




Updated for 8.2.1

Data ONTAP[®] 8.2

Storage Efficiency Management Guide

For 7-Mode



NetApp, Inc.
495 East Java Drive
Sunnyvale, CA 94089
U.S.

Telephone: +1 (408) 822-6000
Fax: +1 (408) 822-4501
Support telephone: +1 (888) 463-8277
Web: www.netapp.com
Feedback: doccomments@netapp.com

Part number: 215-08530_A0
November 2013

Contents

What storage efficiency is	6
Storage efficiency features	6
Storage efficiency with Flash Cache and Flash Pool aggregate	8
Protection against double-disk failure with RAID-DP	9
About RAID-DP protection	9
How Data ONTAP uses RAID to protect your data and data availability	10
Efficient utilization of storage using space management capabilities ...	11
How volume guarantees work with FlexVol volumes	11
How file and LUN reservations work	13
Considerations for setting fractional reserve	13
How Data ONTAP can automatically provide more space for full FlexVol volumes	15
Thin provisioning for greater efficiencies using FlexVol volumes	16
Benefits of using thin provisioning	16
Best practices for using thin provisioning	17
Efficient file recovery with Snapshot copies	18
What Snapshot disk consumption is	18
How Snapshot copies consume disk space	18
What the Snapshot copy reserve is	19
Use of deleted active file disk space	19
Example of what happens when Snapshot copies exceed the reserve	20
Recovery of disk space for file system use	21
What file folding means and how it saves disk space	21
Storage efficiency with deduplication and data compression	22
How to estimate space savings achieved with deduplication and data compression	22
Using deduplication to increase storage efficiency	22
How deduplication works	22
What deduplication metadata is	23
Guidelines for using deduplication	24
Deduplication interoperability with Data ONTAP features	25
Using data compression to increase storage efficiency	26

Guidelines for using data compression on existing data	26
How detecting incompressible data saves system resources	27
Data compression interoperability with Data ONTAP features	27
Best practices for using deduplication and data compression with other Data ONTAP features	28
Thin replication using SnapVault and Volume SnapMirror	29
Disk-to-disk backups using SnapVault	31
How SnapVault backup works	31
Best practices for using SnapVault with other Data ONTAP features	32
Using SnapVault with deduplication and data compression	32
Using SnapVault with volume SnapMirror	33
Efficient data protection using volume SnapMirror	33
How SnapMirror works	34
Best practices for using SnapMirror with other Data ONTAP features	34
Using volume SnapMirror with deduplication	35
Using volume SnapMirror with data compression	35
Using volume SnapMirror with FlexClone volumes	35
Using volume SnapMirror with SnapVault	36
Using qtree SnapMirror with deduplication	36
Using qtree SnapMirror with data compression	36
Using qtree SnapMirror with FlexClone volumes	37
Using FlexClone technology to create efficient copies of volumes, files, and LUNs	38
How FlexClone files and FlexClone LUNs work	39
FlexClone volumes and space guarantees	41
Uses of FlexClone files and FlexClone LUNs	41
Considerations for working with FlexClone files and FlexClone LUNs	42
Space savings achieved by using FlexClone files and FlexClone LUNs	42
Best practices for using FlexClone technology with other Data ONTAP features	43
Using FlexClone volumes with volume SnapMirror	43
Using FlexClone files or FlexClone LUNs with volume SnapMirror	43
Using FlexClone files or FlexClone LUNs with qtree SnapMirror and SnapVault	43
Using FlexClone files or FlexClone LUNs with deduplication	44
Copyright information	45
Trademark information	46

How to send your comments	47
Index	48

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. NetApp 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, RAID-DP, Flash Cache, and Flash Pool aggregate 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 aggregate 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.

Storage efficiency features

NetApp provides features such as SATA disks, Flash Cache, RAID-DP, thin provisioning, Snapshot copy, deduplication, data compression, SnapVault, SnapMirror, Flash Pool aggregate, and FlexClone, which help achieve significant improvements in storage utilization.

When the storage efficiency features are used together, you can maximize the space savings achieved on the storage system. NetApp offers the following technologies to implement storage efficiency:

- RAID-DP protects your data from double-disk failures which significantly increases fault tolerance.
- Flash Cache lowers the cost by reducing the number of drives and enables you to use less expensive drives and achieve the same performance.
- Flash Pool aggregate combines solid-state drives (SSDs) and HDDs in one aggregate. The HDDs provide the storage capacity and the SSDs are used as cache for the volumes that are provisioned on the Flash Pool aggregate.
- Thin provisioning enables you to provision more logical storage space to the hosts than the available physical storage space and allocate storage space on demand instead of provisioning preallocated storage space.
- Snapshot copies are read-only images of 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 aggregate, 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.

Storage efficiency with Flash Cache and Flash Pool aggregate

Flash Cache and Flash Pool aggregate are flash-based storage caching technologies that you can use individually (or together) to increase storage efficiency. Both technologies enable storage system configurations to use fewer or lower cost per gigabyte hard disk drives (HDDs), compared to systems configured only with disk drives, while also improving performance.

Flash Cache uses PCIe cards to provide additional controller-level read cache. It enables you to improve the read response times of application workloads that include a high mix of repetitive random or sequential reads. This is often accomplished while using fewer HDDs, or slower (lower RPM) and lower cost per GB HDDs, thereby reducing storage cost and increasing storage efficiency.

Flash Cache also works with other Data ONTAP storage efficiency features. The space savings that deduplication and FlexClone provide for storage are realized in cache as well; one physical block in cache can represent multiple logical copies of data. Data compression is supported with Flash Cache, however, only uncompressed blocks are cached.

The Flash Pool aggregate combines solid-state drives (SSDs) and HDDs in one aggregate. The HDDs provide the storage capacity and the SSDs are used as cache for the volumes that are provisioned on the Flash Pool aggregate. Flash Pool technology improves performance by reducing response times for repetitive random reads. It also offloads repeated updates of previously written random data to SSDs, thereby decreasing the workload required from much slower HDDs. Both caching capabilities can help reduce the quantity or speed of HDDs needed for many application workloads, which results in increased storage efficiency.

Flash Pool technology works with deduplication, FlexClone, and data compression in the same way as Flash Cache.

The following are the benefits of using Flash Cache and the Flash Pool aggregate:

- Using fewer HDDs, and often also slower and lower cost per GB HDDs, delivers faster read response times when compared to HDD-only configurations.
- Configuring fewer or lower cost per GB HDDs provides the additional efficiency benefits of lower power consumption, lower cooling demand, and lower data center space consumption.
- Flash Cache and Flash Pool cache deduplicated and cloned blocks, thereby increasing their ability to cache more unique data.

For more information about Flash Pool aggregate, see *Technical Report 4070: Flash Pool Design and Implementation Guide* and for more information about Flash Cache, see *Technical Report 3832: Flash Cache Best Practices Guide*.

Related information

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

Flash Cache Best Practices Guide: media.netapp.com/documents/tr-3832.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: In a double-disk failure, 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

[Documentation on the NetApp Support Site: support.netapp.com](http://support.netapp.com)

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 RAID group even if one or two of those drives fail. Parity drives provide redundancy for the data stored in the data drives. If a drive fails (or, for RAID-DP, up to two drives), the RAID subsystem can use the parity drives to reconstruct the data in the drive that failed.

For array LUNs, 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.

Efficient utilization of storage using space management capabilities

The space management capabilities of Data ONTAP enable you to effectively use the storage provided by FlexVol volumes. You can balance the overall available storage space against the storage space required by users and applications.

Data ONTAP enables space management by using the following capabilities:

- **Volume (space) guarantee**
The *volume guarantee*, also called *space guarantee* or *guarantee*, determines, whether or not, space is preallocated in the volume's associated aggregate when the volume is created.
- **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.
- **Fractional reserve**
Fractional reserve, also called *fractional overwrite reserve* or *LUN overwrite reserve*, enables you to control the size of the overwrite reserve for a FlexVol volume. By using this volume attribute correctly, you can maximize your storage utilization.
- **Volume autosize and Snapshot autodelete**
Volume autosize and Snapshot autodelete can either increase the size of a volume or delete Snapshot copies to prevent a volume from running out of space without any manual intervention.

You can use these capabilities together to determine, on a volume-by-volume basis, whether to improve storage utilization, ease of management, or both.

You can also use the storage space management capabilities using OnCommand Unified Manager. For more information about storage space management capabilities using OnCommand Unified Manager, see the *OnCommand Unified Manager Provisioning and Data Protection Online Help for 7-Mode*.

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 space for user data is not preallocated.

The guarantee is an attribute of the volume.

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.

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 Snapshot copies) 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: Thin Provisioning Deployment and Implementation Guide*. You should also be aware that volume guarantees of type `file` will not be supported in a future release of Data ONTAP.

- 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*. The amount of space consumed by volumes with this guarantee type grows as data is added instead of being determined by the initial volume size, which might leave space unused if the volume data does not grow to that size. The maximum size of a volume with a guarantee of `none` is not limited by the amount of free space in its aggregate. It is possible for the total size of all volumes associated with an aggregate to exceed the amount of free space for the aggregate.

Writes to LUNs or files (including space-reserved LUNs and 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: Thin Provisioning Deployment and Implementation Guide* for information about how doing so can affect storage availability.

When space in the aggregate is allocated for a volume or `file` guarantee for an existing volume, that space is no longer considered free in the aggregate, even if the volume is not yet using the space. 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.

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 try to bring that volume back online, if there is insufficient available space in the aggregate to fulfill its guarantee, it will remain offline. You must force the volume online, at which point the volume's guarantee will be disabled.

Related information

Technical Report: Thin Provisioning Deployment and Implementation Guide: media.netapp.com/documents/tr-3965.pdf

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.

Considerations for setting fractional reserve

Fractional reserve, also called *LUN overwrite reserve*, enables you to control the size of the overwrite reserve for reserved LUNs and files in a volume. By using this volume attribute correctly you can maximize your storage utilization, but you should understand how it interacts with other technologies.

The fractional reserve setting is expressed as a percentage; the only valid values are 0 and 100 percent.

Setting fractional reserve to 0 increases your storage utilization. However, an application accessing data residing in the volume could experience a data outage if the volume is out of free space, even with the volume guarantee set to `volume`, when any of the following technologies and Data ONTAP features are in use:

- Deduplication
- Compression
- FlexClone files
- FlexClone LUNs
- Virtual environments

14 | Storage Efficiency Management Guide

If you are using one or more of these technologies with no fractional reserve, and you need to prevent errors due to running out of space, you must use all of the following configuration settings for the volume:

- Volume guarantee of `volume`
- File or LUN reservations `enabled`
- Volume Snapshot copy automatic deletion enabled with a commitment level of `destroy`

Note: If your rate of change is high, in rare cases the Snapshot copy automatic deletion could fall behind, resulting in the volume running out of space, even with all of the required configuration settings in use.

In addition, you can optionally use the volume autogrow capability to decrease the likelihood of volume Snapshot copies needing to be deleted automatically. If you enable the autogrow capability, you must monitor the free space in the associated aggregate. If the aggregate becomes full enough that the volume is prevented from growing, more Snapshot copies will probably be deleted as the free space in the volume is depleted.

If you do not want to monitor aggregate free space or have volume Snapshot copies automatically deleted, you can set the volume's fractional reserve setting to 100. This requires more free space up front, but guarantees that data modification operations will succeed even when the technologies listed above are in use.

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

Volume guarantee	Default fractional reserve	Allowed values
Volume	100	0, 100
None	0	0, 100
File	100	100

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*

Related information

Technical Report: Thin Provisioning Deployment and Implementation Guide: media.netapp.com/documents/tr-3965.pdf

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

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

Data ONTAP uses two methods for automatically providing more space for a FlexVol volume when that volume is nearly full: allowing the volume size to increase, and deleting Snapshot copies (with any associated storage objects). If you enable both of these methods, you can specify which method Data ONTAP should try first.

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 (known as the *autogrow* feature).
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. The increase is automatically triggered based on the amount of data being written to the volume in relation to the current amount of used space and any thresholds set.
- Delete Snapshot copies when the volume is nearly full.
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*.

If you enable both of these methods, you can specify which method Data ONTAP tries first when a volume is nearly full. If the first method does not provide sufficient additional space to the volume, Data ONTAP tries the other method next. By default, Data ONTAP tries to increase the size of the volume first.

Thin provisioning for greater efficiencies using FlexVol volumes

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

The unused aggregate space is available to other 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

FlexVol technology enables you to oversubscribe free space so that your environment can adapt rapidly to your changing business needs.

The benefits of using thin provisioning are as follows:

- Allows storage to be provisioned but 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 NetApp Management Console data protection capability.

Using the NetApp 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

[Documentation on the NetApp Support Site: support.netapp.com](http://support.netapp.com)

Best practices for using thin provisioning

Thin provisioning is effective when the data increase rate in the volume is predictable.

You should be aware of the following 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 NetApp DataMotion 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 *Technical Report 3965: Thin Provisioning Deployment and Implementation Guide*.

Related information

[Thin Provisioning Deployment and Implementation Guide: media.netapp.com/documents/tr-3965.pdf](http://media.netapp.com/documents/tr-3965.pdf)

Efficient file recovery with Snapshot copies

Snapshot copies are read-only images of traditional volumes, FlexVol volumes, FlexClone volumes, or aggregates that capture the state of the file system at a point in time. You use Snapshot copies 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 simultaneously 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:



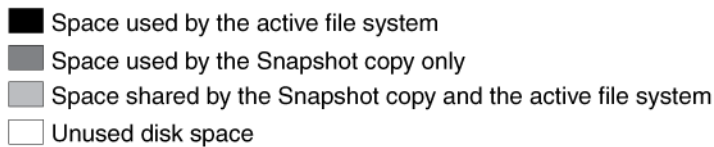
Before any Snapshot copy is created, disk space is consumed by the active file system only.



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.



What the Snapshot copy reserve is

The Snapshot copy reserve sets a specific percent of the disk space for Snapshot copies. For traditional volumes, the default Snapshot copy reserve is set to 20 percent of the disk space. For FlexVol volumes, the default Snapshot copy reserve is set to 5 percent of the disk space.

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.

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/vol0/         3000000 3000000 0      100%
/vol/vol0/.snapshot 1000000 500000 500000 50%
```

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

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

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

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, they begin to occupy the active file system space. This process is called Snapshot spill. When the Snapshot copies continue to occupy the active file system space, the system is in danger of becoming full. If the system becomes full due to Snapshot spill, you can create files only after you delete enough Snapshot copies.

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.

Storage efficiency with deduplication and data compression

Deduplication and data compression are Data ONTAP features that you can use together or independently on a FlexVol volume to achieve optimal space savings. You can provision volumes with deduplication and data compression along with other Data ONTAP features to increase the space savings achieved.

How to estimate space savings achieved with deduplication and data compression

You can use the space savings estimation tool (SSET) to estimate the amount of space savings that can be achieved using deduplication and data compression on a specific set of data.

You can download the SSET from communities.netapp.com/docs/DOC-18699.

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. Storing only unique blocks in the FlexVol volume improves the storage utilization and increases the storage efficiency you can achieve.

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

Related information


[Documentation on the NetApp Support Site: support.netapp.com](https://support.netapp.com)

How deduplication works

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

Each block of data has a digital signature that is compared with all other signatures in a data volume. If an exact block signature match exists, a byte-by-byte comparison is done for all the bytes in the block. Only if all the bytes match, the duplicate block is discarded and its disk space is reclaimed resulting in no data loss.

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



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. An additional copy of the deduplication metadata resides in the volume.

When a volume is moved, the deduplication metadata is also transferred with the volume. If the volume ownership changes, the next time deduplication is run, then 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 8.2, the fingerprints are stored for each physical block, this reduces the amount of space required to store the deduplication metadata.

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

- In a volume, deduplication metadata can occupy up to 4 percent of the total amount of data contained within the volume.
- In an aggregate, deduplication metadata can occupy up to 3 percent of the total physical 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 varying savings percentage on each volume.

The space required in the different volumes for deduplication metadata is as follows:

- 2 GB [$4\% \times (50\% \text{ of } 100 \text{ GB})$] for a 100 GB of logical data with 50 percent savings
- 6 GB [$4\% \times (75\% \text{ of } 200 \text{ GB})$] for a 200 GB of logical data with 25 percent saving
- 3 GB [$4\% \times (25\% \text{ of } 300 \text{ GB})$] for a 300 GB of logical data with 75 percent savings

The aggregate needs a total of 8.25 GB [$(3\% \times (50\% \text{ of } 100 \text{ GB})) + (3\% \times (75\% \text{ of } 200 \text{ GB})) + (3\% \times (25\% \text{ of } 300 \text{ GB})) = 1.5+4.5+2.25= 8.25 \text{ GB}$] of space available in the aggregate for deduplication metadata.

Guidelines for using deduplication

Deduplication runs as a system operation and 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 *Technical Report 3958: Data Compression and Deduplication Deployment and Implementation Guide for 7-Mode*.

Related information

[Data Compression and Deduplication Deployment and Implementation Guide for 7-Mode: media.netapp.com/documents/tr-3958.pdf](http://media.netapp.com/documents/tr-3958.pdf)

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
- OnCommand Unified Manager server
- Volume copy
- Data compression
- FlexClone volumes
- HA pair
- vFiler units

- DataMotion for Volumes
- SnapLock

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.

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.

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

Related information

[Documentation on the NetApp Support Site: support.netapp.com](http://support.netapp.com)

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 is 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 default mode, in which data that is locked in Snapshot copies is 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 operation is complete.
- For a volume that is the source for NDMP copy, dump, volume SnapMirror, qtree 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.

How detecting incompressible data saves system resources

Starting with Data ONTAP 8.2, inline compression enables you to determine if a file is incompressible and, for large files, determine if a compression group within the file is incompressible. Detecting incompressible data enables you to save the system resources used by inline compression to compress files or compression groups that are not compressible.

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
- Qtree 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
- Flash cache cards
- vFiler units
- DataMotion for Volumes
- Flash Pool aggregates

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 deduplication and data compression with other Data ONTAP features

You should be aware that there are best practices for using deduplication and data compression together or independently with other Data ONTAP features.

The following are some of the best practices:

- You should run deduplication or data compression operations before the Snapshot copies are created.
This minimizes the amount of data that is locked in the Snapshot copies.
- You should delete as many Snapshot copies as possible before running data compression on existing data and disable Snapshot copy schedule until compression of existing data is complete.
- SnapRestore retains the space savings achieved on a volume.
During a SnapRestore operation, deduplication metadata is restored.
- If you require only inline compression without incompressible data detection to run on a vault destination volume, then you should set the deduplication and postprocess compression schedule to `manual` on the destination volume.
When you set the schedule to `manual`, the efficiency operations do not run on the volume after each incremental update.
- When you run efficiency operations on existing data in the source volume, you should wait for the efficiency operations to complete before initiating the volume SnapMirror baseline transfer.

For more information about deduplication and data compression best practices, see *Technical Report 3958: Data Compression and Deduplication Deployment and Implementation Guide for 7-Mode*.

Related information

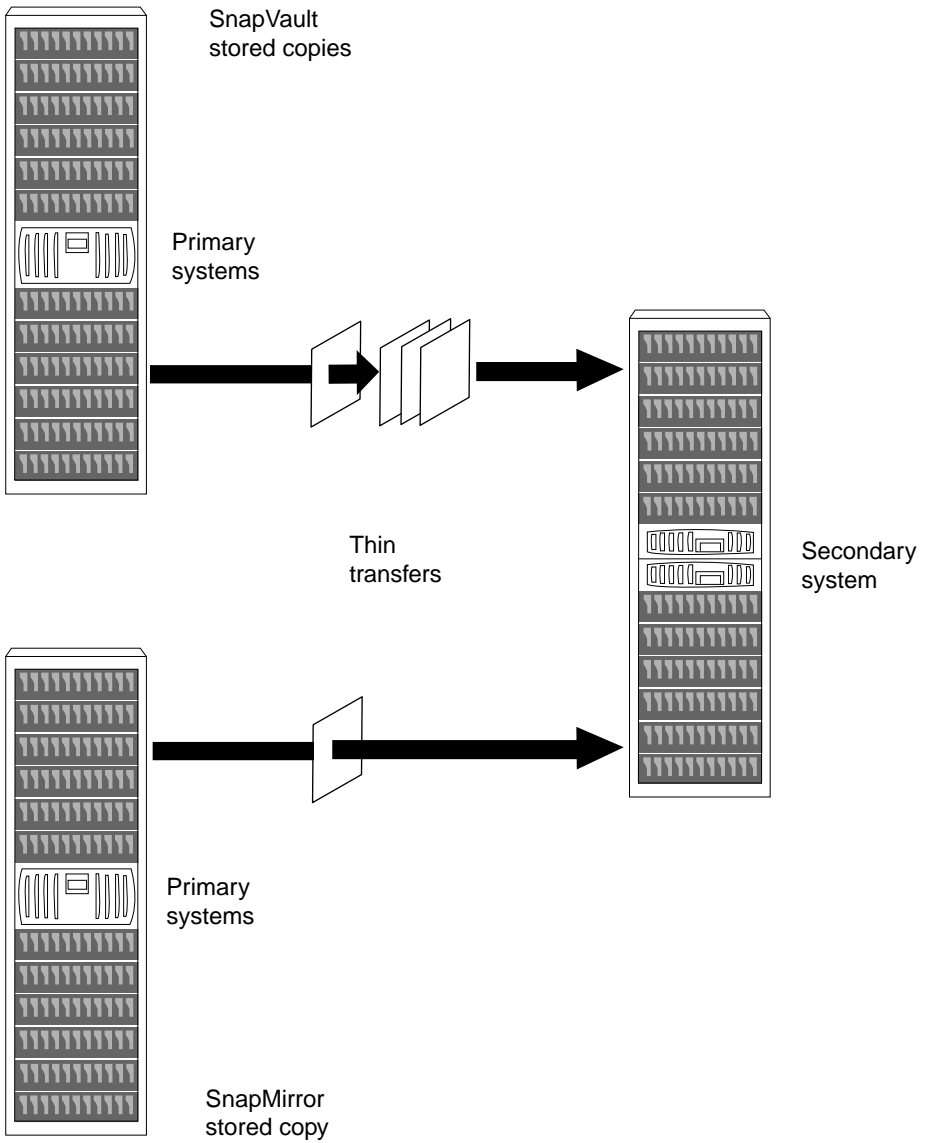
[Data Compression and Deduplication Deployment and Implementation Guide for 7-Mode: *media.netapp.com/documents/tr-3958.pdf*](https://media.netapp.com/documents/tr-3958.pdf)

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*.

Related information

[Documentation on the NetApp Support Site: support.netapp.com](http://support.netapp.com)

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.

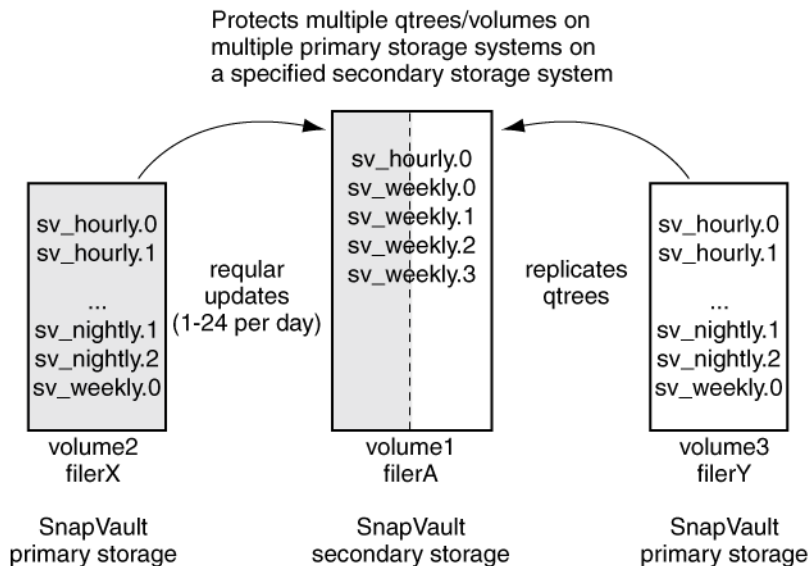
How to make scheduled incremental transfers:

- 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 the SnapVault functionality:



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 *Technical Report 3487: SnapVault Best Practices Guide*.

Related information

SnapVault Best Practices Guide: media.netapp.com/documents/tr-3487.pdf

Using SnapVault with deduplication and data compression

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

You should be aware of the following best practices:

- You should run postprocess compression to compress existing data on the primary system before running the baseline transfers for SnapVault.

- You should use inline 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.

You should be aware of the following 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

[Documentation on the NetApp Support Site: support.netapp.com](http://support.netapp.com)

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, which can be 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 of the following replication methods consists of a pair of operations, one operation each at the source storage system and the destination storage system:

- Volume SnapMirror replication
- Qtree SnapMirror replication
- SnapVault replication

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 *Technical Report 3446: SnapMirror Async Overview and Best Practices Guide*.

Related information

[SnapMirror Async Overview & Best Practices Guide: media.netapp.com/documents/tr-3446.pdf](http://media.netapp.com/documents/tr-3446.pdf)

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.

You must be aware of the following 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 undeduplicated data and temporary metadata are 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.

You should be aware of the following 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 postprocess 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.

You should be aware of the following 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 postprocess 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.

You should be aware of the following best practices:

- You should run inline compression to achieve optimal data compression savings on a qtree SnapMirror destination.
- You should perform qtree SnapMirror updates only after postprocess compression on the source volume is complete.
- When you are using qtree SnapMirror with postprocess compression and deduplication, or only postprocess 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.

You should be aware of the following 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 back up FlexClone volumes using the Snapshot copy of the source from which the clone is created.

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

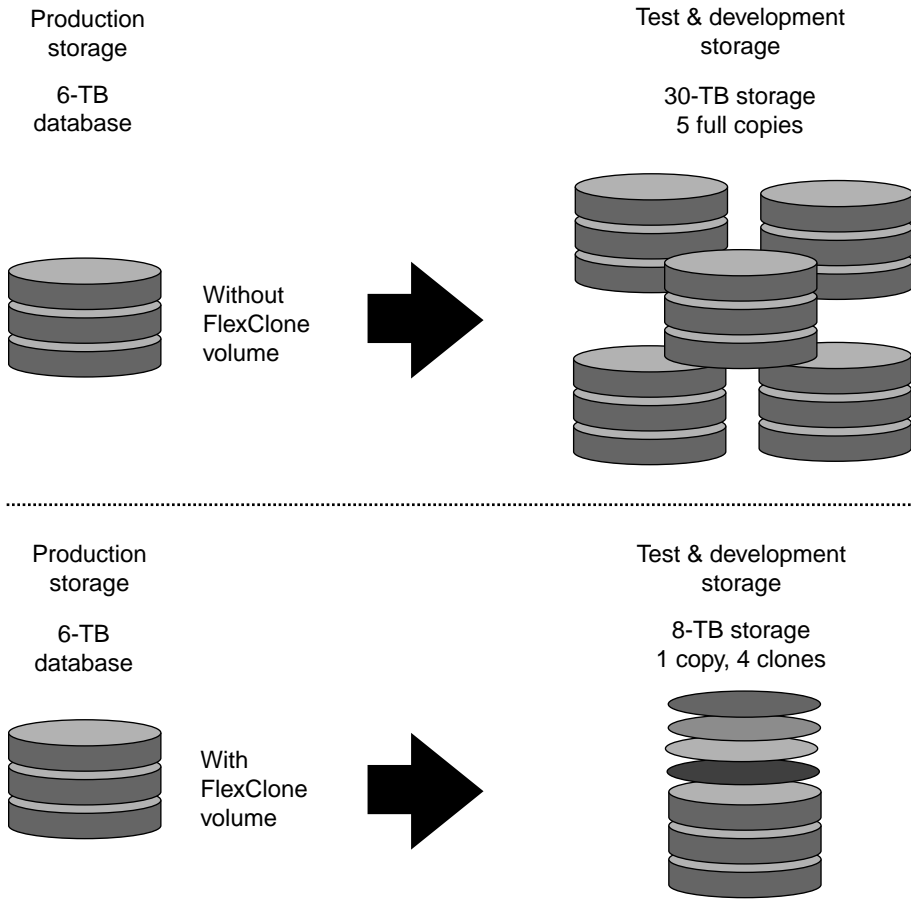
Using FlexClone technology to create efficient copies of volumes, files, and LUNs

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

[Documentation on the NetApp Support Site: support.netapp.com](http://support.netapp.com)

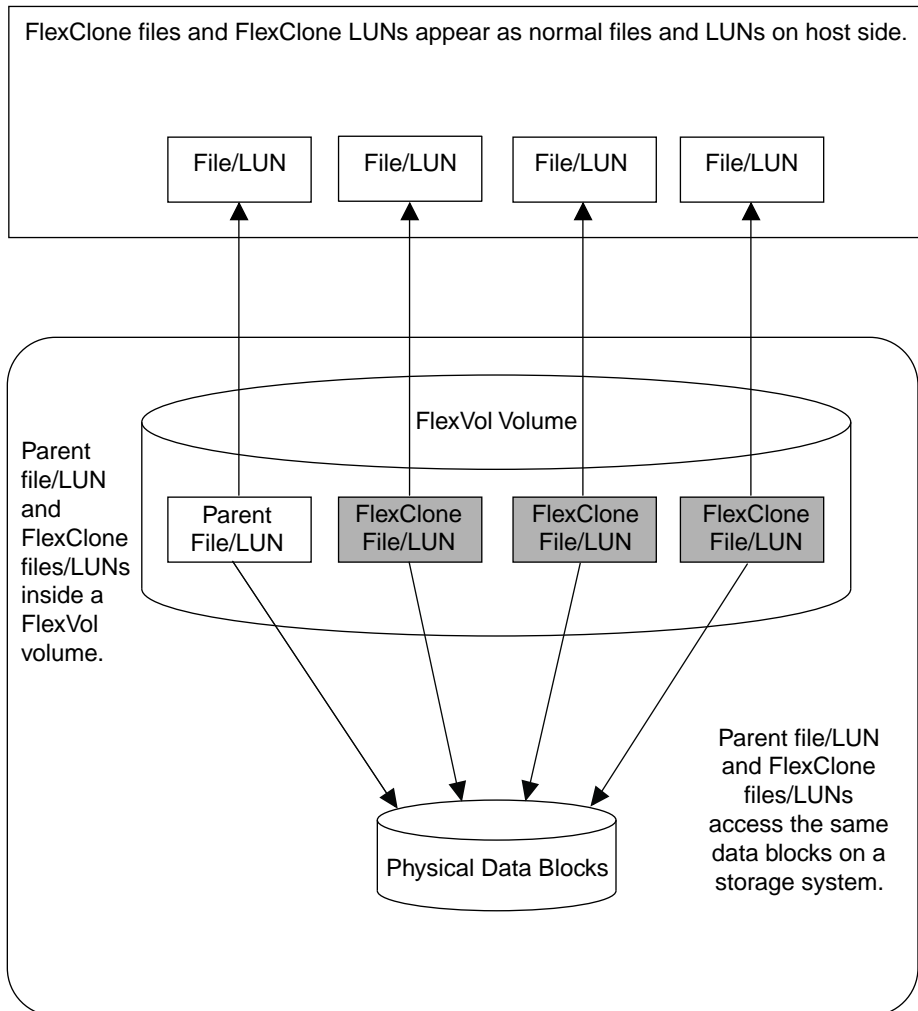
How FlexClone files and FlexClone LUNs work

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-efficient 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.

Note: FlexClone LUNs are not the same as LUN clones. You can create a FlexClone LUN without a backing Snapshot copy of the parent LUN. However, to create a LUN clone, you must have a backing Snapshot copy of the parent LUN. For more information about LUN clones, see the *Data ONTAP SAN Administration Guide for 7-Mode*.

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 working with FlexClone files and FlexClone LUNs

You should keep several considerations in mind when working with 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.
- If a FlexVol volume containing FlexClone files and LUNs has the fractional reserve set to zero, you must follow extra configuration and management requirements to ensure that the applications accessing the FlexVol volume do not receive errors due to lack of space (ENOSPC).

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 *Technical Report 3742: Using FlexClone to Clone Files and LUNs*.

Related information

Using FlexClone to Clone Files and LUNs: media.netapp.com/documents/tr-3742.pdf

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 qtree 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 information

Copyright © 1994–2013 NetApp, Inc. All rights reserved. Printed in the U.S.

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, CyberSnap, Data Center Fitness, Data ONTAP, DataFabric, DataFort, Decru, Decru DataFort, DenseStak, Engenio, Engenio logo, E-Stack, ExpressPod, FAServer, FastStak, FilerView, Flash Accel, Flash Cache, Flash Pool, FlashRay, FlexCache, FlexClone, FlexPod, FlexScale, FlexShare, FlexSuite, FlexVol, FPolicy, GetSuccessful, gFiler, Go further, faster, Imagine Virtually Anything, Lifetime Key Management, LockVault, Mars, 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, Snap Creator, SnapDirector, SnapDrive, SnapFilter, SnapIntegrator, SnapLock, SnapManager, SnapMigrator, SnapMirror, SnapMover, SnapProtect, SnapRestore, Snapshot, SnapSuite, SnapValidator, SnapVault, StorageGRID, StoreVault, the StoreVault logo, SyncMirror, Tech OnTap, The evolution of storage, Topio, VelocityStak, 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.

How to send your comments

You can help us to improve the quality of our documentation by sending us your feedback.

Your feedback is important in helping us to provide the most accurate and high-quality information. If you have suggestions for improving this document, send us your comments by email to doccomments@netapp.com. To help us direct your comments to the correct division, include in the subject line the product name, version, and operating system.

You can also contact us in the following ways:

- NetApp, Inc., 495 East Java Drive, Sunnyvale, CA 94089 U.S.
- Telephone: +1 (408) 822-6000
- Fax: +1 (408) 822-4501
- Support telephone: +1 (888) 463-8277

Index

- A**
- about
 - snapshot copy reserve [19](#)
 - active file system [19](#)
 - autogrow
 - how Data ONTAP can add space for FlexVol volumes automatically [15](#)
 - autosizing
 - how Data ONTAP can add space for FlexVol volumes automatically [15](#)
 - availability
 - how Data ONTAP uses RAID to ensure data [10](#)
- B**
- benefits
 - of thin provisioning [16](#)
- D**
- data
 - how Data ONTAP uses RAID to protect and ensure availability [10](#)
 - data compression
 - best practices [28](#)
 - data compression
 - recommendations [26](#)
 - estimating space savings as a result of [22](#)
 - guidelines [26](#)
 - how detecting incompressible data saves system resources [27](#)
 - increasing storage efficiency [26](#)
 - interoperability with Data ONTAP features [27](#)
 - deduplication
 - best practices [28](#)
 - estimating space savings as a result of [22](#)
 - factors affecting performance [25](#)
 - FlexVol volumes
 - deduplication guidelines [24](#)
 - guidelines for running [24](#)
 - how it works [22](#)
 - increasing storage efficiency [22](#)
 - interoperability with Data ONTAP features [25](#)
 - metadata relocated [23](#)
 - using qtree SnapMirror with [36](#)
 - disk space
 - recovery of [21](#)
- E**
- example
 - when Snapshot copies exceed the reserve [20](#)
- F**
- file folding
 - defined [21](#)
 - file reservations
 - how they work [13](#)
 - file system
 - recovery of disk space for use by the [21](#)
 - FlexClone
 - best practices [43](#)
 - FlexClone files with deduplication
 - FlexClone LUNs with deduplication [44](#)
 - FlexClone files with qtree SnapMirror [43](#)
 - FlexClone files with SnapVault [43](#)
 - FlexClone files with volume SnapMirror [43](#)
 - FlexClone LUNs with qtree SnapMirror [43](#)
 - FlexClone LUNs with SnapVault [43](#)
 - FlexClone LUNs with volume SnapMirror [43](#)
 - FlexClone volumes with volume SnapMirror [43](#)
 - storage efficiency [43](#)
 - FlexClone files
 - how they work [39](#)
 - FlexClone files and FlexClone LUNs
 - considerations [42](#)
 - space saving [42](#)
 - uses [41](#)
 - FlexClone LUNs
 - how they work [39](#)
 - FlexClone volumes
 - space guarantees and [41](#)
 - FlexVol volumes
 - fractional reserve
 - considerations for setting [13](#)
 - how Data ONTAP can automatically add space for [15](#)
 - how volume guarantees work with [11](#)
 - thick provisioning for [11](#)
 - thin provisioning for [11](#)

fractional reserve
 considerations for setting [13](#)

free space
 how Data ONTAP can increase automatically for
 FlexVol volumes [15](#)

G

guarantees
 volume, how they work with FlexVol volumes [11](#)

guidelines
 for running deduplication [24](#)

H

how to achieve storage efficiency
 using NetApp products [6](#)

I

incompressible data
 how detection saves system resources [27](#)

L

LUN reservations
 how they work [13](#)

P

performance
 deduplication, factors affecting [25](#)

protection
 how Data ONTAP uses to RAID for data [10](#)

R

RAID
 how Data ONTAP to protect data and data
 availability [10](#)

RAID-DP
 protection against double-disk failures [9](#)

reservations
 how they work [13](#)

reserves
 considerations for setting fractional [13](#)

S

SnapMirror
 best practices [34](#)
 qtree replication [34](#)
 qtree SnapMirror with data compression [36](#)
 qtree SnapMirror with deduplication [36](#)
 qtree SnapMirror with FlexClone [37](#)
 storage efficiency [34](#)
 volume replication [34](#)
 volume SnapMirror with data compression [35](#)
 volume SnapMirror with deduplication [35](#)
 volume SnapMirror with FlexClone [35](#)
 volume SnapMirror with SnapVault [36](#)

Snapshot copies
 using to increase file-recovery efficiency [18](#)

Snapshot copies exceed the reserve
 example [20](#)

snapshot copy
 snapshot copy reserve [19](#)

SnapVault
 best practices [32](#)
 SnapVault with data compression [32](#)
 SnapVault with deduplication [32](#)
 SnapVault with volume SnapMirror [33](#)
 storage efficiency [32](#)

SnapVault backup
 how it works [31](#)

software efficiency
 achieving by using RAID-DP [9](#)
 achieving using FlexVol volumes [16](#)
 FlexClone technology [38](#)
 increasing by using Snapshot copies [18](#)
 SnapVault [29, 31](#)
 volume SnapMirror [29, 33](#)

space
 how Data ONTAP can automatically add FlexVol
 volume [15](#)

space guarantees
See volume guarantees

space management
 how you use [11](#)

space reservations
See reservations

space savings
 estimating using SSET [22](#)

space savings estimation tool
See SSET

SSET

using to estimate deduplication and data
compression space savings [22](#)

storage efficiency

about [6](#)

data compression [22](#)

deduplication [22](#)

increasing with Flash Cache [8](#)

Increasing with Flash Pool aggregates [8](#)

T

thick provisioning

for FlexVol volumes [11](#)

thin provisioning

benefits of [16](#)

best practices [17](#)

for FlexVol volumes [11](#)

using FlexVol volumes [16](#)

V

volume guarantees

effect on maximum FlexVol volume size [11](#)

how they work with FlexVol volumes [11](#)

volume SnapVault and SnapMirror [29](#)

volumes

fractional reserve

considerations for setting [13](#)

how Data ONTAP can automatically add space for
[15](#)

W

what is storage efficiency [6](#)