

Nondisruptive Operations with Clustered Data ONTAP



Charlotte Brooks
Technical Marketing Engineer

The move to shared infrastructure has made it nearly impossible to schedule downtime to accomplish routine maintenance. A single storage system might be virtualized among many applications, many different stakeholders, or many sets of users, so negotiating a downtime window—when it's possible at all—consumes a huge amount of time and leaves no one happy. In addition, the process of upgrading and replacing equipment as it completes its lifecycle—even if you can do it with minimal disruption—is time consuming, and the planning is complex.

Clustered NetApp® Data ONTAP® is designed to eliminate the planned downtime needed for maintenance operations and lifecycle operations as well as the unplanned downtime caused by hardware and software failures. NetApp's goal is to allow your storage infrastructure to be fully available and resilient at all times without downtime. We think that data should be delivered with the same reliability as any other utility. When you go to the sink and turn on the tap, you expect water.




Figure 1) NetApp nondisruptive operations protect your infrastructure against both planned and unplanned downtime while greatly simplifying maintenance and lifecycle operations.

Our data and network mobility features make it possible to accomplish maintenance and lifecycle tasks without downtime. The benefits you get from nondisruptive operations are significant, and many of them flow straight to the bottom line:

- **Deploy new hardware and/or upgrade software faster.** Ever have new hardware arrive just to have it sitting on the dock or cluttering up a hallway for months while you wait for downtime to upgrade? That's a low return on investment. With nondisruptive operations, waiting is a thing of the past.
- **Higher utilization.** Because you can add new capacity when you need it without waiting for downtime, you can push your NetApp cluster to higher levels of utilization. You no longer have to maintain a big safety net in the form of storage capacity sitting idle.

Explore

Learn More About Clustered Data ONTAP 8.2 and Nondisruptive Operations

In case you missed it, last month NetApp TME Julian Cates reviewed the new 8.2 features and wrote in depth about aggregate relocate, quality of service, and Microsoft® Windows® enhancements. Check out his article, plus a new white paper that explores the business benefits of nondisruptive operations, efficiency, and scalability.

- [What's New in Clustered Data ONTAP 8.2?](#)
 - [Remove IT Constraints and Speed Response to Business Changes](#)
-

- **Simpler operations.** Scheduled downtime often requires spending significant time to take running applications offline. Then you have to restart the applications and verify that everything is still running correctly when the maintenance is completed. Nondisruptive operations let you focus on performing storage tasks without this complexity. The NetApp tools for nondisruptive operations are simple to use and designed for low-stress, repeated execution as needed.

In this article, I'll dig into the tools that NetApp offers for nondisruptive operations, and I'll look at how these capabilities can be used to accomplish important maintenance and lifecycle tasks.

Tools of the Trade

Clustered Data ONTAP enables nondisruptive operations by being resilient to failures and by allowing you to change your storage infrastructure without disruption to facilitate day-to-day operations and maintenance. This is possible because—rather than accessing physical resources directly—in clustered Data ONTAP all data access goes through a logical construct called a Storage Virtual Machine (SVM). As a result, the physical resources used by an SVM can change without necessitating any client-side or host-side changes or disruptions.

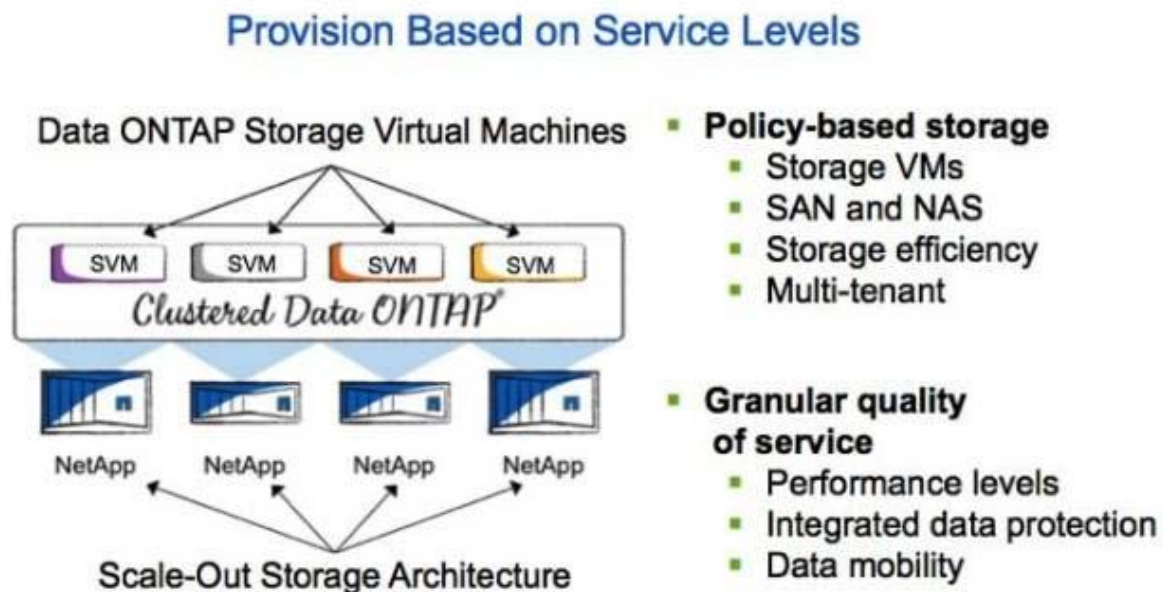


Figure 2) Clustered Data ONTAP Storage Virtual Machines (SVMs) abstract data access from physical hardware for greater flexibility.

Three standard tools make this possible:

- **DataMotion™ for Volumes (vol move).** Allows you to move data volumes from one aggregate to another, on the same or a different cluster node.
- **LIF migrate.** Logical interfaces (LIFs) virtualize the physical interfaces in clustered Data ONTAP. LIF migrate lets you move LIFs from one network port to another, on the same or a different cluster node.
- **Aggregate relocate (ARL).** Allows you to transfer complete aggregates from one controller in an HA pair to the other without data movement.

Used individually and in combination, these tools give you the ability to nondisruptively perform a full range of operations, from moving a volume from faster to slower disk, all the way up to a complete controller and storage technology refresh.

DataMotion for Volumes

DataMotion for Volumes (often referred to as vol move) lets you move a volume within an SVM from one aggregate (the source) to another aggregate (the destination). The destination can be on the same node or any other node in the cluster. Regardless of where the data is moved to, and regardless of the data protocol being served (SAN or NAS), data access is preserved transparently before, during, and after the move for client/host applications.

There are four phases in the volume move process. Once the volume move is initiated, the progression of stages is automatic, but it is important to understand each phase and the underlying activities.

- **Validation phase.** Verifies that the requested vol move is possible by checking available capacity on the destination aggregate as well as other requirements.
- **Setup phase.** A new volume is created on the destination aggregate.
- **Iterative phase.** Data is replicated from the source volume to the destination volume by replicating groups of Snapshot™ copies over the cluster network. After each iteration, the delta between the source and destination is checked to see if it is small enough that a final replication can be completed in the time defined for the cutover phase. I/O from clients and hosts to the source volume is not affected during this phase.
- **Cutover phase.** All I/O access is queued and requests to the source volume are blocked. The final replication transfer is completed and the volume database is updated with the new volume information. Queued I/O is then resumed on the volume at the new location. The cutover completes in a defined "cutover period" that is within an acceptable window of time for the client/host application.

If the cutover cannot complete in the specified cutover time, then the cutover phase aborts and data access resumes on the source volume. Any pending requests that were queued during the cutover attempt are completed and the iterate phase is resumed until conditions exist to retry a cutover.

The cutover window can be defined between 30 and 300 seconds; the default value is 45 seconds. Because this is the most critical part of the process, vol move gives you quite a bit of control over the cutover phase. For instance, you can perform the operation so that the cutover does not occur until you trigger it. This gives you the option to complete the cutover at a time of your choosing.

LIF Migrate

LIF migrate gives you the same ability to move network connections that vol move gives you for data volumes. A LIF is a logical network interface that virtualizes your SAN and NAS network connections. LIFs are tied to an SVM and mapped to physical network ports, interface groups, or VLANs (when tagging is used) on the controller. Because LIFs are virtualized, a LIF address remains the same even when a LIF is migrated to another physical port on the same or a different node within the cluster. NAS LIFs automatically fail over if one of the cluster nodes goes down, working together with storage failover in the HA pair to preserve data access. You can also manually migrate a LIF to another port.

Each cluster node can support a maximum of 262 LIFs, 6 of which are reserved for management and cluster functions. Data LIFs are used to serve data to clients or hosts and are designated as either SAN or NAS. IP-based LIFs (NAS or iSCSI) are assigned IP addresses, and FC-based LIFs are assigned WWPNs. Each SVM requires at least one data LIF. In normal operation, you should limit the number of data LIFs per node to 128 or fewer. That way, if an HA failover occurs, the limit on the partner node won't be exceeded even if it takes over all LIFs from the failing node.

In addition to data LIFs, there are management LIFs for accessing the cluster via the CLI or OnCommand® System Manager, and intercluster LIFs for the cluster interconnect network.

LIF migrate lets you move an IP-based LIF from one physical port or interface group to another. SAN data LIFs

(including iSCSI) do not need to be migrated and do not fail over. Instead, ALUA and MPIO processes on the hosts with initiators are used to optimize paths and handle path failures.

You can use LIF migrate to move all data LIFs (and thus all network traffic) off of a particular node to accomplish hardware maintenance or replacement. Another use case for LIF migrate is the ability to nondisruptively upgrade from an entry-level 2-node switchless cluster (new in clustered Data ONTAP 8.2) to a 2-node switched cluster. LIF migrate makes it possible to move the cluster interconnect LIFs around so that you can introduce the switch without disrupting data flow. Once the switch is in place, you can expand the cluster as needed to create larger configurations.

You can use LIF migrate to move a LIF to a different port on the same node. For instance, you might have a LIF configured on a GbE port. If that LIF requires more bandwidth, you could move it temporarily or permanently to a 10GbE port on the same node.

To find out more about LIFs and other clustered Data ONTAP networking topics, check out [TR-4182: Best Practices for Clustered Data ONTAP Network Configurations](#).

Aggregate Relocate

Aggregate relocate (ARL) is a new feature that was introduced in clustered Data ONTAP 8.2. Because all cluster nodes in clustered Data ONTAP are part of an HA pair (with the exception of single-node clusters), ARL makes it possible to temporarily transfer ownership from one controller in an HA pair to the other to facilitate the upgrade process without moving data.

Using ARL, you can accomplish controller upgrades in significantly less time than it would take to migrate data to other controllers, upgrade the existing controllers, and migrate data back. A recent Tech OnTap® article by Julian Cates, [What's New in Clustered Data ONTAP 8.2?](#) discussed ARL in significant detail, describing how it works and best practices for using it.

Performing Maintenance and Lifecycle Tasks

Now that you understand the basic tools, let's see how they can be used to accomplish both maintenance and lifecycle tasks. Table 1 summarizes many of these tasks and describes the advantages of accomplishing them without disruption.

Table 1) Examples of nondisruptive lifecycle and maintenance operations.

Lifecycle Operation	Benefit
<ul style="list-style-type: none">• Rebalance controller performance, capacity and/or disk performance, and usage	<ul style="list-style-type: none">• Lower capex (up-front outlay)• Better lifecycle management• Eliminate hotspots (better performance, less risk of unplanned outage)
<ul style="list-style-type: none">• Add storage controllers or disk shelves• Add hardware to controllers	<ul style="list-style-type: none">• Improve performance and density• Improve resiliency
<ul style="list-style-type: none">• Upgrade storage controllers, disk shelves, cluster switches	<ul style="list-style-type: none">• Avoid disruptive technology refreshes
Maintenance Operation	
<ul style="list-style-type: none">• Upgrade storage software	<ul style="list-style-type: none">• Leverage new features sooner

<ul style="list-style-type: none"> • Upgrade system, disk, switch, firmware 	<ul style="list-style-type: none"> • Eliminate potential risks
<ul style="list-style-type: none"> • Replace failed controller or component within a controller, such as NIC, HBA, and failed storage components such as cables, drives, I/O modules 	<ul style="list-style-type: none"> • Reduce administrative overhead

Accomplishing Maintenance Tasks

Much stored data will outlive the storage system it's stored on. Over time the software will need to be updated, and the hardware will need to be replaced or repaired.

Software and Firmware Update

Nondisruptive upgrade (NDU) includes the upgrade of both storage system software and storage system firmware. NDU is a comprehensive solution for upgrading:

- Operating system software (Data ONTAP)
- Operating system firmware (BIOS)
- Shelf firmware
- Disk firmware
- Alternate Control Path (ACP) firmware

These tasks are all accomplished in such a way that I/O interruptions are brief. Applications continue to operate without the need for user notification or complicated downtime scheduling. Storage takeover and giveback (using ARL under the covers) together with LIF migrate allow you to perform maintenance operations on one controller in an HA pair at a time, without interrupting data services. You can use the [Upgrade Advisor](#) tool (requires access to the NetApp Support site) in My AutoSupport™ to help you plan nondisruptive upgrades. The tool generates a complete list of the steps required to accomplish an upgrade of your entire cluster.

Prior to clustered Data ONTAP 8.2, upgrades to the operating system had to be performed using the "rolling upgrade" process, which upgraded one HA pair at a time. In larger clusters, this was a time-consuming process. Starting with clustered Data ONTAP 8.2, you also have the option of performing a batch upgrade on clusters with 8 or more nodes. This speeds up the time needed to accomplish an upgrade for large clusters. Batch upgrade allows you to perform upgrade operations on several nodes in parallel, reducing the total time needed to upgrade the entire cluster. A cluster is allowed to run two different versions of clustered Data ONTAP while an upgrade is in progress, but the best practice is to keep the time during which a cluster runs in mixed mode as short as possible. Batch upgrade facilitates that.

Hardware Repair and Replacement

Nondisruptive operations support the nondisruptive repair or replacement of hardware components in the storage subsystem, ranging from disk drives and cables to controllers and shelves. Disk drives are protected by RAID and can be repaired and replaced using standard procedures, typically without requiring the use of the tools described earlier. Many redundant components such as cables can also be replaced after a failure without using these tools.

Accomplishing Lifecycle Operations

Lifecycle operations include activities to balance and optimize capacity and/or performance as well as operations to expand or refresh the technology in your cluster. A clustered Data ONTAP infrastructure is flexible and resilient to the many necessary changes that occur over years of continuous operation.

No matter how good your planning is, situations inevitably arise where some aggregates are short on capacity while others have extra space. It's easy to correct this type of capacity imbalance by using vol move to transfer volumes from crowded aggregates to ones with room.

Performance imbalances can be handled similarly. Volumes that require higher performance can be moved to a more capable controller (in a mixed cluster), to a controller with a lighter load, or to faster media. For example, you might move a volume that requires higher performance from an aggregate of high-capacity disks to one composed of performance disks, or you might move it to a controller that contains Flash Cache™ or a Flash Pool™ aggregate that combines SSDs and HDDs. Conversely, if the performance requirements of a dataset decrease, you might move the associated volume or volumes to an aggregate composed of high-capacity disks.

For clustered Data ONTAP cluster administrators, vol move is a standard, low-stress event, which typically does not require a change request. Vol move allows your IT team to achieve its capacity and performance objectives and optimize operations without breaking the budget, because they can easily move data to the right class of storage for application requirements rather than having to provision and deploy everything on more expensive, higher-performing drives.

Tech Refresh

The "icing on the cake" as far as NetApp nondisruptive operations goes is the ability to perform a complete tech refresh without disruption. The process of replacing storage hardware has always been disruptive, time consuming, and expensive. In fact, [a recent study](#) suggests that migrating data from an old storage array to a new one takes about 5 months to accomplish on average and adds almost 50% to the cost of owning an array.

Clustered Data ONTAP lets you accomplish full hardware refreshes simply and easily without taking data offline, thereby avoiding these hidden costs. Clustered storage systems aren't required to be the same generation or model, so you can replace one FAS platform with another—or change your entire storage infrastructure—without interrupting running applications or busy users. No other storage can do this.

You can use ARL to quickly and conveniently upgrade existing storage controllers, or you can accomplish full upgrades—including drives and shelves—by using vol move. In the latter case, you typically add the new system to the cluster and move the data from the old system to the new system before decommissioning the old system.

This isn't just theoretically possible; full tech refresh has been accomplished many times. For instance, one long-time clustered Data ONTAP user transitioned nondisruptively from a cluster of 20+ FAS6080s to 16 FAS6280s, each with 512GB of Flash Cache storage. Total capacity before and after the refresh remained around 1 petabyte.

The transition was accomplished by working in sets of four. The IT team added four new nodes and moved volumes from four of the old nodes to the new nodes. Then they shut the old nodes down and moved any still-supported disk shelves (some of this hardware was quite old and had to be retired) to the next set of new nodes.

Users reported a noticeable increase in throughput and perceived performance, and the upgrade also reduced ongoing maintenance costs. Most importantly, the whole process occurred without downtime. Note that the availability of ARL would make this type of upgrade even simpler and much quicker than it was at the time this upgrade occurred. Some early adopters have already used ARL to refresh full clusters in a 1-day event, without requiring any data migration or downtime.

Conclusion

Clustered Data ONTAP takes nondisruptive operations to a new level. A few simple tools—vol move, LIF migrate, and aggregate relocate—make it possible to quickly and easily accomplish maintenance and lifecycle tasks that were previously impossible to do without planned downtime and significant disruption. The ability to perform tasks

as needed instead of waiting for infrequent downtime windows means that your storage environment is optimized and risks are greatly reduced.

› **Got opinions about Nondisruptive Operations?**

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



By Charlotte Brooks, Technical Marketing Engineer

Charlotte is a TME for clustered Data ONTAP. In more than 5 years at NetApp, she has written and presented extensively on clustered Data ONTAP and storage management topics. Her current focus is nondisruptive operations and nondisruptive upgrades.

Quick Links

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