




Back to Basics: FlexClone

By Carlos Alvarez


This article is the third installment of Back to Basics, a series of articles that discusses the fundamentals of popular NetApp technologies.

In the IT world, there are countless situations in which it is desirable to create a copy of a dataset: Application development and test (dev/test) and the provisioning of new virtual machines are common examples. Unfortunately, traditional copies don't come free. They consume significant storage capacity, server and network resources, and valuable administrator time and energy. As a result, your operation probably makes do with fewer, less up-to-date copies than you really need.

This is exactly the problem that NetApp FlexClone® technology was designed to solve. FlexClone was introduced in Data ONTAP® 7G to allow you to make fast, space-efficient copies of flexible volumes (FlexVol® volumes) and LUNs. [A previous Tech OnTap® article](#) describes how one IT team used the NetApp rapid cloning capability built on FlexClone technology (now incorporated as part of the [NetApp Virtual Storage Console, or VSC](#)) to deploy a 9,000-seat virtual desktop environment with flexible, fast reprovisioning and using a fraction of the storage that would normally be required. [NetApp uses the same approach for server provisioning](#) in its own data centers.



Traditional Approach



FlexClone®

Explore

More Back to Basics

Learn the fundamentals of NetApp® core technologies. Installments in this series so far:

- [Chapter 1: Thin Provisioning](#)
 - [Chapter 2: Deduplication](#)
 - [Chapter 3: FlexClone](#)
 - [More](#)
-

Figure 1) FlexClone technology versus the traditional approach to data copies.

Using FlexClone technology instead of traditional copies offers significant advantages: It is:

- **Fast.** Traditional copies can take many minutes or hours to make. With FlexClone technology even the largest volumes can be cloned in a matter of seconds.
- **Space efficient.** A clone uses a small amount of space for metadata, and then only consumes additional space as data is changed or added.
- **Reduces costs.** FlexClone technology can cut the storage you need for dev/test or virtual environments by 50% or more.
- **Improves quality of dev/test.** Make as many copies of your full production dataset as you need. If a test corrupts the data, start again in seconds. Developers and test engineers spend less time waiting for access to datasets and more time doing productive work.
- **Lets you get more from your DR environment.** FlexClone makes it possible to clone and fully test your DR processes, or use your DR environment for dev/test without interfering with ongoing replication. You simply clone your DR copies and do dev/test on the clones.
- **Accelerates virtual machine and virtual desktop provisioning.** Deploy tens or hundreds of new VMs in minutes with only a small incremental increase in storage.

Most Tech OnTap readers probably know about the use of FlexClone for cloning volumes. What's less well known is that, starting with Data ONTAP 7.3.1, NetApp also gave FlexClone the ability to clone individual files and improved the capability for cloning LUNs.

This chapter of Back to Basics explores how NetApp FlexClone technology is implemented, the most common use cases, best practices for implementing FlexClone, and more.

How FlexClone Is Implemented in Data ONTAP

Cloning Volumes

FlexClone volumes have all the capabilities of any other FlexVol volume, including the ability to grow, shrink, and be the source of a Snapshot® copy or even another FlexClone volume. The technology that makes this all possible is integral to how Data ONTAP manages storage. NetApp storage systems use the Write Anywhere File Layout (WAFL®) to manage disk storage. Any new data that gets written to the volume doesn't need to go to a specific spot on the disk; it can be written anywhere. WAFL just updates the metadata to integrate the newly written data.

Snapshot copies simply make a copy of the metadata associated with a volume. As data is changed in the parent FlexVol volume, the original data blocks stay associated with the Snapshot copy rather than getting marked for reuse. All metadata updates that occur are just pointer changes.

You can think of a FlexClone volume as a transparent writable layer in front of a Snapshot copy. Because a FlexClone volume is writable, it needs physical space to store data that is written to the clone. A Snapshot copy simply links to existing data that was overwritten in the parent volume; a FlexClone volume stores the data written to it on disk (using WAFL) and then links to the new data as well. The disk space associated with the Snapshot copy and the FlexClone volume is accounted for separately from the data in the parent FlexVol.

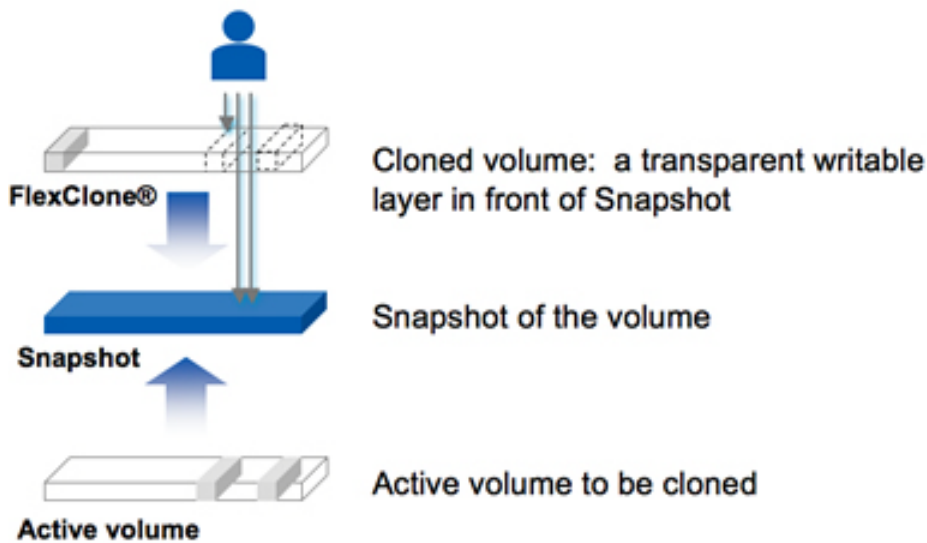


Figure 2) Volume-level cloning.

When a FlexClone volume is created, it needs to know the parent FlexVol volume and also needs a Snapshot copy of the parent to use as a base. The Snapshot copy can already exist, or it can be created automatically. The FlexClone volume gets a copy of the Snapshot metadata and then updates its metadata as the clone volume is created. Creating the FlexClone volume takes just a few moments because the copied metadata is very small compared to the actual data.

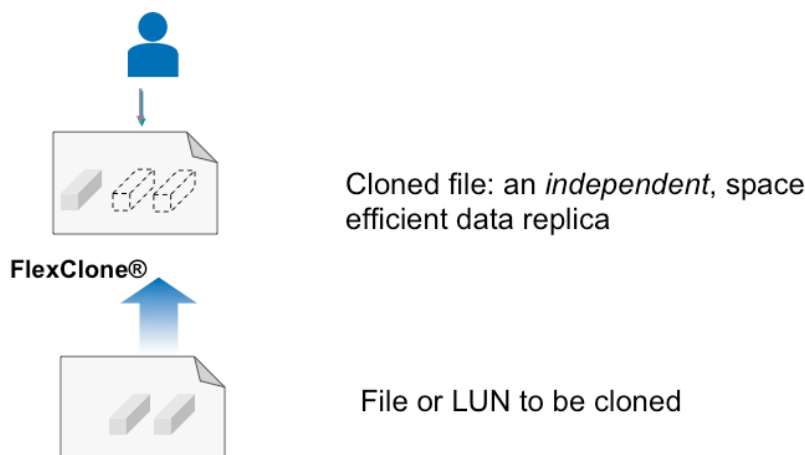
The parent FlexVol volume can change independent of the FlexClone volume, because the Snapshot copy is there to keep track of the changes and prevent the original parent's blocks from being released while the Snapshot copy exists. The Snapshot copy is read-only and can be reused as the base for multiple FlexClone volumes. Space is used very efficiently, since the only new disk space used is associated either with the small amounts of metadata or updates and/or additions to either the parent FlexVol volume or the FlexClone volume.

This approach for cloning volumes can be used to clone volumes that contain LUNs as well. Typically, you need to make sure that the target LUN or LUNs within the volume are in a consistent state before cloning. This and much more is described in detail in [NetApp TR-3347: A Thorough Introduction to FlexClone Volumes](#). However, the volume-level cloning approach has been largely superseded for LUN cloning by the approach described next.

File and LUN Cloning

Starting with Data ONTAP 7.3.1, you can create a clone of a file in a FlexVol volume in a NAS environment or clone a LUN in a SAN environment without the need for a backing Snapshot copy.

Similar to volume-level cloning, cloning of files and LUNs is highly space efficient, since the cloned copies share the same physical space with the source and the space occupied by the initial metadata is negligible. Cloned files or LUNs occupy additional space only as data is overwritten or added to the source or the clone. Clone creation is a fast and time-efficient process since no physical data copying is required.



File or LUN in active volume

Figure 3) File- or LUN-level cloning. (Requires Data ONTAP 7.3.1 or later.)

The process of creating a clone of an existing file or LUN has no impact on client access, either during the creation process or after the cloning is done. Clients can write to the source file or LUN while the cloning process is in progress. Once the cloning process is complete, the clone files or LUNs can be accessed from a client and treated like any other file or LUN. Source files and source LUNs and clone files and clone LUNs can be deleted with no ill effect.

These new abilities, combined with the volume-level cloning capability described above, provide a space- and time-efficient solution for many data center problems that require multiple copies of the same dataset. FlexClone at the volume, file, and LUN levels can be combined to create a powerful time- and space-efficient solution to store redundant datasets, since all the redundant files or LUNs share the same underlying physical storage.

FlexClone can also be used to clone individual files inside a LUN in a SAN environment. Data ONTAP provides an API to enable this capability. However, host-side support is needed to integrate the clone file into the host file system and to make the clone file usable by the client. You can read more about this process and all aspects of file and LUN cloning in [TR-3742: Using FlexClone to Clone Files and LUNs](#).

Use Cases

You can make use of FlexClone technology in almost any situation in which you need a copy of a file, LUN, or volume. Some even use FlexClone technology successfully in their production environments by using [FlexShare®](#) [13] to manage the latency of clones. If you have large databases, you may find FlexClone particularly useful in conjunction with data warehouse operations as well as dev/test.

In this section, I talk about the two most popular use cases:

- Dev/test—volume-level cloning
- Provisioning for virtual environments—file- or LUN-level cloning

Dev/Test

Because making clones has zero impact, you can refresh the cloned production data you use for your development work more often, so that you always test against current rather than stale data. Most shops refresh only once every 90 days.

Also, instead of having all your developers and testers share one or two copies of a test database, you can create a gold copy and clone it multiple times so that each person has his or her own clones to work with. You can do destructive testing without affecting anything outside the clone. When the testing is finished, you simply delete the clone, and you can create a new, clean clone image in a matter of minutes. This approach is described in a recent [Tech OnTap article on development for Oracle11g™](#) that also describes how you can integrate data masking (to eliminate confidential user data) into your development process.

You may have a DR environment that sits more or less idle most of the time. With FlexClone technology, you can clone the DR volumes that correspond to your production environment and put that infrastructure to work. Replication to clone source volumes continues unimpeded while your dev/test team works on clones of production data.

The net result is significant improvements in development and test capabilities that can lead to improvements in application quality and faster application delivery as well as reduced costs. For example, dev/test for a 100GB production database usually starts with a full mirror and then several copies for developers and several for testers. If we conservatively assume that three of each type are needed, the total storage requirement (including the production database) is 800GB. Retaining the full mirror (to avoid any impact on production storage) and using FlexClone for the dev/test copies reduces the total storage requirement to just 260GB—a 67% reduction in storage required. (This assumes that the average change rate in dev/test volumes is about 10%.) Rapid creation and cleanup of clones also means that users waste less time waiting for copies and more time working.

Provisioning for Virtual Environments

Provisioning in virtual server and desktop environments can also benefit from FlexClone technology. Traditional provisioning requires a full copy, and the full process takes 20 to 30 minutes to complete for a typically sized VM. Provisioning with FlexClone reduces that time to about 3 minutes, start to finish.

For VMware® environments, file-level cloning can be used to create clones of VMDK files stored in VMware datastores accessed via NFS. LUN-level cloning can be used when VMDK files are stored inside a VMFS datastore in a LUN that is accessed over either FCP or iSCSI. This capability can be accessed from VMware vCenter™ using the rapid cloning utility integrated into the [NetApp Virtual Storage Console](#). VSC not only handles cloning, it also provides proper configuration and registration with vCenter. VSC also supports redeploying existing virtual machines to bring them up to date with the latest patches or other changes. You can learn more about provisioning and other management tasks in VMware environments with NetApp storage in a [recent Tech OnTap article](#).

To integrate cloning in Microsoft® Hyper-V™ environments, NetApp provides [ApplianceWatch PRO for Microsoft System Center](#). Cloning integration with XenServer is provided by the codeveloped [Citrix StorageLink Adapter for NetApp Data ONTAP](#).

Using FlexClone Technology

A few best practices can help you succeed with FlexClone technology. For full details refer to [NetApp TR-3347: A Thorough Introduction to](#)

[FlexClone Volumes](#) and [TR-3742: Using FlexClone to Clone Files and LUNs](#). A few best practices are summarized here.

For volume-level cloning:

- The maximum number of clones of a volume is 255.
- Be aware of how space reservations work and monitor available space when working with FlexClone volumes. Alerts can be set in many tools to send out notifications when space is getting low.
- When a volume is cloned, the FlexClone volume contains data with the same ownership and permissions as the source volume. Users and applications that can access the FlexClone volume will also be able to access the parent. It is better to use separate user accounts for development/testing and production. That means you should allow access to the FlexClone volume but not the parent. One method is to first mount/map each FlexClone volume on an administrative host, change file permissions and/or ownership to match the authorized development/test users, and then remount the FlexClone volumes to the appropriate servers.
- Don't make any clones of volumes protected with NetApp SnapLock®.
- For manually created Snapshot copies, it is a good idea to use a name that makes it clear that the Snapshot copy is of a clone. Since Snapshot copies cannot be renamed, you need to be aware of any previously existing Snapshot copies that are also being used to back clone volumes.
- Data ONTAP locks any Snapshot copy used to back clone volumes until the clone is either split off or destroyed. Any disk blocks associated with the Snapshot copy volume will remain locked and will not be released for reuse until the Snapshot copy is deleted.
- Don't delete the initial Snapshot copy. The presence of this snapshot helps limit the amount of data that can be changed in a FlexClone volume and speeds up certain client operations on

data within the FlexClone volume.

- Data ONTAP does not automatically delete backing Snapshot copies in the parent when a FlexClone volume is split or destroyed, because it's hard to guess when the Snapshot copy may be needed for future work. It is up to you to review existing Snapshot copies after FlexClone volumes are removed and choose which ones to delete.

For file- and LUN-level cloning:

- A block in WAFL currently can have a maximum of 255 pointers to it. This means that you can clone a file or LUN a maximum of 255 times within a single FlexVol volume. If you create more clones than that, physical copies are made.
- FlexClone files are not automatically space reserved after creation. No matter what the space reservation on the source file is, the clone file has no space reservation. To enable space reservation on a file clone, use the `file reservation` command.
- FlexClone LUNs inherit space guarantee settings from the source. If there is not enough space in a volume to create a clone with the same space guarantee as the source, the cloning process will fail. Note that the source and clone LUN will share blocks on the disk even with space guarantee enabled.
- Quota usage for clones is charged at the logical level. So, the amount of extra space usage charged to the quota for creating a clone is equal to the total logical size of the clone. For example, if you create a clone of a 10GB file, the total used space charged to the quota for the source file and the clone file is 20GB (10GB for the source and 10GB for the clone).
- The effect of exceeding the quota limit due to the creation of a FlexClone instance of a file is different for qtree quotas and user or group quotas. If the total logical used space occupied after a clone is created is more than the allowed tree quota for the qtree, the clone operation will fail.

- If the total logical used space occupied after a clone creation is more than the allowed quota for that user or group, the clone operation will still succeed if the FlexVol volume has enough space to hold the metadata or the data for the clone. However, after the success of the clone operation, the quota for that user or group will be oversubscribed.
- If the file that is the source of a FlexClone operation has access control lists or streams, the ACLs and streams will not be cloned and so will not be present for the clone file. If you want the same ACL-based permission as the source file on the clone file, or attach streams to the clone file, you must do so separately on the clone file after the cloning process is completed.
- The `DU -k` command is extremely useful for identifying unique blocks in a file or LUN clone.

Combining volume-level and file/LUN-level cloning:

- You can combine the volume-level and file/LUN-level approaches to cloning in situations in which you need to create a large number of clones of a single file or LUN. This can be accomplished by making the maximum number of clones of the file or LUN (255) in the original volume and then cloning the volume as many times as necessary to get to the desired number.
- For a large number of file/LUN copies in the same volume, make 255 clones. The 256th clone will be a full copy. You can then clone this copy 255 times and so on until the desired number is reached.
- See TR-3742: Using FlexClone to Clone Files and LUNs[5] for full details of these procedures.

FlexClone with Other NetApp Technologies

When it comes to integration with other NetApp products, NetApp FlexClone technology shares many similarities with [NetApp deduplication](#). This is because both technologies reduce storage use by allowing a single block of storage to have many pointers to it. Here's

how FlexClone works in conjunction with deduplication and a few other NetApp technologies:

- **Deduplication.** A FlexClone volume of a deduplicated volume inherits the savings from deduplication. You can also create a FlexClone of a nondeduplicated volume and deduplicate the clone to realize the savings without the parent copy being touched. In virtual environments, you can use FlexClone technology to create virtual machines very space efficiently, and combine that with deduplication to keep space savings at a maximum over time.
- **Flash Cache.** [Flash Cache](#) provides intelligent caching that accelerates I/O operations. NetApp FlexClone technology increases the likelihood of a cache hit. When a block that is shared by many files or volumes is in Flash Cache, the probability that it will be requested again is much higher. This effect is referred to as cache amplification, and is particularly beneficial with server and desktop virtualization.

SnapMirror. When you use volume SnapMirror® with file-level or LUN-level FlexClone technology, space savings are maintained because, no matter how many copies you have, the clone is only replicated once. With qtree SnapMirror and SnapVault®, space savings are lost, and you end up with full copies of cloned files. Deduplication can be used to recover space on the target.

In some instances, the space-efficient volume clones will contain critical data that warrants replication. Before Data ONTAP 8.0.1 (7-Mode), when a FlexClone volume is replicated using volume SnapMirror, space savings are lost. The FlexClone volume on the target requires capacity equal to the size of the parent volume. Starting with Data ONTAP 8.0.1, when operating in 7-Mode, FlexClone volumes can be replicated using volume SnapMirror without the need for additional capacity on the destination system as long as the parent of the FlexClone volume is also replicated. See [TR-3446: SnapMirror Best Practices](#) for more details.

Conclusion

NetApp FlexClone technology is an important storage efficiency tool that can be used alone or in conjunction with other solutions such as NetApp Flash Cache, deduplication, and others. To learn more about NetApp FlexClone, refer to [NetApp TR-3347: A Thorough Introduction to FlexClone Volumes](#) and [TR-3742: Using FlexClone to Clone Files and LUNs](#).

Got opinions about FlexClone?

Ask questions, exchange ideas, and share your thoughts online in NetApp Communities.



Carlos Alvarez
Sr. Technical Marketing Engineer
NetApp

Carlos has been working for NetApp since 2008, specializing in storage efficiency with deep-dive expertise in deduplication, data compression, and thin provisioning. He regularly provides guidance for integrating the most effective and appropriate NetApp storage efficiency technologies into customer configurations. With over 20 years of industry experience, Carlos has been called on to create numerous implementation guides, technical white papers, reference architectures, best practices, and solutions guides.



Quick Links

[Tech OnTap Community](#)

[Archive](#)

[User Groups](#)
