-node-route-net-delete

cluster-node-route-gateway-add

cluster-node-route-gateway-delete

router-dump

router-dump-active

router-dump-by-evs

These commands are the new way to manipulate the routing table. With multi-tenancy enabled, the route command can no longer modify HNAS routes. The 'router-dump-active' replaces 'dump-router' and 'router-dump-by-evs' replaces 'dump-interface-routes-by-evs' (old names are now aliases for new command names).

smb2-client-side-symlink-handling-

SMB2 protocols support for symlinks was added in 12.2. Whether the HNAS uses this feature, or our legacy method, to present syminks to clients is controlled using the following commands:

smb2-client-side-symlink-handling-default smb2-client-side-symlink-handling-disable smb2-client-side-symlink-handling-enable smb2-client-side-symlink-handling-status

tree-delete-job-submit

Submits a request to delete a directory and all its contents recursively.

Modified CLI commands

nfs-hostname

The nfs-hostname now supports -r or --reset to clear the hostname.

Additionally spaces are 'resolved', the output of a "set" now shows what was actually set, not what was attempted.

Options now changed to -c or --clear

cifs-config-clear

The command now deletes CIFS configuration from a single EVS, previously it deleted it from all EVSes. It requires a valid service vnode to be in context.

cifs-config-list

The command now lists CIFS configuration for a single EVS, previously it listed it for all EVSes. It requires a valid service vnode to be in context. The change is to keep behavior in parallel with its "partner" command cifs-config-clear (above).

span-tune-allocator

span-tune-allocator has been dropped from Dev level to Supervisor level.

virtualization-path-create

New switch "--with-metadata-ext-copy" added to enable a mode for external copy of metadata.

cifs-perf, nfs-perf

The "--type" switch now accepts "throughputcontrol" to display time spent controlling throughput. This is only relevant for throughput-limited models.

route

When multi-tenancy is enabled, the route command can no longer modify HNAS routes.

last-few-packets

Thelast-few-packets packet storing mode toggled by the '-v' switch has been removed. It stored the first 64B of SMB requests for debugging purposes and had been unused for a while.

span-list --sds

When the --sds switch is supplied, span-list now displays more information in the `Set 0:' line: it reports how much space on each stripeset is currently free, and, if the span resides on HDP storage, it says how much of the stripeset consists of space on the vacated-chunks list (which is explained in the hdp man page).

The purpose of the change is to make it easier to determine how to make more space available for file systems on a span that resides on HDP storage. If the stripeset has more space than the (thinly provisioned) HDP pool on which it resides, it will be helpful to add more pool volumes to the HDP pool. Conversely, if the stripesets on a given pool have, in total, *less* free space than that pool, and if (as strongly recommended) the pool is not shared with other spans or clusters or foreign servers, then it'll be necessary to add further stripesets on the same HDP pool in order to use all its available free space.

man shutdown

man shutdown now longer lists the '--app' option. This is a cosmetic change only, and the functionality of *shutdown --app* has not changed. The decision was made to remove the information from the manual page as if used, it places the mercury in a state where it either needs a power cycle or access to a root shell to restart BALI.

irdp

The 'show' option has been extended to also show the network interface on which the routers were discovered.

rip

The 'show' option has been extended to also show the network interface on which the routes were learned.

ftp-cfg, lockd-sync-grant, protocol-character-set

These commands now apply to a Tenant rather than the admin vnode. See multi-tenancy concepts page for more details.

test-route

The 'test-route' command has been extended to allow the outgoing interface to be specified. Also, when multi-tenancy is enabled the source address and the interface are mandatory arguments.

Deleted commands

migration-cloud-stop

This command was replaced with **migration-cloud-schedule-stop-now**.

Documentation

Related documents

NAS Platform product documentation is shipped with your NAS Platform. The Hitachi Data Systems Support Portal (https://portal.hds.com) also contains the most up-to-date documents and troubleshooting information for this product release.

Copyrights and licenses

Notice

Hitachi Data Systems products and services can be ordered only under the terms and conditions of Hitachi Data Systems' applicable agreements. The use of Hitachi Data Systems' products is governed by the terms of your agreements with Hitachi Data Systems.

This product includes software developed by the OpenSSL Project for use in the OpenSSL Toolkit (http://www.openssl.org/). Some parts of ADC use open source code from Network Appliance, Inc. and Traakan, Inc.

Part of the software embedded in this product is gSOAP software. Portions created by gSOAP are copyright 2001-2009 Robert A. Van Engelen, Genivia Inc. All rights reserved. The software in this product was in part provided by Genivia Inc. and any express or implied warranties, including, but not limited to, the implied warranties of merchantability and fitness for a particular purpose are disclaimed. In no event shall the author be liable for any direct, indirect, incidental, special, exemplary, or consequential damages (including, but not limited to, procurement of substitute goods or services; loss of use, data, or profits; or business interruption) however caused and on any theory of liability, whether in contract, strict liability, or tort (including negligence or otherwise) arising in any way out of the use of this software, even if advised of the possibility of such damage.

The product described in this guide may be protected by one or more U.S. patents, foreign patents, or pending applications.

Notice of export controls

Export of technical data contained in this document may require an export license from the United States government and/or the government of Japan. Contact the Hitachi Data Systems Legal Department for any export compliance questions.