



Best Practices for Hitachi NAS Universal Migrator

Francisco Salinas, Global Solution Services



MK-92HNAS047-00

© 2014 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, 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.

Notices and Disclaimer

The performance data contained herein was obtained in a controlled isolated environment. Actual results that may be obtained in other operating environments may vary significantly. While Hitachi Data Systems Corporation has reviewed each item for accuracy in a specific situation, there is no guarantee that the same results can be obtained elsewhere.

All designs, specifications, statements, information and recommendations (collectively, "designs") in this manual are presented "AS IS," with all faults. Hitachi Data Systems Corporation and its suppliers disclaim all warranties, including without limitation, the warranty of merchantability, fitness for a particular purpose and non-infringement or arising from a course of dealing, usage or trade practice. In no event shall Hitachi Data Systems Corporation or its suppliers be liable for any indirect, special, consequential or incidental damages, including without limitation, lost profit or loss or damage to data arising out of the use or inability to use the designs, even if Hitachi Data Systems Corporation or its suppliers have been advised of the possibility of such damages.

This document has been reviewed for accuracy as of the date of initial publication. Hitachi Data Systems Corporation may make improvements and/or changes in product and/or programs at any time without notice. No part of this document may be reproduced or transmitted without written approval from Hitachi Data Systems Corporation.

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.

Document Revision Level

Revision	Date	Description
MK-92HNAS047-00	May 2014	First publication

Contributors

The information included in this document represents the expertise, feedback, and suggestions of a number of skilled practitioners. The author would like to recognize and sincerely thank the following contributors and reviewers of this document.

- HNAS Engineering

Contact

Hitachi Data Systems

2845 Lafayette Street

Santa Clara, California 95050-2627

<https://portal.hds.com>

North America: 1-800-446-0744

Table of Contents

Intended Audience	1
Overview	1
Hitachi NAS Universal Migrator	1
How Universal Migrator functions	3
Availability of Universal Migrator	4
Preparing to use Hitachi NAS Platform Universal Migrator	5
Legacy NAS NFSv3 export settings.....	5
File system block size selection	5
Space requirements for virtualization.....	5
Preparing HNAS exports and shares	6
HNAS virtual volumes	7
Performing a test virtualization.....	7
Universal migrator performance expectations	8
Heavy client access during virtualizations	9
Perpetual virtualization.....	10
Monitoring virtualization and migration performance	11
Caveats and limitations	12
Caveats	12
Snapshots	12
Root squashing while using Universal Migrator	12
Open files	12
Shortnames	12
LNAS character set	13
HNAS bossock fibers	13
Security operations on virtualized files.....	13
LNAS file system capacity tracking while in the virtualization phase.....	13
Changed Directory List (CDL).....	13
Using another HNAS file system as the LNAS	13
Migration of non-NFSv3 accessible data	13
SMB access to virtualized data	13
Nesting and overlapping of exports on the LNAS	13
Hard and soft links	14
Handling failed files during virtualizations and migrations	14
Reverting back to the LNAS.....	14
Interoperability with other Hitachi NAS Platform features.....	15
Frequently asked questions	16
Terminology	17

Intended Audience

This guide is intended for customers, authorized service providers, and Hitachi Data Systems (HDS) personnel.

Overview

The Hitachi NAS Universal Migrator (UM) feature provides customers with a convenient and minimally disruptive method to migrate from their existing NAS system to the Hitachi NAS Platform. The practices and recommendations outlined in this document describe how to best use this feature. For detailed instructions on how to configure and set up Hitachi NAS Platform Universal Migrator, refer to the *Hitachi NAS Platform Data Migrator Administration Guide* located in the help section of the Server Management Unit (SMU) Web Manager GUI.

Hitachi NAS Universal Migrator

Universal Migrator is designed to migrate data from an existing NAS system, referred to as the “Legacy NAS” or LNAS, to a Hitachi NAS Platform (HNAS). Universal Migrator accomplishes the migration by first virtualizing a file share on the existing NAS system and then copying the data to the HNAS WFS-2 file system. The movement of data is performed in the background and is transparent to clients connected to the virtualized exports/shares. The data is available for access immediately after the virtualization process starts.

A high level summary of the data migration process is described below and illustrated in Figure 1.

1. NFSv3 clients are accessing exports on the LNAS
2. HNAS is configured to virtualize export(s) on the LNAS.
3. While HNAS is configured to virtualize exports on the LNAS (step 2), the clients are reconfigured to access the virtualized data via an HNAS NFSv3 export
4. While the clients are accessing the virtualized data, HNAS migrates the data onto the HNAS file system in the background either immediately or per schedule.

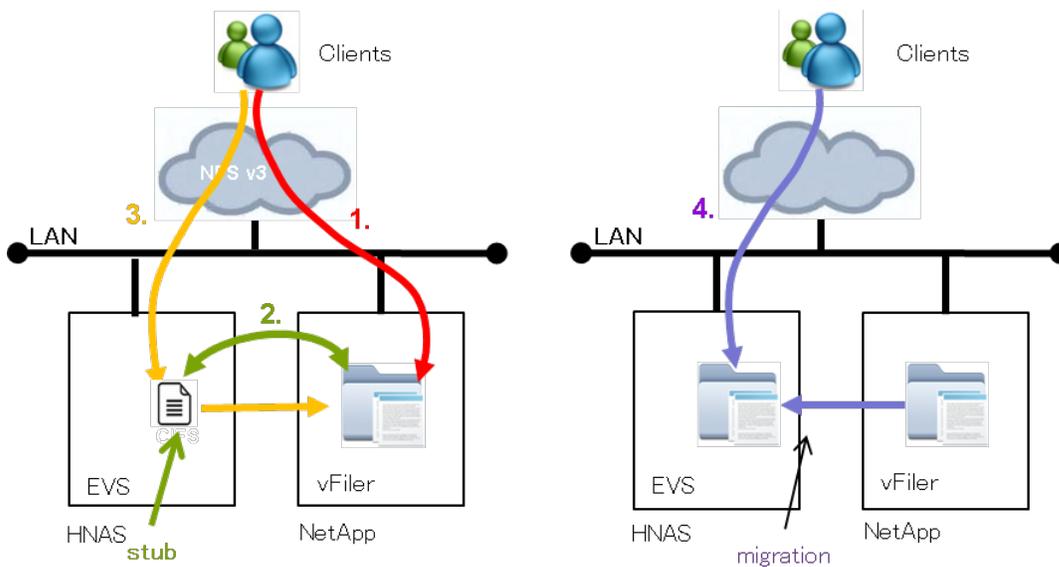


Figure 1. Universal Migrator migration stages.

Hitachi NAS Universal Migrator significantly reduces the long cut-over times associated with traditional data migration using utilities such as *rsync*. The table below lists the typical stages an administrator or consultant would follow to perform a data migration and compares them to stages required to perform a migration using Universal Migrator.

Traditional Migration	Duration	Universal Migrator	Duration
Set source exports read-only	Minutes	Set source exports read-only ¹	Minutes
Perform final incremental rsync	Hour(s)	N/A	-
N/A	-	Start virtualization	Minutes
Cutover to new exports	Minutes	Cutover to new exports	Minutes
N/A	-	Data is migrated in the background	Hours ²
Remove source exports	Minutes	Remove source exports	Minutes
Total Cutover Time	Hours	Total cut-over time	Minutes

¹HNAS must be the only host that has read-write access to the source exports.

²The data migration duration depends on numerous factors, but in the case of Universal Migrator, the cut-over time is not affected; therefore, it is not accounted for in the total cut-over time.

In the table above, the biggest contribution to cut-over time is typically the incremental job. *With Universal Migrator, there is no incremental replication.* The cut-over time then is significantly reduced and the data is migrated transparently in the background.

How Universal Migrator functions

This section describes how Universal Migrator virtualizes and migrates data from a Legacy NAS system.

To virtualize an NFSv3 export, the administrator creates an association to pair an export on the LNAS with a directory on a HNAS file system. For example:

- A Linux file server has four NFSv3 exports, exp1, exp2, exp3 and exp4.
- The administrator creates a new file system on HNAS, and then creates four directories. For example, /lnas-exp1, /lnas-exp2, /lnas-exp3, /lnas-exp4. The directory names are not important but should be something meaningful.
- The administrator then associates each export on the LNAS with a corresponding directory on the HNAS file system. For example:

exp1 -> lnas-exp1

exp2 -> lnas-exp2

exp3 -> lnas-exp3

exp4 -> lnas-exp4

Note: It is also possible to create Virtual Volumes on HNAS, but not required. See [HNAS virtual volumes](#) for more information.

During the creation of the association, HNAS mounts the NFSv3 export as root. At this point, HNAS does not access the data or verify it has permissions. An error message will display if the mount fails. Next, the administrator will create four NFSv3 exports on HNAS, with each export referencing one of the newly created directories. The administrator is now ready to start the virtualization process and cut over to the HNAS exports. Before doing this, the administrator needs to perform the following:

- On the LNAS, the exports are set to read only, with only the IP/hostname of the HNAS EVS(s) having read/write and root access.
- Un-mount the 4 exports on the clients.
- Update DNS; replacing the IP of the LNAS with the IP of the HNAS EVS.
- Start the virtualization process.
- Mount the HNAS exports on the clients.

Note: It is strongly recommended to perform a test virtualization run and validate that no permission, character set, name clashes, or network issues exist before attempting a migration in production.

During the virtualization phase, Universal Migrator discovers all files and directories located on the designated LNAS export path. Universal Migrator leverages multiple threads to walk the LNAS export. This will potentially generate many NFSv3 metadata operations on the LNAS to obtain directory and file detail. The threads are pooled and shared between any other running virtualizations and/or migrations on that HNAS node. The default number of threads is 20 and can be altered at any time. See [Universal Migrator Performance Expectations](#) for details.

As the threads encounter directories, they are migrated to the HNAS file system. When files are encountered, special stubs are created on the HNAS file system. The stub contains all the necessary information required (tracking metadata) to access the file it points to on the LNAS. These stubs are called "IngestedFiles". These look like normal files to the clients. While the virtualization is occurring, it is possible for clients to start accessing the data from an HNAS NFS export and/or SMB share. If a client requests or browses to a directory that has not yet been virtualized, HNAS will pull the required information and create the necessary directories/stubs to process the client request. This process occurs on-demand so that the client is guaranteed to have a fully functioning file system once the virtualization process begins.

While in virtualization mode, HNAS will forward all operations to the LNAS including operations that create new files (which are created on the LNAS).

Once the virtualization phase completes, the migration phase can optionally be started. The exports can also remain in virtualization mode for an unlimited period of time, though this is not recommended by Hitachi Data Systems Global Solutions Services. While in virtualization mode, the administrator could revert back to the LNAS without having to worry about synchronizing data on the LNAS and HNAS because the LNAS holds all the data.

When the migration phase starts, HNAS will begin migrating (copy) files from the LNAS onto the local HNAS file system. The migration is transparent to the clients and the clients maintain full read/write access to files during migration. All writes are forwarded to the LNAS and also written locally until the file is migrated in whole. Once a file is fully migrated, HNAS removes the files tracking metadata.

Note: In the migration phase, security modifying operations are not forwarded to the LNAS

After the migration completes, the associations are removed and the LNAS system can be retired.

Availability of Universal Migrator

Hitachi NAS Platform Universal Migrator is a licensed feature available starting in HNAS release 11.3.3450.

Preparing to use Hitachi NAS Platform Universal Migrator

This section details considerations and best practices for successfully migrating data onto the Hitachi NAS Platform.

Legacy NAS NFSv3 export settings

The table below shows some of the systems tested by Hitachi Data Systems Global Solution Services with Universal Migrator. It lists the required export settings for successful virtualization.

LNAS	Settings				
	sec=sys	rw	sync	no_subtree_check	no_root_squash
HNAS	applicable	applicable	N/A	N/A	applicable
NetApp	applicable	rw=<ip>	N/A	N/A	root=<ip>
Linux	applicable	applicable	applicable	applicable	applicable

Note: Any NAS server that supports NFSv3 is expected to function with Universal Migrator.

File system block size selection

When performing any type of file share migration, the file system block size of the source and destination file systems is important in determining the amount of capacity required on the destination. The file system block size determines how much capacity a file requires to be stored on disk. For example, a 16 KB file would occupy 4 blocks on a file system with a 4 KB block size and use up 16 KB of space plus any additional space needed by the file system for internal metadata (that is, inodes or other similar file system structures). On a file system with a 32 KB block size, that same file would occupy 32 KB of disk space. That is a 50% overhead. For this reason, HNAS offers two file system block sizes: 4 KB and 32 KB.

The 32KB file system block size offers faster transfer rates whereas the 4KB file system block size offers higher OPS/sec. The file system block size also affects Hitachi NAS Primary Deduplication rates and speeds. For more information see the *Hitachi NAS Platform Deduplication Best Practices Guide*. Choose the file system block size that best meets capacity and performance requirements.

Note: The HNAS WFS-2 file system is object-based and does not use inodes nor is it bound to any of the restrictions inodes bring with them.

Space requirements for virtualization

During the virtualization process, a stub is created on the HNAS file system for each file discovered on the LNAS. The stub contains metadata about the virtualization path and other details needed to track it through the virtualization and migration phases. This stub consists of five file system blocks. This is regardless of the HNAS file system block size or the size of the

virtualized file. A stub on a HNAS file system configured to with a 4 KB block size would use $5 * 4 \text{ KB} = 20 \text{ KB}$ of space.

Once a file is fully migrated onto the HNAS file system, during the migration phase of Universal Migrator, the stub and its five blocks are deleted leaving behind the fully migrated file.

For example, a file system containing one million files with an average size of 8 KB would require roughly 19 GB of space during the virtualization phase and would consume around 11.5 GB of space after the migration completes assuming the HNAS file system was configured with a 4 KB block size.

A file system containing one million files with an average size of 512 KB would require approximately 152 GB of space during the virtualization phase and would consume around 519 GB of space after the migration phase completes assuming the HNAS file system was configured with a 32 KB block size.

To estimate the space requirements for the virtualization phase, use the following formula:
Capacity = Number of files x (5 x HNAS file system block size).

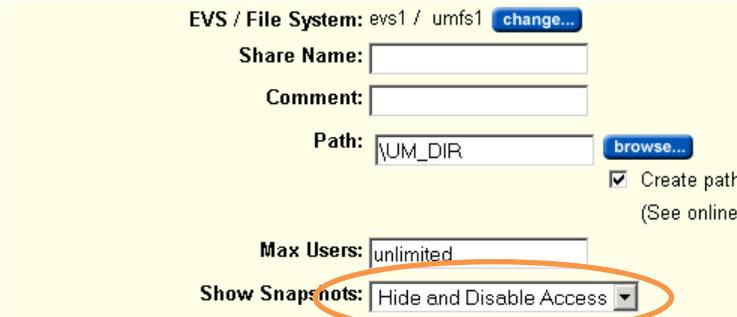
Note: When creating one or more HNAS file systems for use with Universal Migrator, the best practice is to allocate all of the file system capacity upfront. In the HNAS Web Manager GUI, select the “Allocate Now” radio button when creating the file system.

Preparing HNAS exports and shares

Prior to starting the virtualization process, exports and (optionally) SMB shares must be created on the HNAS pointing to the virtualization directory specified when the virtualization association was created. These will be used by the clients to access the virtualized data. When creating the exports/shares, make sure to disable access to HNAS snapshots. See [Caveats and limitations](#) for more details.

On exports, be sure to squash root access (this is the default setting for HNAS exports). See [Caveats and limitations](#) for more details.

Add Share



EVS / File System: evs1 / umfs1 [change...](#)

Share Name:

Comment:

Path: [browse...](#)

Create path
(See online)

Max Users:

Show Snapshots:

Figure 2. HNAS CIFS/SMB share settings.

EVS / File System: evs1 / umfs1 [change...](#)

Export Name:

Path: /UM_DIR [browse...](#)

Path Options

These options only apply when 'path'

Create path if it does not exist. (S

Allow this export path to overlap c

Show snapshots: **Hide and Disable Access**

Local Read Cache: Do not cache files

Transfer to Object Replication Target: Use FS Default (The File System

Access Configuration: * (root_squash)

Figure 3. HNAS NFS export settings.

HNAS virtual volumes

HNAS Virtual Volumes (Vivols) are special directories that can have quotas and a security mode assigned to them. These quotas can grow and shrink. Virtual volumes provide similar functionality to that of NetApp Qtrees. When migrating to HNAS it may be desirable to leverage this feature. Be aware that HNAS snapshots and object replications occur at the file system level when planning to consolidate multiple Qtrees or similar technologies using virtual volumes. An HNAS cluster can have up to 10,000 virtual volumes defined.

Note: HNAS virtual volumes need to be created prior to establishing a Universal Migrator association.

Performing a test virtualization

A test virtualization can be performed at any time with relatively little risk. The test will serve to identify potential security issues, misconfigurations, or other problems. A test virtualization can be performed by:

- Setting the HNAS EVS IP as a read/write host with root access.
- Creating a test file system on HNAS.
- Creating a directory on the test file system.
- Creating an association to an export on the LNAS and the directory on the test file system.
- Initiating the virtualization process
 - **Note:** During the entire virtualization process, the administrator must ensure that no modifying operations occur on the LNAS export.

It is not necessary to create any exports or shares on HNAS. **Caution:** If shares or exports are created, HNAS will forward any modifying operations to the LNAS performed by clients.

Once the test is complete, the virtualization path can be forcefully deleted. HNAS does not clean up directories or stub files created during the virtualization. These should be manually removed.

Universal migrator performance expectations

Performance is an important aspect of any migration. This section provides insight into the factors that can affect performance in the various migration phases and also provides some reference performance data to aid in planning. Universal Migrator performance can be viewed from three perspectives:

- Virtualization phase (export walking and stub creation).
- Migration phase.
- Perpetual virtualization (never entering the migration phase).

Performance of the migration phase and perpetual virtualization is primarily contingent on the capabilities of the LNAS system. Some of the main performance factors are listed below.

- Number and type of spindles behind LNAS export
- Number and type of spindles behind HNAS target file system
- LNAS busy rate
- HNAS busy rate
- LNAS link speed
- HNAS link speed
- Network health
- Average file size

If either of the NAS systems is busy, the virtualization and migration phases may take longer to complete. The virtualization phase places the greatest stress on HNAS. The amount of stress depends on the number of virtualization threads configured. The table below lists the load on a HNAS 3090 during the virtualization phase and not under client load.

Virtualization threads	MMB load	MFB load
20 (Default)	69%	40%
40	69%	60%
60	85%	65%

The overall load on the LNAS, which was another HNAS 3090, was relatively low. The MMB and MFB operated at 8% and 4% busy, respectively. This is expected as the virtualization phase generates relatively low OPS/sec rates. Tests indicate that each virtualization thread generates approximately 100-110 OPS/sec on the LNAS. In the worst case, each of these OPS may translate into 1 IO to the underlying spindles; therefore, plan accordingly.

The migration phase will generate a read load on the LNAS and write load on the HNAS. Additional workload may also be generated by clients accessing the virtualized exports. The migration transfer rate is primarily dependent on the number of threads, the average file size, and the performance capabilities of the source LNAS file system being migrated.

The average file size will have an impact on the migration transfer rates. A source export with many small files will have a high number of files migrated per second and a low MB/s transfer rate. Conversely, larger files will have a high MB/s transfer rate and low files migrated per second. The table below displays sample migration throughput data to aid in planning. HNAS used the default number of threads and neither it nor the LNAS was not under any client load.

File size	Throughput
Small Files <128 KB	160 MB/s
Large Files >100 MB	500 MB/s

Depending on the models of the HNAS and LNAS, the network connectivity may be 1GbE and/or 10 GbE. The virtualization phase can operate at full speed with 1 GbE links. The migration phase will benefit from 10 GbE links. In tests, the migration phase was able to saturate a 1 GbE link.

The health of the network will also affect the data transfer rates. Ideally, the HNAS EVS should be on the same network segment as the LNAS export. Some basic IO tests should be run against the HNAS to validate it is correctly integrated into the network and to ensure the network is healthy.

Virtualization and migration performance can be tuned by modifying the number of threads processing the LNAS exports. Increasing the threads will add additional load to the HNAS head and LNAS. If no performance improvement is observed after tuning, it is likely that a limit on the LNAS or HNAS has been reached. The threads can be adjusted by setting the following environment variable via the HNAS CLI:

set Virtualization-threads *number_of_threads*

This variable does not require a reboot and takes immediate effect.

Note: If the number of threads is being reduced, it may take several minutes for the setting to take effect.

Note: All the stated performance numbers were observed in a controlled test environment and are meant to be used as guidelines only. Actual performance may vary.

Heavy client access during virtualizations

As described earlier, many factors affect performance while in virtualization mode such as the capabilities, number of disks, and network of the LNAS. In addition, HNAS' parallel processing capabilities for virtualized files is significantly lower than that of normal files residing on the HNAS

file system. The HNAS architecture uses hardware FPGA acceleration to improve file serving performance. While in the virtualization phase, virtualized files are only partially hardware accelerated. HNAS must also track the state of virtualized files, and manage all IO operations to the LNAS. Therefore, it is generally expected that performance for virtualized files will be significantly lower. This is especially true when many hundreds of simultaneous client requests are sent to virtualized files.

To ensure optimal client performance during the virtualization phase, one option is to postpone client access until the virtualization phase is complete when high performance is required. Another option is to evaluate the performance needs of the exports being virtualized and schedule migrations so the amount of client load is minimal during the virtualization phase.

For example, an LNAS contains 30 exports spread across 10 file systems. Historical performance information can be used to determine the average number of OPS/sec for each file system and its exports. See the table below which shows an example of this.

File system name	Avg OPS load	File system name	Avg OPS load
vol1	15000	vol11	1800
vol2	8000	vol6	1500
vol5	5000	vol8	500
vol10	4500	vol9	450
vol4	2000	vol3	200

When deciding which file systems and exports to virtualize in a migration event, it would be best to start with the ones having the lightest load. Migration events for higher load exports should be scheduled during a window where the load is typically at its lowest. It is recommended not to have a virtualized load greater than 20 K-30 K OPS/sec per migration event. To migrate the LNAS in the example, a minimum of two migration events would be required: vols 3 through 11 in the first migration event and vol1 and vol2 in the second migration event. The earliest the second event would begin is when all the data from the previous event has been fully migrated onto HNAS.

Performance may not be the only factor in determining which file systems and exports to migrate in a migration window. Other factors such as application dependencies, backup, snapshot, and replication requirements should be considered. The best strategy is typically to find a balance between all the other factors while following the recommendations made as closely as possible.

Perpetual virtualization

Universal Migrator is designed to migrate data from an existing NAS system onto HNAS. During the course of a large migration, some LNAS exports may remain in virtualization phase for an extended period of time. It is also possible to remain in the virtualization phase indefinitely. It is important to understand the caveats of remaining in the virtualization phase.

While in virtualization mode it is not possible to perform backup, replication, and snapshots on the virtualized data. All of these must be managed from the LNAS.

As mentioned in the previous section, performance is limited while in virtualization. Tests indicate that while in virtualization mode, performance of the HNAS head is limited to approximately 30 K-35 K OPS/sec for all virtualized data. Any non-virtualized file systems on the HNAS are not subject to this limit.

The general recommendation is not to remain in virtualization mode for any long period of time.

Monitoring virtualization and migration performance

There are several ways to track the performance of Universal Migrator. Overall performance of the HNAS system can be tracked through the Web GUI or CLI commands. The following is a list of commands to track Universal Migrator performance:

Virtualization-path-stats – tracks response times for a virtualization path.

```
Virtualization path information for target volume: UM-filesystem
Virtualization path for directory: /vivoll
Legacy NAS source URI : nfs://192.168.0.230/vol/voll/home1
Association           : voll-home1-qtrees
Legacy read time:
  Samples=    5217  Total=   17.2 min  Min=3.13 ms  Av=198 ms  Max=3.73 s
  Last reset 2.07 d ago
HNAS write time:
  Samples=    5217  Total=    7.86 s  Min=386 us  Av=1.51 ms  Max=88.3 ms
  Last reset 2.07 d ago
File conversion time:
  Samples=    2856  Total=   16.7 s  Min=2.68 ms  Av=5.85 ms  Max=116 ms
  Last reset 2.07 d ago
```

Virtualization-path-list – shows the state of the virtualization, number of files discovered, and number of files currently being migrated.

```
HNAS:$ virtualization-path-list -t UM-filesystem
Virtualization path for directory: /vivoll
Virtualization path source and status:
  Source URI           : nfs://192.168.0.230/vol/voll/home1
  Association          : voll-home1-qtrees
  Mode                 : Virtualize, Migrate
  Options              : Default
  Status               : Virtualization Completed, Migration In Progress
  Seen Dirs            : 2118
  Processed Dirs       : 2118
  Pending Dirs         : 0
  Failed Dirs          : 0
  Seen Files           : 1788828
  Processed Files      : 966548
  Pending Files        : 822260
  Migrating Files      : 20
  Failed Files         : 0
```

The “Migrating Files” field indicates the concurrent number of files being migrated. If the number of virtualization threads is changed, the change can be observed by examining this field.

Caveats and limitations

This section details information about known limitations and caveats when using Universal Migrator.

Caveats

Snapshots

Universal Migrator requires a stable file system on the LNAS. Universal Migrator has no way of knowing that files are changing underneath it on the source LNAS export. For this reason, you cannot access LNAS snapshots. LNAS snapshot directories must be excluded from the virtualization path using the **virtualization-excluded-directory-add** command.

In addition, snapshots created on the HNAS must not be accessible to the clients while in the virtualization and migration phases. No snapshots should be created during the virtualization phase. Snapshots may be created during the migration phase for backup or replication purposes only. However, clients should not access these snapshots. Client access of these snapshots can lead to IO errors and unexpected results.

After all the migrations are complete for a specific HNAS file system, snapshot access can be allowed.

Note: Special NDMP snapshots (used for backup and replication) may still be extant on the file system. A new backup or replication should be run after the migration process completes so these snapshots are deleted.

Root squashing while using Universal Migrator

During the virtualization and migration process, Universal Migrator maintains a special virtualization directory. It is paramount that clients do not access or modify this directory. All NFS exports belonging to virtualization paths must have root squashing enabled. This can be changed after the migration process has successfully completed.

Note: By default all HNAS NFS exports squash root access.

Open files

Due to the added overhead tracking of virtualized files, the overall number of open files on an HNAS head may be significantly reduced while virtualizing LNAS systems and under client load.

Shortnames

HNAS will not allow creation of files whose longname matches the shortname of an existing file. If HNAS reads a LNAS NFS file name that happens to be the same name as an HNAS shortname that was generated when a separate file was previously virtualized, an error is generated and the file will not be virtualized.

LNAS character set

HNAS supports the latin1 and UTF-8 character sets. If the LNAS is using something different, this can cause file-naming issues during virtualization.

HNAS bossock fibers

Heavy client use of virtualized files can lead to increased bossock fiber utilization, especially if the LNAS or the connection to the LNAS is slow. This can affect other non-virtualized workloads on the HNAS head. Contact support to get assistance limiting bossock fibers on virtualized file systems.

Security operations on virtualized files

During the virtualization phase, HNAS forwards all security operations to the LNAS. During the migration phase, HNAS does not forward any security operations to the LNAS.

LNAS file system capacity tracking while in the virtualization phase

Be aware that HNAS does not generate any warnings or alerts prior to the LNAS file system becoming full. When the LNAS file system is full, clients will not be able to write further and will see an error indicating no free space is available. During the virtualization phase, rely on the LNAS for file system capacity warnings.

Changed Directory List (CDL)

The HNAS file replication and NDMP backup option CDL is not supported on file systems containing virtualized files.

Using another HNAS file system as the LNAS

Due to limitations in the network stack, you cannot create an association between an HNAS file system and an NFS export from the same cluster.

Migration of non-NFSv3 accessible data

Because the only protocol with which the HNAS communicates with the LNAS is NFSv3, any property of a file system object that is not accessible over NFSv3 will not be duplicated on the HNAS. This is mostly of concern to SMB (CIFS) specific properties such as ACLs, Attributes, DOS flags, alternate data streams, and so forth. Transferring of these properties is possible through Hitachi Data Systems Professional Services. Contact your sales team for more details.

Note: This applies not only to pre-existing SMB attributes on the LNAS, but also SMB modifications to IngestedFile objects (stubs) on the HNAS. These will not be reproduced on the LNAS during the migration phase. During the virtualization phase, HNAS will convert SMB security changes to UNIX mode bits and forward them to the LNAS.

SMB access to virtualized data

HNAS supports SMB access to virtualized data. As mentioned in the previous section, HNAS cannot read or transfer any non-NFSv3 data to the LNAS during the virtualization process. SMB access can be provided to Windows clients provided user mapping is in place so HNAS can translate SMB users to NFS UID's and GID's. Access to the files will be based on NFS security. For more information on how HNAS handles multiprotocol access, see the *Hitachi NAS Operating System 6.x Security Model Guide*.

Nesting and overlapping of exports on the LNAS

Nesting or overlapping exports on the LNAS used for associations should be avoided and can lead to unpredictable behavior.

Hard and soft links

Hard links should not span multiple LNAS exports. This can lead to unpredictable behavior. HNAS migrates soft links like any other file.

Handling failed files during virtualizations and migrations

There are numerous reasons that can cause a virtualization or migration to fail. The most likely causes are permissions and network-related issues. The best practice is to dry run the virtualization process to ensure proper connectivity and permissions to data on the LNAS. Use the `remote-nfs-errors` CLI command to display any errors encountered during the virtualization.

```
Count  Command  Status          FSBStatus
-----  -
4      SetAttr  nfserr_acces
```

To see a list of failed files, use the `virtualization-path-files` CLI command with the `--list-failed` option.

Reverting back to the LNAS

Reverting back to LNAS will vary in complexity depending on the phase the migration is currently in when the decision to revert is made. To streamline this process, a plan and any necessary systems should be in place to expedite the process. The table below offers some advice on how to revert back to the LNAS.

Protocol	Phase	Recommended action
NFSv3	Virtualization	Delete virtualization paths. Remove files/directories created by the virtualization process. Point clients back to LNAS.
NFSv3	Partial Migration	Use <code>rsync</code> to transfer files stored on HNAS back to the LNAS*
NFSv3	Completed Migration	Traditional data migration, using a tool such as <code>rsync</code> . If the LNAS is an HNAS, replication can also be used.

See the *Hitachi NAS Platform Data Migrator Administration Guide* for detailed instructions on how to recover from backups, snapshots, and when using replication on the LNAS and HNAS.

Interoperability with other Hitachi NAS Platform features

The table below lists interoperability with other HNAS features. The table displays what is supported in the various stages of the migration.

Feature	Virtualization Phase	Migration Phase	Post Migration
Hitachi NAS Snapshots	No	Partial *1	Yes
Hitachi NAS File Clone	No	Partial *1	Yes
Hitachi NAS Object Replication	No*2	No*2	Yes
Hitachi NAS File Replication	No	Partial *1	Yes
Hitachi NAS Primary Deduplication	No	Partial *1*3	Yes
Hitachi NAS Data Migrator and Data Migrator to Cloud	No	Partial *1	Yes

1 - Only applies to data that has been migrated to HNAS.

2 – While the HNAS does not prevent its use, this is unsupported and undesirable results may occur if used.

3 - Not officially tested but expected to work.

Frequently asked questions

Question: Can Universal Migrator migrate CIFS data?

Answer: Native CIFS support is planned for a future release. Transferring of CIFS data is possible via a Hitachi Data Systems Professional Services engagement. Contact your sales team for more details.

Question: My LNAS has CIFS shares that provide Windows clients access to UNIX files. Will Windows clients still be able to access UNIX files while they are being virtualized/migrated with Universal Migrator?

Answer: Yes, create a share on HNAS pointing to the virtualization directory. User mapping will need to be configured on HNAS ahead of time. Also, disable snapshot access on the HNAS share. The Windows clients can access the share once the virtualization phase is started.

Question: Can object replication be used while in virtualization mode?

Answer: No, it is not supported while in virtualization/migration mode.

Question: Can Universal Migrator be used to migrate VMware VM's?

Answer: No. Instead, use Storage vMotion.

Question: Is client access affected while a file is being migrated onto HNAS?

Answer: No

Question: Does Universal Migrator scan source exports to make sure it has the correct permissions it needs?

Answer: No. It is recommended to perform a test run of virtualization and migration to make sure it works with no disruption to production.

Question: Is NFSv4 or NFSv4.1 supported?

Answer: Not currently.

Question: Why would virtualization or migration fail?

Answer: Name clashes, permissions, or network problems are likely causes.

Question: What firmware does the LNAS need if it is an HNAS?

Answer: Any supported firmware level.

Question: Does the client have to remount their exports?

Answer: Yes. The virtualization process will change the file handles requiring a remount.

Question: Does HNAS migrate hard and soft links?

Answer: Yes. Note that hard links should not span exports on the LNAS.

Terminology

- **Association** – The configured connection between a mount point on the **LNAS** and a directory in a file system on the **HNAS**. An association is also referred to as a virtualization path.
- **Excluded directories** – The user-defined list of directories that will not be created on the virtualized HNAS file system.
- **HNAS** – The Hitachi NAS system that will hold the data, after migration.
- **IngestedFile** – An object (directory or file) in the HNAS file system that is in the process of being virtualized or migrated. As such, it contains extra metadata (compared to a **regular file/directory**), but incomplete or absent user data. In some ways (principally user data read/write access) an **IngestedFile** can be thought of as behaving similarly to an XVL. The extra metadata is necessary for two reasons:
 1. To keep track of the current virtualization/migration state of the file.
 2. To maintain a link to the **LNAS** version of the file.
- **LNAS** – The pre-existing "legacy" NAS system that holds the data to be migrated.
- **Migration** – The process of bringing user data for an object from the **LNAS** to the HNAS.
- **Regular file/directory** – The normal type of object (a file or directory) in the HNAS WFS file system.
- **Virtualization** – The process of discovering an object on the **LNAS**, and creating an **IngestedFile** to represent the LNAS object in the HNAS file system.
- **XVL** – External cross volume link.
- **EVS** – Enterprise Virtual Server, is a logical entity that owns file systems, has its own IP addresses, shares, and exports. An EVS either leverages the global configuration or can have an individual security and networking configuration. Individual security is provided by the EVS Security license.

Hitachi Data Systems

Corporate Headquarters

2845 Lafayette Street
Santa Clara, California 95050-2639
U.S.A.

www.hds.com

Regional Contact Information

Americas

+1 408 970 1000

info@hds.com

Europe, Middle East, and Africa

+44 (0)1753 618000

info.emea@hds.com

Asia Pacific

+852 3189 7900

hds.marketing.apac@hds.com



MK-92HNAS047-00