

IBM System Storage N series



Data ONTAP 8.1 Storage Efficiency Management Guide for 7-Mode

Contents

Preface	1
About this guide.	1
Supported features	1
Websites	1
Getting information, help, and service	2
Before you call	2
Using the documentation	2
Hardware service and support	3
Firmware updates	3
How to send your comments	3
 What storage efficiency is	 5
The IBM N series advantage.	5
 Storage efficiency with SATA storage disks and Flash Cache	 7
 Storage efficiency with Flash Pool	 9
 Protection against double disk failure with RAID-DP	 11
About RAID-DP protection	11
How Data ONTAP uses RAID to protect your data and data availability	12
 Using space management to balance competing requirements	 13
How volume guarantees work with FlexVol volumes	13
How file and LUN reservations work	15
How fractional reserve works	16
How Data ONTAP can automatically provide more space for full FlexVol volumes	17
 Thin provisioning for greater efficiencies using FlexVol volumes.	 19
Benefits of using thin provisioning	20
Best practices for using thin provisioning	20
Storage space management using OnCommand Unified Manager	21
Automate thin provisioning using the N series Management Console provisioning capability.	22
 Efficient file recovery with Snapshot copies	 23
What Snapshot disk consumption is	23
How Snapshot copies consume disk space.	23
What the Snapshot copy reserve is	24
Use of deleted active file disk space	25
Example of what happens when Snapshot copies exceed the reserve	25
Recovery of disk space for file system use	26
What file folding means and how it saves disk space	27
 Using deduplication to increase storage efficiency.	 29

How deduplication works	29
What deduplication metadata is	30
Guidelines for using deduplication	31
Performance considerations for deduplication	31
Deduplication interoperability with Data ONTAP features	32
Best practices for using deduplication with other Data ONTAP features	33
Using data compression to increase storage efficiency	35
Guidelines for using data compression on existing data	36
Data compression interoperability with Data ONTAP features	36
Best practices for using data compression with other Data ONTAP features	37
Thin replication using SnapVault and Volume SnapMirror	39
Disk-to-disk backups using SnapVault	40
How SnapVault backup works.	41
Types of SnapVault deployment	42
What basic SnapVault deployment is	42
Primary to secondary to tape backup variation	43
Primary to secondary to SnapMirror variation	44
Best practices for using SnapVault with other Data ONTAP features	44
Using SnapVault with deduplication and data compression	45
Using SnapVault with volume SnapMirror.	45
Efficient data protection using volume SnapMirror	45
How SnapMirror works	46
Best practices for using SnapMirror with other Data ONTAP features.	47
Using volume SnapMirror with deduplication	47
Using volume SnapMirror with data compression	47
Using volume SnapMirror with FlexClone volumes.	48
Using volume SnapMirror with SnapVault.	48
Using qtree SnapMirror with deduplication	48
Using qtree SnapMirror with data compression	49
Using qtree SnapMirror with FlexClone volumes	49
Using FlexClone technology to create efficient copies of volumes, files, and LUNs	51
How FlexClone files and FlexClone LUNs work.	52
FlexClone volumes and space guarantees	54
Uses of FlexClone files and FlexClone LUNs	54
Considerations for planning FlexClone files or FlexClone LUNs	54
Space savings achieved by using FlexClone files and FlexClone LUNs	55
Best practices for using FlexClone technology with other Data ONTAP features	55
Using FlexClone volumes with volume SnapMirror.	56
Using FlexClone files or FlexClone LUNs with volume SnapMirror	56
Using FlexClone files or FlexClone LUNs with qtree SnapMirror and SnapVault	56
Using FlexClone files or FlexClone LUNs with deduplication	56
Copyright and trademark information	59
Trademark information	60
Notices	63

Index 65

Preface

About this guide

This document applies to IBM N series systems running Data ONTAP, including systems with gateway functionality. If the term *7-Mode* is used in the document, it refers to Data ONTAP operating in 7-Mode, which has the same features and functionality found in the prior Data ONTAP 7.1, 7.2, and 7.3 release families.

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.

Supported features

IBM System Storage N series storage systems are driven by NetApp Data ONTAP software. Some features described in the product software documentation are neither offered nor supported by IBM. Please contact your local IBM representative or reseller for further details.

Information about supported features can also be found on the N series support website (accessed and navigated as described in Websites).

Websites

IBM maintains pages on the World Wide Web where you can get the latest technical information and download device drivers and updates. The following web pages provide N series information:

- A listing of currently available N series products and features can be found at the following web page:
www.ibm.com/storage/nas/

- The IBM System Storage N series support website requires users to register in order to obtain access to N series support content on the web. To understand how the N series support web content is organized and navigated, and to access the N series support website, refer to the following publicly accessible web page:

www.ibm.com/storage/support/nseries/

This web page also provides links to AutoSupport information as well as other important N series product resources.

- IBM System Storage N series products attach to a variety of servers and operating systems. To determine the latest supported attachments, go to the IBM N series interoperability matrix at the following web page:

www.ibm.com/systems/storage/network/interophome.html

- For the latest N series hardware product documentation, including planning, installation and setup, and hardware monitoring, service and diagnostics, see the IBM N series Information Center at the following web page:

publib.boulder.ibm.com/infocenter/nasinfo/nseries/index.jsp

Getting information, help, and service

If you need help, service, or technical assistance or just want more information about IBM products, you will find a wide variety of sources available from IBM to assist you. This section contains information about where to go for additional information about IBM and IBM products, what to do if you experience a problem with your IBM N series product, and whom to call for service, if it is necessary.

Before you call

Before you call, make sure you have taken these steps to try to solve the problem yourself:

- Check all cables to make sure they are connected.
- Check the power switches to make sure the system is turned on.
- Use the troubleshooting information in your system documentation and use the diagnostic tools that come with your system.
- Refer to the N series support website (accessed and navigated as described in Websites) for information on known problems and limitations.

Using the documentation

The latest versions of N series software documentation, including Data ONTAP and other software products, are available on the N series support website (accessed and navigated as described in Websites).

Current N series hardware product documentation is shipped with your hardware product in printed documents or as PDF files on a documentation CD. For the latest N series hardware product documentation PDFs, go to the N series support website.

Hardware documentation, including planning, installation and setup, and hardware monitoring, service, and diagnostics, is also provided in an IBM N series Information Center at the following web page:

publib.boulder.ibm.com/infocenter/nasinfo/nseries/index.jsp

Hardware service and support

You can receive hardware service through IBM Integrated Technology Services. Visit the following web page for support telephone numbers:

www.ibm.com/planetwide/

Firmware updates

IBM N series product firmware is embedded in Data ONTAP. As with all devices, ensure that you run the latest level of firmware. Any firmware updates are posted to the N series support website (accessed and navigated as described in Websites).

Note: If you do not see new firmware updates on the N series support website, you are running the latest level of firmware.

Verify that the latest level of firmware is installed on your machine before contacting IBM for technical support.

How to send your comments

Your feedback helps us to provide the most accurate and high-quality information. If you have comments or suggestions for improving this document, please send them by email 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 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.

Storage efficiency includes using technologies such as thin provisioning, Snapshot copy, deduplication, data compression, FlexClone, thin replication with SnapVault and volume SnapMirror, and Flash Pool which help to increase storage utilization and decrease storage costs.

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 configured within Flash Pool or with Flash Cache and RAID-DP technology, increase efficiency without affecting performance and resiliency.

Technologies such as thin provisioning, Snapshot copy, deduplication, data compression, thin replication with SnapVault and volume SnapMirror, and FlexClone offer better savings. You can use these technologies individually or together to achieve maximum storage efficiency.

The IBM N series advantage

IBM N series has a rich set of features such as SATA disks, Flash Cache, RAID-DP, thin provisioning, Snapshot copy, deduplication, data compression, SnapVault, SnapMirror, Flash Pool, and FlexClone, which help achieve significant improvements in storage utilization. When used together, these technologies help in increasing storage efficiency.

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

- RAID-DP protects your data from double-disk failures which significantly increases fault tolerance.
- Flash Cache improves performance for workloads that are random read-intensive without adding additional disk drives.

- Flash Pool reduces the number of disk reads and writes, which improves performance in environments where high amount of shared blocks are repeatedly read or randomly written.
- Thin provisioning enables you to maintain a common unallocated storage space that is readily available to other applications as needed.
- Snapshot copies are read-only images of traditional volumes, FlexVol volumes, or aggregates that capture the state of the storage system at a point-in-time.
Snapshot copies take up zero space when it is created, and help back up and restore data.
- Deduplication saves storage space by eliminating redundant data blocks within a FlexVol volume.
With Flash Cache and Flash Pool, the deduplication savings are retained in the cache that exists in the disk, which increases the amount of data stored in the cache.
- Data compression reduces the physical capacity that is required to store data on the storage system by compressing data chunks within a FlexVol volume.
- SnapVault is a disk-based storage backup feature of Data ONTAP.
SnapVault enables data stored on multiple systems to be backed up as read-only Snapshot copies to a central secondary system quickly and effectively.
- SnapMirror is a disaster recovery solution, which enables you to replicate data from specified source volumes or qtrees to specified destination volumes and qtrees.
This solution helps in minimizing downtime costs due to a production site failure and enables you to centralize the backup of data to disks from multiple data centers.
- FlexClone allows you to instantly create writable virtual copies of data volumes, files, and LUNs, which do not consume storage space.
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 storage efficiently, storing only data that changes between the parent and the clone. This feature also enables you to save space, power, and cost. Additionally, FlexClones 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 (except RAID-DP and SATA), and you can apply all the preceding efficiencies to your non-IBM N series systems.

Storage efficiency with SATA storage disks and Flash Cache

SATA disks along with Flash Cache provide low-cost and high-capacity storage solutions without sacrificing performance.

You can combine Flash Cache with SATA disks, which provide a low-cost alternative to Fibre Channel disks, ensure less power consumption, and better cooling due to less number of physical drives. The performance and resiliency are maintained by using Flash Cache and RAID-DP.

Flash Cache 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 by using less number of disks or lower cost disks.

Flash Cache works with other Data ONTAP features to provide the following benefits:

- Deduplication blocks are cached similar to any other block and only one copy is kept in the Flash Cache, similar to how the blocks are stored on the disk.

This characteristic enables you to use Flash Cache more effectively.

- Flash Cache is fully interoperable with FlexClone technology and data access from FlexClone volumes is accelerated by caching only one copy of FlexClone data block, regardless of how many FlexClone volumes share that block.
- Flash Cache works at a level below the protocols and can store data for FlexCache volumes.

Note: Read performance of compressed data remains the same with or without Flash Cache cards.

Storage efficiency with Flash Pool

Flash Pool technology enables you to combine solid-state disks (SSDs) and traditional hard disk drives (HDDs) to form a single aggregate. Flash Pool caches random reads and writes, but does not cache sequential I/O operations. Storage efficiency features deduplication and FlexClone are supported in Flash Pool.

You can combine deduplication and FlexClone to increase the cache memory space. Although the amount of physical cache memory available in the SSDs is fixed, the amount of data stored on the SSDs is more due to the space savings achieved by deduplication and FlexClone.

Because the I/O operations of compressed data is performed sequentially, blocks of data that contain compressed data are not stored in the SSDs. However, the metadata and uncompressed data continue to benefit from the throughput and latency improvements provided by Flash Pool.

For more information about Flash Pool, see the *Flash Pool Design and Implementation Guide*.

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.



Flash Pool Design and Implementation Guide: media.netapp.com/documents/tr-4070.pdf

Protection against double disk failure with RAID-DP

RAID-DP is a standard Data ONTAP feature that protects your data from double-disk failures or single-disk failures with media errors. This feature significantly increases the fault tolerance of RAID groups in Data ONTAP over traditional single-parity RAID.

When considering the standard mean time to data loss (MTTDL) formula with RAID-DP versus single-parity RAID, RAID-DP is more reliable.

RAID-DP enables effective deployment of SATA based storage systems by providing increased resiliency and the ability to deploy larger RAID groups when compared to single parity RAID. 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 for 7-Mode*.

Note: It is likely that one disk failed before the other and 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.

Related information:

 IBM N series support website: www.ibm.com/storage/support/nseries

About RAID-DP protection

If the aggregate RAID groups are configured as RAID-DP (default RAID type), Data ONTAP reconstructs the data from up to two failed drives within a single RAID group. RAID-DP is the default RAID type for all aggregates.

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

- There is a double-disk failure within a RAID group.
- There is a single-disk failure or a media error on the disk, or both instances within a RAID group.

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 disk failure in a RAID-DP group, Data ONTAP replaces the failed disk with a spare disk and uses row parity data to reconstruct the failed disk.

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

How Data ONTAP uses RAID to protect your data and data availability

Understanding how RAID protects your data and data availability can help you administer your storage systems more effectively.

For native storage, Data ONTAP uses RAID-DP (double-parity) or RAID Level 4 (RAID4) protection to ensure data integrity within a group of disks even if one or two of those disks fail. Parity disks provide redundancy for the data stored in the data disks. If a disk fails (or, for RAID-DP, up to two disks), the RAID subsystem can use the parity disks to reconstruct the data in the drive that failed.

For third-party storage, Data ONTAP stripes data across the array LUNs using RAID0. The storage arrays, not Data ONTAP, provide the RAID protection for the array LUNs that they make available to Data ONTAP.

Using space management to balance competing requirements

The space management capabilities of Data ONTAP enable 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:

- Volume (space) guarantee
The *volume guarantee*, also called *space guarantee* or just *guarantee*, determines how much space for the volume is preallocated from the volume's associated aggregates when the volume is created. The guarantee is an attribute of the volume.
- Reservations
Reservations, also called *space reservations*, *file reservations*, or *LUN reservations*, determine whether space for a particular file or LUN is preallocated from the volume. Reservations are an attribute of the file or LUN.
- Fractional reserve
Fractional reserve, also called *fractional overwrite reserve*, enables the administrator to reduce the amount of space preallocated for over-writes to space-reserved files and LUNs in the volume. It is expressed as a percentage.
- Automatic free space preservation
Automatic free space preservation can either increase the size of a volume or delete Snapshot copies to prevent a volume from running out of space—all without operator intervention.

These capabilities are used together to enable you to determine, on a volume-by-volume basis, whether to emphasize storage utilization or ease of management or something in-between.

How volume guarantees work with FlexVol volumes

Volume guarantees (sometimes called *space guarantees*) determine how space for a volume is allocated from its containing aggregate—whether the space is preallocated for the entire volume or for only the reserved files or LUNs in the volume, or whether no space at all is preallocated for the volume.

The guarantee is an attribute of the volume. It is persistent across storage system reboots, takeovers, and givebacks. Volume guarantee types can be **volume** (the default type), **file**, or **none**.

- A guarantee type of **volume** allocates space in the aggregate for the volume when you create the volume, regardless of whether that space is used for data yet.

This approach to space management is called *thick provisioning*. The allocated space cannot be provided to or allocated for any other volume in that aggregate. This is the default type.

When you use thick provisioning, all of the space specified for the volume is allocated from the aggregate at volume creation time. The volume cannot run out of space before the amount of data it contains (including snapshots) reaches the size of the volume. However, if your volumes are not very full, this comes at the cost of reduced storage utilization.

- A guarantee type of **file** allocates space for the volume in its containing aggregate so that any reserved LUN or file in the volume can be completely rewritten, even if its blocks are being retained on disk by a Snapshot copy.

However, writes to any file in the volume that is not reserved could run out of space.

Before configuring your volumes with a guarantee of **file**, you should refer to Technical Report 3965.

- A guarantee of **none** allocates space from the aggregate only as it is needed by the volume.

This approach to space management is called *thin provisioning*. Writes to LUNs or files (including space-reserved files) contained by that volume could fail if the containing aggregate does not have enough available space to accommodate the write.

If you configure your volumes with a volume guarantee of **none**, you should refer to Technical Report 3965 for information about how doing so can affect storage availability.

When space in the aggregate is allocated for the guarantee for an existing volume, that space is no longer considered free in the aggregate. Operations that consume free space in the aggregate, such as creation of aggregate 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 allocated to another volume.

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

Note: Guarantees are honored only for online volumes. If you take a volume offline, any allocated 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 guarantee, you must use the force option, and the volume's guarantee is disabled.

Note: These technical reports contain 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 3965: Thin Provisioning Deployment and Implementation Guide



Technical Report 3483: Thin Provisioning in a NetApp SAN or IP SAN Enterprise Environment

How file and LUN reservations work

When reservations are enabled for one or more files or LUNs, Data ONTAP reserves enough space in the volume so that writes to those files or LUNs do not fail because of a lack of disk space.

Reservations are an attribute of the file or LUN; they are persistent across storage system reboots, takeovers, and givebacks. Reservations are enabled for new LUNs by default, but you can create a file or LUN with reservations disabled or enabled.

When a volume contains one or more files or LUNs with reservations 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 reservations enabled continue to succeed.

You can enable reservations for files and LUNs contained by volumes with volume guarantees of any value. However, if the volume has a guarantee of **none**, reservations do not provide protection against out-of-space errors.

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.

How fractional reserve works

Fractional reserve, also called *LUN overwrite reserve*, enables you to control the size of the overwrite reserve for a volume. By using this volume attribute correctly, you can maximize your storage utilization.

The fractional reserve setting is expressed as a percentage.

Generally, you set fractional reserve to zero. In SAN environments, if you are not using automatic Snapshot deletion, you might set the fractional reserve setting to a non-zero value to provide LUN overwrite reserve as a safeguard to ensure that enough space is reserved in the volume for data that is being overwritten between Snapshot copies.

If the fractional reserve setting is 100, providing full LUN overwrite reserve, you might not be able to create more Snapshot copies or use other block-sharing capabilities, but you can always overwrite any data in your LUN, even if block-sharing capabilities are in use.

Setting fractional reserve to less than 100 percent causes LUN overwrite reserve for that volume to be reduced to that percentage. Writes to the space-reserved files and LUNs in that volume are no longer unequivocally guaranteed when block-sharing capabilities are in use.

Reducing the fractional reserve percentage does not affect the size of a LUN. You can write data to the entire size of the LUN.

The default value and allowed values for the fractional reserve setting depends on the guarantee of the volume:

- For volumes with a guarantee of **volume**, the default value is **100**, and the allowed values are **0** to **100**, inclusive.
- For volumes with a guarantee of **none**, the default value is **0**, and the allowed values are **0** to **100**, inclusive.
- For volumes with a guarantee of **file**, the fractional reserve setting cannot be changed; it is always **100**.

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

Example

If you create a 500-GB reserved LUN in a volume with a guarantee type of **volume**, Data ONTAP ensures that 500 GB of free space always remains available for that LUN to handle writes to the LUN.

If you later set fractional reserve to 0 for the LUN's containing volume and then create a Snapshot copy, the blocks used in the LUN are locked, and Data ONTAP does not reserve any space in the volume for future overwrites to the LUN. Any subsequent overwrites to the LUN could fail due to insufficient free space in the volume, unless you configure the volume to automatically provide more free space, and provide sufficient free space in the aggregate.

For more information about using fractional reserve, see the following Technical Reports:

- *TR-3965: Thin Provisioning Deployment and Implementation Guide*
- *TR-3483: Thin Provisioning in a NetApp SAN or IP SAN Enterprise Environment*

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 3965: Thin Provisioning Deployment and Implementation Guide



Technical Report 3483: Thin Provisioning in a NetApp SAN or IP SAN Enterprise Environment

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

Data ONTAP uses two methods for automatically making more space available for a FlexVol volume when that volume is nearly full: allowing the volume size to increase, and deleting Snapshot copies.

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 (try_first option set to **volume_grow**).

This method is useful if the volume's containing aggregate has enough space to support a larger volume. You can configure Data ONTAP to increase the size in increments and set a maximum size for the volume.

- Delete Snapshot copies when the volume is nearly full (try_first option set to **snap_delete**).

For example, you can configure Data ONTAP to 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 Data ONTAP to delete first—your oldest or newest Snapshot copies. You can also determine when Data ONTAP should begin deleting Snapshot copies—for example, when the volume is nearly full or when the volume's Snapshot reserve is nearly full.

For more information about deleting Snapshot copies automatically, see the *Data ONTAP Data Protection Online Backup and Recovery Guide for 7-Mode*.

Thin provisioning for greater efficiencies using FlexVol volumes

With thin provisioning, when you create volumes and LUNs for different purposes in a given aggregate, you do not actually allocate any space for those volumes in advance. The space is allocated as data is written to the volumes.

The unused aggregate space is available to other thin provisioned volumes and LUNs. 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 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, 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 400 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 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.

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 the N series Management Console data protection capability.

Using the N series Management Console provisioning capability, 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 Guide to Common Workflows for Administrators*.

Related information:



IBM N series support website: www.ibm.com/storage/support/nseries

Best practices for using thin provisioning

You should be aware of the best practices when you use thin provisioning with other Data ONTAP features.

The following are the best practices:

- You can achieve maximum storage efficiency by using thin provisioning with FlexClone technology.
- You should add disks to increase space in the aggregates, until the maximum aggregate size limit for the system is reached.

At this stage, you can use deduplication, volume, or LUN migration to a different aggregate to add more space.

- You should not use thin provisioning for root volumes.
- You should use thin provisioning when sizing backup destination volumes using SnapVault or disaster recovery destination volumes using SnapMirror.

You should set the volume size to the maximum limit for the platform or to the size of the parent aggregate. This operation ensures that volume sizing on the primary storage system is handled automatically.

- You should use IBM N series Data Motion for vFiler and MultiStore to handle out-of-space conditions with no downtime, which enables thin provisioning to handle critical applications.

For more information about thin provisioning related best practices, see the *Thin Provisioning Deployment and Implementation Guide*.

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 3965: Thin Provisioning Deployment and Implementation Guide

Storage space management using OnCommand Unified Manager

The OnCommand Unified Manager suite consists of OnCommand Unified Manager Core and Host Packages which include front-end user interfaces (such as the OnCommand console, the N series Management Console, and the Operations Manager console) and back-end server such as the DataFabric Manager server.

The N series Management Console provisioning capability 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 adding volumes or qtrees to a dataset on existing storage and newly provisioned storage
- 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
- Dashboard views of resource status, and space status, and error and warning messages to be issued when user-set space usage and overcommitment thresholds are exceeded
- Storage services support that allows you to create preconfigured bundles of protection policies, provisioning policies, resource pools, and vFiler templates (for vFiler unit creation and attachment), that you can assign to a set of data in accordance with that data's provisioning and protection needs

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

- Discover 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
- Run Data ONTAP CLI commands simultaneously on multiple systems
- Configure role-based access control (RBAC)
- Manage host users, user groups, domain users, local users, and host roles

For more information about the Operations Manager console, see the *OnCommand Unified Manager Operations Manager Administration Guide*.

Automate thin provisioning using the N series Management Console provisioning capability

Using the N series Management Console provisioning capability, you can take advantage of thin provisioning and resource pooling to achieve better utilization of 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 management, and enables flexible and efficient use of the storage systems.

Efficient file recovery with Snapshot copies

Snapshot copies are read-only images of traditional volumes, FlexVol volumes, or aggregates that capture the state of the file system at a point in time. Snapshot copies are used for backup and disaster recovery.

Snapshot copies consume minimal storage space because only changes to the active file systems are written. You can back up your volumes more frequently and 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 requirements. This schedule creates Snapshot copies automatically and deletes old Snapshot copies after a predetermined amount of time.

SnapManager uses Snapshot technology to integrate Exchange backup, restore APIs, and Volume Shadow Copy Service (VSS). SnapManager provides an integrated data management solution for Microsoft Exchange that enhances the availability, scalability, and reliability of Exchange databases. SnapManager provides rapid online backup and restoration of databases, along with local or remote backup set mirroring for disaster recovery.

For more information about Snapshot technology, see the *Data ONTAP Data Protection Online Backup and Recovery Guide for 7-Mode*.

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:



Before any Snapshot copy is created, disk space is consumed by the active file system only.



After a Snapshot copy is created, the active file system and Snapshot copy point to the same disk blocks. The Snapshot copy does not use extra disk space.



After *myfile.txt* is deleted from the active file system, the Snapshot copy still includes the file and references its disk blocks. That is why deleting active file system data does not always free disk space.

- Space used by the active file system
- Space used by the Snapshot copy only
- Space shared by the Snapshot copy and the active file system
- Unused disk space

What the Snapshot copy reserve is

The Snapshot copy reserve sets a specific percent of the disk space for Snapshot copies. By default, the Snapshot copy reserve is 20 percent of the disk space. However, for a FlexVol volume, the Snapshot copy reserve is set to 5 percent by default. The active file system cannot consume the Snapshot copy reserve space, but the Snapshot copy reserve, if exhausted, can use space in the active file system.

Managing the Snapshot copy reserve involves the following tasks:

- Ensuring that enough disk space is allocated 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

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.

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/vol10/    3000000 3000000 0      100%
/vol/vol10/.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/vol10/    3000000 2900000 100000 97%
/vol/vol10/.snapshot 1000000 600000 400000 60%
```

Example of what happens when Snapshot copies exceed the reserve

Because there is no way to prevent Snapshot copies from consuming disk space greater than the amount reserved for them, 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.

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.


Using deduplication to increase storage efficiency

Deduplication is a Data ONTAP feature that reduces the amount of physical storage space required by eliminating duplicate data blocks within a FlexVol volume.

Deduplication works at the block level on an 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, a byte-by-byte comparison is done for all the bytes in the block, and the duplicate block is discarded and its disk space is reclaimed.

You can configure deduplication operations to run automatically or according to a schedule. You can deduplicate new and existing data, or only new data. You cannot enable deduplication on the root volume.

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



For more information about deduplication, see the *Data ONTAP Storage Management Guide for 7-Mode* .

 IBM N series support website: www.ibm.com/storage/support/nseries

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 volume with existing data, it scans all the blocks in the volume and creates a digital fingerprint for each of the blocks. Each of the fingerprints is compared to all the other fingerprints within the volume. If two fingerprints are found to be identical, a byte-by-byte comparison is done for all data within the block. If the byte-by-byte comparison detects identical data, the pointer to the data block is updated, and the duplicate block is removed.

Note: When deduplication is run on a volume with existing data, it is best to configure deduplication to scan all the blocks in the volume for better space savings.

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

The deduplication metadata includes the fingerprint file and change logs. Fingerprints are the digital signatures for every 4-KB data block in a FlexVol volume.

The deduplication metadata contains two change log files. When deduplication is running, the fingerprints of the new data blocks from one change log file are merged into the fingerprint file, and the second change log file stores the fingerprints of the new data that is written to the volume during the deduplication operation. The roles of the change log files are reversed when the next deduplication operation is run.

In Data ONTAP 8.0.1, the deduplication metadata is located within the aggregate. Starting with Data ONTAP 8.1, two copies of deduplication metadata are maintained per volume. A copy of the deduplication metadata resides in the volume and another copy is in the aggregate. The deduplication metadata in the aggregate is used as the working copy for all the deduplication operations. The change log entries are appended in the deduplication metadata copy residing in the volume.

When a volume is moved, the deduplication metadata is also transferred with the volume. If the volume ownership changes and the next time deduplication is run, the deduplication metadata which resides in the aggregate is created automatically by using the copy of deduplication metadata in the volume. This method is a faster operation than creating a new fingerprint file.

Starting with Data ONTAP 8.1.1, the change log file size limit is set to 1 percent of the volume size. The change log file size is relative to the volume size limit and the space assigned for change log files in a volume is not reserved.

Deduplication metadata can occupy up to 7 percent of the total logical data contained within the volume, as follows:

- In a volume, deduplication metadata can occupy up to 4 percent of the total logical data contained within the volume.
- In an aggregate, deduplication metadata can occupy up to 3 percent of the total logical data contained within the volume.

You can use the **df -A *aggrname*** command to check the available space in an aggregate and the **df *path*** command to check the available space in a volume. For more information about these commands, see the man pages.

Example

A 2-TB aggregate has four volumes, each 400 GB in size, in the aggregate. You need three volumes to be deduplicated, with 100 GB of data, 200 GB of data, and 300 GB of data. In this case, the volumes need 4 GB, 8 GB, and 12 GB of space, and the aggregate needs a total of 18 GB ((3% of 100 GB) + (3% of 200 GB) + (3% of 300 GB) = 3+6+9=18 GB) of space available in the aggregate.

Guidelines for using deduplication

Deduplication is a low-priority background process that consumes system resources when the deduplication operation is running on FlexVol volumes.

If the data does not change often in a volume, it is best to run deduplication less frequently. If you run multiple concurrent deduplication operations on a storage system, these operations lead to a higher consumption of system resources. It is best to begin with fewer concurrent deduplication operations. Increasing the number of concurrent deduplication operations gradually enables you to better understand the impact on the system.

Performance considerations for deduplication

Various 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 can affect the performance of deduplication:

- 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
- Disk types (for example, ATA/FC, and RPM of the disk)

For more information about performance aspects of deduplication, see the technical report *TR-3958: Data Compression and Deduplication Deployment and Implementation Guide for 7-Mode*.

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 3958: Data Compression and Deduplication Deployment and Implementation Guide for 7-Mode

Deduplication interoperability with Data ONTAP features

When you use deduplication, you should be aware of the features supported by deduplication and how they work with deduplication.

The following features are supported by deduplication:

- Snapshot copies
- Volume SnapMirror
- Qtree SnapMirror
- SnapVault
- SMTape backup
- SnapRestore
- Stretch and fabric-attached MetroCluster configurations
- DataFabric Manager server
- Volume copy
- Data compression
- FlexClone volumes
- HA pair
- vFiler units
- Data Motion for Volumes

You can enable extents on deduplicated volumes. You can perform read reallocation to improve the file layout and the sequential read performance of a deduplicated volume.

However, deduplicated volumes cannot be replicated using synchronous SnapMirror and semi-synchronous SnapMirror.

Best practices for using deduplication with other Data ONTAP features

You should be aware of the best practices when you use deduplication with other Data ONTAP features.

For more information about deduplication related best practices, see the *Data Compression and Deduplication Deployment and Implementation Guide for 7-Mode*.

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 3958: Data Compression and Deduplication Deployment and Implementation Guide for 7-Mode

Using data compression to increase storage efficiency

Data compression is a Data ONTAP feature that enables you to reduce the physical capacity that is required to store data on storage systems by compressing the data blocks within a FlexVol volume. You can use data compression only on volumes that are created on 64-bit aggregates.

You can use data compression on primary, secondary, and tertiary storage tiers.

Data compression enables you to store more data in less space. Further, you can use data compression to reduce the time and bandwidth required to replicate data during volume SnapMirror transfers.

Data compression can save space on regular files or LUNs. However, file system internal files, Windows NT streams, and volume metadata are not compressed.

Data compression can be done in the following ways:

- In-line compression

If in-line compression is enabled on a volume, the subsequent writes to the volume are compressed before they are written to the volumes.
- Post-process compression

If post-process compression is enabled on a volume, the new data writes to the volume are not compressed initially, but are rewritten as compressed data to the volume when data compression is run. The post-process compression operation runs as a low-priority background process.

If both in-line and post-process compression are enabled, then post-process compression compresses only the blocks on which in-line compression was not run. This includes blocks that were bypassed by in-line compression such as small, partial compression group overwrites.

Note: You cannot enable data compression on the storage system root volumes or on the volumes that are contained within 32-bit aggregates.

For more information about data compression, see the *Data ONTAP Storage Management Guide for 7-Mode*.

 IBM N series support website: www.ibm.com/storage/support/nseries

Guidelines for using data compression on existing data

You can achieve maximum space savings on FlexVol volumes by running data compression before creating Snapshot copies. After data compression is enabled, by default only the new data that is written to the disk is compressed. However, you can compress the data that existed on the disk prior to enabling compression.

When existing data is compressed, by default the data containing shared blocks and data that is locked in a Snapshot copy are bypassed.

You should be aware of the following guidelines when using the data compression operation:

- For source FlexVol volumes with Snapshot copies that cannot be deleted (for example, volumes created by FlexClone, SnapMirror, LUN clone, and dump operations), you can compress existing data in the default mode where data that is locked in Snapshot copies are bypassed.
- You can disable Snapshot copy creation temporarily until the data compression operation is complete.
- You should disable the deduplication schedule until the data compression operations is complete.
- For a volume that is the source for NDMP copy, dump, volume SnapMirror, qtrees SnapMirror, or SnapVault relationships, you should run the data compression operation before the initial baseline is created.
- You should disable volume SnapMirror and SnapVault schedules until the data compression scanning is complete.
- If you create a FlexClone volume when decompression is active on the parent volume, then the decompression operation does not run on the cloned volume.

Data compression interoperability with Data ONTAP features

When you use data compression, you should be aware of the features supported by data compression and how they work with data compression.

The following features are supported by data compression:

- Snapshot copies
- Volume SnapMirror
- Qtrees SnapMirror
- SnapVault
- SMTape backup
- SnapLock
- Volume-based SnapRestore

- Single file SnapRestore
- Stretch and fabric-attached MetroCluster configurations
- Volume copy
- Aggregate copy
- Deduplication
- FlexClone volumes
- FlexClone files
- HA pair
- Performance Acceleration Module or Flash cache cards
- vFiler units
- Data Motion for Volumes
- Flash Pool

Compressed volumes cannot be replicated using synchronous SnapMirror and semi-synchronous SnapMirror. Read reallocation and extents are not supported on compressed volumes.

Best practices for using data compression with other Data ONTAP features

You should be aware of the best practices when you use data compression with other Data ONTAP features.

For more information about data compression related best practices, see the *Data Compression and Deduplication Deployment and Implementation Guide for 7-Mode*.

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:



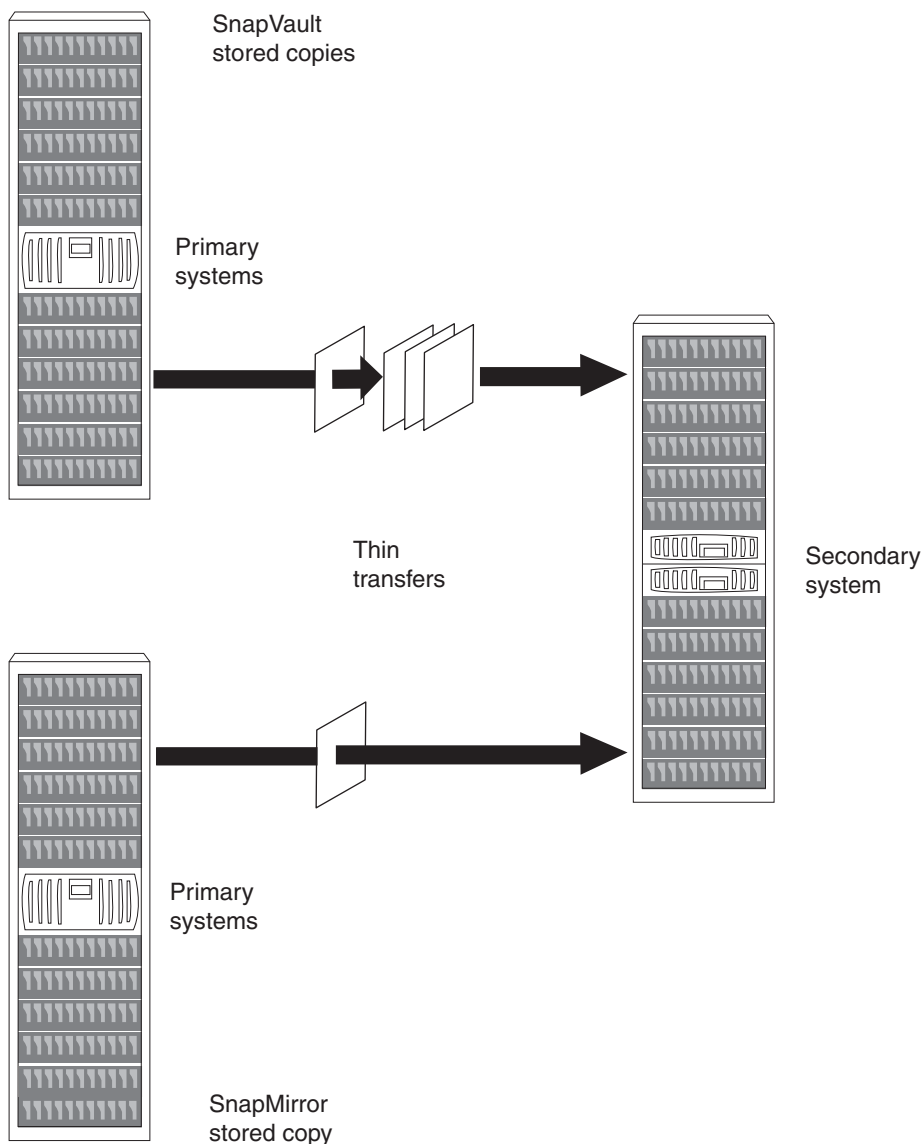
Technical Report 3958: Data Compression and Deduplication Deployment and Implementation Guide for 7-Mode

Thin replication using SnapVault and Volume SnapMirror

Thin replication refers to the copying of changed blocks to a secondary site during backup and disaster recovery. Snapshot copies can be backed up and replicated to a secondary site using SnapVault and volume SnapMirror. Thin replication increase storage efficiency by transferring only changed data blocks after the baseline copy is created.

You can combine deduplication and data compression with thin replication technologies to achieve greater savings. Retaining deduplication and data compression savings during volume SnapMirror transfers result in network bandwidth savings. Further, you can use SnapMirror network compression or compression feature of Open Systems SnapVault with these technologies to reduce the size of the replicated data and save 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:



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 duplicate copies that result from other backup technologies, such as tape.

Further, deduplication and data compression facilitates space reduction at the source and destination systems.

For more information about SnapVault and deduplication, see the *Data ONTAP Data Protection Online Backup and Recovery Guide for 7-Mode*.

 IBM N series support website: www.ibm.com/storage/support/nseries

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:

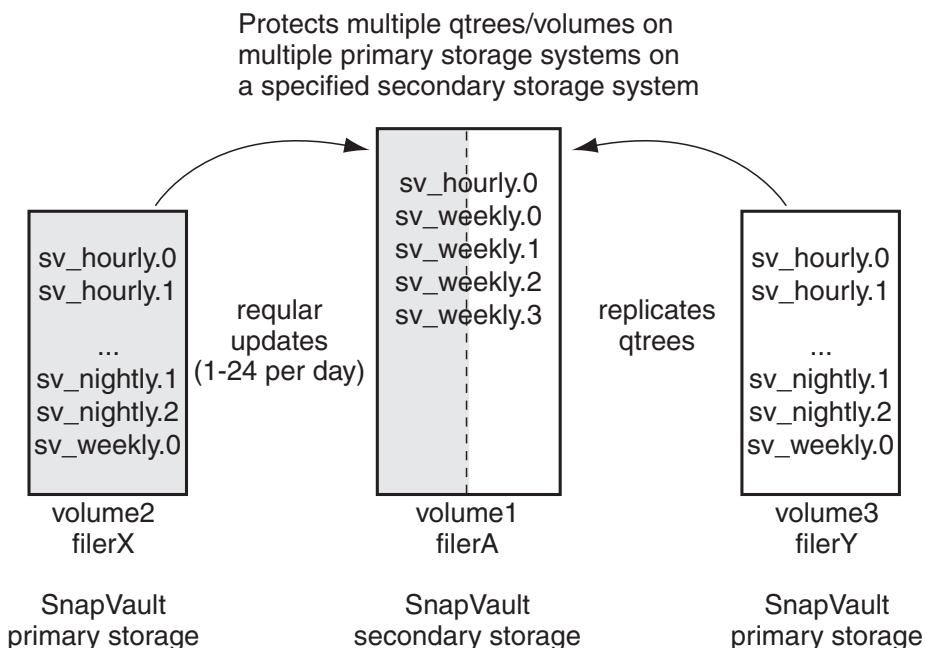


Figure 1. SnapVault functionality

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

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-qtree data, 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 NearStore system.

The following figure shows a 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 2. 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.

Best practices for using SnapVault with other Data ONTAP features

You should be aware of the best practices when you use SnapVault with other Data ONTAP features.

For more information about SnapVault related best practices, see the *SnapVault Best Practices Guide*.

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 3487: SnapVault Best Practices Guide

Using SnapVault with deduplication and data compression

You can use both deduplication and data compression with SnapVault to achieve better space savings.

The following are the best practices:

- You should run post-process compression to compress existing data on the primary system before running the baseline transfers for SnapVault.
- You should use in-line compression to achieve data compression savings on a SnapVault destination volume without affecting the Snapshot space.
- You can independently configure the deduplication and data compression schedules on a source volume, because the schedules are not tied to the SnapVault update schedule.

However, for a destination volume the deduplication and data compression schedules are tied to the SnapVault update schedule.

Using SnapVault with volume SnapMirror

You can protect the SnapVault secondary system by using volume SnapMirror, which enables you to replicate all the SnapVault destination volumes to another SnapVault secondary system at a remote site.

The following are the best practices:

- You can use SnapMirror to create a Snapshot copy of the SnapVault destination volume which is writeable.
- When you set up SnapVault schedules, you should be aware of the number of Snapshot copies that are retained for the volume, including Snapshot copies for SnapVault and SnapMirror.

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 remote storage systems 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 and data compression, it consumes minimal storage space and saves on network bandwidth. When volume SnapMirror is combined with deduplication and data compression, any savings on the SnapMirror source volume are inherited over the network and on the destination volume.

For more information about volume SnapMirror, see the *Data ONTAP Data Protection Online Backup and Recovery Guide for 7-Mode* .

Related information:

 IBM N series support website: www.ibm.com/storage/support/nseries

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 performs the following operations:

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 and 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.

Best practices for using SnapMirror with other Data ONTAP features

You should be aware of the best practices when you use SnapMirror with other Data ONTAP features.

For more information about SnapMirror related best practices, see the *SnapMirror Async Overview and Best Practices Guide*.

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 3446: SnapMirror Async Overview & Best Practices Guide

Using volume SnapMirror with deduplication

If you enable deduplication on the source volume, the deduplication savings are retained on the destination volume. The deduplication savings are extended to bandwidth savings because volume SnapMirror transfers only unique blocks.

The following are the best practices:

- You should replicate a deduplicated source volume to achieve deduplication savings on the destination volume.
You cannot run deduplication independently on the destination volume.
For example, a 100 GB source volume is deduplicated and now consumes 50 GB achieving 50% storage savings. The volume SnapMirror replication transfers only 50 GB of data, and the destination volume consumes only 50 GB data.
- You should ensure that the destination volume size is the same as or more than the source volume.
- You should start the volume SnapMirror transfer of a deduplicated volume only after the deduplication operation is complete, thereby ensuring that unduplicated data, and temporary metadata is not transferred over the network.

Using volume SnapMirror with data compression

Data compression savings are retained throughout the volume SnapMirror transfer, and this reduces network bandwidth usage.

The following are the best practices:

- You should start the volume SnapMirror transfer of a compressed volume only after the data compression operation is complete.

- You should run the baseline transfer for volume SnapMirror on the primary storage system only after post-process compression is complete.

Using volume SnapMirror with FlexClone volumes

When you create a FlexClone volume on a volume SnapMirror destination volume, Data ONTAP locks the Snapshot copy of the parent clone. You cannot delete the Snapshot copy, and therefore the clone is protected.

The following are the best practices:

- You should create FlexClone volumes on the destination volume from manually created Snapshot copies and not from scheduled Snapshot copies.
- If you want to create FlexClone volumes on the destination volume from SnapDrive created rolling Snapshot copies, you can perform one of the following actions:
 - Perform a SnapMirror update after the FlexClone volume is created. This update creates a softlock on the corresponding Snapshot copy. The next SnapDrive update creates another rolling Snapshot copy without deleting the Snapshot copy that has a softlock.
 - Rename the rolling Snapshot copy, created by SnapDrive before creating the FlexClone volume. This ensures that SnapDrive does not delete the renamed Snapshot copy.
- You should not create FlexClone volumes on the destination volume by using the volume SnapMirror Snapshot copies. If you have to create FlexClone volumes from the volume SnapMirror Snapshot copies, then use the latest SnapMirror Snapshot copy.
- You should not delete the Snapshot copies on the source volume, if a FlexClone volume exists on the destination volume that corresponds to the Snapshot copy.

Using volume SnapMirror with SnapVault

SnapVault creates backup of a specific Snapshot copy from the volume SnapMirror destination using the `-s` option.

However, you cannot create Snapshot copies on the volume SnapMirror destination, because the destination volume is read-only. SnapVault transfers the Snapshot copies from the volume SnapMirror destination.

You should schedule a SnapVault backup when the volume SnapMirror updates are not scheduled to run, which ensures that there are no volume SnapMirror update failures.

Using qtree SnapMirror with deduplication

You can run deduplication independently on the source and destination volumes. If a source volume is deduplicated, then a qtree which is replicated

does not retain the space savings on the destination volume and cannot achieve bandwidth savings on the source volume.

If you want to achieve space savings on the destination qtree, then you should run deduplication independently on the destination volume.

For example, a 100 GB source volume with one source qtree of 100 GB data is deduplicated and now consumes 50 GB. The qtree SnapMirror replication sends 100 GB data, and the destination qtree consumes 100 GB of data. If deduplication is run independently on the destination, then the consumption is reduced to 50 GB.

The best practice for using qtree SnapMirror with deduplication and post-process compression, or only deduplication, is to allow qtree SnapMirror to use minimum number of Snapshot copies that are required.

Using qtree SnapMirror with data compression

You can run data compression independently on the source and destination storage systems.

The following are the best practices:

- You should run in-line compression to achieve optimal data compression savings on a qtree SnapMirror destination.
- You should perform qtree SnapMirror updates only after post-process compression on the source volume is complete.
- When you are using qtree SnapMirror with post-process compression and deduplication, or only post-process compression, you should allow qtree SnapMirror to use minimum number of Snapshot copies that are required.

Using qtree SnapMirror with FlexClone volumes

Qtree SnapMirror does not maintain the same Snapshot copies of the volume on the source and destination systems. Due to this characteristic, when a FlexClone volume is created from a Snapshot copy on the qtree SnapMirror destination, a lock is not created for the Snapshot copy on the source volume.

The advantage of qtree SnapMirror is that a FlexClone volume can exist on the SnapMirror destination system for a long time without any space implications on the source system.

When you create a FlexClone volume from a Snapshot copy on the qtree SnapMirror destination system, it does not create a lock on the corresponding Snapshot copy in the source volume.

The following are the best practices:

- You should not delete the SnapMirror Snapshot copies on the source or destination system, because the SnapMirror update requires a common SnapMirror Snapshot copy.
- You can backup FlexClone volumes using the Snapshot copy of the source from which the clone is created.

The backed up Snapshot copy is assigned a hard-lock and cannot be deleted. If you delete the FlexClone volume, then the lock is removed.



Using FlexClone technology to create efficient copies of volumes, files, and LUNs

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

You can use FlexClone technology to save space, power, and cost. Additionally, FlexClone has the same high performance as their parent volumes.

For more information about flexible clones, see the *Data ONTAP Storage Management Guide for 7-Mode* .

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



Related information:

 IBM N series support website: www.ibm.com/storage/support/nseries

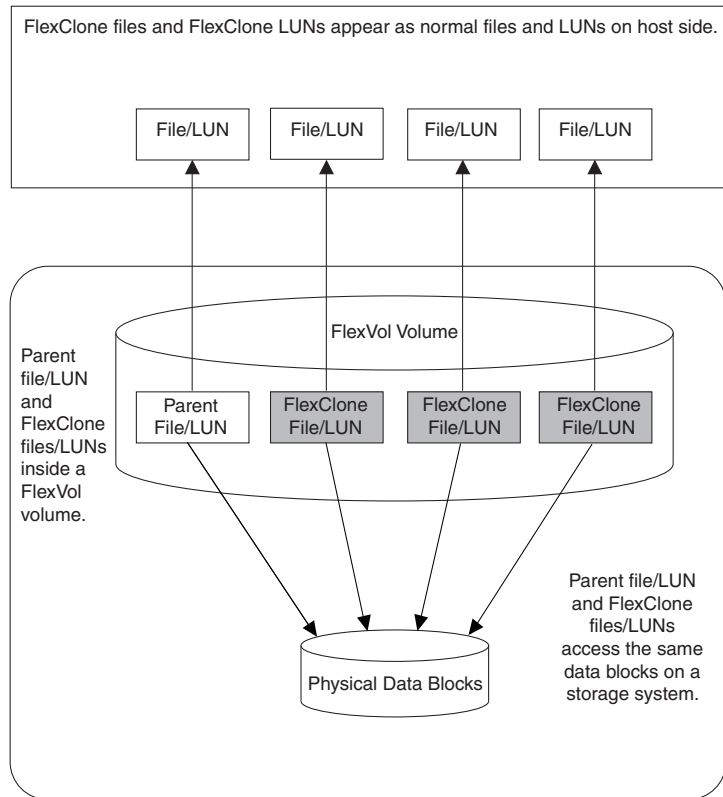
How FlexClone files and FlexClone LUNs work

FlexClone files and FlexClone LUNs share the same physical data blocks with their parent files and LUNs present in FlexVol or FlexClone volumes, and occupy negligible space in the form of metadata.

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 LUN in a SAN environment.

The cloned copies are highly space- and time-efficient because the cloning operation does not copy physical blocks of data. When you write new data to a parent or clone, then the entity on which new data is written starts occupying extra storage space.

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:



The cloning operation is instantaneous and has no impact on client access to the parent file or LUN. Clients that are accessing the parent file or LUN do not experience any disruption or outage. Clients can perform all operations on FlexClone files and FlexClone LUNs as they can on standard files and LUNs.

You can create a maximum of 32,767 FlexClone files or FlexClone LUNs from a parent file or LUN without creating a physical copy of the parent entity. If you try to create more than 32,767 clones, Data ONTAP automatically creates a new physical copy of the parent file or LUN.

The total logical size of all FlexClone files and FlexClone LUNs in a FlexVol volume is 640 TB, and is independent of the maximum size limit for the volume.

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.

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 for 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 containing the parent files and LUNs.
- 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 **sis** attribute is added to a FlexVol volume when a FlexClone file or FlexClone LUN is created for the first time.
- When clients write new data either to a FlexClone file or FlexClone LUN, or the parent file or parent LUN, then the new data occupies additional storage space.
- If you delete the FlexClone files or LUNs, the parent files or LUNs are not affected.
Deleting a parent file or LUN has no impact on the FlexClone files or FlexClone LUNs.
- If you create FlexClone files or LUNs from a Snapshot copy, you cannot create new Snapshot copies until the cloning process is complete.

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% $((10/70)*100)$.

Best practices for using FlexClone technology with other Data ONTAP features

You should be aware of the best practices when you use FlexClone volumes, FlexClone files, and FlexClone LUNs with other Data ONTAP features.

For more information about FlexClone related best practices, see the technical report *Using FlexClone to Clone Files and LUNs*.

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 3742: Using FlexClone to Clone Files and LUNs

Using FlexClone volumes with volume SnapMirror

When a volume SnapMirror transfer is in progress, you must not start a new clone operation even after completion of the first clone operation on the same FlexVol volume. Here, the clone operation uses the temporary Snapshot copies.

You must start the clone operation by using temporary Snapshot copies on the source FlexVol volume only after the volume SnapMirror transfer is complete.

Using FlexClone files or FlexClone LUNs with volume SnapMirror

For a volume SnapMirror operation to be successful, you must ensure that the size of the destination volume is reduced to less than the allowed limit for the platform you are using.

If the size of the source FlexVol volume in a volume SnapMirror relationship is less than the allowed volume size limit and the destination volume size is larger than the FlexClone limit for that platform, then creating a FlexClone file or FlexClone LUN on the source volume causes the next volume SnapMirror update to fail.

Using FlexClone files or FlexClone LUNs with qtrees SnapMirror and SnapVault

Qtree SnapMirror and SnapVault work at the logical file level, because of this the FlexClone files and FlexClone LUNs are transferred to the destination volume as different physical files and LUNs and then stored on the destination volume as different files and LUNs.

You should ensure that the size of the destination volume is equal to the sum of the size of all logical files and LUNs present on the source volume. You can use the **df** command to estimate the used and saved space on the destination volume. The size of the destination volume must be equal to or greater than the sum of the used and saved space.

Using FlexClone files or FlexClone LUNs with deduplication

You can create FlexClone files or FlexClone LUNs on FlexVol volumes in which deduplication is enabled. The **-I** option of the **clone** command enables change logging on the FlexVol volume.

Enabling change logging slows down the clone creation process. You must enable change logging only if there is a huge amount of data overwrites to the source file or LUN and you want to share these new data blocks during the deduplication operation.

Copyright and trademark information

Copyright © 1994–2012 NetApp, Inc. All rights reserved. Printed in the U.S.

Portions copyright © 2012 IBM Corporation. All rights reserved. Printed in the U.S.

U.S. Government Users Restricted Rights—Use, duplication, or disclosure restricted by GSA ADP Schedule Contract with IBM Corporation.

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. 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

NetApp, the NetApp logo, Network Appliance, the Network Appliance logo, Akorri, ApplianceWatch, ASUP, AutoSupport, BalancePoint, BalancePoint Predictor, Bycast, Campaign Express, ComplianceClock, Cryptainer, CryptoShred, Data ONTAP, DataFabric, DataFort, Decru, Decru DataFort, DenseStak, Engenio, Engenio logo, E-Stack, FAServer, FastStak, FilerView, FlexCache, FlexClone, FlexPod, FlexScale, FlexShare, FlexSuite, FlexVol, FPolicy, GetSuccessful, gFiler, Go further, faster, Imagine Virtually Anything, Lifetime Key Management, LockVault, Manage ONTAP, MetroCluster, MultiStore, NearStore, NetCache, NOW (NetApp on the Web), Onaro, OnCommand, ONTAPI, OpenKey, PerformanceStak, RAID-DP, ReplicatorX, SANscreen, SANshare, SANtricity, SecureAdmin, SecureShare, Select, Service Builder, Shadow Tape, Simplicity, Simulate ONTAP, SnapCopy, SnapDirector, SnapDrive, SnapFilter, SnapLock, SnapManager, SnapMigrator, SnapMirror, SnapMover, SnapProtect, SnapRestore, Snapshot, SnapSuite, SnapValidator, SnapVault, StorageGRID, StoreVault, the StoreVault logo, SyncMirror, Tech OnTap, The evolution of storage, Topio, vFiler, VFM, Virtual File Manager, VPolicy, WAFL, Web Filer, and XBB are trademarks or registered trademarks of NetApp, Inc. in the United States, other countries, or both.

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 www.ibm.com/legal/copytrade.shtml.

Apple is a registered trademark and QuickTime is a trademark of Apple, Inc. in the United States and/or other countries. Microsoft is a registered trademark and Windows Media is a trademark of Microsoft Corporation in the United States and/or other countries. RealAudio, RealNetworks, RealPlayer, RealSystem, RealText, and RealVideo are registered trademarks and RealMedia, RealProxy, and SureStream are trademarks of RealNetworks, Inc. in the United States and/or other countries.

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

NetApp, Inc. is a licensee of the CompactFlash and CF Logo trademarks.

NetApp, Inc. 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:

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.

Index

C

configuration of thin provisioning with Provisioning Manager 22

D

data compression 35
 best practices 37
 interoperability with Data ONTAP features 36
data compression scanner
 data compression scanner
 recommendations 36
 guidelines 36
deduplication
 best practices 33
 FlexVol volumes
 deduplication guidelines 31
 guidelines for running 31
 how it works 30
 increasing storage efficiency using 29
 interoperability with Data ONTAP features 32
 metadata relocated 30
 performance considerations 31
disk space
 recovery of 26

F

file folding
 defined 27
file reservations
 how they work 15
file system
 recovery of disk space for use by the 26
Flash Cache 7
FlexClone
 best practices 55
 FlexClone files with deduplication
 FlexClone LUNs with deduplication 57
 FlexClone files with qtree SnapMirror 56
 FlexClone files with SnapVault 56
 FlexClone files with volume SnapMirror 56
 FlexClone LUNs with qtree SnapMirror 56
 FlexClone LUNs with SnapVault 56
 FlexClone LUNs with volume SnapMirror 56
 FlexClone volumes with volume SnapMirror 56
 storage efficiency 55
FlexClone files
 how they work 52
FlexClone files and FlexClone LUNs
 considerations 55

FlexClone files and FlexClone LUNs (*continued*)

 space saving 55
 uses 54

FlexClone LUNs
 how they work 52

FlexClone volumes
 space guarantees and 54

FlexVol volumes
 automatically adding space for 17
 how fractional reserve works on 16
 how volume guarantees work with 14
 thick provisioning for 14
 thin provisioning for 14
 try_first volume option 17

fractional reserve
 how it works 16

free space
 automatically increasing 17

G

guidelines
 for running deduplication 31

H

hardware efficiency
 Flash Cache 7
 Flash Pool 9
 SATA drives 7
how to achieve storage efficiency 5
 using IBM N series products 5

I

IBM solutions to industry storage challenges 5

L

LUN reservations
 how they work 15

R

RAID-DP 11
 protection against double disk failures 11
reservations
 how they work 15

S

SATA disks 7
SnapMirror 45
 best practices 47
 qtree replication 46
 qtree SnapMirror with data compression 49

- SnapMirror (*continued*)
 - qtree SnapMirror with deduplication 49
 - qtree SnapMirror with FlexClone 49
 - storage efficiency 47
 - volume replication 46
 - volume SnapMirror with data compression 47
 - volume SnapMirror with deduplication 47
 - volume SnapMirror with FlexClone 48
 - volume SnapMirror with SnapVault 48
- Snapshot copies 23
- SnapVault 41
 - basic deployment 42
 - best practices 44
 - secondary to SnapMirror deployment 44
 - SnapVault with data compression 45
 - SnapVault with deduplication 45
 - SnapVault with volume SnapMirror 45
 - storage efficiency 44
- SnapVault deployment
 - introduction 42
- software efficiency
 - FlexClone technology 51
 - FlexVol volumes 19
 - RAID-DP 11
 - Snapshot copy 23
 - SnapVault 39, 41
 - volume SnapMirror 39, 45
- space
 - increasing for full FlexVol volumes 17
- space guarantees
 - See* volume guarantees
- space management
 - DataFabric Manager 21
 - overview 13
- space reservations
 - See* reservations
- storage efficiency
 - about 5
 - data compression 35
 - N series Management Console 21
 - Operations Manager 21

T

- tape backup of a SnapVault secondary 43
- thick provisioning
 - for FlexVols 14
- thin provisioning
 - automation using Provisioning Manager 22
 - benefits 20
 - best practices 20
 - for FlexVols 14
 - using FlexVol volumes 19
- try_first volume option 17

V

- volume guarantees
 - how they work with FlexVol volumes 14
- volume SnapVault and SnapMirror 39
- volumes
 - automatically adding space for 17
 - how fractional reserve works on 16

W

- what is storage efficiency 5



NA 210-05701_A0, Printed in USA

GA32-1047-03

