



Technical Report

NetApp Deduplication for FAS and V-Series Deployment and Implementation Guide

Carlos Alvarez, NetApp
Feb 2011 | TR-3505 Version 8

ABSTRACT

This technical report covers NetApp® deduplication for FAS and V-Series. The report describes in detail how to implement and use deduplication and provides information on best practices, operational considerations, and troubleshooting.

This information is useful for both NetApp and channel partner sales and services field personnel who need to understand details in order to deploy solutions with specific applications that include deduplication.

TABLE OF CONTENTS

1	DEPLOYMENT AND IMPLEMENTATION GUIDE	6
2	INTRODUCTION AND OVERVIEW OF NETAPP DEDUPLICATION	6
3	NETAPP DEDUPLICATION FOR FAS AND V-SERIES	7
3.1	DEDUPLICATED VOLUMES	8
3.2	DEDUPLICATION METADATA	9
3.3	GENERAL DEDUPLICATION FEATURES	10
4	CONFIGURATION AND OPERATION	10
4.1	OVERVIEW OF REQUIREMENTS.....	11
4.2	INSTALLING AND LICENSING DEDUPLICATION	12
4.3	COMMAND SUMMARY	12
4.4	INTERPRETING SPACE USAGE AND SAVINGS.....	13
4.5	DEDUPLICATION QUICK START.....	14
4.6	END-TO-END DEDUPLICATION EXAMPLES	14
4.7	CONFIGURING DEDUPLICATION SCHEDULES.....	19
5	SIZING FOR PERFORMANCE AND SPACE EFFICIENCY.....	21
5.1	BEST PRACTICES.....	21
5.2	PERFORMANCE.....	21
5.2.1	PERFORMANCE OF THE DEDUPLICATION OPERATION.....	22
5.2.2	IMPACT ON THE SYSTEM DURING DEDUPLICATION	22
5.2.3	I/O PERFORMANCE OF DEDUPLICATED VOLUMES	23
5.2.4	PAM AND FLASH CACHE CARDS.....	23
5.3	SPACE SAVINGS	24
5.3.1	TYPICAL SPACE SAVINGS BY DATASET	24
5.3.2	SPACE SAVINGS ON EXISTING DATA	25
5.3.3	DEDUPLICATION METADATA OVERHEAD.....	25
5.3.4	SPACE SAVINGS ESTIMATION TOOL	26
5.4	LIMITATIONS.....	27
5.4.1	GENERAL CAVEATS	27
5.4.2	MAXIMUM FLEXIBLE VOLUME SIZE.....	27
5.4.3	MAXIMUM SHARED DATA LIMIT FOR DEDUPLICATION.....	30
5.4.4	MAXIMUM TOTAL DATA LIMIT FOR DEDUPLICATION.....	30
5.4.5	NUMBER OF CONCURRENT DEDUPLICATION PROCESSES.....	32
6	DEDUPLICATION WITH OTHER NETAPP FEATURES.....	33
6.1	MANAGEMENT TOOLS	33
6.2	DATA PROTECTION	33

6.2.1	SNAPSHOT COPIES	33
6.2.2	SNAPRESTORE.....	33
6.2.3	VOLUME SNAPMIRROR.....	34
6.2.4	QTREE SNAPMIRROR.....	35
6.2.5	SNAPVAULT	36
6.2.6	OPEN SYSTEMS SNAPVAULT (OSSV)	37
6.2.7	SNAPMIRROR SYNC	38
6.2.8	SNAPLOCK.....	38
6.3	CLUSTERING TECHNOLOGIES.....	38
6.3.1	DATA ONTAP CLUSTER-MODE	38
6.3.2	ACTIVE-ACTIVE CLUSTER CONFIGURATION	39
6.3.3	METROCLUSTER.....	40
6.4	OTHER NETAPP FEATURES	40
6.4.1	QUOTAS.....	40
6.4.2	FLEXCLONE VOLUMES	40
6.4.3	FLEXCLONE FILES.....	41
6.4.4	64-BIT AGGREGATE SUPPORT	42
6.4.5	FLEXCACHE	42
6.4.6	NONDISRUPTIVE VOLUME MOVEMENT	42
6.4.7	NETAPP DATA MOTION.....	42
6.4.8	PERFORMANCE ACCELERATION MODULE AND FLASH CACHE CARDS	43
6.4.9	SMTAPE	43
6.4.10	DUMP	43
6.4.11	NONDISRUPTIVE UPGRADES.....	43
6.4.12	NETAPP DATAFORT ENCRYPTION.....	44
6.4.13	READ REALLOCATION (REALLOC)	44
6.4.14	VOL COPY COMMAND	44
6.4.15	AGGREGATE COPY COMMAND	45
6.4.16	MULTISTORE (VFILER)	45
6.4.17	SNAPDRIVE	45
6.4.18	LUNS	45
7	DEDUPLICATION BEST PRACTICES WITH SPECIFIC APPLICATIONS	49
7.1	VMWARE BEST PRACTICES	49
7.1.1	VMFS DATASTORE ON FIBRE CHANNEL OR ISCSI: SINGLE LUN.....	51
7.1.2	VMWARE VIRTUAL DISKS OVER NFS/CIFS	52
7.1.3	DEDUPLICATION ARCHIVE OF VMWARE.....	53

7.2	MICROSOFT SHAREPOINT BEST PRACTICES	54
7.3	MICROSOFT SQL SERVER BEST PRACTICES.....	54
7.4	MICROSOFT EXCHANGE BEST PRACTICES.....	54
7.5	LOTUS DOMINO BEST PRACTICES.....	55
7.5.1	DOMINO ATTACHMENT AND OBJECT SERVICE (DAOS)	55
7.5.2	DOMINO ENCRYPTION.....	55
7.5.3	DOMINO QUOTAS	55
7.5.4	DOMINO PERFORMANCE	55
7.6	ORACLE BEST PRACTICES	56
7.7	TIVOLI STORAGE MANAGER BEST PRACTICES.....	56
7.8	SYMANTEC BACKUP EXEC BEST PRACTICES	56
7.9	BACKUP BEST PRACTICES	57
8	8 TROUBLESHOOTING	57
8.1	DEDUPLICATION WILL NOT RUN	57
8.1.1	CHECK DEDUPLICATION LICENSING	57
8.2	MAXIMUM VOLUME AND DATA SIZE LIMITS	58
8.2.1	MAXIMUM VOLUME SIZE LIMIT.....	58
8.2.2	MAXIMUM SHARED DATA LIMIT FOR DEDUPLICATION.....	60
8.2.3	MAXIMUM TOTAL DATA LIMIT	60
8.3	DEDUPLICATION SCANNER OPERATIONS TAKING TOO LONG TO COMPLETE.....	62
8.4	LOWER-THAN-EXPECTED SPACE SAVINGS	62
8.5	SYSTEM SLOWDOWN	63
8.5.1	UNEXPECTEDLY SLOW READ PERFORMANCE CAUSED BY ADDING DEDUPLICATION.....	63
8.5.2	UNEXPECTEDLY SLOW WRITE PERFORMANCE CAUSED BY ADDING DEDUPLICATION	64
8.5.3	SYSTEM RUNNING MORE SLOWLY SINCE ENABLING DEDUPLICATION	65
8.6	REMOVING SPACE SAVINGS (UNDO).....	65
8.6.1	UNDEDUPLICATING A FLEXIBLE VOLUME.....	66
8.6.2	UNDEDUPLICATING A FLEXIBLE VOLUME.....	67
8.6.3	REVERTING A FLEXIBLE VOLUME TO BE ACCESSIBLE IN OTHER DATA ONTAP RELEASES	68
8.7	WHERE TO COLLECT TROUBLESHOOTING INFORMATION.....	68
8.7.1	LOCATION OF LOGS AND ERROR MESSAGES	68
8.7.2	UNDERSTANDING DEDUPLICATION ERROR MESSAGES	68
8.7.3	UNDERSTANDING OPERATIONS MANAGER EVENT MESSAGES.....	69
8.8	ADDITIONAL DEDUPLICATION REPORTING.....	69
8.8.1	DEDUPLICATION REPORTING WITH SIS STATUS.....	69
8.9	DEDUPLICATION REPORTING WITH SIS STAT.....	70

8.10 WHERE TO GET MORE HELP	73
8.10.1 CONTACT INFORMATION FOR SUPPORT	73
8.10.2 INFORMATION TO GATHER BEFORE CONTACTING SUPPORT	73
9 ADDITIONAL READING AND REFERENCES	73
10 VERSION TRACKING	75

LIST OF TABLES

Table 1) Overview of deduplication requirements.	11
Table 2) Deduplication commands.	12
Table 3) Interpreting df -s results.....	13
Table 6) Deduplication quick start.	14
Table 7) Typical deduplication space savings.	24
Table 8) Maximum supported volume sizes for deduplication.....	29
Table 9) Maximum total data limit in a deduplicated volume.	31
Table 10) Maximum total data limit example for deduplication.....	32
Table 11) Maximum supported volume sizes for deduplication.....	59
Table 12) Example illustrating maximum total data limit - deduplication only.	60
Table 13) Maximum total data limit in a deduplicated volume.....	61

LIST OF FIGURES

Figure 1) How NetApp deduplication works at the highest level.	7
Figure 2) Data structure in a deduplicated volume	8
Figure 3) VMFS data store on Fibre Channel or iSCSI - single LUN3.	55
Figure 4) VMware virtual disks over NFS/CIFS.....	56
Figure 5) Archive of VMware with deduplication.	57

1 DEPLOYMENT AND IMPLEMENTATION GUIDE

This is the deployment and implementation guide for NetApp deduplication. This document is publically available online at <http://media.netapp.com/documents/tr-3505.pdf>.

2 INTRODUCTION AND OVERVIEW OF NETAPP DEDUPLICATION

Despite the introduction of less-expensive ATA disk drives, one of the biggest challenges for companies today continues to be the cost of storage. There is a desire to reduce storage consumption (and therefore storage cost per megabyte) by reducing the number of disks it takes to store data. NetApp deduplication is a key component of NetApp's storage efficiency technologies, which enable users to store the maximum amount of data for the lowest possible cost.

This document focuses on NetApp deduplication for FAS and V-Series. NetApp deduplication is a process that can be triggered when a threshold is reached, scheduled to run when it is most convenient, or run as part of an application. It will remove duplicate blocks in a volume or LUN.

This section is an overview of how deduplication works for FAS and V-Series systems.

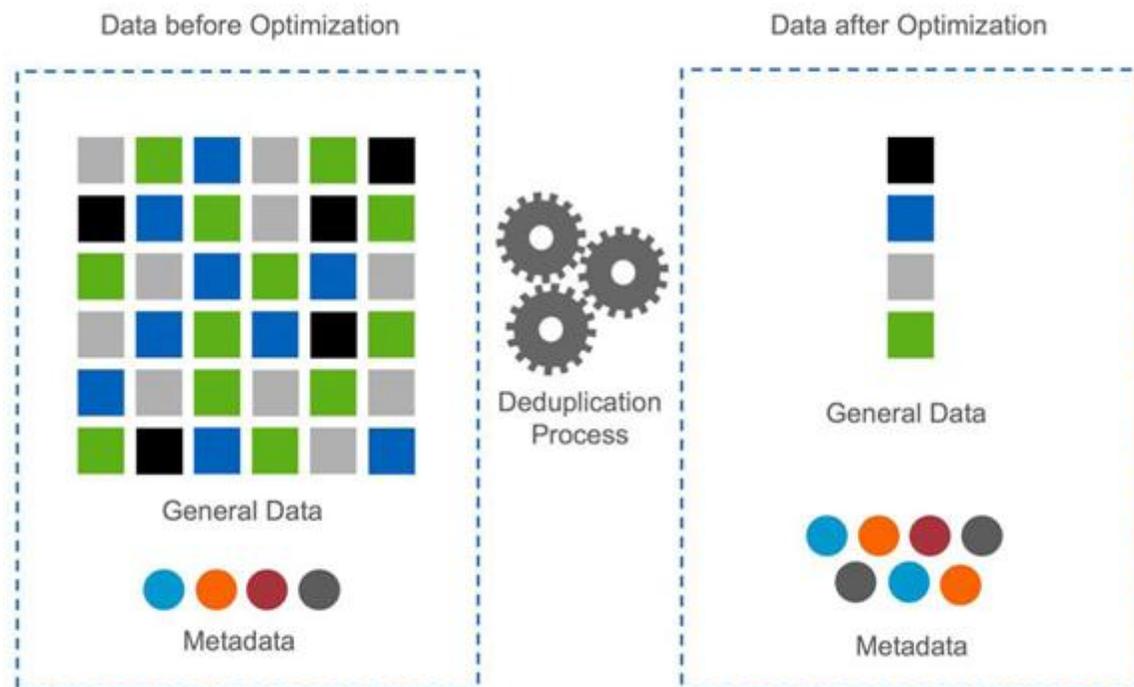
Notes:

1. Whenever references are made to deduplication in this document, the reader should assume we are referring to NetApp deduplication for FAS and V-Series.
2. The reader should assume the same information applies to both FAS and V-Series systems, unless otherwise noted.
3. NetApp deduplication for VTL is not within the scope of this technical report.

3 NETAPP DEDUPLICATION FOR FAS AND V-SERIES

Part of NetApp's storage efficiency offerings, NetApp deduplication for FAS provides block-level deduplication within the entire flexible volume on NetApp storage systems. Beginning with Data ONTAP® 7.3, V-Series also supports deduplication. NetApp V-Series is designed to be used as a gateway system that sits in front of third-party storage, allowing NetApp storage efficiency and other features to be used on third-party storage.

Figure 1) How NetApp deduplication works at the highest level.



Essentially, deduplication stores only unique blocks in the flexible volume and creates a small amount of additional metadata in the process. Notable features of deduplication include:

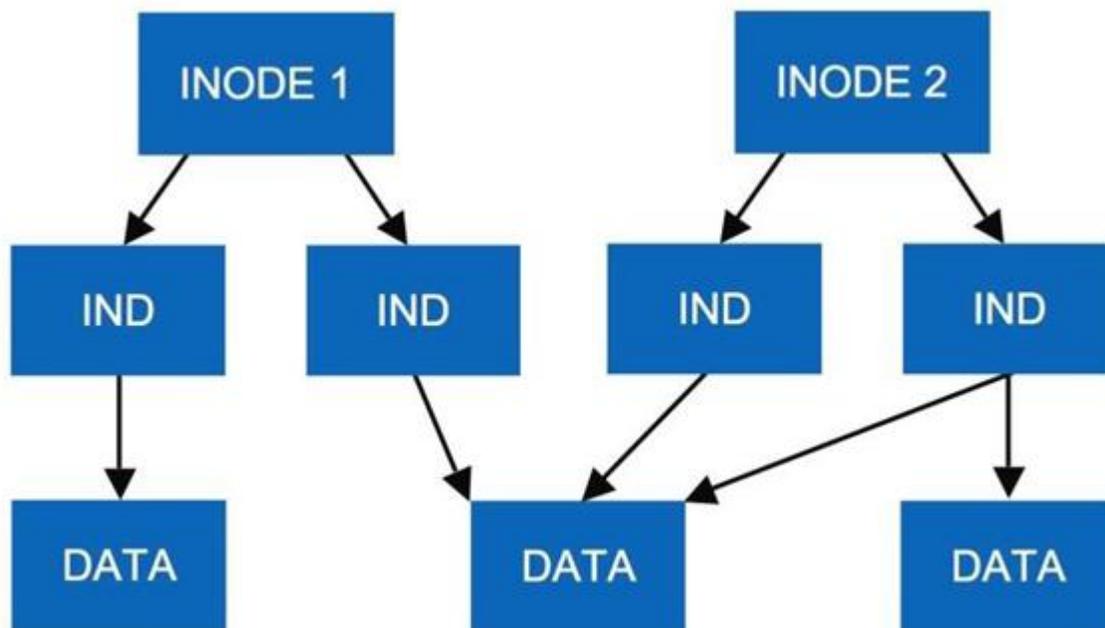
- It works with a high degree of granularity: that is, at the 4KB block level.
- It operates on the active file system of the flexible volume. Any block referenced by a Snapshot® copy is not made “available” until the Snapshot copy is deleted.
- It is a background process that can be configured to run automatically, scheduled, or run manually through the command line interface (CLI), NetApp Systems Manager, or NetApp Provisioning Manager.
- It is application transparent, and therefore it can be used for deduplication of data originating from any application that uses the NetApp system.
- It is enabled and managed by using a simple CLI or GUI.
- It can be enabled and can deduplicate blocks on flexible volumes with new and existing data.

In summary, this is how deduplication works. Newly saved data on the FAS system is stored in 4KB blocks as usual by Data ONTAP. Each block of data has a digital fingerprint, which is compared to all other fingerprints in the flexible volume. If two fingerprints are found to be the same, a byte-for-byte comparison is done of all bytes in the block, and, if there is an exact match between the new block and the existing block on the flexible volume, the duplicate block is discarded, and its disk space is reclaimed.

3.1 DEDUPLICATED VOLUMES

A deduplicated volume is a flexible volume that contains shared data blocks. Data ONTAP supports shared blocks in order to optimize storage space consumption. Basically, in one volume, there is the ability to have multiple references to the same data block, as shown in Figure 2.

Figure 2) Data structure in a deduplicated volume.



In Figure 2, the number of physical blocks used on the disk is 3 (instead of 5), and the number of blocks saved by deduplication is 2 (5 minus 3). In this document, these are referred to as used blocks and saved blocks.

Each data block has a block count reference that is kept in the volume metadata. As additional indirect blocks (“IND” in Figure 2) point to the data, or existing ones stop pointing to it, this value is incremented or decremented accordingly. When no indirect blocks point to a data block, it is released.

The NetApp deduplication technology allows duplicate 4KB blocks anywhere in the flexible volume to be deleted, as described in the following sections.

The maximum sharing for a block is 255. This means, for example, that if there are 500 duplicate blocks, deduplication would reduce that to only 2 blocks. Also note that this ability to share blocks is different from the ability to keep 255 Snapshot copies for a volume.

3.2 DEDUPLICATION METADATA

The core enabling technology of deduplication is fingerprints. These are unique digital “signatures” for every 4KB data block in the flexible volume.

When deduplication runs for the first time on a flexible volume with existing data, it scans the blocks in the flexible volume and creates a fingerprint database, which contains a sorted list of all fingerprints for used blocks in the flexible volume.

After the fingerprint file is created, fingerprints are checked for duplicates, and, when found, first a byte-by-byte comparison of the blocks is done to make sure that the blocks are indeed identical. If they are found to be identical, the block’s pointer is updated to the already existing data block, and the new (duplicate) data block is released.

Releasing a duplicate data block entails updating the indirect inode pointing to it, incrementing the block reference count for the already existing data block, and freeing the duplicate data block.

In real time, as additional data is written to the deduplicated volume, a fingerprint is created for each new block and written to a change log file. When deduplication is run subsequently, the change log is sorted, its sorted fingerprints are merged with those in the fingerprint file, and then the deduplication processing occurs.

There are two change log files, so that as deduplication is running and merging the new blocks from one change log file into the fingerprint file, new data that is being written to the flexible volume is causing fingerprints for these new blocks to be written to the second change log file. The roles of the two files are then reversed the next time that deduplication is run. (For those familiar with Data ONTAP usage of NVRAM, this is analogous to when it switches from one half to the other to create a consistency point.)

Note: When deduplication is run for the first time on an empty flexible volume, it still creates the fingerprint file from the change log.

Here are some additional details about the deduplication metadata:

- There is a fingerprint record for every 4KB data block, and the fingerprints for all the data blocks in the volume are stored in the fingerprint database file.
- Fingerprints are not deleted from the fingerprint file automatically when data blocks are freed. When a threshold of 20% new fingerprints is reached, the stale fingerprints are deleted. This can also be done by a manual operation from the command line.
- In Data ONTAP 7.2.X, all the deduplication metadata resides in the flexible volume.
- Starting with Data ONTAP 7.3.0, part of the metadata resides in the volume, and part of it resides in the aggregate outside the volume. The fingerprint database and the change log files that are used in the deduplication process are located outside of the volume in the aggregate and are therefore not captured in Snapshot copies. This change enables deduplication to achieve higher space savings. However, some other temporary metadata files created during the deduplication operation are still placed inside the volume. These temporary metadata files are deleted once the deduplication operation is complete. These temporary metadata files can get locked in Snapshot copies if the Snapshot copies are created during a deduplication operation. The metadata files remain locked until the Snapshot copies are deleted.
- During an upgrade from Data ONTAP 7.2 to 7.3, the fingerprint and change log files are moved from the flexible volume to the aggregate level during the next deduplication process following the upgrade. During the deduplication process where the fingerprint and change log files are being moved from the volume to the aggregate, the `sis status` command displays the message “Fingerprint is being upgraded.”
- In Data ONTAP 7.3 and later, the deduplication metadata for a volume is located outside the volume, in the aggregate. During a reversion from Data ONTAP 7.3 to a pre-7.3 release, the deduplication

metadata is lost during the revert process. To obtain optimal space savings, use the `sis start -s` command to rebuild the deduplication metadata for all existing data. If this is not done, the existing data in the volume retains the space savings from deduplication run prior to the revert process; however, any deduplication that occurs after the revert process applies only to data that was created after the revert process. It does not deduplicate against data that existed prior to the revert process. The `sis start -s` command can take a long time to complete, depending on the size of the logical data in the volume, but during this time the system is available for all other operations. Before using the `sis start -s` command, make sure that the volume has sufficient free space to accommodate the addition of the deduplication metadata to the volume. The deduplication metadata uses 1% to 6% of the logical data size in the volume.

- The deduplication configuration files are located inside the volume, not the aggregate; therefore, the configuration will never be reset, including VSM, unless the user runs the “`sis undo`” command.

3.3 GENERAL DEDUPLICATION FEATURES

Deduplication is enabled on a per flexible volume basis. It can be enabled on any number of flexible volumes in a storage system. Deduplication can be scheduled to run in one of four different ways:

- Scheduled on specific days and at specific times
- Manually, by using the command line
- Automatically, when 20% new data has been written to the volume
- Automatically on the destination volume, when used with SnapVault®

Only one deduplication scanner process can run on a flexible volume at a time.

Up to eight concurrent deduplication processes can run concurrently on the same NetApp storage system. Individually you can have a maximum of eight concurrent deduplication processes on eight volumes within the same NetApp storage system.

Beginning with Data ONTAP 7.3.1, deduplication checkpoint restart allows a deduplication process that was interrupted to continue from the last checkpoint. Prior to Data ONTAP 7.3.1, an interrupted deduplication process would result in a restart of the entire deduplication process.

4 CONFIGURATION AND OPERATION

This section discusses the requirements for using deduplication, how to configure it, and various aspects of managing it. Although it discusses some basic things, in general it assumes both that the NetApp storage system is already installed and running and that the reader is familiar with basic NetApp administration.

4.1 OVERVIEW OF REQUIREMENTS

Table 1) Overview of deduplication requirements.

Requirement	Deduplication
Hardware	<p>NearStore® R200</p> <p>FAS2000 series</p> <p>FAS3000 series</p> <p>FAS3100 series</p> <p>FAS3200 series</p> <p>FAS6000 series</p> <p>IBM N5000 series</p> <p>IBM N7000 series</p> <p>Note: Starting with Data ONTAP 7.3, the V-Series systems corresponding to the NetApp FAS systems and IBM N series Gateway systems listed above are also supported.</p>
Minimum Data ONTAP version required	Data ONTAP 7.2.5.1 (7-Mode only for 8.0.X)
Licenses required	A-SIS NearStore license (required for versions of Data ONTAP prior to 8.0)
Volume type supported	FlexVol® only, no traditional volumes
Maximum volume size	For Data ONTAP 8.0.1 the limit is 16TB for all platforms for deduplication. For earlier releases see “Maximum Flexible Volume Size” in later section.
Supported protocols	All

Some additional considerations with regard to maximum volume sizes include:

- Once an upgrade is complete, the new maximum volume sizes for Data ONTAP are in effect. The volumes can be grown to the new maximum volume size at that point.
- When considering a downgrade or revert, NetApp highly recommends consulting NetApp Global Services for best practices.
- During a reversion to an earlier version of Data ONTAP with smaller volume limits, volumes should be within the limits of the lower version of Data ONTAP. When downgrading to 7.2.X, you cannot simply resize the volume. Instead, a new flexible volume that is within the maximum volume size limits must be created, and the data must be moved to that volume before the downgrade.
- If a downgrade occurs from 7.3.1 to 7.3.0, where the volume size was greater than the 7.3.0 volume size limit, the volume goes offline. This would also apply if downgrading from 8.0.1 to 8.0 or to 7.3.x. If this happens, contact NetApp Global Services for assistance in bringing the volume back online.

4.2 INSTALLING AND LICENSING DEDUPLICATION

Deduplication is included with Data ONTAP 7.2.5.1 and later. Beginning with Data ONTAP 8.0.1, deduplication is available in 7-Mode but not Cluster-Mode. To use deduplication, it just needs to be licensed. To add the license use the following command:

```
license add <license key>
```

The deduplication license is a no charge license.

Note: NetApp deduplication on V-Series is only supported with block checksum (BCS), not zone checksum (ZCS). For more information, refer to [TR-3461: V-Series Best Practice Guide](#).

4.3 COMMAND SUMMARY

The following sections describe the various deduplication commands.

Table 2) Deduplication commands.

Command	Summary
sis on <vol/volname >	Enables deduplication on the flexible volume specified.
sis start <vol/volname >	Begins the deduplication process on the flexible volume specified.
sis start -s <vol/volname >	Begins the deduplication process on the flexible volume specified and performs a scan of the flexible volume to process existing data. This option is typically used upon initial configuration and deduplication on an existing flexible volume that contains undeduplicated data. (There's no need to use this option on a volume that has just been created and doesn't contain any data.)
sis start -sp <vol/volname >	Begins the deduplication process on the flexible volume specified by using the existing checkpoint information, regardless of the age of the checkpoint information. This option should only be used with the -s option.
sis start -d <vol/volname >	Deletes the existing checkpoint information. This option is used to delete checkpoint information that is still considered valid. By default, checkpoint information is considered invalid after 24 hours.
sis status [-l] <vol/volname >	Returns the current status of deduplication for the specified flexible volume. The -l option displays a long list.
sis config [-s sched] <vol/volname >	Creates an automated deduplication schedule. When deduplication is first enabled on a flexible volume, a default schedule is configured, running it each day of the week at midnight. If the auto option is used, deduplication is triggered when 20% new data is written to the volume. Starting with Data ONTAP 7.3.1, the 20% threshold can be adjusted by using the auto@num option, where num is a two-digit number to specify the percentage.
sis stop <vol/volname >	Suspends an active deduplication process on the flexible volume.

Command	Summary
<code>sis off <vol/volname ></code>	Deactivates deduplication on the flexible volume specified. This means that there will be no additional change logging or deduplication operations, but the flexible volume remains a deduplicated volume, and the storage savings are preserved. If this command is used, and then deduplication is turned back on for this flexible volume, the flexible volume must be rescanned with the <code>sis start -s</code> command.
<code>sis check <vol/volname ></code>	Verifies and updates the fingerprint database for the flexible volume specified; includes purging stale fingerprints (diag mode only).
<code>sis stat <vol/volname ></code>	Displays the statistics of flexible volumes that have deduplication enabled (diag mode only).
<code>sis undo <vol/volname ></code>	Reverts a deduplicated volume to a normal flexible volume (diag or advanced mode only).
<code>sis revert_to <ver> <vol/volname></code>	Converts the sis metafiles to appropriate lower version format (currently 7.2 or 7.3). Note: When no volume name is provided, <code>revert_to</code> runs on all volumes with deduplication enabled.
<code>sis revert_to <ver> -delete <vol/volname></code>	Deletes the original metafiles that existed for release 8.0.X, resulting in no metafiles being created when <code>revert_to</code> is run. When no volume is provided, the <code>-delete</code> option will be applied to all volumes that are reverted by the <code>revert_to</code> command.
<code>sis revert_to <ver> -cleanup <vol/volname></code>	Deletes the metafiles of the lower version that were created by previously running <code>sis revert_to</code> . Useful if user decides not to revert to previous release after all. When no volume is provided, the <code>-cleanup</code> option will delete all existing metafiles that are the format of the release that was specified when running the <code>revert_to</code> command.
<code>df -s <vol/volname ></code> <code>df -s <volname></code>	Shows space savings from deduplication as well as actual physical used capacity per volume.

4.4 INTERPRETING SPACE USAGE AND SAVINGS

The `df` command shows savings as well as actual physical space used per volume. The following is a summary of the various displayed fields:

Table 3) Interpreting `df -s` results.

<code>df -s</code>	Description
Filesystem	Volume name.
Used	Amount of physical space used on the volume.
Saved	Number of blocks saved by deduplication.

df -s	Description
%Saved	Percentage of space saved from deduplication.

To determine the used (logical space used) capacity you would add the values "Used" + "Saved" from the df -s output.

4.5 DEDUPLICATION QUICK START

This section describes the steps necessary to enable and configure deduplication on a FlexVol volume.

Table 4) Deduplication quick start.

	New Flexible Volume	Existing Flexible Volume
Add deduplication license	<code>license add <a-sis license code></code>	
Enable deduplication on volume	<code>sis on /vol/<volname></code>	
Run initial deduplication scan	NA	<code>sis start -s /vol/<volname></code>
Create, modify, delete deduplication schedules	<code>sis config -s [sched] /vol/<volname></code>	
Manually run deduplication (if not using schedules)	<code>sis start /vol/<volname></code>	
Monitor status of deduplication	<code>sis status /vol/<volname></code>	
Monitor deduplication savings	<code>df -s /vol/<volname></code>	

4.6 END-TO-END DEDUPLICATION EXAMPLES

This section describes the steps necessary to enable and configure deduplication on a FlexVol volume using Data ONTAP 8.0.1 or earlier.

The first example describes the process of creating a new flexible volume and then configuring, running, and monitoring deduplication.

The second example describes the process of enabling deduplication on an already existing flexible volume that contains data. We then run deduplication to get savings on the already existing data on disk.

Note: The steps are spelled out in detail, so the process appears much longer than it would be in the real world.

EXAMPLE ONE: CREATING A NEW VOLUME AND ENABLING DEDUPLICATION

This example creates a place to archive several large data files. The destination NetApp storage system is called fas6070c-ppe02.

1. Add deduplication licenses to the destination storage system.

```
fas6070-ppe02> license add <deduplication license key>
```

2. Create a flexible volume (no larger than the 16TB limit).

```
fas6070-ppe02> vol create volArchive aggrTest 200g
Creation of volume 'volArchive' with size 200g on containing aggregate
'aggrTest' has completed.
```

3. Enable deduplication on the flexible volume (`sis on`), and verify that it is turned on. The `vol status` command shows the attributes for flexible volumes.

After you turn deduplication on, Data ONTAP lets you know that if this were an existing flexible volume that already contained data before deduplication was enabled, you would want to run `sis start -s`. In this example it's a brand new flexible volume, so that's not necessary.

```
fas6070-ppe02> sis on /vol/volArchive
SIS for "/vol/volArchive" is enabled.
Already existing data could be processed by running "sis start -s
/vol/volArchive".
```

```
fas6070-ppe02> vol status volArchive
      Volume State           Status           Options
      volArchive online      raid_dp, flex
                               sis
                               Volume UUID: 70191fea-854c-11df-806c-00a098001674
                               Containing aggregate: 'aggrTest'
```

4. Another way to verify that deduplication is enabled on the flexible volume is to check the output from running `sis status` on the flexible volume.

```
fas6070-ppe02> sis status /vol/volArchive
Path           State           Status           Progress
/vol/volArchive Enabled         Idle             Idle for 00:01:04
```

5. Turn off the default deduplication schedule.

```
fas6070-ppe02> sis config /vol/volArchive
Path           Schedule           Minimum Blocks Shared
/vol/volArchive sun-sat@0          1
```

```
fas6070-ppe02> sis config -s - /vol/volArchive
```

```
fas6070-ppe02> sis config /vol/volArchive
Path           Schedule           Minimum Blocks Shared
/vol/volArchive -                   1
```

6. Mount the flexible volume and copy data into the new archive directory flexible volume.
7. Examine the flexible volume. Use the `df -s` command to examine the space used and the space saved. Note that no savings have been achieved so far by simply copying data to the flexible volume, even though deduplication is turned on. Since deduplication was enabled, all the new blocks have

had their fingerprints written to the change log file. Until deduplication is actually run, the duplicate blocks will not be removed.

```
fas6070-ppe02*> df -s
Filesystem      used      saved      %saved
/vol/volArchive/ 140090168    0          0%
```

8. Manually run deduplication on the flexible volume. This causes the change log to be processed, fingerprints to be sorted and merged, and duplicate blocks to be found.

```
fas6070-ppe02> sis start /vol/volArchive
The SIS operation for "/vol/volArchive" is started.
```

9. Use `sis status` to monitor the progress of deduplication.

```
fas6070-ppe02> sis status /vol/volArchive
Path           State      Status      Progress
/vol/volArchive Enabled    Active      43 GB Searched
```

```
fas6070-ppe02> sis status /vol/volArchive
Path           State      Status      Progress
/vol/volArchive Enabled    Active      99 GB Searched
```

```
Path           State      Status      Progress
/vol/volArchive Enabled    Active      132 GB Searched
```

```
fas6070-ppe02> sis status /vol/volArchive
Path           State      Status      Progress
/vol/volArchive Enabled    Active      18 GB (28%) Done
```

```
fas6070-ppe02> sis status /vol/volArchive
Path           State      Status      Progress
/vol/volArchive Enabled    Active      50 GB (76%) Done
```

```
fas6070-ppe02> sis status /vol/volArchive
Path           State      Status      Progress
/vol/volArchive Enabled    Idle        Idle for 00:00:05
```

10. When `sis status` indicates that the flexible volume is once again in the “Idle” state, deduplication has finished running, and you can check the space savings it provided in the flexible volume.

```
fas6070-ppe02> df -s
Filesystem      used      saved      %saved
/vol/volArchive/ 73185512 108777800    60%
```

11. Adjust the deduplication schedule as required in your environment. Refer to the section on [configuring deduplication schedule](#) for more specifics.

That is all there is to it.

EXAMPLE TWO: ENABLING DEDUPLICATION ON AN EXISTING VOLUME

This example enables deduplication on an existing flexible volume with data preexisting on the volume. While it is not necessary, this example includes the steps involved if you wanted to deduplicate the data already on disk, in addition to the new writes to disk. The destination NetApp storage system is called fas6070c-ppe02, and the volume is /vol/volExisting.

- 1) Add deduplication licenses to the destination storage system.

```
fas6070-ppe02> license add <deduplication license key>
```

- 2) Enable deduplication on the flexible volume (sis on), and verify that it is turned on. The *vol status* command shows the attributes for flexible volumes.

```
fas6070-ppe02> sis on /vol/volExisting
SIS for "/vol/volExisting" is enabled.
Already existing data could be processed by running "sis start -s
/vol/volExisting".
```

```
fas6070-ppe02> vol status volExisting
      Volume State      Status      Options
volExisting online      raid_dp, flex  nosnap=on,
                               sis          guarantee=none,
                               fractional_reserve=0
      Volume UUID: d685905a-854d-11df-806c-00a098001674
      Containing aggregate: 'aggrTest'
```

- 3) Turn off the default deduplication schedule.

```
fas6070-ppe02> sis status /vol/volExisting
Path              State      Status      Progress
/vol/volExisting  Enabled   Idle        Idle for 00:02:47
fas6070-ppe02> sis config /vol/volExisting
Path              Schedule      Minimum Blocks Shared
/vol/volExisting  sun-sat@0    1
fas6070-ppe02> sis config -s - /vol/volExisting
fas6070-ppe02> sis config /vol/volExisting
Path              Schedule      Minimum Blocks Shared
/vol/volExisting  -            1
```

- 4) Examine the flexible volume. Use the *df -s* command to examine the storage consumed and the space saved.

```
fas6070-ppe02>df -s
Filesystem      used      saved      %saved
/vol/volExisting/ 181925592 0          0
```

At this time only new data will have fingerprints created. From here if you only want to deduplicate new data from this point, you could skip to step 12.

If you want to deduplicate the preexisting data on disk, use the following additional steps:

- 5) Disable the deduplication schedule.

```
fas6070-ppe02> sis config /vol/volExisting
Path              Schedule      Minimum Blocks Shared
/vol/volExisting  sun-sat@0    1
fas6070-ppe02> sis config -s - /vol/volExisting
```

```
fas6070-ppe02> sis config /vol/volExisting
Path                               Schedule           Minimum Blocks Shared
/vol/volExisting                   -                 1
```

Note: Steps 6, 7, and 11 pertain to Snapshot copies and are optional, but should be used whenever possible to maximize savings. See the section on using deduplication with Snapshot copies for details.

6) Record the current Snapshot schedule for the volume. Disable the snap schedule.

```
fas6070-ppe02> snap sched volExisting
Volume volExisting: 5 7 10
fas6070-ppe02> snap sched volExisting 0 0 0
```

7) Delete as many Snapshot copies as possible.

```
fas6070-ppe02> snap list volExisting
Volume volExisting
working...
```

%/used	%/total	date	name
23% (0%)	20% (0%)	Jul 18 11:59	snap.1
24% (1%)	20% (0%)	Jul 18 23:59	snap.2
26% (3%)	23% (2%)	Jul 19 11:59	snap.3
26% (0%)	23% (0%)	Jul 19 23:59	snap.4
27% (2%)	24% (1%)	Jul 20 11:59	snap.5

```
fas6070-ppe02> snap delete volExisting snap.1
fas6070-ppe02> Tue Jul 20 16:20:06 EDT [waf1.snap.delete:info]: Snapshot copy
sn      ap.1 on volume volExisting NetApp was deleted by the Data ONTAP
function snapcmd      _delete. The unique ID for this Snapshot copy is (3,
6768).
```

```
fas6070-ppe02> snap delete volExisting snap.2
fas6070-ppe02> Tue Jul 20 16:20:10 EDT [waf1.snap.delete:info]: Snapshot copy
sn      ap.2 on volume volExisting NetApp was deleted by the Data ONTAP
function snapcmd      _delete. The unique ID for this Snapshot copy is (2,
6760).
```

```
fas6070-ppe02> snap delete volExisting snap.3
fas6070-ppe02> Tue Jul 20 16:20:15 EDT [waf1.snap.delete:info]: Snapshot copy
sn      ap.3 on volume volExisting NetApp was deleted by the Data ONTAP
function snapcmd      _delete. The unique ID for this Snapshot copy is (4,
6769).
```

```
fas6070-ppe02> snap list volExisting
Volume volExisting
working...
```

%/used	%/total	date	name
12% (0%)	10% (0%)	Jul 19 23:59	snap.4
13% (1%)	10% (1%)	Jul 20 11:59	snap.5

8) Manually run the duplication process during a low system usage time. The deduplication process with the `-s` option causes the fingerprint database to be recreated for all existing data within the volume, and then removes the duplicates.

```

fas6070-ppe02> sis start -s /vol/volExisting
The file system will be scanned to process existing data in /vol/volExisting.
This operation may initialize related existing metafiles.
Are you sure you want to proceed (y/n)? y
The SIS operation for "/vol/volExisting" is started.
fas6070-ppe02> Wed Jun 30 13:22:09 EDT [waf1.scan.start:info]: Starting SIS
volume scan on volume volExisting.

```

9) Use `sis status` to monitor the progress of deduplication.

```

fas6070-ppe02> sis status /vol/volExisting
Path                State      Status      Progress
/vol/volExisting    Enabled   Active      29 GB Scanned

fas6070-ppe02> sis status /vol/volExisting
Path                State      Status      Progress
/vol/volExisting    Enabled   Active      33 GB (52%) Done

fas6070-ppe02> sis status /vol/volExisting
Path                State      Status      Progress
/vol/volExisting    Enabled   Idle        Idle for 00:04:01

```

10) When `sis status` indicates that the flexible volume is once again in the “Idle” state, deduplication has finished running, and you can check the additional space savings it provided in the flexible volume.

```

fas6070-ppe02> df -s
Filesystem          used      saved          %saved
/vol/volExisting/  73054628  108956860      60%

```

11) Reconfigure the Snapshot schedule for the volume.

```

snap sched volExisting 5 7 10

```

12) Adjust the deduplication schedule as required for your environment. Refer to the section on [configuring deduplication schedule](#) for more specifics.

That’s all there is to it.

4.7 CONFIGURING DEDUPLICATION SCHEDULES

It’s best to set up a schedule for deduplication so that you don’t have to run it manually each time. This section describes how to configure schedules with deduplication from the CLI. You can also use Systems Manager or Provisioning Manager to configure this using policies.

The `sis config` command is used to configure and view deduplication schedules for flexible volumes. Here is the usage syntax:

```

sis help config
sis config [[-s schedule] | [ -m minimum_blocks_shared] ] <path> ...]
    - Sets up, modifies and retrieves schedule and minimum blocks
    shared value of SIS volumes.

```

Run with no arguments, `sis config` returns the schedules for all flexible volumes that have deduplication enabled. The following example shows the four different formats that the reported schedules can have.

```

sis config
Path          Schedule
/vol/dvol_1  -
/vol/dvol_2  23@sun-fri
/vol/dvol_3  auto
/vol/dvol_4  sat@6

```

The meaning of each of these schedule types is as follows:

- On flexible volume dvol_1, deduplication is not scheduled to run.
- On flexible volume dvol_2, deduplication is scheduled to run every day from Sunday to Friday at 11 p.m.
- On flexible volume dvol_3, deduplication is set to autoschedule. This means that deduplication is triggered by the amount of new data written to the flexible volume, specifically when there are 20% new fingerprints in the change log.
- On flexible volume dvol_4, deduplication is scheduled to run at 6 a.m. on Saturday.

When the `-s` option is specified, the command sets up or modifies the schedule on the specified flexible volume. The schedule parameter can be specified in one of four ways:

```

[day_list][@hour_list]
[hour_list][@day_list]
-
auto

```

The `day_list` specifies which days of the week deduplication should run. It is a comma-separated list of the first three letters of the day: sun, mon, tue, wed, thu, fri, sat. The names are not case sensitive. Day ranges such as mon-fri can also be used. The default `day_list` is sun-sat.

The `hour_list` specifies which hours of the day deduplication should run on each scheduled day. The `hour_list` is a comma-separated list of the integers from 0 to 23. Hour ranges such as 8-17 are allowed.

Step values can be used in conjunction with ranges. For example, 0-23/2 means "every 2 hours." The default `hour_list` is 0: that is, midnight on the morning of each scheduled day.

If "-" is specified, there is no scheduled deduplication operation on the flexible volume.

The `auto` schedule causes deduplication to run on that flexible volume whenever there are 20% new fingerprints in the change log. This check is done in a background process and occurs every hour. Beginning with Data ONTAP 7.3.1, this 20% threshold can be set to a customized value.

When deduplication is enabled on a flexible volume for the first time, an initial schedule is assigned to the flexible volume. This initial schedule is sun-sat@0, which means "once every day at midnight."

To configure the schedules shown earlier in this section, the following commands would be issued:

```

sis config -s - /vol/dvol_1
sis config -s 23@sun-fri /vol/dvol_2

```

```
sis config -s auto /vol/dvol3
sis config -s sat@6 /vol/dvol_4
```

5 SIZING FOR PERFORMANCE AND SPACE EFFICIENCY

This section discusses the deduplication behavior that you can expect. It includes best practices, space savings estimates and tools, volume size limits, general caveats, overhead, and performance. Information in this report comes from testing, observations, and knowledge of how deduplication functions.

5.1 BEST PRACTICES

This section describes best practices and lessons learned from internal tests and customer deployments.

- Deduplication consumes system resources and can alter the data layout on disk. Due to the application's I/O pattern and the effect of deduplication on the data layout, the read and write I/O performance can vary. The space savings and the performance impact depend on the application and the data contents.
- NetApp recommends carefully considering and measuring the performance impact of deduplication, as well as the sizing aspects, in a test environment before deploying deduplication.
- If there is only a small amount of new data, run deduplication infrequently, because there is very little benefit in running it frequently in such a case. The frequency for running deduplication depends on the rate of change of the data in the flexible volume.
- The more concurrent deduplication scanner processes you're running, the more system resources are consumed.
- Given the previous two items, the best option is to do *one* of the following:
 - Stagger the deduplication schedule for the flexible volumes so that deduplication processes run on alternate days, reducing the possibility of running too many concurrent processes.
 - Use Auto mode so that deduplication runs only when significant additional data has been written to each flexible volume. (This tends to naturally stagger out when deduplication runs in smaller environments.)
 - Run deduplication manually.
 - Run deduplication nightly to minimize the amount of new data to be deduplicated, requiring less time to complete.
- If NetApp Snapshot copies are required, run the deduplication process before creating them to minimize the amount of space used before the data gets locked into a Snapshot copy. Make sure that the deduplication process has completed before creating the Snapshot copy. If a Snapshot copy is created before the deduplication process has completed, it is likely to result in lower space savings.
- For deduplication to run properly, you need to leave some free space for the deduplication metadata. For information about how much extra space to leave in the volume and in the aggregate, see section 5.3, [Deduplication Metadata Overhead](#).

5.2 PERFORMANCE

This section discusses the performance aspects of deduplication.

Since deduplication is part of Data ONTAP, it is tightly integrated with the NetApp WAFL[®] (Write Anywhere File Layout) file structure. Because of this, deduplication is optimized to perform with high efficiency. It is able to leverage the internal characteristics of Data ONTAP to create and compare digital fingerprints, redirect data pointers, and free up redundant data blocks.

However, the following factors can affect the performance of the deduplication processes and the I/O performance of deduplicated volumes:

- The application and the type of dataset being used
- The data access pattern (for example, sequential versus random access, the size and pattern of the I/O)
- The amount of duplicate data, the amount of total data, and the average file size
- The nature of the data layout in the volume
- The amount of changed data between deduplication runs
- The number of concurrent deduplication processes running
- The number of volumes that have deduplication enabled on the system
- The hardware platform—the amount of CPU/memory in the system
- The amount of load on the system
- Disk types ATA/FC/SAS, and the RPM of the disk
- The number of disk spindles in the aggregate

Because of these factors, NetApp recommends that performance with deduplication should be carefully measured in a test setup and taken into sizing consideration before deploying deduplication in performance-sensitive solutions. For information about testing deduplication performance in a customer environment, see [TR-3849: NetApp Deduplication for FAS and V-Series Performance and Savings Testing in Customer Environments](#).

5.2.1 PERFORMANCE OF THE DEDUPLICATION OPERATION

The performance of the deduplication operation itself varies widely depending on the factors previously described, and this determines how long it takes this low-priority, background process to finish running.

On a FAS6080 system with no other load on the system, deduplication performance can be as high as 120MB/sec for a single deduplication session. With multiple sessions running concurrently, deduplication performance can be as high as 180MB/sec for the same scenario. This total bandwidth is divided across the multiple sessions, and each session gets a fraction of the aggregated throughput. There are many factors that can affect the deduplication throughput. To get an idea of how long it takes for a deduplication process to complete, suppose that the deduplication process is running on a flexible volume at a conservative rate of 100MB/sec. If 1TB of new data has been added to the volume since the last deduplication update, this deduplication operation takes about 2.5 to 3 hours to complete. (There are no configurable parameters that can tune the deduplication process; that is, the priority of this background process in Data ONTAP is fixed.)

This scenario is merely an example. Deduplication typically completes much faster following the initial scan, when it is run nightly. Running deduplication nightly can minimize the amount of new data to be deduplicated, requiring less time to complete.

5.2.2 IMPACT ON THE SYSTEM DURING DEDUPLICATION

The deduplication process runs as a low-priority background process on the system. However, it can still affect the performance of user I/O and other applications running on the system.

The number of deduplication processes that are running and the phase that each process is running in can affect the performance of other applications running on the system. Up to eight concurrent background deduplication processes can run concurrently on the same NetApp storage system. If there is an attempt to run an additional deduplication process beyond the maximum, the deduplication process is queued and automatically started when a running deduplication process completes.

Here are some observations about running deduplication on a FAS3050 system:

- With eight deduplication processes running, and no other processes running, deduplication uses 15% of the CPU in its least invasive phase, and nearly all of the available CPU in its most invasive phase.
- When one deduplication process is running, there is 0% to 15% performance degradation on other applications.
- With eight deduplication processes running, there might be a 15% to more than 50% performance penalty on other applications running on the system.

5.2.3 I/O PERFORMANCE OF DEDUPLICATED VOLUMES

WRITE PERFORMANCE OF A DEDUPLICATED VOLUME

The impact of deduplication on the write performance of a system is a function of the hardware platform that is being used, as well as the amount of load that is placed on the system.

For deduplicated volumes, if the load on a system is low—that is, for systems in which the CPU utilization is around 50% or lower—there is a negligible difference in performance when writing data to a deduplicated volume, and there is no noticeable impact on other applications running on the system. On heavily used systems, however, where the system is nearly saturated, the impact on write performance can be expected to be around 15% for most NetApp systems. The performance impact is more noticeable on higher end systems than on lower end systems. On the FAS6080 system, this performance impact has been as much as 7% for sequential writes and can be as much as 35% for random writes. Testing has not yet been completed on the new FAS6200 models.

Note: The deduplication numbers are for FC drives. If ATA drives are used in a system, the performance impact would be greater.

Because the performance impact varies, it should be tested before implementing into production.

READ PERFORMANCE OF A DEDUPLICATED VOLUME

When data is read from a deduplication-enabled volume, the impact on the read performance varies depending on the difference between the deduplicated block layout and the original block layout. There is minimal impact on random reads.

Because deduplication alters the data layout on the disk, it can affect the performance of sequential read applications such as dump source, qtree SnapMirror®, or SnapVault source, SnapVault restore, and other sequential read-heavy applications. This impact is more noticeable in Data ONTAP releases earlier than Data ONTAP 7.2.6 and Data ONTAP 7.3.1 with datasets that contain blocks with many repeating patterns (such as applications that preinitialize data blocks to a value of zero). Data ONTAP 7.2.6 and Data ONTAP 7.3.1 introduced specific optimizations, referred to as *intelligent cache*, that improve the performance of these workloads to be close to the performance of nondeduplicated datasets. This is useful in many scenarios and has been especially successful in virtualized environments. In addition, the Performance Acceleration Modules (PAM) and Flash Cache cards are also deduplication aware and use intelligent caching.

5.2.4 PAM AND FLASH CACHE CARDS

The original PAM card (RAM-based, 16GB per module) was available starting with Data ONTAP 7.3.0. The Flash Cache card (PAM II, Flash-based at 256 or 512GB per module) is available as of Data ONTAP 7.3.2.

In environments with high amounts of shared (deduplicated) blocks that are read repeatedly, the PAM and Flash Cache card can significantly reduce the number of disk reads, thus improving the read performance.

The amount of performance improvement with the PAM or Flash Cache card depends on the amount of shared blocks, the access rate, the active dataset size, and the data layout.

Adding a PAM or Flash Cache card to a system does not increase the deduplication maximum volume size for that system.

The PAM and Flash Cache cards have provided significant performance improvements in VMware® environments. These advantages are further enhanced when combined with shared block technologies, such as NetApp deduplication or NetApp FlexClone® technology.

For more information about the PAM and Flash Cache cards, refer to [TR-3705: NetApp and VMware VDI Best Practices](#).

5.3 SPACE SAVINGS

This section discusses the storage savings for deduplication.

Comprehensive testing with various datasets has been performed to determine typical space savings in different environments. These results were obtained in three ways:

- Running deduplication on various production datasets in NetApp
- NetApp systems deployed in the real world running deduplication
- NetApp personnel and end users running a simulation tool on various datasets (See section 5.4, [Space Savings Estimation Tool](#), for information about how to use this tool.)

5.3.1 TYPICAL SPACE SAVINGS BY DATASET

Table 5 summarizes typical space savings.

Table 5) Typical deduplication space savings.

Dataset Type	Application Type	Deduplication Only
File services/IT infrastructure		30%
Virtual servers and desktops		70%
Database	Oracle® OLTP	0%
	Oracle DW	15%
	SQL Server®	20%
E-mail, collaborative	Exchange 2003/2007	3%
	Exchange 2010	15%
Engineering data		30%
Geoseismic		3%
Archival data		25%
Backup data		95%

These results are based on internal testing and customer feedback and are considered realistic and typically achievable. Savings estimates can be validated for an existing environment by using the [Space Savings Estimation Tool](#) (SSET), as discussed in section 5.4.

Note: Nonrepeating archival data such as image files and encrypted data is generally not considered a good candidate for deduplication.

Note: The deduplication space savings in Table 5 result from deduplicating a dataset one time, with the following exception. In cases where the data is being backed up or archived over and over again, the realized storage savings get better and better, achieving 20:1 (95%) in many instances. The backup case also assumes that the backup application is maintaining data coherency with the original, and that the data's block alignment will not be changed during the backup process. If these criteria are not true, then there can be a discrepancy between the space savings recognized on the primary and secondary systems.

5.3.2 SPACE SAVINGS ON EXISTING DATA

A major benefit of deduplication is that it can be used to deduplicate existing data in the flexible volumes. It is realistic to assume that there will be Snapshot copies—perhaps many—of the existing data.

When you first run deduplication on a flexible volume, the storage savings will probably be rather small or even nonexistent. As previous Snapshot copies expire, some small savings are realized, but they too are likely to be low.

During this period of old Snapshot copies expiring, it is typical for new data to be created on the flexible volume with new Snapshot copies being created. The storage savings might continue to stay low.

When the last Snapshot copy that was created before deduplication was run is deleted, the storage savings should increase noticeably.

Therefore the question is when to run deduplication again in order to achieve maximum capacity savings. The answer is that deduplication should be run, and allowed to complete, before the creation of each and every Snapshot copy; this provides the most storage savings benefit. However, depending on the flexible volume size and possible performance impact on the system, this might not always be advisable.

5.3.3 DEDUPLICATION METADATA OVERHEAD

This section discusses the storage overhead that deduplication introduces. Although deduplication can provide substantial storage savings in many environments, a small amount of storage overhead is associated with it. This should be considered when sizing the flexible volume.

The total storage used by the deduplication metadata files is approximately 1% to 6% of the total data in the volume. Total data = used space + saved space, as reported when using `df -s` (that is, the size of the data before it is deduplicated). So for 1TB of total data, the metadata overhead would be approximately 10GB to 60GB. The breakdown of the overhead associated with the deduplication metadata is as follows:

- There is a fingerprint record for every 4KB data block, and the fingerprint records for all of the data blocks in the volume are stored in the fingerprint database file. An overhead of less than 2% is associated with this database file.
- The size of the deduplication change log files depends on the rate of change of the data and on how frequently deduplication is run. This accounts for less than 2% overhead in the volume.
- Finally, when deduplication is running, it creates some temporary files that could account for up to 2% of the size of the volume. These temporary metadata files are deleted when the deduplication process has finished running.

In Data ONTAP 7.2.X, all of the deduplication metadata files just described reside in the volume, and this metadata is therefore captured and locked in the Snapshot copies of the volume as well. Starting with Data ONTAP 7.3, part of the metadata still resides in the volume, and part of it resides in the aggregate outside of the volume. The fingerprint database and the change log files are located outside of the volume

in the aggregate and are therefore not captured in Snapshot copies. This change enables deduplication to achieve higher space savings. However, the other temporary metadata files created during the deduplication operation are still placed inside the volume. These temporary metadata files are deleted when the deduplication operation is complete. However, if Snapshot copies are created during a deduplication operation, these temporary metadata files can get locked in Snapshot copies, and they remain there until the Snapshot copies are deleted.

The guideline for the amount of extra space that should be left in the volume and aggregate for the deduplication metadata overhead is as follows:

If you are running Data ONTAP 7.2.X, you will need up to 6% of the total data to be stored in the deduplicated volume available in the volume.

For example, if 100GB of data is to be deduplicated, then there should be 6GB worth of available space in the volume.

Note: The amount of space required for deduplication metadata depends on the amount of data being deduplicated in the volume, and not on the size of the volume.

If you are running Data ONTAP 7.3.X, there will both volume and aggregate space requirements, as follows:

- **Volume deduplication overhead.** For each volume with deduplication enabled, up to 2% of the logical amount of data written to that volume is required in order to store volume deduplication metadata.
- **Aggregate deduplication overhead.** For each aggregate that contains any volumes with deduplication enabled, up to 4% of the logical amount of data contained in all of those volumes with deduplication enabled is required in order to store the aggregate deduplication metadata.

For example, if 100GB of data is to be deduplicated in a single volume, then there should be 2GB of available space in the volume and 4GB of space available in the aggregate. As a second example, consider a 2TB aggregate with four volumes, each 400GB in size, in the aggregate. Three volumes are to be deduplicated, with 100GB of data, 200GB of data, and 300GB of data, respectively. The volumes need 2GB, 4GB, and 6GB of space, respectively, and the aggregate needs a total of 24GB ((4% of 100GB) + (4% of 200GB) + (4% of 300GB)) = 4+8+12 = 24GB) of space available in the aggregate.

Note: The amount of space required for deduplication metadata depends on the amount of data being deduplicated in the volumes, and not on the size of the volumes or the aggregate.

5.3.4 SPACE SAVINGS ESTIMATION TOOL

The actual amount of data space reduction depends on the type of data. For this reason, the Space Savings Estimation Tool (SSET 3.0) should be used to analyze the actual dataset to determine the effectiveness of deduplication.

When executed, SSET crawls through all the files in the specified path and estimates the space savings that will be achieved with deduplication. Although actual deduplication space savings might deviate from what the estimation tool predicts, use and testing so far indicate that in general, the actual results are within +/-5% of the space savings that the tool predicts.

OVERVIEW OF SSET

SSET is available to NetApp employees and NetApp partners. It performs nonintrusive testing of the dataset to determine the effectiveness of deduplication.

This tool is intended for use only by NetApp personnel to analyze data at current or prospective NetApp users. By installing this software, the user agrees to keep this tool and any results from this tool confidential between them and NetApp.

The deduplication Space Savings Estimation Tool is available for Linux[®] and Microsoft[®] Windows[®] systems, which have the data available locally or use CIFS/NFS. For complete usage information, see the SSET readme file.

LIMITATIONS OF SSET

SSET runs on either a Linux system or a Windows system.

It is limited to evaluating a maximum of 2TB of data. If the given path contains more than 2TB, the tool processes the first 2TB of data, indicates that the maximum size has been reached, and displays the results for the 2TB of data that it processed. The rest of the data is ignored.

The tool is designed to examine data either that is available locally or that uses NFS/CIFS only. The data does not need to reside on a NetApp storage system for SSET to perform an analysis.

For more information about SSET, see the SSET readme file. The SSET tool, including the readme file, can be downloaded by NetApp personnel and NetApp partners from the NetApp Field Portal.

5.4 LIMITATIONS

This section discusses what is supported and what is not supported, and the dos and don'ts of using deduplication. Some of this information is covered elsewhere in this and other technical reports as well.

5.4.1 GENERAL CAVEATS

Deduplication metadata (fingerprint file and change logs) is not deduplicated.

Other metadata, such as directory metadata, is also not deduplicated. Therefore, for heavily replicated directory environments with a large number of small files (for example, Web space), the space savings that can be achieved might be low.

Backup of the deduplicated volume using NDMP is supported, but there is no space optimization when the data is written to tape because it's a logical operation. (This could actually be considered an advantage, because in this case the tape does not contain a proprietary format.) To preserve the deduplication space savings on tape, NetApp SMTape is recommended.

Quotas are based on logical space usage. Therefore, when deduplication is used in an environment where quotas are used, the quotas cannot be oversubscribed on a volume. For example, a user with a quota limit of 1TB can't store more than 1TB of data in a deduplicated volume, even if the data has been deduplicated and fits into less than 1TB of physical space on the storage system.

Storage administrators can use the saved space as desired.

Only data in the active file system is deduplicated. Data pointed to by Snapshot copies that were created before deduplication was run is not released until the Snapshot copy is deleted or expires. For more information about deduplication and Snapshot copies, refer to the section called "Snapshot Copies" in this document.

5.4.2 MAXIMUM FLEXIBLE VOLUME SIZE

Deduplication has a maximum flexible volume size limit. If deduplication is to be used on a volume, that volume must be within the size limit. If an attempt is made to enable deduplication on a volume that is

larger than the maximum volume size, it fails with an appropriate error message. If an upgrade occurs to a version of Data ONTAP that has larger maximum flexible volume size limits, then the larger size limits will be in effect once the upgrade is complete.

The maximum volume size limit varies based on the platform and Data ONTAP version. This is an important consideration if the flexible volumes are moved to a different platform with a smaller maximum flexible volume size.

Table 6 shows the deduplication maximum flexible volume size limits (including any snap reserve space) for the different NetApp storage system platforms and versions of Data ONTAP. For versions of Data ONTAP before 7.3.1, if a volume ever gets larger than its maximum flexible volume size limit and is later shrunk to a smaller size, deduplication cannot be enabled on that volume.

Table 6) Maximum supported volume sizes for deduplication.

	Data ONTAP 7.2.5.1–7.3.0	Data ONTAP 7.3.1.1 and Later 7.3.x	Data ONTAP 8.0	Data ONTAP 8.0.1
FAS2020/N3300	0.5TB	1TB	Not supported	Not supported
FAS2040/N3400	Not supported	3TB ¹	3TB	16TB
FAS2050/N3600	1TB	2TB	Not supported	Not supported
R200	4TB	4TB	Not supported	Not supported
FAS3020/N5200	1TB	2TB	Not supported	Not supported
FAS3040/N5300	3TB	4TB	4TB	16TB
FAS3050/N5500	2TB	3TB	Not supported	Not supported
FAS3070/N5600	6TB	16TB	16TB	16TB
FAS3140/N6040	3TB	4TB	4TB	16TB
FAS3160/N6060	6TB	16TB	16TB	16TB
FAS3170/N6070	10TB	16TB	16TB	16TB
FAS3210	Not supported	4TB ²	Not supported	16TB
FAS3240	Not supported	16TB ²	Not supported	16TB
FAS3270	Not supported	16TB ²	Not supported	16TB
FAS6030/N7600	10TB	16TB	16TB	16TB
FAS6040/N7700	10TB	16TB	16TB	16TB
FAS6070/N7800	16TB	16TB	16TB	16TB
FAS6080/N7900	16TB	16TB	16TB	16TB
FAS6210	Not supported	Not supported	Not supported	16TB
FAS6240	Not supported	Not supported	Not supported	16TB
FAS6280	Not supported	Not supported	Not supported	16TB

¹ Deduplication on this platform is only supported with Data ONTAP 7.3.2 or later. Prior versions are not supported.

² Only supported with Data ONTAP 7.3.5 or later. Prior versions are not supported.

5.4.3 MAXIMUM SHARED DATA LIMIT FOR DEDUPLICATION

In addition to the maximum supported volume sizes, there is also a maximum shared data limit for deduplication. Recall that deduplication achieves savings by keeping a single copy of a block and replacing all duplicates of that block with pointers to the single block. As a result of deduplication, much of the logical data is stored as pointers to a reduced number of physical blocks written to the volume. It is very common for a single block to have many pointer references. The logical data that ends up being stored as pointers is called *shared data*. The maximum shared data limit per volume for deduplication is 16TB, regardless of the platform type. Once this limit is reached, any additional new data written to the volume is not deduplicated. Writes to the volume continue to work successfully.

5.4.4 MAXIMUM TOTAL DATA LIMIT FOR DEDUPLICATION

Both the maximum supported volume sizes and the maximum shared data limit for deduplication limit the maximum data size for a flexible volume. The sum of the two becomes the maximum total (logical data) size limit for a deduplicated volume. For example, in a FAS3140 system that can have a deduplicated volume of up to 4TB in size, 20TB of data can be stored; that is, 4TB + 16TB = 20TB.

Table 7) Maximum total data limit in a deduplicated volume.

	Data ONTAP 7.2.5.1–7.3.0	Data ONTAP 7.3.1.1 and Later 7.3.X	Data ONTAP 8.0	Data ONTAP 8.0.1
FAS2020/N3300	16.5TB	17TB	Not supported	Not supported
FAS2040/N3400	Not supported	19TB ³	19TB	32TB
FAS2050/N3600	17TB	18TB	Not supported	Not supported
R200	20TB	20TB	Not supported	Not supported
FAS3020/N5200	17TB	18TB	Not supported	Not supported
FAS3040/N5300	19TB	20TB	20TB	32TB
FAS3050/N5500	18TB	19TB	Not supported	Not supported
FAS3070/N5600	22TB	32TB	32TB	32TB
FAS3140/N6040	19TB	20TB	20TB	32TB
FAS3160/N6060	22TB	32TB	32TB	32TB
FAS3170/N6070	26TB	32TB	32TB	32TB
FAS3210	Not supported	20TB ⁴	Not supported	32TB
FAS3240	Not supported	32TB ⁴	Not supported	32TB
FAS3270	Not supported	32TB ⁴	Not supported	32TB
FAS6030/N7600	26TB	32TB	32TB	32TB
FAS6040/N7700	26TB	32TB	32TB	32TB
FAS6070/N7800	32TB	32TB	32TB	32TB
FAS6080/N7900	32TB	32TB	32TB	32TB
FAS6210	Not supported	Not supported	Not supported	32TB
FAS6240	Not supported	Not supported	Not supported	32TB
FAS6280	Not supported	Not supported	Not supported	32TB

³ Deduplication on this platform is only supported with Data ONTAP 7.3.2 or later. Prior versions are not supported.

⁴ Only supported with Data ONTAP 7.3.5 or later. Prior versions are not supported.

For example, consider the following theoretical example as it progressively adds more data to a volume. This example assumes that deduplication can provide ~99% savings on this dataset.

The sample system is a FAS3140 running Data ONTAP 7.3.2. Keep in mind that this system has a volume size limit of 4TB and a shared data limit of 16TB, yielding a maximum logical data size of 20TB.

Table 8) Maximum total data limit example for deduplication.

Action	Amount of Shared Data on Volume	Physical Space Usage	Total (Logical) Data Size
4TB volume is created and populated with 4TB of data.	0	4TB	4TB
Deduplication is run.	4TB	1MB	4TB
An additional 4TB is written to the volume.	4TB	4TB	8TB
Deduplication is run again.	8TB	1MB	8TB
An additional 4TB is written to the volume.	8TB	4TB	12TB
Deduplication is run again.	12TB	1MB	12TB
An additional 4TB is written to the volume.	12TB	4TB	16TB
Deduplication is run again.	16TB	1MB	16TB
An additional 4TB is written to the volume.	16TB	4TB	20TB
Deduplication is run again.	16TB	4TB	20TB
Attempt to write additional data to the volume.	New writes fail because the volume is full.		

In this example, when the limit of 16TB of shared data on the volume is reached, any additional writes are not deduplicated. The logical data size is the sum of the amount of shared data and the physical space use on the volume.

5.4.5 NUMBER OF CONCURRENT DEDUPLICATION PROCESSES

It is possible to run a maximum of eight concurrent deduplication processes on eight different volumes on the same NetApp storage system.

If another flexible volume is scheduled to have deduplication run while 8 deduplication scanner processes are already running, deduplication for this additional flexible volume is queued. For example, suppose that a user sets a default schedule (`sun-sat@0`) for 10 deduplicated volumes. Eight are run at midnight, and the remaining 2 are queued. As soon as one of the 8 current deduplication processes completes, one of the queued ones starts; when another deduplication process completes, the second queued one starts. The next time that deduplication is scheduled to run on these same 10 flexible volumes; a round-robin paradigm is used so that the same volumes aren't always the first ones to run.

With Data ONTAP 7.2.X, for manually triggered deduplication runs, if 8 deduplication processes are already running when a command is issued to start another one, the request fails and the operation is not queued. However, starting with Data ONTAP 7.3, the manually triggered deduplication runs are also queued if 8 deduplication operations are already running (including the `sis start -s` command).

6 DEDUPLICATION WITH OTHER NETAPP FEATURES

This section discusses how deduplication will interact with other NetApp products and features.

NetApp deduplication for FAS and V-Series is designed to provide space and cost savings. NetApp deduplication works as a postprocess. Deduplication removes duplicate blocks in a data volume. This section details deduplication interoperability with other NetApp solutions. It gives details regarding support, limitations, and best practices.

When we refer in this guide to deduplication we are referring to NetApp deduplication for FAS and V-Series. NetApp deduplication is only supported on a minimum of Data ONTAP version 7.2.5.1.

6.1 MANAGEMENT TOOLS

Provisioning Manager, Operations Manager, and Protection Manager 3.8 and later support deduplication.

Provisioning Manager 3.8 and later provides the ability to define deduplication policies and configure deduplication across multiple systems from a single management system.

Operations Manager 3.8 and later provides the ability to monitor and report the effects of deduplication across multiple systems from a single management system.

Protection Manager 3.8 and later seamlessly provides optimized management of SnapVault schedules with deduplication.

For additional information about Provisioning Manager, Operations Manager, and Protection Manager, refer to [TR-3440: Operations Manager, Protection Manager, and Provisioning Manager Sizing Guide](#).

6.2 DATA PROTECTION

6.2.1 SNAPSHOT COPIES

Snapshot copies lock blocks on disk that cannot be freed until the Snapshot copy expires or is deleted. On any volume, once a Snapshot copy of data is made, any subsequent changes to that data temporarily require additional disk space, until the Snapshot copy is deleted or expires. The same is true with deduplication-enabled volumes.

Volumes using deduplication do not see the savings initially if the blocks are locked by Snapshot copies. This is because the blocks are not freed until the lock is removed.

Some best practices to achieve the best space savings from deduplication-enabled volumes that contain Snapshot copies include:

- Run deduplication before creating new Snapshot copies.
- Limit the number of Snapshot copies you maintain.
- If possible, reduce the retention duration of Snapshot copies.
- Schedule deduplication only after significant new data has been written to the volume.
- Configure appropriate reserve space for the Snapshot copies.

6.2.2 SNAPRESTORE

SnapRestore® functionality is supported with deduplication, and it works in the same way with or without deduplication. When you initiate a SnapRestore operation on a FlexVol volume, the restored data retains the original space savings.

If you run Data ONTAP 7.3, consider the following points:

- Starting with Data ONTAP 7.3, the deduplication metadata files (the fingerprint database and the change log files) do not get restored when SnapRestore is executed, because they are located outside the volume in the aggregate. In this case, after the SnapRestore operation, there is not a fingerprint database file in the active file system for the data. This data, however, retains the original space savings.
- After SnapRestore, if deduplication is enabled on the volume, any new data written to the volume continues to be deduplicated. However, the deduplication process obtains space savings in the new data only and does not deduplicate between the new data and the restored data. To run deduplication for all the data in the volume (and thus obtain greater space savings), use the `sis start -s` command. This command builds the fingerprint database for all the data in the volume. Be aware that, depending on the size of the logical data in the volume, this process can take a long time to complete.

6.2.3 VOLUME SNAPMIRROR

Volume SnapMirror allows you to back up your data to another location for disaster recovery purposes. Deduplication is supported with volume SnapMirror. Volume SnapMirror operates at the physical block level; thus when deduplication is enabled on the source, the data sent over the wire for replication is also deduplicated and therefore the savings are inherited at the destination. This can significantly reduce the amount of network bandwidth required during replication. The combinations of volume SnapMirror and deduplication shown in Table 9 are supported.

Table 9) Supported deduplication configurations for volume SnapMirror.

Source Volume	Destination Volume	
	Dedupe Enabled	Dedupe Disabled
Dedupe enabled	Supported	Not supported
Dedupe disabled	Not supported	Supported

To run deduplication with volume SnapMirror:

- Both source and destination storage systems should have the same deduplication licenses.
- Both source and destination systems should use an identical release of Data ONTAP. If this is not possible, then the destination must be a newer version than the source.
- Deduplication can only be managed on the source system—the flexible volume at the destination system inherits all the efficiency attributes and storage savings.
- Shared blocks are transferred only once, so deduplication reduces network bandwidth usage.
- The volume SnapMirror update schedule is not tied to the deduplication schedule.
- Maximum volume size limits for deduplicated volumes are constrained to the lower limit between the source and the destination systems.
- Volume SnapMirror typically sees smaller transfer sizes as a result of deduplication space savings. However, if deduplication is enabled for the first time on a volume with existing data, file buffer trees are relinked. Because all data transferred with volume SnapMirror must be deswizzled on the destination as well as the relinked buffer trees, there might be an increase in deswizzling time when using deduplication on the volume SnapMirror source. When configuring volume SnapMirror and deduplication, it is important to consider the deduplication schedule and the volume SnapMirror schedule. As a best practice, start volume SnapMirror transfers of a deduplicated volume after deduplication is complete (that is, not in the middle of the deduplication process). This is to avoid sending undeduplicated data and additional temporary metadata files over the network. If the

temporary metadata files in the source volume are locked in Snapshot copies, they also consume extra space in the source and destination volumes.

THE IMPACT OF MOVING DEDUPLICATION METADATA FILES OUTSIDE THE VOLUME

Starting with Data ONTAP 7.3, most of the deduplication metadata resides in the aggregate outside the volume. Therefore, it does not get captured in Snapshot copies, and also volume SnapMirror does not replicate this data. This provides additional network bandwidth savings. However, some temporary metadata files are still kept inside the volume and are deleted when the deduplication operation is complete. If Snapshot copies are created during the deduplication operation, these temporary metadata files are locked in Snapshot copies. A volume SnapMirror update that is initiated during a deduplication process transfers these temporary metadata files over the network. To prevent this extra data from being replicated, schedule the volume SnapMirror updates to take place after the deduplication operation has finished running on the source volume.

In case of a disaster at the primary location, you might need to break the volume SnapMirror relationship and have the volume SnapMirror destination start serving data. In this case, there will not be a fingerprint database file at the destination system for the existing data on the destination volume. However, the existing data retains the space savings from the deduplication operations performed earlier on the original volume SnapMirror source. Also, the deduplication process continues for new data being written to the volume and creates the fingerprint database for this new data. The deduplication process obtains space savings in the new data only and doesn't deduplicate between the new data and the old data. To run deduplication for all the data in the volume (and thus obtain greater space savings), use the `sis start -s` command. This command builds the fingerprint database for all the data in the volume. It is typically only needed if the original primary will be offline for an extended amount of time. Depending on the size of the logical data in the volume, this process might take a long time to complete.

Important: Before using the `sis start -s` command, make sure that both the volume and the aggregate containing the volume have sufficient free space to accommodate the addition of the deduplication metadata. For information about how much extra space to leave for the deduplication metadata, see “Deduplication Metadata Overhead” section of [TR-3505: Deduplication for FAS and V-Series Deployment and Implementation Guide](#).

6.2.4 QTREE SNAPMIRROR

Qtree SnapMirror allows you to back up your data to another location for disaster recovery purposes. Deduplication is supported with qtree SnapMirror. The combinations of qtree SnapMirror and deduplication shown in Table 10 are supported.

Table 10) Supported deduplication configurations for qtree SnapMirror.

Source Volume	Destination Volume	
	Dedupe Enabled	Dedupe Disabled
Dedupe enabled	Supported	Supported
Dedupe disabled	Supported	Supported

To run deduplication with qtree SnapMirror:

- Remember to consider that deduplication space savings are always lost during the network transfer; therefore, there are no network bandwidth savings.
- Both source and destination storage systems run deduplication independently of each other.
- The deduplication schedule is not tied to the qtree SnapMirror update. In other words, the deduplication process does not automatically start at the completion of the qtree SnapMirror transfer.

As a best practice, NetApp recommends performing qtree SnapMirror updates after the deduplication process on the source volume has finished running. If a qtree SnapMirror update occurs while the deduplication process is running on the source volume, then, in addition to the transfer of the changed data blocks, some redundant data blocks might also get transferred to the destination. If deduplication is not running on the destination volume, then the redundant data that is transferred occupies extra storage space on the destination volume.

NetApp also recommends that if deduplication is used on the source volume, then it should also be used on the destination volume. However, you are not restricted using deduplication on the source volume only.

Typically, there are only a couple of qtree SnapMirror base Snapshot copies on the destination storage system. If Snapshot copies are not retained long term, they are constantly rotated out and the deduplicated blocks are freed as the Snapshot copies roll off.

If users want to keep Snapshot copies long term (as a replacement for SnapVault or for other reasons such as the ability to have writable, reverse, or resync copies in the event of a disaster), it is possible that the deduplicated data can be locked in Snapshot copies for longer periods of time, which reduces the deduplication storage savings. This situation can arise when users create Snapshot copies manually or by using the `snap sched` command.

The best practice when using qtree SnapMirror with deduplication is to let qtree SnapMirror use the minimum number of Snapshot copies it requires (essentially, keep the latest version).

6.2.5 SNAPVAULT

The behavior of deduplication with SnapVault is similar to the behavior with qtree SnapMirror, except for the following points:

- The deduplication schedule is tied to the SnapVault schedule on the destination system. The deduplication schedule on the source is not tied to the SnapVault update schedule, and it can be configured independently on a volume, just like qtree SnapMirror.
- Every SnapVault update (baseline or incremental) kicks off the deduplication process on the destination after the archival Snapshot copy is created.
- Deduplication results in logical-level changes to the deduplicated blocks. This means that both qtree SnapMirror and SnapVault recognize these as changed blocks and include them in their data transfers to the destination volume. As a result, qtree SnapMirror and SnapVault transfers that follow a `sis start -s` command are likely to be much larger than normal. NetApp recommends running deduplication on the primary if possible before running baseline transfers for qtree SnapMirror and SnapVault. For preexisting qtree volume SnapMirror or SnapVault relationships, it is important to take into consideration the big surge of data involved in the transfer and to plan accordingly.
- The archival Snapshot copy is replaced with a new one after deduplication has finished running on the destination. (The name of this new Snapshot copy is the same as that of the archival copy, but the creation time of this copy is changed.)
- The deduplication schedule on the destination cannot be configured manually, and the `sis start` command is not allowed either. However, the `sis start -s` command can be run manually on the destination.

- The SnapVault update is not tied to the deduplication operation; that is, a subsequent incremental update is allowed to run while the deduplication process on the destination volume from the previous backup is still in progress. In this case, the deduplication process continues to run, but the archival Snapshot copy does not get replaced after deduplication has finished running.
- When using SnapVault with deduplication, the maximum volume sizes for deduplication for the primary and secondary are independent of one another. Volumes on each of the systems must abide by their respective maximum volume size limits.
- Use Protection Manager 3.8 or later to manage deduplication with SnapVault for optimal performance.

For additional information about SnapVault, refer to [TR-3487: SnapVault Design and Implementation Guide](#).

For additional information about Protection Manager, refer to [TR-3710: Protection Manager Best Practice Guide](#).

6.2.6 OPEN SYSTEMS SNAPVAULT (OSSV)

Because SnapVault transfers are logical, any deduplication savings on disk are lost as a result of the data transfer to the destination. To reduce the bandwidth consumption of SnapVault data transfers, it is possible to transfer the data in a compressed state using OSSV, referred to as link compression. Link compression will compress the data before sending it across the network, and will decompress the data once the transfer is complete and received at the destination system. For additional information about link compression, refer to TR-3466: [Open Systems SnapVault Best Practice Guide](#).

Deduplication is supported with OSSV. Deduplication space savings can be achieved on the OSSV destination system by enabling deduplication. As an example, if your destination volume has deduplication enabled and utilizes OSSV, it would work as follows:

- Source volume contains data. The data transfer request will result in the data being sent to OSSV for transfer from the host system.
- Prior to sending the data over the network, OSSV will use its built-in link compression to compress the data.
- OSSV compresses the data and transfers compressed data over the wire.
- When the data is received on the OSSV destination system, it is decompressed in preparation to be written to disk.
- Deduplication will run on the destination after the data transfer completes.

Consider these additional points when implementing OSSV with deduplication:

- In the case of OSSV, if a single block is changed in a file, the full file is transferred to the destination. This is different from SnapVault, where only the changed block would be transferred. As a result of this difference, OSSV will typically see greater space savings from deduplication than SnapVault.
- OSSV can run in either a changed block or full file transfer mode as documented in [TR-3466: Open Systems SnapVault Best Practices Guide](#). Running OSSV in changed block mode will provide savings similar to SV.
- Deduplication will achieve space savings on the destination for the baseline transfer, and each incremental will yield deduplication space savings, which will result in the total savings on the destination volume remaining somewhat consistent. For example, in the case of a home directory if you see a space savings of 30% for the baseline transfer, then the space savings following additional transfers would typically remain at about 30% over time.

6.2.7 SNAPMIRROR SYNC

SnapMirror Sync mode is not supported with deduplication.

6.2.8 SNAPLOCK

Deduplication is fully supported with SnapLock® beginning with Data ONTAP 7.3.1, including both enterprise and compliance modes.

For Data ONTAP 8.0.X, deduplication is not supported with SnapLock 8.0 or 8.0.1.

When implementing SnapLock and NetApp deduplication for FAS, consider the following points:

- A SnapLock volume with files committed to “write once, read many” (WORM) can be deduplicated. Capacity savings are similar to savings in which the files are not committed to WORM. Both deduplication and subsequent undeduplication do not result in any changes to the SnapLock attributes or the WORM behavior of the volume or the file.
- Deduplication is applied across WORM, WORM append, and non-WORM (normal) files.
- Volume restore from a Snapshot copy is permitted only on SnapLock enterprise volumes. When a volume restore occurs from a Snapshot copy with deduplicated data, the file system returns to the state at which the Snapshot copy was created, including the state of deduplication, and to the WORM status of the volume and the files.
- File folding continues to function, regardless of the WORM and deduplication status of the files.
- When locking Snapshot copies on a SnapVault secondary, a Snapshot copy is permanent. This means that it can be deleted only after a retention period. No archive Snapshot copy is created on the secondary until deduplication is complete. If deduplication is still running when the next transfer attempts to begin, then the next transfer is delayed. Therefore, deduplication on the SnapVault secondary can result in the disruption of the transfer schedule on the primary. Avoiding the mismatched schedule allows optimal capacity savings to be recognized.
- Autocommit functions regardless of the deduplication status of the files.
- When using qtree SnapMirror, deduplication must be run separately on the source and destination. The WORM property is carried forward by qtree SnapMirror. Switching on WORM or deduplication on either end has no effect on the qtree SnapMirror transfers. Undoing deduplication also has no effect when done on either the source or the destination.
- When using volume SnapMirror, the WORM property of the files is carried forward by volume SnapMirror. Deduplication needs to be run only on the primary. Volume SnapMirror allows the secondary to inherit the deduplication. Undoing deduplication can be done only after breaking the volume SnapMirror relationship.
- To revert to a previous release on a system hosting a volume that is deduplicated and has WORM data on it, deduplication must first be undone (undeduplicated).
- If you are reverting to a previous release that does not support deduplication with SnapLock volumes, prior to Data ONTAP 7.3.1, you must first run the `sis undo` command. If this command is not run prior to the revert operation, an error message is displayed stating that `sis undo` must be performed.

For additional details about SnapLock, refer to TR-3263: [WORM Storage on Magnetic Disks Using SnapLock Compliance and SnapLock Enterprise](#).

6.3 CLUSTERING TECHNOLOGIES

6.3.1 DATA ONTAP CLUSTER-MODE

Deduplication is not supported in Data ONTAP 8.0.X Cluster-Mode.

6.3.2 ACTIVE-ACTIVE CLUSTER CONFIGURATION

Active-active configurations are supported with deduplication. It is recommended that both nodes in the cluster run the same version of Data ONTAP and have the same deduplication licenses installed.

Deduplication works on each node independently.

Writes to the flexible volume have fingerprints written to the change log. Upon failover or giveback to the partner node, deduplication change logging continues to happen. Deduplication operations will stop on the failed node during failover and resume after giveback, using the updated change log. The resumed deduplication processes start at the times scheduled for each volume, or they can be started manually. A maximum of eight concurrent deduplication operations are allowed on each node on an active-active configuration.

Starting with Data ONTAP 7.3, the following commands are supported for partner volumes in takeover mode:

```
sis status, sis stat, sis on, sis off
```

Starting with Data ONTAP 7.3, for SnapVault with Symantec™ NetBackup™, block sharing is supported for partner volumes in takeover mode.

NetApp recommends there are no active deduplication processes during the planned takeover or planned giveback:

1. Perform the planned takeover or giveback during a time when deduplication processes are not scheduled to run.
2. Determine whether any deduplication processes are active and stop them until the planned takeover or giveback is complete.

You can use the `sis status` command to determine whether the status of deduplication is Active or Idle. If a deduplication process is running, then the status of deduplication will be Active.

On a system with deduplication enabled, the output of the `sis status` command is similar to the following:

Path	State	Status	Progress
/vol/v460	Enabled	Idle	Idle for 00:12:30
/vol/v461	Enabled	Active	521MB Scanned
/vol/v462	Enabled	Active	489MB Scanned

You can use the `sis stop` command to abort the active SIS operation on the volume and the `sis start` command to restart it.

For information about deduplication, see the "[Data ONTAP Storage Management Guide](#)" and the `sis(1)` man page.

Deduplication is a licensed option. NetApp recommends having both nodes in an active-active controller configuration licensed the same for deduplication.

Deduplication does not add any overhead in an active-active configuration other than additional disk I/O. For additional information about active-active configurations, refer to TR-3450: [Active/Active Controller Configuration Overview and Best Practice Guidelines](#).

6.3.3 METROCLUSTER

Deduplication is supported in both fabric and stretch MetroCluster™ starting in Data ONTAP 7.2.5.1 and 7.3.1 and later.

When using MetroCluster with deduplication, consider the following points:

- Deduplication has an impact on CPU resources as a result of extra disk write operations. The increase is due to writing to two plexes. On most platforms the impact is less than 10%. This impact is experienced on low-end systems (for example, FAS3000 systems) more than on high-end systems (for example, FAS6000 systems).
- In takeover mode, writes to partner flexible volumes are change logged. The deduplication process does not run on the partner flexible volumes while in takeover mode. Upon giveback, data in the change logs is processed and data gets deduplicated.
- In takeover mode, change logging continues until the change log is full. This can occur if the node remains in takeover mode for a long period of time, such as during a disaster. All data continues to be accessible regardless of change log availability.
- A node in takeover mode takes over the servicing of I/Os targeted at the partner volumes, as well as its change logging. As a result, additional system resources are consumed, which might require that the system workload be adjusted.
- Only a subset of deduplication commands for the partner volumes is available in takeover mode. These commands are `sis status`, `sis stat`, `sis on`, `sis off`.
- Deduplication must be licensed on both nodes.

For additional information about MetroCluster, refer to TR-3548: [MetroCluster Design and Implementation Guide](#).

6.4 OTHER NETAPP FEATURES

6.4.1 QUOTAS

For deduplicated files, the logical (undeduplicated) size is charged against the quotas. There are several advantages to this scheme as opposed to charging quotas based on the physical (deduplicated) size of the file:

- This is in line with the general design principle of making deduplication transparent to the end user.
- It is easier for system administrators to manage quotas. They can maintain a single quota policy across all volumes, whether or not deduplication is enabled on it.
- There are no out-of-space failures when data is being sent using SnapMirror from a volume with deduplication enabled to a destination volume that has deduplication disabled.

6.4.2 FLEXCLONE VOLUMES

FlexClone technology instantly replicates data volumes and datasets as transparent, virtual copies—without requiring additional storage space. FlexClone volumes support deduplication. When a FlexClone volume (cloned volume) is created:

- If the parent FlexClone volume has duplication enabled, the new volume inherits the savings.
- The cloned volume inherits the deduplication configuration of the parent volume, such as the deduplication schedule.

Starting with Data ONTAP 7.3, the deduplication metadata files (the fingerprint database and the change log files) do not get cloned, because they are located outside the volume in the aggregate. In this case, there is no fingerprint database file in the cloned volume for the data that came from the parent. However, the data in the cloned volume inherits the space savings of the original data. The deduplication process also continues for any new data written to the clone and creates the fingerprint database for the new data. However, the deduplication process obtains space savings in the new data only and does not deduplicate between the new data and the old data. To run deduplication for all the data in the cloned volume (and thus obtain higher space savings), use the `sis start -s` command. Depending on the size of the logical data in the volume, this process can take a long time to complete.

SPACE SAVINGS TESTING WITH FLEXCLONE VOLUMES

You can use a FlexClone volume to determine the possible deduplication savings on a volume. Using FlexClone volumes to test deduplication allows you to run tests with minimal storage overhead and without affecting your production data.

To create these test environments you would first create a FlexClone volume of the parent volume. This will take up negligible space on the aggregate.

Scenario 1: Calculating Deduplication Savings on a Normal Volume

To determine the deduplication savings you would run `sis start -s` on the clone. This will deduplicate the data within the clone volume only, and not modify the data on the parent volume. To determine the savings achieved by deduplication, see the `%saved` column in the output from running the `df -s` command on the clone to determine the amount of savings, as you would for a volume not created by FlexClone.

CREATING A TEST ENVIRONMENT WITH FLEXCLONE VOLUMES

You can use a FlexClone volume to quickly setup a test environment to test deduplication on a volume without modifying the production volumes.

To create these test environments with minimal storage space and without modifying your production data you would:

Scenario 1: Create Test Environment Adding Deduplication to a Normal Volume

To create a test volume that adds deduplication while minimizing the space usage you would create a FlexClone volume of a parent volume and then run `sis start -s` on the clone. This will deduplicate the data within the clone volume only, and not modify the data on the parent volume. From here you can run your tests against the deduplicated clone.

VOLUME SPLITTING

When a cloned volume is split from the parent volume, all of the deduplicated data in the clone that was part of the parent volume (that is, not including the data that was written to the cloned volume after the clone was created) gets undeduplicated after the volume split operation. If deduplication is running on the cloned volume, this data gets deduplicated again in subsequent deduplication operations on the volume.

6.4.3 FLEXCLONE FILES

Beginning with Data ONTAP 7.3.1, in addition to FlexClone volumes, FlexClone files and LUNs are available and are allowed on deduplicated volumes.

Deduplication can be used to regain capacity savings on data that was copied using FlexClone at the file or LUN level and that has been logically migrated (that is, with qtree SnapMirror, SnapVault, NDMPdump, and so on).

Creating FlexClone files or LUNs by using the `-l` option enables change logging, which can result in optimal deduplication performance and space savings. However, there can be a significant trade-off in cloning performance due to the additional overhead of change logging.

6.4.4 64-BIT AGGREGATE SUPPORT

Beginning with Data ONTAP 8.0, deduplication supports both 64-bit and 32-bit aggregates. Deduplication is supported with both 32-bit and 64-bit aggregates.

If data is migrated from a volume in a 32-bit aggregate to a volume in a 64-bit aggregate using a logical migration method, such as NDMPcopy, NetApp SnapVault, or qtree SnapMirror, the data is undeduplicated during the process. Once the transfer is complete you can run deduplication again to regain the deduplication savings.

If data is migrated between 64-bit aggregates using volume SnapMirror, the deduplication space savings will be inherited at the destination, as when migrated between two 32-bit aggregates.

6.4.5 FLEXCACHE

When using FlexCache, data will not be stored on the caching systems as deduplicated.

6.4.6 NONDISRUPTIVE VOLUME MOVEMENT

Nondisruptive volume movement is supported with deduplication; however, the deduplication metadata will not be transferred. This means that although the deduplication savings will still exist for the transferred data after the transfer completes, new deduplication operations will not deduplicate against the data that was transferred. To realize deduplication savings across transferred and new data, NetApp recommends that you run `sis start -s` to recreate the fingerprint database after the volume movement completes.

6.4.7 NETAPP DATA MOTION

Data Motion is supported with deduplication. However, the combination of having to manage resources associated with nondisruptive migration and metadata for deduplication could result in a small possibility of degradation to the performance of client applications. This might happen only if Data Motion is performed under high load (greater than 60% CPU). Therefore, Data Motion should not be attempted under such conditions.

NetApp best practice recommendation is that NetApp Data Motion be performed during off peak hours or periods of lower load to enable the fastest migration times and hence minimal impact. In this initial release, it is recommended that customers actively monitor the system performance during Data Motion cutover operation for systems that have deduplication enabled (on either the source or destination system). If excessive load is encountered during the migration, NetApp Data Motion can be aborted by the storage administrator, while maintaining access to data.

For details of the recommended monitoring during the NetApp Data Motion cutover period for systems running deduplication, please consult the following NOW™ article:

<https://now.netapp.com/Knowledgebase/solutionarea.asp?id=kb56838>.

After a successful Data Motion cutover, the deduplicated volume remains deduplicated on the destination array, but before another deduplication process can run on the volume, the deduplication fingerprint database must be rebuilt. The Data Motion software automatically starts rebuilding the fingerprint database after a successful migration.

For more information on Data Motion, refer to [TR-3814: Data Motion Best Practices](#).

6.4.8 PERFORMANCE ACCELERATION MODULE AND FLASH CACHE CARDS

Beginning with Data ONTAP 7.3, NetApp has had Performance Acceleration Module (PAM) and Flash Cache (PAM II) cards. In environments in which there are shared blocks that are read repeatedly, Flash Cache cards can help reduce the number of disk reads, thus improving read performance.

The amount of performance improvement with the Flash Cache card depends on the duplication rate, the access rate, the active dataset size, and the data layout. Adding a Flash Cache card to a system does not increase the deduplication maximum volume size for that system.

The Flash Cache card has provided significant performance improvements in VMware VDI environments. The advantages provided by NetApp Flash Cache cards are further enhanced when combined with other shared block technologies, such as NetApp deduplication or FlexClone.

For additional information about the PAM card, refer to TR-3705: [NetApp and VMware VDI Best Practices](#).

6.4.9 SMTAPE

Deduplication supports SMTape backups. Backup to a tape through the SMTape engine preserves the data format of the source volume on the tape; therefore, the savings are maintained. Because the data is backed up as deduplicated it can only be restored to a volume that supports deduplication.

6.4.10 DUMP

Deduplication supports backup to a tape using NDMP. The data sent from the source volume is undeduplicated and then written to the tape in undeduplicated format.

6.4.11 NONDISRUPTIVE UPGRADES

Major and minor nondisruptive upgrades are supported with deduplication.

NetApp recommends that no active deduplication operations be running during the nondisruptive upgrade:

- Perform the Data ONTAP upgrade during a time when deduplication processes are not scheduled to run.
- Determine whether any deduplication processes are active and halt them until the Data ONTAP upgrade is complete.

You can use the `sis status` command to determine whether the status of a deduplication is Active or Idle. On a system with deduplication enabled, the output of the `sis status` command is similar to the following:

Path	State	Status	Progress
/vol/v460	Enabled	Idle	Idle for 00:12:30
/vol/v461	Enabled	Active	521 MB Scanned
/vol/v462	Enabled	Active	489 MB Scanned

- You can use the `sis stop` command to abort an active `sis` process on the volume and the `sis start` command to restart it.

For specific details for doing a nondisruptive upgrade on your system refer to Upgrade Advisor in [AutoSupport](#)[™] if you have AutoSupport enabled; otherwise, refer to the [release notes](#) for the version of Data ONTAP to which you are upgrading.

6.4.12 NETAPP DATAFORT ENCRYPTION

Encryption removes duplicate data, so no savings are expected from deduplication. Because encryption can be run at the share level, it is possible to create a flexible volume where only part of the data on the volume is encrypted. If deduplication is run on such a volume, negligible capacity savings are expected on the encrypted data, but it is still possible to achieve savings on the rest of the volume effectively.

6.4.13 READ REALLOCATION (REALLOC)

Because read reallocation does not predictably improve the file layout and the sequential read performance when used on deduplicated volumes, performing read reallocation on deduplicated volumes is not supported.

For workloads that perform mostly random writes and long sequential reads on volumes not using deduplication, read reallocation could be used to improve the file layout and the sequential read performance.

When you enable read reallocation, Data ONTAP analyzes the parts of the file that are read sequentially. If the associated blocks are not already largely contiguous, Data ONTAP updates the file layout by rewriting those blocks to another location on disk. The rewrite improves the file layout, thus improving the sequential read performance the next time that section of the file is read. However, read reallocation might result in more storage use if Snapshot copies are used. It might also result in a higher load on the storage system. If you want to enable read reallocation but storage space is a concern, you can enable read reallocation on FlexVol volumes using the `space_optimized` option. This option conserves space but can slow read performance through the Snapshot copies. Therefore, if fast read performance through Snapshot copies is a high priority, do not use `space_optimized`.

A read reallocation scan does not rearrange deduplicated blocks. Instead, for files to benefit from read reallocation, they should be stored on volumes that are not enabled for deduplication.

6.4.14 VOL COPY COMMAND

When deduplicated data is copied by using the volume copy command, the copy of the data at the destination location inherits all of the deduplication attributes and storage savings of the original data.

Reads to the deduplicated data on the destination are always allowed even if deduplication is not licensed on the destination system.

Starting with Data ONTAP 7.3, some of the deduplication metadata files do not get copied by the `vol copy` command, because they are located outside the volume in the aggregate. In this case, there is no fingerprint database file in the destination volume for the data. However, the data retains the space savings. The deduplication process also continues for any new data written to the destination volume and creates the fingerprint database for the new data. The deduplication process obtains space savings in the new data only and does not deduplicate between the new data and the old data. To run deduplication for all the data in the destination volume (and thus obtain higher space savings), use the `sis start -s`

command. Depending on the size of the logical data in the volume, this process can take a long time to complete.

6.4.15 AGGREGATE COPY COMMAND

When deduplicated or compressed data is copied by using the `aggrcopy` command, the copy of the data at the destination location inherits all of the deduplication attributes and storage savings of the original data. Aggregate copy is also the only way to copy data and maintain the layout of data on disk and enable the same performance with deduplicated data.

6.4.16 MULTISTORE (VFILER)

Starting with Data ONTAP 7.3, deduplication is supported with MultiStore[®]. In Data ONTAP 7.3, the deduplication commands are available only in the CLI of vFiler0; however, they allow any volume to be included in the command arguments, regardless of which vFiler[®] unit the volume is associated with. Beginning with Data ONTAP 7.3.1, the deduplication commands are available in the CLI of each vFiler unit, allowing each vFiler unit to be configured from within itself. When configuring deduplication for MultiStore with Data ONTAP 7.3.1 or later, consider the following:

- NetApp deduplication is supported at the volume level only; it is not supported at the qtree level.
- A vFiler unit can enable deduplication on a volume only if it owns the volume.
- If deduplication is run on a volume, then all qtrees within that volume are deduplicated.
- Standard vFiler unit behavior states that if a vFiler unit owns a volume, then any qtree within that volume cannot be owned by any other vFiler unit. With that in mind, the following is true:
 - If deduplication is enabled on a volume, then all qtrees within the volume must be owned by the vFiler unit that owns that volume.
 - If a vFiler unit owns a volume where deduplication is enabled, then any qtree within that volume cannot be owned by any other vFiler unit.
- vFiler0 is a special case, because it is a master vFiler unit:
 - vFiler0, as a master vFiler unit, has access to all resources on the system, and is thus allowed to own a volume that contains qtrees that are owned by other vFiler units.
 - By default, if any storage is not owned by a vFiler unit within the hosting system, then it is owned by vFiler0.
 - vFiler0 is able to run deduplication on a volume that is not owned by a vFiler unit on the system.
 - As a master vFiler unit, vFiler0 can run deduplication commands on a volume inside any vFiler unit in the system.

6.4.17 SNAPDRIVE

Beginning with Data 7.3.4, NetApp has introduced new functionality in SnapDrive[®] for Windows. This might reduce the savings from deduplication; however, the overall system savings will still benefit from the effects of SnapDrive capacity reduction. The effects on the overall space savings is not significant, and will usually not warrant any special actions. If SnapDrive is being used, and there is less deduplication space savings than expected, it is recommended that NetApp support be contacted for further analysis.

6.4.18 LUNS

When using deduplication on a file-based (NFS/CIFS) environment, deduplication is straightforward. As duplicate blocks are freed, they are marked as available. In both cases, as blocks of free space become available, the NetApp system recognizes these free blocks and makes them available to the volume.

Deduplication on a block-based (FCP/iSCSI) LUN environment are slightly more complicated. This is because of the space guarantees and fractional reservations often used by LUNs. With space guarantees, for instance, a 500GB LUN consumes exactly 500GB of physical disk space. If the data in the LUN is reduced through deduplication, the LUN still reserves the same physical space capacity of 500GB, and the space savings are not apparent to the user.

LUN space guarantees and fractional reserves can be configured to control how the NetApp system uses freed blocks. By varying the values of certain parameters, freed blocks can be returned to the LUN overwrite reserve, the volume free pool, the aggregate free pool, or a combination.

Note: Although deduplication will allow volumes to effectively store more than 16TB worth of data, the system should not be configured with more than 16TB worth of guaranteed LUN space in a single volume.

This section describes five common examples of LUN configurations and deduplication behavior, as summarized in Table 11.

Table 11) Summary of LUN configuration examples (as described in the text).

	A (Default)	B	C	D	E
LUN space guarantee value	Yes	Yes	Yes	No	No
Volume fractional reserve value	100	1–99	0	Any	Any
Volume thin provisioned?	No	No	No	No	Yes
After deduplication	Fractional overwrite reserve	Fractional overwrite reserve	Volume free pool	Volume free pool	Aggregate free pool

DEFINITIONS

Fractional overwrite reserve: The space that Data ONTAP reserves will be available for overwriting blocks in a LUN when the space guarantee = Yes. Behavior of the fractional reserve space parameter with deduplication is the same as if a Snapshot copy was created in the volume and blocks are being overwritten.

Volume free pool: Refers to the free blocks in the parent volume of the LUN. These blocks can be assigned anywhere in the volume as needed.

Aggregate free pool: Refers to the free blocks in the parent aggregate of the LUN. These blocks can be assigned anywhere in the aggregate as needed.

LUN CONFIGURATION EXAMPLES

Configuration A: Default LUN Configuration

The default configuration of a NetApp LUN follows. (The best practice for all NetApp LUNs is to turn the controller Snapshot copy off, delete all scheduled Snapshot copies, and set the snap reserve to 0.)

1. LUN space reservation value = on Default = on
2. Volume fractional reserve value = 100 Default = 100%
3. Volume guarantee = volume Default = volume

- | | |
|----------------------------|-----------------------|
| 4. Snap reserve = 0% | Default = 20% |
| 5. Autodelete = off | Default = off |
| 6. Autosize = off | Default = off |
| 7. Try_first = volume_grow | Default = volume_grow |

Description: When a LUN containing default values is deduplicated, no apparent savings are observed by the storage administrator because the LUN by default was space reserved when it was created and the fractional reserve was set to 100% in the volume. Any blocks freed through deduplication are allocated to the fractional reserve area. This configuration means that overwrite to the LUN should never fail, even if it is overwritten entirely.

Pros and cons: The advantage of this configuration is that Snapshot copies consume less space when blocks in the active file system are no longer being used. As a result, this volume can hold more Snapshot copies. The disadvantage of this configuration is that free blocks are not returned to either the free volume pool or the free aggregate pool. Moreover, there is no direct space saving in the active file system. In fact, this configuration could consume more space in the volume due to new indirect blocks being written if no Snapshot copies exist in the volume and the Snapshot schedule is turned off.

Note: If Snapshot copies are turned off for the volume (or no copy exists in the volume), this is *not* a recommended configuration or volume for deduplication.

Configuration B: LUN Configuration for Shared Volume Space Savings

If the user wants to apply the freed blocks to both the fractional overwrite reserve area and the volume free pool, this can be accomplished with the following configuration:

1. LUN space reservation value = on
2. Volume fractional reserve value = any value from 1–99
3. Volume guarantee = volume
4. Snap reserve = 0%
5. Autodelete = off
6. Autosize = off
7. Try_first = volume_grow

Description: The only difference between this configuration and configuration A is that the amount of space reserved for overwrite is based on the fractional reserve value set for the volume. As a result, this configuration splits the free blocks between the fractional overwrite reserve and the volume free space. For instance, if the fractional reserve value is set to 25, 25% of the freed blocks go into the fractional overwrite reserve and 75% of the freed blocks are returned to the volume free pool.

Pros and cons: The advantage of this configuration is that the overwrite space reserve does not increase for every block being deduplicated. Freed blocks are split between the volume free pool and the fractional reserve. The disadvantage of this configuration is that overwrites to the LUN beyond the fractional reserve capacity might fail because freed blocks might have already been allocated. Another disadvantage of this configuration is that freed blocks stay in the parent volume and cannot be provisioned to any other volumes in the aggregate.

Note: If Snapshot copies are turned off for the volume (or if no Snapshot copy exists in the volume) and the percentage of savings due to deduplication is less than the fractional reserve, then NetApp does *not* recommend this configuration or volume for deduplication.

Configuration C: LUN Configuration for Maximum Volume Space Savings

If the user wants to apply the freed blocks to the volume free pool, this can be accomplished with the following configuration:

1. LUN space reservation value = on
2. Volume fractional reserve value = 0
3. Volume guarantee = volume
4. Snap reserve = 0%
5. Autodelete = off
6. Autosize = off
7. Try_first = volume_grow

Description: The only difference between this configuration and configuration B is that the value of the fractional reserve is set to zero. As a result, this configuration "forces" all the freed blocks to the volume free pool and no blocks are set aside for the fractional reserve.

Pros and cons: The advantage of this configuration is that all the freed blocks are returned to the volume free pool. The disadvantage is that the chance of overwrite failure is higher than with configurations A and B because no freed blocks are assigned to the fractional overwrite area.

Configuration D: LUN Configuration for Maximum Volume Space Savings

If the user wants to apply the freed blocks to the volume free pool, this can be accomplished with the following configuration:

1. LUN space reservation value = off
2. Volume fractional reserve value = any value from 0–100
3. Volume guarantee = volume
4. Snap reserve = 0%
5. Autodelete = off
6. Autosize = off
7. Try_first = volume_grow

Description: The difference between this configuration and configuration C is that the LUN is not space reserved. With LUN space guarantees off, the value for the volume fractional reserve is ignored for all LUNs in this volume. From a deduplication perspective, there is no difference between this and the previous configuration, and all freed blocks go to the volume free pool.

Pros and cons: From a deduplication perspective, this configuration has the same advantages and disadvantages as configuration C.

Configuration E: LUN Configuration for Maximum Aggregate Space Savings

In many cases, the user might prefer to reclaim all freed blocks from the volume and return these blocks to the aggregate free pool. This is accomplished with the following configuration:

1. LUN space reservation value = off
2. Volume fractional reserve value = any value from 0–100
3. Volume guarantee = none

4. Snap reserve = 0%
5. Autodelete = on
6. Autosize = on
7. Try_first = volume_grow

Description: This configuration "forces" the free blocks out of the volume and into the aggregate free pool, where the blocks can be reprovisioned for any other volumes in the aggregate.

Pros and cons: The advantage of this configuration is that it provides the highest efficiency in aggregate space provisioning. It also uses the thin provisioning features of Data ONTAP, volume autosize, and Snapshot autodelete to help administer the space in the solution.

The disadvantage of this configuration is that it requires the storage administrator to monitor the free space available in the aggregates. With volume autosize and Snapshot autodelete turned on, the volume grows first if space is available in the aggregate; if not, then Snapshot copies are deleted.

7 DEDUPLICATION BEST PRACTICES WITH SPECIFIC APPLICATIONS

This section discusses how deduplication will interact with some popular applications. Best practice guides for various applications are available and should be referenced when building a solution.

This technical report describes some specific deduplication best practices to keep in mind for some common applications.

7.1 VMWARE BEST PRACTICES

VMware environments deduplicate extremely well. However, while working out the VMDK and datastore layouts, keep the following points in mind.

Operating system VMDKs deduplicate extremely well because the binary files, patches, and drivers are highly redundant between virtual machines (VMs). Maximum savings can be achieved by keeping these in the same volume.

Application binary VMDKs deduplicate to varying degrees. Duplicate applications deduplicate very well; applications from the same vendor commonly have similar libraries installed and deduplicate somewhat successfully; and applications written by different vendors don't deduplicate at all.

When deduplicated, application datasets have varying levels of space savings and performance impact based on application and intended use. Careful consideration is needed, just as with nonvirtualized environments, before deciding to keep the application data in a deduplicated volume.

Transient and temporary data such as VM swap files, pagefiles, and user and system temp directories do not deduplicate well and potentially add significant performance pressure when deduplicated. Therefore NetApp recommends keeping this data on a separate VMDK and volume that are not deduplicated. For more information on page files, refer to <http://media.netapp.com/documents/tr-3749.pdf>.

NetApp Data ONTAP 7.2.6 and 7.3.1 introduce a performance enhancement referred to as *intelligent cache*. Although it is applicable to many different environments, intelligent caching is particularly applicable to VM environments, where multiple blocks are set to zero as a result of system initialization. These zero blocks are all recognized as duplicates and are deduplicated very efficiently. The warm cache extension enhancement provides increased sequential read performance for such environments, where there are very large amounts of deduplicated blocks. Examples of sequential read applications that

benefit from this performance enhancement include NDMP, NetApp SnapVault, some NFS-based applications, and dump. This performance enhancement is also beneficial to the boot-up processes in VDI environments.

The expectation is that about 30% space savings will be achieved overall. This is a conservative figure, and in some cases users have achieved savings of up to 80%. The major factor that affects this percentage is the amount of application data. New installations typically deduplicate extremely well, because they do not contain a significant amount of application data.

Important: In VMware environments, the need for proper partitioning and alignment of the VMDKs is extremely important (not just for deduplication). VMware must be configured so that the VMDKs are aligned on NetApp WAFL (Write Anywhere File Layout) 4K block boundaries as part of a standard VMware implementation. To learn how to prevent the negative performance impact of LUN/VMDK misalignment, read [TR-3747: Best Practices for File System Alignment in Virtual Environments](#), [TR-3428: NetApp and VMware Best Practices Guide](#), or [TR-3749: NetApp and VMware vSphere Storage Best Practices](#). Also note that the applications in which performance is heavily affected by deduplication (when these applications are run *without* VMware) are likely to suffer the same performance impact from deduplication when they are run *with* VMware.

A deduplication and VMware solution on NFS is easy and straightforward. Combining deduplication and VMware with LUNs requires a bit more work.

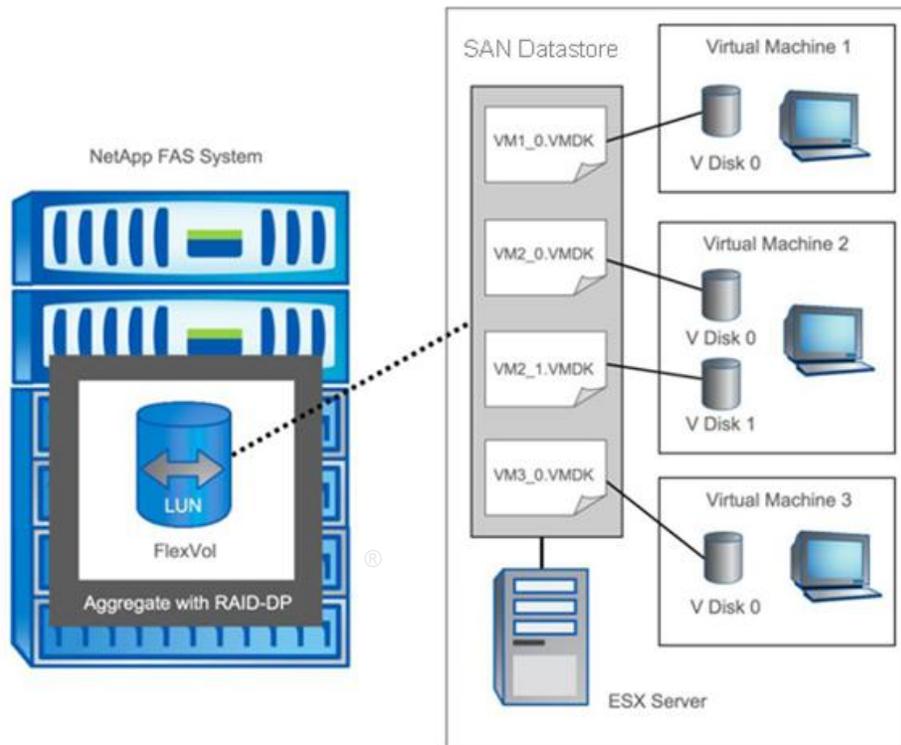
The following subsections describe the different ways that VMware can be configured. The concepts covered in this document are general in nature. Specific configuration settings provided in a separate NetApp solution-specific document should be given precedence over this document for that specific solution.

For more information about NetApp storage in a VMware environment, see [TR-3428: NetApp and VMware Virtual Infrastructure 3 Storage Best Practices](#) or [TR-3749: NetApp and VMware vSphere Storage Best Practices](#).

7.1.1 VMFS DATASTORE ON FIBRE CHANNEL OR ISCSI: SINGLE LUN

This is the default configuration, and it's the way that a large number of VMware installations are done today. Deduplication occurs across the numerous VMDKs.

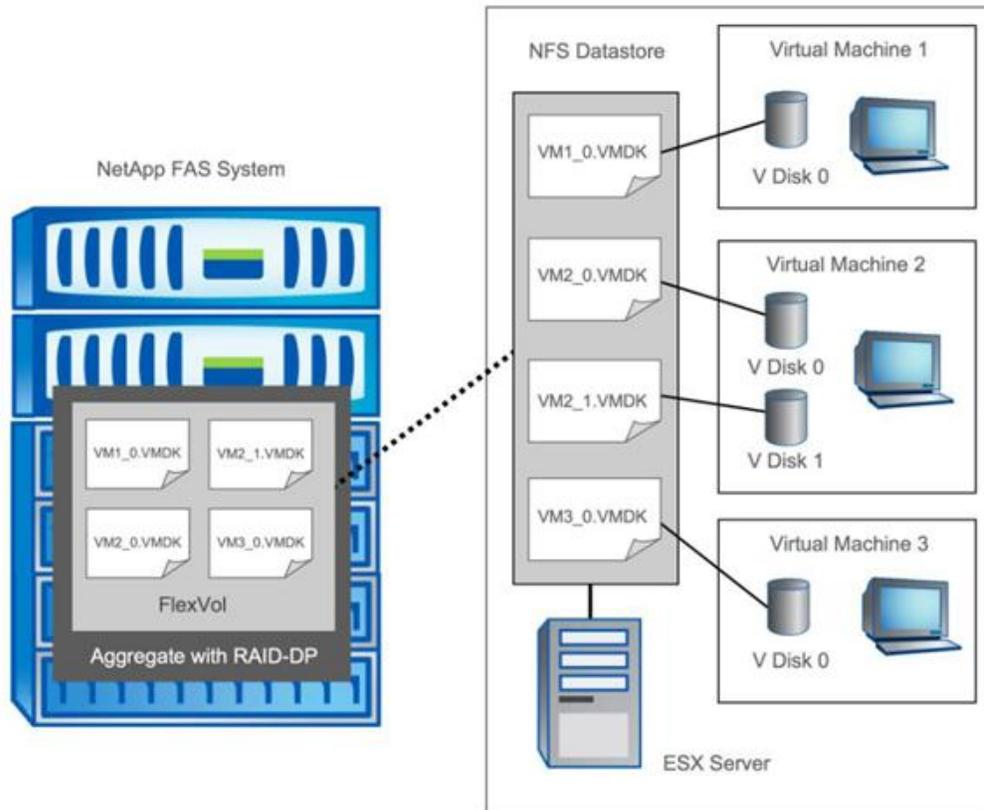
Figure 3) VMFS datastore on Fibre Channel or iSCSI: single LUN3.



7.1.2 VMWARE VIRTUAL DISKS OVER NFS/CIFS

This is a new configuration that became available starting with VMware 3.0. It is the easiest to configure and allows deduplication to provide the most space savings.

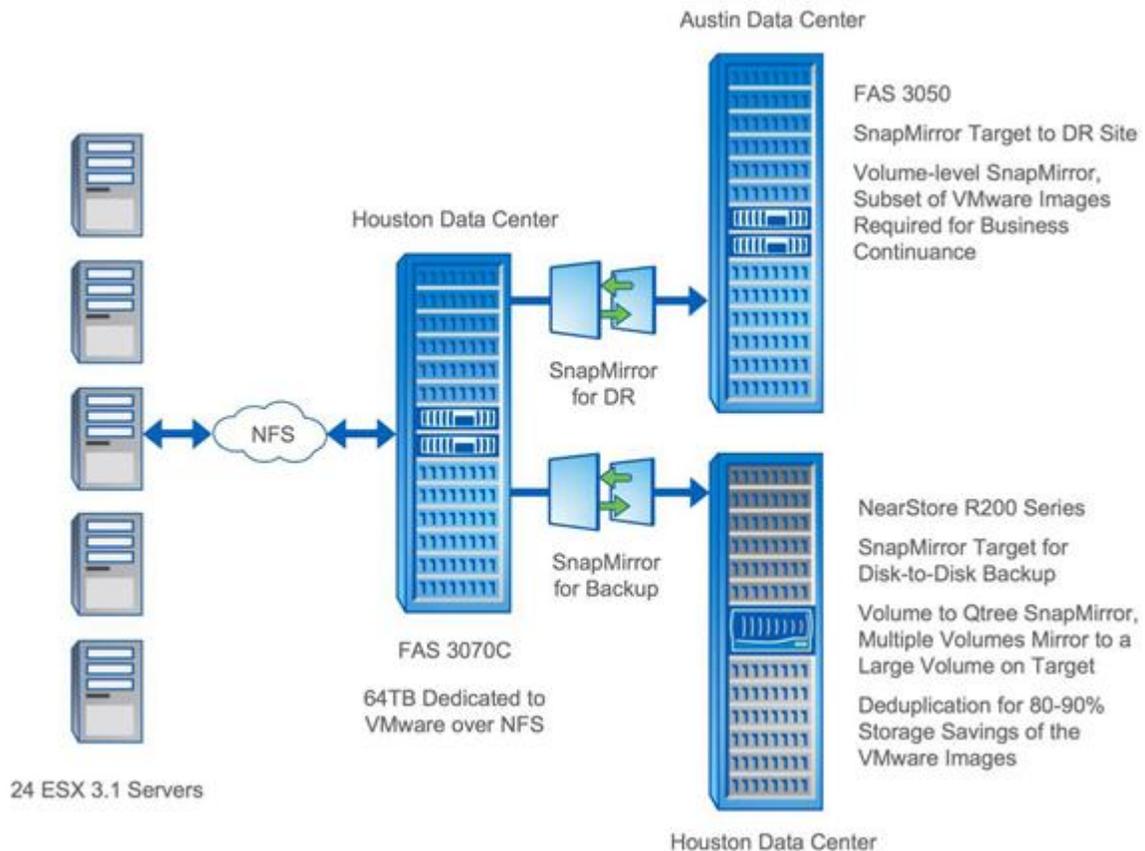
Figure 4) VMware virtual disks over NFS/CIFS.



7.1.3 DEDUPLICATION ARCHIVE OF VMWARE

Deduplication has proven very useful in VMware archive environments, as shown in Figure 5.

Figure 5) Archive of VMware with deduplication.



Specifications for the example shown in Figure 5:

- In this environment, VMware is configured to run over NFS.
- This environment uses approximately 1,800 clone copies of the master VMware image. These images are used to create virtual machines for primary applications and for test and development purposes.
- All 1,800 clone copies (~32TB) are stored on a FAS3070 in the Houston data center.
- The data is mirrored to the remote site in Austin for disaster recovery.
- Once per hour, the FAS3070 images are transferred to a NetApp NearStore R200 by using NetApp SnapMirror.
- Deduplication is run nightly on the R200, and the VMware images are reduced in size by 80% to 90%.

7.2 MICROSOFT SHAREPOINT BEST PRACTICES

If Microsoft SharePoint® and NetApp deduplication for FAS will be used together, consider the following points:

- Make sure that space is available in the volume before using the `sis on` command. If this command is used on a flexible volume that already has data and is completely full, it fails. Up to 6% of the total data size is needed for deduplication of metadata files.
- Deduplication is transparent to SharePoint. The block-level changes are not recognized by SharePoint, so the SharePoint database remains unchanged in size, even though there are capacity savings at the volume level.

7.3 MICROSOFT SQL SERVER BEST PRACTICES

Deduplication can provide space savings in Microsoft SQL Server environments, but proper testing should be done to determine the savings for your environment. The Space Savings Estimation Tool (SSET 3.0) can be used to estimate the amount of savings that would be achieved with deduplication.

Our testing has shown deduplication savings in the range of 15%. It is recommended to run the SSET tool against your data to better estimate the dedupe savings that your environment can achieve.

A Microsoft SQL Server database will use 8KB block sizes. Each SQL Server block is composed of two NetApp 4K blocks. Recall that deduplication works at the 4K block level. So, although Microsoft SQL Server will place a unique header at the beginning of each SQL Server block (corresponding to one 4K block), the rest of the blocks might still contain duplicates. This means that deduplication might provide significant savings when comparing the 4KB blocks within the volume.

7.4 MICROSOFT EXCHANGE BEST PRACTICES

If Microsoft Exchange and NetApp deduplication will be used together, consider the following points:

- In some Exchange environments, extents are enabled to improve the performance of database validation. Enabling extents does not rearrange blocks on disk that are shared between files by deduplication on deduplicated volumes. Enabling extents does not predictably optimize sequential data block layout when used on deduplicated volumes, so there is no reason to enable extents on deduplicated volumes.
- Beginning with Microsoft Exchange 2010, single instancing storage will no longer be available. NetApp deduplication for FAS and V-Series provides significant savings for primary storage running Exchange 2010. Similar savings can be expected with Microsoft Exchange 2007 if Exchange single instancing is disabled. NetApp recommends running the SSET tool on your environment to better estimate the deduplication savings that your environment can achieve.

For additional details about Exchange, refer to [TR-3578: Microsoft Exchange Server 2007 Best Practices Guide](#) or [TR-3824: Storage Efficiency and Best Practices for Microsoft Exchange Server 2010](#).

7.5 LOTUS DOMINO BEST PRACTICES

The deduplication space savings that you can expect will vary widely with the type (e-mail, applications, and so on) and frequency of data in your environment. In the NetApp test lab, running the space savings estimation tool (SSET) on a fresh install of Domino indicates a deduplication space savings of approximately 22%. This is before any actual data is managed. NetApp customers using Domino have reported anywhere from 25% to 60% deduplication savings in their Domino environment. It is recommended to run the SSET tool on your environment to better estimate the deduplication savings that could be achieved on your environment.

For additional Lotus Domino best practices, refer to [TR-3843: Storage Savings with Domino and NetApp Deduplication](#).

7.5.1 DOMINO ATTACHMENT AND OBJECT SERVICE (DAOS)

Domino 8.5 introduces a new feature called Domino Attachment and Object Service (DAOS). DAOS shares data identified as identical between multiple databases. Typically, in a Domino environment a separate and complete copy of each attached document is kept in each database that receives it. With DAOS the server saves a reference to each attached file in an internal repository and refers back to the same file from multiple documents and/or databases.

DAOS provides space savings. NetApp deduplication is still expected to provide space savings when DAOS is enabled, but it is anticipated that space savings will be reduced compared to an environment where DAOS is not used.

7.5.2 DOMINO ENCRYPTION

If Domino database encryption is enabled for all or the majority of databases, it should be anticipated that the deduplication space savings will be very small. This is because encrypted data is by its nature unique.

7.5.3 DOMINO QUOTAS

Domino quotas are not affected by deduplication. A mailbox with a limit of 1GB cannot store more than 1GB of data in a deduplicated volume even if the data consumes less than 1GB of physical space on the storage system.

7.5.4 DOMINO PERFORMANCE

The performance of deduplication varies for based on types of data, amounts of duplicate data, average file sizes, storage platforms, types of disks, and number of disk spindles in the aggregate.

Because of these factors NetApp recommends that the performance impact be carefully considered and taken into sizing considerations before deploying deduplication into production environments. NetApp recommends that you carefully test features such as deduplication in a test lab before introducing them into a production environment.

7.6 ORACLE BEST PRACTICES

Deduplication can provide savings in Oracle environments, but proper testing should be done to determine the savings for your environment. The Space Savings Estimation Tool (SSET 3.0) can be used to estimate the amount of savings that would be achieved with deduplication.

Deduplication savings are dependant upon the Oracle configurations.

A typical Oracle data warehouse or data mining database will typically use 16KB or 32KB block sizes, made of multiple NetApp 4K blocks. Although Oracle will place a unique identifier at the beginning of each Oracle block, and a near-unique identifier at the end of each Oracle block, the rest of the blocks within each Oracle block might still contain duplicates. This means that deduplication might provide significant savings when comparing the 4KB blocks within the volume.

Oracle OLTP databases typically use an 8KB block size. Oracle will once again place a unique identifier at the beginning and end of each Oracle block, making the 4K blocks unique. Testing has shown that these environments typically do not have a significant amount of duplicate blocks, and do not show very much deduplication savings.

One additional case to consider is if a table space is created or extended. In this case Oracle will initialize the blocks, and commit many of them in the same transaction. This will result in the creation of duplicate blocks, allowing deduplication to provide savings. However, as these blocks are filled with incoming data the space savings will shrink.

7.7 TIVOLI STORAGE MANAGER BEST PRACTICES

If IBM Tivoli Storage Manager (TSM) and NetApp deduplication will be used together, consider the following points:

- Deduplication savings with TSM will not be optimal because TSM does not block-align data when it writes files out to its volumes. The net result is that there are fewer duplicate blocks available to deduplicate.
- TSM compresses files backed up from clients to preserve bandwidth. Data that is compressed by using Tivoli storage manager compression does not usually yield much additional savings when deduplicated by the NetApp storage system.
- TSM client-based encryption results in data with no duplicates. Encrypted data does not usually yield good savings when deduplicated.
- TSM's progressive backup methodology backs up only new or changed files, which reduces the number of duplicates, since there are not multiple full backups to consider.

7.8 SYMANTEC BACKUP EXEC BEST PRACTICES

If Symantec Backup Exec™ and NetApp deduplication for FAS will be used together, consider the following point:

- Deduplication savings with Backup Exec will not be optimal because Backup Exec does not block-align data when it writes files out to its volumes. The net result is that there are fewer duplicate blocks available to deduplicate.

7.9 BACKUP BEST PRACTICES

There are various ways to backup your data. Here are a few deduplication considerations for your backup environment:

- To make sure of optimal backup throughput it is a best practice to:
 - Make sure deduplication operations initiate only after your backup completes.
 - Deduplication operations on the destination volume complete prior to initiating the next backup.
- If backing up data from multiple volumes to a single volume you might achieve significant space savings from deduplication beyond that of the deduplication savings from the source volumes. This is because you are able to run deduplication on the destination volume which could contain duplicate data from multiple source volumes.
- If you are backing up data from your backup disk to tape consider using SMTape to preserve the deduplication savings. Utilizing NDMP to tape will not preserve the deduplication savings on tape.

8 TROUBLESHOOTING

This section discusses basic troubleshooting methods and common pitfalls to be considered when working with deduplication.

8.1 DEDUPLICATION WILL NOT RUN

8.1.1 CHECK DEDUPLICATION LICENSING

Make sure that deduplication is properly licensed.

```
fas6070-ppe02> license
... a_sis <license key>
...
```

If the platform is not an R200 and you are using a version of Data ONTAP prior to 8.0, make sure that the NearStore option is also properly licensed:

```
fas3070-rtp01> license ... a_sis <license key> nearstore_option <license key> ...
```

If licensing is removed or has expired, no additional deduplication can occur, and no `sis` commands can run. However, the flexible volume remains a deduplicated volume, the existing storage savings are kept, and all data is usable.

Before removing the deduplication license, you must first disable deduplication on all the flexible volumes in the system by using the `sis off` command. If you attempt to remove the deduplication license without first disabling deduplication, you will receive a warning message asking you to disable this feature.

Note: Any volume deduplication that occurred before removing the license remains unchanged.

8.2 MAXIMUM VOLUME AND DATA SIZE LIMITS

8.2.1 MAXIMUM VOLUME SIZE LIMIT

Deduplication has limits on maximum supported volume sizes. These limits vary depending on platform and Data ONTAP version, and will be the same on each system regardless of whether the volumes reside in standard 32-bit aggregates or the new 64-bit aggregates available beginning with Data ONTAP 8.0. Table 12 shows a quick reference to volume size limits.

Table 12) Maximum supported volume sizes for deduplication.

	Data ONTAP 7.2.5.1–7.3.0	Data ONTAP 7.3.1.1 and Later 7.3.X	Data ONTAP 8.0	Data ONTAP 8.0.1
FAS2020/N3300	0.5TB	1TB	Not supported	Not supported
FAS2040/N3400	Not supported	3TB ⁵	3TB	16TB
FAS2050/N3600	1TB	2TB	Not supported	Not supported
R200	4TB	4TB	Not supported	Not supported
FAS3020/N5200	1TB	2TB	Not supported	Not supported
FAS3040/N5300	3TB	4TB	4TB	16TB
FAS3050/N5500	2TB	3TB	Not supported	Not supported
FAS3070/N5600	6TB	16TB	16TB	16TB
FAS3140/N6040	3TB	4TB	4TB	16TB
FAS3160/N6060	6TB	16TB	16TB	16TB
FAS3170/N6070	10TB	16TB	16TB	16TB
FAS3210	Not supported	4TB ⁶	Not supported	16TB
FAS3240	Not supported	16TB ⁶	Not supported	16TB
FAS3270	Not supported	16TB ⁶	Not supported	16TB
FAS6030/N7600	10TB	16TB	16TB	16TB
FAS6040/N7700	10TB	16TB	16TB	16TB
FAS6070/N7800	16TB	16TB	16TB	16TB
FAS6080/N7900	16TB	16TB	16TB	16TB
FAS6210	Not supported	Not supported	Not supported	16TB
FAS6240	Not supported	Not supported	Not supported	16TB
FAS6280	Not supported	Not supported	Not supported	16TB

⁵ Deduplication on this platform is only supported with Data ONTAP 7.3.2 or later. Prior versions are not supported.

⁶ Only supported with Data ONTAP 7.3.5 or later. Prior versions are not supported.

8.2.2 MAXIMUM SHARED DATA LIMIT FOR DEDUPLICATION

In addition to the maximum supported volume sizes, there is also a maximum shared data limit for deduplication. Recall that deduplication achieves savings by keeping a single copy of a block and replacing all duplicates of that block with pointers to the single block. As a result of deduplication, much of the logical data will now be stored as pointers to a reduced number of physical blocks written to the volume. It is very common for a single block to have many pointer references. The logical data that ends up being stored as pointers is called shared data. The maximum shared data limit per volume for deduplication is 16TB, regardless of the platform type. Once this limit is reached, any additional new data written to the volume will not be deduplicated. Writes to the volume will continue to work successfully until the volume is completely full.

8.2.3 MAXIMUM TOTAL DATA LIMIT

For a deduplicated volume the sum of the maximum supported volume sizes and the maximum shared data limit for deduplication is the maximum total data limit for a flexible volume.

Consider the following theoretical example as it progressively adds more data to a volume. This example assumes deduplication can provide ~99% savings on this dataset.

The sample system is a FAS3040 running Data ONTAP 7.3.2. Keep in mind this system has a volume size limit of 4TB and a shared data limit of 20TB.

Table 13) Example illustrating maximum total data limit: deduplication only.

Action	Amount of Shared Data on Volume	Physical Space Usage	Total Data Size (Logical)
4TB volume is created and populated with 4TB of data.	4TB	4TB	4TB
Deduplication is run.	4TB	1MB	4TB
An additional 4TB is written to the volume.	8TB	4TB	8TB
Deduplication is run again.	8TB	1MB	8TB
An additional 4TB is written to the volume.	12TB	4TB	12TB
Deduplication is run again.	12TB	1MB	12TB
An additional 4TB is written to the volume.	16TB	4TB	16TB
Deduplication is run again.	16TB	1MB	16TB
An additional 4TB is written to the volume.	16TB	4TB	20TB
Deduplication is run again.	16TB	4TB	20TB
Attempt to write additional data to the volume.	New writes will fail since the volume is full.		

In this example after we reached 16TB of shared data on the volume any additional writes will not be deduplicated.

Table 14) Maximum total data limit in a deduplicated volume.

	Data ONTAP 7.2.5.1–7.3.0	Data ONTAP 7.3.1.1 and Later 7.3.X	Data ONTAP 8.0	Data ONTAP 8.0.1
FAS2020/N3300	16.5TB	17TB	Not supported	Not supported
FAS2040/N3400	Not supported	19TB ⁷	19TB	32TB
FAS2050/N3600	17TB	18TB	Not supported	Not supported
R200	20TB	20TB	Not supported	Not supported
FAS3020/N5200	17TB	18TB	Not supported	Not supported
FAS3040/N5300	19TB	20TB	20TB	32TB
FAS3050/N5500	18TB	19TB	Not supported	Not supported
FAS3070/N5600	22TB	32TB	32TB	32TB
FAS3140/N6040	19TB	20TB	20TB	32TB
FAS3160/N6060	22TB	32TB	32TB	32TB
FAS3170/N6070	26TB	32TB	32TB	32TB
FAS3210	Not supported	20TB ⁸	Not supported	32TB
FAS3240	Not supported	32TB ⁸	Not supported	32TB
FAS3270	Not supported	32TB ⁸	Not supported	32TB
FAS6030/N7600	26TB	32TB	32TB	32TB
FAS6040/N7700	26TB	32TB	32TB	32TB
FAS6070/N7800	32TB	32TB	32TB	32TB
FAS6080/N7900	32TB	32TB	32TB	32TB
FAS6210	Not supported	Not supported	Not supported	32TB
FAS6240	Not supported	Not supported	Not supported	32TB
FAS6280	Not supported	Not supported	Not supported	32TB

⁷ Deduplication on this platform is only supported with Data ONTAP 7.3.2 or later. Prior versions are not supported.

⁸ Only supported with Data ONTAP 7.3.5 or later. Prior versions are not supported.

For Data ONTAP versions prior to 7.3.1, from a deduplicated volume size limit perspective, a volume cannot exceed the size limit for the entire life of the volume. That is, if a volume is larger than the maximum size and is then shrunk, you still cannot enable deduplication on that volume. If you need to run deduplication on a volume that was (at some point in time) larger than the maximum supported size, you can do so by creating a new volume and migrating the data from the old volume to the newly created volume. On versions of Data ONTAP 7.3.1 and later you can resize the volume to within the maximum dedupe volume size limit and then enable dedupe on that same volume.

Here is an example of the message displayed if the volume is, or has been, too large to enable deduplication:

```
london-fs3> sis on /vol/projects  
Volume or maxfiles exceeded max allowed for SIS: /vol/projects
```

8.3 DEDUPLICATION SCANNER OPERATIONS TAKING TOO LONG TO COMPLETE

When the deduplication process begins, it will sort and search the fingerprint database, and then remove all verified duplicates. This process can be time consuming, depending on the amount of data to be processed. In some cases other factors might play a key role, and should be considered accordingly.

When running a deduplication scanner process, verify that there are enough resources available. Remember that the deduplication scanner process is a low-priority process and will give up resources for other processes. Operations Manager is a good starting point since it can provide good resource information to see how much CPU, memory, and I/O bandwidth are available.

Another consideration is the number of simultaneous deduplication processes that are running on a single system. Make sure to follow the best practice of using a gradual approach to determine how many simultaneous deduplication scanner processes you can run safely in your environment before saturating your system resources.

A common practice is to run deduplication scanner processes at off-peak hours. If deduplication scanners do not complete during off-peak hours, you might consider stopping the deduplication scanner processes during peak hours and resuming during off-peak hours.

NetApp strongly recommends that performance testing be performed prior to using deduplication in your environment. For information on performance testing at a customer environment, see [TR-3849: NetApp Deduplication for FAS and V-Series Performance and Savings Testing in Customer Environments](#).

8.4 LOWER-THAN-EXPECTED SPACE SAVINGS

If you are not seeing significant savings when using deduplication, consider the following factors.

There might be little duplicate data within the volume

Run the Space Savings Estimation Tool (SSET) against the dataset to get an idea of the amount of duplicate data within the dataset.

There might not be enough space for deduplication to run

On Data ONTAP versions prior to 7.3, each flexible volume should have 6% of the total data's worth of free space.

On Data ONTAP versions 7.3 or later, the aggregate should have 4% (fingerprint + change logs) of the total data's worth of free space for all deduplicated flexible volumes, and each flexible volume should have 2% of the total data's worth of free space.

For additional details about the overhead associated with the deduplication metadata files, see the section "Deduplication Metadata Overhead" within this document.

The LUN configuration might be masking the deduplication savings

Different LUN configurations will cause freed blocks to be returned to different logical pools within the storage system. Refer to the LUN section of this document for information about deduplication with LUNs and how space savings are recognized with different configurations.

Snapshot copies

Snapshot copies lock blocks in place by design, so that they are available for data recovery. This locking mechanism will not allow blocks that are freed by deduplication to be returned to the free pool until the locks expire or are deleted. If you are using deduplication with Snapshot copies, consider the following:

- If possible, run the deduplication scanner to completion before creating a Snapshot copy.
- Use the `snap list` command to see what Snapshot copies exist and the `snap delete` command to remove them. Alternatively, wait for the Snapshot copies to expire and the space savings to appear.

For more information about using deduplication with Snapshot copies, see the "Snapshot Copies" section of this document.

Misaligned data can result in reduced or little space savings

This issue is a bit harder to track down. If you believe this might be the issue, it is best to contact NetApp Support for assistance.

Misaligned blocks are a possibility if best practices are not followed when configuring a virtualization environment. For more information on VMware best practices, see [TR-3749: NetApp and VMware vSphere Storage Best Practices](#).

8.5 SYSTEM SLOWDOWN

8.5.1 UNEXPECTEDLY SLOW READ PERFORMANCE CAUSED BY ADDING DEDUPLICATION

The information in this section is provided assuming that basic proof-of-concept testing has been performed prior to running in production to understand what performance to expect when deduplication is enabled on the NetApp system. NetApp strongly recommends that performance testing be performed prior to using deduplication in your environment. See [TR-3849: NetApp Deduplication for FAS and V-Series Performance and Savings Testing in Customer Environments](#) for more information on how to perform testing in a test environment.

In older versions of Data ONTAP, pre-7.2.6 in the 7.2 codeline and pre-7.3.1 in the 7.3 codeline, there can be a reduction in sequential read performance. Some examples of applications that use sequential reads are NDMP, qtree SnapMirror, SnapVault, some NFS apps, and so on. In addition to sequential reads, these older versions of Data ONTAP might also see degraded performance when reading a volume with a high amount of deduplication savings.

In either cases NetApp recommends an upgrade to Data ONTAP 7.2.6 or later, or Data ONTAP 7.3.1 or later. These two releases introduced intelligent caching, which resulted in the cache being deduplication aware. This means that blocks that are being shared as a result of deduplication will now be cached, resulting in a performance boost for read requests. Intelligent cache also applies to the Performance Acceleration Module (PAM) and Flash Cache (PAM II). Refer to the section about using deduplication with other NetApp features for information about using deduplication with PAM and Flash Cache cards.

In the case of random reads, there is usually not much impact on performance, if at all, from random reads of deduplicated data.

If slow read performance continues to be an issue, NetApp highly recommends that NetApp Support be contacted for in-depth troubleshooting. In many cases there are other factors such as misconfigured applications or conflicting policies that can be easily fixed to regain acceptable performance. See the “Contact Information for Support” section of this document for contact information and data collection guidance.

8.5.2 UNEXPECTEDLY SLOW WRITE PERFORMANCE CAUSED BY ADDING DEDUPLICATION

The information in this section is provided assuming that basic proof-of-concept testing has been performed to understand what performance to expect when deduplication is enabled on the NetApp system. NetApp strongly recommends that performance testing be performed prior to using deduplication in your environment. See [TR-3849: NetApp Deduplication for FAS and V-Series Performance and Savings Testing in Customer Environments](#) for more information on how to perform testing in a test environment.

Enabling deduplication on a volume will cause the creation of deduplication metadata (fingerprints) as data is written to the volume. The deduplication metadata is a standard part of the deduplication process. The creation of the metadata is not typically an issue on systems that have available resources for the deduplication process; however, if the storage system resources are used by other applications, then even the slightest impact can have a visible impact on the storage system.

If write performance appears to be degraded, check the NetApp system resources (CPU, memory, and I/O) to make sure they are not saturated. This can be easily done with NetApp Operations Manager or SANscreen®, if available. If system resources are indeed saturated, you can consider disabling deduplication to see if performance resumes. Remember, turning off deduplication will not undo deduplicated blocks already on disk. Disabling deduplication will stop the creation of deduplication metadata and stop any future deduplication from running. All deduplication savings achieved prior to that point will continue to exist.

It is also worth noting that write performance will vary based on different platforms. So it is not typically a good approach to compare write performance results across different NetApp platforms. Write performance can also be affected by using slower disk drives. Thus, comparing FC drives, or SAS drives, to SATA drives can give different results.

If slow write performance continues to be an issue, NetApp highly recommends that NetApp Support be contacted for in-depth troubleshooting. In many cases there are other factors such as misconfigured applications or conflicting policies that can be easily fixed to regain acceptable performance. See the “Contact Information for Support” section of this document for contact information and data collecting guidance.

8.5.3 SYSTEM RUNNING MORE SLOWLY SINCE ENABLING DEDUPLICATION

The information in this section is provided assuming that basic proof-of-concept testing has been performed to understand what performance to expect when deduplication is enabled on the NetApp system. NetApp strongly recommends that performance testing be performed prior to using deduplication in your environment to fully understand this impact. For information on how to run a test with deduplication, see [TR-3849: NetApp Deduplication for FAS and V-Series Performance and Savings Testing in Customer Environments](#). For more details, see the performance section of this document.

Deduplication has been known to compound system performance issues associated with misalignment. For instance, there could be a misconfiguration that has been causing the NetApp system to run at less than optimum performance levels without being noticed. A good example is the case of misaligned blocks in a VMware ESX environment. See [TR-3749: NetApp and VMware vSphere Storage Best Practices](#) for additional details. When deduplication is enabled in this environment, it would be common to see noticeable degradation in system performance, because deduplication will effectively cause the misconfigured blocks to be accessed more often since they will now be shared and more likely to be accessed. In this case, other troubleshooting efforts outlined within this document would not resolve this issue. The best approach for this scenario is to contact NetApp Support. The Support team would be able to analyze the system and recommend a remedy to resolve the performance issue while maintaining the space savings.

Another common performance caveat occurs when too many deduplication scanner processes are run simultaneously on a single system. The max number of simultaneous deduplication scanner processes that can be run on a single storage system is eight. Running eight simultaneous deduplication scanner processes will use significant resources, and can reduce application performance. The number of simultaneous deduplication scanner processes that will run on a system should be reevaluated as additional applications and processes are run on the storage systems. In such a case you can consider running fewer deduplication scanner processes simultaneously. The best approach would be to rerun the performance testing on the system to understand how deduplication scanners will run once the additional workload had been added to it. This performance testing would likely entail reducing the number of simultaneous deduplication scanners processes to one for a better understanding of the performance effects, and then systematically increasing the number of simultaneous deduplication processes until a system performance threshold is reached.

Although it is relatively easy to undeduplicate a flexible volume, the process is resource intensive and can take a significant amount of time to complete. As an alternative, you can try simply to disable deduplication. This would stop any future deduplication from occurring. All space savings from deduplication run prior to their being disabled would continue to exist. Undoing deduplication should be considered only if recommended by NetApp Support, following their analysis of the environment to ascertain whether the source of the problem is being properly addressed. If you must undo deduplication, it can be done while the flexible volume is online, as described in the "Removing space savings" section within this document.

If slow system performance continues to be an issue, NetApp highly recommends that NetApp Support be contacted for in-depth troubleshooting. In many cases there are other factors such as misconfigured applications or conflicting policies that can be easily fixed to regain acceptable performance. See the "Where to get more help" section of this document for Support contact information and data collecting guidance.

8.6 REMOVING SPACE SAVINGS (UNDO)

NetApp recommends contacting NetApp Support prior to undoing deduplication on a volume to make sure that removing the space savings is really necessary. In many cases system performance can be restored by finding the true source of the degradation, which often can be unrelated to deduplication.

It is relatively easy to undeduplicate a flexible volume and turn it back into a regular flexible volume. This can be done while the flexible volume is online, as described in this section.

Note: The sis undo operation takes system resources and thus should be scheduled to run during low usage times.

8.6.1 UNDEDUPLICATING A FLEXIBLE VOLUME

To remove deduplication from a volume you must first turn off deduplication on the flexible volume. To do this you would use the command:

```
sis off <vol_name>
```

This command stops fingerprints from being written to the change log as new data is written to the flexible volume. If this command is used, and then deduplication is turned back on for this flexible volume, the flexible volume must be rescanned with the `sis start -s` command.

Next, if you want to remove the deduplication savings by recreating the duplicate blocks in the flexible volume, you would use the following command (the `sis undo` command is available only in Diag mode):

```
priv set diag
```

```
sis undo <vol_name>
```

This command will recreate the duplicate blocks and delete the fingerprint and change log files.

Here is an example of undeduplicating a flexible volume:

```
fas6070-ppe02*> df -s /vol/volHomeNormal
Filesystem          used          saved          %saved
/vol/volHomeNormal/ 83958172    98060592         54%
```

```
r200-rtp01> sis status /vol/ volHomeNormal
Path                State      Status  Progress
/vol/VolHomeNormal  Enabled   Idle    Idle for 11:11:13
```

```
fas6070-ppe02> sis off /vol/volHomeNormal
SIS for "/vol/volHomeNormal" is disabled.
```

```
fas6070-ppe02> priv set diag
```

Warning: These diagnostic commands are for use by NetApp personnel only.

```
fas6070-ppe02*> sis undo /vol/volHomeNormal
fas6070-ppe02*> Mon Jul  5 15:52:13 EDT [waf1.scan.start:info]: Starting SIS
volume scan on volume volHomeNormal.
```

```
fas6070-ppe02*> sis status /vol/volHomeNormal
Path                State      Status  Progress
/vol/volHomeNormal  Disabled  Undoing  90 GB Processed
```

```
fas6070-ppe02> sis status /vol/volHomeNormal
No status entry found.
```

```
fas6070-ppe02> df -s /vol/volHomeNormal
Filesystem          used          saved          %saved
/vol/volHomeNormal/ 182008592         0           0%
```

Note: If `sis undo` starts processing and then there is not enough space to undeduplicate, it stops, sends a message about insufficient space, and leaves the flexible volume deduplicated. Use `df -r` to find out how much free space you really have, and then remove some data from the volume or delete some Snapshot copies to provide the needed free space.

8.6.2 UNDEDUPLICATING A FLEXIBLE VOLUME

To remove both deduplication savings from a volume you must first turn off deduplication on the flexible volume. To do this you would use the commands:

```
sis off <vol_name>
```

This command stops new data that is written to the flexible volume from having its fingerprints written to the change log. For more details refer to the section about undeduplicating a flexible volume in this document.

Use the following command to remove deduplication savings and recreate the duplicate blocks (the `sis undo` command is available only in Diag mode):

```
sis undo <vol_name>
```

Here is an example of undeduplicating a flexible volume:

```
fas6070-ppe02*> df -s /vol/volExisting
Filesystem      used      saved   %saved
/vol/volExisting/  73054628  108956860    60%
```

```
fas6070-ppe02*> vol status /vol/volExisting
      Volume State      Status      Options
volExisting online      raid_dp, flex  nosnap=on,
sis      guarantee=none,
      fractional_reserve=0
      Volume UUID: d685905a-854d-11df-806c-00a098001674
      Containing aggregate: 'aggrTest'
      Volinfo mode: 7-mode
```

```
fas6070-ppe02> sis off /vol/volExisting
SIS for "/vol/volExisting" is disabled.
```

```
fas6070-ppe02> vol status volExisting
      Volume State      Status      Options
volExisting online      raid_dp, flex  nosnap=on, guarantee=none,
sis      fractional_reserve=0
      Volume UUID: d685905a-854d-11df-806c-00a098001674
      Containing aggregate: 'aggrTest'
```

```
fas6070-ppe02> priv set diag
```

Warning: These diagnostic commands are for use by NetApp personnel only.

```
fas6070-ppe02*> sis undo /vol/volExisting
```

```
fas6070-ppe02*> Mon Jul 5 17:01:23 EDT [wafl.scan.start:info]: Starting SIS
volume scan on volume volExisting.
```

```
fas6070-ppe02> sis status /vol/volExisting
```

Path	State	Status	Progress
/vol/volExisting	Disabled	Undoing	65 GB Processed

```
fas6070-ppe02> sis status /vol/volExisting
```

```
No status entry found.
```

```
fas6070-ppe02> df -s /vol/volExisting
```

Filesystem	used	saved	%saved
/vol/volExisting/	182011480	0	0%

8.6.3 REVERTING A FLEXIBLE VOLUME TO BE ACCESSIBLE IN OTHER DATA ONTAP RELEASES

Once you have completely removed deduplication savings from your volume's active file system and Snapshot copies, you can revert the volume to make it accessible to other Data ONTAP releases. To do this you can use the `revert_to` command for deduplication.

To do this you would use the commands:

```
sis revert_to <ver> <vol/vol_name>
```

8.7 WHERE TO COLLECT TROUBLESHOOTING INFORMATION

This section provides guidance for collecting system information for deduplication. This information can be very useful for troubleshooting.

8.7.1 LOCATION OF LOGS AND ERROR MESSAGES

The location of the deduplication log file is:

```
/etc/log/sis
```

8.7.2 UNDERSTANDING DEDUPLICATION ERROR MESSAGES

Deduplication error messages with explanations:

Registry errors:	Check if vol0 is full (only in Data ONTAP 7.2.X).
Metafile op errors:	Check if the deduplicated volume is full (in Data ONTAP 7.2.X).
Metafile op errors:	Check if the deduplicated aggregate is full (in Data ONTAP 7.2.X).
License errors:	Check if the license is installed.
Change log full error:	Perform a sis start operation to empty the change log metafile when finished.

Beginning with Data ONTAP 7.2.5.1 of the 7.2 codeline and Data ONTAP 7.3 of the 7.3 codeline, additional details are included in the sis logs. The additional information includes detailed information about how blocks and fingerprints are processed.

Example:

```
Fri Aug 15 00:43:49 PDT /vol/max1 [sid: 1218783600] Stats (blks
gathered 0,finger prints sorted 603349248,dups found 0,new dups found
4882810,blks deduped 0,finger prints checked 0,finger prints deleted 0)
```

This example reveals the following information:

Total number of new blocks created since the last deduplication process ran = 0

Total number of fingerprint entries (new + preexisting) that were sorted for this deduplication process = 603349248

Total number of duplicate blocks found = 0

Total number of new duplicate blocks found = 4882810

Total number of duplicate blocks that were deduplicated = 0

Total number of fingerprints checked for stale condition = 0

Total number of stale fingerprints deleted = 0

8.7.3 UNDERSTANDING OPERATIONS MANAGER EVENT MESSAGES

A new event has been introduced beginning with Operations Manager 3.8, "Deduplication: Volume is over deduplicated." By default, this event is triggered when the size of the deduplicated volume will not be large enough to fit all data if the data is undeduplicated by using the `sis undo` command.

Users can change the default threshold settings in Operations Manager to make it a very high value so that the event does not get triggered. The triggering of this event does not change the deduplication operation in any way. The event is simply a notification that this condition has occurred.

For a complete list of events, refer to the [Operations Manager Administration Guide 4.0](#).

8.8 ADDITIONAL DEDUPLICATION REPORTING

8.8.1 DEDUPLICATION REPORTING WITH SIS STATUS

Basic status for deduplication can be collected using the `sis status` command. The following are some of the most common questions that can be addressed with the `sis status` command:

- What command shows how long the last deduplication operation ran?
- What command shows how much data was fingerprinted in the changelog during the last deduplication operation?
- What command shows the deduplication schedule?

The following is an example of the output from the `sis status -l` command, followed by definitions for each value. This information corresponds to the last deduplication operation.

```
nss-u> sis status -l /vol/test
Path:                /vol/test
State:               Enabled
Status:              Idle
Progress:            Idle for 00:00:33
Type:                Regular
Schedule:            sun-sat@0
Last Operation Begin: Thu Sep  4 17:45:23 GMT 2008
Last Operation End:   Thu Sep  4 17:46:36 GMT 2008
Last Operation Size:  2048 MB
```

Last Operation Error: -

Path: Absolute path of the volume.

State: Shows if deduplication is enabled or disabled for the volume.

Status: Shows the current state of the deduplication process: Active, Idle, Initializing, Undoing.

Progress: Shows how long the deduplication state has been in the current phase.

Type: Shows the type of the deduplicated volume: Regular, SnapVault.

Schedule: Shows the deduplication schedule for the volume.

Last Operation Begin: The time when the last deduplication process (operation) began.

Last Operation End: The time when the last deduplication process (operation) ended.

Last Operation Size: The amount of new data that was processed during the last deduplication process (operation).

Last Operation Error: The error that occurred, if any, during the last deduplication process (operation).

8.9 DEDUPLICATION REPORTING WITH SIS STAT

For detailed status information, you can use `priv set diag` and then use the `sis stat -l` command for long, detailed listings. Here are details about the `sis stat` command:

When the volume name is omitted, the command is executed for all known SIS volumes.

- `-l` lists all the details about the volume.
- `-b` shows the disk space usage and saved disk space in number of blocks.
- `-v` shows the disk space usage and saved disk space in number of bytes; if the `stat` command was executed without any option, it runs with the `iv` option only.
- `-g` lists all information about the deduplication block and change log buffer status.
- `-lv` generates reference histograms that can be used for troubleshooting, if instructed by NetApp Global Support.
- The block-sharing histogram indicates how many shared blocks are contiguous (adjacent to one another).
- The refcount histogram shows the total number of refcounts. That is, it shows the number of blocks with one reference pointer to them, then the number of blocks with two references to them, then the number of blocks with three references to them, and so on.

The following is an example of output for the `sis stat -lv` command that is available in Diag mode, followed by definitions for each value. This information is for the entire life of the volume, as opposed to the last deduplication operation as reported by the `sis status -l` command.

```
netappl*> priv set diag; sis stat -lv /vol/test
Path:                /vol/test
Allocated:           32604 KB
Shared:              8228 KB
```

```

Saving:                2088924 KB
%Saved:                98 %
Unclaimed:            8228 KB

Max Refcount:         255

Total Processed:      4194304 KB
Total Process Time:   00:01:48
SIS Files:            1
Succeeded Op:         2
Started Op:           2
Failed Op:            0
Stopped Op:           0
Deferred Op:          0
Succeeded Check Op:  1
Failed Check Op:     0
Total Sorted Blocks: 1572864
Same Fingerprint:    1048575
Same FBN Location:   0
Same Data:            345
Same VBN:             0
Mismatched Data:     0
Max Reference Hits:   3770
Staled Recipient:    0
Staled Donor:         3
File Too Small:      0
Out of Space:        0
FP False Match:      0
Delino Records:      0

```

Block Sharing Histogram:

```

0 1 2 3 4 5 6 7
0 1 0 0 0 0 0 0
<16 <24 <32 <40 <48 <56 <64 UNUSED
0 0 0 0 0 0 0 0
<128 <192 <256 <320 <384 <448 <512 UNUSED
0 0 0 0 0 0 0 0
<1024 <1536 <2048 <2560 <3072 <3584 <4096 UNUSED
0 0 1 0 74 0 0 0
<8192 <12288 <16384 <20480 <24576 <28672 <32768 UNUSED
0 0 0 0 0 0 0 0
<65536 <98304 <131072 <163840 <196608 <229376 <262144 >262144
0 0 0 0 0 0 0 0

```

Reference Count Histogram:

```

0 1 2 3 4 5 6 7
0 0 0 0 0 0 0 0
<16 <24 <32 <40 <48 <56 <64 UNUSED
0 0 0 0 0 0 0 0
<128 <192 <256 <320 <384 <448 <512 UNUSED
0 1 895 0 0 0 0 0
<1024 <1536 <2048 <2560 <3072 <3584 <4096 UNUSED
0 0 0 0 0 0 0 0
<8192 <12288 <16384 <20480 <24576 <28672 <32768 UNUSED
0 0 0 0 0 0 0 0
<65536 <98304 <131072 <163840 <196608 <229376 <262144 >262144
0 0 0 0 0 0 0 0

```

nss-u164*>

Path: Absolute path of the volume.

Allocated: Total allocated KB in the dense volume.

Shared: Total space shared by doing deduplication.

Saving: Total saving due to deduplication.

%Saved: Percentage of saved space over allocated space.

Unclaimed: Space in KB that was shared and now not referred by anybody.

Max Refcount: Maximum number of references to a shared block.

Total Processed: Space in KB processed by the deduplication engine.

Total Processed Time: Total time taken for the deduplication engine to process the total processed amount of space.

SIS Files: The number of files that share blocks either with their own blocks or with other files.

Succeeded Op: Number of deduplication operations that have succeeded.

Started Op: Number of deduplication operations that have started.

Failed Op: Number of deduplication operations that were aborted due to some failure.

Stopped Op: Number of deduplication operations that have been stopped.

Deferred Op: Number of deduplication operations that have been deferred because of hitting the watermark of number of concurrent deduplication processes running.

Succeeded Check Op: Number of fingerprint check operations that have succeeded.

Failed Check Op: Number of fingerprint check operations that have failed.

Total Sorted Blocks: Number of blocks fingerprinted and sorted based on the fingerprints.

Same Fingerprint: Total number of blocks that have the same fingerprints.

Same FBN Location: Number of deduplications that did not happen because the donor and recipient blocks have the same block number in the same file.

Same Data: Number of blocks that have matching fingerprints and the same data.

Same VBN: Number of files that have the same VBN in their buftrees.

Mismatched Data: Number of blocks that have the same fingerprints but mismatched data. This counter is not persistent across volumes.

Max Reference Hits: Number of blocks that are maximally shared.

Staled Recipient: Total number of recipient inodes' blocks that were not valid during deduplication. This is the count for the whole volume.

Staled Donor: Number of donor inodes' blocks that were stale during deduplication. This is the count for the whole volume.

Out of Space: Number of deduplication operations that were aborted due to lack of space.

FP False Match: Number of blocks that have fingerprint match, but the data does not match. This is a persistent per-volume counter.

Delino records: Number of records associated with the deleted inode or stale inode.

8.10 WHERE TO GET MORE HELP

8.10.1 CONTACT INFORMATION FOR SUPPORT

For additional support, contact one of the following:

- Your local account team
- Systems engineer
- Account manager
- [NetApp Global Services](#)
- [NOW](#) (NetApp on the Web)
- 888.4.NETAPP (United States and Canada)
- 00.800.44.NETAPP (EMEA/Europe)
- +800.800.80.800 (Asia/Pacific)

8.10.2 INFORMATION TO GATHER BEFORE CONTACTING SUPPORT

Deduplication Information

The following deduplication commands and logs provide useful information for troubleshooting the root cause of deduplication issues. This information is very useful when working with NetApp Global Support.

- `priv set diag; sis stat -lv /vol/<vol_name>`
- `priv set diag; sis status -l /vol/<vol_name>`
- `priv set diag; sis check -c /vol/<vol_name>`
- `snap list <vol_name>`
- `df -sh <vol_name>`
- `df -rh <vol_name>`
- `df -h <vol_name>`
- All dedupe logs located in `/etc/log/ - sis, sis.0, sis.1, sis.2, sis.3, sis.4, sis.5`
- Inspect the EMS logs for the time when the issue is seen

9 ADDITIONAL READING AND REFERENCES

- TR-3440: Operations Manager, Protection Manager, and Provisioning Manager Sizing Guide
<http://media.netapp.com/documents/tr-3440.pdf>
- TR-3446: SnapMirror Async Best Practices Guide

- <http://media.netapp.com/documents/tr-3446.pdf>
- TR-3742: Using FlexClone to Clone Files and LUNs
<http://media.netapp.com/documents/tr-3742.pdf>
- TR-3487: SnapVault Best Practice Guide
<http://media.netapp.com/documents/tr-3487.pdf>
- TR-3483: Thin Provisioning in a NetApp SAN or IP SAN Enterprise Environment
<http://media.netapp.com/documents/tr-3483.pdf>
- TR-3548: MetroCluster Design and implementation
<http://media.netapp.com/documents/tr-3548.pdf>
- TR-3263: WORM Storage on Magnetic Disks Using SnapLock Compliance and SnapLock Enterprise
<http://media.netapp.com/documents/tr-3263.pdf>
- TR-3450: Active-Active Controller Configuration Overview and Best Practice Guidelines
<http://media.netapp.com/documents/tr-3450.pdf>
- TR-3814: Data Motion Best Practices
<http://media.netapp.com/documents/tr-3814.pdf>
- TR-3466: Open Systems SnapVault Best Practice Guide
<http://media.netapp.com/documents/tr-3466.pdf>
- TR-3747: Best Practices for File System Alignment in Virtual Environments
<http://media.netapp.com/documents/tr-3747.pdf>
- TR-3705: NetApp and VMware VDI Best Practices
<http://media.netapp.com/documents/tr-3705.pdf>
- TR-3770: 2,000-Seat VMware View on NetApp Deployment Guide Using NFS: Cisco Nexus Infrastructure
<http://media.netapp.com/documents/tr-3770.pdf>
- TR-3428: NetApp and VMware Virtual Infrastructure 3 Storage Best Practices
<http://media.netapp.com/documents/tr-3428.pdf>
- TR-3749: NetApp and VMware vSphere Storage Best Practices
<http://media.netapp.com/documents/tr-3749.pdf>
- TR-3702: NetApp Storage Best Practices for Microsoft Virtualization
<http://media.netapp.com/documents/tr-3702.pdf>
- WP-7053: The 50% Virtualization Guarantee Program Technical Guide
<http://media.netapp.com/documents/wp-7053.pdf>
- TR-3694: NetApp and Citrix XenServer 4.1: Building a Virtual Infrastructure from Server to Storage
<http://media.netapp.com/documents/tr-3694.pdf>
- TR-3584: Microsoft Exchange 2007 Disaster Recovery Model Using NetApp Solutions
<http://media.netapp.com/documents/tr-3584.pdf>
- TR-3651: Microsoft Exchange 2007 SP1 Continuous Replication Best Practices Guide
<http://media.netapp.com/documents/tr-3651.pdf>
- TR-3824: Storage Efficiency and Best Practices for Microsoft Exchange Server 2010
<http://media.netapp.com/documents/tr-3824.pdf>
- TR-3712: Oracle VM and NetApp Storage Best Practices Guide
<http://media.netapp.com/documents/tr-3712.pdf>

10 VERSION TRACKING

Before Version 6		No tracking provided
Version 6	Mar 2009	Update for Data ONTAP 7.3.1
Version 7	Jan 2010	Optimization updates and updates for Data ONTAP 7.3.2
Version 8	Dec 2010	Updates for Data ONTAP 8.0.1 and slight format update

NetApp provides no representations or warranties regarding the accuracy, reliability or serviceability of any information or recommendations provided in this publication, or with respect to any results that may be obtained by the use of the information or observance of any recommendations provided herein. The information in this document is distributed AS IS, and the use of this information or the implementation of any recommendations or techniques herein is a customer's responsibility and depends on the customer's ability to evaluate and integrate them into the customer's operational environment. This document and the information contained may be used solely in connection with the NetApp products discussed in this document.



www.netapp.com

© 2011 NetApp, Inc. All rights reserved. No portions of this document may be reproduced without prior written consent of NetApp, Inc. Specifications are subject to change without notice. NetApp, the NetApp logo, Go further, faster, AutoSupport, Data ONTAP, FlexCache, FlexClone, FlexVol, MetroCluster, MultiStore, NearStore, NOW, RAID-DP, SANscreen, SnapDrive, SnapLock, SnapMirror, SnapRestore, Snapshot, SnapVault, vFiler, and WAFL are trademarks or registered trademarks of NetApp, Inc. in the United States and/or other countries. Microsoft, SharePoint, SQL Server, and Windows are registered trademarks of Microsoft Corporation. Linux is a registered trademark of Linus Torvalds. Oracle is a registered trademark of Oracle Corporation. Symantec, Backup Exec, and NetBackup are trademarks of Symantec Corporation. VMware is a registered trademark and vSphere is a trademark of VMware, Inc. All other brands or products are trademarks or registered trademarks of their respective holders and should be treated as such.
TR-3505-0211