

Hitachi Unified Storage (HUS) File Module NAS Operating System SU 12.3 and SMU 12.3 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. Server update (SU) 12.3.3826.07 contains the latest defect fixes, only. Features in the previous HNAS version 12.3.3826.03c are included in these release notes, as well as their compatibility with HNAS server platforms, and whether or not a license must be purchased to obtain the functionality. The HUS File Module NAS operating system, which includes SU 12.3.3826.07 and SMU 12.3.3826.04, supports the HUS File Module M1 and M2 models. NAS operating system, which includes server update SU 12.3.3826.07 and SMU 12.3.3826.04, 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

Intended audience.....	1
Additions or changes	1
Getting help	1
Performance Improvement after a File Clone is deleted.....	4
HUVM/GAD: Support for automatic selection of preferred paths	4
DM2CD: Support for using Aggregate ports	5
Bitmap readahead: Improvements to write performance in a File System.....	6
Northern NSS API Support	6
Active Directory Support for SMU	7
AES Support for SMB Authentication.....	7
1PiB File System Support	8
Native SNMPv3 support.....	8
Support for new providers in Data Migrator to Cloud.....	8
Support for Block on next generation VSPGxxx storage.....	9
Improved Anti Virus and Bossock Fiber Behavior	9
Important note on downgrading from 12.x versions.....	9
Note on supported AMS storage arrays	10
Note on data spillage from tier 1 to tier 0	10
Note on using Hitachi Dynamic Provisioning on HNAS before and after v12.1	10
Important note on SMU code release 12.3.3826.04	11
SMU, server, and cluster compatibility	11
Upgrade Path Flowchart	13
Important note on upgrading from older versions of HNAS: 7.x, 8.x, 10.x	13
Licensing.....	14
Fixes in SU 12.3.3826.07	15
Fixes in SU 12.3.3826.03a	17
Important known issue on the global-symlinks-add CLI command	18
Related documents.....	21
Copyrights and licenses	21

Document history

The following changes have been made to this document.

Revision	Description
RN-92USF024-00	Initial publication for server update (SU) 12.3.3826.03a, and SMU 12.3.3826.02
RN-92USF024-01	Updated for SU 12.3.3826.03c
RN-92USF024-01	Updated to include defects fixes in 12.3.3826.07

New features at a glance

SU 12.3.3826.07 contains no new features; however, the following table lists the features included in the previous HNAS version 12.3.3826.03, compatibility with HNAS server platforms, and whether or not a license must be purchased to obtain the functionality. These features are explained in greater detail in this document.

Feature	HNAS 30x0 G1	HNAS 30x0 G2	4040	4060	4080	4100	Licensed
Performance Improvement after a File Clone is deleted	✓	✓	✓	✓	✓	✓	
HUVM/GAD: Support for automatic selection of preferred paths	✓	✓	✓	✓	✓	✓	
DM2CD: Support for using Aggregate ports	✓	✓	✓	✓	✓	✓	
Bitmap readahead: Improvements to write performance in a File System	✓	✓	✓	✓	✓	✓	
Northern NSS API Support	✓	✓	✓	✓	✓	✓	
Active Directory Support for SMU	✓	✓	✓	✓	✓	✓	

Enabling and disabling file system capacity and free space reporting based on virtual volume quotas	✓	✓	✓	✓	✓	✓	
AES Support for SMB Authentication	✓	✓	✓	✓	✓	✓	
1PiB File System Support				✓	✓	✓	
Native SNMPv3 support	✓	✓	✓	✓	✓	✓	
Support for new providers in Data Migrator to Cloud	✓	✓	✓	✓	✓	✓	
Support for Block on next generation VSPGxxx storage			✓	✓	✓	✓	

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. In addition to the new features, SU 12.3.3826.03c includes improvements to Anti-Virus and Bossock Fiber behavior. Details are explained below, in a new topic at the end of this New Features section.

Note: The 3080/3090 series is *not* supported on the VSPGxxx.

Performance Improvement after a File Clone is deleted

During a file backup process a file clone is created. Multiple backups – and multiple clones – of the same file can exist. For example, a live file may have a clone that was created a day ago, and another clone that was created two days ago, and another clone that was created three days ago, etc.

Deletion of any clone is always performed the same way except for the deletion of the most recent clone and/or the oldest clone. Deletion of the most recent clone and/or the oldest clone is a much more complex operation within the HNAS server.

Previously, the complex operation of deleting the most recent clone might have caused performance problems when accessing the live clone (while deleting the most recent clone). This enhancement addresses these performance problems so that file backup will not cause system-performance issues.

HUVM/GAD: Support for automatic selection of preferred paths

The VSP G1000 and VSP GX00(series) can be configured to use a feature which manages the primary/secondary relationship in a mirror configuration, internally, exposing a virtual

Global-Access Device (GAD) volume to the host (HNAS). To HNAS, the GAD volume presents as a normal simplex SD. Failover of the mirror relationship is handled internally by the array and is thus transparent to HNAS.

GAD volumes in a metro cluster can be accessed by either node in the cluster. However, attempting to write to the GAD volume using a path to the array that is physically hosting the secondary volume of the pair causes increased overhead, as the write request will have to be forwarded to the array that hosts the primary volume.

To avoid this overhead, HNAS has implemented preferred path support, so that HNAS will use the preferred path as presented from the array to access the array that hosts the primary volume. The main implication of implementing GAD preferred path support is that there is no requirement to use the `sdpath` command to force HNAS to use the optimal path to the storage.

DM2CD: Support for using Aggregate ports

Release 12.3 features new functionality that leveragages the HNAS aggregate ports for data transfer between file systems and cloud targets. Management ports are still offered, but aggregate ports provide better and faster performance.

- The key interface is setting up the networking, which is not different from setting up the aggregates network.
- The default configuration should be sufficient in most cases. The default mode, for customers who never used DM2C before, will be to send data via aggregate ports with zero copy for local HCP cloud target.
- For customers that have cloud setup configured in a current version and have upgraded to 12.3 firmware, the default will continue to be management ports. They have an option to switch to use aggregate ports, but before doing that they should configure the network properly and then execute the following command:

```
migration-cloud-service-mode-set -a
```

Or, if interested in local HCP migrations, use the same command with `-z` option:

```
migration-cloud-service-mode-set -z
```

- In some special cases, for example when multi-tenancy feature or vlans are used, switching back to management ports may be required to be able to run DM2C. Then the following command should be executed:

```
migration-cloud-service-mode-set -m
```

- The system should be smart enough to switch to management ports when multi-tenancy is enabled; however, vlans are not checked.

Bitmap readahead: Improvements to write performance in a File System

The HNAS server has an internal logic to optimally select FreeSpace to write to, whenever a write request is sent to a file system. Extremely full regions of the file system take a long time to search for FreeSpace, leading to possible checkpoint dead time.

In this release, improvements are made to reduce that dead time and to improve write performance. This enhancement is recommended for customers with write performance issues on file systems with areas that are extremely full. Checkpoint dead time and write performance improve using this enhancement.

Note: In a prior version, bitmap readahead only worked on single stripesets. In 12.3, bitmap readahead works on multiple stripesets.

Configuration and use

Users can check for dead time with the following command:

```
checkpoint-stats -s -f <file_system>
```

If the CLI output is "time for the FSA_CREATE_CHECKPOINT to complete" >1s, you can check for a dead time cache miss with:

```
fsa-file-system-stats -f <file_system>
```

If the CLI output shows "GET_BLOCK_IN_DEAD_TIME_CACHE_MISS", then bitmap readahead may help.

This feature is disabled by default. To enable it, perform the following:

```
sdg-allocate-control-mode --mode READ_AHEAD -f <file_system>
```

Also, set cache bias to "auto" (default is large files):

```
cache-bias [--small-files | --large-files | --auto]
```

Northern NSS API Support

NSS (Northern Storage Suite) is a Windows-based software solution that can be used on any SMB compliant device. It enables organizations to control the use of storage resource, specifically the way in which users exploit this resource. It can monitor and control storage on multiple devices across the organization.

The API is intended to integrate NSS with HNAS and is to be used to facilitate:

- **Quotas** - to restrict the amount of storage used by individuals. Although quotas can be set on the HNAS, some organizations have a wider scope for their quota approach and so need to have a system which can access all storage locations, aggregate an individual's usage, and monitor for compliance with the quota.
- **File-blocking/file-filtering** - Prevent creation of and/or access to certain objects which the organization wishes to block for whatever reason. E.g. prevention of adding music files onto servers which would cause the organization to be in breach of copyright.

Active Directory Support for SMU

Active Directory is an LDAP-compliant hierarchical database of objects. It is very popular in enterprise environments and is becoming a de facto standard for user authentication.

Once Active Directory connection settings and groups have been configured for the SMU, it will allow logins from enabled users who supply their Active Directory name and password. This is typically the same name and password that the user would use to log into Windows and other enterprise applications. Unlike SMU local and RADIUS user names, Active Directory user names are case-insensitive. Active Directory passwords are case-sensitive and cannot be changed from the SMU; they are maintained in the Active Directory server.

There are a number of benefits for SMU users. The administrator does not need to maintain a separate set of user details, because the SMU can just make use of the Active Directory enterprise user database. Users can login using their usual name and password instead of having to remember a separate set of credentials for the SMU. And instead of configuring access for individual users, the SMU administrator just has to specify the Active Directory groups whose members have login rights.

Enabling and disabling file system capacity and free space reporting based on virtual volume quotas

The file system capacity and free space reporting for virtual volume quotas option supports thin provisioning within a virtual volume. When this option is enabled and a virtual volume quota is created, capacity/free space counts returned to clients are derived solely from the virtual volume quota. This affects only those clients that have mounted an export or share within a virtual volume.

You may want to enable this option when data migration is configured. In this scenario, the primary file system could ingest more data than it has capacity itself for. You can define a quota for a virtual volume based on available capacity of migration target(s) and enable this feature so that the capacity defined by the quota is reported to protocol clients rather than the primary file system capacity/free space.

Note: This option is disabled by default.

AES Support for SMB Authentication

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
- AES256-CTS-HMAC-SHA1-96

1PiB File System Support

The maximum size of a WFS-2 file system is increased to 1PiB, but note that a 1PiB file system is currently *only* supported on an HDP storage pool. This feature adds support for file systems up to 1PiB in size, up from 256TiB. It is a WFS-2-only feature; WFS-1 file systems continue to follow the 256TiB limit. This feature is automatically enabled in SU 12.3, running on 4060/4080/4100 models.

This feature does not require any new configuration and is used per file systems in earlier releases. Manual and auto expansion allow existing file systems to grow bigger than 256 TiB, up to 1PiB. This improvement has no impact to file system performance, but note that file system utilities (`checkfs`, for example) will take proportionally longer to complete on large file systems.

Note: Any file systems larger than 256 TiB will not be visible in older releases.

Native SNMPv3 support

SNMPv3 defines a secure version of SNMP compared to the previously supported SNMPv1 and SNMPv2c. SNMPv3 adds support for user based authentication and encryption to achieve secure access to the management information held on the HNAS server. SNMPv1 and SNMPv2c continue to be available but cannot be enabled at the same time as SNMPv3.

All notifications are sent using SNMPv1 traps regardless of the configured SNMP protocol version.

Support for new providers in Data Migrator to Cloud

Microsoft Azure and Hitachi Cloud Services have been added as new cloud providers for this release.

Microsoft Azure, you must have:

- An established Microsoft Azure storage account.
- The credentials of a user with read/write permissions to the target.

Hitachi Cloud Services, you must have:

- Fully qualified domain name of the Hitachi Cloud Services namespace for the account credentials
- The credentials of a user with read/write permissions to the target.

Support for Block on next generation VSPGxxx storage

Support for block storage on new VSPGxxx series arrays (G200, G400 and G600) by the HNAS 4xxx series platforms is enabled by default for this release.

Note: The 3080/3090 series is not supported on the VSPGxxx.

Improved Anti Virus and Bossock Fiber Behavior

Virus scanning and file filtering necessarily involve calls to a remote third-party server; such a system can take many seconds to return a result, locking up Bossock fibers for the duration. The fiber-exhaustion-protector-percentage limits the number of fibers that may be blocked, waiting for responses from a third-party sever. The entire pool of fibers is capable of waiting for responses from a third-party, but the total allowed to be doing so at any moment is limited.

In previous versions, only 2 fibers were reserved for this purpose. In this version, this reservation has been increased—from two to approximately 1/4 of all fibers. The value is field-adjustable, but at least two are always reserved.

No fiber is given a specific role or privilege; the ability of any fiber to block, waiting for a response is gated only by the number of fibers already waiting for responses at that moment.

Important considerations to read before installation

Please read the following sections before installing and using SU 12.3.3826.07.

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-file systems

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 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 dividing 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.

Important note on SMU code release 12.3.3826.04

This release uses SMU code release 12.3.3826.04, and it requires the CentOS 6.2 operating system. Before you install the 12.3.3826.04 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 12.3.3826.03c 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.

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

SMU	CentOS Version	Supported Server Releases	Number of Supported Servers / Cluster	
			SMU 300	SMU 400
		HNAS OS		
SMU 8.1	4.8.1	7.0, 8.2	5	Not supported
SMU 10.1, 10.2	6.2	8.1, 8.2, 10.0, 10.2	5	5
SMU 11.x	6.2	8.1, 8.2, 10.2, 11.x	5	5
SMU 12.x	6.2	8.1, 8.2, 10.2, 11.x, 12.x	5	5

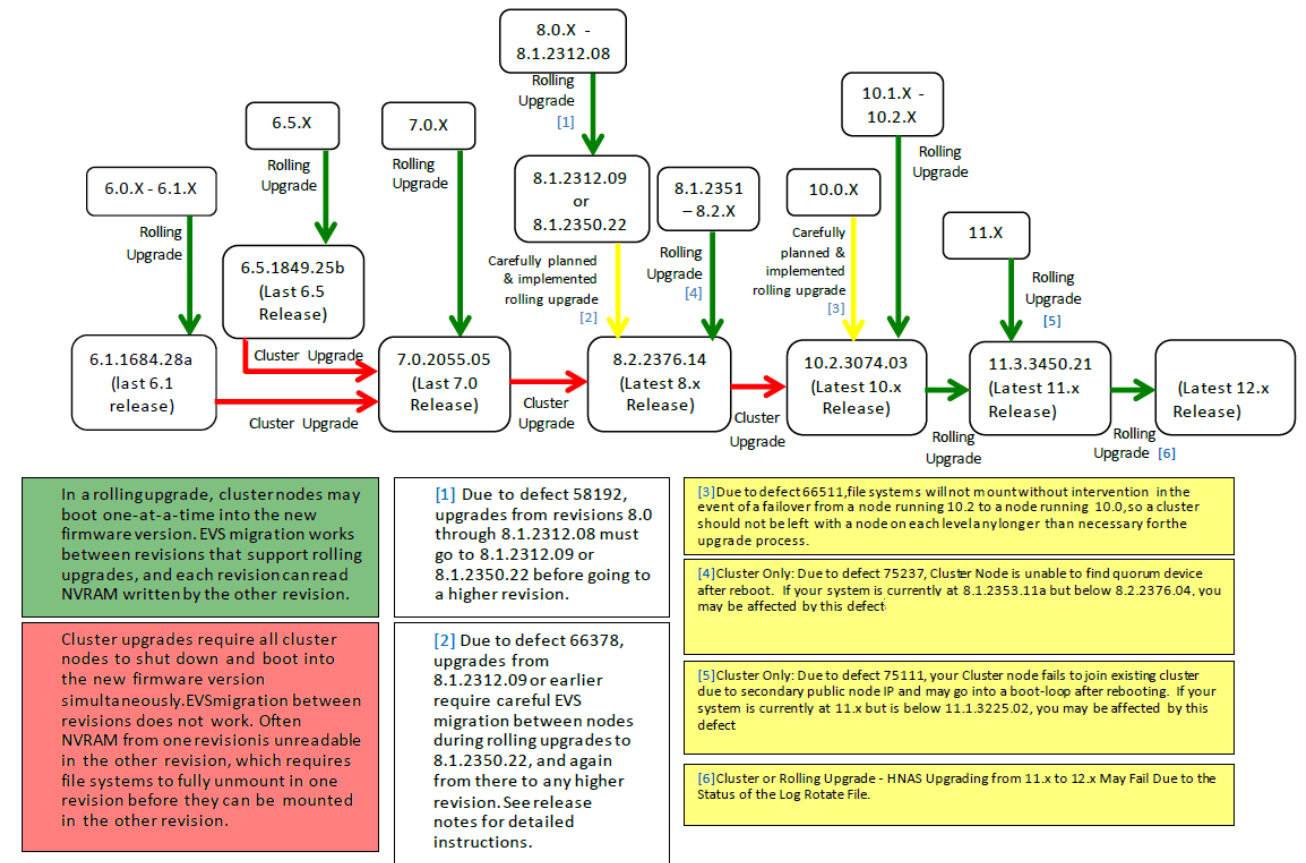
Notes on installing, upgrading and downgrading

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](#)

Important note on upgrading from older versions of HNAS: 7.x, 8.x, 10.x

If you are upgrading to version 12.3 from these earlier versions of HNAS, please note that there are critical steps, including rolling upgrades, that must be followed in a precise sequence to correctly upgrade to version 12.3. Please refer to the corresponding release notes of each earlier version for details on rolling upgrades. Additionally, please consult with your HDS representative for assistance in upgrading from earlier versions of HNAS.

Licensing

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

In order to obtain license keys for new, licensed features, note that:

- 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 SU 12.3.3826.07

Issue ID	Severity	Summary	Explanation
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.
94052	C	Start of “assertion failure” panic in ‘fsb/dedupe/main/lib/DedupeQueue.cpp’	This issue is resolved. Manually disabling and enabling the Dedupe service looks like the most common trigger, and should be avoided. Restarting and mounting the file system will clear the inconsistent dedupe config.
1077991	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.
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.
108124	C	fatal - (now) (now) WLOG/T2_WLOG_REG/wlog_stuck_block – caused by corrupted buffer manager free queue in NI	This issue is resolved. Fixed NI VLSI buffers initialization.
108402	C	vlan-convert-config.rb fails.	This issue is resolved. Previously, when the 12.x VLAN script checked for the minimum supported HNAS version, it failed because the script was written to check versions ending in numbers, only. It expected a dot-separated numeric version, so it was failing to parse the software version due to the “a” at the end of the version number. The script is now corrected to check letters in the version string.
98861	C	LongTermSizeBasedPool-33554440 pool has been tuned to its maximum capacity, which is 0 items (0 B).	This issue is resolved. This fixes a problem in the handling of SMB2 and SMB3 SessionSetup requests which could allow

			large numbers of unauthenticated sessions to build up, eventually causing a FailedToAllocate panic.
56700	C	Node became wedged when truncating a 570GB file to 0.	This issue is resolved. Feature improves performance of truncate to zero for files with length greater than 'background-truncate-to-zero-threshold' (default value 5242880 blocks which is 160G in 32K file system and 20G in a 4K file system). Feature enabled by default in v12.3 onwards. Feature can be controlled from 'background-truncate-to-zero' config.
109170	C	CSpnegoServerCodec::ExtractNegTokenInit() should check all "mech" types for a match, not just the first one	This issue is resolved. The HNAS will now negotiate a mutually supported authentication mechanism when the client prioritizes mechanisms that the HNAS does not support.
108622	C	bytesUsed showing incorrect value when truncating an existing file	This issue is resolved. The cached value of disk space occupied by an object will now be correctly updated when a file is truncated.
102724	C	warning assert WDIR/T2_WDIR_MAIN/severe_event_name_crc_mismatch from WdirM1 (MFB1): No debug info regards FS or volume.	This issue is resolved. 'severe_event_name_crc_mismatch' warning was being seen on the console, with no indication of why it was happening, or which files and file systems were affected. This has now been fixed, so the request which triggered the error is dumped to the console.
109918	C	Mustn't attempt encroachment I/O on unopened secondary SDs.	This issue is resolved. Fixed a possible deadlock caused by 'sd-list --has-good-encroachment' on systems with mirrored storage (e.g. TrueCopy or HUR).
100821	C	Segmentation fault in "Bluestone bossock" while upgrading one node from 11.3.3450.18 to 12.3.3813.00.	This issue is resolved.

Fixes in SU 12.3.3826.03

Issue ID	Severity	Summary	Explanation
83869	C	fatal asserts SSI/T3_SI_CACHECTRL/vm_table_lookup_in_infinite_loop SSI/T3_SI_CACHECTRL/sim_inq_volume_offset_beyond_end_of_volume while fs-checkpoint-health runs	This issue is resolved. <code>fs-checkpoint-health</code> will not panic if bitmap block pointers are outside file system range.
100209	C	HSR scheduler/policies lock ordering violation deadlock.	This issue is resolved. Fixed a possible deadlock when deleting an object replication policy with at least 1 schedule when the object replication scheduler is running at the same time.
97376	C	Node1 Failed to allocate 1.31 MB (1376264 B) of heap.	This issue is resolved. The scalability of the SNMP TCP connection doobrey has been improved.
78807	C	MMB1 .../main/libs/bluestone/core/lock.cpp:987: in function int BS::rtl_destroy(const char*)	This issue is resolved. Improved handling of race conditions during NDMP tape backups.
93509	C	"assertion failure" panic in function static void DLM::DLM::registerResourceOwner(DLM::ResourceOwner&)	This issue is resolved. Fixed an issue where restarting an EVS in quick succession could (rarely) cause a crash while mounting pseudo file systems. The fix also applies to file systems.
100994	C	"invalid netmask" panic in function void net::ipv4::ensureContiguousNetmask triggered by wildcard in cifs-dc filter	This issue is resolved. <code>cifs-dc</code> addfilter now only supports well-formed subnet wildcards. Non-compliant configuration is silently discarded on upgrade.
103629	C	CHD-1 became unresponsive and had to be reset (RST button)	This issue is resolved. Code modified so that deadlock can no longer occur.
83952	C	Sparse Metadata: Able to create sparse file beyond 256TiB limit.	This issue is resolved. It is now impossible to create sparse files beyond 256TiB limit, even if an incorrect error to the NFS client is returned.
100293	C	Lost cluster after eth1 IP change.	This issue is resolved. Adding a pnode IP address on ag ports is no longer allowed.
96770	C	SIGSEGV in net::lACP::AggregatorContext::Agg	Fix to the <code>agg-upgrade</code> command, which could previously cause valid aggregation configuration to be removed and trigger a panic when additional aggregation configuration changes were made.
77801	C	Dedupe: Object Replication -File System: A shared file system lock has been hogged for an excessive time by "HSRChangedObjectReader"	This issue is resolved. Object replication now suspends and resumes the file system context on the source file system when iterating over large changed regions.

Important known issue on the global-symlinks-add CLI command

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.
97524	C	Recovered fatal on MMB: "assertion failure" panic at /home/builder/angel/3450.10/main/fsb/viruscan/Manager.cpp:434: in function FSBStatus VirusScan::Manager::registerServer(VirusScan::ServerPtr).	This issue is resolved. Fixes an assertion with RPC based AV clients that write to the named pipe using SMB1.
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.
97753	C	Virtualization: unbounded allocation can cause a crash if HNAS is under a lot of pressure.	This issue is resolved. Fix prevents panic caused by running out of memory. It is advisable not to use as unbounded allocation with Universal Migrator.
98773	C	Unable to delete snapshots; ERR_OBJ_OBJ_INQ_REUSE_COUNT_TOO_LOW	This issue is resolved. Snapshot deletion will no longer fail with ERR_OBJ_OBJ_INQ_REUSE_COUNT_TOO_LOW, on 6.x-formatted file systems.
100500	C	The SMB2 server calls into the file system without a file system context: smb::v2::FileSystemContext::begin(smb::v2::FileSystemContextResources const&	This issue is resolved. A cause of server resets when accessing files via SMB has been corrected.
106470	B	fatal assert SI/T2_SI_REG/perx_stats_fatal_error from SiM2	This issue is resolved. 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.

Important known issue on the global-symlinks-add CLI command

When `global-symlinks-add` is set, the data cannot be found in the original location, after rename.

The command, `fsm set cifs-dfs-relax-server-validation true`, is required as the workaround. The HNAS documentation will be enhanced in SU 12.4 to update the CLI reference and man page for the FSM SET <parameter>, to disable strict checking of GET_DFS_REFERRAL requests.

New, modified, and deleted CLI commands

New commands

span-hdp-thickly-provisioned
migration-cloud-service-mode-set
migration-cloud-service-mode-show
cifs-keytab-list
span-encode-tiers, span-decode-tiers
span-csv-list, span-csv-file systems, sd-csv-with-hitachi-ids, filesystem-mark-unmounted-benignly
fs-space-reporting-based-solely-on-vivol-quota
for-each-vivol
tree-delete-job-list
tree-delete-job-status
tree-delete-job-abort
tree-delete-job-reschedule
file-filtering-on
file-filtering-off
file-filtering-state
file-filtering-reset
file-filtering-deregister
file-filtering-server-list
file-filtering-statistics-show
file-filtering-statistics-reset
file-filtering-failures-show
file-filtering-failures-reset
snmpv3-user-add
snmpv3-user-list
snmpv3-user-delete
snmpv3-user-delete-all

Modified commands

cifs-name

connection

cifs-share

audit-log-perf

tdp-profile

touch

span-create, span-expand, autoserver

autoserver

ntp

filesystem-create, filesystem-expand, format

profile

Commands for CoW and Thin Image

Mirroring commands now set secondary LU tiers automatically

file-filtering

file-filtering-perf

file-filtering-failures

span-delete

protocol-character-set

thread-stack-trace (bt)

span-create and span-expand now treat SD groups as optional

licensekey syntax/usage made clearer

snmp-protocol

sdg-allocate-control-mode

domain-mappings-modify

migration-cloud-account-create

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

© 2011-2015 Hitachi, Ltd. All rights reserved.

No part of this publication may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying and recording, or stored in a database or retrieval system for any purpose without the express written permission of Hitachi, Ltd.

Hitachi, Ltd., reserves the right to make changes to this document at any time without notice and assumes no responsibility for its use. This document contains the most current information available at the time of publication. When new or revised information becomes available, this entire document will be updated and distributed to all registered users.

Some of the features described in this document might not be currently available. Refer to the most recent product announcement for information about feature and product availability, or contact Hitachi Data Systems Corporation at <https://portal.hds.com>.

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

Hitachi is a registered trademark of Hitachi, Ltd., in the United States and other countries. Hitachi Data Systems is a registered trademark and service mark of Hitachi, Ltd., in the United States and other countries.

Archivas, BlueArc, Dynamic Provisioning, Essential NAS Platform, HiCommand, Hi-Track, ShadowImage, Tagmaserve, Tagmasoft, Tagmasolve, Tagmastore, TrueCopy, Universal Star Network, and Universal Storage Platform are registered trademarks of Hitachi Data Systems Corporation.

AIX, AS/400, DB2, Domino, DS8000, Enterprise Storage Server, ESCON, FICON, FlashCopy, IBM, Lotus, OS/390, RS6000, S/390, System z9, System z10, Tivoli, VM/ESA, z/OS, z9, zSeries, z/VM, z/VSE are registered trademarks and DS6000, MVS, and z10 are trademarks of International Business Machines Corporation.

All other trademarks, service marks, and company names in this document or website are properties of their respective owners.

Microsoft product screen shots are reprinted with permission from Microsoft Corporation.

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.