



IBM System Storage N series
**Data ONTAP 7.3 Storage Efficiency
Management Guide**

Contents

Copyright information	7
Trademark information	9
About this guide	11
Audience	11
Terminology	12
Keyboard and formatting conventions	13
Special messages	14
How to send your comments	14
What storage efficiency is	15
The IBM N series advantage	15
Storage efficiency with SATA storage disks and the Performance Acceleration Module	17
Protection against double disk failure with RAID-DP	19
What RAID-DP protection is	19
How RAID-DP protection works	20
How space management works	21
What kind of space management to use	21
What space guarantees are	22
What space reservation is	23
What fractional reserve is	24
Reasons to set fractional reserve to zero	25
How Data ONTAP can automatically provide more free space for full volumes	26
Thin provisioning for greater efficiencies using FlexVol volumes	27
Benefits of using thin provisioning	28
Storage space management using DataFabric Manager	29
Automate thin provisioning using Provisioning Manager	29
Efficient file recovery with Snapshot copies	31
What Snapshot disk consumption is	31
How Snapshot copies consume disk space	31
What the Snapshot copy reserve is	32
Use of deleted active file disk space	33
Snapshot copies can exceed reserve	34

Recovery of disk space for file system use	34
What file folding means and how it saves disk space	35
Space savings with deduplication	37
How deduplication works	38
What deduplication metadata is	38
Guidelines for using deduplication	39
Performance considerations for deduplication	39
How deduplication works with other features and products	40
Deduplication and Snapshot copies	40
Deduplication and volume SnapMirror	41
Deduplication and qtree SnapMirror	42
Deduplication and SnapVault	42
Deduplication and SnapRestore	43
Deduplication and DataFabric Manager 3.8	43
Deduplication and volume copy	44
Deduplication and FlexClone volumes	45
Deduplication and an active/active configuration	45
Deduplication and VMware	46
Deduplication and MultiStore	46
Thin replication using SnapVault and Volume SnapMirror	49
Disk-to-disk backups using SnapVault	51
What data gets backed up and restored through SnapVault	51
Types of SnapVault deployment	52
How SnapVault backup works	54
How SnapVault backup works for open systems	55
Guidelines for creating a SnapVault relationship	56
About LUN clones and SnapVault	56
Efficient data protection using volume SnapMirror	57
How SnapMirror works	58
Applications of SnapMirror	59
Recommended actions while using SnapMirror	60
SnapMirror deployment variations	60
Data replication from one destination to another in a series (cascading)	61
Virtual data replication with FlexClone technology	63
About FlexClone volumes	63
FlexClone volumes and space guarantees	64

About FlexClone files and FlexClone LUNs	65
How FlexClone files and FlexClone LUNs work	65
Uses of FlexClone files and FlexClone LUNs	67
Considerations when planning FlexClone files or FlexClone LUNs	67
Space savings achieved by using FlexClone files and FlexClone LUNs	68
Storage efficiency on non-IBM N series storage with IBM gateway	69
Restructuring IT with Storage Efficiency calculator	70
How to quantify storage utilization using the Storage Efficiency calculator	70
Index	73

Copyright and trademark information

Copyright information

Copyright ©1994 - 2010 Network Appliance, Inc. All rights reserved. Printed in the U.S.A.

Portions copyright © 2010 IBM Corporation. All rights reserved.

US Government Users Restricted Rights - Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.

No part of this document covered by copyright may be reproduced in any form or by any means—graphic, electronic, or mechanical, including photocopying, recording, taping, or storage in an electronic retrieval system—without prior written permission of the copyright owner.

References in this documentation to IBM products, programs, or services do not imply that IBM intends to make these available in all countries in which IBM operates. Any reference to an IBM product, program, or service is not intended to state or imply that only IBM's product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any of IBM's or NetApp's intellectual property rights may be used instead of the IBM or NetApp product, program, or service. Evaluation and verification of operation in conjunction with other products, except those expressly designated by IBM and NetApp, are the user's responsibility.

No part of this document covered by copyright may be reproduced in any form or by any means—graphic, electronic, or mechanical, including photocopying, recording, taping, or storage in an electronic retrieval system—without prior written permission of the copyright owner.

Software derived from copyrighted NetApp material is subject to the following license and disclaimer:

THIS SOFTWARE IS PROVIDED BY NETAPP "AS IS" AND WITHOUT ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE, WHICH ARE HEREBY DISCLAIMED. IN NO EVENT SHALL NETAPP BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

NetApp reserves the right to change any products described herein at any time, and without notice. NetApp assumes no responsibility or liability arising from the use of products described herein, except as expressly agreed to in writing by NetApp. The use or purchase of this product does not convey a license under any patent rights, trademark rights, or any other intellectual property rights of NetApp.

The product described in this manual may be protected by one or more U.S.A. patents, foreign patents, or pending applications.

RESTRICTED RIGHTS LEGEND: Use, duplication, or disclosure by the government is subject to restrictions as set forth in subparagraph (c)(1)(ii) of the Rights in Technical Data and Computer Software clause at DFARS 252.277-7103 (October 1988) and FAR 52-227-19 (June 1987).

Trademark information

IBM, the IBM logo, and ibm.com are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both. A complete and current list of other IBM trademarks is available on the Web at <http://www.ibm.com/legal/copytrade.shtml>

Linux is a registered trademark of Linus Torvalds in the United States, other countries, or both.

Microsoft, Windows, Windows NT, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.

UNIX is a registered trademark of The Open Group in the United States and other countries.

NetApp; the NetApp logo; the Network Appliance logo; Cryptainer; Cryptoshred; DataFabric; Data ONTAP; Decru; Decru DataFort; FAServer; FilerView; FlexCache; FlexClone; FlexShare; FlexVol; FPolicy; gFiler; Go further, faster; Manage ONTAP; MultiStore; NearStore; NetCache; NOW (NetApp on the Web); ONTAPI; RAID-DP; SANscreen; SecureShare; Simulate ONTAP; SnapCopy; SnapDrive; SnapLock; SnapManager; SnapMirror; SnapMover; SnapRestore; SnapValidator; SnapVault; Spinnaker Networks; Spinnaker Networks logo; SpinAccess; SpinCluster; SpinFlex; SpinFS; SpinHA; SpinMove; SpinServer; SpinStor; StoreVault; SyncMirror; Topio; vFiler; VFM; and WAFL are registered trademarks of NetApp, Inc. in the U.S.A. and/or other countries. Network Appliance, Snapshot, and The evolution of storage are trademarks of NetApp, Inc. in the U.S.A. and/or other countries and registered trademarks in some other countries. The StoreVault logo, ApplianceWatch, ApplianceWatch PRO, ASUP, AutoSupport, ComplianceClock, DataFort, Data Motion, FlexScale, FlexSuite, Lifetime Key Management, LockVault, NOW, MetroCluster, OpenKey, ReplicatorX, SecureAdmin, Shadow Tape, SnapDirector, SnapFilter, SnapMigrator, SnapSuite, Tech OnTap, Virtual File Manager, VPolicy, and Web Filer are trademarks of NetApp, Inc. in the U.S.A. and other countries. Get Successful and Select are service marks of NetApp, Inc. in the U.S.A.

All other brands or products are trademarks or registered trademarks of their respective holders and should be treated as such.

Network Appliance is a licensee of the CompactFlash and CF Logo trademarks.

Network Appliance NetCache is certified RealSystem compatible.

Notices

This information was developed for products and services offered in the U.S.A.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe on any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not give you any license to these patents. You can send license inquiries, in writing to:

IBM Director of Licensing
IBM Corporation
North Castle Drive
Armonk, N.Y. 10504-1785
U.S.A.

For additional information, visit the web at:
<http://www.ibm.com/ibm/licensing/contact/>

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law:

INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM web sites are provided for convenience only and do not in any manner serve as an endorsement of those web sites. The materials at those web sites are not part of the materials for this IBM product and use of those web sites is at your own risk.

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Any performance data contained herein was determined in a controlled environment. Therefore, the results obtained in other operating environments may vary significantly. Some measurements may have been made on development-level systems and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some measurement may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

If you are viewing this information in softcopy, the photographs and color illustrations may not appear.

About this guide

You can use your product more effectively when you understand this document's intended audience and the conventions that this document uses to present information.

Note: In this document, the term *gateway* describes IBM N series storage systems that have been ordered with gateway functionality. Gateways support various types of storage, and they are used with third-party disk storage systems—for example, disk storage systems from IBM, HP®, Hitachi Data Systems®, and EMC®. In this case, disk storage for customer data and the RAID controller functionality is provided by the back-end disk storage system. A gateway might also be used with disk storage expansion units specifically designed for the IBM N series models.

The term *filer* describes IBM N series storage systems that either contain internal disk storage or attach to disk storage expansion units specifically designed for the IBM N series storage systems. Filer storage systems do not support using third-party disk storage systems.

This guide documents the IBM features and functionalities that help to significantly improve storage utilization. It covers

- Hardware efficiencies such as SATA high-density disks and the Performance Acceleration Module.
- Software efficiencies such as Snapshot copies, FlexVol volumes, FlexClone volumes, SnapMirror, SnapVault, and Deduplication.
- High-level tools such as the Storage Efficiency calculator and the Restructuring IT with Storage Efficiency calculator.

Next topics

[Audience](#) on page 11

[Terminology](#) on page 12

[Keyboard and formatting conventions](#) on page 13

[Special messages](#) on page 14

[How to send your comments](#) on page 14

Audience

This document is written with certain assumptions about your technical knowledge and experience.

This document is for system administrators and system architects who are familiar with the storage system equipment and who need to perform the following tasks:

- Understand storage efficiency technologies.
- Plan long-term roadmaps for deploying storage.
- Make architectural decisions about storage.

- Recommend, qualify, and deploy features that contribute to storage efficiency.

Terminology

To understand the concepts in this document, you might need to know how certain terms are used.

Storage terms

LUN (logical unit number)	Refers to a logical unit of storage identified by a number.
native disk	Refers to a disk that is sold as local storage for storage systems that run Data ONTAP software.
native disk shelf	Refers to a disk shelf that is sold as local storage for storage systems that run Data ONTAP software.
storage controller	Refers to the component of a storage system that runs the Data ONTAP operating system and controls its disk subsystem. Storage controllers are also sometimes called <i>controllers</i> , <i>storage appliances</i> , <i>appliances</i> , <i>storage engines</i> , <i>heads</i> , <i>CPU modules</i> , or <i>controller modules</i> .
storage system	Refers to the hardware device running Data ONTAP that receives data from and sends data to native disk shelves, third-party storage, or both. Storage systems that run Data ONTAP are sometimes referred to as <i>filers</i>, <i>appliances</i>, <i>storage appliances</i>, <i>gateways</i>, or <i>systems</i>. Note: The term <i>gateway</i> describes IBM N series storage systems that have been ordered with gateway functionality. Gateways support various types of storage, and they are used with third-party disk storage systems—for example, disk storage systems from IBM, HP®, Hitachi Data Systems®, and EMC®. In this case, disk storage for customer data and the RAID controller functionality is provided by the back-end disk storage system. A gateway might also be used with disk storage expansion units specifically designed for the IBM N series models. The term <i>filer</i> describes IBM N series storage systems that either contain internal disk storage or attach to disk storage expansion units specifically designed for the IBM N series storage systems. Filer storage systems do not support using third-party disk storage systems.
third-party storage	Refers to the back-end storage arrays, such as IBM, Hitachi Data Systems, and HP, that provide storage for storage systems running Data ONTAP.

Keyboard and formatting conventions

You can use your product more effectively when you understand how this document uses keyboard and formatting conventions to present information.

Keyboard conventions

Convention	What it means
The IBM NAS support site	Refers to http://www.ibm.com/storage/support/nas/ .
<i>Enter, enter</i>	- Used to refer to the key that generates a carriage return; the key is named Return on some keyboards. - Used to mean pressing one or more keys on the keyboard and then pressing the Enter key, or clicking in a field in a graphical interface and then typing information into the field.
hyphen (-)	Used to separate individual keys. For example, Ctrl-D means holding down the Ctrl key while pressing the D key.
type	Used to mean pressing one or more keys on the keyboard.

Formatting conventions

Convention	What it means
<i>Italic font</i>	- Words or characters that require special attention. - Placeholders for information that you must supply. For example, if the guide says to enter the <code>arp -d <i>hostname</i></code> command, you enter the characters "arp -d" followed by the actual name of the host. - Book titles in cross-references.
Monospaced font	- Command names, option names, keywords, and daemon names. - Information displayed on the system console or other computer monitors. - Contents of files. - File, path, and directory names.
Bold monospaced font	Words or characters you type. What you type is always shown in lowercase letters, unless your program is case-sensitive and uppercase letters are necessary for it to work properly.

Special messages

This document might contain the following types of messages to alert you to conditions that you need to be aware of.

Note: A note contains important information that helps you install or operate the system efficiently.

Attention: An attention notice contains instructions that you must follow to avoid a system crash, loss of data, or damage to the equipment.

How to send your comments

Your feedback is important in helping us provide the most accurate and high-quality information. If you have comments or suggestions for improving this document, send us your comments by e-mail to starpubs@us.ibm.com. Be sure to include the following:

- Exact publication title
- Publication form number (for example, GC26-1234-02)
- Page, table, or illustration numbers
- A detailed description of any information that should be changed

What storage efficiency is

Storage efficiency includes technologies such as FlexVol volume, Snapshot copy, deduplication, SnapVault, SnapMirror, and FlexClone, which help to increase storage utilization and decrease storage costs.

Storage efficiency enables you to store the maximum amount of data for the lowest cost and accommodates rapid data growth while consuming less space. IBM N series strategy for storage efficiency is based on the built-in foundation of storage virtualization and unified storage provided by its core Data ONTAP operating system and Write Anywhere File Layout (WAFL) file system.

The unified storage architecture allows you to efficiently consolidate a storage area network (SAN), network-attached storage (NAS), and secondary storage on a single platform.

High-density disk drives, such as serial advanced technology attachment (SATA) drives mitigated with RAID-DP technology, provide increased efficiency and read performance.

Technologies such as FlexVol volume, Snapshot copy, deduplication, SnapVault, SnapMirror, and FlexClone offer dramatic cost savings and performance improvements. You can use these technologies together to improve storage utilization.

The IBM N series advantage

IBM N series has a rich set of features such as SATA disks, the Performance Acceleration Module, RAID-DP, FlexVol, Snapshot copies, deduplication, SnapVault, SnapMirror, and FlexClone, which help to achieve significant improvements in storage utilization. When used together, these technologies help to achieve increased performance.

IBM N series offers the following technologies to implement storage efficiency:

- SATA disks provide an attractive price point to reduce storage costs and optimize system performance. You can incorporate key features and products, which increase data reliability and availability, in the SATA deployment to implement a reliable storage solution.
- The Performance Acceleration Module improves performance for workloads that are random read-intensive without adding additional disk drives. This read cache helps to reduce latency and improve I/O throughput.
- RAID-DP is a double-parity RAID6 implementation that protects against dual disk drive failures.
- Thin provisioning enables you to maintain a common unallocated storage space that is readily available to other applications as needed. It is based on the FlexVol technology.
- Snapshot copies are a point-in-time, read-only view of a data volume, which consumes minimal storage space. Two Snapshot copies created in sequence differ only by the blocks added or changed in the time interval between the two. This block incremental behavior limits the associated consumption of storage capacity.

- Deduplication saves storage space by eliminating redundant data blocks within a FlexVol volume.
- SnapVault is a centralized and cost-effective solution for replicating Snapshot copies to inexpensive secondary storage. This technology copies only the data blocks that changed since the last backup, and not the entire files, so backups run more frequently and use less storage capacity.
- SnapMirror is a flexible solution for replicating data over local area, wide area, and Fibre Channel networks. It can serve as a critical component in implementing enterprise data protection strategies. You can replicate your data to one or more storage systems to minimize downtime costs in case of a production site failure. You can also use SnapMirror to centralize the backup of data to disks from multiple data centers.
- FlexClone technology copies data volumes, files, and LUNs as instant virtual copies. A Flexclone volume, file, or LUN is a writable point-in-time image of the FlexVol volume or another FlexClone volume, file, or LUN. This technology enables you to use space efficiently, storing only data that changes between the parent and the clone. Use of this feature results in savings in space, power, and dollars. Additionally, clones have the same high performance as their parent.
- The unified architecture integrates multiprotocol support to enable both file-based and block-based storage on a single platform. With gateway systems, you can virtualize your entire storage infrastructure under one interface, and you can apply all the preceding efficiencies to your non-IBM N series systems.

Storage efficiency with SATA storage disks and the Performance Acceleration Module

SATA disks are low-cost, high-capacity storage solutions best suited for secondary storage applications, such as online backup and archiving. You can combine the Performance Acceleration Module with SATA disks to provide a low-cost alternative to Fibre Channel disks without sacrificing read performance.

The Performance Acceleration Module enables you to optimize the performance of mainstream storage platforms with workloads that are random read intensive, such as the file services. This intelligent read-cache technology helps to reduce latency and improve I/O throughput without adding more disk drives.

The Performance Acceleration Module works with deduplication to reduce the amount of redundant data it keeps in cache.

Protection against double disk failure with RAID-DP

RAID-DP is a standard Data ONTAP feature that protects your data from double-disk failures during reconstruction. Understanding how RAID-DP provides this protection can help you administer your storage systems more effectively.

RAID-DP offers protection against either two failed disks within the same RAID group or against a single-disk failure followed by a bad block or bit error before reconstruction is complete. This protection is equivalent to traditional hardware disk mirroring but requires 43 percent fewer spindles.

RAID-DP protection makes a SATA disk a viable option for your enterprise storage. You can use the less-expensive SATA disks without worrying about data loss and also lower your storage acquisition costs. For more information about RAID-DP, see the *Data ONTAP Storage Management Guide*.

Note:

- The time to reconstruct data from two failed disks is slightly less than the time to reconstruct data from a single-disk failure.
- It is highly likely that one disk failed before the other and at least some information has already been re-created with traditional row parity. RAID-DP automatically adjusts for this occurrence by starting recovery where two elements are missing from the second disk failure.

Next topics

[What RAID-DP protection is](#) on page 19

[How RAID-DP protection works](#) on page 20

Related information

[IBM NAS documentation and support site](http://www.ibm.com/storage/support/nas) — <http://www.ibm.com/storage/support/nas>

What RAID-DP protection is

If an aggregate is configured for RAID-DP protection, Data ONTAP reconstructs the data from one or two failed disks within a RAID group and transfers that reconstructed data to one or two spare disks as necessary.

RAID-DP provides double-parity disk protection when the following conditions occur:

- There is a single-disk or double-disk failure within a RAID group.
- There are media errors on a block when Data ONTAP is attempting to reconstruct a failed disk.

The minimum number of disks in a RAID-DP group is three: at least one data disk, one regular parity disk, and one double-parity (or dParity) disk.

If there is a data-disk or parity-disk failure in a RAID-DP group, Data ONTAP replaces the failed disk in the RAID group with a spare disk and uses the parity data to reconstruct the data of the failed disk on the replacement disk. If there is a double-disk failure, Data ONTAP replaces the failed disks in the RAID group with two spare disks and uses the double-parity data to reconstruct the data of the failed disks on the replacement disks.

RAID-DP is the default RAID type for all aggregates.

How RAID-DP protection works

RAID-DP employs both the traditional RAID4 horizontal parity and diagonal parity structures to significantly increase the data fault tolerance from failed disks.

RAID-DP protection adds a secondary parity disk to each RAID group in an aggregate or traditional volume. A RAID group is an underlying construct on which aggregates and traditional volumes are built. Each traditional RAID4 group has a minimum of four data disks and one parity disk; aggregates and volumes contain one or more RAID4 groups. Whereas the parity disk in a RAID4 volume stores row parity across the disks in a RAID4 group, the additional RAID-DP parity disk stores diagonal parity across the disks in a RAID-DP group. With these two parity stripes in RAID-DP, one the traditional horizontal and the other diagonal, data protection is obtained even if two disk failures occur in the same RAID group.

How space management works

The space management capabilities of Data ONTAP allow you to configure your storage systems to provide the storage availability required by the users and applications accessing the system, while using your available storage as effectively as possible.

Data ONTAP enables space management using the following capabilities:

- Space guarantees
- Space reservations
- Fractional reserve
- Automatic free space preservation

Next topics

[What kind of space management to use](#) on page 21

[What space guarantees are](#) on page 22

[What space reservation is](#) on page 23

[What fractional reserve is](#) on page 24

[Reasons to set fractional reserve to zero](#) on page 25

[How Data ONTAP can automatically provide more free space for full volumes](#) on page 26

What kind of space management to use

The type of space management you should use depends on many factors, including your tolerance for out-of-space errors, whether you plan to overcommit your aggregates, and your rate of data overwrite.

The following table can help you determine which space management capabilities best suit your requirements.

Note: LUNs in this context refer to the LUNs that Data ONTAP serves to clients, not to the array LUNs used for storage on a storage array.

If...	Then use...	Typical usage	Notes
You want management simplicity	FlexVol volumes with a space guarantee of volume OR Traditional volumes	NAS file systems	This is the easiest option to administer. As long as you have sufficient free space in the volume, writes to any file in this volume will always succeed.

If...	Then use...	Typical usage	Notes
Writes to certain files must always succeed You want to overcommit your aggregate	FlexVol volumes with a space guarantee of <code>file</code> OR Traditional volume AND space reservation enabled for files that require writes to succeed	LUNS Databases	This option enables you to guarantee writes to specific files.
You need even more effective storage usage than file space reservation provides You actively monitor available space on your volume and can take corrective action when needed Snapshot copies are short-lived Your rate of data overwrite is relatively predictable and low	FlexVol volumes with a space guarantee of <code>volume</code> OR Traditional volume AND Space reservation on for files that require writes to succeed AND Fractional reserve < 100%	LUNs (with active space monitoring) Databases (with active space monitoring)	With fractional reserve <100%, it is possible to use up all available space, even with space reservations on. Before enabling this option, be sure either that you can accept failed writes or that you have correctly calculated and anticipated storage and Snapshot copy usage.
You want to overcommit your aggregate You actively monitor available space on your aggregate and can take corrective action when needed	FlexVol volumes with a space guarantee of <code>none</code>	Storage providers who need to provide storage that they know will not immediately be used Storage providers who need to allow available space to be dynamically shared between volumes	With an overcommitted aggregate, writes can fail due to insufficient space.

What space guarantees are

Space guarantees on a FlexVol volume ensure that writes to a specified FlexVol volume or writes to files with space reservations enabled do not fail because of lack of available space in the containing aggregate.

Space guarantee is an attribute of the volume. It is persistent across storage system reboots, takeovers, and givebacks. Space guarantee values can be `volume` (the default value), `file`, or `none`.

- A space guarantee of `volume` reserves space in the aggregate for the volume. The reserved space cannot be allocated to any other volume in that aggregate.

The space management for a FlexVol volume that has a space guarantee of volume is equivalent to a traditional volume.

- A space guarantee of file reserves space in the aggregate so that any file in the volume with space reservation enabled can be completely rewritten, even if its blocks are being retained on disk by a Snapshot copy.
- A FlexVol volume that has a space guarantee of none reserves no extra space for user data; writes to LUNs or files contained by that volume could fail if the containing aggregate does not have enough available space to accommodate the write.

Note: Because out-of-space errors are unexpected in a CIFS environment, do not set space guarantee to none for volumes accessed using CIFS.

When space in the aggregate is reserved for space guarantee for an existing volume, that space is no longer considered free space. Operations that consume free space in the aggregate, such as creation of Snapshot copies or creation of new volumes in the containing aggregate, can occur only if there is enough available free space in that aggregate; these operations are prevented from using space already committed to another volume.

When the uncommitted space in an aggregate is exhausted, only writes to volumes or files in that aggregate with space guarantees are guaranteed to succeed.

Note: Space guarantees are honored only for online volumes. If you take a volume offline, any committed but unused space for that volume becomes available for other volumes in that aggregate. When you bring that volume back online, if there is not sufficient available space in the aggregate to fulfill its space guarantees, you must use the force (-f) option, and the volume's space guarantees are disabled. When a volume's space guarantee is disabled, the word (disabled) appears next to its space guarantees in the output of the `vol status` command.

What space reservation is

When space reservation is enabled for one or more files or LUNs, Data ONTAP reserves enough space in the volume (traditional or FlexVol) so that writes to those files or LUNs do not fail because of a lack of disk space.

Note: LUNs in this context refer to the LUNs that Data ONTAP serves to clients, not to the array LUNs used for storage on a storage array.

For example, if you create a 100-GB space reserved LUN in a 500-GB volume, that 100 GB of space is immediately allocated, leaving 400 GB remaining in the volume. In contrast, if space reservation is disabled on the LUN, all 500 GB in the volume remain available until writes are made to the LUN.

Space reservation is an attribute of the file or LUN; it is persistent across storage system reboots, takeovers, and givebacks. Space reservation is enabled for new LUNs by default, but you can create a LUN with space reservations disabled or enabled. After you create the LUN, you can change the space reservation attribute by using the `lun set reservation` command. You can change the space reservation for files by using the `file reservation` command.

When a volume contains one or more files or LUNs with space reservation enabled, operations that require free space, such as the creation of Snapshot copies, are prevented from using the reserved space. If these operations do not have sufficient unreserved free space, they fail. However, writes to the files or LUNs with space reservation enabled will continue to succeed.

What fractional reserve is

Fractional reserve is a volume option that enables you to determine how much space Data ONTAP reserves for Snapshot copy overwrites for LUNs, as well as for space-reserved files when all other space in the volume is used.

The fractional reserve setting defaults to 100%, but you can use the `vol options` command to set fractional reserve to any percentage from zero to 100.

It is best to use the autodelete function, but there may occasionally be circumstances under which fractional reserve can be used, including:

- When Snapshot copies cannot be deleted
- When preserving existing Snapshot copies is more important than creating new ones

Fractional reserve can be used on the following types of volumes:

- Traditional volumes
- FlexVol volumes with a space guarantee of `volume`

Note: If the `guarantee` option for a FlexVol volume is set to `file`, then fractional reserve for that volume is set to 100 percent and is not adjustable. If the `guarantee` option for a FlexVol volume is set to `none`, then fractional reserve for that volume can be set to the desired value. For the vast majority of configurations, you should set fractional reserve to zero when the `guarantee` option is set to `none` because it greatly simplifies space management.

If fractional reserve is set to 100%, when you create space-reserved LUNs, you can be sure that writes to those LUNs will always succeed without deleting Snapshot copies, even if all of the space-reserved LUNs are completely overwritten.

Setting fractional reserve to less than 100 percent causes the space reservation held for all space-reserved LUNs in that volume to be reduced to that percentage. Writes to the space-reserved LUNs in that volume are no longer unequivocally guaranteed, which is why you should use `snap autodelete` or `vol autogrow` for these volumes.

Fractional reserve is generally used for volumes that hold LUNs with a small percentage of data overwrite.

Note: If you are using fractional reserve in environments in which write errors due to lack of available space are unexpected, you must monitor your free space and take corrective action to avoid write errors. Data ONTAP provides tools for monitoring available space in your volumes.

Note: Reducing the space reserved for overwrites (by using fractional reserve) does not affect the size of the space-reserved LUN. You can write data to the entire size of the LUN. The space reserved for overwrites is used only when the original data is overwritten.

Example

If you create a 500-GB space-reserved LUN, then Data ONTAP ensures that 500 GB of free space always remains available for that LUN to handle writes to the LUN.

If you then set fractional reserve to 50 for the LUN's containing volume, then Data ONTAP reserves 250 GB, or half of the space it was previously reserving for overwrites with fractional reserve set to 100. If more than half of the LUN is overwritten, then subsequent writes to the LUN could fail due to insufficient free space in the volume.

Note: When more than one LUN in the same volume has space reservations enabled, and fractional reserve for that volume is set to less than 100 percent, Data ONTAP does not limit any space-reserved LUN to its percentage of the reserved space. In other words, if you have two 100-GB LUNs in the same volume with fractional reserve set to 30, one of the LUNs could use up the entire 60 GB of reserved space for that volume.

See the *Technical Report* on thin provisioning for detailed information on using fractional reserve.

Note: This technical report contains information about NetApp products that IBM licenses and in some cases customizes. Technical reports might contain information about models and features that are not supported by IBM.

Related information

Technical Report: Thin Provisioning in a NetApp SAN or IP SAN Enterprise Environment - <http://media.netapp.com/documents/tr3483.pdf>

Reasons to set fractional reserve to zero

Setting fractional reserve on a volume to zero and managing your free space manually, rather than having Data ONTAP manage it for you, gives you more control over space utilization.

You can maximize your space utilization by exercising more control over how much space is allocated for overwrites. The cost is extra management time and the risk of encountering out-of-space errors on a space-reserved file or LUN.

For example, setting fractional reserve to zero might be helpful for a Qtree SnapMirror destination of a space-reserved LUN. (For Volume SnapMirror, the fractional reserve of the target is always the same as the source.) Depending on the rate of change (ROC) of your data, how much free space is available in the volume, and whether you are taking snapshots of the volume containing the target, it might make sense to reduce the fractional reserve to zero for the target.

How Data ONTAP can automatically provide more free space for full volumes

Data ONTAP can automatically make more free space available for a FlexVol volume when that volume is nearly full. You can choose to make the space available by first allowing the volume size to increase, or by first deleting Snapshot copies.

You enable this capability for a FlexVol volume by using the `vol options` command with the `try_first` option.

Data ONTAP can automatically provide more free space for the volume by using one of the following methods:

- Increase the size of the volume when it is nearly full.
This method is useful if the volume's containing aggregate has enough space to support a larger volume. You can increase the size in increments and set a maximum size for the volume.
- Delete Snapshot copies when the volume is nearly full.
For example, you can automatically delete Snapshot copies that are not linked to Snapshot copies in cloned volumes or LUNs, or you can define which Snapshot copies you want to delete first—your oldest or newest Snapshot copies. You can also determine when to begin deleting Snapshot copies—for example, when the volume is nearly full or when the volume's Snapshot reserve is nearly full.

You can choose which method (increasing the size of the volume or deleting Snapshot copies) you want Data ONTAP to try first. If the first method does not provide sufficient extra free space to the volume, Data ONTAP will try the other method next.

Thin provisioning for greater efficiencies using FlexVol volumes

With thin provisioning, when you create volumes for different purposes in a given aggregate, you do not actually allocate any space for those volumes in advance. The space is allocated only when the application host needs it.

The unused aggregate space is available for the thinly provisioned volumes to expand or for creating new volumes. By allowing as-needed provisioning and space reclamation, thin provisioning can improve storage utilization and decrease storage costs.


A FlexVol volume can share its containing aggregate with other FlexVol volumes. Therefore, a single aggregate is the shared source of all the storage used by the FlexVol volumes it contains. Flexible volumes are no longer bound by the limitations of the disks on which they reside. A FlexVol volume is a pool of storage that can be sized based on how much data you want to store in it, rather than on the size of your disk. This flexibility enables you to maximize the performance and capacity utilization of the storage systems. Because FlexVol volumes can access all available physical storage in the system, dramatic improvements in storage utilization are possible.

Example

A 500-GB volume is allocated with only 100 GB of actual data; the remaining 400 GB allocated has no data stored in it. This unused capacity is assigned to a business application, even though the application might not need all 500 GB until later. The allocated but unused 400 GB of excess capacity is temporarily wasted.

With thin provisioning, the storage administrator provisions 500 GB to the business application but uses only 100 GB for the data. The difference is that with thin provisioning, the unused 400 GB is still available to other applications. This approach allows the application to grow transparently, and the physical storage is fully allocated only when the application truly needs it. The rest of the storage remains in the free pool to be used as needed. Storage administrators can set thresholds, so they are alerted when more disks need to be added to the pool.

The following illustration compares thin provisioning with traditional provisioning:

**Next topics**

[Benefits of using thin provisioning](#) on page 28

[Storage space management using DataFabric Manager](#) on page 29

Benefits of using thin provisioning

The FlexVol technology enables you to oversubscribe the free space to adapt rapidly to the changing business needs.

The benefits of using thin provisioning are as follows:

- Allows storage to be provisioned just like traditional storage, but it is not consumed until data is written.
- Storage-provisioning time is greatly reduced, because you can create the storage for an application quickly without depending on the actual physical space available.
- Through notifications and configurable threshold values, you can plan your procurement strategies well in advance and have enough storage for thin provisioned volumes to grow.
- You can set aggregate overcommitment thresholds by using Protection Manager. Using Provisioning Manager, you can also set policies for provisioning, exporting, and managing your space requirements. For more information about aggregate overcommitment threshold values and provisioning policies, see the *Provisioning Manager and Protection Manager Administration Guide*.

Related information

[IBM NAS documentation and support site](http://www.ibm.com/storage/support/nas) — <http://www.ibm.com/storage/support/nas>

Storage space management using DataFabric Manager

Provisioning Manager and Operations Manager are part of an add-on management suite that integrates thin provisioning for cross-system configurations.

The Provisioning Manager provides the following capabilities:

- Provisioning policies that manage provisioning and exporting of storage.
- Automatic provisioning of a dataset when you assign a provisioning policy to it.
- Periodic checking of provisioned storage for conformance to the policy.
- Manual controls for resizing volumes and for deleting old Snapshot copies on existing storage and newly provisioned storage.
- Migration of datasets and vFiler units to new storage systems.
- Deduplication to eliminate duplicate data blocks to reduce storage space.

You can use Operations Manager for the following day-to-day activities on storage systems:

- Monitor the device or the object health, the capacity utilization, and the performance characteristics of a storage system.
- View or export reports.
- Configure alerts and thresholds for event management.
- Group devices, vFiler units, host agents, volumes, qtrees, and LUNs

For more information about Provisioning Manager and Operations Manager, see *Provisioning Manager and Protection Manager Printable Help* and *Operations Manager Administration Guide*.

Related information

IBM NAS documentation and support site — <http://www.ibm.com/storage/support/nas>

Automate thin provisioning using Provisioning Manager

With Provisioning Manager, you can automate thin provisioning and eliminate the need to provision extra storage space. You can use resource pools for more efficient aggregation of resources.

Provisioning Manager enables you to take advantage of thin provisioning, and resource pooling to get the highest possible level of storage efficiency from your storage resources. You can pool your resources based on attributes such as performance, cost, physical location, or availability.

By grouping related resources into a pool, you can treat the pool as a single unit for monitoring, provisioning, reporting, and role-based access control (RBAC). This approach simplifies the management of these resources and allows for a more flexible and efficient use of the storage.

For more information about Provisioning Manager, see the *Provisioning Manager and Protection Manager Printable Help*.

Related information

IBM NAS documentation and support site - <http://www.ibm.com/storage/support/nas>

Efficient file recovery with Snapshot copies

Snapshot copies consume minimal storage space because only changes to the active file systems are written. You can back up your volumes more frequently.

A Snapshot copy is a read-only image of a traditional volume, a FlexVol volume, or an aggregate that captures the state of the file system at a point in time. Two Snapshot copies created in sequence differ only by the blocks added or changed in the time interval between the two. This block incremental behavior limits the consumption of associated storage capacity.

When a Snapshot copy is created, the data is not copied from one location to another. As a result, you can store up to 255 Snapshot copies at one time on each volume without any performance overhead.

Data ONTAP provides a default Snapshot copy schedule for each volume. You can configure the schedule to fit your needs. This schedule creates Snapshot copies automatically and deletes old Snapshot copies after a predetermined amount of time.

For more information about Snapshot copies, see the *Data ONTAP Data Protection Online Backup and Recovery Guide*.

Next topics

[What Snapshot disk consumption is](#) on page 31

[How Snapshot copies consume disk space](#) on page 31

[What the Snapshot copy reserve is](#) on page 32

[What file folding means and how it saves disk space](#) on page 35

Related information

[IBM NAS documentation and support site](http://www.ibm.com/storage/support/nas) — <http://www.ibm.com/storage/support/nas>

What Snapshot disk consumption is

Data ONTAP preserves pointers to all the disk blocks currently in use at the time the Snapshot copy is created. When a file is changed, the Snapshot copy still points to the disk blocks where the file existed before it was modified, and changes are written to new disk blocks.

How Snapshot copies consume disk space

Snapshot copies minimize disk consumption by preserving individual blocks rather than whole files. Snapshot copies begin to consume extra space only when files in the active file system are changed or deleted. When this happens, the original file blocks are still preserved as part of one or more Snapshot copies.

In the active file system the changed blocks are rewritten to different locations on the disk or removed as active file blocks entirely. As a result, in addition to the disk space used by blocks in the

modified active file system, disk space used by the original blocks is still reserved to reflect the status of the active file system before the change.

The following illustration shows disk space usage for a Snapshot copy.




Figure 2: How Snapshot copies consume disk space before and after you delete a file *myfile.txt*.

What the Snapshot copy reserve is

The Snapshot copy reserve sets a specific percent of disk space for Snapshot copies. By default, the Snapshot copy reserve is 20 percent of disk space. The Snapshot copy reserve can be used only by Snapshot copies, not by the active file system.

If the active file system runs out of disk space, any disk space still remaining in the Snapshot copy reserve is not available for use by the active file system.

Note: Although the active file system cannot consume disk space reserved for Snapshot copies, Snapshot copies can exceed the Snapshot copy reserve and consume disk space normally available to the active file system.

Managing the Snapshot copy reserve involves the following tasks:

- Ensuring that enough disk space is set aside for Snapshot copies so that they do not consume active file system space

- Keeping disk space consumed by Snapshot copies below the Snapshot copy reserve
- Ensuring that the Snapshot copy reserve is not so large that it wastes space that could be used by the active file system

Next topics

[Use of deleted active file disk space](#) on page 33

[Snapshot copies can exceed reserve](#) on page 34

[Recovery of disk space for file system use](#) on page 34

Use of deleted active file disk space

When enough disk space is available for Snapshot copies in the Snapshot copy reserve, deleting files in the active file system frees disk space for new files, while the Snapshot copies that reference those files consume only the space in the Snapshot copy reserve.

About this task

If Data ONTAP created a Snapshot copy when the disks were full, deleting files from the active file system does not create any free space because everything in the active file system is also referenced by the newly created Snapshot copy. Data ONTAP has to delete the Snapshot copy before it can create any new files.

Example

The following example shows how disk space being freed by deleting files in the active file system ends up in the Snapshot copy:

If Data ONTAP creates a Snapshot copy when the active file system is full and there is still space remaining in the Snapshot reserve, the output from the `df` command—which displays statistics about the amount of disk space on a volume—is as follows:

Filesystem	kbytes	used	avail	capacity
/vol/vol0/	3000000	3000000	0	100%
/vol/vol0/.snapshot	1000000	500000	500000	50%

If you delete 100,000 KB (0.1 GB) of files, the disk space used by these files is no longer part of the active file system, so the space is reassigned to the Snapshot copies instead.

Data ONTAP reassigns 100,000 KB (0.1 GB) of space from the active file system to the Snapshot reserve. Because there was reserve space for Snapshot copies, deleting files from the active file system freed space for new files. If you enter the `df` command again, the output is as follows:

Filesystem	kbytes	used	avail	capacity
/vol/vol0/	3000000	2900000	100000	97%
/vol/vol0/.snapshot	1000000	600000	400000	60%

Snapshot copies can exceed reserve

There is no way to prevent Snapshot copies from consuming disk space greater than the amount reserved for them; therefore, it is important to reserve enough disk space for Snapshot copies so that the active file system always has space available to create new files or modify existing ones.

Example

Consider what happens in the following example if all files in the active file system are deleted. Before the deletion, the `df` output is as follows:

Filesystem	kbytes	used	avail	capacity
/vol/vol0/	3000000	3000000	0	100%
/vol/vol0/.snapshot	1000000	500000	500000	50%

After the deletion, the `df` command generates the following output:

Filesystem	kbytes	used	avail	capacity
/vol/vol0/	3000000	2500000	500000	83%
/vol/vol0/.snapshot	1000000	3500000	0	350%

The entire 3,000,000 KB (3 GB) in the active file system is still being used by Snapshot copies, along with the 500,000 KB (0.5 GB) that was being used by Snapshot copies before, making a total of 3,500,000 KB (3.5 GB) of Snapshot copy data. This is 2,500,000 KB (2.5 GB) more than the space reserved for Snapshot copies; therefore, 2.5 GB of space that would be available to the active file system is now unavailable to it. The post-deletion output of the `df` command lists this unavailable space as used even though no files are stored in the active file system.

Recovery of disk space for file system use

Whenever Snapshot copies consume more than 100% of the Snapshot reserve, the system is in danger of becoming full. In this case, you can create files only after you delete enough Snapshot copies.

Example

If 500,000 KB (0.5 GB) of data is added to the active file system, a `df` command generates the following output:

Filesystem	kbytes	used	avail	capacity
/vol/vol0	3000000	3000000	0	100%
/vol/vol0/.snapshot	1000000	3500000	0	350%

As soon as Data ONTAP creates a new Snapshot copy, every disk block in the file system is referenced by some Snapshot copy. Therefore, no matter how many files you delete from the active file system, there is still no room to add any more. The only way to recover from this situation is to delete enough Snapshot copies to free more disk space.

What file folding means and how it saves disk space

File folding describes the process of checking the data in the most recent Snapshot copy, and if this data is identical to the Snapshot copy currently being created, by referencing the previous Snapshot copy instead of taking up disk space writing the same data in the new Snapshot copy.

File folding saves disk space by sharing unchanged file blocks between the active version of the file and the version of the file in the latest Snapshot copy, if any.

The system must compare block contents when folding a file, so file folding might affect system performance.

If the folding process reaches a maximum limit on memory usage, it is suspended. When memory usage falls below the limit, the processes that were halted are restarted.

Space savings with deduplication

Deduplication is an optional feature of Data ONTAP that significantly improves physical storage space by eliminating duplicate data blocks within a FlexVol volume.

Deduplication works at the block level on the active file system, and uses the WAFL block-sharing mechanism. Each block of data has a digital signature that is compared with all other signatures in a data volume. If an exact block match exists, the duplicate block is discarded and its disk space is reclaimed.

You can configure deduplication operations to run automatically or on a schedule. You can deduplicate new and existing data, or only new data, on a FlexVol volume.

To use deduplication, you must install the deduplication licenses.

Deduplication removes data redundancies, as shown in the following illustration:




Figure 3: How deduplication removes data redundancies

Next topics

[How deduplication works](#) on page 38

[What deduplication metadata is](#) on page 38

[Guidelines for using deduplication](#) on page 39

[Performance considerations for deduplication](#) on page 39

[How deduplication works with other features and products](#) on page 40

How deduplication works

Deduplication operates at the block level within the entire FlexVol volume, eliminating duplicate data blocks and storing only unique data blocks.

Data ONTAP writes all data to a storage system in 4-KB blocks. When deduplication runs for the first time on a FlexVol volume with existing data, it scans all the blocks in the FlexVol volume and creates a digital fingerprint for each of the blocks. Each of the fingerprints is compared to all other fingerprints within the FlexVol volume. If two fingerprints are found to be identical, a byte-for-byte comparison is done for all data within the block. If the byte-for-byte comparison detects identical fingerprints, the pointer to the data block is updated, and the duplicate block is freed.

Deduplication runs on the active file system. Therefore, as additional data is written to the deduplicated volume, fingerprints are created for each new block and written to a change log file. For subsequent deduplication operations, the change log is sorted and merged with the fingerprint file, and the deduplication operation continues with fingerprint comparisons as previously described.

What deduplication metadata is

Deduplication uses fingerprints, which are digital signatures for every 4-KB data block in a FlexVol volume. The fingerprint database and the change logs form the deduplication metadata.

In Data ONTAP 7.3 and later, the fingerprint database and the change logs used by the deduplication operation are located outside the volume and in the aggregate. Therefore, the deduplication metadata is not included in the FlexVol volume Snapshot copies.

This approach enables deduplication to achieve higher space savings than in Data ONTAP 7.2. However, some of the temporary metadata files created during the deduplication operation are still placed inside the volume and are deleted only after the deduplication operation is complete. The temporary metadata files, which are created during a deduplication operation, can be locked in the Snapshot copies. These temporary metadata files remain locked until the Snapshot copies are deleted.

While deduplication can provide substantial space savings, a percentage of storage overhead is associated with it, which you should consider when sizing a FlexVol volume.

The deduplication metadata can occupy up to 6 percent of the total logical data of the volume, as follows:

- Up to 2 percent of the total logical data of the volume is placed inside the volume.
- Up to 4 percent of the total logical data of the volume is placed in the aggregate.

Guidelines for using deduplication

You must remember certain guidelines about system resources and free space when using deduplication.

The guidelines are as follows:

- Deduplication is a background process that consumes system resources while it is running. If the data does not change very often in a FlexVol volume, it is best to run deduplication less frequently. Multiple concurrent deduplication operations running on a storage system lead to a higher consumption of system resources.
- You must ensure that sufficient free space exists for deduplication metadata in the volumes and aggregates.
- You cannot increase the size of a volume that contains deduplicated data beyond the maximum supported size limit, either manually or by using the autogrow option.
- You cannot enable deduplication on a volume if it is larger than the maximum volume size. However, starting with Data ONTAP 7.3.1, you can enable deduplication on a volume after reducing its size within the supported size limits.

Performance considerations for deduplication

Certain factors affect the performance of deduplication. You should check the performance impact of deduplication in a test setup, including sizing considerations, before deploying deduplication in performance-sensitive or production environments.

The following factors affect the performance of deduplication:

- Application and the type of data used
- The data access pattern (for example, sequential versus random access, the size and pattern of the input and output)
- The amount of duplicate data, the amount of total data, and the average file size
- The nature of data layout in the volume
- The amount of changed data between deduplication operations
- The number of concurrent deduplication operations
- Hardware platform (system memory and CPU module)
- Load on the system (for example, MBps)
- Disk types (for example, ATA/FC, and RPM of the disk)

How deduplication works with other features and products

You must keep certain considerations in mind when using deduplication with other features.

Next topics

[*Deduplication and Snapshot copies*](#) on page 40

[*Deduplication and volume SnapMirror*](#) on page 41

[*Deduplication and qtree SnapMirror*](#) on page 42

[*Deduplication and SnapVault*](#) on page 42

[*Deduplication and SnapRestore*](#) on page 43

[*Deduplication and DataFabric Manager 3.8*](#) on page 43

[*Deduplication and volume copy*](#) on page 44

[*Deduplication and FlexClone volumes*](#) on page 45

[*Deduplication and an active/active configuration*](#) on page 45

[*Deduplication and VMware*](#) on page 46

[*Deduplication and MultiStore*](#) on page 46

Deduplication and Snapshot copies

You can run deduplication only on the active file system. However, this data can get locked in Snapshot copies created before you run deduplication, resulting in reduced space savings.

Data can get locked in Snapshot copies in two ways:

- If the Snapshot copies are created before the deduplication operation is run.
You can avoid this situation by always running deduplication before Snapshot copies are created.
- When the Snapshot copy is created, a part of the deduplication metadata resides in the volume and the rest of the metadata resides in the aggregate outside the volume. The fingerprint files and the change-log files that are created during the deduplication operation are placed in the aggregate and are not captured in Snapshot copies, which results in higher space savings. However, some temporary metadata files that are created during a deduplication operation are still placed inside the FlexVol volume; these files are deleted after the deduplication operation is complete. These temporary metadata files can get locked in Snapshot copies if the copies are created during a deduplication operation. The metadata remains locked until the Snapshot copies are deleted.

To avoid conflicts between deduplication and Snapshot copies, you should follow these guidelines:

- Run deduplication before creating new Snapshot copies.
- Remove unnecessary Snapshot copies stored in deduplicated volumes.
- Reduce the retention time of Snapshot copies stored in deduplicated volumes.
- Schedule deduplication only after significant new data has been written to the volume.
- Configure appropriate reserve space for the Snapshot copies.

- If snap reserve is 0, you should turn off the schedule for automatic creation of Snapshot copies (which is the case in most LUN deployments).

Deduplication and volume SnapMirror

You can use volume SnapMirror to replicate a deduplicated volume.

When using volume SnapMirror with deduplication, you must consider the following information:

- You need to enable both the deduplication and SnapMirror licenses.
- You can enable deduplication on the source system, the destination system, or both systems. The deduplication and the near-line functionality licenses must be enabled on the primary storage system (source). Starting with Data ONTAP 7.3.1, neither the near-line functionality license nor the deduplication license is needed on the destination storage system.

Note: A deduplication license is not required on the destination storage system. However, if the primary storage system is not available and the secondary storage system becomes the new primary, deduplication must be licensed on the secondary storage system for deduplication to continue. Therefore, you might want to license deduplication on both storage systems.

You can enable, run, and manage deduplication only from the primary storage system. However, the FlexVol volume in the secondary storage system inherits all the deduplication attributes and storage savings through SnapMirror.

- The shared blocks are transferred only once. Therefore, deduplication also reduces the use of network bandwidth.
- If the source and destination volumes are on different storage system models, they might have different maximum volume sizes. The lower maximum applies.

When creating a SnapMirror relationship between two different storage system models, you should ensure that the maximum volume size with deduplication is set to the lower maximum volume size limit of the two models.

For example, an N5300 system supports a maximum volume size with deduplication of 4 TB, and an N5600 system supports 16 TB. When establishing a volume SnapMirror relationship between the N5300 and N5600 systems, you should ensure that the volume SnapMirror relationship is established for 4 TB. After a failover, if the volume size on the N5600 system is increased to more than 4 TB, you will not be able to perform a new baseline transfer or resynchronize the storage systems, because the N5300 system has a maximum volume size with deduplication of 4 TB.

- The volume SnapMirror update schedule does not depend on the deduplication schedule.

When configuring volume SnapMirror and deduplication, you should coordinate the deduplication schedule and the volume SnapMirror schedule. You should start volume SnapMirror transfers of a deduplicated volume after the deduplication operation is complete. This schedule prevents the sending of undeduplicated data and additional temporary metadata files over the network. If the temporary metadata files in the source volume are locked in Snapshot copies, these files consume extra space in the source and destination volumes.

Starting with Data ONTAP 7.3.1, volumes whose size has been reduced to within the limit supported by deduplication can be part of the SnapMirror primary storage system and the secondary storage system.

Deduplication and qtree SnapMirror

You can use deduplication for volumes that use qtree SnapMirror.

In Data ONTAP 7.3 and later, deduplication operations are supported with qtree SnapMirror. Qtree SnapMirror does not automatically initiate a deduplication operation at the completion of every individual qtree SnapMirror transfer. You can set up a deduplication schedule independent of your qtree SnapMirror transfer schedule.

When using qtree SnapMirror with deduplication, you must consider the following information:

- You need to enable both the deduplication and SnapMirror licenses.
Note: You can enable deduplication on the source system, the destination system, or both systems.
- Even when deduplication is enabled on the source system, duplicate blocks are sent to the destination system. Therefore, no network bandwidth savings is achieved.
- To recognize space savings on the destination system, you should run deduplication on the destination after the qtree SnapMirror transfer is complete.
- You can set up a deduplication schedule independently of the qtree SnapMirror schedule. For example, on the destination system, the deduplication process does not start automatically after qtree SnapMirror transfers are finished.
- Qtree SnapMirror recognizes deduplicated blocks as changed blocks. Therefore, when you run deduplication on an existing qtree SnapMirror source system for the first time, all the deduplicated blocks are transferred to the destination system. This process might result in a transfer several times larger than the regular transfers.

When using qtree SnapMirror with deduplication, you should ensure that qtree SnapMirror uses only the minimum number of Snapshot copies that it requires. To ensure this minimum, you should retain only the latest Snapshot copies.

Deduplication and SnapVault

The deduplication feature is integrated with the SnapVault secondary license. This feature increases the efficiency of data backup and improves the use of secondary storage.

The behavior of deduplication with SnapVault is similar to the behavior of deduplication with qtree SnapMirror, with the following exceptions:

- Deduplication is also supported on the SnapVault destination volume.
- The deduplication schedule depends on the SnapVault update schedule on the destination system. However, the deduplication schedule on the source system does not depend on the SnapVault update schedule, and it can be configured independently on a volume.

- Every SnapVault update (baseline or incremental) starts a deduplication process on the destination system after the archival Snapshot copy is taken.
- A new Snapshot copy replaces the archival Snapshot copy after deduplication finishes running on the destination system. (The name of this new Snapshot copy is the same as that of the archival copy, but the Snapshot copy uses a new timestamp, which is the creation time.)
- You cannot configure the deduplication schedule on the destination system manually or run the `sis start` command. However, you can run the `sis start -s` command on the destination system.
- The SnapVault update does not depend on the deduplication operation. A subsequent incremental update is allowed to continue while the deduplication operation on the destination volume from the previous backup is still in progress. In this case, the deduplication operation continues; however, the archival Snapshot copy is not replaced after the deduplication operation is complete.
- The SnapVault update recognizes the deduplicated blocks as changed blocks. Thus, when deduplication is run on an existing SnapVault source for the first time, all saved space is transferred to the destination system. The size of the transfer might be several times larger than the regular transfers. Running deduplication on the source system periodically will help prevent this issue for future qtree SnapMirror transfers. You should run deduplication before the SnapVault baseline transfer.

Note: You can run a maximum of eight concurrent deduplication operations on a system. This number includes the deduplication operations linked to SnapVault volumes and those that are not linked to SnapVault volumes. The N5500 and N3400 platforms support a maximum of five concurrent deduplication operations.

Deduplication and SnapRestore

Starting with Data ONTAP 7.3, the metadata created during a deduplication operation is located in the aggregate. Therefore, when you initiate a SnapRestore operation on a volume, the metadata is not restored to the active file system. The restored data, however, retains the original space savings.

After a SnapRestore operation, if deduplication is enabled on the volume, any new data written to the volume continues to be deduplicated. However, space savings is obtained for only the new data.

To run deduplication for all the data on the volume, you should use the `sis start -s` command.

This command builds the fingerprint database for all the data in the volume. The amount of time this process takes depends on the size of the logical data in the volume. Before using the `sis start -s` command, you must ensure that the volume and the aggregate containing the volume have sufficient free space for the deduplication metadata.

Deduplication and DataFabric Manager 3.8

Starting with Data ONTAP 7.3.1, deduplication is supported with Protection Manager, Provisioning Manager, and Operations Manager in DataFabric Manager 3.8.

Deduplication and Protection Manager in DataFabric Manager 3.8

In releases earlier than DataFabric Manager 3.8, Protection Manager waits for an active deduplication operation to complete, before renaming the Snapshot copies. While Protection

Manager waits, it does not allow clients to list the Snapshot copies or restore from them. Therefore, in releases prior to DataFabric Manager 3.8, the use of deduplication with Protection Manager is not optimal.

However, this limitation is removed in DataFabric Manager 3.8.

Deduplication and Provisioning Manager in DataFabric Manager 3.8

With Provisioning Manager in DataFabric Manager 3.8, you can enable the provisioning policies to support all three modes of deduplication, namely, on-demand deduplication, automated deduplication, and scheduled deduplication.

For more information about using deduplication with Provisioning Manager and Protection Manager, see the *Provisioning Manager and Protection Manager Administration Guide*.

Deduplication and Operations Manager in DataFabric Manager 3.8

Deduplication is integrated into Operations Manager in DataFabric Manager 3.8.

You can configure deduplication on the system and generate reports or graphs summarizing space savings for file and LUN clones.

For more information about using deduplication with Operations Manager, see the *Operations Manager Administration Guide*.

Related information

IBM NAS documentation and support site - <http://www.ibm.com/storage/support/nas>

Deduplication and volume copy

Volume copy is a method of copying both data in the active file system and data in storage systems from one volume to another. The source and destination volumes must both be FlexVol volumes.

When deduplicated data is copied by using the `vol copy` command, the copy of the data at the destination inherits all the deduplication attributes and storage savings of the source data.

Starting with Data ONTAP 7.3, the metadata created during a deduplication operation (fingerprint files and changelog files) are located outside the volume in the aggregate. Therefore, when you run the volume copy operation on a volume, the fingerprint files and change-log files are not restored to the active file system. After a volume copy operation, if deduplication is enabled on the volume, any new data written to the volume continues to be deduplicated. However, space savings is only obtained for the new data.

To run deduplication for all the data on the volume, you should use the `sis start -s` command.

This command builds the fingerprint database for all the data in the volume. The amount of time this process takes depends on the size of the logical data in the volume. Before using the `sis start -s` command, you must ensure that the volume and the aggregate containing the volume have sufficient free space for deduplication metadata.

Deduplication and FlexClone volumes

Deduplication is supported on FlexClone volumes. FlexClone volumes are writable clones of a parent FlexVol volume.

The FlexClone volume of a deduplicated volume is a deduplicated volume. The cloned volume inherits the deduplication configuration of the parent volume (for example, deduplication schedules).

The FlexClone volume of a non-deduplicated volume is a non-deduplicated volume. If you run deduplication on a clone volume, the clone is deduplicated, but the original volume remains non-deduplicated.

Starting with Data ONTAP 7.3, the metadata created during a deduplication operation (fingerprint files and change-log files) are located outside the volume in the aggregate; therefore, they are not cloned. However, the data retains the space savings of the original data.

Any new data written to the destination volume continues to be deduplicated and fingerprint files for the new data are created. Space savings is only obtained for the new data.

To run deduplication for all the data on the cloned volume, you should use the `sis start -s` command. The time the process takes to finish depends on the size of the logical data in the volume.

When a cloned volume is split from the parent volume, deduplication of all data in the clone that was part of the parent volume is undone after the volume-split operation. However, if deduplication is running on the clone volume, the data is deduplicated in the subsequent deduplication operation.

Deduplication and an active/active configuration

You can activate deduplication in an active/active configuration.

The maximum number of concurrent deduplication operations allowed on each node of an active/active configuration is eight. If one of the nodes fails, the other node takes over the operations of the failed node. In takeover mode, the working node continues with its deduplication operations as usual. However, the working node does not start any deduplication operations on the failed node.

Note:

- The N5500 and N3400 platforms support a maximum of five concurrent deduplication operations.
- Change logging for volumes with deduplication continues for the failed node in takeover mode. Therefore, deduplication operations can be performed on data written during takeover mode after the failed node is active, and there is no loss in space savings. To disable change logging for volumes that belong to a failed node, you can turn off deduplication on those volumes. You can also view the status of volumes with deduplication for a failed node in takeover mode.

Deduplication and VMware

You can run deduplication in VMware environments for efficient space savings.

While planning the Virtual Machine Disk (VMDK) and data store layouts, you should follow these guidelines.

- Operating system VMDKs deduplicate efficiently because the binary files, patches, and drivers are highly redundant between virtual machines. You can achieve maximum savings by keeping these VMDKs in the same volume.
- Application binary VMDKs deduplicate to varying degrees. Applications from the same vendor commonly have similar libraries installed; therefore, you can achieve moderate deduplication savings. Applications written by different vendors do not deduplicate at all.
- Application datasets when deduplicated have varying levels of space savings and performance impact based on the application and intended use. You should carefully consider what application data needs to be deduplicated.
- Transient and temporary data, such as VM swap files, pagefiles, and user and system temp directories, does not deduplicate well and potentially adds significant performance impact when deduplicated. Therefore, it is best to keep this data on a separate VMDK and volume that are not deduplicated.

Application data has a major effect on the percentage of storage savings achieved with deduplication. New installations typically achieve large deduplication savings.

Note: In VMware environments, proper partitioning and alignment of the VMDKs is important. Applications whose performance is impacted by deduplication operations are likely to have the same performance impact when you run deduplication in a VMware environment.

Deduplication and MultiStore

Deduplication commands are available in all the `vfiler` contexts. Deduplication support on vFiler units allows users to reduce redundant data blocks within vFiler units.

You can enable deduplication only on FlexVol volumes in a vFiler unit. Deduplication support on vFiler units ensures that volumes owned by a vFiler unit are not accessible to another vFiler unit. Deduplication also supports disaster recovery and migration of vFiler units. If you enable deduplication on the volume in the source vFiler unit, the destination vFiler unit inherits all deduplication attributes.

You must license deduplication on the primary storage system. It is best that you also license deduplication on the secondary storage system. These licenses ensure that deduplication operations can continue without any disruption in case a failure causes the secondary vFiler unit to become the primary storage system.

To use the deduplication feature, you should activate the following licenses on the storage system :

- `multistore`
- `a_sis`

You can run deduplication commands using the RSH or SSH protocol. Any request is routed to the IP address and IP space of the destination vFiler unit.

Note: During an online migration of a vFiler unit, the following deduplication operations are not allowed on volumes that are owned by vFiler units:

- `sis start`
- `sis start -s`
- `sis on`
- `sis off`
- `sis config -s`

For more information about disaster recovery and migration of vFiler units, see the *Data ONTAP MultiStore Management Guide*.

Thin replication using SnapVault and Volume SnapMirror

Thin replication refers to the copying of data to another facility during backup and disaster recovery. Snapshot copies can be backed up and replicated to another facility using SnapVault and Volume SnapMirror. These two technologies increase storage efficiency by transferring only changed data blocks after the baseline copy is created.

You can combine deduplication with these thin replication technologies to achieve greater savings by eliminating redundant data. Further, you can use compression with these technologies to reduce the size of the replicated data, saving network bandwidth. The data on a SnapVault or volume SnapMirror source is compressed, transferred over the network, and then uncompressed on the destination before being written to the disk.

The following figure illustrates how SnapVault and SnapMirror store data using thin transfers:




Figure 4: Thin transfers using SnapVault and Volume SnapMirror

Next topics

[Disk-to-disk backups using SnapVault](#) on page 51

[Efficient data protection using volume SnapMirror](#) on page 57

Disk-to-disk backups using SnapVault

SnapVault uses network bandwidth efficiently, because it transfers only the blocks that changed since the last Snapshot copy. It automatically eliminates the duplication that results from other backup technologies, such as tape.

Further, deduplication facilitates space reduction at the source and destination systems and during the data transfer between the two systems.

For more information about SnapVault and Deduplication, see the *Data ONTAP Data Protection Online Backup and Recovery Guide*.

Next topics

What data gets backed up and restored through SnapVault on page 51

Types of SnapVault deployment on page 52

How SnapVault backup works on page 54

How SnapVault backup works for open systems on page 55

Guidelines for creating a SnapVault relationship on page 56

About LUN clones and SnapVault on page 56

Related information

IBM NAS documentation and support site — <http://www.ibm.com/storage/support/nas>

What data gets backed up and restored through SnapVault

The data structures that are backed up and restored through SnapVault depend on the primary system.

- On systems running Data ONTAP, the qtree is the basic unit of SnapVault backup and restore. SnapVault backs up specified qtrees on the primary system to associated qtrees on the SnapVault secondary system. If necessary, data is restored from the secondary qtrees back to their associated primary qtrees.
- On open systems storage platforms, the directory is the basic unit of SnapVault backup. SnapVault backs up specified directories from the native system to specified qtrees in the SnapVault secondary system.
If necessary SnapVault can restore an entire directory or a specified file to the open systems platform.

Note: You can back up the qtrees from multiple primary systems, or directories from multiple open systems storage platforms, to associated qtrees on a single SnapVault secondary volume.

The illustration shows the backup of qtrees and directories on different systems to a single secondary volume.




Figure 5: Primary and secondary systems

Types of SnapVault deployment

You can deploy SnapVault in three ways as per business requirements.

- Basic SnapVault deployment
- Primary to secondary to tape backup variation
- Primary to secondary to SnapMirror variation

Next topics

[What basic SnapVault deployment is](#) on page 52

[Primary to secondary to tape backup variation](#) on page 53

[Primary to secondary to SnapMirror variation](#) on page 54

What basic SnapVault deployment is

The basic SnapVault backup system deployment consists of primary system and a secondary system.

Primary storage systems: Primary systems are the platforms that run Data ONTAP and open systems storage platforms to be backed up.

- On primary systems, SnapVault backs up primary qtree data, non-qtrees, and entire volumes, to qtree locations on the SnapVault secondary systems.
- Supported open systems storage platforms include Windows servers, Solaris servers, AIX servers, Red Hat Linux servers, SUSE Linux servers, and HP-UX servers. On open systems storage platforms, SnapVault can back up directories to qtree locations on the secondary system. SnapVault can restore directories and single files. For more information, see the *Open Systems SnapVault Installation and Administration Guide*.

Secondary storage system: The SnapVault secondary system is the central disk-based unit that receives and stores backup data from the system as Snapshot copies. Any system can be configured as a SnapVault secondary system; however, the recommended hardware platform is a near-line system.

The following figure shows a basic SnapVault deployment.




Figure 6: Basic SnapVault deployment

Primary to secondary to tape backup variation

A common variation to the basic SnapVault backup deployment adds a tape backup of the SnapVault secondary system.

This deployment can serve two purposes:

- It enables you to store an unlimited number of network backups offline while keeping the most recent backups available online in secondary storage. This can help in the quick restoration of data. If you run a single tape backup off the SnapVault secondary storage system, the storage platforms are not subject to the performance degradation, system unavailability, and complexity of direct tape backup of multiple systems.
- It can be used to restore data to a SnapVault secondary system in case of data loss or corruption on that system.

Note: Some UNIX attributes are not preserved using this method; notably, UNIX access control lists (ACLs).

The following figure shows a basic SnapVault deployment with tape backup.




Figure 7: Basic SnapVault deployment with tape backup

Primary to secondary to SnapMirror variation

In addition to the basic SnapVault deployment, you can replicate the SnapVault secondary using SnapMirror. This protects the data stored on the SnapVault secondary against problems with the secondary system itself.

The data backed up to SnapVault secondary storage is replicated to a SnapMirror destination.

If the secondary system fails, the data mirrored to the SnapMirror destination can be converted to a secondary system and used to continue the SnapVault backup operation with minimum disruption.

How SnapVault backup works

Backing up qtrees using SnapVault involves starting the baseline transfers, making scheduled incremental transfers, and restoring data upon request.

How to start the baseline transfers:

- In response to command-line input, the SnapVault secondary system requests initial base transfers of qtrees specified for backup from a primary storage volume to a secondary storage volume. These transfers establish SnapVault relationships between the primary and secondary qtrees.
- Each primary system, when requested by the secondary system, transfers initial base images of specified primary qtrees to qtree locations on the secondary system.

How to make scheduled incremental transfers:

- Each primary system, in response to command-line input, creates sets of scheduled SnapVault Snapshot copies of the volumes containing the qtrees to be backed up. For tracking purposes, you might name according to frequency, for example, `sv_hourly`, `sv_nightly`, and so on.

For each Snapshot set, SnapVault saves the number of primary storage Snapshot copies you specify and assigns each Snapshot a version number (0 for most current, 1 for second most recent, and so on).

- The SnapVault secondary system, in response to command-line input, carries out a specified set of scheduled data transfer and Snapshot actions. For each of its secondary qtrees on a given volume, SnapVault retrieves, from the Snapshot data of each corresponding primary qtree, the incremental changes to the primary qtrees made since the last data transfer. Then SnapVault creates a volume Snapshot copy of the changes in the secondary qtrees. For each transfer and Snapshot set, SnapVault saves the number of secondary storage Snapshot copies that you specify and assigns each Snapshot copy a version number (0 for most current, 1 for second most recent, and so on).

Restoration upon request:

- If data needs to be restored to the primary system, SnapVault transfers the specified versions of the qtrees back to the primary system that requests them.

The following diagram illustrates SnapVault functionality.




Figure 8: SnapVault functionality

How SnapVault backup works for open systems

The process of SnapVault backup of open systems platform directories involves starting the baseline transfers, scheduling incremental transfers, and restoring data upon request.

Starting the baseline transfers:

- In response to CLI input, the SnapVault secondary system requests from an open systems platform initial baseline image transfers of directories specified for backup. These transfers establish SnapVault relationships between the open systems platform directories and the SnapVault secondary qtrees.

- Each open systems platform, when prompted by the secondary system, transfers initial base images of specified directories to qtree locations on the secondary system.

Note: There are no primary-side Snapshot copies in Open Systems SnapVault.

Scheduling incremental transfers:

- The SnapVault secondary storage system, in response to CLI input, follows a set of scheduled data transfers (to which, for tracking purposes, you can assign names like “sv_hourly,” “sv_nightly,” and so on).
To each secondary qtree on a given volume, from a corresponding primary directory on the open systems storage platform, SnapVault transfers the files that have been added or modified since the previous data transfer.
For each set of scheduled data transfers, SnapVault creates a set of incremental Snapshot copies that capture the changes to the secondary qtrees after each transfer. For each set of Snapshot copies, the SnapVault secondary system saves the number of secondary storage Snapshot copies you specify and assigns each Snapshot copy in the set a version number (0 for most current, 1 for second most recent, and so on).
- Restore upon request: If directory or file data needs to be restored to the open systems storage platform, SnapVault retrieves the data from one of the retained Snapshot copies and transfers the data back to the open systems storage platform that requests it.
For more information, see the *Open Systems SnapVault Installation and Administration Guide*.

Guidelines for creating a SnapVault relationship

You need to follow certain guidelines when creating a SnapVault relationship.

When creating a SnapVault relationship, remember the following guidelines for volumes and qtrees.

- Establish a SnapVault relationship between volumes that have the same `vol_lang` settings.
- After you establish a SnapVault relationship, do not change the language assigned to the destination volume.
- Avoid white space (spaces and tab characters) in names of source and destination qtrees.
- Do not rename volumes or qtrees after establishing a SnapVault relationship.
- The qtree cannot exist on the secondary system before the baseline transfer.

About LUN clones and SnapVault

A LUN clone is a space-efficient copy of another LUN. Initially, the LUN clone and its parent share the same storage space. More storage space is consumed only when one LUN or the other changes.

In releases prior to Data ONTAP 7.3, SnapVault considers each LUN clone as a new LUN. Therefore, during the initial transfer of the LUN clone, all data from the clone and the backing LUN is transferred to the secondary system.

Note: LUNs in this context refer to the LUNs that Data ONTAP serves to clients, not to the array LUNs used for storage on a storage array.

For descriptions of data backup and restore on volumes containing LUNs, see the *Data ONTAP Block Access Management Guide for iSCSI and FC*.

Starting with Data ONTAP 7.3, SnapVault can transfer LUN clones in an optimized way by using SnapDrive for Windows. To manage this process, SnapDrive for Windows creates two Snapshot copies:

- Backing Snapshot copy, which contains the LUN to be cloned
- Backup Snapshot copy, which contains both the LUN and the clone

Modes of transfer

Starting with Data ONTAP 7.3, a SnapVault transfer with LUN clones can run in two modes:

- In non-optimized mode, a LUN clone is replicated as a LUN. Therefore, a LUN clone and its backing LUN get replicated as two separate LUNs on the destination. SnapVault does not preserve space savings that come from LUN clones.
- In optimized mode, a LUN clone is replicated as a LUN clone on the destination. Transfers of LUN clones to the secondary system in optimized mode are possible only with SnapDrive for Windows.

These modes apply to newly created LUN clones. On successive update transfers, only the incremental changes are transferred to the destination in both modes.

Note: A single relationship must either be optimized or non-optimized. Switching between the two modes is not allowed.

Efficient data protection using volume SnapMirror

Volume SnapMirror provides an easy-to-administer replication solution that makes efficient use of available network bandwidth by transferring only changed blocks. If a disaster occurs, businesses can access data from a replica on a remote storage system for uninterrupted operation.

When mirroring asynchronously, SnapMirror replicates Snapshot copies from a source system to a destination system. When an update occurs, a new Snapshot copy is created and is compared against the previous Snapshot copy to determine the changes since the last update. Only the new and changed blocks are sent to the destination system. At the destination system, the changed blocks are merged with the existing data blocks resulting in a full mirror copy of the source system.

Because SnapMirror is based on the Snapshot copy technology and also seamlessly integrates with deduplication, it consumes minimal storage space and saves on network bandwidth. When volume SnapMirror is combined with deduplication, any savings on the SnapMirror source volume are inherited at the destination volume.

For more information about volume SnapMirror, see the *Data ONTAP Data Protection Online Backup and Recovery Guide*.

Next topics

[How SnapMirror works](#) on page 58

[Applications of SnapMirror](#) on page 59

[Recommended actions while using SnapMirror](#) on page 60

[SnapMirror deployment variations](#) on page 60

[Data replication from one destination to another in a series \(cascading\)](#) on page 61

Related information

[IBM NAS documentation and support site](http://www.ibm.com/storage/support/nas) — <http://www.ibm.com/storage/support/nas>

How SnapMirror works

SnapMirror replicates data from a source volume or qtree to a partner destination volume or qtree, respectively, by using Snapshot copies. Before using SnapMirror to copy data, you need to establish a relationship between the source and the destination.

You can specify a SnapMirror source and destination relationship between volumes or qtrees by using one of the following options.

- The `/etc/snapmirror.conf` file
- The `snapmirror.access` option
- The `/etc/snapmirror.allow` file

The SnapMirror feature does the following:

1. Creates a Snapshot copy of the data on the source volume.
2. Copies it to the destination, a read-only volume or qtree.
3. Updates the destination to reflect incremental changes on the source, as per the schedule you specify.

The result of this process is an online, read-only volume or qtree that contains the same data as the source at the time of the most recent update.

Each volume SnapMirror replication, qtree SnapMirror replication, or SnapVault replication consists of a pair of operations. There is one operation each at:

- The source storage system
- The destination storage system

Therefore, if a storage system is the source for one replication and the destination for another replication, it uses two replication operations. Similarly, if a storage system is the source as well as the destination for the same replication, it uses two replication operations.

Applications of SnapMirror

SnapMirror is used to replicate data. Its qualities make SnapMirror useful in several scenarios, including disaster recovery, data backup, and data restoration.

You can copy or use the data stored on a SnapMirror destination. The additional advantages of SnapMirror make it useful in data retrieval situations such as those described in the following table.

Situation	How to use SnapMirror
Disaster recovery: You want to provide immediate access to data after a disaster has made a qtree, volume, or system unavailable.	You can make the destination writable so clients can use the same data that was on the source volume the last time data was copied.
Disaster recovery testing: You want to test the recovery of data and restoration of services in the event of a disaster.	You can use FlexClone technology on the SnapMirror destination, and test for disaster recovery, without stopping or pausing other replication operations.
Data restoration: You want to restore lost data on a qtree or volume source from its mirrored qtree or volume SnapMirror partner.	You can temporarily reverse the roles for the source and destination qtrees or volumes and copy the mirrored information back to its source.
Application testing: You want to use an application on a database, but you want to test it on a copy of the database in case the application damages the data.	You can make a copy of the database to be used in the application testing to ensure that the data on the source cannot be lost.
Load balancing: A large number of users need read-only access to a qtree or volume.	You can copy the data in a qtree or volume to multiple volumes or systems to distribute the load.
Off-loading tape backups: You need to reserve all processing and networking resources on a system for serving NFS and CIFS requests.	After copying data on the source system, you can back up the data in the destination to tape. This means that the source system does not have to allocate resources for performing backups.
Access to remote data: Users who need read access to a volume are distributed over a large geographical area.	You can copy the source volume to other systems that are geographically closer to the users. Users accessing a local system can read the data using less resource time than if they connected to a distant system.

Recommended actions while using SnapMirror

While using SnapMirror, you can increase the efficiency of data copying by performing certain actions. This includes the staggering of Snapshot copy schedules and SnapMirror update schedules.

- To optimize performance, stagger your Snapshot copy update schedules so that SnapMirror activity does not begin or end at the exact minute a snap sched command operation attempts to create a Snapshot copy.

If the SnapMirror feature is scheduled to perform Snapshot copy management at the same time as a snap sched activity, then the Snapshot copy management operations scheduled using the snap sched command might fail with syslog messages: "Skipping creation of hourly snapshot" and "Snapshot already exists."

- For optimum SnapMirror volume replication performance, ensure that the SnapMirror source volume and destination volume contain disks of the same size, organized in the same RAID configuration.
 - If the SnapMirror source and destination are FlexVol volumes, the RAID configurations do not make a difference.
 - If the SnapMirror source and destination are qtrees, volume size and configuration do not make any difference.

SnapMirror deployment variations

There are several variations possible while deploying SnapMirror. These variations allow you to customize the solution to suit your requirements.

Source to destination to tape variation: A common variation to the basic SnapMirror backup deployment adds a tape backup of the destination volume. By running a tape backup off the SnapMirror destination volume (as shown in the following illustration), you do not subject the heavily-accessed source volume to the performance degradation and complexity of a direct tape backup.




Figure 9: SnapMirror deployment: Source to destination to tape

Source to tape to destination variation: A SnapMirror deployment that supports SnapMirror replication over low-bandwidth connections accommodates an initial mirroring between a source and

destination volume using physically-transported tape (as shown in the following illustration). After the large base Snapshot copy has been replicated, smaller, incremental Snapshot copy updates can be carried out over a low-bandwidth connection.




Figure 10: SnapMirror deployment: Source to tape to destination

Cascading destinations variation: A variation on the basic SnapMirror deployment and function involves a writable source volume replicated to multiple read-only destinations. The function of this deployment is to make a uniform set of data available on a read-only basis to users from various locations throughout a network and to allow for updating that data uniformly at regular intervals.

Note: The cascade deployment (as shown in the following illustration) is supported for volume SnapMirror only. It is not supported for qtree SnapMirror.




Figure 11: SnapMirror deployment: Cascade

Data replication from one destination to another in a series (cascading)

You can replicate data from a SnapMirror destination to another system using SnapMirror.

Therefore, a system that is a destination for one SnapMirror relationship can act as the source for

another SnapMirror relationship. This is useful when you need to copy data from one site to many sites.

Instead of replicating data from a single source to each of the destinations, you can replicate data from one destination to another destination, in a series. This is referred to as cascading.

Note: You can replicate data from a destination volume in the same way you replicate from a writable source volume.

Virtual data replication with FlexClone technology

A FlexClone volume, file, or LUN is a writable, point-in-time image of the FlexVol volume or another FlexClone volume, file, or LUN. This technology enables you to use space efficiently, storing only data that changes between the parent and the clone.

Flexible clones result in savings in space, power, and dollars. Additionally, clones have the same high performance as their parent volumes.

Next topics

[*About FlexClone volumes*](#) on page 63

[*FlexClone volumes and space guarantees*](#) on page 64

[*About FlexClone files and FlexClone LUNs*](#) on page 65

[*How FlexClone files and FlexClone LUNs work*](#) on page 65

[*Uses of FlexClone files and FlexClone LUNs*](#) on page 67

[*Considerations when planning FlexClone files or FlexClone LUNs*](#) on page 67

[*Space savings achieved by using FlexClone files and FlexClone LUNs*](#) on page 68

About FlexClone volumes

FlexClone volumes provide writeable volume copies that use only as much storage as necessary to hold new data. They can be created instantaneously without interrupting access to the parent FlexVol volume. Additionally, they use the space very efficiently, storing only data that changes between the parent and the clone.

A FlexClone volume is initialized with a Snapshot copy and updated continuously thereafter.

For more information about flexible clones, see the *Data ONTAP Storage Management Guide*.

The following figure illustrates the space savings of test and development storage without and with FlexClone volumes.




Figure 12: Space savings of test and development storage without and with FlexClone volumes

Related information

IBM NAS documentation and support site — <http://www.ibm.com/storage/support/nas>

FlexClone volumes and space guarantees

A FlexClone volume inherits its initial space guarantee from its parent volume. For example, if you create a FlexClone volume from a parent volume with a space guarantee of `volume`, then the FlexClone volume's initial space guarantee will be `volume` also. You can change the FlexClone volume's space guarantee.

For example, suppose that you have a 100-MB FlexVol volume with a space guarantee of `volume`, with 70 MB used and 30 MB free, and you use that FlexVol volume as a parent volume for a new FlexClone volume. The new FlexClone volume has an initial space guarantee of `volume`, but it does

not require a full 100 MB of space from the aggregate, as it would if you had copied the volume. Instead, the aggregate needs to allocate only 30 MB (100 MB minus 70 MB) of free space to the clone.

If you have multiple clones with the same parent volume and a space guarantee of volume, they all share the same shared parent space with each other, so the space savings are even greater.

Note: The shared space depends on the existence of the shared Snapshot copy (the base Snapshot copy that was used to create the FlexClone volume). If you delete this shared Snapshot copy, you lose the space savings provided by the FlexClone volume.

About FlexClone files and FlexClone LUNs

FlexClone files and FlexClone LUNs are writable, space-efficient clones of parent files and parent LUNs.

The Data ONTAP block-sharing mechanism is used for creating FlexClone files and LUNs. Clones use a small amount of storage space to store their metadata. Clones share the data blocks of their parent files and parent LUNs and occupy negligible storage space until clients write new data either to the parent file or LUN, or to the clone.

You can create FlexClone files and LUNs in the same FlexVol volume as their parent files and LUNs.

Clients can perform all normal file and LUN operations on both parent entities and clone entities.


How FlexClone files and FlexClone LUNs work

Creating FlexClone files or FlexClone LUNs is highly space-efficient and time-efficient because the cloning operation does not involve physically copying any data.

You can create a clone of a file that is present in a FlexVol volume in a NAS environment, and you can also clone a complete LUN without the need of a backing Snapshot copy in a SAN environment.

The cloned copies initially share the same physical data blocks with their parents and occupy negligible extra space in the storage system for their initial metadata.

The following illustration shows the parent files or LUNs and FlexClone files or LUNs accessing the same data blocks on the storage system. On the host side, the parent files or LUNs and FlexClone files or LUNs appear as normal files and LUNs.



Unlike FlexClone volumes and LUN clones, the FlexClone files and FlexClone LUNs do not depend on a backing Snapshot copy. However, by default the cloning operation creates a temporary Snapshot copy of the FlexVol volume in which the cloning operation is being carried out. The temporary Snapshot copy is deleted immediately after a FlexClone file or LUN is created. You can stop the creation of a temporary Snapshot copy by using the `-n` option of the `clone start` command, but you should do so only when you are certain that no writes will happen to the parent file or LUN during the cloning operation.

The cloning operation has no impact on client access to the parent file or LUN, either during the creation of clones or after the cloning operation is complete. Clients that are accessing the parent file or LUN do not experience any disruption or outage during the cloning operation. Clients can write to the source file or LUN while the cloning operation is in progress. Once the cloning operation is

complete, clients see the FlexClone files or FlexClone LUNs as normal files and LUNs. Clients can perform all normal operations on them as they can on standard files and LUNs.

When clients write new data to a parent or clone, then the entity on which new data is written starts occupying extra storage space.

Uses of FlexClone files and FlexClone LUNs

FlexClone files and FlexClone LUNs can help save time and storage space in a variety of situations.

You can quickly create space-efficient copies of your data by using FlexClone files and FlexClone LUNs in the following situations:

- When you need to deploy, upgrade, or redeploy thousands of standardized virtual desktops or servers
- When you need to test video, sound, or image processing applications
You can use the cloned files for testing the applications.
- When you need to boot servers in a server farm
You can create FlexClone LUNs of the parent boot LUN, then use the FlexClone LUN to boot a server in a server farm.

Considerations when planning FlexClone files or FlexClone LUNs

You should keep several considerations in mind when planning how to deploy FlexClone files and FlexClone LUNs.

- You can create FlexClone files and LUNs only in the same FlexVol volume as the parent files and LUNs.
- The following hardware platforms support FlexClone files and FlexClone LUNs:
 - N3300, N3400, and N3600 series
 - N5000 series
 - N6040, N6060, or N6070
 - N7600, N7700, N7800, or N7900
- You can create a FlexClone file or LUN only of a file or LUN that is part of the active file system. If you want to clone a file or LUN inside a Snapshot copy, you must first restore the entity to the active file system.
- You can clone a complete file, sub-file, LUN, or sub-LUN.
To clone a sub-file or sub-LUN, you should know the block range of the parent entity and clone entity.

- The time required for creating a FlexClone file or FlexClone LUN depends on the size of the parent file or LUN.
- The `sis` attribute is added to a FlexVol volume when a FlexClone file or FlexClone LUN is created for the first time.

FlexVol volumes with deduplication enabled also show the `sis` attribute when you run the `vol status` command.

Space savings achieved by using FlexClone files and FlexClone LUNs

You can use the `df -s` command to view the amount of storage space saved by creating FlexClone files and FlexClone LUNs. When you create a clone, you save the amount of space that is occupied by its parent.

Note: If you run the `df -s` command on a FlexVol volume with deduplication enabled, the output displays the space saved by both deduplication and FlexClone files or FlexClone LUNs.

Example

If you have a FlexVol volume of 100 GB with 50 GB used space and then create a file of 10 GB and a clone of it, the total used physical space is about 60 GB (50 GB + 10 GB for file and its clone). If the clone were a full physical copy, you would be using 70 GB (50 GB + 10 GB for file + 10 GB for the clone). Therefore, you saved space of 10 GB by creating a FlexClone file. Your savings are 14% $((10/70)*100)$.

Storage efficiency on non-IBM N series storage with IBM gateway

The IBM N series system with gateway functionality is a network-based solution that virtualizes tiered, heterogeneous storage arrays. Customers who have other vendors' storage systems and disk shelves can take advantage of all the IBM N series storage-saving features by using the gateway product line.

The IBM gateway is a controller that you place in front of a third-party SAN. It unifies block-storage and file-storage networking paradigms under a common architecture, so that your entire storage infrastructure can be virtualized through one interface.

To use the gateway on third-party SANs, you must create LUNs using the third-party disks and RAID configurations. When the third-party LUNs are available, the gateway uses them to create aggregates that can contain NAS volumes, SAN LUNs, or both. The final step is the host level migration of the data to the gateway volumes or LUNs.

The following figure illustrates a typical gateway:




Figure 13: N series gateway system

Next topics

[*Restructuring IT with Storage Efficiency calculator*](#) on page 70

[*How to quantify storage utilization using the Storage Efficiency calculator*](#) on page 70

Restructuring IT with Storage Efficiency calculator

The Restructuring IT with Storage Efficiency (RISE) calculator has approximately 50 input fields that are used to validate cost and capacity savings after evaluating your environment.

The RISE calculator does the following:

- Analyzes techniques for storage-capacity reduction.
- Validates storage CAPEX / OPEX cost savings.
- Establishes priorities and assists in developing a plan for optimal storage efficiency.

For every question in the survey, a default response is provided, based on typical customer values or acknowledged U.S. industry standards. If you do not know the answer to any question, you can use the default value.

Note: Defaults are not as accurate as the actual data obtained from the evaluation site.

The RISE calculator is divided into several sections.

The first section, indicated by the yellow tabs, provides an area to enter the actual information that describes the evaluation site. The second section, indicated by the green tab, represents the summary output resulting from an analysis of the current environment and factoring for storage efficiency. A third section, indicated by the blue tabs, provides a background analysis of the data used to calculate the summary information.

The best way to use this tool is to begin with the baseline data in Section One. As you begin to enter data, you should examine the summary output. As more input cells are filled in with known data, the results in the summary output become more clearly defined. In this manner, you can start with a small subset of site data and build on it as more information becomes available.

Note: The RISE calculator is only available from your IBM representative or authorized IBM partner representative.

How to quantify storage utilization using the Storage Efficiency calculator

The Storage Efficiency calculator is an interactive tool that helps to quantify the savings that you might expect from using IBM N series products.

The Storage Efficiency calculator calculates the power and cooling savings. It also generates a detailed report that contains the following information:

- Estimated primary storage capacity savings

- Estimated storage capacity savings by storage tier
- Total space, cooling, power, and dollars saved

Related information

Storage Efficiency calculator: www.secalc.com/ibm

Index

A

access to remote data
using SnapMirror 59

B

backups
using SnapMirror 59

C

configuration of thin provisioning with Provisioning
Manager 29

D

deduplication
FlexVol volume
maximum size 39
maximum size with deduplication 39
maximum volume size 39
metadata relocated 38
with FlexClone 45
with qtree SnapMirror 42
with SnapRestore 43
with SnapVault 42
with vFiler units 46
with volume copy 44
with volume SnapMirror 41
deduplication with FlexClone 45
deduplication with SnapRestore 43
deduplication with volume copy 44
disaster recovery
using SnapMirror 59

F

file folding
defined 35
FlexClone
with deduplication 45
FlexClone files and FlexClone LUNs
about 65
considerations 67

hardware platform support 67
how 65
space saving 68
uses 67

FlexClone technology 63

FlexClone volumes
space guarantees and 64

FlexVol volumes
automatically adding space for 26

fractional reserve
about 24
and space management 21

free space
automatically increasing 26

H

hardware efficiency
Performance Acceleration Module 17
SATA drives 17
how to achieve storage efficiency
using IBM products 15
HP-UX
backing up with SnapVault 52
See also open systems platform 52

I

IBM solutions to industry storage challenges 15

L

load balancing
using SnapMirror 59
LUN clones
description 56
modes of transfer 56

N

non-optimized mode 57

O

open systems platform

- non-filer systems supported for SnapVault backup 52

- optimized mode 57

P

- Performance Acceleration Module 17

- primary storage system

- Data ONTAP and open systems platforms 52

Q

- qtree SnapMirror with deduplication 42

R

- RAID-DP

- protection against double disk failures 19

- Restructuring IT with Storage Efficiency calculator 70

- RISE calculator 70

S

- SATA disks 17

- secondary storage system

- role in a basic SnapVault deployment 52

- SnapDrive for Windows 57

- SnapMirror

- cascading 61

- qtree replication 58

- recommendations 60

- uses 59

- volume replication 58

- SnapRestore

- with deduplication 43

- Snapshot copies 31, 40

- SnapVault

- basic deployment 52

- secondary to SnapMirror deployment 54

- with deduplication 42

- software efficiency

- FlexClone 63

- FlexClone technology 63

- FlexVol volumes 27

- RAID-DP 19, 20

- Snapshot copy 31

- SnapVault 49, 51

- volume SnapMirror 49, 57

- Solaris

- backing up with SnapVault 52

- space guarantees

- about 22

- space management option 21

- space management

- choosing 21

- DataFabric Manager 29

- how it works 21

- space reservations

- about 23

- space management option 21

- storage efficiency

- about 15

- Operations Manager 29

- Provisioning Manager 29

- storage efficiency calculator

- quantify storage utilization 70

- storage efficiency with IBM N series systems 69

- storage utilization using the storage efficiency calculator 70

T

- tape backup of a SnapVault secondary 53

- thin provisioning

- automation using Provisioning Manager 29

- benefits 28

- using FlexVol volumes 27

V

- vFiler unit with deduplication 46

- volume copy

- with deduplication 44

- volume SnapMirror with deduplication 41

- volume SnapVault and SnapMirror 49

- volumes

- automatically adding space for 26

W

- what is storage efficiency 15

- Windows 2000

- backing up with SnapVault 52



NA 210-04805_A0, Printed in USA

GA32-0773-00

