

# What's New in Clustered Data ONTAP 8.2?



**Julian Cates**

Technical Marketing Engineer

Over the last year and a half or so, the Tech OnTap® newsletter has devoted many articles to the clustered Data ONTAP® operating system. The technologies in clustered Data ONTAP help you increase IT agility and prepare your data center for the future by creating a flexible, software-defined storage (SDS) infrastructure with virtualized storage services, application self-service, and the capability to incorporate multivendor hardware. You can read more about SDS in [this month's article from Vaughn Stewart](#).

Every release of clustered Data ONTAP enhances existing capabilities and adds new features to deliver:

- Nondisruptive operations that eliminate the need for planned downtime
- Maximum efficiency so you can accomplish more with less storage hardware and less time spent managing that hardware
- Seamless scalability so you can start small and grow big without taking data offline for disruptive tech refreshes

In this article, I provide an overview of new features we've added to clustered Data ONTAP 8.2, and then I dig into some of the most significant capabilities in more detail:

- Quality of service (QoS)
- Data-in-place controller upgrades
- New features for Microsoft® Windows®

## New Features

---

Clustered Data ONTAP 8.2 offers a tremendous number of enhancements and new features. Table 1 highlights many of the new features.

**Table 1)** New features of clustered Data ONTAP 8.2.

Feature	Advantage
<b>Nondisruptive Operations and Data Protection</b>	
<ul style="list-style-type: none"><li>• Data-in-place online controller</li></ul>	Upgrade storage controllers without migrating data or taking it offline. (Described in more detail later.)

---

## Explore

### Learn More About Clustered Data ONTAP and Version 8.2

- [Now the Fun Really Begins...](#) (blog), NetApp CTO, Jay Kidd
- [Liberate Your Business with Clustered Data ONTAP 8.2](#) (blog)
- [New, Innovative Version of Clustered Data ONTAP](#) (blog)
- [TR-3982: NetApp Clustered Data ONTAP 8.2: An Introduction](#)
- See what customers are saying about clustered Data ONTAP ([video](#))
- [GetSuccessful: Clustered Data ONTAP and Virtualization](#) ([video](#))

### Liberate Your Business from IT Constraints Webcast Series

This three-part series provides in-depth insights into clustered Data ONTAP.

- "Maintain 24/7 Availability for Your Enterprise Applications," June 27, 2013
- "Get Breakthrough Efficiency in Virtual and Private Cloud Environments," July 18, 2013
- "Start Small, Grow Big with a Unified Scale-Out Infrastructure," August 8, 2013

[▶ Register now!](#)

---

upgrades	
<ul style="list-style-type: none"> <li>Volume-level storage-efficient SnapVault® D2D backup</li> </ul>	Back up to secondary or remote storage while preserving deduplication to reduce network bandwidth and the amount of storage required.
<b>Efficiency and Management</b>	
<ul style="list-style-type: none"> <li>Quality of Service (QOS)</li> </ul>	Limit the resources that can be consumed by a storage virtual machine (SVM, formerly referred to as a Vserver), volume, file, or LUN. (Described in more detail later.)
<b>Scalability</b>	
<ul style="list-style-type: none"> <li>Larger SAN clusters</li> </ul>	SAN and mixed-SAN/NAS clusters now support up to 8 nodes.
<ul style="list-style-type: none"> <li>Single-node and switchless 2-node clusters</li> </ul>	Start small and scale cost effectively up to 8 nodes and 23PB (SAN) or 24 nodes and 69PB (NAS).
<ul style="list-style-type: none"> <li>Infinite Volume enhancements</li> </ul>	<a href="#">Infinite Volume</a> allows you to create a single large volume spanning numerous controllers. With 8.2, you can now create numerous Infinite Volumes in a single cluster and use both NFS and CIFS protocols.
<ul style="list-style-type: none"> <li>Intracuster FlexCache® software</li> </ul>	FlexCache lets you build out a caching architecture within your storage cluster to accelerate the speed of parallel software builds, animation rendering, EDA, seismic analysis, and financial market simulations running over NFS.
<ul style="list-style-type: none"> <li>Expanded limits</li> </ul>	<p>Many of the limits in clustered Data ONTAP were dramatically expanded in version 8.2 to make the platform even more scalable. This includes support for:</p> <ul style="list-style-type: none"> <li>Up to 100,000 NFS clients</li> <li>Aggregates up to 400TB in size</li> <li>12,000 volumes in a 24-node NAS cluster</li> <li>49,000 LUNs in an 8-node SAN cluster</li> </ul>
<b>Windows Enhancements</b> (described in more detail later)	
<ul style="list-style-type: none"> <li>SMB 3.0</li> </ul>	Enhanced nondisruptive operations for Hyper-V™ and other capabilities
<ul style="list-style-type: none"> <li>ODX</li> </ul>	Offloads data transfers from hosts to

	NetApp® storage
<ul style="list-style-type: none"> <li>• Auto Location</li> </ul>	A capability unique to NetApp that enables clients to have the most direct path to storage
<ul style="list-style-type: none"> <li>• BranchCache v2</li> </ul>	Allows Windows clients to cache data locally to increase performance, particularly over WAN connections
<ul style="list-style-type: none"> <li>• File access auditing</li> </ul>	Supports monitoring, evidence, compliance, and recovery needs
<b>Other</b>	
<ul style="list-style-type: none"> <li>• IPv6</li> </ul>	Provides support for protocols and SNMP

## Enhance Workload Management with Quality of Service

---

Clustered Data ONTAP uses Storage Virtual Machines (SVMs, formerly referred to as Vservers) to decouple data access from physical storage devices. A NetApp storage cluster can be subdivided into distinct SVMs, each governed by its own rights and permissions. SVMs are used to securely isolate individual tenants—for instance, in a service provider environment—or individual applications, workgroups, business units, and so on. Because an SVM isn't tied to particular physical resources, you can adjust its resources without disruption.

Each application or tenant typically has its own SVM, and that SVM can be managed by the application owner or tenant. (Single-tenant environments can operate in a single SVM if desired.) Application-driven storage services, available through our OnCommand® plug-ins and APIs, allow application owners to automatically provision, protect, and manage data through the application management tools they are already familiar with.

### Clustered Data ONTAP 8.2 Adds Workload Management

Any time you put numerous workloads on a storage system or storage cluster there is the possibility that excessive activity from one workload can affect other workloads. This is especially true in multi-tenant environments, such as those of service providers, where you may have little or no idea what a particular tenant is doing with the storage you provide to the company. That's why we've added quality of service (QoS) to Data ONTAP 8.2; it's part of the base operating system, so no separate licensing is required.

QoS workload management allows you to define service levels by creating policies that control the resources that can be consumed by storage objects such as volumes, LUNs, and files (including VMDKs) or SVMs to manage performance spikes and improve customer satisfaction. Limits are defined in terms of MB/sec or I/O operations per second (IOPS). MB/sec limits work best for workloads that handle large block I/O, while IOPS limits are best for transactional workloads.

QoS gives you the ability to consolidate many workloads or tenants on a cluster without fear that the most important workloads will suffer or that activity in one tenant partition will affect another.

### Using QoS

Attention to a few best practices will help you get the best results with QoS.


QoS currently works with clusters of up to eight nodes. You can set limits on different types of storage objects in the same cluster, but you can't nest limits on objects. For instance, if you set a limit on a volume, you can't also set limits on LUNs or files inside that volume. Similarly, if you set a limit on an SVM, you can't place further limits

on storage objects inside the SVM. Every object in the SVM is rolled up and covered under the policy on the SVM.

QoS is applied by creating policy groups and applying limits to each policy group. For instance, a policy group can contain a single SVM, numerous SVMs, or a collection of volumes used by an application. In virtual environments, a policy group can contain one or more VMDK files or LUNs containing datastores. The limit applied to a policy group is a combined limit for all the objects the policy group contains. The scheduler actively controls work so that resources are apportioned fairly to all objects in a group.

Note that the objects need not be on the same cluster node, and if an object moves, the policy limit remains in effect. You can set a limit on a policy group in terms of IOPS or MB/s, but not both.

When a policy group reaches its throughput limit, throttling happens at the protocol layer. Additional I/O is queued and does not impact other cluster resources. To the application or end user, reaching the limit is very similar in behavior to the behavior you would expect when a physical storage system approaches its performance limit.



**Figure 1)** A policy group contains a collection of storage objects such as SVMs, volumes, LUNs, or files. The limit on a policy group applies to all the objects in that group collectively.

QoS is managed by the cluster administrator, and it cannot be delegated to a tenant or an application owner managing an SVM.

## Upgrade Controllers Without Moving Data

---

Eventually, the time will come when your storage controllers need to be upgraded to newer hardware. Previous versions of clustered Data ONTAP facilitated this process by allowing you to nondisruptively move active data off an HA pair, upgrade its controllers, and then move the data back. This data-motion, or **vol move**, process has proved tremendously popular for accomplishing upgrades and a variety of other maintenance and management functions.

Clustered Data ONTAP 8.2 simplifies and accelerates the upgrade process further by allowing you to accomplish controller upgrades without moving data using a new process called aggregate relocate, or ARL. Because all cluster nodes in clustered Data ONTAP are part of an HA pair (with the exception of single-node clusters), aggregate relocate makes it possible to simply pass active aggregates 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.

## How ARL Works

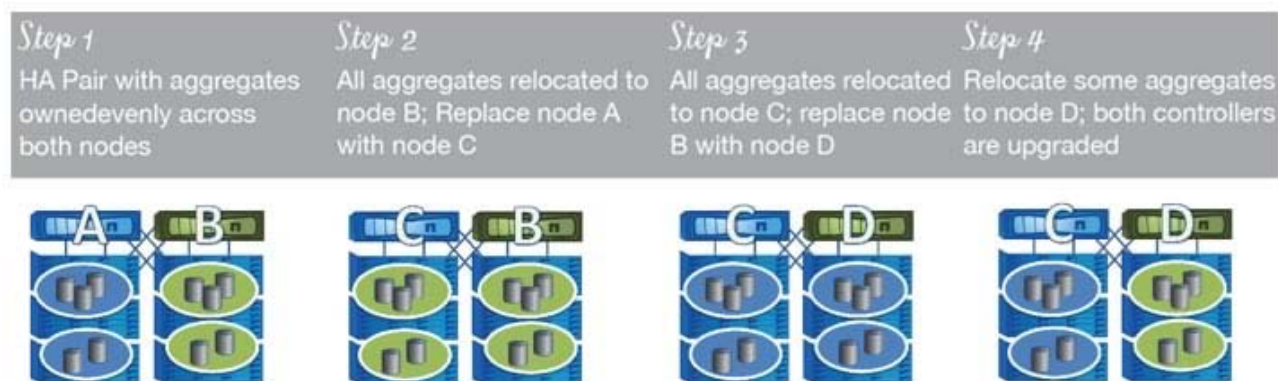
The ARL process proceeds in several phases.

- **Validation Phase:** In this phase the condition of the source and destination nodes is checked as well as the aggregates to be relocated.
- **Precommit Phase:** Processing that needs to be performed before the relocation is executed is done in this phase. This includes preparing the aggregate(s) to be relocated, to set flags, and to transfer certain noncritical subsystem data. The processing that occurs in this phase can be simply reverted or cleaned up if needed.
- **Commit Phase:** The processing associated with relocating the aggregate to the destination node is done in this phase. Once the commit phase has been entered, the ARL cannot be aborted. The processing that takes place in this stage is performed within a period of time acceptable to the client or host application. The aggregate is only offline momentarily as control is passed from the source node to the destination node. This time never exceeds 60 seconds and is usually less. Thirty seconds is typical.

The ARL process can be aborted if the checks that occur during the validation phase or precommit phase are not satisfied for some reason. A series of cleanup processes reverts any activity that occurred if the process is aborted.

An overview of the steps in a controller upgrade is given below. Note that logical interfaces (LIFs) are "virtualized" network interfaces that NAS clients or SAN hosts use to access network storage.

1. Use ARL to migrate aggregates from node A to node B.
2. Migrate data logical interfaces (LIFs) from node A to node B.
3. Disable storage failover (SFO) for the HA pair.
4. In two-node clusters, move epsilon (used to maintain a quorum of nodes within a cluster) to node B.
5. Replace node A with node C (new controller) and execute all setup, disk reassign, and licensing administration for node C.
6. Use ARL to relocate aggregates from node B to node C.
7. Migrate data LIFs from node B to node C.
8. Replace node B with node D (new controller) and execute all setup, disk reassign, and licensing administration for node D.
9. Enable SFO. Migrate selected aggregates and LIFs to node D.



**Figure 2)** Typical controller upgrade steps using ARL.

Note that while ARL is being performed, HA for the pair is disabled. In storage controllers with Flash Cache™ intelligent caching, cached data is not transferred between caches when an aggregate is moved; the cache on the destination node takes time to warm with data from the transferred aggregate. You may need additional rack



space to accommodate the new controllers if they are larger than the controllers they replace.

## Best Practices for ARL

Attention to a few best practices will help you succeed with aggregate relocate.

- Controllers that have internal disks (the FAS2200 series, for example) within the controller chassis require the use of `vol move` to physically relocate the data on the internal drives to other storage.
- You can use ARL to upgrade to a storage controller running the same version of Data ONTAP or a later version. (New hardware may sometimes be supported by a later software version.) If the upgrade results in a mixed-version state, the other nodes in the cluster should be updated to the same release as soon as possible. A node running a later version of Data ONTAP cannot relocate aggregates to a node running an older version.
- Don't use ARL while another HA pair in a cluster is in a failover state.
- You can issue aggregate relocate jobs in parallel. However, it may be necessary to consider volume limits on the destination node. With numerous ARLs running in parallel, validation checks may not identify situations in which the volume limit would be exceeded. It's best to run ARL jobs in sequence if you are near this limit.
- ARL jobs and `vol` moves should be executed separately when the two jobs use the same resources.
- Set client and host retry limits to a minimum of 60 seconds and set protocol retry windows to 120 seconds for protocols that will support it.

## New Features for Microsoft Windows

---

Clustered Data ONTAP 8.2 introduces a range of new features to enhance the experience of Windows environments.

### Support for SMB 3.0

CIFS uses the underlying SMB protocol for network file sharing in Windows environments. SMB 3.0 adds new features to previous revisions of the SMB protocol (SMB 2.0 and 2.1) that improve nondisruptive operations (NDO) and other operations in Windows environments.

Continuous availability shares (CA) provide improved availability for Microsoft Windows Hyper-V. With earlier versions of the SMB protocol, clients had to reconnect to storage if a storage controller failover event occurred. With CA, file handles are persistent, allowing uninterrupted service during brief network outages and storage failovers.

When a storage failover occurs, *witness protocol* is used to alert clients to proactively move requests to the surviving storage node.

For client-side failovers, clients running SMB 3.0 specify an *application instance ID* when a file is opened. This ID is then maintained by the appropriate nodes on the NetApp cluster for the life of the file handle. If one client fails, the surviving client can use the ID to reclaim access to the file.

**Offloaded Data Transfer (ODX).** This new SMB 3.0 feature allows Windows clients to use NetApp storage to accomplish data copying, reducing the load on the host and network. This capability works within the same volume, between volumes on the same node, and between volumes on different nodes.

When a user initiates a file copy, the file is opened in the appropriate SVM, resulting in a token for the data to be copied. This token is then passed to the destination SVM with instructions to copy the data. The copy is initiated and the client is notified when it completes.

Whenever possible, the file to be copied is cloned rather than physically copied to save storage and reduce the time required to complete the operation.

**SMB Auto Location.** This NetApp exclusive feature optimizes data access by redirecting client requests to the logical interface (LIF) on the node hosting the volume. Because volumes can be moved dynamically in a NetApp cluster, it is possible that situations will occur in which clients access volumes through LIFs located on other cluster nodes.

With auto location, if a cluster node receives an SMB request for a volume on another node, it will refer the client to the IP address of a LIF on that node to satisfy future requests, optimizing the data path and minimizing latency.

**BranchCache.** BranchCache allows data to be cached locally on the client or a dedicated cache server to improve read performance, especially in situations in which data is accessed over a WAN. With BranchCache, numerous clients can all cache data and share access to that data between themselves (mediated by the NetApp cluster acting as a content server for the CIFS share) or a hosted cache server can be configured that pulls cached data from individual clients. BranchCache can be configured within each VSM on "all-shares" or "per-share."

## FPolicy

FPolicy enables control and configuration of file policies. Partner applications can use this framework to connect to NetApp storage to monitor and control file access operations. FPolicy was first introduced in Data ONTAP 6.4 and is now available for clustered Data ONTAP for the first time.

FPolicy provides various use cases like file blocking, quota management, file access auditing, and archiving. Common use cases include:

- **File archiving.** Archived files are restored upon access by the client.
- **File access monitoring.** Any file operation may be recorded.
- **File auditing.** This use case maintains access records for file objects.
- **File blocking.** Files such as videos and music can be blocked from storage.
- **Quota on directories.** This use case provides fine-grained control of quotas on directories.

## File Access Auditing

File access auditing gives you the ability to:

- **Monitor** access to protected resources and take action when necessary.
- **Prove** that a security event did (or did not) occur in situations in which evidence may be required.
- **Comply** with legal record-keeping requirements.
- **Recover** by using information from detailed auditing to roll back to a known good state.

Clustered Data ONTAP can log both successful and unsuccessful access to data as defined by NTFS ACLs and NFSv4 audit ACLs. Events are generated and presented at the SVM level, and access event records are generated across several distributed nodes.

## Conclusion

---

Clustered Data ONTAP version 8.2 offers numerous new features that continue to extend the capabilities of NetApp storage to make your storage environment more efficient and more scalable while eliminating the need for planned downtime. QoS lets you create a flexible policy framework to control the amount of resources that can be consumed by a given workload, and aggregate relocate lets you easily and quickly upgrade storage controllers with data in place. New capabilities for Microsoft Windows environments give you more options than ever before to

enhance Windows availability and control and audit file access.

---

› **Got opinions about  
Clustered Data OnTap 8.2?**

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



**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](#)
-