



Technical FAQ

Scalable SAN

Clustered Data ONTAP 9

Michael Peppers, NetApp

October 2016

Version	Date	Document Version History
Version 1.0	March 2013	Updated for clustered Data ONTAP 8.2
Version 1.1	November 2013	Updated for clustered Data ONTAP 8.2.1
Version 2.0	October 2014	Updated for clustered Data ONTAP 8.3
Version 3.0	July 2015	Updated for clustered Data ONTAP 8.3.1
Version 4.0	February 2016	Updated for clustered Data ONTAP 8.3.2
Version 5.0	June 2016	Updated for ONTAP 9
Version 6.0	October 2016	Updated for ONTAP 9.1

TABLE OF CONTENTS

1 Scalable SAN Concepts and Overview	6
How many nodes can participate in a cluster running clustered Data ONTAP 8.3 if SAN protocols are licensed?	6
Can SAN and NAS protocols be configured in the same cluster?	7
Can SAN and NAS protocols be configured in the same SVM?	7
Can a subset of the total number of cluster nodes be licensed for a SAN protocol?	7
Can quality of service (QoS) policies be used to control SAN traffic?	7
How many LIFs should an SVM have?	7
How are SVM, volume, and LUN limits changing for SAN in clustered Data ONTAP 8.3?	7
Do FAS2200 series and 2500 series storage controllers support block protocols?	8
2 SAN Logical Interfaces	8
2.1 Overview	8
What is a logical interface?	8
What is the difference between direct and indirect I/O for SAN protocols?	8
How do hosts accessing LUN data using a SAN protocol know which paths are direct and which are indirect?	8
Can ALUA assign a path priority or grouping?	8
Does indirect Fibre Channel traffic use the virtual interconnect adapter?	9
Do SAN LIFs use failover groups?	9
What happens to LIFs during a cluster failover or storage failover event?	9
Can SAN and NAS protocols be configured on the same LIF?	9
Can SAN and NAS LIFs be configured to use the same physical port?	9
How many LIFs can be configured for each physical port?	9
Is a direct-attach configuration possible?	9
2.2 Fibre Channel LIFs	9
How do Fibre Channel LIFs present themselves to a fabric?	9
Which Fibre Channel switches can use the NPIV facility?	9
How are the WWPNs of Fibre Channel LIFs determined?	9
Is it necessary to include a target port's physical WWPN in a Fibre Channel zone?	10
Is zoning by port address rather than by WWN supported with clustered Data ONTAP?	10
How can I change the initiator/target status of a Fibre Channel adapter that supports both types of operation?	10

How can I change the FC/CNA personality of a converged network adapter (CNA)?.....	10
Can ports on a UTA have different personalities (FC vs. Ethernet)?	10
What speeds does the UTA2 support?	10
How can I view the data center bridging (DCB) status of a node's unified target adapter (UTA)?	10
2.3 iSCSI LIFs.....	11
Do iSCSI LIFs have a designated partner LIF on a node's HA partner, as in 7-Mode?	11
Is it possible to use interface groups (ifgrps) with SAN LIFs?	11
Why is it recommended that you leverage MPIO instead of ifgrps to provide redundancy and manage multiple physical paths within a single storage controller?	11
Is IPv6 supported with iSCSI LIFs?.....	11
Are iSCSI TLVs supported in clustered Data ONTAP?	11
3 Scalable SAN Logical Design Considerations	11
How many LIFs should an SVM have?	11
How are limits for the maximum number of connected hosts calculated?	12
What are the prerequisites for nondisruptively moving volumes containing LUNs?	12
Is it possible to convert a LUN from an iSCSI to Fibre Channel LUN or vice-versa?	12
4 Host-Specific Configurations	12
Is the host utilities kit required for hosts accessing clustered Data ONTAP using a SAN protocol?	12
Is the Data ONTAP DSM required when using Microsoft Windows 2003?	13
Is the Data ONTAP DSM required when using Windows 2008, Windows 2008 R2, or Windows 2012?	13
What is required to use NetApp SnapDrive with clustered Data ONTAP?	13
What additional operating system support is available in clustered Data ONTAP 8.3?	13
Are advanced storage feature such as thin provisioning and copy offload, supported in clustered Data ONTAP using SAN protocols?	13
What is copy offload?.....	13
What is Microsoft ODX, and how does copy offload support it?.....	13
What is hardware-assisted locking?.....	14
Hardware-assisted locking allows hosts to lock ranges of blocks, or extents, instead of the entire LUN to allow write operations to function in parallel while guaranteeing the consistency of the data in the LUN.	14
What is space reclamation?	14
Space reclamation allows the host to communicate to the storage controller which blocks have been deleted and can be freed from the storage controller's containing volume.	14
Does running space reclamation on a LUN in a deduplicated volume have a performance penalty?	14
What is accelerated zeroing?	14
What is thresholding?	15
What is resource exhaustion reporting?	15
Which vStorage APIs for array integration (VAAI) are supported in clustered Data ONTAP 8.2 and later using SAN protocols?	15
Which Microsoft advanced storage features are supported in clustered Data ONTAP 8.2 and later using SAN protocols?	15
Which Red Hat advanced storage features are supported in clustered Data ONTAP 8.2 and later using SAN protocols?	15
Which Symantec advanced storage features are supported in clustered Data ONTAP 8.2 and later using SAN protocols?	15
Can Cisco UCS mini be directly connected to a clustered Data ONTAP storage array?	15
5 New Features in Clustered Data ONTAP 8.3.....	15
5.1 Foreign LUN Import	16

What is foreign LUN import?	16
Which third-party arrays are supported by FLI?	16
Does FLI require downtime?	16
Does FLI require a cutover period?	16
Does FLI require a FlexArray license?	16
Are LUNs imported with FLI NetApp native LUNs?	16
Does FLI require a Professional Services (PS) engagement?	17
Can I virtualize third-party storage arrays using FLI?	17
What length of downtime window is required for FLI?	17
How many concurrent imports does FLI support?	17
What is the maximum size LUN that FLI can import?	17
When should FLI 7-Mode-to-cDOT be used instead of other transition tools?	17
Can I use foreign LUN Import (FLI) with All Flash FAS (AFF)?	17
5.2 Selective LUN Mapping	17
What is selective LUN mapping, and what does it do?	17
What is meant by a node's locality when referring to SLM?	18
What about mappings imported from 8.2 or previous versions of clustered Data ONTAP?	18
Can I have more than one HA pair reporting permanently?	18
Can I still use SLM and portsets concurrently?	18
Do I have to use SLM for LUNs in clustered Data ONTAP?	18
Would I ever want to have more than the owning HA pair reporting paths?	18
Why is a host rescan required after modifying SLM reporting nodes?	18
Where can I find host rescan steps or an action plan that will remove old paths and discover new ones? ..	18
5.3 DataMotion for LUNs	18
What is DataMotion for LUNs?	18
Does DataMotion for LUNs use the same underlying engine developed to support VMware® Virtual Volumes (VVols)?	19
Can I throttle DataMotion for LUN operations?	19
Can I move LUNs between SVMs?	19
What interactions does a LUN move have with storage efficiency?	19
What interactions does a LUN move operation have with a vol move operation?	19
What happens to the Snapshot copies of a LUN that is moved to another volume?	19
What is meant by "promote early" or "promote late" when using LUN copy or move commands?	20
What is the OnDemand engine?	20
5.4 Other Features and Improvements	20
What are the cluster resiliency enhancements added to clustered Data ONTAP 8.3?	20
What are the copy offload enhancements in clustered Data ONTAP 8.3?	20
6 New Features in Clustered Data ONTAP 9	21
What is igroup ping?	21
What is fast failover?	21
7 New Features in Clustered Data ONTAP 9.1	21
What is the maximum SAN Cluster size in ONTAP 9.1?	21
What application provisioning workflows are available in ONTAP 9.1 SAN?	21
Where do you check source array FLI support?	21
Can copy offload (ODX) be used to move data between clusters?	22

Can ONTAP be configured to respond to iSCSI queries with a FQDN instead of an IP and if so, why would this be useful?.....	22
References.....	22

Table 1) New in clustered Data ONTAP 8.3.

Section	Feature Name	Summary
5.1	Foreign LUN import	LUNs from other storage arrays may be imported to a Data ONTAP® environment using a storage controller's Fibre Channel initiator ports.
5.2	Selective LUN mapping	LUN mappings by default restrict a LUN's paths to the node on which its data resides and its HA partner.
5.3	DataMotion™ for LUNs	LUNs can be moved or copied independently of their container volume. LUNs immediately appear at the destination, and data is copied in the background.

1 Scalable SAN Concepts and Overview

Storage systems running the NetApp® clustered Data ONTAP operating system are referred to as nodes. These nodes are aggregated together into a clustered system. The nodes in the cluster communicate with each other continuously, coordinate cluster activities, and move data transparently from node to node using redundant paths to a dedicated cluster network consisting of two 10-Gigabit Ethernet switches.

Although the basic unit of a cluster is the node, with the exception of the single-node cluster, nodes are added to the cluster as part of a high-availability (HA) pair. As with Data ONTAP operating in 7-Mode, HA pairs enable high availability by communicating with each other over an HA interconnect (separate from the dedicated cluster network) and by maintaining redundant connections to the HA pair's disks. Also, like Data ONTAP 7-Mode, disks are not shared between HA pairs, though shelves may contain disks that belong to either member of an HA pair.

Clustered systems are administered on a per-cluster rather than a per-node basis, and data is served from one or more storage virtual machines, or SVMs. Each SVM is configured to own storage in the form of volumes provisioned from a physical aggregate. SVMs also own logical interfaces (LIFs) assigned to either a physical Ethernet network or a Fibre Channel target port. Logical disks (LUNs) are created inside an SVM's volumes and mapped to hosts in order to provide them with storage space. SVMs are node independent and cluster based: they can use physical resources such as aggregates (by creating FlexVol® volumes) or network ports (in the form of LIFs) anywhere in the cluster.

When an SVM is first created and a block protocol (FC or iSCSI) is enabled, the SVM gets either a Fibre Channel worldwide name (WWN) or an iSCSI qualified name (IQN), respectively. This identifier is used irrespective of which physical node is being addressed by a host, with Data ONTAP enabling the SCSI target ports on all the cluster nodes to work together to present a virtual distributed SCSI target to hosts accessing block storage.

In practice, this means that, no matter with which physical node a host communicates, it is communicating with the same SCSI target. This method of access presents new opportunities for data resiliency and mobility, but also has implications for best practices when accessing data using block protocols in clustered Data ONTAP.

How many nodes can participate in a cluster running clustered Data ONTAP 8.3 if SAN protocols are licensed?

Answer: Refer to Table 2 for supported node counts.

Table 2) Node count limits with SAN protocols.

Clustered Data ONTAP Version	Min. Node Count	Max. Node Count
Data ONTAP 8.1	2	4
Data ONTAP 8.1.1	2	6
Data ONTAP 8.2	1	8
Data ONTAP 8.3	1	8
ONTAP 9	1	8
ONTAP 9.1	1	12

Can SAN and NAS protocols be configured in the same cluster?

Answer: Yes.

Can SAN and NAS protocols be configured in the same SVM?

Answer: Yes.

Can a subset of the total number of cluster nodes be licensed for a SAN protocol?

Answer: No. Protocol licenses are clusterwide.

Can quality of service (QoS) policies be used to control SAN traffic?

Answer: Yes. In clustered Data ONTAP 8.2, QoS can be used to control traffic to a specific LUN, volume, or SVM.

How many LIFs should an SVM have?

Answer: It is recommended that an SVM have 2 target LIFs per node. In order to limit paths across the cluster, SLM can be used on a per-LUN basis (see section 5.1).

How are SVM, volume, and LUN limits changing for SAN in clustered Data ONTAP 8.3?

Answer: Refer to Table 3.

Note that limits are not the same for all controller models and cluster sizes. Refer to the limits section of the [SAN Configuration Guide](#) for additional details.

Table 3) Object limits for Data ONTAP clusters licensing SAN protocols.

Resource	Data ONTAP 8.1.1		Data ONTAP 8.2		Data ONTAP 8.3	
SVMs	64	256	125	750	125	1,000
Volumes	500	2,000	1,000	6,000	1,000	8,000
LUNs	2,048	8,192	8,192	49,152	12,288	98,304
LIFs/port	16		16*		32	

Note: Clustered Data ONTAP 8.2.1 increased the LIFs/port limit from 16 to 32.

Do FAS2200 series and 2500 series storage controllers support block protocols?

Answer: Yes. Single-node FAS2200 series clusters support Fibre Channel protocols when they have an FC mezzanine card installed.

Multi-node FAS2200 series clusters can support iSCSI over their on-board Ethernet interfaces when they have a 10GbE mezzanine card installed for the cluster interconnect network.

The 2520 will support iSCSI only; the other two controllers in the 25xx series will support iSCSI, FCoE, and FCP using onboard UTA2 ports, which can be configured to either 10GbE, or 4, 8, or 16Gb FCP. The FAS25xx series removed the mezzanine card that was present on the FAS22xx series it replaces. For more information, refer to the hardware universe and SAN configuration guide.

Can I use direct attach with tape devices?

Answer: Yes. In general, all supported tape drives are considered supported on all platforms, as long as the required physical interface (SCSI LVD/SE/HVD or FC) is available on the platform. For a full list of supported configurations, refer to the [Interoperability Matrix Tool](#).

2 SAN Logical Interfaces

2.1 Overview

What is a logical interface?

Answer: A logical interface, or LIF, is a logical construct that connects the logical storage system of an SVM to a physical resource such as an Ethernet network port or Fibre Channel target adapter. This allows it to serve data to clients or hosts, accept administrative commands, perform data replication, and access network resources such as name services.

A storage area network (SAN) LIF serves data to hosts using a block protocol such as Fibre Channel, Fibre Channel over Ethernet, or iSCSI.

What is the difference between direct and indirect I/O for SAN protocols?

Answer: Direct I/O occurs when both the target SAN LIF and the LUN data being accessed reside on the same cluster node. Indirect I/O occurs when the SAN LIF and the LUN data being accessed are on different nodes. Indirect I/O has to use the back-end cluster network in order to move blocks from the target LIF to the target LUN because they are on different nodes. This is another reason why each node hosting an SVM should have two or more LIFs for each SVM they host.

How do hosts accessing LUN data using a SAN protocol know which paths are direct and which are indirect?

Answer: They use the Asymmetrical Logical Unit Access (ALUA) extension to the SCSI protocol. Every path to a LIF will identify itself as either active/optimized (direct) or active/non-optimized (indirect), depending on whether or not the LIF and data being accessed occupy the same cluster node. ALUA has also been implemented in Data ONTAP running in 7-Mode and is common in server-class host operating systems.

Can ALUA assign a path priority or grouping?

Answer: No. The ALUA extension only specifies the status of a path as optimized or non-optimized and cannot otherwise prioritize paths. Some host-side multipathing implementations can separate paths into subgroups and assign different path selection priorities.

Does indirect Fibre Channel traffic use the virtual interconnect adapter?

Answer: No. This is the way indirect Fibre Channel traffic is handled when Data ONTAP is running in 7-Mode. In clustered Data ONTAP, the cluster network is used to forward indirect traffic to the appropriate cluster node.

Do SAN LIFs use failover groups?

Answer: No. SAN LIFs do not move from physical port to physical port upon failure the way that NAS LIFs do. Instead, host-based MPIO is relied upon to select another path.

What happens to LIFs during a cluster failover or storage failover event?

Answer: LIFs on a cluster node that is taken over or offline are marked as down by host-side MPIO. The down node's HA partner, after takeover, indicates using ALUA that its paths are direct until giveback has been completed.

Can SAN and NAS protocols be configured on the same LIF?

Answer: No. Because NAS LIFs are expected to migrate from port to port and SAN LIFs are expected to remain in place, a LIF that is assigned the iSCSI protocol cannot have any other protocols assigned.

Can SAN and NAS LIFs be configured to use the same physical port?

Answer: Yes.

How many LIFs can be configured for each physical port?

Answer: Clustered Data ONTAP 8.3 maintains the same 32 LIF per port limit first reached in clustered Data ONTAP 8.2.1.

Is a direct-attach configuration possible?

Answer: No. The failover capability of clustered Data ONTAP requires an active path to a node's HA partner. Furthermore, in the case of Fibre Channel configurations, direct connections do not support the N_Port ID virtualization (NPIV) facility (see section 2.2).

2.2 Fibre Channel LIFs

How do Fibre Channel LIFs present themselves to a fabric?

Answer: Fibre Channel LIFs use the Fibre Channel protocol's NPIV facility in order to allow more than one LIF to log into the fabric from a single physical N_Port.

Which Fibre Channel switches can use the NPIV facility?

Answer: All clustered Data ONTAP configurations listed in the [Interoperability Matrix Tool](#) can support NPIV. In general, all 4Gb-capable switches can support NPIV; for some switches NPIV might need to be enabled before it is operational. See the switch manufacturer's documentation for details.

How are the WWPNs of Fibre Channel LIFs determined?

Answer: Fibre Channel LIFs have two fixed and two cluster-determined sections, in the form of 20:xx:00:a0:98:yy:yy:yy.

- The 20 prefix is meant to signify a virtual fabric port.
- The xx section is incremented starting from 00 as LIFs are created in a cluster. This value can eventually be reused if enough LIFs are created and then deleted.
- The 00:a0:98 section is NetApp's organizationally unique identifier (OUI).

- The `yy:yy:yy` section is a unique cluster identifier.

Is it necessary to include a target port's physical WWPN in a Fibre Channel zone?

Answer: No. Although this is necessary when running Data ONTAP in 7-Mode, in clustered Data ONTAP this WWPN logs into the Fibre Channel fabric but does not present a SCSI target service. It should not be included in any zones. Only an SVM's LIF's WWPN should be included in order for it to be accessed.

Is zoning by port address rather than by WWN supported with clustered Data ONTAP?

Answer: No. Because data is sent and received using an SVM's LIF WWPN instead of the underlying physical port WWPN, only WWN-based zoning allows granular control over which target WWNs are part of a zone.

How can I change the initiator/target status of a Fibre Channel adapter that supports both types of operation?

Answer: You can use the `node hardware unified-connect modify -type` command to change the status of an FC adapter from initiator to target and vice versa. Note that in order to load the appropriate device driver after making a change, a node reboot is required.

How can I change the FC/CNA personality of a converged network adapter (CNA)?

Answer: You can use the `node hardware unified-connect modify -mode` command to change the personality of these adapters from 16Gb FC to 10GbE CNA and vice versa. Note that in order to load the appropriate device driver after making a change, a node reboot is required.

Can ports on a UTA have different personalities (FC vs. Ethernet)?

Answer: UTA cards and onboard ports capable of supporting both FC and Ethernet personalities come in pairs, with both ports sharing a single application integrated circuit, or ASIC. For this reason, both ports must be set to the same personality.

For example, a UTA2 card in slot 3 of a storage controller would be designated as 3a and 3b (if FC) or e3a and e3b (if Ethernet), and any changes to personality must be to both ports.

As another example, a FAS8040-series storage controller has 4 onboard UTA2 ports, designated as 0e, 0f, 0g, and 0h when using the FC personality, and e0e, e0f, e0g, and e0h when using the Ethernet personality. Ports 0e and 0f are one pair that must have the same personality. Ports 0g and 0h are the other pair.

Answer: When configured to use the FC personality, each port in a port pair can be configured as either an initiator or as a target. Keep in mind that an appropriate SFP will need to be inserted that matches the personality: for FC a 16Gb or 8Gb SFP would need to be inserted, The Ethernet personality requires a 10GbE SFP.

What speeds does the UTA2 support?

Answer: When using the Ethernet personality, speed is set at 10GB. When set to the Fibre Channel personality, the UTA2 will autonegotiate to 4, 8, or 16 Gb Fibre Channel, as appropriate.

How can I view the data center bridging (DCB) status of a node's unified target adapter (UTA)?

Answer: The DCB status of a node's UTA is visible at the node shell. Use the `run -node <nodename>` command to enter the node shell and then the `dcb show` command to view DCB-related information.

2.3 iSCSI LIFs

Do iSCSI LIFs have a designated partner LIF on a node's HA partner, as in 7-Mode?

Answer: No. In clustered Data ONTAP, iSCSI and Fibre Channel share the same pathing model. Hosts accessing LUN data using the iSCSI protocol must have active sessions to both a node and its HA partner at the same time.

Is it possible to use interface groups (ifgrps) with SAN LIFs?

Answer: LIFs that have the iSCSI protocol can be created using an interface group for their physical port.

Why is it recommended that you leverage MPIO instead of ifgrps to provide redundancy and manage multiple physical paths within a single storage controller?

Answer: Unlike file protocol LIFs, block protocol LIFs do not fail over in the event of a takeover or giveback. Therefore, the only benefit to ifgrps in an iSCSI environment is to allow port aggregation/trunking. Because clustered Data ONTAP uses a FCP style distributed target model, it really isn't necessary to aggregate ports for redundancy, if you follow LIF placement guidelines you can allow host based MPIO to manage paths and load balancing.

Is IPv6 supported with iSCSI LIFs?

Answer: Yes. In clustered Data ONTAP 8.2 and later, IPv6 is supported with iSCSI.

Are iSCSI TLVs supported in clustered Data ONTAP?

Answer: Yes, with an X1143 adapter in a CNA personality mode. If iSCSI TLVs have been enabled and configured on the connected switch, the DCB status of the adapter will identify bandwidth as being reserved for the iSCSI application. To view the adapter's DCB status, enter the node shell using the `run -node <nodename>` command and then use the `dcb show` command.

```
NETAPP-CLUS::> node run -node NETAPP-CLUS-01
NETAPP-CLUS-01> dcb show
```

Interface	PGID	Priority	Application	Bandwidth
-----	----	-----	-----	-----
e3a				
	0	0 1 2 5 6 7	unassigned	20%
	1	3	FCoE	50%
	2	4	iSCSI	30%
e3b				
	0	0 1 2 5 6 7	unassigned	20%
	1	3	FCoE	50%
	2	4	iSCSI	30%

3 Scalable SAN Logical Design Considerations

How many LIFs should an SVM have?

Answer: In order to serve data during a cluster failover or a storage failover event, a host **must** be connected to both a node and its HA partner.

For overall cluster design purposes, at a minimum, every SVM should have LIFs for every network or fabric to which each node is connected. For example, in a 4-node cluster connected to two separate Fibre Channel fabrics, if each node has a target adapter on each fabric, both of those target adapters should have an associated LIF, for a total of 8 LIFs in that SVM.

However, there is an upper limit to the total number of paths a given host operating system can easily support. In larger clusters, limiting the number of paths will improve path manageability. This can be accomplished by configuring an SVM to only contain LIFs and volumes on a subset of the total number of nodes, or by using portsets.

How are limits for the maximum number of connected hosts calculated?

Answer: Prior to Data ONTAP 8.2.1, this limit was calculated assuming that every node in the cluster had a connection to every host accessing cluster resources using a block protocol, and that indirect paths also counted toward this limit. For instance, if the limit for a 4-node cluster using Fibre Channel is 256 hosts, this assumes that that host is connected over 2 redundant fabrics to each of the cluster nodes.

In Data ONTAP 8.2.1 and later, this limit refers to the sum of Fibre Channel I_T (initiator-target) nexuses (I_T_N) or iSCSI sessions, whether they refer to a direct or indirect path, and it no longer assumes that a host is necessarily connected to every node in a SAN-licensed cluster. There are a maximum of 256 IP LIFs per node and a maximum of 512 FCP LIFs per node (keep in mind that best practice is to use two LIFs per node for redundancy). Additionally, there are limits to the number of FC ITNs; 512 per physical port, and 8,192 FC ITNs per node. For iSCSI there is a maximum of 8,192 sessions per node. Finally, both FC and iSCSI share a 32 LIF per port limit, noted in Table 3. Many of the limits are derived during extensive prerelease QA and scalability testing.

The current count of nexuses and sessions can be retrieved using the `fc initiator show` and `iscsi session show` commands for their respective protocols.

What are the prerequisites for nondisruptively moving volumes containing LUNs?

Answer: Even if the destination node does not have a LIF that can provide a direct path after the volume, access to the LUN will not be disrupted. However, the direct path is always the most efficient path, and so care should be taken that the destination node has an available path or paths connecting the host accessing a volume's LUN or LUNs.

Is it possible to convert a LUN from an iSCSI to Fibre Channel LUN or vice-versa?

Answer: LUNs and LUN mappings on a storage system running Data ONTAP are protocol-independent. Instead, igroups are designated as either `fc`, `iscsi`, or `mixed`, and that determines whether FC WWPNS, iSCSI IQNs, or both are permitted to be added to it.

4 Host-Specific Configurations

Is the host utilities kit required for hosts accessing clustered Data ONTAP using a SAN protocol?

Answer: The host utilities kit is required to set host-specific values for optimal operation and continued operation during a cluster or storage failover event as well as to provide the standard set of utilities. Although it might be possible to manually make the settings/registry changes the host utilities kits would do for you, some changes might be difficult to perform manually, and the likelihood that one or more of the settings might be misconfigured is high. In the event of a problem, it is much easier to know that the host utilities were installed and have correctly made the required settings changes than it is to verify each was performed manually. Finally, the host utilities will install a common and useful set of utilities that can be used for both reporting and troubleshooting.

There are three OS exceptions to the preceding requirement:

- There is no ESX host utilities kit for ESXi™.
- There is a requirement to install the host utilities kit for Linux®. The host utilities kit will not change any setting; however, it will install a useful set of utilities.

- There is no need to install the host utilities kit on Windows Server® instances that have the Data ONTAP DSM installed. Because the Data ONTAP DSM makes all of the required setting and registry changes during installation, the host utilities kit is redundant.

Is the Data ONTAP DSM required when using Microsoft Windows 2003?

Answer: Yes. Microsoft® Windows® 2003 does not have a native MPIO implementation that supports ALUA.

Windows 2003 is not supported with clustered Data ONTAP 8.3.

Is the Data ONTAP DSM required when using Windows 2008, Windows 2008 R2, or Windows 2012?

Answer: No. The native MPIO solution (Microsoft Device-Specific Module, or MSDSM) is also supported. Installing the MPIO feature is required, even if using the iSCSI protocol. However, the Data ONTAP DSM sets host-specific timeout values, permits fine-grained path selection, and contains a Server Manager snap-in management pane that shows more detail about the attached NetApp storage than does the Windows native MPIO solution.

What is required to use NetApp SnapDrive with clustered Data ONTAP?

Answer: SnapDrive® technology is scoped to access and administer storage at the SVM rather than at the cluster level. Therefore, SnapDrive requires that SVMs be accessible using a LIF configured to allow management traffic and that there is an unlocked SVM-scoped administrative account available (`vsadmin` by default). Note that the SVM Setup Wizard in NetApp OnCommand® System Manager automatically guides administrators through this process.

What additional operating system support is available in clustered Data ONTAP 8.3?

Answer: Clustered Data ONTAP 8.3 supports a wide variety of operating systems. For detailed supportability information, refer to the [Interoperability Matrix Tool](#).

Are advanced storage features such as thin provisioning and copy offload, supported in clustered Data ONTAP using SAN protocols?

Answer: Yes. However, any given advanced storage features must be supported both by the host operating system and by clustered Data ONTAP to be available.

Clustered Data ONTAP will advertise a LUN as being capable of supporting advanced features and will accept SCSI commands if a LUN has the Space Allocation setting enabled before the LUN is discovered by a SCSI bus scan on the host operating system.

What is copy offload?

Answer: Copy offload is a feature that increases the efficiency of copy operations.

Traditional copy operations require the host to read data from the storage system and write the data back to the storage system.

If copy offload is supported, a host may initiate a copy operation and allow Data ONTAP to copy the data from one place to another without transferring the data between the host and the storage over the network. Copy offload may also leverage FlexClone® volumes to create a read/write copy very quickly and with almost no additional storage consumption.

What is Microsoft ODX, and how does copy offload support it?

Answer: Clustered Data ONTAP supports Microsoft Offloaded Data Transfer (ODX), which uses tokenized operations to move data locally on storage devices. ODX is a copy offload procedure that removes the server initiating the copy from the copy path. Using ODX, a Windows Server instance

initiates a LUN or CIFS copy from either a CIFS share or a LUN and to either another CIFS share or LUN; this triggers a direct copy operation within or between volumes and controllers. This is different than a standard copy operation in which the initiating device is part of the critical copy path. By allowing for a storage-to-storage copy (storage source to storage target), ODX is able to reduce network bandwidth usage and remove latencies associated with the three-way storage source to initiating host to storage target traditional copy.

The clustered Data ONTAP 8.3 improvement for ODX is direct copy, which will allow for the potential to skip using the point-in-time (PiT) copy. This depends on how the file being copied was opened and if the source of the copy offload is a CIFS share. Note that a PiT file will still be created even though direct copy does not need it. This is done to make the token manager state machine simple. If all the copies done from a volume are direct copies, then the PiT file will not be used and will not consume any space on the volume. ODX is available whether or not the LUN has space allocation enabled.

What is hardware-assisted locking?

Answer: Hardware-assisted locking is a feature that increases efficiency in environments with shared access to one or more LUNs, such as a host cluster.

When writes occur in a shared LUN environment, a mechanism must be in place so that two or more writes cannot occur simultaneously to the same blocks on the LUN to prevent the corruption of data. Traditionally, when a host performs a write to a LUN, the entire LUN is reserved, or locked, while the write occurs, preventing other hosts from issuing writes to the entire LUN until the reservation is released.

Hardware-assisted locking allows hosts to lock ranges of blocks, or extents, instead of the entire LUN to allow write operations to function in parallel while guaranteeing the consistency of the data in the LUN.

What is space reclamation?

Answer: Space reclamation is a feature that increases the efficiency of thinly provisioned environments.

A LUN contains a host file system, and the host maintains a map of the file system that records which blocks contain active file system data and which blocks contain deleted data that can be overwritten. When blocks are deleted in the host file system, the block map is updated to mark the blocks as available to receive new writes. However, with traditional SCSI commands, the storage system cannot distinguish between host file system deletions and any other write operation.

Space reclamation allows the host to communicate to the storage controller which blocks have been deleted and can be freed from the storage controller's containing volume.

Does running space reclamation on a LUN in a deduplicated volume have a performance penalty?

Answer: Yes. In order to free the blocks the host says it is no longer using the clustered Data ONTAP has to un-deduplicate the blocks that are being zeroed, that process does require some processor intensive operations to complete.

What is accelerated zeroing?

Answer: Accelerated zeroing is a feature that increases the efficiency of writing large ranges of zeroes to the LUN, such as when provisioning new eager-zeroed virtual disks.

Traditional zeroing required the host to send a series of writes to the storage system with a pattern of zeroes in the SCSI payload. As a result, if the host provisioned a new 10GbE eager-zeroed virtual disk, 10GbE of writes would be sent across the network.

Accelerated zeroing allows a host to instruct the storage to write zeroes to a range of blocks, which vastly reduces the amount of traffic that must be sent across the network.

What is thresholding?

Answer: Thresholding increases the stability of thinly provisioned environments. This feature allows the storage system to send an alert to the connected host when the available space within the LUN's containing volume is below a predefined threshold. This allows the host operating system to log the error in the system log and alert the user.

What is resource exhaustion reporting?

Answer: Resource exhaustion reporting increases the stability of thinly provisioned environments.

In a traditional thinly provisioned storage infrastructure, if a thinly provisioned LUN ran out of space for new writes, the LUN would become read only, and new writes from the host would fail, even though the host file system reports the presence of free storage space. This often results in a host going offline completely.

With resource exhaustion reporting, the storage system warns the host that new writes cannot occur, which allows the host operating system to handle the out-of-space condition gracefully.

Which vStorage APIs for array integration (VAAI) are supported in clustered Data ONTAP 8.2 and later using SAN protocols?

Answer: VAAI v1 and v2 are supported starting in clustered Data ONTAP 8.2. Supported features include hardware-assisted locking, inter- and intravolume copy offload, accelerated zeroing, space reclamation, and resource exhaustion reporting.

Which Microsoft advanced storage features are supported in clustered Data ONTAP 8.2 and later using SAN protocols?

Answer: Supported features include thin provisioning, space reclamation with and without SnapDrive for Windows, resource exhaustion reporting, thresholding, and copy offload. Copy offload is supported for copies within the same datastore, across datastores within the same node, across datastores between nodes, and between CIFS and SAN protocols.

Which Red Hat advanced storage features are supported in clustered Data ONTAP 8.2 and later using SAN protocols?

Answer: Supported features include thin provisioning, space reclamation, and resource exhaustion reporting.

Which Symantec advanced storage features are supported in clustered Data ONTAP 8.2 and later using SAN protocols?

Answer: Symantec™ Storage Foundation environments are not supported, except by approved product variance request (PVR). With a PVR, Symantec Storage Foundation environments will only be supported on two-node clusters. This limitation does not change in 8.3.

Can Cisco UCS mini be directly connected to a clustered Data ONTAP storage array?

Answer: Yes, if the UCS is set to FC-switching mode. Check the IMT to verify that you are on a support version of UCS firmware. Keep in mind that directly connecting forgoes the redundancy you would get connecting to a switch fabric.

5 New Features in Clustered Data ONTAP 8.3

This section answers questions about SAN-specific functionality introduced with the release of clustered Data ONTAP 8.3.

5.1 Foreign LUN Import

What is foreign LUN import?

Answer: Foreign LUN Import (FLI) offline is a feature added to the data migration tool portfolio included with the clustered Data ONTAP (cDOT) storage operating system, beginning with version 8.3. Clustered Data ONTAP 8.3.1 includes enhancements to FLI, and now incorporates both online and 7-Mode to clustered Data ONTAP transition (7-Mode-to-cDOT) workflows.

FLI leverages FlexArray technology to allow a storage controller running clustered DATA ONTAP 8.3 (or later) to act as an initiator and access data hosted on third-party arrays. After a LUN that appears in the storage controller's disk list is marked as "foreign," it can be used as a template to create a destination LUN (with alignment calculated appropriately based on the LUN's data) in a clustered Data ONTAP volume and then be copied block for block into that LUN.

Table 4) New in clustered Data ONTAP 8.3.

Feature	Disruption Window	Type of Workflow	Available in
FLI offline	Hours	Migration	8.3
FLI online	Minutes	Migration	8.3.1
FLI 7-Mode-to-cDOT	Minutes	Transition	8.3.1

Which third-party arrays are supported by FLI?

Answer: Use the [FlexArray portion of the Interoperability Support Matrix](#) to determine which arrays can be used as a source for FLI. If an array isn't listed, FLI-only support may be provided on the basis of a Product Variance Request (PVR). There is no need to check foreign array support for FLI 7-Mode-to-cDOT: no NetApp 7-Mode arrays will be listed, because none are on the FlexArray support matrix. However, all Data ONTAP 7-Mode arrays with any version of Data ONTAP 7-Mode that support FCP are supported for FLI 7-Mode-to-cDOT.

Does FLI require downtime?

Answer: Yes. For FLI offline, the source LUN must be masked only to the clustered Data ONTAP node's initiator, and the destination LUN must be offline while data is copied. Hosts consuming the LUN can't access it for the duration of the import and optional verification. With clustered Data ONTAP 8.3.1 two new FLI workflows are available; FLI online and FLI 7-Mode-to-cDOT, both of these FLI workflows require a small disruption window after the LUN import relationship is created where the host(s) consuming the LUN(s) need to do a LUN rescan to point them to the new target. This disruption window is minutes (duration of the rescan). Once the host sees and mounts the new LUN, LUN access is restored and there are no further disruptions required, with the exception of an optional LUN verification after the import completes.

Does FLI require a cutover period?

Answer: Yes. Hosts accessing a LUN that has been imported to clustered Data ONTAP will need to discover the LUN at its new target WWN (if using Fibre Channel) or target portal (if using iSCSI).

Does FLI require a FlexArray license?

Answer: No. FLI functionality is included in clustered Data ONTAP.

Are LUNs imported with FLI NetApp native LUNs?

Answer: Yes.

Does FLI require a Professional Services (PS) engagement?

Answer: It depends. In clustered Data ONTAP 8.3, FLI requires a PS or partner PS engagement. The requirement for a PS or partner PS engagement is dropped with clustered Data ONTAP 8.3.1; however, it is still recommended that PS or partner PS be engaged to help scope, plan, and train in support of any customer-performed FLI.

Can I virtualize third-party storage arrays using FLI?

Answer: No, FLI is meant for import purposes only. FLI leverages FlexArray functionality to access data on other Fibre-Channel capable arrays, but in doing so copies the data to a LUN inside a Data ONTAP volume, rather than “virtualizing” it. In order to continuously host Data ONTAP data on another array, a FlexArray license is required, and the LUNs dedicated to storing data for a storage controller running Data ONTAP must be created and designated for only that purpose.

What length of downtime window is required for FLI?

Answer: This will vary due to a number of variables including the LUN being imported and the speed that writes can be processed on the target node. A test LUN imports can help characterize the performance to be expected. In QA testing, speeds of 1.5TB/hour were observed; however, for planning purposes, we suggest using 500GB/hour per node.

How many concurrent imports does FLI support?

Answer: FLI supports up to 32 concurrent imports per node. However, for best performance, NetApp recommends no more than 16 concurrent imports per node.

What is the maximum size LUN that FLI can import?

Answer: 6TB is the maximum size. This limitation is based on maximum size drive limits in clustered Data ONTAP. Currently the largest drive clustered Data ONTAP supports is 6TB. If you attempt to access a LUN or drive larger than 6TB in clustered Data ONTAP, the LUN will be marked as “broken” and its label and designation will be immutable.

Note: ONTAP 8.3.2 increased the max disk size supported on FAS from 6 to 8TB, however, FlexArray and therefore FLI didn't match this increase. They continue to enforce a 6TB maximum size.

When should FLI 7-Mode-to-cDOT be used instead of other transition tools?

Answer: NetApp recommends using the FLI 7-Mode-to-cDOT transition workflow instead of the more common 7-Mode Transition Tool (7MTT) when the LUN(s) being transitioned are also hosted on a 32-bit aggregate and/or the LUN(s) are misaligned.

Can I use foreign LUN Import (FLI) with All Flash FAS (AFF)?

Clustered Data ONTAP 8.3.2 includes AFF support for FLI via PVR. Once the PVR has been approved you will be provided with a diag level command that will enable FLI on an AFF system.

5.2 Selective LUN Mapping

What is selective LUN mapping, and what does it do?

Answer: Selective LUN mapping (SLM) is a new default mapping behavior for LUNs created in clustered Data ONTAP 8.3. LUNs created in 8.3 will be mapped by default to the hosting HA pair, so target ports on other nodes in the cluster will not present the LUN. In 8.2, LUNs were mapped to the entire cluster, and target ports could be limited using portsets. Think of SLM as a node-level mapping and portsets as a port-level mapping.

What is meant by a node's locality when referring to SLM?

Answer: A LUN's "local" nodes are the node where its data is kept and its HA partner node.

What about mappings imported from 8.2 or previous versions of clustered Data ONTAP?

Answer: The SLM feature is available but not enabled for LUNs first mapped with clustered Data ONTAP 8.2, so a LUN's previous configuration, with all nodes reporting the LUN as present, is preserved after an upgrade.

Can I have more than one HA pair reporting permanently?

Answer: Yes. Note that LUN mappings must always include both nodes in an HA pair.

Do I still need to use portsets?

Answer: Portsets are still available and have the same functionality as they did in previous versions of clustered Data ONTAP.

Can I still use SLM and portsets concurrently?

Answer: Yes. The LUN will be mapped to the nodes (node-level mapping). And the portset can limit the target ports on each mapped node (port-level mapping). This is useful for nodes that have several target LIFs or environments where an application-to-port association is desired.

Do I have to use SLM for LUNs in clustered Data ONTAP?

Answer: Every LUN mapping has an associated list of reporting nodes. The list would usually be limited to the node owning the LUN and its ha-partner, but could include more than just that ha-pair. SLM will be very useful in reducing the number of paths reported that need to be managed by your hosts.

Would I ever want to have more than the owning HA pair reporting paths?

Answer: Yes. There might be circumstances such as frequent data mobility operations between HA pairs where managing changes to reporting nodes would result in too-frequent configuration changes. Or in organizations where storage and host administration responsibilities do not overlap, the coordination involved to do path management using SLM might be too time consuming.

Why is a host rescan required after modifying SLM reporting nodes?

Answer: A rescan is required for the host to pick up changes and determine which controllers are now advertising paths to a given LUN.

Where can I find host rescan steps or an action plan that will remove old paths and discover new ones?

Answer: There is a customer-accessible [KB article](#) with detailed steps and output for every supported host OS. These steps have been validated by SAN interoperability engineering.

5.3 DataMotion for LUNs

What is DataMotion for LUNs?

Answer: DataMotion for LUNs allows individual LUNs to be moved and copied nondisruptively between volumes in a cluster in such a way as to permit hosts to access the LUN on its destination immediately while data is transferred over the cluster interconnect network in the background.

DataMotion for LUNs is similar to DataMotion for volumes in that the operation takes place nondisruptively in the background with ALUA changing path status in order to reflect a LUN's new

location. Unlike DataMotion for volumes, which initiates a cutover to the destination only when all data has been transferred to the destination, however, DataMotion for LUNs performs the cutover at once.

There are two operations associated with DataMotion for LUNs: move and copy:

- Move operations can only occur within the same SVM, though the source and destination volumes can be on any node within the cluster. If necessary, the SVM signals a path state change using ALUA as soon as the destination LUN is prepared to accept data. The move is nondisruptive to hosts accessing the LUN.
- Copy operations can occur between any two volumes within an SVM. LUN copy can also specify a different SVM's volume as a destination if the operation is performed by a cluster administrator. LUN copy is nondisruptive insofar as the LUN is available for access at the destination as soon as it has been created.

Can I use DataMotion for LUNs to move or copy a LUN to a volume where other LUNs already exist?

Answer: Yes. A destination can have other LUNs.

Does DataMotion for LUNs use the same underlying engine developed to support VMware® Virtual Volumes (VVols)?

Answer: Yes. However, DataMotion for LUNs commands such as `lun move` and `lun copy` should not be used on VVol endpoints.

Can I throttle DataMotion for LUN operations?

Answer: Yes, the throughput (that is, MB/s, GB/s) can be throttled for each individual LUN move/copy operation. The throttle can be set to another value while the operation is in flight.

Can I move LUNs between SVMs?

Answer: A LUN move is not possible between SVMs. A LUN mapping must be created on the destination SVM before data can be accessed. Additionally, other SVMs will have a different Fibre Channel target WWN or iSCSI target portal identifier and will be seen by the host as an independent target.

A LUN copy between SVMs is permitted if performed by a cluster administrator.

What interactions does a LUN move have with storage efficiency?

Answer: LUN copy operations with multiple source and destination LUNs are efficient in transit, because deduplicated blocks shared by more than one source LUN will be read only once. However, after LUNs arrive at the destination, they will not be storage efficient on disk, and deduplication or compression must be run to recoup space in the destination volume. Destination volumes that have storage efficiency enabled will log incoming blocks as needing a deduplication or compression scan during the next background run.

What interactions does a LUN move operation have with a vol move operation?

Answer: Vol move is not allowed on the destination volume while a LUN copy or move operation is active. However, the source volume can be relocated by vol move during a LUN copy or move operation, but it will be paused during the vol move cutover.

What happens to the Snapshot copies of a LUN that is moved to another volume?

Answer: Sub-volume moves or copies will not copy volume-level data such as Snapshot® copies or deduplication and compression. After the move or copy is completed, compression and deduplication scans can be started to regain any storage efficiencies lost during the move or copy.

Until new Snapshot copies have been initiated on the destination volume, it is a best practice to retain existing Snapshot copies on the source volume for disaster recovery and backup purposes. It is also possible to use LUNs in the source volume's Snapshot copies as LUN copy sources. If multiple copies of a LUN are all part of the same LUN copy operation, they will be transported efficiently to the destination and after a deduplication scan will share blocks in the same way that SIS-based LUN clones do.

What is meant by “promote early” or “promote late” when using LUN copy or move commands?

Answer: When moving or copying a LUN, the notion of early or late promotion has to do with when the cluster advertises the new LUN. In the case of LUN move, the default is “promote early.” This means that the new LUN is advertised to hosts as soon as the destination LUN is created. All reads and writes are routed to the LUN's new location while a background operation copies existing LUN data from the source to the destination. Snapshot operations will be blocked until the new LUN is fully framed and the move/copy reaches “moving data” status.

The alternative is late promotion (default for LUN copy), in which the cluster promotes the new LUN location after the LUN has been fully framed. Fully framed means that the new LUN exists as a fully functional LUN, but its data is missing/hasn't arrived yet. A consequence of being fully framed is that Snapshot copies are no longer blocked.

When moving or copying a LUN, the promoting “late” or early is controlled by the command line LUN move/LUN copy command using the argument `-promote-late true` for LUN move and `-promote-early true` for LUN copy. If not specified, then the default for the operation will be taken:

Table 5) Data Mobility defaults.

Operation	Default	Default Argument
LUN copy	promote late	-promote-early false
LUN move	promote early	-promote-late false

What is the OnDemand engine?

Answer: The OnDemand engine enables actions, such as LUN move, to appear to complete instantly. This allows traffic to begin on the newly created destination and allows ALUA to signal path changes, while clustered Data ONTAP transfers the actual data in the background.

5.4 Other Features and Improvements

What are the cluster resiliency enhancements added to clustered Data ONTAP 8.3?

Answer: Clustered Data ONTAP 8.3 includes cluster resiliency enhancements that harden controllers to better withstand losses of back-end network connectivity or networking errors by extending the outage window before a cluster or node in a cluster is declared “out of quorum.” This will improve resiliency significantly and will result in fewer cases in which a node or cluster goes out of quorum, requiring a reboot (usually a takeover/giveback) to regain quorum.

What are the copy offload enhancements in clustered Data ONTAP 8.3?

Answer: Copy offload enhancements introduce support for Windows ODX. [Please see the relevant answer above.](#)

6 New Features in Clustered Data ONTAP 9

What is igroup ping?

Answer: igroup ping is a ONTAP 9 enhancement that allows a storage administrator to check to make sure that igroup members can be verified on a protocol appropriate method. FC igroups use FC ping, and iSCSI igroups use ICMP echo request/replies.

What is fast failover?

Answer: Fast Failover is an ONTAP 9 enhancement that applies to SAN-Optimized All Flash FAS controllers that are configured using prescribed SAN configurations. AFF's configured this way can expect to failover much more rapidly then previous non-optimized models. Unplanned failovers will complete in between 2-15 seconds. Planned failovers complete in less then 2-10 seconds. These times are measured from the surviving controller and will normally be reflected in the IO resume times seen on impacted hosts. Most hosts will resume IO within the same times (unplanned 2-15, planned 2-10) however, some host stacks, especially those that aren't properly tuned could take longer. See the table below for several failover time comparisons.

Platform	Planned Takeover	Unplanned Takeover
All Flash FAS running prescriptive SAN configurations	2-10 seconds	2-15 seconds
All Flash FAS	15 seconds	30 seconds
FAS with Flash Pool™ or SSD aggregates	30 seconds	60 seconds
MetroCluster™	60 seconds	120 seconds

7 New Features in Clustered Data ONTAP 9.1

What is the maximum SAN Cluster size in ONTAP 9.1?

Answer: ONTAP 9.1 increased the maximum cluster size for 8 to 12 nodes. The current cluster limit of 98,304 LUNs remains, which means that the maximum number of LUNs per node in a 12-node cluster is 8192.

What application provisioning workflows are available in ONTAP 9.1 SAN?

Answer: ONTAP 9.1 application configuring workflows were added so that now SAN application provisioning workflows are are:

- SAN Oracle Single
- SAN Oracle RAC
- SAN SQL Server
- SAN Virtual Server Infrastructure
- SAN Virtual Desktop
- SAN SAP HANA

Where do you check source array FLI support?

Answer: There is a new FLI IMT that should be used to check whether a source array is supported. You can find the FLI IMT by searching for FLI in the IMT and following the returned links for online or offline.

Can copy offload (ODX) be used to move data between clusters?

Answer: Yes, beginning in ONTAP 9 you can use ODX to transfer LUNs between clusters, SAN only, this isn't supported with SMB.

Can ONTAP be configured to respond to iSCSI queries with a FQDN instead of an IP and if so, why would this be useful?

Answer: ONTAP can be configured to respond to an iSCSI query with a FQDN instead of an IP, which may be useful if there is a firewall that NATs addresses.

References

The following references were used in this FAQ:

Note: All linked product documentation is for versions compatible with clustered Data ONTAP 8.2.

- Interoperability Matrix Tool
<http://mysupport.netapp.com/matrix/>
- SAN Configuration Guide
<http://support.netapp.com/documentation/productlibrary/index.html?productID=30092>
- SAN Administration Guide
<http://support.netapp.com/documentation/productlibrary/index.html?productID=30092>
- Data ONTAP DSM for Windows MPIO Documentation
<http://support.netapp.com/documentation/docweb/index.html?productID=60927>
- SnapDrive for Windows
<http://support.netapp.com/documentation/docweb/index.html?productID=60386>
- SnapDrive for Linux
<http://support.netapp.com/documentation/docweb/index.html?productID=60397>
- TR-3863: 10-Gigabit Ethernet FAQ v2.0
<https://fieldportal.netapp.com/Core/DownloadDoc.aspx?documentID=83857&contentID=65493>

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.