HITACHI
Inspire the Next

Hitachi Data Systems

File System Snapshot Operational Best Practice

Author: Denis Kornilov

Hitachi Data Systems

MK-92HNAS068-00

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.

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.

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, Genithrough Inc. All rights reserved. The software in this product was in part provided by Genithrough 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.

Document Revision Level

Revision	Date	Description
1.3	17 March 2015	Document created by HNAS Engineering
1.4	4 June 2015	Format and edit by GSS Technical Publications

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 (listed alphabetically):

- Stuart Bell
- Gaurav Sharma
- Amit Brahmbhatt
- Dan Burrows
- Richard Powell

Table of Contents

Intended audience

The content presented in this document assumes hands-on knowledge of HNAS administration.

About this document

This document provides operational guidance on HNAS file system snapshots. This document is supplemental to the *Snapshot Administration Guide* and *HNAS operational envelope collateral*, and the rest of the HNAS product documentation suite. This document is not intended to be a procedural guide or a replacement for *HNAS Snapshot Administration Guide*.

Introduction

File system snapshots (or simply *snapshots* within this document) are point-in-time read-only representations of HNAS file systems. Snapshots preserve the contents of the blocks which would otherwise be lost due to modifying operations (for example, writes, creates, deletions, and dedupe). In other words, snapshots preserve the data, so changes can be written elsewhere in the live file system

Use case

Snapshots solve the problem of maintaining consistency within a backup or file replication; specifically, during a live file system backup/replication, users may continue to modify its component files, resulting in backup copies that may not provide a consistent set. Since a snapshot provides a frozen image of the file system, a backup/replication copy of a snapshot (rather than of the live file system) provides a usable, consistent backup.

When more than one snapshot is involved, it is possible to perform incremental backup/replication of data changed since the last snapshot.

From an application point of view, snapshots provide a crash-consistent view of data. In order to provide application-consistent view, additional integration with the application can be scripted, where snapshots are initiated after the application is either quiesced or placed into a backup mode.

Snapshots can also be used to roll back the live file system to a previous point in time, often beyond the timeframe provided by the roll-back mechanism of file system checkpoints.

Feature description

Snapshots can be created by rule, manually, or by an application such as NDMP backup software, SMU (file and object replication, data migrator), SRA adapter for HNAS, VSS, CFN, dedupe, and online `checkfs`.

Snapshots can be deleted manually, by rule, or upon exceeding the cluster-wide maximum retention period.

Snapshot retention is controlled by the *snapshot rule's queue size*, which represents the number of snapshots kept until the system automatically deletes the oldest snapshot. For NDMP backup, the maximum retention can also be capped by the number of days (default is 7 and the maximum is 40/80 in versions 11.3 and 12.1).

There is a limit of 1024 snapshots per file system. There is no limit on how often you can create a snapshot.

Access point into a snapshot is `.snapshot` for NFS clients and `~snapshot` for SMB clients.

Snapshots can be made hidden in exports/shares but are still accessible or they can be made both hidden and inaccessible.

HNAS provides the latest snapshot capability through `.snapshot/.latest` (for NFS clients) or `~snapshot/.latest` (for SMB clients). This view automatically changes as new snapshots are taken.

There are two modes of operation for the latest snapshot feature related to how `fileids` are reported when a new snapshot "rolls into" the latest snapshot view. By default, `fileids` are not preserved. This can be toggled by "`fsm set latest-snapshot-preserve-fileids true`".

Snapshot's freeable space can be reported using Snapshot GUI or `snapshot-list --freeable-space` command.

Operation

For schedule-based snapshots, a snapshot rule must be created first, defining the number of snapshots to keep (retention policy). One or more schedules can then be associated with that rule.

During NDMP backup, SRA or file replication, a snapshot will be created by default and deleted after the next backup by default.

Automatic snapshot creation during backup can be turned off globally on HNAS or overridden with the NDMP environment variable **`NDMP_BLUEARC_TAKE_SNAPSHOT`** configured within backup application and passed to HNAS.

After initially configured, day-to-day operation of snapshots generally involves monitoring file system space utilized by snapshots and tuning the snapshot retention policy, if necessary.

Note that snapshot deletion does not result in instantaneous reclamation of free space—snapshot deletion is a background process. After a snapshot is deleted, an info-level event is posted in the HNAS event log.

HNAS supports individual reporting of snapshot size (or size of freeable space, should a snapshot be deleted). Refer to Management Interface on page 3.

Aggregate snapshot usage can be monitored using `df` HNAS CLI command or *File system* details page on the SMU.

For file systems that support dedupe, the snapshot space `fs-usage` threshold setting is not applicable, and both the `livefs` and `totalfs` threshold settings apply to the total file system usage. All usage is considered live. Alerts will be issued when this usage exceeds the `livefs` thresholds and when it exceeds the total thresholds. Alert thresholds for snapshot usage will be ignored for such file systems.

Note: In this case, it is recommended to monitor free space over time, and ensure there is enough to support the ongoing/estimated rate of change and the snapshot policy.

Management interface

Snapshots can be managed through SMU GUI and through the HNAS CLI

HNAS SMU GUI

Data Protection section contains the following pages relevant to snapshots:

Snapshots	List attributes (including snapshot size) and manually create/delete/rename snapshots
Snapshot rules	Create snapshot retention policy and snapshot schedules
File replication	Associate replication policy with snapshot rule
File replication	Rules define/override automatic snapshot creation/deletion
NDMP History & Snapshots	Rame as previous plus set maximum snapshot retention period

Storage Management>File Systems contains the following page relevant to snapshots:

- **File System Details**: Show actual snapshot usage and warning controls

HNAS CLII

The following HNAS CLI commands are relevant to snapshots:

CLI command	*Description*
`snapshots`	Brief introduction into snapshots
`snapshot-create`	Create a snapshot with the supplied name or by invoking a rule
`snapshot-delete`	Delete the named or all snapshots associated with a rule
`snapshot-delete-all`	Delete all snapshots on a file system
`snapshot-deleter-queue`	Show work on the deleter queue
`snapshot-deletion`	Show/set whether a volume is modifiable by the deleter
`snapshot-list`	Display details of a file system's snapshots, including their size
`snapshot-recover-fs`	Recover a file system from a snapshot
`snapshot-rename`	Rename a manually created snapshot
`snapshot-rule-create`	Create a snapshot rule
`snapshot-rule-delete`	Delete a snapshot rule and associated snapshots/schedules
`snapshot-rule-list`	Display details of the snapshot rules
`snapshot-rule-modify`	Modify a snapshot rule
`fs-browse-snapshots`	Set/set snapshot directory browsing configuration globally

CLI command	*Description*
`kill-snapshots`	Quickly removes all snapshots from a file system (use under support supervision)
`ndmp-snapshot-options`	Display and/or set NDMP snapshot options
`fs-usage`	Set/list events when the file system hits certain level

Troubleshooting

Most troubleshooting related to snapshot operation is done on the basis of events reported by HNAS server.

The following table shows HNAS events in 11.3.3450.21 code which are relevant to block-based snapshots.

Event ID	*Event*	*Severity*	*Error Message*	*Error Meaning*	*Remediate Action*
7166	SnapshotConfig Unreadable	Severe	The snapshots on file system %s could not be loaded.	When the file system was mounted the existing snapshots could not be validated. Snapshots are disabled on this file system until the file system is fixed.	Under support guidance, run `kill-snapshots` or `fixfs` - note that either will delete any existing snapshots on the file system.
7176	SnapshotDelete FatalCorruption Encountered	Severe	Deletion of snapshot (%s) encountered file system inconsistency and has been aborted.	The deletion of a snapshot could not be completed. Snapshot deletion is now suspended on this file system.	Contact your support provider.
7163	SnapshotRuled CreateFailed	Warning	Failed to create snapshot by rule (%s).	An attempt to create a snapshot by rule failed. The file system, rule name and reason are indicated in parentheses.	Refer to the `snapshot-create` man page for an explanation of failure conditions. If necessary, contact your support provider.
7164	SnapshotCreate Failed	Warning	Failed to create snapshot (%s).	An attempt to create a snapshot failed. The file system and reason are indicated in parentheses.	Refer to the `snapshot-create` man page for an explanation of failure conditions. If necessary, contact your support provider.

Event ID	Event	Severity	Error Message	Error Meaning	Remediate Action
7169	SnapshotDelete BlocksListValidator Incorrect	Warning	Deletion of snapshot (%s) encountered an incorrect snapshot blocks list validator.	The deletion of a snapshot completed with warnings. This may indicate a problem with the file system.	Contact your support provider.
7179	SnapshotDelete InUseByReplication	Warning	The snapshot (%s) is marked as in-use by %u replication session(s) and cannot be deleted.	The snapshot cannot be deleted as it is required by one or more replication sessions.	Wait until the replication completes or abort the replication before deleting the snapshot.
7160	SnapshotCreated	Info	Snapshot created (%s).	A snapshot has been created. The snapshot name and file system are indicated in parentheses.	No action required.
7161	SnapshotDelete Started	Info	Snapshot marked for deletion (%s).	A snapshot has been marked for deletion. New requests to the snapshot will not be granted. The snapshot name and file system are indicated in parentheses.	No action required.
7162	SnapshotDelete Finished	Info	Snapshot deleted (%s).	A snapshot has been deleted. Disk space consumed by the snapshot has been freed. The snapshot name and file system are indicated in parentheses.	No action required.

Performance tuning

Under most circumstances, there is no additional tuning required for snapshot operation. In some cases, however, snapshot deletion process may need to be tuned for more or less aggressive performance, and only under Support guidance.

Block freeing is constrained by two variables (no reboot is required, if changed):

- `fsm set max-unlink-blocks-per-checkpoint`

  This is not snapshot specific, and the default is 0 “no limit.”

- `fsm set snapshot-max-unlink-blocks-per-checkpoint`

  This is snapshot-specific, and the default is 100,000.

The number of file systems that can have active deletions of block-based snapshots at any time is constrained by the following **`telc`** variable (a reboot is required to take effect):

- `block_based_workers_count`

The default is 16 file systems

To check the status of the snapshot delete queue use:

- `snapshot-deleter-queue [<file system>] [--verbose]`

This displays the snapshot deleter's current work queue, including snapshots undergoing or queued for deletion.

Interaction with other HNAS features

Dedupe

Dedupe is a disk-modifying operation of freeing up file system blocks. The block to be freed up will be preserved in the snapshot, so it will not be free. Deduplicated space will only be realized when all snapshots referencing deduped content are deleted (however initiated).

The ‘`df`’ space reporting command output may not show the space savings after a dedupe run due to blocks still used by snapshot.

In a deduped file system, it is not possible to separate the references to a block made through snapshots and through the live file system. The snapshot block usage reporting for deduped file systems have the following characteristics:

- The `snapshot-list` command reports the logical (hydrated) block usage of a snapshot
- Deleting a snapshot on a deduped file system may not free all the reported space when the snapshot is deleted
- The space used by snapshots can exceed the physical size of the file system. To avoid confusion, the ‘df’ command’s Snapshot field will state "NA" for file systems that support dedupe
- Running `kill-snapshots` on a file system that supports dedupe will leak all snapshot blocks. The `fs-reclaim-leaked-space` command will recover most of the leaked

space, but `fixfs` will be needed to recover all the lost block snapshot column to display "NA"

Note: Always use `fixfs` under support guidance. The file system must be offline before running `fixfs`, which may also take a very long time to complete.

Dedupe process also creates its own file system snapshots, used by dedupe for the static and consistent point-in-time view of the file system, while the dedupe process is running. These automatically created snapshots are named "*Dedupe full indexing*" or "*Dedupe incremental indexing*", depending on the type of dedupe job run. Once the dedupe process is finished, the dedupe-created file system snapshot is automatically deleted.

Virtual Infrastructure Integrator (V2I)

V2I mechanism provides file-based clones (backups), which are similar to snapshots. File system snapshots provide crash-consistent view of VMs, similar to V2I clones in a non-quiesced mode. File-based clones (backups) are more space efficient than file system snapshots because the VM files alone are captured at a specific point in time instead of every file on the file system.

Backup, SRA, or file replication

By default, snapshots are a byproduct of NDMP backup, SRA, or file replication, unless backup or replication snapshots were explicitly disabled.

Data Migrator to Cloud (DM2C) and Hitachi Content Platform (HCP)

HNAS Release 11.2 introduced advanced support for snapshot integration when archiving to HCP using DM2C Option. This feature is only available with DM2C and is not supported in the classic Data Migrator.

Important: In release 11.2 and later, DM2C keeps its own journal of what XVLs have been deleted from the primary file system where a snapshot of the XVL is still present.

Once a snapshot is removed on HNAS, DM2C is triggered to check the journal to see if any objects should be removed from HCP, if no more instances of the XVL exist in any other snapshot. This way, DM2C will always clean up objects on the HCP side at the appropriate time, and there is no reliance on coordination between DM2C and a native pruning mechanism on HCP.

If DM2C determines a snapshot has been restored along with the XVLs for a file marked for deletion on HCP, DM2C will reverse-migrate the files and will return the data to HNAS, so it will not be deleted from HCP. This feature is automatically installed with HNAS v11.2 and no special configuration information is required.

Best Practices

- Monitor snapshot usage on a regular basis.
- If snapshot usage is larger than desired, adjust retention policy by reducing appropriate snapshot queue size.
- Adjust the frequency of snapshots to meet your recovery point objective.
- If V2I backup is utilized, disable file system level snapshots when backing up virtual machines.

- If dedupe is enabled and running, weigh benefits of snapshots versus dedupe space reclamation.

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