



Technical FAQ

Clustered Data ONTAP 8.1.x, 8.2.x, and 8.3.x

NetApp
March 2016

Unless otherwise noted, this document uses “clustered Data ONTAP®” to mean clustered Data ONTAP 8.1 and any subsequently available release. In this document we point out specific cases in which a subsequent release is required, such as clustered Data ONTAP 8.2.x or clustered Data ONTAP 8.3.x.

TABLE OF CONTENTS

1 Overview.....	8
What are the key features of clustered Data ONTAP?	8
What is new in clustered Data ONTAP 8.2.x?	8
What is new in clustered Data ONTAP 8.3.x?	9
Is Data ONTAP 8.3 a clustered ONTAP only release?	11
What is a cluster? Is it the same as an HA pair?	11
What are the physical components and the logical components in a cluster?	11
Is a node a single controller or an HA pair?	11
How do the nodes in the cluster communicate?	12
What is an SVM?	12
Can an SVM serve SAN and NAS?	12
How is an SVM different from a vFiler unit?	12
Does clustered Data ONTAP provide secure multi-tenancy?	12
Can an SVM contain aggregates?	12
Can I have multiple SVMs?	13
What is a namespace?	13
How many files can exist in a namespace? How many volumes?	13
What is an SVM root volume?	13
Is an SVM root volume the same as a node root volume?	13
It appears that an SVM can stretch across multiple nodes in the cluster. Does this mean that data is striped across the cluster?	13
Can I change the language setting of an SVM after it is created?	14
Why is clustered Data ONTAP referred to as virtualizing network and storage?	14
2 Hardware Support.....	15
Of how many nodes can a cluster consist?	15
Can I have a 24-node cluster but designate 6 nodes for SAN?	15
Which storage systems are supported?	15
Can I mix various storage systems?	15
Is FAS2240 supported?	15
Does clustered Data ONTAP support third-party arrays through the use of FlexArray (formerly V-Series)?	15
Are NetApp disks required for a FlexArray (formerly V-Series) installation?	16
Is Data ONTAP Edge available for clustered Data ONTAP?	16
Can an HA pair share a single shelf?	16
Which add-in cards can I use?	16
Are HA pairs required?	16
Is a cluster interconnect switch required?	16
Does it make sense to run a two-node HA pair cluster?	16
Is Flash Cache supported?	16
Is PCS supported?	16

Is Flash Pool supported?	17
Are SSD configurations supported?	17
Which platforms support Infinite Volumes?	17
Do Infinite Volumes support Flash Cache and Flash Pool?	17
Is NetApp Storage Encryption (NSE) supported?	17
Is a simulator available?	17
Is a FlexPod configuration available?	17
What is Advanced Drive Partitioning?	17
When a system is upgraded to clustered ONTAP 8.3, will Advanced Drive Partitioning be enabled?	17
What is Root/Data Advanced Disk Partitioning?	17
What configurations support Root/Data Advanced Disk Partitioning?	18
What is SSD Partitioning for Flash Pool™ Cache?	18
What are the maximum Flash Pool cache size increases in clustered ONTAP 8.3?	18
What is All-Flash FAS?	18
What does the AFF personality do?	18
What are the key differentiators of AFF?	18
What performance optimizations exist for All-Flash FAS?	19
How was performance optimized for All-Flash FAS?	19
What sizing considerations should be taken into account for AFF?	19
What is the maximum storage efficiency savings with inline compression?	19
Why is post-process deduplication used in conjunction with inline deduplication?	19
What other enhancements to Flash Pool and Flash Cache exist in clustered ONTAP 8.3?	20
Where can I find more information about supported configurations?	20
3 Software Support.....	20
Can I mix versions of clustered Data ONTAP in the same cluster?	20
Are qtrees supported?	20
Are qtree-level export policies supported?	20
Is there Kerberos support for AES 128 and 256?	21
What is an Infinite Volume?	21
Can a cluster contain multiple Infinite Volumes?	21
Is an in-place Infinite Volume to flexible volume conversion process available?	21
Can an Infinite Volume be a SnapMirror source for a flexible volume destination?	21
What is QoS?	21
Does QoS support FlexArray (formerly V-Series)?	21
What performance enhancements are introduced with clustered ONTAP 8.3?	21
Which legacy platforms benefit from these performance enhancements?	21
Does clustered ONTAP support Vmware VVOLs?	21
Does clustered ONTAP support inline zero write detection and elimination?	22
Does clustered ONTAP support FlexClone for server virtualization (SVI)?	22
What is Cloud ONTAP?	22
Does Cloud ONTAP support Encryption?	22

What is NetApp Private Storage (NPS)?	22
4 Storage Subsystem	22
How many disks are required for a node root aggregate?	22
What are the aggregate size limits?	22
Are 64-bit aggregates supported?	23
What are the volume size limits?	23
What is the maximum number of volumes?	23
Where can I find more information about the storage subsystem?	23
How many files can be in a flexible volume?	23
What are the size requirements for the node and SVM root volumes?	23
What is the Automated Workload Analyzer (AWA)?	23
Does clustered ONTAP 8.3 support 32-bit aggregates?	24
5 Licenses	24
Do all the nodes in a cluster require the same licenses?	24
If a cluster contains V-Series or FlexArray-enabled nodes, does V-Series or FlexArray need to be licensed on every node?	24
Where can I find information about licenses?	24
6 Installation, Upgrades, and Transition	25
Where can I find information on making the transition from 7-Mode?	25
What is Copy Free Transition (CFT)?	25
Can I upgrade from a previous version of clustered Data ONTAP?	25
How do I upgrade an entire cluster?	25
Can I add more nodes to the cluster?	25
Can I remove nodes from the cluster?	25
Can I expand 32-bit aggregates to 64-bit aggregates?	25
Can two clusters be merged?	26
What is System Setup?	26
What is an automated nondisruptive upgrade (NDU)?	26
7 Networking	27
What is a LIF?	27
When a NAS LIF migrates, does client access immediately move to the destination network port?	27
Can a SAN LIF migrate?	27
How many ports can host a LIF?	27
How many LIFs can I create?	27
How many LIFs can I assign to a port?	27
What types of LIFs exist?	27
Can a data LIF support more than one data protocol?	28
What is the cluster interconnect?	28
For what is the cluster interconnect used?	28
What protocol is used for the cluster interconnect?	28

What is the maximum distance a single cluster can span?	28
Can I directly connect nodes instead of using an interconnect switch?	28
Are redundant 10GbE cluster interconnect connections required?	28
Is the transition from a switchless to a switched cluster nondisruptive?	29
Are switch upgrades nondisruptive?	29
Is IPv6 supported?	29
Can customers use their own management switches?	29
Is load balancing among data ports supported?	29
Is SNMP supported?	29
Where can I get more information about the cluster interconnect?	29
How many simultaneous TCP connections are supported?	29
What are Broadcast domains?	29
What are subnet objects?	30
What happened to port roles in clustered ONTAP 8.3?	30
How do I manage routing tables in clustered ONTAP?	30
What are IPspaces?	30
8 Storage Protocols.....	30
Which storage protocols are supported for flexible volume access?	30
Which storage protocols are supported for Infinite Volume access?	30
Do any client-specific requirements exist for pNFS?	31
Do any host-specific requirements exist to support ALUA or MPIO?	31
Do CIFS clients recognize clustered Data ONTAP junctions?	31
Are SMB 3.0 continuously available (CA) shares supported?	31
Is off-board antivirus supported?	31
Is access-based enumeration supported?	31
Is BranchCache supported?	31
Is Volume Shadow Copy Service (VSS) supported?	31
Is file access auditing supported?	31
Is FPolicy supported?	31
Is FTP, SFTP, or HTTP supported?	32
Is NDMP supported?	32
Is it possible to initiate an NDMPcopy from a clustered Data ONTAP system?	32
Is WebDAV supported?	32
How do I control NAS access from individual clients or groups of clients?	32
Why does the client-side “showmount” command return only “/”?	32
Where can I get detailed information on NFS support?	32
What is offline Foreign LUN import (FLI)?	32
What enhancements exist to FLI in Data ONTAP 8.3.1?	33
What is “online FLI”?	33
What is selective LUN mapping?	33
Is Dynamic Access Control (DAC) supported?	33

Is Microsoft Management Console (MMC) supported?	33
Besides DAC and MMC support, what new CIFS features were introduced in clustered ONTAP 8.3? ...	33
Where can I find information about SAN interoperability and supported configurations?	33
9 Nondisruptive Operations	35
What is NDO?.....	35
What is vol move?	35
What is LUN move?.....	35
Does vol move require a SnapMirror license?	35
Is vol move automatic?	35
Can I move a volume while a SnapMirror operation is running?.....	35
Can I move volumes from one SVM to another SVM?	35
Can clones or mirrored volumes be moved using vol move?	35
What is LIF migrate?	36
Can CIFS clients survive a LIF migrate?	36
Can I change the hardware within the cluster nondisruptively using volume move?	36
Can I upgrade the controllers in an HA pair nondisruptively without using volume move?	37
Can ARL be used to downgrade controllers?	37
Can I add and remove storage shelves nondisruptively?	37
Can data LIFs fail over automatically to other nodes?	37
Can a LIF be configured to be tied to a volume so that if the volume is moved, the LIF also moves?	37
How do failover and giveback times compare to those with 7-Mode?	37
10 Data Protection and High Availability.....	38
Are Snapshot copies supported?.....	38
How many Snapshot copies can I create?	38
Is SnapMirror supported?	38
What is SnapMirror for Storage Virtual Machines?	38
What is SnapMirror for Storage Virtual Machines' "Identity Preserve Mode"?	38
What is SnapMirror for Storage Virtual Machines' "Identity Discard Mode"?	39
Can SnapMirror relationships be throttled?	39
Can I mirror between clusters using SnapMirror?.....	39
Are synchronous SnapMirror relationships supported?	39
Does the destination of a SnapMirror relationship have to be upgraded before the source?	39
What are the advantages of version flexible SnapMirror?	39
What do I need to know about version flexible SnapMirror?	39
What is SnapMirror network compression?	40
Can multiple SnapMirror relationships in the same cluster use overlapping IP addresses or subnets? ..	40
Where can I find information about SnapMirror?	40
Is SnapVault available?	40
Does SnapVault provide granular restoration of files and LUNs?.....	40
Can I mirror 7G/7-Mode systems to clustered Data ONTAP systems?	40
Can I mirror clustered Data ONTAP systems to 7G/7-Mode systems?	40

Can I mirror entire SVMs?	40
Is MetroCluster available?	40
Are two-node MetroCluster configurations available?	41
Is stretch MetroCluster available?	41
Is SyncMirror supported?	41
Is SnapLock software supported?	41
Is NDMP supported?	41
Where can I get more information about data protection?	41
11 Storage Efficiency	42
What storage efficiency features are supported?	42
Where can I get more information about storage efficiency features?	42
12 Management.....	43
How do I manage a clustered Data ONTAP system?	43
Are APIs available?	43
Must I manage each node individually?.....	43
Can event management system events be routed to different locations, depending on their severity?...	43
Can I measure latency?.....	43
Can I back up a cluster's configuration?	43
Can I use Operations Manager?.....	43
What information is available for storage administrators with 7-Mode experience?	44
13 Other Concepts and More Information	45
Does a technical report exist that I can give to my customer?.....	45
Where can I find more information?.....	45
What are the differences between 7-Mode and clustered Data ONTAP?.....	45
Are other clustered Data ONTAP FAQs available?	45

LIST OF TABLES

Table 1) SVM and vFiler unit differences.....	12
Table 2) Volume types supported for vol move.	36

1 Overview

What are the key features of clustered Data ONTAP?

Answer:

- Transparent, nondisruptive data movement between any type of storage on any node
- Unified architecture: namespace for NAS volumes, container for SAN LUNs
- NetApp® Snapshot™ technology
- Mirroring within and between clusters
- Transparent node failover
- Fully integrated management system
- Balanced scalable performance
- Scale nondisruptively from small to huge
- Scalable up to 24 nodes in a single NAS system
- Scalable to 4 (clustered Data ONTAP 8.1), 6 (clustered Data ONTAP 8.1.1), or 8 (clustered Data ONTAP 8.2.x and above) nodes in a SAN-enabled system
- Single system to manage
- Intelligent caching with Flash Cache™ or Flash Pool™
- No client code required: uses standard NFS and CIFS clients and MPIO/ALUA for SAN
- Highly reliable: built for nondisruptive operation
- WAFL® file system, RAID-DP® technology, active-active high-availability (HA) controllers
- Storage efficiencies including thin provisioning, deduplication, and compression
- FlexClone® volumes
- Qtrees for quotas
- Snapshot technology, asynchronous mirroring, NDMP

What is new in clustered Data ONTAP 8.2?

Answer:

- Volume-level, storage-efficient SnapVault® D2D backup
- Quality of service (QoS) workload management
- Quota management and file auditing using FPolicy® feature
- More volumes per controller and larger aggregates
- Larger SAN clusters and support for more LUNs
- Single-node and switchless 2-node clusters
- Encryption of data at rest
- Online controller hardware upgrade while keeping data in place
- More NetApp Infinite Volume capabilities and use cases
- IPv6 for protocols and SNMP
- Support for Hyper-V® on SMB 3.0 Continuously Available shares

What is new in clustered Data ONTAP 8.2.1?

Answer:

- CIFS features such as ABE, local users and groups, SMB signing, SMB 3, BranchCache, SSL-secured LDAP communication, and multiple CIFS domain searches for user mapping
- Support for SQL Server on SMB 3.0 Continuously Available shares
- Off-board antivirus for CIFS
- Nondisruptive shelf removal
- Automated Flash Pool cache sizing with Automated Workload Analyzer (AWA)
- Qtree-level export policies for NFSv3
- Nondisruptive aggregate updates from 32 bit to 64 bit without a PVR or additional disks
- Active Directory® authentication for administrators without a CIFS protocol license
- More export policies per cluster
- Flash Pool cache size increase from 400GB to 800GB per HA pair on entry platforms (applies to both clustered Data ONTAP and 7-Mode)

What is new in clustered Data ONTAP 8.3?

Answer:

- Support for MetroCluster: a synchronous mirroring solution that consists of two distinct, two-node clusters up to 200KM apart
- Support for SyncMirror
- All-flash FAS optimization, including optimized random read performance at high load
- Advanced drive partitioning for increased storage efficiency of all-flash FAS and entry platforms
- Advanced drive partitioning for Flash Pool, offering increased cache allocation flexibility between the flash component of Flash Pool aggregates, and reduction of parity and spare SSD cost for some Flash Pool cache configurations
- Support for larger Flash Pool cache configurations (4 times larger on most platforms)
- Automated Nondisruptive Upgrades (NDU): a three-command automated process for cluster upgrades
- DataMotion for LUNs: the ability to move a LUN from one volume in a cluster to another volume, with near-instantaneous cutover
- Offline Foreign LUN Import (FLI): support for importing a third-party LUN into clustered ONTAP
- Support for Data ONTAP version independent SnapMirror relationships (requires 8.3+ on the source and destination)
- Support for VM granular management (VVols)
- CIFS enhancements, including performance enhancements for file services workloads, support for Dynamic Access Control (DAC), and support for Microsoft Management Console (MMC) management of shares.
- NFS enhancements, including Qtree exports for NFSv4 and NFSv4.1, Kerberos support for AES 128 and 256, “showmount” support, and support for 1024 Unix groups.
- Support for Cloud ONTAP, a clustered ONTAP instance that runs directly on public cloud infrastructure
- Support for IPspaces, which allow SVMs in the same cluster to have overlapping subnets and IP addresses
- Support for broadcast domains and subnet objects: constructs that simplify clustered ONTAP networking
- On-box System Manager, which eliminates the need to download and install System Manager separately and maintain a specific Java version on clients
- On-demand AutoSupport collection, using a secure, outbound-only connection

- Support for optional network compression of SnapMirror transfers
- Cluster fan-in and fan-out ratio increases: 255:1 fan-in, 16:1 fan-out
- IPv6 enhancements, including support for IPspaces and cluster peering relationships
- Increased performance for most SnapMirror and SnapVault transfers and updates
- Selective LUN mapping: per-LUN mapping to HA pairs without using portsets
- Removal of the NetApp shelf requirement for FlexArray-enabled systems
- Performance enhancements for FlexClone for SVI: faster deployment, cloning, recomposition, and patching of virtual machines
- Release-aligned System Setup enhancements

What is new in clustered Data ONTAP 8.3.1?

Answer:

- Inline compression for primary workloads with All Flash FAS (AFF): performance equal to or better than Data ONTAP 8.3 running without compression in most cases
- Enhancements to Foreign LUN Import (FLI): the import process can occur while the LUN is in use. A brief disruption is required to re-target the host at the Data ONTAP system
- SnapMirror for Storage Virtual Machines (Storage Virtual Machine DR or SVM DR): a solution for mirroring the configuration of an SVM, not just its data
- Cluster Peering from multiple IPspaces: multiple replication relationships on the same cluster can use the same IP addresses or subnets
- Usability enhancements: support for the banner and Message of the Day (MOTD); enhancements for audit log management, including the ability to forward audit logs to a syslog server; support of LDAP and NIS user authentication for cluster access; new Flash Pool cache policies that align with similar Flash Cache policies; and Automated Workload Analyzer (AWA) volume-level reporting
- Protocol enhancements: support for dynamic DNS; support for Windows NFSv3 clients; support for SMB encryption for data transfers over SMB 3.0; support for mapping Windows users from a different domain a guest UNIX user; support for mapping all members of the Windows administrators group to root; enhancements to SQL Server (DBCC workarounds removed) and Hyper-V (enhancements to TRIM support) over SMB
- Availability enhancements: support for a LIF sufficiency checker that helps prevent configurations where a LIF attempts to fail over to an ineligible port; checks to prevent creation of a nonroot volume in the root aggregate; support for displaying license entitlement risk status
- Cloud ONTAP encryption: enables customers to encrypt data in the public cloud while retaining control of the encryption keys; delivers Data ONTAP software encryption with aggregate granular controls
- Support for 2-node MetroCluster Configurations

What is new in clustered Data ONTAP 8.3.2?

Answer:

- Opportunistic inline deduplication for primary workloads
- Quality of Service (QoS) support for up to 24 node NAS clusters by PVR
- 12,000 QoS policy groups by PVR
- Improved performance with a compressible random I/O larger than 8KB
- In-place adaptive compression: compresses ingested uncompressed data that is not compressed inline. For example, through the 7MTT to All Flash FAS or to a volume move from a FAS node to an All Flash FAS node
- Reduced AFF SAN latency for most workloads on the order of 100-150us

- Copy free transition (CFT): transition from Data ONTAP operating in 7-mode to clustered Data ONTAP with a procedure similar to a head swap
- System Manager enhancements: cluster performance charts, manual IP address assignment, SyncMirror support, inline deduplication support, volume move details, MirrorVault relationship management, initial cluster configuration support, SVM dashboard
- Support for logging and forwarding authentication failures
- Support for FCoIP (Fibre Channel over IP) with MetroCluster
- LDAP Signing and Sealing available through PVR
- All Flash FAS FLI by PVR
- Simplified System Manager workflow for NAS All Flash FAS
- HTTP support by FPVR
- 300 km MetroCluster configurations

Is Data ONTAP 8.3 a clustered ONTAP only release?

Answer: Yes. Data ONTAP 8.3 is a clustered ONTAP only release. 7-mode support will continue in the 8.2 release line. NetApp will provide updates, bug fixes or patches as deemed necessary and commercially reasonable for its 7-mode deployments for 3 years beyond End of Availability (EoA) of FAS8000 and FAS2500 platforms in Data ONTAP 8.2.x releases.

What is a cluster? Is it the same as an HA pair?

Answer: A cluster refers to a group of nodes running clustered Data ONTAP 8.x. In clustered Data ONTAP 8.2, a cluster can consist of 1 to 24 nodes. In clustered Data ONTAP 8.1 and clustered Data ONTAP 8.1.1, a cluster can consist of 2 to 24 nodes.

We formerly used the term "cluster" to refer to an HA pair in a 7G/7-Mode context. This term should no longer be used. HA pair is the correct term for this going forward.

What are the physical components and the logical components in a cluster?

Answer: Clustered Data ONTAP systems consist of logical and physical entities. Physical entities make up the cluster. Logical entities make up the storage virtual machines (SVMs) that provide data access using the cluster.

Physical:

- Nodes
- Ports
- Aggregates
- Disks

Logical:

- SVMs
- Logical interfaces (LIFs)
- Volumes
- LUNs

Is a node a single controller or an HA pair?

Answer: A "node" refers to a single instance of clustered Data ONTAP. Each HA pair contains two nodes.

How do the nodes in the cluster communicate?

Answer: Nodes in a clustered Data ONTAP system communicate through a dedicated 10 Gigabit Ethernet (10GbE) interconnect.

Starting with clustered Data ONTAP 8.2, single-node clusters without an Ethernet interconnect are supported. Clustered Data ONTAP 8.2 also supports two-node switchless clusters.

The networking section in this document provides additional information about the cluster interconnect.

What is an SVM?

Answer: An SVM is a storage virtual machine. It can be thought of as a secure virtual storage system that manages resources, including volumes and LIFs. An SVM is called a Vserver in the clustered Data ONTAP command line.

The physical resources of the cluster are not bound to any particular SVM. An SVM's volumes can be moved to different physical aggregates without disrupting client access. Similarly, an SVM's LIFs can be moved to different physical network ports without disrupting client access.

Can an SVM serve SAN and NAS?

Answer: Yes. An SVM can serve both SAN and NAS concurrently, just as a 7-Mode storage system can.

How is an SVM different from a vFiler unit?

Answer: Several similarities between vFiler® units in 7-Mode and clustered Data ONTAP SVMs exist. Both allow a layer of storage virtualization, and both allow administrative control of just that virtual storage unit. However, some differences exist between the two.

Table 1) SVM and vFiler unit differences.

Clustered Data ONTAP SVM	7-Mode MultiStore vFiler Unit
Required. Clustered Data ONTAP needs at least one SVM defined in order to serve data.	Optional. A controller running Data ONTAP in 7-Mode can serve data without a vFiler unit defined.
Serves a single namespace. Multiple flexible volumes can be accessed using a single CIFS share or NFS export.	Volumes must be exported individually.
Uses resources from one or many nodes within the cluster.	Bound to the resources of a single controller.
Supports any clustered Data ONTAP data protocol, including FCP and FCoE.	Limited to NFS, CIFS, and iSCSI.
Provides fully delegated role-based access control (RBAC).	Provides limited administration delegation.

Does clustered Data ONTAP provide secure multi-tenancy?

Answer: Yes. SVMs can be isolated in their own separate VLANs using separate physical ports or using VLAN tagging. Users of one SVM might or might not be granted access to another SVM. Administration can be delegated to SVM administrators independently as well.

Can an SVM contain aggregates?

Answer: No. An SVM manages volumes. However, a cluster administrator can limit which aggregates an SVM can use for volume creation.

Can I have multiple SVMs?

Answer: Yes. Clustered Data ONTAP supports up to a maximum of 1,000 SVMs per cluster, depending on the number of controllers in the cluster, the number of ports in each controller, and the protocols enabled. Details are in the “System Administration Guide” and TR-4160i, “Secure Multi-Tenancy in Clustered Data ONTAP.”

What is a namespace?

Answer: In clustered Data ONTAP, FlexVol® volumes containing NAS data are junctioned into the owning SVM in a hierarchy. This hierarchy presents NAS clients with a unified view of the storage, regardless of the physical location of FlexVol volumes inside the cluster. A volume must be junctioned into a namespace in order to be reachable by NAS clients.

Junctions allow each FlexVol volume to be browsable like a directory or folder. NFS clients can access multiple FlexVol volumes using a single mount point. CIFS clients can access multiple FlexVol volumes using a single CIFS share. The NAS namespace consists of the hierarchy of FlexVol volumes within a single SVM as presented to the NAS clients.

For more information about namespaces, see TR-4129, “Namespaces in Clustered Data ONTAP.”

How many files can exist in a namespace? How many volumes?

Answer: Because a namespace consists of multiple flexible volumes, the maximum number of files in a namespace is the product of the number of volumes in the namespace and the maximum number of files in each of those volumes. The same goes for capacity. Both are influenced by the controller types that make up your cluster.

A single namespace can contain all the volumes in a cluster: up to 12,000 in a very large, high-end cluster.

For more information about high-file-count environments, see TR-3537, “High-File-Count Environment Best Practices.”

What is an SVM root volume?

Answer: Each SVM has a root volume that serves as the top of SVM’s namespace. Other volumes are junctioned beneath the SVM root volume. An SVM root volume typically contains only junctions and directories. The volumes junctioned under the SVM root volume may contain junctions to more volumes and so on.

Is an SVM root volume the same as a node root volume?

Answer: A node root volume is associated with a particular physical cluster node. It is bound to that cluster node. It contains configuration files, logs, and other files associated with a node’s normal operation. A node root volume is part of the physical cluster infrastructure. It is not associated with an SVM, does not hold user data, and does not contain junctions to other volumes.

An SVM root volume is not bound to any physical cluster node. It can be nondisruptively moved between cluster nodes like any other flexible volume in an SVM. It is part of the virtual storage infrastructure that serves data to NAS clients and SAN hosts.

It appears that an SVM can stretch across multiple nodes in the cluster. Does this mean that data is striped across the cluster?

Answer: No. Data is not striped across nodes within the cluster. An SVM uses various physical resources, potentially distributed across many nodes, to serve data. Each flexible volume, however, resides on a single aggregate.

Can I change the language setting of an SVM after it is created?

Answer: Yes. Starting with clustered Data ONTAP 8.2, the administrator can change an existing SVM's language, and volumes in a single SVM can have different language settings.

A volume's language setting cannot be changed after the volume is created. This is also true of the SVM root volume. Changing an SVM's language changes the default language for new volumes, but it does not change the SVM's root volume language or the language of existing volumes.

In clustered Data ONTAP 8.1.x, an SVM's language setting cannot be changed, and all volumes in the SVM must share a language setting with the SVM.

In almost all cases, volumes should use the C.UTF-8 language setting. Detailed recommendations can be found in TR-4160, "Secure Multi-Tenancy in Clustered Data ONTAP."

Why is clustered Data ONTAP referred to as virtualizing network and storage?

Answer: The software components of clustered Data ONTAP that provide client-facing protocol access are separated from the software components that access the data.

This separation allows an SVM to be abstracted from the hardware it uses and allows independent mobility of an SVM's LIFs and flexible volumes. An SVM is virtualized because its storage and network connections are not permanently bound to any node or group of nodes.

2 Hardware Support

Of how many nodes can a cluster consist?

Answer: The maximum size of a cluster depends on a number of factors, including the clustered Data ONTAP version, the controller types that make up the cluster, and whether or not SAN is enabled.

For a detailed breakdown of supported cluster hardware and software configurations, see the “[Clustered Data ONTAP Platform Mixing Rules Technical FAQ](#),” available on the Field Portal.

Can I exceed the node limit to perform a hardware refresh?

Answer: The node limit cannot be exceeded to perform lifecycle operations. DataMotion™ for Volumes (that is, vol move) and the new functionality of aggregate relocate (ARL) can be used in clustered Data ONTAP 8.2 and later to nondisruptively upgrade controllers without exceeding the node limit.

Can I have a 24-node cluster but designate 6 nodes for SAN?

Answer: No. SAN-enabled clustered Data ONTAP systems can scale up to 4 (clustered Data ONTAP 8.1), 6 (clustered Data ONTAP 8.1.1), or 8 (clustered Data ONTAP 8.2) nodes if SAN is enabled, with or without NAS.

Running SAN on a subset of nodes in a larger cluster is not supported.

Which storage systems are supported?

Answer: Information about supported platforms and configurations is in the NetApp Hardware Universe: <http://hww.netapp.com/home.aspx>.

Clustered Data ONTAP 8.2 does not support FAS2040, FAS3040, FAS3070, FAS6070, or FAS6030.

Clustered Data ONTAP 8.3 does not support FAS6080, FAS6040, FAS3170, FAS3160, FAS3140, or FAS3210

Can I mix various storage systems?

Yes. Clusters can be heterogeneous, so long as all the nodes in the cluster support the same version of clustered Data ONTAP and the cluster platform mixing rules are followed.

Both members of each HA pair must be the same controller type.

See the “Clustered Data ONTAP Platform Mixing Rules Technical FAQ” on the Field Portal for more information.

Is FAS2240 supported?

Answer: Yes. Clustered Data ONTAP supports FAS2240. Single-node FAS2240 clusters support an 8GB Fibre Channel mezzanine card. Multinode FAS2240 configurations require a 10GbE mezzanine card for use as a cluster interconnect. Single-node FAS2240 clusters do not require a 10GbE mezzanine card.

Does clustered Data ONTAP support third-party arrays through the use of FlexArray (formerly V-Series)?

Answer: Yes. Consult the Support Matrix at <http://support.netapp.com/matrix> for details on third-party array support.

The FAS8000 series is supported starting with clustered ONTAP 8.2.1. The FAS8000 series of controllers does not have a corresponding V-Series product line. Support for array virtualization is available in the FAS8000 product line through the use of the FlexArray software option.

Are NetApp disks required for a FlexArray (formerly V-Series) installation?

Answer: Clustered ONTAP 8.3 can boot using a foreign LUN, and is therefore not subject to the NetApp disk requirement. Clustered Data ONTAP 8.2.X and earlier systems require a shelf of NetApp disks for the root volume.

Is Data ONTAP Edge available for clustered Data ONTAP?

Answer: Yes. See the July 2014 Data ONTAP Edge Release Update on Field Portal for additional details.

Can an HA pair share a single shelf?

Answer: Yes. However, NetApp does not generally recommend this type of configuration and it should be required only in certain smaller installations.

Which add-in cards can I use?

Answer: See the NetApp Hardware Universe: <http://hwu.netapp.com/home.aspx>.

Are HA pairs required?

Answer: Clustered Data ONTAP 8.2 supports single-node clusters. All clustered Data ONTAP 8.2 systems with two or more nodes require HA pairing for all nodes.

Clustered Data ONTAP 8.1 and clustered Data ONTAP 8.1.1 do not support single-node clusters. Every node in a clustered Data ONTAP 8.1 or clustered Data ONTAP 8.1.1 system must be a member of an HA pair.

Is a cluster interconnect switch required?

Answer: Clustered Data ONTAP 8.2 supports direct 10GbE connections in two-node switchless clusters. Clustered Data ONTAP 8.2 also supports single-node clusters, which do not require a cluster interconnect switch.

All clustered Data ONTAP 8.1 and clustered Data ONTAP 8.1.1 systems require a cluster interconnect switch.

Does it make sense to run a two-node HA pair cluster?

Answer: Yes. All the benefits of clustered Data ONTAP, including nondisruptive operations with volume move, LIF migrate, multipath ALUA/MPIO for SAN, online controller hardware upgrade, and so on, are available with a two-node system.

Is Flash Cache supported?

Answer: Yes. More information about Flash Cache, including operating modes, cache warming, and interaction with other clustered Data ONTAP features, is in the Flash Cache technical FAQ and TR-3832, "Flash Cache Best Practices Guide."

Is PCS supported?

Answer: Yes. Predictive cache statistics (PCS) are supported. They are available as node-scoped commands similar to those provided in 7-Mode.

Clustered ONTAP 8.2.1 also supports the Automated Workload Analyzer (AWA), which can be used for Flash Pool sizing.

Is Flash Pool supported?

Answer: Yes. Information about Flash Pool is in the Flash Pool technical FAQ and in [TR-4070, “Flash Pool Design and Implementation Guide.”](#)

Are SSD configurations supported?

Answer: Yes. See the “Pure SSD and Flash Pool with DS424X Disk Shelf” product bulletin on the Field Portal for more information.

Which platforms support Infinite Volumes?

Answer: In Data ONTAP 8.2, Infinite Volume is supported on FAS or V-Series 3000, 6000, or 8000 series systems. 2000 series systems are not supported with Infinite Volume.

Do Infinite Volumes support Flash Cache and Flash Pool?

Answer: Yes. Both are supported in clustered Data ONTAP 8.2. See TR-4178, “Infinite Volume Deployment and Implementation Guide,” and the “Infinite Volume Technical Collateral Binder” on the Field Portal for details.

Is NetApp Storage Encryption (NSE) supported?

Answer: Yes. See the “NetApp Storage Encryption FAQ” available on the Field Portal for more information.

Is a simulator available?

Answer: Yes. A clustered Data ONTAP simulator is available for download from the NetApp Support (formerly NOW®) site.

Is a FlexPod configuration available?

Answer: Yes. See TR-4114, “VMware vSphere 5.0 on FlexPod Clustered Data ONTAP Deployment Guide,” for details.

What is Advanced Drive Partitioning?

Answer: “Advanced Drive Partitioning” is a technology NetApp is releasing with Data ONTAP 8.3 that allows a single physical drive to be shared by multiple aggregates: two for Root/Data Advanced Disk Partitioning, four for SSD Partitioning for Flash Pool cache. Each aggregate using the drive is permitted to use a system-defined range of blocks on that drive. Initiators from both controller nodes in an HA pair can access the drive concurrently, if needed.

NetApp supports advanced drive partitioning for three configurations: Root/Data Partitioning for all-flash FAS, and entry platforms, and the cache portion of Flash Pool aggregates. Note that “advanced drive partitioning” was previously referred to as “disk slicing”, a term that has been retired.

When a system is upgraded to clustered ONTAP 8.3, will Advanced Drive Partitioning be enabled?

Answer: Not automatically. Eligible systems can be upgraded to clustered ONTAP 8.3, but Advanced Drive Partitioning will not be enabled by default. Partitioning requires a low-level re-initialization of the drives, and as such, the system must first be evacuated using a nondisruptive procedure like DataMotion for Volumes or DataMotion for LUNs.

What is Root/Data Advanced Disk Partitioning?

Answer: Planned takeover and giveback operations in clustered Data ONTAP operate on one aggregate at a time. This reduces the takeover and giveback time required for the volumes and LUNs on each

individual aggregate. This means that an aggregate must be dedicated to the node root volume so it can be given back first, in order to bring the node online to give back the aggregates containing user data.

Root/Data Advanced Disk Partitioning frees the size of the node root aggregate from the constraints imposed by the physical storage media. When Root/Data Advanced Disk Partitioning is used, a single set of disks can contain both the node root aggregate and a single data aggregate. This increases storage efficiency for those configurations that support Root/Data Advanced Disk Partitioning.

What configurations support Root/Data Advanced Disk Partitioning?

Answer: Root/Data Advanced Disk Partitioning is supported for All-Flash FAS and entry configurations. Up to 24 SSDs per node can be shared for root and data aggregates in an All-Flash FAS configuration.

Root/Data SSD partitioning will be the default configuration for All-Flash FAS8000 and FAS2500 series systems shipped with Data ONTAP 8.3.

What is SSD Partitioning for Flash Pool™ Cache?

Answer: Advanced Drive Partitioning for Flash Pool Cache provides better SSD capacity efficiency by sharing SSDs across multiple Flash Pool™ aggregates. Unlike Root/Data Advanced Disk Partitioning, SSD Partitioning for Flash Pool Cache divides each drive into four partitions, rather than two.

SSD Partitioning for Flash Pool Cache reduces the parity and spare SSD cost with up to 4 Flash Pool™ aggregates. It can boost cache performance in some cases by allowing for more data SSDs per cache.

SSD Partitioning for Flash Pool Cache is supported for SSDs only in Flash Pool™ aggregates. It does not provide partitioning of HDDs in Flash Pool™ aggregates. SSD partitioning is not supported on all-SSD (aka all-flash) aggregates.

What are the maximum Flash Pool cache size increases in clustered ONTAP 8.3?

Answer: The maximum Flash Pool cache sizes per HA pair are four times larger for most platforms in clustered ONTAP 8.3. See the Flash Pool FAQ for additional details.

What is All-Flash FAS?

Answer: AFF is an 8000 Series HA pair with SSDs only, running 8.3 GA or later with the flash-optimized personality enabled. A FAS HA pair configured with SSDs running any other Data ONTAP configuration is not an AFF system.

AFF8000 systems are different from FAS8000 systems that are configured with SSDs. The differentiation is in Data ONTAP configuration settings. In general, converting a FAS8000 with SSDs to a FAS8000 is NOT supported. If there is a compelling business reason for such a conversion, a PVR for a one-way, one-time only conversion can be submitted for approval.

What does the AFF personality do?

Answer: In Data ONTAP 8.3, the AFF personality disables the addition of spinning media. In Data ONTAP 8.3.1, AFF systems have inline compression and zero-block deduplication enabled by default.

What are the key differentiators of AFF?

- Proven, audited performance with SPC-1 publication (http://www.storageperformance.org/results/benchmark_results_spc1_top-ten)
- Unified multiprotocol access: FC, iSCSI, NFS and CIFS
- Nondisruptive data mobility (LUNs, volumes, Vvols) between flash, hybrid, and HDD-only storage tiers (all can exist in the same cluster)
- SnapMirror for DR and SnapVault for long-term snapshot retention supported on all media types
- All-flash synchronous mirroring with MetroCluster

- Native data mobility to a public cloud with cloud ONTAP or NPS
- High IOPS throughput at consistently low latency

What performance optimizations exist for All-Flash FAS?

Answer: Clustered ONTAP 8.3 contains performance and enhancements for All-Flash FAS. The primary goal of these enhancements is performance resiliency: continued low latency when a large number of I/O operations are demanded from the system.

WAFL is a write-optimized disk layout, and spinning media performs well on sequential I/O, so the focus of all-flash performance optimization in clustered ONTAP 8.3 is random read operations. Clustered ONTAP 8.3 can deliver up to a 50%+ improvement in IOPS at a 1ms latency with 8KB random read operations on an all-flash FAS8080EX over 8.2.X.

Clustered ONTAP 8.3.1 further improves AFF performance, and delivers inline compression for primary workloads, with performance similar to or better than 8.3 without compression.

Clustered Data ONTAP 8.3.2 GA will automatically enable inline deduplication on NetApp® All Flash FAS volumes that are created with NetApp Data ONTAP® 8.3.2 GA. Data ONTAP 8.3.2 supports “in-place” adaptive compression to compress ingested data that isn’t compressed inline (for example, through 7MTT or vol move). Inline adaptive compression processing is more parallelized in Data ONTAP 8.3.2. The performance with compressible random 16KB and 32KB I/Os now looks more like it did with Data ONTAP 8.3 and no inline compression. That is the case for random 8KB I/Os when comparing 8.3.1 and 8.3, so with 8.3.2, it applies with larger I/Os as well.

Data ONTAP 8.3.2 implements an auto-assignment policy for AFF that implements the performance best practice of assigning 12 SSDs per shelf to each node in the HA pair. Thin provisioning is the default volume policy (no space guarantee) for all new volumes (SAN and NAS) in Data ONTAP 8.3.2.

How was performance optimized for All-Flash FAS?

Answer: Data ONTAP 8.3 bypasses most of WAFL on the return path from read operations. Data ONTAP 8.3.1 bypasses most of RAID on the return path from most read operations, unless reconstruction or recovery is required.

What sizing considerations should be taken into account for AFF?

Answer: With All-Flash arrays in general, the controllers must be powerful because flash media is no longer the bottleneck it once was with spinning media. Therefore, controllers must be sized appropriately.

This information is available in SPM. This is especially important when a workload is write-heavy as write operations require more CPU when inline compression is enabled.

What is the maximum storage efficiency savings with inline compression?

Answer: The maximum storage efficiency savings with inline compression *alone* is 2:1. If every 8K group compresses to a 4K block, the effective capacity of the system doubles.

With benchmarks designed to mimic real-world database workloads, storage efficiency savings have been closer to 1.8:1 with inline compression alone.

These efficiency savings reflect inline compression only, and do not take into account other storage efficiency features like Snapshot copies or RAID-DP.

Why is post-process deduplication used in conjunction with inline deduplication?

Answer: Inline deduplication in 8.3.2 deduplicates a subset of the blocks in the volume, not always all of them. This is because the block fingerprint hash is located in memory, and is subject to size constraints. Therefore post process or always-on deduplication is still used in conjunction with inline deduplication in Data ONTAP 8.3.2.

What other enhancements to Flash Pool and Flash Cache exist in clustered ONTAP 8.3?

- An increase in the maximum Flash Cache per node on the 8080 EX to 12TB from 8TB in 8.2.2.
- Random writes are the Flash Cache automatically, which will improve the cache hit rate and improves the performance of random read after random write workloads.
- Flash Pool has a new cluster shell CLI, removing one of the few remaining reasons to use the node shell.
- Configuration of cache policy options has been simplified by grouping them in pairs.
- The 16KB cutoff for overwrite insertion into the SSD cache portion of a Flash Pool has been removed, allowing overwrites of all sizes to be placed in the Flash Pool cache, when Flash Pool is operating in SSD write cache mode.
- Compressed blocks can be cached in Flash Pool SSDs.
- Flash pool caches can now cache read-only volume contents, like SnapMirror and SnapVault destinations.
- The Flash Pool cache capacity reserve was decreased, resulting in a 13% increase in cached data capacity.

Where can I find more information about supported configurations?

Answer: See the NetApp Hardware Universe: <http://hwwu.netapp.com/home.aspx>.

3 Software Support

Can I mix versions of clustered Data ONTAP in the same cluster?

Answer: The same cluster can run multiple versions of clustered Data ONTAP while the cluster is being upgraded. NetApp does not recommend using mixed versions for an extended period of time.

No more than two versions of clustered Data ONTAP can run in a cluster at a time.

The two versions must be no more than one major release apart. For example, you cannot upgrade a cluster directly from 8.0 to 8.2. You must upgrade the entire cluster from 8.0 to 8.1 to 8.2.

Features available in the higher version of clustered Data ONTAP will not be available until the entire cluster is upgraded.

Are qtrees supported?

Answer: Yes, qtrees are supported. However, no qtree-level replication, deduplication, or QoS is provided. Qtrees can be subject to user and group quotas. Qtrees can be mounted directly by a client using a “deep mount” into the containing SVM’s namespace.

In clustered Data ONTAP 8.2 and 8.1.x, qtrees always inherit the export policy of the parent volume. Qtrees in clustered Data ONTAP 8.2.1 inherit the export policy of the parent volume by default, but also support qtree-level export policies for NFSv3.

Are qtree-level export policies supported?

Answer: Yes, NFSv3 qtree-level export policies are supported starting with clustered data ONTAP 8.2.1. NFSv4 and NFSv4.1 qtree-level export policies are supported starting with clustered ONTAP 8.3.

Since DataMotion for volumes, QoS, and replication are managed at the flexible volume level, NetApp recommends tree-level export policies in clustered Data ONTAP only for 7-Mode transition scenarios in

which too many 7-Mode qtrees exist for a qtree-to-volume transition. See the [Transition Binder](#) on the Field Portal for more information about making the transition from 7-Mode to clustered Data ONTAP.

Is there Kerberos support for AES 128 and 256?

Answer: Yes, starting with clustered ONTAP 8.3.

What is an Infinite Volume?

Answer: An Infinite Volume is a single volume that can use storage on multiple controllers and can reach up to 20PB in size. See TR-4178, “Infinite Volume Deployment and Implementation Guide,” and the “Infinite Volume Technical Collateral Binder” on the Field Portal.

Can a cluster contain multiple Infinite Volumes?

Answer: Yes. Starting with Data ONTAP 8.2, no special limits are placed on the number of SVMs using Infinite Volume, flexible volumes, or both. However, NetApp recommends that the number of SVMs using Infinite Volume does not exceed the number of nodes in the cluster.

Is an in-place Infinite Volume to flexible volume conversion process available?

Answer: No. Host-side utilities should be used.

Can an Infinite Volume be a SnapMirror source for a flexible volume destination?

Answer: No. However, SnapMirror® relationships between two Infinite Volumes are supported.

What is QoS?

Answer: QoS stands for “quality of service.” QoS is a mechanism introduced in clustered Data ONTAP 8.2 that allows an administrator to limit the number of I/O operations per second or raw throughput (MB/s) directed to a policy group that could consist of a single SVM or a group of LUNs, flexible volumes, or files within an SVM.

Does QoS support FlexArray (formerly V-Series)?

Answer: Yes. QoS is supported on FlexArray.

What performance enhancements are introduced with clustered ONTAP 8.3?

Answer: The primary goal of the performance enhancements in clustered ONTAP 8.3 is performance resiliency: continued low latency when a large number of I/O operations are demanded from the system.

These enhancements include improvements to random read performance for All-Flash FAS. They also contain performance improvements for CIFS file services workloads and data protection relationships, like SnapMirror and SnapVault.

Which legacy platforms benefit from these performance enhancements?

Answer: High-end legacy platforms have comparable performance to 8.1.x and 8.2.1. Entry and midrange legacy platforms have comparable SAN performance, but up to ~15% lower NAS performance compared to 8.1.x.

Does clustered ONTAP support Vmware VVOLs?

Answer: Yes. Clustered ONTAP 8.3 supports vDisk restore, move, and copy on demand.

Does clustered ONTAP support inline zero write detection and elimination?

Answer: Clustered ONTAP 8.3 delivers inline zero write detection and elimination. When this feature is enabled and blocks containing zeros are written, they will be detected and replaced with space and time efficient zeroing. This will make sure that zero writes (like disk zeroing) do not consume system resources like disk bandwidth and space, enabling them to be better utilized for user IO and operations.

Does clustered ONTAP support FlexClone for server virtualization (SVI)?

Answer: Yes. Clustered ONTAP supports FlexClone for SVI. Clustered ONTAP 8.3 includes a number of enhancements to rapidly deploy, clone, recompose, and patch virtual machines, offers more consistent performance using QoS integration, and offers a faster delete engine for FlexClone files and LUNs.

What is Cloud ONTAP?

Answer: Cloud ONTAP is a single clustered Data ONTAP instance running directly on the cloud infrastructure, rather than on a dedicated FAS platform. It is available starting with clustered ONTAP 8.3.

Cloud ONTAP provides NetApp's industry leading data management in the cloud. Clustered Data ONTAP has far superior snapshots, replication, storage efficiency, data protection, and protocol support than native cloud storage.

Cloud ONTAP is subject to the availability of the cloud itself, which is not as high as can be achieved with a purpose built appliance in a well-run data center.

Does Cloud ONTAP support Encryption?

Answer: Yes. Cloud ONTAP in Data ONTAP 8.3.1 allows customers to encrypt data in a public cloud while maintaining exclusive access to their encryption keys.

Aggregate-level encryption with cloud ONTAP shares a common certificate and key management framework with NetApp Storage Encryption (NSE).

What is NetApp Private Storage (NPS)?

Answer: NetApp private storage is for those customers who need their data in a low-latency colocation facility next to the cloud. NPS allows them to use the flexible compute resources of the cloud, while maintaining control of their data and achieving the full value of clustered Data ONTAP.

4 Storage Subsystem

How many disks are required for a node root aggregate?

Answer: By default, node root volumes exist in a three-disk RAID-DP aggregate. Two-disk RAID 4 root aggregates are supported but not recommended in clustered Data ONTAP 8.2.

A node root aggregate should contain only the node root volume. It should not contain flexible volumes that house user data. A nonroot volume in the node root aggregate is unsupported and likely to experience an outage during a takeover or giveback event.

Clustered ONTAP 8.3 supports Root/Data Advanced Disk Partitioning on entry platforms and All-Flash FAS configurations, which allows disks to be shared between a root and data aggregate.

What are the aggregate size limits?

Answer: Aggregate size limits vary by platform and clustered Data ONTAP version. Information on aggregate size limits is in the NetApp Hardware Universe: <http://hwwu.netapp.com/home.aspx>.

Are 64-bit aggregates supported?

Answer: Yes, 64-bit aggregates are supported and are now the default aggregate type. They allow for advanced features like compression, Flash Pool™, storage-efficient vaulting, and larger aggregate and volume sizes.

Clustered ONTAP 8.3 supports 64-bit aggregates only. Before an upgrade to clustered ONTAP 8.3, all aggregates must be upgraded to 64-bit, and any snapshots that were created when the containing volume was in a 32-bit aggregate must be either deleted or allowed to age off the system.

What are the volume size limits?

Answer: Flexible volume size limits vary by platform and clustered Data ONTAP version. Information on flexible volume size limits is in the NetApp Hardware Universe: <http://hww.netapp.com/home.aspx>.

What is the maximum number of volumes?

Answer: The maximum number of flexible volumes varies depending on the controllers that make up the cluster. A cluster can host no more than 12,000 flexible volumes. More information is available in the Storage Subsystem Technical FAQ.

Where can I find more information about the storage subsystem?

Answer: Much more detail is in the Storage Subsystem Technical FAQ.

How many files can be in a flexible volume?

Answer: File system limits are covered in TR-3537, “High-File-Count Environment Best Practices.”

What are the size requirements for the node and SVM root volumes?

Answer: When a node is initialized, a root aggregate containing a root volume is created. The root aggregate uses three disks, and these are dedicated to clustered Data ONTAP.

No data volumes may be placed on the root aggregate. This is because root and data aggregates are handled differently during giveback operations. Additional information on root volumes is in the storage subsystem technical FAQ.

Each SVM has one root volume, which resides in a data aggregate. An SVM root volume is created by default when the SVM is created. The SVM root volume can be resized as needed.

What is the Automated Workload Analyzer (AWA)?

Answer: The Automated Workload Analyzer is an efficiency feature that makes it easier for customers, partners, and PS personnel to estimate the performance gains that will be achieved by adding Flash Pool.

AWA determines the cacheable portion of random read-write workloads for an aggregate. It recommends the appropriate Flash Pool cache size and provides the expected workload handled by SSD. It is usable when Flash Cache or Flash Pool is installed on the system.

When sizing a system for Flash Pool with SPM, you must provide several inputs that you may not know well or at all, including read-write mix, percentage of random overwrites, and the cache size needed to effectively hold the working set.

AWA streamlines this process by automatically capturing these inputs from a FAS or V-Series system that is running Data ONTAP 8.2.1.

AWA works on both clustered Data ONTAP and 7-Mode systems running Data ONTAP 8.2.1.

Does clustered ONTAP 8.3 support 32-bit aggregates?

Answer: No. 32-bit aggregates are not supported in clustered ONTAP 8.3. That also means that 32-bit Snapshots, SnapMirror, and SnapVault relationships are not supported in 8.3.

Starting with Data ONTAP 8.2.1, in-place, nondisruptive expansion of aggregates from 32-bit to 64-bit is supported without requiring additional disks. Anything below 8.1.4P4 and 8.2.0 requires a PVR to perform the nondisruptive upgrade.

Only the active file systems can be expanded to 64-bit using the nondisruptive expansion procedure. Snapshot® copies either need to be deleted manually or the customer can wait for them to expire and be deleted automatically after the containing aggregate is expanded to 64-bit.

The “snap list” command can be used on 8.2.1 to identify 32-bit snapshots. There are code checks in 8.3 that will prevent an upgrade if 32-bit data is present on the system.

There are several strategies customers with long-term snapshot copies and vault relationships available in our transition collateral on field portal. See the Transition SE Presentation for more details.

5 Licenses

Do all the nodes in a cluster require the same licenses?

Answer: Yes. Licensing is symmetrical in clustered Data ONTAP. With the exception of V-Series and FlexArray licensing, every node in a cluster must share the same set of licensed features.

If a cluster contains V-Series or FlexArray-enabled nodes, does V-Series or FlexArray need to be licensed on every node?

Answer: No. If a cluster contains both V-Series or FlexArray-enabled FAS systems and FAS systems without FlexArray, FlexArray and V-Series licenses are required only on the nodes that require that functionality.

Where can I find information about licenses?

Answer: For information about licenses in clustered Data ONTAP, visit the Field Portal and search for “FAQ Data ONTAP Licensing.” The licensing FAQ includes information about how to obtain evaluation licenses.

6 Installation, Upgrades, and Transition

Where can I find information on making the transition from 7-Mode?

Answer: Information on the 7-Mode to clustered Data ONTAP transition is available in the Clustered Data ONTAP Transition Collateral Binder on the Field Portal.

What is Copy Free Transition (CFT)?

Answer: CFT, introduced with clustered Data ONTAP 8.3.2, allows a 7mode HA pair to transition to clustered Data ONTAP using a procedure similar to a head swap. This process requires a new HA pair running clustered Data ONTAP, and the 7-mode transition tool. Copy free transition allows data to remain in-place on the existing disk shelves, rather than mirroring to a new system.

Can I upgrade from a previous version of clustered Data ONTAP?

Answer: Yes, assuming your hardware supports the target version of clustered ONTAP. Nondisruptive upgrades are supported.

How do I upgrade an entire cluster?

Answer: This is done using a rolling upgrade. You can NDU each HA pair in turn to the new release until all nodes run the release.

While the overall cluster upgrade is in progress, your cluster will operate in "mixed mode" because two versions of Data ONTAP are running in the cluster. This is a supported configuration. However, only two different versions can run in the cluster (not three or more). The cluster is not intended to remain in mixed mode for an extended period of time.

Clustered Data ONTAP 8.2 supports batch cluster upgrades. Batch upgrades allow multiple nodes to be upgraded at the same time, reducing the amount of time it takes to upgrade a cluster.

Clustered Data ONTAP 8.3 supports automated NDU, a three-step process for upgrading a cluster. The logic that drives automated NDU is delivered in clustered ONTAP 8.3, so automated NDU is available to get from 8.3 to a subsequent release, not to get from 8.2 to 8.3.

Certain features introduced in new releases will not be available until all the nodes have been upgraded. For more information, see the Upgrade Advisor and the nondisruptive upgrade technical FAQ.

Can I add more nodes to the cluster?

Answer: Yes. You can add more nodes nondisruptively to a cluster, provided you do not exceed the maximum number of nodes.

Can I remove nodes from the cluster?

Answer: Yes, nodes can be removed from the cluster without affecting client access. You need to move any volumes and LIFs off those nodes before removing the nodes.

The nodes can be rejoined into the cluster later, but they will be seen as new nodes and hence will not maintain any previously defined configuration or data.

Can I expand 32-bit aggregates to 64-bit aggregates?

Answer: Yes. In-place expansion of 32-bit aggregates to 64-bit aggregates is available and supported. Starting with clustered Data ONTAP 8.2.1, expansion is available as an admin-level CLI command that can be run without a PVR or additional disks.

For additional details, see TR-3978, "In-Place Expansion of 32-Bit Aggregates to 64-Bit Overview and Best Practices."

Can two clusters be merged?

Answer: Two clusters cannot be merged into a single cluster such that data is preserved from both clusters. You need to copy the data from one cluster to the other using SnapMirror, SnapVault, or a client-based method. You can then wipe clean the cleared-out cluster and join its hardware to the remaining cluster as new nodes.

What is System Setup?

Answer: System Setup is an off-box GUI setup utility designed to simplify the initial configuration of a clustered Data ONTAP system. It streamlines the initial configuration of the cluster and performs the initial setup to conform to best practices.

System Setup 3.0 is aligned with the clustered ONTAP 8.3 RC1 release. It supports single and multinode switched and switchless configurations for the current FAS platforms, including single node and multi-node configurations, and All-flash FAS. It supports clustered ONTAP 8.3 networking, including IPspaces, broadcast domains, and subnets. It also supports the 4 port cluster interconnect option for FAS8000 series systems.

What is an automated nondisruptive upgrade (NDU)?

Answer: Clustered ONTAP 8.3 adds support for automated, nondisruptive software upgrades using three commands. These commands will bring the clustered ONTAP package (first obtained from support.netapp.com) into the cluster, do a validation of the cluster to see it is prepared for the upgrade, and then do the actual upgrade. Underneath there still will be downloads, takeovers, and givebacks, but the cluster infrastructure will drive the process.

The administrator will be able to see the progress, pause or resume an upgrade, and see the cluster update history.

Automated NDU is available to get from 8.3 to a subsequent release. The code to run the automated upgrades is in 8.3, so a traditional approach will be required to get from 8.2 to 8.3. Plan for version-dependent snapmirror relationships manually before running the automated NDU. Also, do any post-upgrade checks manually, as they are not automated.

7 Networking

What is a LIF?

Answer: A LIF is a logical interface. LIFs are logical resources tied to an SVM. A NAS LIF has an IP address. A SAN LIF has a WWPN. A LIF is hosted by a physical network port on a node.

This separation of the physical port and the LIF allows a NAS LIF to migrate nondisruptively to another physical port on any node within the cluster.

When a NAS LIF migrates, does client access immediately move to the destination network port?

Answer: Yes. When a LIF migrates, clustered Data ONTAP uses a gratuitous ARP to immediately notify every node on the LAN of the LIF's new location.

Can a SAN LIF migrate?

Answer: No. Nondisruptive SAN LIF migration is not required, because host-side ALUA MPIO is used to provide continuous access to a LUN through an alternate LIF when the LUN is moved to another node or a SAN LIF is taken offline.

SAN LIFs must be taken offline before they are moved to another port in the cluster.

How many ports can host a LIF?

Answer: A LIF can exist on one network port, VLAN-tagged interface, or interface group at a time. A single network port, VLAN-tagged interface, or interface group can house multiple LIFs at once. A LIF can migrate to a different network port, VLAN-tagged interface, or interface group if required.

How many LIFs can I create?

Answer: Up to 256 LIFs can be hosted on a node at one time; however, no more than 128 LIFs per node should be used to account for failover.

You should account for failover scenarios when one node might become responsible for another node's LIFs. It is also important to account for cluster interconnect LIFs, node management LIFs, cluster management LIFs, and intercluster LIFs.

How many LIFs can I assign to a port?

Answer: There is no limit to the number of LIFs that can be placed on a single port for NAS; however, the per-node limit of 256 LIFs cannot be exceeded.

Eight SAN LIFs can use a single port in clustered Data ONTAP 8.1. Sixteen SAN LIFs can use a single port in clustered Data ONTAP 8.2. Thirty-two SAN LIFs can use a single port in clustered Data ONTAP 8.2.1.

What types of LIFs exist?

Answer: Five different types of LIFs exist:

- **Cluster management.** Used to manage the cluster. A cluster contains one cluster management LIF.
- **Node management.** Used to manage individual nodes.
- **Cluster.** Used for the cluster interconnect. Each node in the cluster contains two cluster interconnect LIFs.
- **Data.** Used to serve data to clients/hosts. These LIFs are designated either SAN or NAS. At least one is required for each SVM. A best practice is to have at least one LIF per node per SVM for

NAS access, two LIFs per node per SVM for iSCSI access, and two LIFs per fabric per node per SVM for FC access.

- **Intercluster.** Used for mirroring between clusters. These are SVM specific. At least one intercluster LIF is required per node for each SVM that will be replicated between clusters.

Can a data LIF support more than one data protocol?

Answer: Yes, for NFS and CIFS. However, the data protocols supported by the LIF cannot be changed after the LIF is created, so plan accordingly. If your SVM provides both NFS and CIFS access, consider enabling both protocols when you create new LIFs on the SVM.

iSCSI requires dedicated LIFs.

What is the cluster interconnect?

Answer: The cluster interconnect is a dedicated, redundant 10GbE infrastructure. Only the approved switches that have been qualified for use with clustered Data ONTAP can be used. The cluster interconnect infrastructure must not be shared for any other purpose.

In clustered Data ONTAP 8.2 and higher, a two-node cluster can be directly connected without a switch.

For what is the cluster interconnect used?

Answer: The cluster interconnect is used for:

- Indirect access to data on any node from any other node in the cluster
- Synchronizing configurations among nodes in the cluster
- Intracluster mirroring
- Volume movement (DataMotion for Volumes)
- LUN movement (DataMotion for LUNs) in clustered ONTAP 8.3

What protocol is used for the cluster interconnect?

Answer: A proprietary network protocol called SpinNP is used on the cluster interconnect. SpinNP uses the IP protocol.

What is the maximum distance a single cluster can span?

Answer: Any HA pair in a cluster can be no more than 300 meters (cable distance) away from the cluster interconnect switch. Any two HA pairs can therefore be 600 meters apart if each is 300 meters from the cluster interconnect switch.

The two nodes that make up an HA pair are connected to the same shelves and the same HA interconnect, so they are subject to additional range constraints.

Can I directly connect nodes instead of using an interconnect switch?

Answer: Yes. Two-node switchless clusters are supported in clustered Data ONTAP 8.2. Larger clusters require an interconnect switch.

Are redundant 10GbE cluster interconnect connections required?

Answer: Yes, in most configurations. However, switchless 2240 clusters can be configured with a single 10GbE interconnect.

Configuration guidelines should always be followed to achieve the highest degree of failure resistance and cluster interconnect bandwidth.

Is the transition from a switchless to a switched cluster nondisruptive?

Answer: Yes. A nondisruptive procedure to add an interconnect switch is available for two-node switchless clusters. Single-node clusters do not support a nondisruptive conversion to an HA pair.

Are switch upgrades nondisruptive?

Answer: Yes. Redundant switches allow nondisruptive upgrades when the prescribed upgrade procedures are followed.

Is IPv6 supported?

Answer: Yes. IPv6 data LIFs are supported starting with clustered Data ONTAP 8.2. See the “File Access and Protocols Management Guide” for information about data protocol support over IPv6.

Clustered ONTAP 8.3 includes additional IPv6 enhancements, including IPv6 inter-cluster peering, SnapMirror over IPv6, IPv6 support for MetroCluster, USGv6, and most requirements for IPv6 Ready Logo Phase 2, except for Stateless Address Auto-Configuration.

Can customers use their own management switches?

Answer: Yes. Ports dedicated to management can be connected to the customer’s existing management switches, or new management switches can be purchased through NetApp.

However, only the approved switches that have been qualified for use with clustered Data ONTAP can be used as cluster interconnect switches.

Is load balancing among data ports supported?

Answer: Yes. DNS load balancing between data ports is supported. See the “Network Management Guide” for protocol support and additional details.

Is SNMP supported?

Answer: Yes. Clustered Data ONTAP and the cluster interconnect switch both support SNMP. MIBs are available on support.netapp.com.

Where can I get more information about the cluster interconnect?

Answer: You will find more information in the clustered Data ONTAP networking technical FAQ.

How many simultaneous TCP connections are supported?

Answer: The maximum number of TCP connections per node varies by platform and clustered Data ONTAP release. See the “NFS Technical FAQ” on the Field Portal for more information.

What are Broadcast domains?

Answer: Broadcast domains are a new construct introduced with clustered Data ONTAP 8.3. All the ports in a broadcast domain are connected to the same VLAN on the switch, either by virtue of its physical switch port, or by virtue of VLAN tagging. This means that a logical interface that exists on any one of the ports in that broadcast domain should have the same network connectivity on any other port in that broadcast domain.

The use of broadcast domains allows clustered ONTAP 8.3 to configure failover groups for you automatically. It also makes certain LIFs will not failover to ports that aren't connected to the same network. It will help you ensure that the home port, current port, and failover group of an interface should have the same Broadcast Domain, and that all the ports being used in the same broadcast domain have the same MTU value.

What are subnet objects?

Answer: Subnet objects are introduced with clustered Data ONTAP 8.3. A subnet object is an administrator-specified IP address range, gateway, and netmask. A subnet object lives inside of a broadcast domain. The IP address range can also be a list of IP addresses or a list of IP ranges. The subnet object does not have to include every IP address in an actual IP subnet, it can be a subset of those IP addresses, if needed.

Once a subnet object is defined, LIFs can be created using that subnet object. The LIF will be configured using the next available IP address from the subnet object, and it will use the subnet's gateway and netmask automatically.

What happened to port roles in clustered ONTAP 8.3?

Answer: Although most ports will continue to serve the same function, like data ports and cluster ports, ports now belong to a broadcast domain, instead of having a role.

How do I manage routing tables in clustered ONTAP?

Answer: Routing tables in clustered ONTAP 8.2 and below manage network routes using "routing groups". Clustered ONTAP 8.3 introduces the "network route" command hierarchy.

What are IPspaces?

Answer: IPspaces are introduced with clustered ONTAP 8.3. They allow multiple SVMs in the same cluster to have overlapping subnets and IP addresses.

They may be used by service providers that cannot control subnet assignments, enterprise environments where the subnet assignment is not controlled by a single entity, transition of 7-mode vFiler environments that contain overlapping IP addresses, and situations where application stacks need to be cloned without re-configuring IP addresses on the storage.

A single Ipspace can contain multiple SVMs. All the SVMs in a cluster are placed in the "Default" Ipspace after an upgrade from 8.2, since 8.2 does not allow overlapping IP addresses. A cluster can also have up to one Ipspace per SVM. IPspaces support IPv6 as well as IPv4.

8 Storage Protocols

Which storage protocols are supported for flexible volume access?

Answer: Flexible volumes in clustered Data ONTAP can be accessed using a variety of NAS protocols. LUNs housed inside flexible volumes can be accessed using SAN protocols. Protocol support includes:

- CIFS SMB 1.0, 2.0, 2.1, 2.2, 3.0*
- NFSv2,* NFSv3, NFSv4, and NFSv4.1 (with pNFS)
- iSCSI
- FCP
- FCoE

*SMB 3.0 support was added in clustered Data ONTAP 8.2.

*Clustered Data ONTAP 8.2 does not support NFSv2.

Which storage protocols are supported for Infinite Volume access?

Answer: Infinite Volumes can be accessed using NFSv3, NFSv4.1, pNFS, and SMB 1.0.

Do any client-specific requirements exist for pNFS?

Answer: The clients need to be running a version of Linux® that supports pNFS.

Do any host-specific requirements exist to support ALUA or MPIO?

Answer: Specific operating systems are required. See the Interoperability Matrix Tool for details:
<http://support.netapp.com/matrix/>.

Do CIFS clients recognize clustered Data ONTAP junctions?

Answer: Starting with clustered Data ONTAP 8.2, Windows® clients accessing a share with SMB 2.X and above will interpret junctions. In previous versions of clustered Data ONTAP, Windows clients treated clustered Data ONTAP junctions as regular directories.

When a Windows client interprets a clustered Data ONTAP junction, it evaluates the space available under the junction, rather than the space available directly under the CIFS share.

Starting with clustered Data ONTAP 8.2, if SMB 2.X or greater is used, it is no longer a requirement that the CIFS share report enough room for a copy operation to any of the junctions beneath it. If the area beneath a clustered Data ONTAP junction has enough room, the copy operation should succeed.

Are SMB 3.0 continuously available (CA) shares supported?

Answer: Yes. Clustered Data ONTAP 8.2 supports CA shares with SMB 3.0 for Hyper-V. Clustered Data ONTAP 8.2.1 supports CA shares with SMB 3.0 for both Hyper-V and Microsoft SQL Server®.

Is off-board antivirus supported?

Answer: Yes. Starting with clustered Data ONTAP 8.2.1, off-board antivirus is supported for CIFS access.

Multiple antivirus servers can interoperate with the clustered Data ONTAP system at the same time, which allows antivirus scanning to scale along with your access requirements. Clustered Data ONTAP 8.2.1 also supports the use of in-memory access cache on the antivirus servers in order to speed up access.

Clustered Data ONTAP 8.2.1 will support the McAfee, Symantec™, and TrendMicro antivirus vendors. Symantec support will not be immediately available in the clustered Data ONTAP 8.2.1 RC time frame. It is expected in the Q1 CY2014 time frame.

Is access-based enumeration supported?

Answer: Yes. Clustered Data ONTAP 8.2 supports access-based enumeration (ABE) for CIFS.

Is BranchCache supported?

Answer: Yes. Clustered Data ONTAP 8.2 supports BranchCache for CIFS.

Is Volume Shadow Copy Service (VSS) supported?

Answer: Yes. Clustered Data ONTAP 8.2 supports VSS.

Is file access auditing supported?

Answer: Yes, starting with clustered Data ONTAP 8.2. See the “File Access and Protocols Management Guide” for details.

Is FPolicy supported?

Answer: Yes. Clustered Data ONTAP 8.2 supports FPolicy for NFS and CIFS.

Is FTP, SFTP, or HTTP supported?

Answer: No. Clustered Data ONTAP does not support data access over FTP, SFTP, or HTTP.

Is NDMP supported?

Answer: Yes. Clustered Data ONTAP supports NDMP. See the “Data Protection Tape Backup and Recovery Guide” for more information.

Is it possible to initiate an NDMPcopy from a clustered Data ONTAP system?

Answer: No. An NDMP dump should be initiated using an external host utility.

Is WebDAV supported?

Answer: No.

How do I control NAS access from individual clients or groups of clients?

Answer: NAS access from clients can be controlled using export policies. Export policies are conceptually similar to access control specifications in an `/etc/exports` file. They are used to control access to volumes from individual clients or groups of clients. Export policies can be used to control access over both NFS and CIFS. The same export policy can be reused to control access to multiple volumes.

Clustered Data ONTAP 8.2.1 and higher support the use of export policies on qtrees as well as volumes. Qtree-level export policies can be used to control NFSv3 access in 8.2.1, and NFSv4 and NFSv4.1 in 8.3.

Qtree-level export policies are only recommended for 7-Mode transition scenarios in which too many 7-Mode qtrees exist for a qtree-to-volume transition. See the Transition Binder on the Field Portal for more information about making the transition from 7-Mode to clustered Data ONTAP.

In most circumstances, volume-level export policies should not be used to control CIFS access, because file- and share-level security are sufficient and a third layer of security can cause troubleshooting difficulty. In clustered Data ONTAP 8.2, export policies can be enabled and disabled in a CIFS-enabled SVM using the `vserver cifs options modify -is-exportpolicy-enabled` command.

For more information about export policies, see TR-4067, “Clustered Data ONTAP NFS Implementation Guide,” or the “File Access and Protocols Management Guide.”

Why does the client-side “showmount” command return only “/”?

Answer: Since NFS clients need only mount the top of an SVM’s namespace to reach any volume in the namespace (with the appropriate export policies and permissions), clustered ONTAP provides a single top-level export, “/”.

Some applications require showmount support in order to view all the volumes in the namespace. Starting with clustered ONTAP 8.3, an SVM can be configured to respond to a “showmount” command with a list containing all the volumes in the namespace.

Where can I get detailed information on NFS support?

Answer: For detailed information on NFS support, including NFSv4, see the “NFS Technical FAQ” or TR-4067, “Clustered Data ONTAP NFS Implementation Guide.”

What is offline Foreign LUN import (FLI)?

Answer: Offline FLI, available with clustered ONTAP 8.3, allows a NetApp partner or NetApp PS to import a LUN from a 3rd party storage system directly into clustered ONTAP. It will work on any system in the FlexArray support matrix.

FLI does not require a DTA or a special license, but it does require planned downtime. The LUN being imported cannot be in use while it is being imported.

What enhancements exist to FLI in Data ONTAP 8.3.1?

Answer: In Data ONTAP 8.3.1, a LUN can be online, servicing I/O while it is imported into Data ONTAP. This means the outage window can be reduced to minutes or seconds, as the only downtime required is to switch over to target the LUN on Data ONTAP.

Third party storage arrays must be on the FlexArray support matrix. FLI in 8.3.1 does not require NetApp or partner PS, but their engagement is highly recommended, particularly for planning, as planning the actions for the cutover window and ensuring that the target configuration conforms to the interoperability matrix are of critical importance.

What is “online FLI”?

Answer: “Online FLI” is the term previously used to refer to the FLI enhancements in 8.3.1. However, as a cutover is required, the operation is not entirely online, so in order to prevent confusion, this term is not used.

What is selective LUN mapping?

Answer: Clustered ONTAP 8.3 delivers selective LUN mapping. Rather than present indirect paths from every node in a cluster for a LUN, you can optionally create a mapping that will limit paths to the LUN's local HA pair, or to a subset of HA pairs.

This is simpler than using portsets at an igroup level. The selective mapping can be changed at any time. A rescan at the host level is necessary to discover new paths and discard unused ones.

Is Dynamic Access Control (DAC) supported?

Answer: Yes, starting with clustered ONTAP 8.3, DAC is supported.

DAC provides another level of security beyond file and share security. It is a feature that Microsoft introduced in Windows 2012. NetApp is the first storage vendor to support DAC.

With DAC, a user can now be identified not only by their Windows group membership but also who the user is (full time, contractor etc.), what department they might work in, what country they are located in, what device they are connecting from or what data they are accessing.

DAC can also be used to apply classifications to data. These classifications can then be used to apply access controls to limit who is authorized to access that data.

Is Microsoft Management Console (MMC) supported?

Answer: Yes, starting with clustered ONTAP 8.3, the MMC can be used to manage CIFS shares. The MMC can also be used to manage sessions and open files, but in order to do so, the administrator must connect to a LIF on the node that houses the session or file.

Besides DAC and MMC support, what new CIFS features were introduced in clustered ONTAP 8.3?

Answer: Other features introduced with clustered ONTAP 8.3 include support for NetBIOS aliases, various options to control security, home directories, audit enhancement, Fpolicy pass-through read, Unix to Windows file name translation, GPO security policy support, TRIM certification, and enhancements to ODX (Copy-offload).

Where can I find information about SAN interoperability and supported configurations?

Answer: See the Interoperability Matrix Tool (IMT): <http://support.netapp.com/matrix/>.

9 Nondisruptive Operations

What is NDO?

Answer: NDO stands for nondisruptive operations. For clustered Data ONTAP systems, NDO provides continuous data availability to client applications. The main areas of focus for NDO are hardware resiliency, hardware and software lifecycle operations, and hardware and software maintenance operations.

What is vol move?

Answer: Vol move, formally referred to as DataMotion for Volumes, is a feature in clustered Data ONTAP. Vol move allows a flexible volume to be moved nondisruptively from its current aggregate to a different aggregate within the cluster. The source and destination aggregates can be on any node and of any disk type. Starting with clustered Data ONTAP 8.2, a volume cannot be moved from a 64-bit aggregate to a 32-bit aggregate. Sufficient space must exist in the destination aggregate to create a volume of the same size as the source.

Volumes cannot be moved between SVMs. SnapMirror can be used to copy a volume from one SVM to another.

Vol move preserves native data protection relationships, Snapshot copies, and storage efficiency.

Vol move is supported for all SAN and NAS protocols. For more information on DataMotion for Volumes for clustered Data ONTAP systems, refer to TR-3975, “DataMotion for Volumes Overview.”

What is LUN move?

Answer: LUN move, formally known as DataMotion for LUNs, is a new feature in clustered ONTAP 8.3.

LUN move provides LUN-granular nondisruptive operations. When a LUN is moved to another volume inside the cluster, the cutover is near-instantaneous. Writes are processed at the destination, while reads are pulled over from the source as needed. This immediately reduces load on the source system.

When a LUN is moved, efficiencies gained from deduplication will be lost until the deduplication scanner runs again at the destination node. Snapshots of the LUN will remain on the source volume. LUNs cannot be moved between SVMs.

Does vol move require a SnapMirror license?

Answer: No, vol move does not require a SnapMirror license.

Is vol move automatic?

Answer: Vol move is currently not an automatic feature. It can be initiated by the cluster administrator using the controller CLI or the System Manager GUI.

Can I move a volume while a SnapMirror operation is running?

Answer: Yes. Starting with clustered Data ONTAP 8.2, vol move is supported while a SnapMirror operation is in progress.

Can I move volumes from one SVM to another SVM?

Answer: No. Volumes cannot be moved between SVMs. SnapMirror can be used to mirror volumes across SVMs.

Can clones or mirrored volumes be moved using vol move?

Answer: This depends on the type of volume. See Table 3.

Table 2) Volume types supported for vol move.

Volume Type	Vol Move Supported
Flexible volume	Yes
DP mirror	Yes
LS mirror	No
FlexClone parent	Yes, but a temporary volume is left on the origin aggregate for the clones
FlexClone clone	Yes, but the clone loses its space savings on the destination aggregate

What is LIF migrate?

Answer: NAS LIFs can be manually moved nondisruptively to different ports on the same node or to a different node. LIF migrate is a key utility for customer use cases that require the storage administrator to perform lifecycle maintenance operations on a node without disrupting client access. LIF migrate can be used in combination with volume move and HA operations to accomplish common lifecycle and maintenance operations nondisruptively.

LIF migrate allows LIFs to easily be returned to their home ports using the controller CLI or the System Manager GUI.

Can CIFS clients survive a LIF migrate?

Answer: LIF migrate is nondisruptive for SMB 2.0 and later versions of the protocol. For SMB 2.0 and 2.1, the use of durable file handles allows the session to reconnect without any impact on the application. However, durable file handles are exclusive to a single-file open, and the nondisruptive behavior may be affected when multiple users access the same file. For additional information on CIFS NDO best practices and caveats, refer to TR-4100.

Starting with clustered Data ONTAP 8.2, CIFS clients running a supported application, continuously available shares, and SMB 3.0 will experience nondisruptive behavior for LIF migrate, removing any caveats that exist for SMB 2.0 and 2.1. These CIFS clients have a persistent open state that facilitates the nondisruptive behavior. This behavior also extends to failover, giveback, and aggregate relocate.

See TR-4100, “Nondisruptive Operations and SMB File Shares for Clustered Data ONTAP,” for more information.

Can I change the hardware within the cluster nondisruptively using volume move?

Answer: Yes. Clustered Data ONTAP allows data and LIFs to be evacuated from nodes in the cluster through the use of volume move and LIF migrate, allowing the physical hardware to be exchanged. All upgrades can be done nondisruptively provided that both of the following are true:

- Sufficient free space exists in the cluster to host the volumes on the nodes to be removed.
- A sufficient number of nodes exist to host the LIFs on the nodes to be removed (see the networking section, later in this FAQ).

One way to make this happen is to first add the new hardware to the cluster and move the volume from the nodes to be evacuated to the new nodes. You can then delete the evacuated nodes from the cluster.

Can I upgrade the controllers in an HA pair nondisruptively without using volume move?

Answer: Yes. Aggregate relocate (ARL) is available starting in clustered Data ONTAP 8.2. ARL can be used to upgrade cluster nodes nondisruptively without using volume move. Controller upgrades using ARL do not require free space to house volumes while an HA pair is being upgraded. Internal disks in entry-level systems must be evacuated before a controller upgrade with ARL.

Can ARL be used to downgrade controllers?

Answer: No. Controller downgrades with ARL are not supported. Consider using volume move instead.

Can I add and remove storage shelves nondisruptively?

Answer: Yes. New storage shelves can be added to a controller nondisruptively. Starting with clustered Data ONTAP 8.2.1, shelves can be removed from a running system nondisruptively, using nondisruptive shelf removal (NDSR).

NDSR requires a supported version of clustered Data ONTAP, qualified storage shelves, Storage Enclosure Services (SES) version 2, and MPHA or dual-path shelf connections. Shelves can be removed from the top, middle, or bottom of the shelf stack.

All the disks in the shelf must be marked as unowned prior to shelf removal. Before marking the disks as unowned, any nonroot aggregates using the drives in the shelf must be evacuated using volume move and the aggregates must be destroyed. Node root aggregates can be moved to other disks using the `disk replace` command.

For more information, see the “Storage Subsystem Technical FAQ” on the Field Portal.

Can data LIFs fail over automatically to other nodes?

Answer: Yes. If a node fails, its storage is taken over by its HA partner and its LIFs fail over to different ports on other nodes according to the failover rules defined for each LIF.

Can a LIF be configured to be tied to a volume so that if the volume is moved, the LIF also moves?

Answer: No. Volumes and LIFs are components of an SVM and not mapped to each other; therefore, a volume move does not automatically trigger a LIF migrate. This can be done manually if required.

How do failover and giveback times compare to those with 7-Mode?

Answer: For customers running within optimal limits, failover is expected to be from 15 through 45 seconds. Most customers can expect unplanned failover times of 30 seconds or less. Approximately 90% of our customer environments are within optimal limits. However, it is still a best practice to configure client-side timeouts to withstand a 120-second failover.

For customers that push the system limits by running with the maximum number of spindles, utilizing the maximum system capacity, using a large number of FlexVol volumes on a single node, or consistently running with CPU utilization greater than 50%, unplanned failover times may be longer than 45 seconds. One hundred twenty seconds is the recommended client-side timeout setting. This recommendation is consistent with best practices for 7-Mode installations.

Giveback is handled differently in clustered Data ONTAP than it is in 7-Mode. During giveback in clustered Data ONTAP, the root aggregate is sent home first so the node can be assimilated back into the cluster. During this time all data aggregates continue to serve data using the partner node.

After the home node is ready to serve data, each data aggregate is given back to the home node serially. Each data aggregate may take up to 30 seconds to complete giveback.

The total giveback time is the time it takes for the root aggregate to do a giveback, the time it takes to assimilate the node back into the cluster, and the time it takes to give back each aggregate. The entire

giveback operation may be lengthy, but any single aggregate will only be unavailable for up to 30 seconds.

Planned takeover in clustered Data ONTAP 8.2 is similar to giveback in previous versions of clustered Data ONTAP. In these planned takeovers, aggregates are taken over one by one, reducing the total amount of time each aggregate is unavailable. Unplanned takeovers in clustered Data ONTAP 8.2 behave the same way as in previous versions of clustered Data ONTAP.

Recent versions of Data ONTAP (7-Mode and clustered Data ONTAP) are configured by default to perform a giveback after a panic-induced takeover.

10 Data Protection and High Availability

Are Snapshot copies supported?

Answer: Yes. Snapshot copies are supported in clustered Data ONTAP.

How many Snapshot copies can I create?

Answer: You can create 255 per volume.

Is SnapMirror supported?

Answer: Yes. Asynchronous volume SnapMirror is supported both within the cluster as well as between clusters. The replication is at the volume level of granularity and is also known as a data protection (DP) mirror. Qtree SnapMirror is not available for clustered Data ONTAP.

What is SnapMirror for Storage Virtual Machines?

Answer: SnapMirror for Storage Virtual Machines, also known as Storage Virtual Machine Disaster Recovery, or SVM DR is introduced with Data ONTAP 8.3.1. SnapMirror for Storage Virtual Machines is designed to mirror not just the data inside an SVM, but the configuration of the SVM. This includes the SVM's namespace, QoS policies, name mapping configurations, and other aspects of the SVM.

The goal of SnapMirror for SVMs is simplicity. When a replication relationship is configured between SVMs, it eliminates the need to maintain replication relationships for each individual volume inside the SVMs. Change management between the two SVMs is handled automatically. SnapMirror for SVMs uses familiar SnapMirror commands in the CLI, or can be managed via GUI using Workflow Automation (WFA).

SnapMirror for Storage Virtual Machines can be configured in two different modes, depending on the business requirements: Identity Preserve and Identity Discard.

What is SnapMirror for Storage Virtual Machines' "Identity Preserve Mode"?

Answer: When a SnapMirror for Storage Virtual Machines relationship is configured in Identity Preserve mode, network configuration information is replicated to the destination SVM. This includes name service configuration (such as DNS), Network configuration (such as LIFs), NFS configuration (such as export policies), and CIFS configuration (such as the SVM's identity as a CIFS server).

SAN configuration is not replicated in either Identity Preserve or Identity Discard mode. Data ONTAP still replicates the LUN to the destination, but not the igroups or mappings for that LUN.

Identity Preserve mode is designed for situations where the source and destination SVM share the same layer 2 network connectivity to the outside world. It is also used in situations where the DR system resides in a different networks, as long as the same domain is used. In those situations, DR switch-over requires network LIF and routing changes.

What is SnapMirror for Storage Virtual Machines' "Identity Discard Mode"?

Answer: When configured in Identity Discard mode, SnapMirror for Storage Virtual Machines does not replicate network configuration information. The network configuration is configured independently on the destination SVM while the remainder of the mirroring relationship is inherited from the source SVM.

Identity Discard mode is designed for instances where the destination SVM has different network connectivity requirements, and in instances where having a read-only copy online at the destination is required.

As with Identity Preserve mode, LUNs are replicated, but SAN configuration is not.

Can SnapMirror relationships be throttled?

Answer: Yes. SnapMirror relationships can be throttled to a specific transfer rate using the `snapmirror modify -throttle` command.

Can I mirror between clusters using SnapMirror?

Answer: Clustered Data ONTAP 8.2 supports asynchronous volume SnapMirror mirroring between clusters (intercluster) or in the same cluster (intracluster).

Are synchronous SnapMirror relationships supported?

Answer: No. Only asynchronous SnapMirror mirroring is supported. Clustered ONTAP 8.3 provides MetroCluster for synchronous mirroring.

Does the destination of a SnapMirror relationship have to be upgraded before the source?

Answer: Yes. SnapMirror relationships require that the destination ONTAP cluster be upgraded before the source.

Clustered ONTAP 8.3 provides the option of version flexible snapmirror relationships, which either the source or the destination to be upgraded first, while the mirroring relationship continues to be updated.

The default SnapMirror engine will also be available for relationships that require RPOs less than two hours and file counts greater than 10 million.

What are the advantages of version flexible SnapMirror?

Answer: Version flexible SnapMirror allows either the source or the destination of a SnapMirror relationship to be upgraded first. It is available starting with clustered ONTAP 8.3.

It provides other advantages, as well. With an added SnapVault license, SnapMirror version flexible replication can retain asymmetric snapshots on the destination system. That means customers can perform a DR to a backup snapshot instead of a DR to the latest snapshot.

Also, with a SnapVault license, the same secondary copy can maintain snapshots of the primary system that can be used for both SnapMirror and SnapVault. That eliminates the need for two distinct snapshot copies on the secondary system, one for DR and one for Backup.

What do I need to know about version flexible SnapMirror?

Answer: Workload sizing need to be considered before deploying the SnapMirror version flexible replication feature. As a rule of thumb, ensure that RPOs are equal or higher than two hours, and the file count is lower than 10 million.

A re-baseline will be required to convert a version flexible snapmirror to the default relationship style

Version flexible relationships between 8.2 and 8.3 will not be possible, since 8.2 does not support version flexible relationships.

What is SnapMirror network compression?

Answer: Network compression for SnapMirror relationships is available starting with clustered ONTAP 8.3. That means those sites that do not have a WAN acceleration appliance can make the best of their available bandwidth for SnapMirror initialization and updates.

This comes at the cost of increased CPU on the controllers involved.

Can multiple SnapMirror relationships in the same cluster use overlapping IP addresses or subnets?

Answer: Yes, starting with Data ONTAP 8.3.1.

In cDOT 8.3.1, cluster peering relationships can be created from any user-created (non-default) IPspace. This capability allows cluster administrators to create separate networks with overlapping subnets for the purposes of peering on a per-tenant basis.

In other words, in 8.3.1, multiple replication relationships on the same cluster can use the same IP addresses.

The primary customer use case for this feature is backup and DR-as-a-service for service providers, where tenants maintain IP address assignment.

Inter-cluster LIFs can be created in any IPspace and will handle only the peering traffic from SVMs in the same IPspace. At least one inter-cluster LIF must be created per node per IPspace to allow for maximum reachability to SVMs in any IPspace.

Where can I find information about SnapMirror?

Answer: For more information, including the number of cluster peering relationships supported in different versions of clustered Data ONTAP, see the “SnapMirror FAQ for Clustered Data ONTAP,” available on the Field Portal.

Is SnapVault available?

Answer: Yes. SnapVault is available in clustered Data ONTAP 8.2 and above. Intercluster SnapVault is supported. SnapVault relationships between clustered Data ONTAP and 7-Mode Data ONTAP are not supported.

Does SnapVault provide granular restoration of files and LUNs?

Answer: Yes, starting with clustered ONTAP 8.3.

Can I mirror 7G/7-Mode systems to clustered Data ONTAP systems?

Answer: No. Replication from 7G/7-Mode systems to clustered Data ONTAP systems is not supported.

Can I mirror clustered Data ONTAP systems to 7G/7-Mode systems?

Answer: No. Replication from clustered Data ONTAP systems to 7G or 7-Mode systems is not supported.

Can I mirror entire SVMs?

Answer: Yes, starting with clustered Data ONTAP 8.3.1. SnapMirror for Storage Virtual Machines (also called Storage Virtual Machine Disaster Recovery or SVM DR) can be used for this purpose.

Is MetroCluster available?

Answer: Yes, starting with clustered ONTAP 8.3. MetroCluster is a synchronous mirroring solution consisting of two separate, fabric-attached clusters up to 200 km apart. In addition to sitewide DR

capability, it also offers local HA for servicing the installation without switching over to an alternate site. More information can be found in the “MetroCluster Technical FAQ” available on Field Portal.

Are two-node MetroCluster configurations available?

Answer: Yes, starting with Data ONTAP 8.3.1, three different two-node MetroCluster configurations are available.

Stretch MetroCluster is a campus-wide (up to 500 meters), two node synchronous mirroring solution that is available to customers with optical cable infrastructure. It does not require switches or bridges.

Stretch MetroCluster with ATTO bridges is also a two node, campus wide (up to 500m with 2G, 150m with 8G) synchronous mirroring solution. It is available to customers with FC infrastructure. With FC infrastructure, ATTO bridges are required and included.

Two-node Fabric MetroCluster is analogous to what was provided in Data ONTAP 8.3, but with two single-node clusters, one at each site. It allows synchronous mirroring between metropolitan areas, up to 200km apart, with dark fibre operating as an ISL. Bridges and switches are required and included.

Starting with clustered Data ONTAP 8.3.2, Fibre Channel over IP is supported with MetroCluster. This allows inter-switch links using Cisco MDS 9250i switches.

Is stretch MetroCluster available?

Answer: Yes. Two-node stretch MetroCluster is available starting with Data ONTAP 8.3.1

Is SyncMirror supported?

Answer: Local SyncMirror will be supported for aggregate replication within an HA pair in a non MetroCluster configured cluster. It will not require a PVR. The aggregate is written to 2 plexes, owned by the same node. FC-VI is only supported in MetroCluster configurations, therefore local SyncMirror is restricted to HA distances.

Is SnapLock software supported?

Answer: No. SnapLock® software is not currently supported. Customers requiring SnapLock functionality should consider Data ONTAP 8.2 operating in 7-Mode.

Is NDMP supported?

Answer: Yes. A list of backup vendor applications that support clustered Data ONTAP 8.2 is available at <http://www.netapp.com/us/solutions/a-z/backup-to-tape/backup-to-tape-ndmp.html>.

Where can I get more information about data protection?

Answer: Many more details are available in the SnapMirror FAQ for clustered Data ONTAP 8.2.

11 Storage Efficiency

What storage efficiency features are supported?

Answer: Clustered Data ONTAP supports a variety of storage efficiency features, including:

- Deduplication
- Compression (inline compression is available starting with Data ONTAP 8.3.1)
- Thin provisioning
- Efficient Snapshot copies
- RAID-DP
- FlexClone

Where can I get more information about storage efficiency features?

Answer: Additional information about the storage efficiency features available in clustered Data ONTAP is available in [TR-3966](#)

12 Management

How do I manage a clustered Data ONTAP system?

Answer: Clustered Data ONTAP systems can be managed in a variety of ways:

- CLI: SSH to cluster, node, or SVMs
- GUI: OnCommand® System Manager
- GUI: OnCommand Unified Manager
- GUI: OnCommand Performance Manager
- GUI: OnCommand Workflow Automation
- GUI: OnCommand Insight
- API: NetApp management SDK

More information about the OnCommand product family is at netapp.com/OnCommand.

Are APIs available?

Answer: Yes. A software development kit is available. It provides developer access to a bundle of clustered Data ONTAP APIs. The APIs available for clustered Data ONTAP differ greatly from those available for 7-Mode Data ONTAP.

Must I manage each node individually?

Answer: No. One of the major benefits of clustered Data ONTAP is the ability to manage the cluster as a single entity through the cluster management interface. After a node is joined to the cluster, its components can be managed within the cluster context.

Can event management system events be routed to different locations, depending on their severity?

Answer: Yes. Events can be sent to e-mail addresses, SNMP trap hosts, or syslog servers, depending on the type and severity of the event. See the “Managing Event Messages” section of the “System Administration Guide for Cluster Administrators” for details.

Can I measure latency?

Answer: Yes. In clustered Data ONTAP 8.2, latency metrics are available for SVMs, protocols, volumes, and nodes, among others. To find these metrics, browse the `statistics show-periodic` command directory in the CLI.

Can I back up a cluster’s configuration?

Answer: Yes. A cluster’s configuration can be backed up using commands available in the `system configuration backup` command directory in the CLI. Note that these commands will back up a cluster’s configuration, not its data.

Can I use Operations Manager?

Answer: Clustered Data ONTAP is supported in OnCommand Unified Manager 5.0 and 5.1.

OnCommand Unified Manager 5.2 is required for clustered Data ONTAP 8.2 support.

What information is available for storage administrators with 7-Mode experience?

Answer: NetApp provides express guides, courses, technical reports, CLI mapping, and more. For more information, search the Field Portal for “Clustered Data ONTAP for the 7-Mode Administrator Portal.”

13 Other Concepts and More Information

Does a technical report exist that I can give to my customer?

Answer: Yes. A technical report that gives an overview of clustered Data ONTAP is available at <http://www.netapp.com/us/library/technical-reports/tr-3982.html>.

Where can I find more information?

Answer: Clustered Data ONTAP collateral, including the TRs and FAQs referenced in this FAQ, is collected at the Field Portal Data ONTAP 8.2 landing page at <https://fieldportal.netapp.com/viewcontent.asp?html=682>.

What are the differences between 7-Mode and clustered Data ONTAP?

Answer: A detailed [7-Mode to clustered Data ONTAP feature matrix](#) is available on the Field Portal.

Are other clustered Data ONTAP FAQs available?

Answer: Yes. The Clustered Data ONTAP FAQs Binder is available in the Field Portal; it includes the following FAQs.

- Storage QoS Technical FAQ
- NFS Technical FAQ
- SnapVault in Clustered Data ONTAP 8.2 Technical FAQ
- SnapMirror FAQ for Clustered Data ONTAP
- Scalable SAN in Clustered Data ONTAP 8.2 Technical FAQ
- Infinite Volume Technical FAQ
- Data Compression and Deduplication FAQ
- Multiprotocol FAQ
- Nondisruptive Operations Technical FAQ
- Nondisruptive Upgrade Technical FAQ
- Flash Pool Technical FAQ
- Clustered Data ONTAP Platform Mixing Rules Technical FAQ
- 32-Bit to 64-Bit in-Place Aggregate Expansion FAQ
- Clustered Data ONTAP Networking FAQ
- NetApp Flash Cache Technical FAQ

Refer to the [Interoperability Matrix Tool \(IMT\)](#) on the NetApp Support site to validate that the exact product and feature versions described in this document are supported for your specific environment. The NetApp IMT defines the product components and versions that can be used to construct configurations that are supported by NetApp. Specific results depend on each customer's installation in accordance with published specifications.

NetApp provides no representations or warranties regarding the accuracy, reliability, or serviceability of any information or recommendations provided in this publication, or with respect to any results that may be obtained by the use of the information or observance of any recommendations provided herein. The information in this document is distributed AS IS, and the use of this information or the implementation of any recommendations or techniques herein is a customer's responsibility and depends on the customer's ability to evaluate and integrate them into the customer's operational environment. This document and the information contained herein may be used solely in connection with the NetApp products discussed in this document.

Go further, faster®



www.netapp.com

© 2013 NetApp, Inc. All rights reserved. No portions of this document may be reproduced without prior written consent of NetApp, Inc. Specifications are subject to change without notice. NetApp, the NetApp logo, Go further, faster, DataMotion, Data ONTAP, Flash Cache, Flash Pool, FlexCache, FlexClone, FlexPod, FlexVol, FPolicy, MetroCluster, MultiStore, NOW, OnCommand, RAID-DP, SnapLock, SnapMirror, Snapshot, SnapVault, SyncMirror, vFiler, and WAFL are trademarks or registered trademarks of NetApp, Inc. in the United States and/or other countries. Windows, SQL Server, Active Directory, and Hyper-V are registered trademarks of Microsoft Corporation. Linux is a registered trademark of Linus Torvalds. VMware and vSphere are registered trademarks of VMware, Inc. Symantec is a trademark of Symantec Corporation. All other brands or products are trademarks or registered trademarks of their respective holders and should be treated as such.