

Storage Efficiency in Clustered Data ONTAP



Julian Cates

Technical Marketing Engineer

The clustered Data ONTAP® architecture is a highly virtualized storage operating system that abstracts physical storage into a set of storage virtual machines (SVMs) in order to deliver policy-based storage services for SAN and NAS with native multi-tenancy. The Tech OnTap® newsletter has included several recent articles on clustered Data ONTAP focused on the capabilities of the latest release, 8.2. The first was a [summary of new features](#) and the second article was a deep dive into [nondisruptive operations](#).

Clustered Data ONTAP makes both your storage and your IT staff more productive and efficient, so you can scale your storage infrastructure without scaling your IT organization. Management processes scale so that having twice as much storage no longer means having to do twice as much work. A common set of features and procedures simplifies complex tasks so your IT staff can focus on solving higher-level problems and your operational efficiency grows as scale increases.

With clustered Data ONTAP, you can consolidate and share the same infrastructure for workloads or tenants with different performance, capacity, and security requirements. You can also achieve storage cost reductions with comprehensive storage efficiency.

In this article, I dig into the suite of storage efficiency technologies in clustered Data ONTAP. Used alone or, especially, together, these technologies can significantly reduce the total amount of storage you need. Reducing the amount of storage you have deployed obviously reduces your capital expenses, but it also has significant benefits in terms of power, cooling, data center space, and administrative overhead.

Explore

Measuring Storage Efficiency

NetApp provides a variety of tools to help you understand storage efficiency. Be sure and check out the following.

- [Efficient IT Calculator](#): Estimate how much storage you can save using NetApp storage efficiency features.
- [NetApp Systems Manager](#): This tool provides a simple graphical interface that lets you enable storage efficiency features and see how much storage you're saving.
- [My AutoSupport](#): Among its many capabilities, the My AutoSupport™ tool includes a storage efficiency dashboard that lets you see how much storage you're saving. (requires NetApp Support login.) Read more about [My AutoSupport](#) in Tech OnTap.

Back to Basics

Learn the fundamentals of key NetApp technologies by reading the Tech OnTap Back to Basics series:

- [Chapter 1: Thin Provisioning](#)
- [Chapter 2: Deduplication](#)
- [Chapter 3: FlexClone](#)
- [Chapter 4: Volume SnapMirror](#)
- [Chapter 5: RAID-DP](#)
- [Chapter 6: Data Compression](#)
- [Chapter 7: SnapProtect](#)
- [Chapter 8: V-Series](#)
- [Chapter 9: SnapManager](#)




Figure 1) Using multiple storage efficiency technologies can compound storage savings and reduce overall storage costs.

In this article I also look at what's new or different with all the familiar Data ONTAP technologies that contribute to storage efficiency, especially SnapVault® technology, which was just added with clustered Data ONTAP 8.2.

Clustered Data ONTAP Versus 7-Mode

For the most part, the storage efficiency features in clustered Data ONTAP are similar to the Data ONTAP 7G and Data ONTAP 8 7-Mode capabilities you may already be familiar with. If you've been reading the Tech OnTap Back to Basics series over the years, you probably already know a lot about NetApp features such as thin provisioning; deduplication; FlexClone®, SnapMirror®, and RAID-DP® technologies; and compression. If you're not familiar with some or all of these features, the following Back to Basics chapters are a good place to start.

- Chapter 1: [Thin Provisioning](#)
- Chapter 2: [Deduplication](#)
- Chapter 3: [FlexClone](#)
- Chapter 4: [Volume SnapMirror](#)
- Chapter 5: [RAID-DP](#)
- Chapter 6: [Data Compression](#)

No chapter has (yet) been written on SnapVault. It's worth mentioning that thin provisioning, deduplication, compression, and RAID-DP are all free. When you purchase a NetApp FAS or V-Series system, they are included.


A few of the storage efficiency technologies such as RAID-DP and Snapshot are pretty much identical between 7-Mode and clustered Data ONTAP, so I don't spend any time on those. For the others, I explain what's new and different, starting with SnapVault.

SnapVault: New in Clustered Data ONTAP 8.2

The most important new storage efficiency capability in clustered Data ONTAP is the addition of SnapVault technology with thin replication. SnapVault is the NetApp solution for space-efficient disk-to-disk backup. It performs asynchronous replication using NetApp Snapshot copies of a primary volume. Once a baseline copy is created, incremental backups are all that are required; these are very efficient since only blocks that have

changed since the last Snapshot copy was taken are copied.

You can convert a SnapMirror to a SnapVault relationship if needed. The major operational difference between SnapMirror and SnapVault is that SnapVault allows you to have different retention schedules for Snapshot copies on the primary and secondary volume. (Snapshot copies serve as incremental backups.) You can retain Snapshot copies that have been released from primary storage on secondary storage. You can keep up to 251 Snapshot copies per volume.



Existing 7-Mode Capabilities	New in Clustered Data ONTAP 8.2
Snapshot based replication over any distance	SLA based Snapshot management policies with intelligent tagging
Forever Incremental Backups	Data protection policies follow volume moves
Multiple Recovery Points	Improved Performance
Fast and Simple Single File Recovery	Preserves primary storage efficiencies on-disk and over-wire
Take Advantage of High-Density SATA	Convert SnapMirror to SnapVault

Figure 2) New features in SnapVault with clustered Data ONTAP 8.2.

As with [volume SnapMirror in 7-Mode](#), SnapVault works at the volume level and preserves storage efficiency from the source over the wire and on the secondary volume. If the primary volume is deduplicated and/or compressed, the secondary volume will be in the same state. Enabling these features on the primary also makes each backup smaller and reduces network bandwidth requirements. You no longer have to rerun deduplication and compression on the destination, which can decrease backup times, sometimes significantly, because you are sending fewer physical blocks over the wire during backup.

Optionally, you can have different storage efficiency settings on the primary and secondary. For instance, you might have deduplication on the primary with deduplication and compression enabled on the secondary.

SnapVault advantages include:

- Fast, efficient backups that make it possible to do backups more frequently for greater levels of protection
- Reliable restores including end-user browse and restore, full-volume restore, and single-file restores
- Creation of usable replicas of your backup data using FlexClone
- Simple management using System Manager 3.0 and Unified Manager 6.0
- Ability to configure SnapVault relationships using OnCommand® WorkFlow Automation 2.1

Fan-In, Fan-Out, and Cascade Configurations

Fan In. SnapVault can be used in a variety of configurations. For instance, a popular configuration is to "fan in" from multiple remote locations to a single centralized location. In clustered Data ONTAP 8.2, up to seven clusters can fan in to a single cluster. In other words, volumes from some or all nodes in seven clusters can be backed up

to a single storage cluster. Each primary volume has its own corresponding secondary volume on the target. Note that this is slightly different than fan-in in 7-Mode, with which you can have multiple qtrees from different sources all backing up to the same volume.

Fan Out. For any single volume, you can fan out to up to four separate targets. These can be either SnapMirror or SnapVault relationships.

Cascade. You can create cascades in which a primary volume is replicated to a secondary, the secondary is then replicated again, and so on. You can mix SnapMirror and SnapVault (one instance only) in such a cascade.

Setup

To configure SnapVault for use with clustered Data ONTAP, separate storage clusters and SVMs have to be configured in a peer relationship. At least one intercluster logical interface must be configured per cluster node, and replication occurs between SVMs. You can also configure SnapVault to work within a single cluster (intracluster SnapVault) if you wish. In that case, SVM peering may or may not be required depending on the location of the source and destination volumes; cluster peering is not required.

You can create the destination volume with the autogrow option on if you wish. In that way the volume will be able to grow to accommodate the backups as needed as long as there is space in the aggregate.

For more details on using SnapVault with clustered Data ONTAP, refer to [TR-4183](#).

Other Storage Efficiency Technologies

Now I dig into what's new or different with thin provisioning, deduplication, compression, FlexClone, and SnapMirror in clustered Data ONTAP, again focusing on the latest 8.2 release.

Thin Provisioning

NetApp thin provisioning lets you present more logical storage to hosts or users than you actually have in your physical storage pool. Instead of allocating space upfront, storage space is dynamically allocated to each volume or LUN as data is written. In most configurations, free space is also released back to your common storage pool when data in the volume or LUN is deleted.

Benefits include:

- You are saved from having large amounts of allocated but unused storage.
- Utilization is increased, reducing the amount of storage you need.
- Planning is simplified. Multiple volumes share the same pool of free storage, so there are fewer storage pools to manage.

Thin provisioning in clustered Data ONTAP is identical to it in 7-Mode operation. When you relocate a volume using **vol move** (see last month's article on [nondisruptive operations](#) for details on that), the thin provisioning settings on the volume move with it.

Deduplication and Compression

I describe deduplication and compression in the same section because they are complementary technologies. Compression requires that deduplication be enabled. There are a number of enhancements to these technologies in clustered Data ONTAP 8.1 and 8.2.

NetApp® deduplication improves efficiency by locating identical 4K blocks of data in a flexible volume (FlexVol® volume) and replacing them with references to a single shared block after performing a byte-level verification check

(which makes sure that blocks are identical before removing them, thereby removing issues related to hash collisions). Deduplication reduces storage capacity requirements by eliminating redundant blocks of data.

I'm pretty sure you already have a general idea of what compression does. The key with NetApp compression is that NetApp has developed a way to provide both transparent inline and postprocess data compression in software while mitigating the impact on storage system resources.

The benefits of deduplication and compression technologies are similar:

- They work on both primary and secondary storage.
- They are transparent to running applications.
- Both can be applied to new data or previously written data.
- Both can run during off-peak hours.

When you do a **vol move**, efficiency is retained for both deduplication and compression. In other words, nothing is lost and you don't have to rerun the operations on the destination. A moved volume retains all of its deduplication and compression settings, so you don't have to do anything to have storage efficiency operations continue on the volume in the new location.

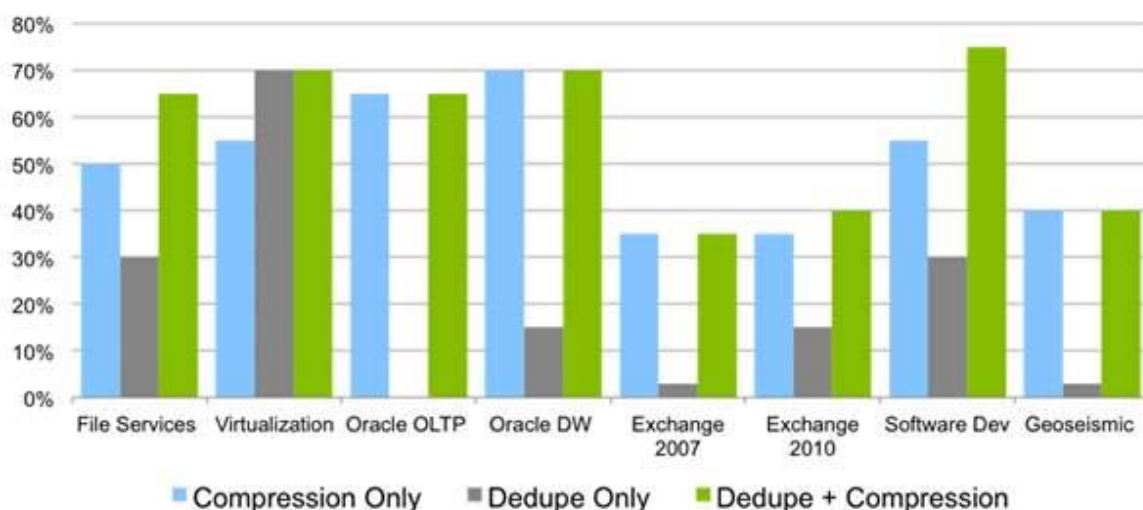


Figure 3) Typical savings due to compression and deduplication for various applications.

Enhancements as of Data ONTAP 8.1. Starting in Data ONTAP 8.1, some significant enhancements have been made (these apply to both 7-Mode and clustered Data ONTAP).

- License requirements for deduplication and compression were removed.
- Significant performance improvements were made to both technologies.
- Deduplication metadata follows a volume. As a result, deduplication operations continue in the face of a failover or when a vol move is performed.
- Enabling compression on a volume enables postprocess compression by default. Postprocess compression shares the same scheduler with deduplication: Compression runs first followed by deduplication. You can also enable inline compression.
- Volume size limits were removed, the logical data size limit was increased, and the block sharing limit was raised from 255 to 32K. (In other words, a deduplicated block can have up to 32K references.)

Specifically for clustered Data ONTAP, we added volume efficiency policies that give you the ability to schedule jobs at a specific time (not just on the hour) and let you include the duration so you can schedule a stop time. If compression and/or deduplication don't complete in the allotted time, a checkpoint is created and the operation continues at the next scheduled time. We also simplified management so that policy changes can affect multiple

volumes so you don't have to modify each volume individually. Set policies are retained when a volume is moved.

Enhancements as of Data ONTAP 8.2. Data ONTAP 8.2 adds some further enhancements.

- The amount of **deduplication metadata has been reduced**. It's now a maximum of 7% of the physical data size (the actual physical blocks in the volume) rather than based on logical data size (the size the volume would be without deduplication). Reducing the amount of metadata means that deduplication runs faster.
- The maximum **logical data size has been increased to 640TB**.
- **Priorities** have been added to volume efficiency policies for clustered Data ONTAP.
 - **Best-effort** (default). Volume efficiency operations compete for resources on an equal footing with user workloads and other system processes not running in the background. This priority allows deduplication and compression to complete most quickly, but there may be some impact on the performance of client I/O.
 - **Background**. Volume efficiency operations use leftover resources not used by user workloads and other system processes not running in the background. As a result, they may take longer to complete but should have less effect on client I/O.

Incompressible Data Detection. Incompressible data detection has been added to decrease the impact of inline compression.

- Inline compression attempts will stop for files that are less than 500MB (this default value is tunable) if a compression group within the file is found to have less than 25% savings.
- For files that are 500MB or larger (again this value is tunable), inline compression will attempt to compress the first 4K of each compression group. If there is at least 25% savings, compression will continue for that compression group; otherwise, the compression group will be written uncompressed.
- Postprocess compression is not affected by these changes and will attempt to compress any compression groups that were skipped.

FlexClone

In the IT world, there are countless situations in which it is desirable to create a copy of a dataset. Unfortunately, copies don't come free; they consume a lot of storage space and copy operations can take a long time to complete. NetApp FlexClone technology lets you clone an existing volume, file, or LUN in a matter of minutes while consuming only incremental amounts of additional storage.

In addition to being fast and space efficient, FlexClone offers a number of other advantages.

- It can reduce the storage needed for dev/test or virtual environments by 50% or more.
- It can improve dev/test quality by providing as many test datasets as necessary and making your team more productive while consuming only minimal incremental storage capacity.
- Virtual machine provisioning is accelerated.
- You can clone and fully test a DR environment using only incremental storage capacity.

There are some operations on FlexClone volumes that you need to be careful with in clustered Data ONTAP environments.

- If you use **vol move** to move a FlexClone volume to a new aggregate, you lose the space savings since the clone can no longer share blocks with its parent, which is still on the source aggregate.
- If you use **vol move** to move a volume that has clones (a parent volume), the volume is moved and client I/O redirects to the destination aggregate. However, the original volume remains on the source aggregate in restricted mode (no client I/O). This is so that all FlexClone volumes on the aggregate can still access their shared blocks. The original volume remains until all clones are split, moved, or deleted.

- File and LUN clones move with the parent volume.

SnapMirror

The thin replication capability provided by NetApp SnapMirror software has been the preferred technology for replication and disaster recovery in a wide variety of NetApp storage environments for years. This is because of its proven efficiency, simplicity, and modest cost compared to other DR solutions. SnapMirror advantages include:

- Efficient, block-level updates that reduce time and network bandwidth requirements.
- Storage efficiency (deduplication, compression) is maintained over the network and on the secondary volume.
- Data can be flexibly replicated between dissimilar NetApp storage nodes using one-to-one, one-to-many, many-to-one, or many-to-many replication topologies (with asynchronous replication).
- In conjunction with FlexClone you can use DR data for dev/test, data mining, or other functions, and you can fully test your DR plan without affecting production and ongoing replication.
- You can use SnapMirror within a single cluster or between clusters.

If you're familiar with [volume SnapMirror](#) running on 7-Mode storage systems, you will find SnapMirror in clustered Data ONTAP very familiar; storage efficiency savings on the primary (from deduplication and/or compression) are inherited by the secondary copy.

New capabilities such as load-sharing mirrors, which allow you to replicate a volume to multiple nodes in the same cluster to accelerate read-only workloads, are also available for use with clustered Data ONTAP.

As with SnapVault, when configuring SnapVault from one cluster to another you have to configure at least one intercluster logical interface per cluster node, and replication occurs between SVMs in a peer relationship. You can find complete details on using SnapMirror with clustered Data ONTAP in [TR-4015](#).

Conclusion

Storage efficiency technologies in clustered Data ONTAP are more than just an extension of the familiar 7-Mode capabilities. NetApp is working hard to add new and greater functionality to allow you to accomplish even more—with less effort and less storage—in your clustered storage environment.



By Julian Cates, Technical Marketing Engineer

In his current role as a TME, Julian focuses on clustered Data ONTAP and secure multi-tenancy. In his nine-year career at NetApp, he has served in pre- and post-sales roles in Systems Engineering and Professional Services as well as technical marketing roles for joint VMware® and NetApp solutions.

Quick Links

- › [Tech OnTap Community](#)
 - › [Archive](#)
 - › [User Groups](#)
 - › [PDF](#)
-