HITACHI
Inspire the Next

Hitachi NAS Platform Data Migrator to Cloud Best Practices Guide

Francisco Salinas, Global Services Engineering

Hitachi Data Systems

MK-92HNAS045-01

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.

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-92HNAS045-00	March 2014	First publication
MK-92HNAS045-01	October 2014	Revision 1, replaces and supersedes MK-92HNAS045-00

Contact

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

North America: 1-800-446-0744

Contributor

Global Services Engineering would like to recognize and sincerely thank the following contributor of this document for their expertise, feedback, and suggestions:

- Dan Burrows

Table of Contents

Intended audience

The intended audience for this guide is Hitachi Data Systems (HDS) customers, employees, and partners.

Overview

Data Migrator to Cloud, introduced in HNAS release 11.1, allows files hosted on the HNAS server to be transparently migrated to cloud storage, providing the benefits associated with both local and cloud storage. Data Migrator to Cloud combines External Cross-Volume Link (XVL) technology in HNAS with cloud storage targets such as the Hitachi Content Platform, Hitachi Cloud Services, and Amazon S3.

A public cloud (such as Amazon S3, Hitachi Cloud Services – Content Archive) is provided by external entities hosting storage at their facility and paid for on a per-use basis. A private cloud is purchased and controlled by the end user. Data Migrator to Cloud supports both and the user can decide which model best suits their business needs. In both cases, it can transparently access the cloud storage directly to view and download data. In both scenarios data is protected both in-flight and at-rest regardless of where the physical storage is hosted.

Data Migrator to Cloud

Data Migrator to Cloud was designed to leverage private and public clouds It also provides integration with the Hitachi Content Platform (HCP). Data Migrator to Cloud provides the following functionality:

- HCP Namespace verification
  - HNAS will validate best practice settings on HCP to ensure optimal configuration. HNAS will check that settings such as retention are disabled and verify the Data Access user permissions are sufficient.
- Migration Performance to HCP
  - The migration engine is multi-threaded which aids in file system walking performance and migration job duration. In HNAS release 11.1 all threads connect to single HCP node. In HNAS release 11.2, HNAS connects to multiple HCP nodes which improves upload performance.
- Writing/modifying migrated data
  - In order for data migration to HCP and cloud targets to be transparent to users, the users need the ability to modify migrated files. With Data Migrator to Cloud, users can seamlessly modify files that have been migrated.
- Delayed Deletion
  - HNAS release 11.2 introduces support for multiple versions of a migrated file. If a file is migrated to a destination (HCP, Hitachi Cloud Services, or Amazon S3) and is subsequently changed, a snapshot will reference the original version of the file and the active file system will reference the latest version. HNAS uses its own versioning functionality and does not use the capability provided by HCP in order to provide a consistent functionality across all of the supported private and public cloud destinations.
- Deletion of stub files

  - When the user deletes a file that has migrated to HCP or one of the other supported cloud targets, the delete is propagated to the target. With 11.2 and beyond, HNAS snapshots protect stubbed data. A file that has been stubbed will not be removed from the target until all XVL references, including snapshots references to the archive file, have been removed.

- SSL (HTTPS) Support
  - Support for encrypted communications to public and private clouds.
  - As of HNAS release 11.3, HTTP is now the default protocol used for communication with local HCP targets.
- Full CLI support
  - Data Migrator to Cloud can be fully managed via the HNAS CLI.
- Fully implemented in the server
  - Data Migrator to Cloud is implemented on the HNAS server itself. It has no dependencies on the System Management Unit (SMU) and NDMP.
- Data Migrator to Cloud can be fully configured via the SMU GUI
- Data Migrator to Cloud allows multiple file systems to share the same HCP namespace and directory by assigning each file system a unique identifier (UUID)
- Data Migrator to Cloud preserves the migrated files path and name for easier browsing on the cloud target.

There are also some differences when compared to Classic Data Migrator.

- Reverse migration of any cloud migrated files can only be initiated through the CLI.
- The HNAS node eth0 and eth1 1GbE network interfaces are used for communication to the cloud target.

Data Migrator to Cloud Architecture

Apcellerator and the CloudApp

Data Migrator to Cloud uses the HNAS Appcellerator infrastructure. This infrastructure allows storage management applications to run on the Linux Platform within HNAS. Appcellerator provides storage management applications on the HNAS Linux Platform access to HNAS WFS-2 file systems. Data Migrator to Cloud is an application that runs on the HNAS Linux Platform known as the CloudApp. The CloudApp consists of the tree walker, uploader, and a web server. The tree walker scans the file system. The uploader sends data to the cloud target, and web server performs caching and read ahead. A simplified diagram of this architecture is visible in Figure 1.

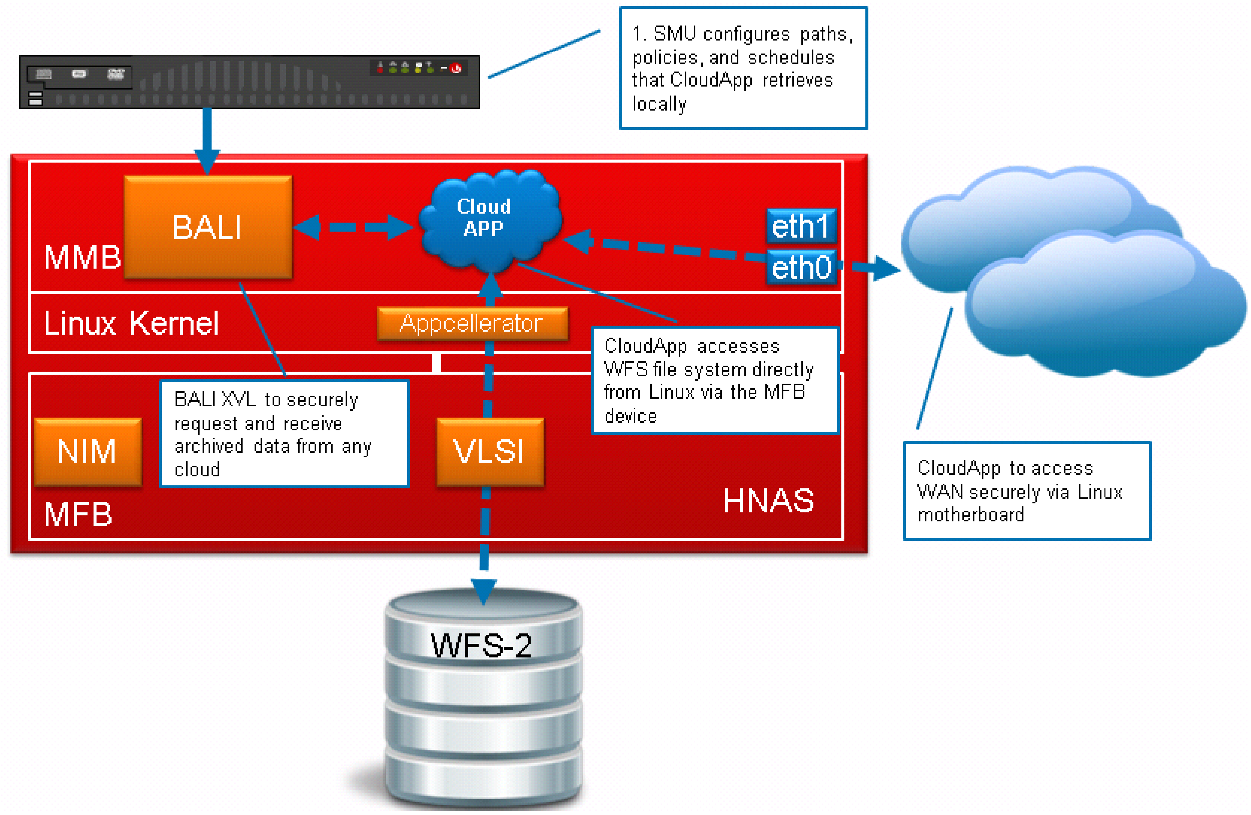
1. SMU configures paths, policies, and schedules that CloudApp retrieves locally
BALI
MMB
Cloud APP
eth1
eth0
Linux Kernel
Appcellerator
CloudApp accesses WFS file system directly from Linux via the MFB device
BALI XVL to securely request and receive archived data from any cloud
NIM
VLSI
MFB
HNAS
CloudApp to access WAN securely via Linux motherboard
WFS-2

Figure 1 - Data Migrator to Cloud Architecture

Networking

Data Migrator to Cloud uses the HNAS Linux Platform to encrypt (optional for local HCP targets starting in HNAS 11.3), send and retrieve cloud data. There are two 1GbE network interfaces available on the Linux Platform, eth0 and eth1. Eth1 is normally used for the private network which provides quorum, heartbeat and management interfaces. On the Systems Management Unit (SMU), eth1 is normally connected to this private network and eth0 is typically not used. All of these connections form the HNAS backend network.

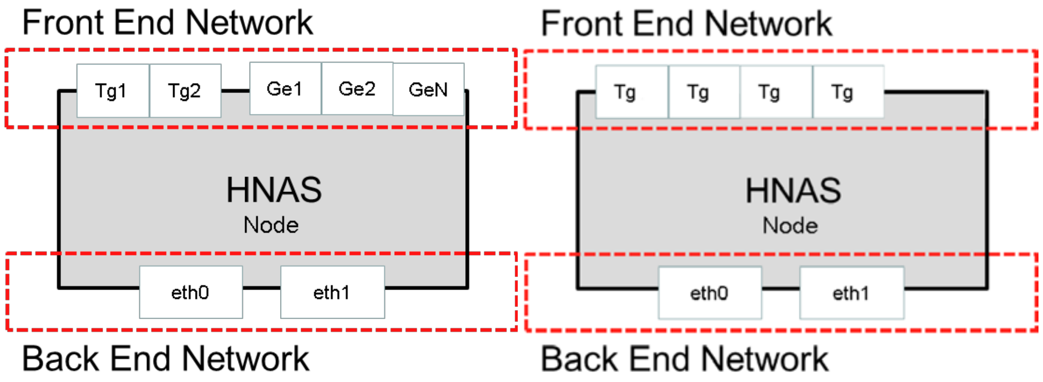

Figure 2 - HNAS 30x0 and 4xx0 Network Interfaces

HNAS requires that at least one of the backend network interfaces on an HNAS node be configured for cloud connectivity, which we will call the cloud interface. In the event of a node failure, cloud data would be accessible from the surviving node. The best practice is to configure both backend network interfaces (eth0 & eth1) on a node for cloud connectivity.

To configure both cloud interfaces, two public IP addresses are required per HNAS node a public IP address for the Admin EVS, and an additional public IP address for the SMU In a two node cluster, six IP addresses are required.

To connect to a cloud target, HNAS needs to have DNS and network routes configured accordingly. The HNAS system must be able to resolve the Fully Qualified Domain Name (FQDN) of the cloud target. It also needs to be able to route to the target from the cloud interfaces. When connecting to HCP as a cloud target, the best practice is to connect HNAS and HCP on the same network subnet. For detailed instructions on configuring these interfaces, DNS, and routing, please see the *Data Migrator Administrator's Guide*.

The SMU should be configured to manage the HNAS cluster using a public IP address on the same subnet as the HNAS eth0 and eth1 public IP addresses. Eth1 on the SMU is not protected by a firewall and is designed to be used in a closed private network. Therefore, the eth0 interface will need to be configured to manage the HNAS cluster. See Figure 3.

Note: The Admin EVS is not depicted in Figure 3.

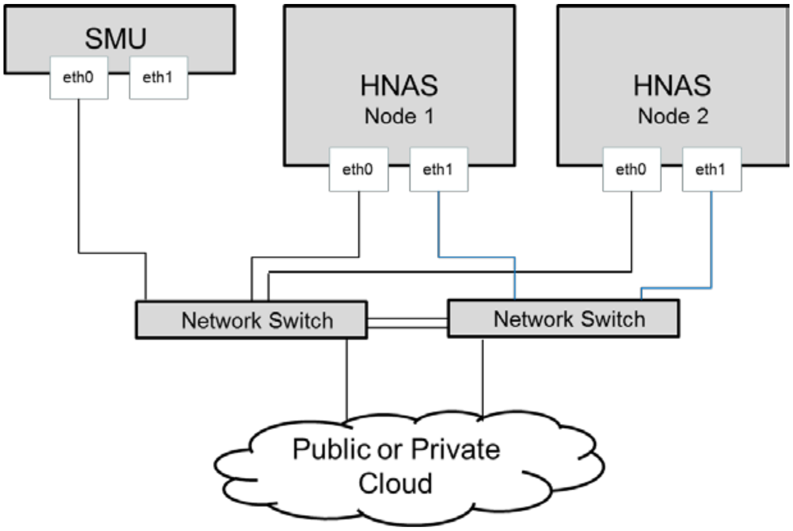

Figure 3 - Backend Network Connectivity

Migrated data identifier on the cloud target.

Data Migrator to Cloud uses a special universally unique identifier (UUID) for each file system that is migrated to a cloud target. This allows the user to migrate data from multiple source file systems to the same directory on a cloud target without have to worry about name collisions.

Though not required, it is possible to create individual folders (Cloud Destinations) named after the source file system for easier identification of data on the cloud if several file systems are migrating data to the same target. Take the following limits into consideration when implementing the above recommendation: the maximum number of Cloud Accounts is 20, and the maximum number of Cloud Destinations is 40. For example, when HCP is the cloud target, it is possible to connect to the HCP namespace and browse the migrated data. See the HCP section on page 15 for screenshots.

For chargeback, consider creating one namespace for each department or group and set the file system cloud migration to that namespace. This makes chargeback reporting easier.

Data Migrator to Cloud Performance

Data Migrator to Cloud has been enhanced over Classic Data Migrator. Specifically, with the implementation of the tree walker, and the multi-threaded uploader mentioned previously in this document, upload performance has been improved. In HNAS release 11.2, further enhancements to the performance of the uploader is implemented by managing multiple concurrent connections to the HCP node(s); Performance tests indicate Data Migrator to Cloud offers improved migration (upload) performance when compared to Classic Data Migrator. In some cases up to 2x, but results will vary.

The table below contains performance information comparing HTTP and HTTPS on an HNAS 3090 running 11.3. The test used 4 clients each with 128 x 1GB files. The VDbench program with 16 threads was used to read the migrated files.

Direction	Classic DM	DM to Cloud HTTP	DM to Cloud HTTPS
Migration Large Files (1GB)	31.7MB/s	103.2MB/s	43.5MB/s
Read Large Files (1GB)	49.3MB/s	43.4MB/s	30.4MB/s

Data Migrator to Cloud migration performance can be determined by examining the migration log file.

Using Data Migrator to Cloud

To allow a user to modify a file that has been migrated, `xvl-auto-recall-on-modify` must be set to True.

When setting up Data Migrator to Cloud using Amazon S3 as a target, the username and password required to create the cloud account on HNAS can be setup in Security Credentials page of the Amazon Accounts page. The **Access Key** and **Secret Access Key** are also required. See Creating an IAM account in Amazon on page 9.

For detailed information on setting up Data Migrator to Cloud, see the *Data Migrator Administration Guide*.

Creating an IAM account in Amazon

For Amazon S3, you must create an Identify and Access Management (IAM) account.

1. Go to https://console.aws.amazon.com/iam/ and log in with your user name and password. Refer to http://docs.aws.amazon.com/IAM/latest/UserGuide/Using_SettingUpUser.html#Using_CreateUser_console for more information.
2. When creating a user, enable the **Generate an access key for each user** (this is enabled by default) check box and then click **Create** to generate the access and secret keys. Refer to http://docs.aws.amazon.com/IAM/latest/UserGuide/ManagingCredentials.html for more information.
3. Click **Download Credentials** and then save the access keys to your local machine. You will need this information when you create a cloud account.
4. Click **Close** to exit.
5. Highlight and select the newly added IAM user account to open the users page.
6. Click **Attach User Policy** and select **Amazon S3 Full Access** (you may have to scroll down the page).
7. Click **Apply Policy**.
8. When you create an Amazon cloud account, provide the access and secret keys just created

When specifying the cloud destination, specify the bucket and then a folder within the bucket. See Figures 4 and 5.

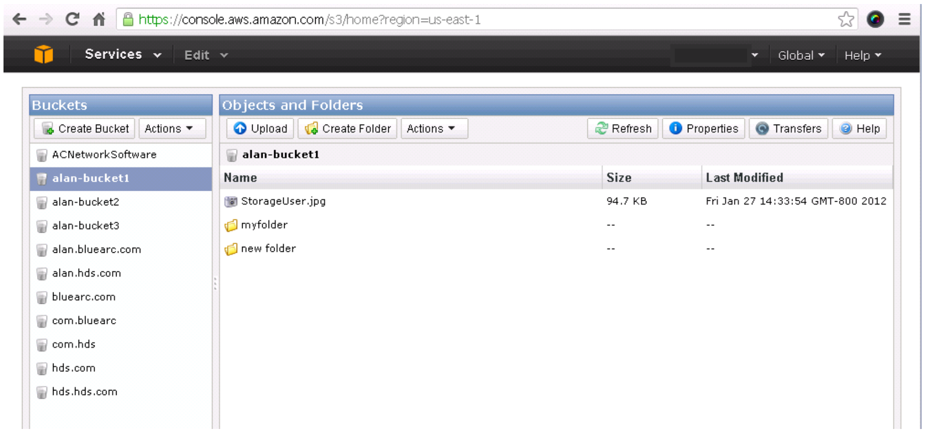

- ACNetworkSoftware
- alan-bucket1
- alan-bucket2
- alan-bucket3
- alan.bluearc.com
- alan.hds.com
- bluearc.com
- com.bluearc
- com.hds
- hds.com
- hds.hds.com

alan-bucket1

Name	Size	Last Modified
StorageUser.jpg	94.7 KB	Fri Jan 27 14:33:54 GMT-800 2012
myfolder	--	--
new folder	--	--

Figure 4 - Amazon S3 Folders

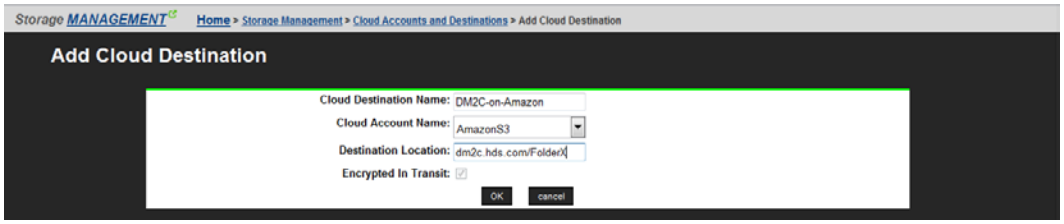
Storage MANAGEMENT
Home > Storage Management > Cloud Accounts and Destinations > Add Cloud Destination
Add Cloud Destination
OK
cancel

Cloud Destination Name: DM2C-on-Amazon
Cloud Account Name: AmazonS3
Destination Location: dm2c.hds.com/FolderX
Encrypted In Transit: ☑

Figure 5 - Amazon S3 as a Destination

With HCP 6.0 and higher, the permissions are required for the cloud account user are shown in Figure 6

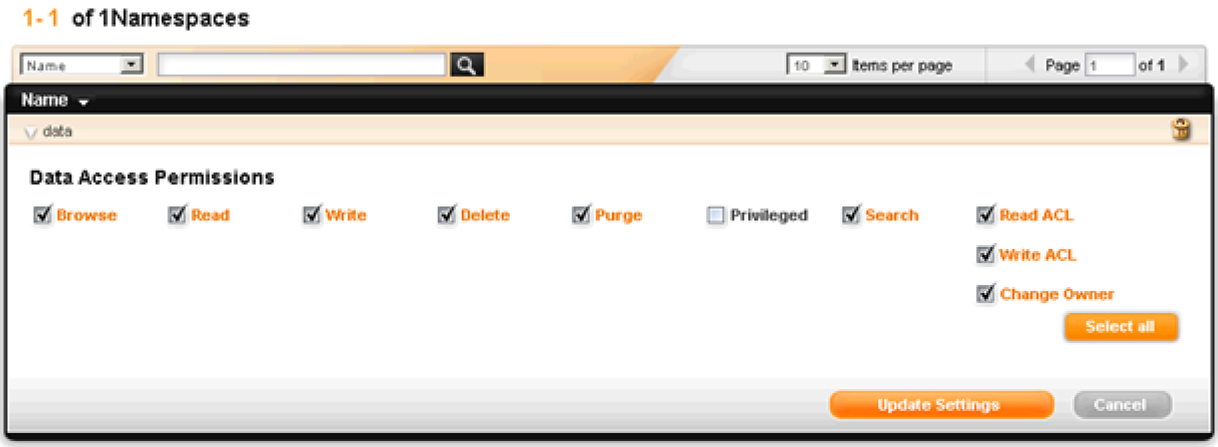

Data Access Permissions

☑ Browse ☑ Read ☑ Write ☑ Delete ☑ Purge □ Privileged ☑ Search ☑ Read ACL
☑ Write ACL
☑ Change Owner

Figure 6 - HCP 6 Cloud Account Permissions

Delayed Deletion

Starting with HNAS release 11.2, Data Migrator to Cloud supports integration with HNAS file system level snapshots. HNAS file system snapshots can preserve changed files that reside on an HCP target for which a snapshot existed prior to the change. This is known as delayed deletion. HNAS uses its own versioning implementation to support delayed deletion. It is possible to enable HCP versioning in HNAS 11.2; however, the best practice is to leave it disabled, which is the default in HCP. Some issues were discovered with versioning in HNAS 11.2, and these issues have been resolved in HNAS 11.3. To reiterate, it is recommended to ensure that HCP versioning is disabled.

Interoperability with other features

Classic Data Migrator

The environment variable `xvl-auto-recall-on-read` can be used with both although setting `xvl-auto-recall-on-read` for Data Migrator to Cloud is not necessary.

Note: If xvl-auto-recall-on-read is set with Object Replication, the replication will trigger a reverse migration of the cloud migrated files.

Upgrading Classic Data Migrator XVL to Data Migrator to Cloud XVL functionlity

There are two options available to users:

(1) - Leave the existing Classic Data Migrator XVL's in place and have all future migrations use the Data Migrator to Cloud.

(2) - Use option 1, and convert the existing Classic Data Migrator XVLs to Data Migrator to Cloud XVLs. This option is available through GSS only.

For those who are upgrading to Data Migrator to Cloud from Classic Data Migrator, all existing Classic XVL files will continue to function as they have been. Converting the Classic XVL files to the new Data Migrator to Cloud XVL format is ONLY required to take advantage of the new versioning or preserving of deleted files in a snapshot once the customer upgrades to HNAS 11.2 or newer. The conversion will not make changes to existing snapshots and the versioning advantages will only be applicable for new snapshots. For files that are migrated with Data Migrator to Cloud, users will benefit from the new capabilities for any new snapshots taken after upgrading to HNAS v11.2 or newer.

Note: Existing customers who are switching the Data Migrator to Cloud will need to reconfigure their backend network. See Networking on page 7 and contact Hitachi GSS for more information.

Instructions for Option 1:

- Make the necessary changes to the backend network to support Cloud connectivity
- Establish connection to the Cloud target (that is, Add a Cloud Account and Destination)
- Forcefully delete the Classic Data Migrator external path
- Add a new external path using the Cloud target and Destination setup earlier

File Replication

Hitachi NAS File Replication is the recommended method to replicate file systems that contain data migrated to the cloud.

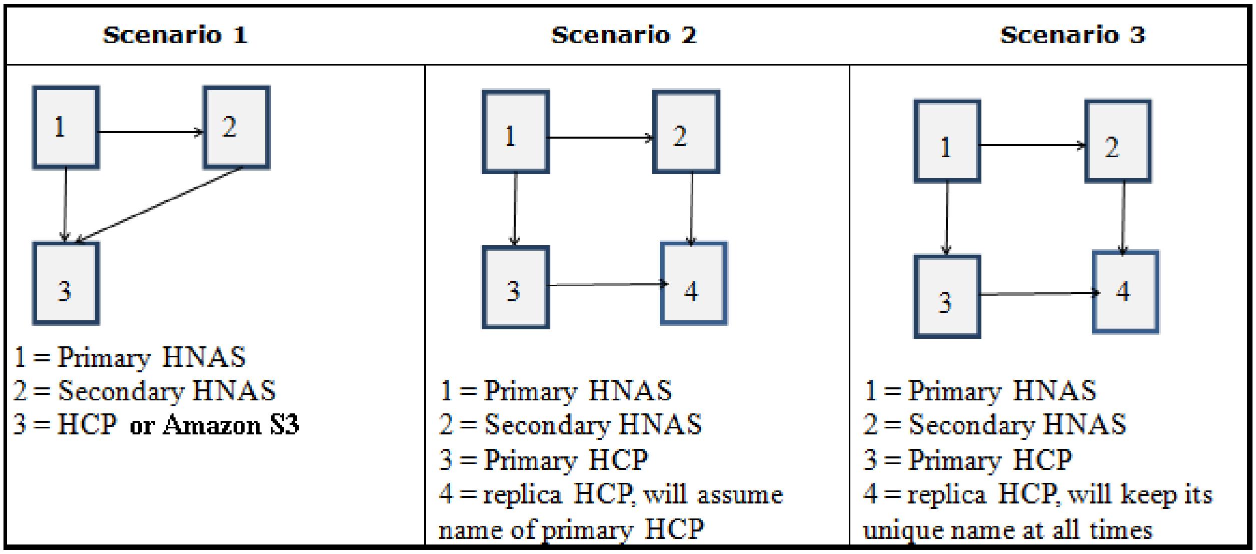

Figure 7 - File Replication Scenario's

Scenario 1 Illustrates replicating file systems between HNAS clusters, both of which point to a single HCP or Amazon S3 target.

Warning! In this scenario, both HNAS clusters/entities map to the same HCP or Amazon S3 target. With HNAS file replication it is possible to access the secondary file system(s) at any time. It is strongly recommended to keep the destination file system syslocked to avoid unintentional deletion of data on the HCP system.

Scenario 2 Illustrates replicating file systems between HNAS clusters, where each cluster points to a HCP. The HCP replicates migrated data and also performs a DNS failover so that the secondary HCP maintains the same name resolution as the primary system.

Warning! In this scenario, HCP uses a DNS failover capability. Due to the way the HCP failover functionality operations, the secondary HNAS will also point to the primary HCP. With HNAS file replication it is possible to access the secondary file system(s) at any time. It is strongly recommended to keep the destination file system syslocked to avoid unintentional deletion of data on the HCP system.

Scenario 3 Illustrates replicating file systems between HNAS clusters, where each cluster points to a HCP. The HCP's replicate migrated data and maintain their own unique name resolution.

HCP replication is neither real time nor is it point-in-time. For this reason, you will be unable to guarantee a recovery point objective for any data migrated to HCP.

For instructions on configuring HNAS in any of the above scenarios, **fully** read the *Data Migrator Administration Guide*, *HNAS 11.2 Release Notes,* and *Replication and Disaster Recovery Guide*.

Object Replication

Hitachi NAS Object Replication, when used to replicate a file system that contains data migrated to cloud, will follow the links and replicate the data to the destination file system. You cannot change this behavior. Alternatively, use options specific to Hitachi NAS File Replication to properly handle the XVL files.

Environment Variable	Behavior of File Replication	Behavior of Object Replication
xvl-auto-recall-on-modify (Not enabled by default)	Does not trigger reverse migration	Does not trigger reverse migration
xvl-auto-recall-on-read (Not enabled by default)	Does not trigger reverse migration	Triggers reverse migration

Snapshots

When using Data Migrator to Cloud, snapshots that are created on the primary file system do not have corresponding snapshots created on the cloud target.

To preserve snapshot protection on migrated files for cloud migrations, you must ensure that snapshots are taken on the cloud target. Any snapshots on the cloud target are not managed, used, or accessed by the HNAS server.

When a snapshot is accessed, and the snapshot contains a file with an external cross volume link (XVL), HNAS will attempt to follow the link, if the file on the cloud target no longer exists, an error will be returned.
Space may not be reclaimed immediately if a file that is migrated is preserved in a snapshot. The space will be eventually reclaimed when the snapshot with the full file ages out and is deleted.

Primary Deduplication

Any file that has been fully or partially deduplicated on a Data Migrator to Cloud source file system will be rehydrated upon migration to the cloud target. If the file is recalled, HNAS will attempt to dedupe the data during the next deduplication job for that file system as it normally would for any newly written data.

IMPORTANT

Certain HNAS 3080/3090 deployment configurations may not have sufficient free MMB memory space to support both Data Migrator to Cloud **and** Primary Deduplication.

Starting in HNAS release 11.3, the user will receive a warning that the MMB memory is over committed. If the user wants to leverage Data Migrator to Cloud and/or Primary Deduplication, they may need to reduce the amount heap memory. See the man page of `mmb-memory` for more details.

NDMP Backup

HNAS NDMP offers several variables to control how migrated data is handled during backups and restores.

If set to the value “Y”, the variable NDMP_BLUEARC_EXCLUDE_MIGRATED will cause all migrated files to be excluded from the backup. Neither the data nor the XVL metadata would be copied. The default value is “n”, which will make the backup copy the migrated files to the backup target.

The variable NDMP_BLUEARC_INCLUDE_ONLY_MIGRATED has the opposite effect. When set to “y”, all files will be excluded from the backup except the files that have been internally migrated. The default value is “n”.

Note: if the `xvl-auto-recall-on-read` environment variable is enabled, an NDMP job will not cause the migrated files to be recalled.

Character Sets

For versions **12.1 and below**, HNAS uses the Latin-1 as its default character set. To support special characters, change this to UTF-8, otherwise files with special characters will not migrate. Note: If the character set is incorrect, the migration will fail and will not be able to progress further than the file with special characters. To set the correct character set, run the following:

```
protocol-character-set --all UTF-8
```

Virtual Server Security (VSS) a.k.a. Secure EVSes

Data Migrator to Cloud uses the Admin EVS to resolve names for cloud targets. The VSS feature is designed to isolate an EVS into its own security domain. If a Secure EVS is configured with its own separate DNS and own HCP target that is not resolvable by the DNS used by the Admin EVS, then cloud account creation will fail. To ensure there are no issues, the Admin EVS must be able to resolve the name of the cloud targets.

Hitachi Content Platform View of XVL's

When using HCP as cloud target, the HCP administrator can optionally browse the namespace. To browse the namespace, login to the Namespace Console. Figure 8 shows the view in HCP when using Classic Data Migrator. Figure 9 shows the view of HCP when using Data Migrator to Cloud. With Data Migrator to Cloud, a directory with the UUID of the source file system is created which then holds any migrated data for that file system. This is visible in Figure 9, the directory starting with 743A2…

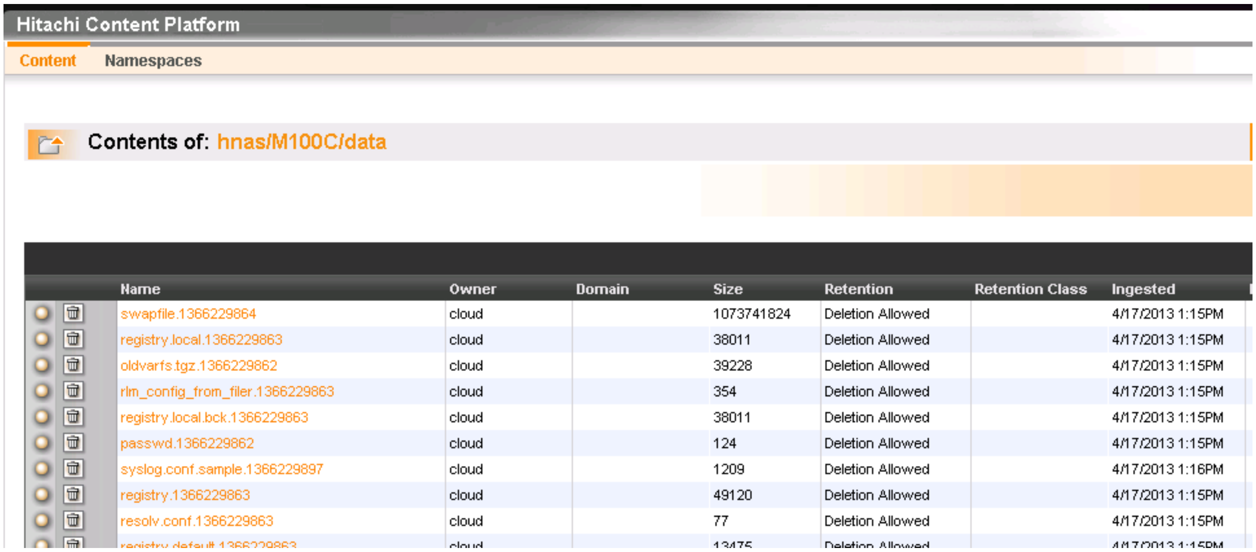

Contents of: hnas/M100C/data

Name	Owner	Domain	Size	Retention	Retention Class	Ingested
swapfile.1366229864	cloud		1073741824	Deletion Allowed		4/17/2013 1:15PM
registry.local.1366229863	cloud		38011	Deletion Allowed		4/17/2013 1:15PM
oldvarfs.tgz.1366229862	cloud		39228	Deletion Allowed		4/17/2013 1:15PM
rlm_config_from_filer.1366229863	cloud		354	Deletion Allowed		4/17/2013 1:15PM
registry.local.bck.1366229863	cloud		38011	Deletion Allowed		4/17/2013 1:15PM
passwd.1366229862	cloud		124	Deletion Allowed		4/17/2013 1:15PM
syslog.conf.sample.1366229897	cloud		1209	Deletion Allowed		4/17/2013 1:16PM
registry.1366229863	cloud		49120	Deletion Allowed		4/17/2013 1:15PM
resolv.conf.1366229863	cloud		77	Deletion Allowed		4/17/2013 1:15PM
registry.default.1366229863	cloud		13475	Deletion Allowed		4/17/2013 1:15PM

Figure 8 - Classic Data Migrator view in HCP

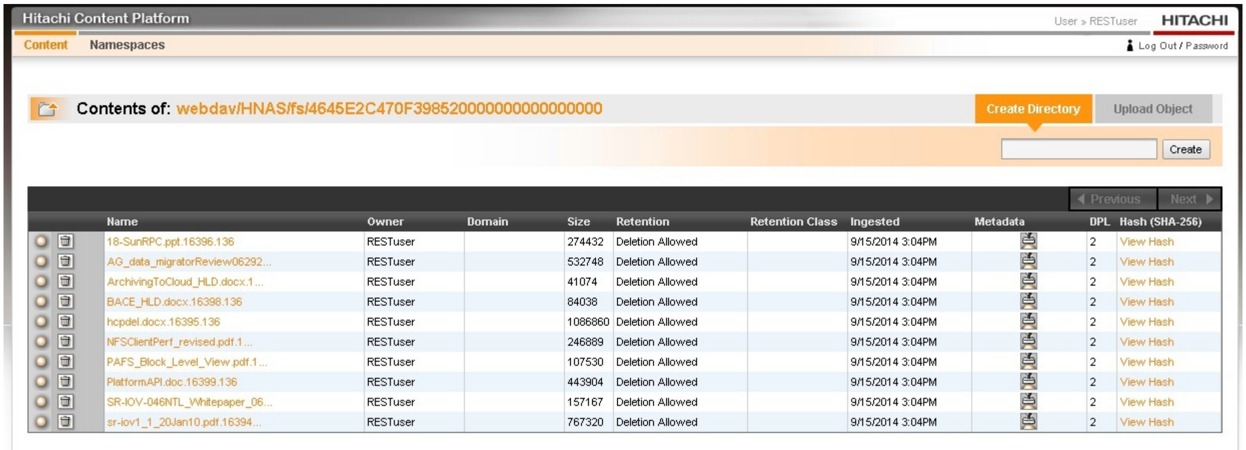

Contents of: webdav/HNAS/fs/4645E2C470F398520000000000000000

Name	Owner	Domain	Size	Retention	Retention Class	Ingested	Metadata	DPL	Hash (SHA-256)
18-SunRPC.ppt.16396.136	RESTuser		274432	Deletion Allowed		9/15/2014 3:04PM		2	View Hash
AG_data_migratorReview06292...	RESTuser		532748	Deletion Allowed		9/15/2014 3:04PM		2	View Hash
ArchivingToCloud_HLD.docx.1...	RESTuser		41074	Deletion Allowed		9/15/2014 3:04PM		2	View Hash
BACE_HLD.docx.16398.136	RESTuser		84038	Deletion Allowed		9/15/2014 3:04PM		2	View Hash
hcpdel.docx.16395.136	RESTuser		1086860	Deletion Allowed		9/15/2014 3:04PM		2	View Hash
NFSClientPerf_revised.pdf.1...	RESTuser		246889	Deletion Allowed		9/15/2014 3:04PM		2	View Hash
PAFS_Block_Level_View.pdf.1...	RESTuser		107530	Deletion Allowed		9/15/2014 3:04PM		2	View Hash
PlatformAPI.doc.16399.136	RESTuser		443904	Deletion Allowed		9/15/2014 3:04PM		2	View Hash
SR-IOV-046NTL_Whitepaper_06...	RESTuser		157167	Deletion Allowed		9/15/2014 3:04PM		2	View Hash
sr-iov1_1_20Jan10.pdf.16394...	RESTuser		767320	Deletion Allowed		9/15/2014 3:04PM		2	View Hash

Figure 9 - Data Migrator to Cloud view in HCP

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

Hitachi Data Systems