



Technical FAQ

All Flash FAS

ONTAP 9.4

Scott Bell, NetApp
Version 2.1
May 2018

Abstract

This document answers frequently-asked technical questions about Netapp® AFF systems. The FAQ covers the requirements and architecture that define AFF systems, as well as configuration recommendations, performance expectations, sizing guidelines, and interoperability support.

AFF8000 systems are supported with ONTAP® 8.3 GA and later releases. AFF A200, AFF A300, AFF A700, and AFF A700s systems are supported with ONTAP 9.1 GA and later releases. AFF A220 and AFF A800 systems are supported with ONTAP 9.4 GA and later releases.

Version History

Version	Date	Document Version History
1.0	5/2017	Initial revision
1.1	8/2017	Revised reallocation verbiage in Performance section
1.2	9/2017	Updated RAID group size information in Table 5
1.2.1	9/2017	Added 7.6TB SSD shelf support and updates to Table 1
2.0	11/2017	Consolidated AFF8000 Series and AFF A-Series content into one document
2.0.1	11/2017	Updated Figure 1
2.1	5/2018	ONTAP 9.4, AFF A220, and AFF A800 updates

TABLE OF CONTENTS

Version History	2
1 Overview	6
How are AFF systems different from FAS systems?	6
Can a FAS system configured with SSDs be converted to an AFF system?	6
Can an AFF system be converted to a FAS system?	6
What are the benefits of using an AFF system?	6
What applications and workloads are best-suited for AFF systems?	6
What break-fix support is available for the SSDs and other hardware used on AFF systems?	6
2 Requirements	7
What AFF systems are currently available?	7
What AFF systems were previously available?	7
What is the minimum, required ONTAP release for AFF systems?	7
What is the recommended ONTAP release for all AFF systems?	7
Which AFF systems provide the best performance?	7
What SSDs and storage enclosures are configurable with a new AFF system?	7
What SSDs and storage enclosures are supported with AFF systems?	7
3 Architecture	9
How does ONTAP deliver high performance and flash media endurance with AFF systems?	9
Does ONTAP include features that are specifically benefit AFF systems?	9
What key features are available with AFF systems running ONTAP 9?	9
What new features in ONTAP 9.2 releases benefit AFF systems?	10
What features in ONTAP 9.3 releases benefit AFF systems?	10
What new features in ONTAP 9.4 releases benefit AFF systems?	10
What ONTAP features enable storage efficiency on AFF systems?	11
What is Advanced Drive Partitioning?	11
What is the benefit of ADP?	11
What are the available types of ADP?	11
How is partitioning of root aggregates on AFF systems enabled?	11
What does the configuration of an AFF system with RD2 partitioning look like with ONTAP 9?	11
What is the maximum number of SSDs using RD2 partitioning on an AFF systems?	12
What size are the root partitions in half shelf, single-shelf, 1.5 shelf, and 2 shelf configurations?	12
What type of flash media is used in the SSDs shipped with AFF systems?	12
How can SSD media usage be monitored on an AFF system?	13
4 Configuration	14
What are the best practice recommendations for configuring storage on AFF HA pairs?	14
What configuration is recommended when a HA pair's nodes require different capacity?	15
What are the RAID group size recommendations for AFF systems?	15
What is the RAID group configuration recommendation when data partitions and unpartitioned SSDs are provisioned in the same aggregate?	15

What are the hot spare policy recommendations for AFF HA pairs?	16
What are the recommended data aggregate and RAID group configurations new AFF systems?	16
What are the recommended RAID group configurations when adding 12 SSDs to an AFF system?	16
What is the maximum number of shelves supported on an AFF A-Series HA pair?	17
What is the maximum number of shelves supported on an AFF8000 Series HA pair?	18
Can SSDs of different capacity be provisioned in the same aggregate?	18
Can SSDs of different capacity be configured in the same shelf?	18
Can shelves with different capacity SSDs be configured in the same stack?	18
What is inline adaptive compression?	18
Is inline adaptive compression enabled on AFF systems?	19
Can data in a volume that has not been previously compressed inline be adaptively compressed?	19
What is inline volume deduplication?	19
Is inline volume deduplication enabled on AFF systems?	19
What is always-on deduplication?	19
What is inline zero block deduplication?	19
Is inline zero block deduplication enabled on AFF systems?	19
What is inline data compaction?	19
Is inline data compaction enabled on AFF systems?	20
Is thin provisioning of volumes recommended on AFF systems?	20
What is inline aggregate deduplication?	20
Is inline aggregate deduplication supported on NVE-protected volumes?	20
What is aggregate background deduplication?	20
Is aggregate background deduplication supported on NVE protected volumes?	20
What are the configuration recommendations for an AFF cluster used as a SAN-only storage?	20
5 Performance	21
What is the performance of AFF systems?	21
How do AFF systems perform with VMware virtual desktop deployments?	21
How does AFF performance compare to hybrid FAS systems configured with Flash Cache™ or Flash Pool?	21
How full can an aggregate get and still sustain excellent performance?	21
Are ONTAP reallocation methods recommended with AFF systems?	21
6 Interoperability	22
Can AFF HA pairs be configured in an ONTAP cluster that contains FAS HA pairs?	22
Is an AFF HA pair supported in ONTAP MetroCluster™ configurations?	22
Is an AFF HA pair supported in ONTAP MetroCluster IP configurations?	22
Can an AFF system be configured as a NSE system?	22
What happens with compressed and deduplicated blocks when volume move is used to migrate data non-disruptively from a volume on an AFF node to a volume on a FAS node?	22
Will uncompressed blocks on a source FAS volume be adaptively compressed inline when volume move is used to transfer data to an AFF volume?	22
Is storage efficiency affected when a volume on an AFF system is the source of SnapMirror or SnapVault transfer and the destination is a volume on a FAS system?	23

Is 7-Mode Transition tool (7MTT) supported to move volumes on a 7-Mode FAS system to an AFF system?	23
Is FlexArray supported with AFF systems?	23
Is disk sanitization supported for SSDs in an AFF system?	23
Does Flash Cache work in an AFF system?.....	23
7 Additional Resources.....	24
8 Contact Us.....	25

LIST OF TABLES

Table 1) SSDs supported with AFF systems.	8
Table 2) Major features for AFF systems with ONTAP 9.....	9
Table 3) Major features for AFF systems with ONTAP 9.2.....	10
Table 4) Major features for AFF systems with ONTAP 9.3.....	10
Table 5) Major new features for AFF systems with ONTAP 9.4	10
Table 6) AFF storage efficiency features.....	11
Table 7) Data aggregate RAID group and spare recommendations for new systems running ONTAP 9.....	16
Table 8) Data aggregate RAID group configurations after adding 24 SSDs to a system running ONTAP 9.1	17
Table 9) Data aggregate RAID group configurations after adding a 24 SSDs to a system running ONTAP 9.2	17
Table 10) Maximum number of shelves supported per HA pair.....	18
Table 11) Maximum number of shelves supported per HA pair.....	18

LIST OF FIGURES

Figure 1) AFF SSD assignment and RD2 partitioning of 12 SSDs with ONTAP 9	12
Figure 2) AFF SSD assignment and RD2 partitioning of 24 SSDs with ONTAP 9	12
Figure 3) AFF SSD assignment and RD2 partitioning of 48 SSDs with ONTAP 9.2	12
Figure 4) A sample output	13
Figure 5) Split-shelf SSD assignment of a 96-drive AFF A300 system	14
Figure 6) SAS cabling for an AFF A300 system with 2 shelves attached.....	15
Figure 7) Space savings with inline data compaction.....	20

1 Overview

How are AFF systems different from FAS systems?

Controllers: Must be configured as a High Availability (HA) pair or Netapp MetroCluster™ (MCC) configuration

- Storage: Solid-State Drive (SSD)
- Personality: Flash-optimized personality must be enabled. Enablement is done on new AFF systems at time of manufacturing. Inline storage efficiencies are enabled by default.

Can a FAS system configured with SSDs be converted to an AFF system?

Answer: No. Conversion of a FAS system to an AFF system is not generally supported.

Can an AFF system be converted to a FAS system?

Answer: No. Converting an AFF system into a FAS system is not generally supported. Customers who require the flexibility of using both SSDs and HDDs in the same HA pair should purchase a FAS system instead.

What are the benefits of using an AFF system?

Answer: The primary benefit of using an AFF system is consistent low latency performance, typically 1 millisecond or less, for random read-write workloads, including at high IOPS workloads. The key advantages AFF systems have when compared to competitive all-flash arrays are:

- SAN and NAS support
- Scale-out and scale-up expansion flexibility for performance and capacity
- Nondisruptive data mobility across storage tiers in a cluster that consists of AFF and FAS HA pairs
- Netapp Data Fabric and public cloud integration
- Rich enterprise data management features of the NetApp ONTAP operating system
- Application integration with leading hypervisors (e.g. VMware, Citrix, Microsoft) and database applications (e.g. Oracle, SAP, Microsoft SQL Server)
- Proven enterprise-level system reliability, availability and support

What applications and workloads are best-suited for AFF systems?

Answer: Application workloads for which consistent low latency random reads and writes are important are a great fit for AFF systems. Many database applications require very fast response times, and this is also critical for virtual desktop and server deployments that support many virtual machines. Two database application examples are radiological imaging (MRI, X-ray, CT scan) and geophysical exploration.

What break-fix support is available for the SSDs and other hardware used on AFF systems?

Answer: All AFF hardware is covered by the standard NetApp 3-year support for new system purchases. Customers have the option of purchasing extended support for an additional 3 years, for 6 total years of support, that covers all hardware, including SSDs. There are not any limitations on extended support for SSDs based on workload type.

Pricing information about the extended support option is available in the [AFF Sales FAQ](#) on Field Portal.

2 Requirements

What AFF systems are currently available?

Answer: The following AFF A-Series systems are currently available from NetApp.

- AFF A-Series: AFF A220, AFF A200, AFF A300, AFF A700s, AFF A700, and AFF A800.

What AFF systems were previously available?

Answer: The AFF8000 Series was previously available before replacement by the AFF A-Series.

- AFF8000 Series systems: AFF8020, AFF8040, AFF8060, and AFF8080 EX

Note: AFF8020 is EOA as of June 2016. AFF8060 is EOA as of November 2016. AFF8040 and AFF8080 EX are EOA as of October 2017.

What is the minimum, required ONTAP release for AFF systems?

Answer: The minimum, required release depends on the AFF system.

- ONTAP 9.4 GA is the minimum, required release for AFF A220 and AFF A800 controllers.
- ONTAP 9.1 GA is the minimum, required release for AFF A200, AFF A300, AFF A700s, and AFF A700 controllers.
- ONTAP 8.3 GA is the minimum, required release for AFF8000 Series controllers.

What is the recommended ONTAP release for all AFF systems?

Answer: ONTAP 9.4 is the currently recommended release for all AFF platforms.

Which AFF systems provide the best performance?

Answer: The following systems are the best performing of their series due to the number of CPU cores and available system memory.

- AFF A-Series: AFF A700s, AFF A700, AFF A800
- AFF8000 Series: AFF8060, AFF8080 EX

What SSDs and storage enclosures are configurable with a new AFF system?

Answer: DS224C with 960GB, 3.8TB, 3.8TB NetApp Storage Encryption (NSE), and 7.6TB are configurable with all available models.

15.3TB drives are standard offerings on AFF A220, AFF A300, AFF A700s and AFF A700 systems. They are also available for purchase with the AFF A200, pending sales order approval.

800GB NSE drives are available for purchase with AFF A200, AFF A220, and AFF A300 systems only.

What SSDs and storage enclosures are supported with AFF systems?

Answer: The SSD part numbers supported with AFF systems, along with their respective marketing and right-sized used storage capacities are listed in Table 1. New systems are sold with all SSD models except 200GB (X466A-R6, X466B-R6), 400GB (X438A) and 1.6TB (X439A).

Table 1) SSDs supported with AFF systems.

SSD Marketing Part Number	Marketed Capacity	Right-sized Used Capacity	Supported Storage Enclosure
X446A-R6 ¹ X446B-R6 ¹	200 GB	190,532 MiB	DS2246 DS224C
X438A-R6 ¹	400 GB	381,304 MiB	
X440A-R6 X440B-R6	800 GB NSE	762,847 MiB	
X447A-R6	800 GB	762,847 MiB	
X371A-R6	960 GB	915,456 MiB	
X365A-R6	1.6 TB NSE	1,525,935 MiB	
X439A-R6 ¹	1.6 TB	1,525,935 MiB	
X356A-R6 X357A-R6	3.8 TB	3,662,580 MiB	
X358A-R6	3.8 TB NSE	3,662,580 MiB	
X319A	7.6TB	7,325,400 MiB	DS224C
X670A-R6	15.3 TB	14,650,368 MiB	

Note: ¹Supported but not offered when configuring a new AFF system.

3 Architecture

How does ONTAP deliver high performance and flash media endurance with AFF systems?

Answer: ONTAP has provided good write performance, storage efficiency, and media endurance with SSDs over many releases. These benefits are the result of the log-structured file system architecture. More recently, feature enhancements in ONTAP 9 dramatically increased random read performance (both IOPS and latency), while adding inline storage efficiency capabilities.

Streamlined read paths and read-ahead algorithms in ONTAP keep read latencies low. Write operations and SSD firmware are tightly coupled and data layouts are well aligned, resulting in high IOPS and low latency. Incoming I/O operations from hosts and clients are acknowledged immediately after data is placed in memory buffers and protected by NVRAM.

Coalescing (batching) incoming write I/Os in memory into larger RAID stripe writes to storage, a core feature of ONTAP from its first release, results in efficient writes to SSDs (low write amplification) and aligns with how SSDs write to free space and manage wear-leveling. This produces good performance and flash media durability.

Does ONTAP include features that are specifically benefit AFF systems?

Answer: Yes. ONTAP 8.3 and later releases include performance, storage efficiency, and other features specifically designed for AFF systems. Although some features are also available on FAS systems configured with SSD aggregates, NetApp Flash Pool™ aggregates, or HDD-only aggregates, AFF systems get the greatest benefit from these enhancements. Other features are available only with AFF systems.

What key features are available with AFF systems running ONTAP 9?

Answer: The key features in ONTAP 9 and the associated benefits, are listed in Table 2.

Table 2) Major features for AFF systems with ONTAP 9

New Feature	Benefit
Inline adaptive compression	Overall storage efficiency
Inline volume deduplication	
Inline data compaction	
Root-Data-Data (RD2) partitioning	Storage efficiency increase ≈17% more usable capacity
NetApp Volume Encryption (NVE)	Enables software-based encryption for Data At-Rest
Onboard Key Manager (OKM)	Key management for data at-rest encryption
FlexGroup	Scale-out NAS container
Quality of Service	Bully workload management

What new features in ONTAP 9.2 releases benefit AFF systems?

Answer: The major features in ONTAP 9.2 and the associated benefits, are listed in Table 3.

Table 3) Major features for AFF systems with ONTAP 9.2

New Feature	Benefit
Inline aggregate deduplication	Overall storage efficiency
QoS minimum throughput policy	Performance management: SAN volumes only
Fabric Pool	Auto-tiering of NetApp Snapshot™ data to lower cost storage
Shared spare SSD for data partitions up to 48 SSD	Storage efficiency: ≈2.5% more usable capacity

What features in ONTAP 9.3 releases benefit AFF systems?

Answer: The major new features in ONTAP 9.3, and the associated benefits, are listed in Table 4.

Table 4) Major features for AFF systems with ONTAP 9.3

New Feature	Benefit
Background aggregate deduplication scanner for newly written data	Increased overall storage efficiency
Adaptive QoS policies	Minimum throughput and limits adapt automatically
Read acceleration for SAN workloads	Increased read performance for FC and iSCSI workloads
iSCSI parallelization	Performance increase for iSCSI processing
External key management for NVE	Additional security options for data at-rest

What new features in ONTAP 9.4 releases benefit AFF systems?

Answer: The major new features in ONTAP 9.4, and the associated benefits, are listed in Table 5.

Table 5) Major new features for AFF systems with ONTAP 9.4

New Feature	Benefit
Background aggregate deduplication scanner for existing data	Increased overall storage efficiency for data already existing on disk post-upgrade to ONTAP 9.4
Adaptive QoS policies	Individual maximum throughputs and policies applied at file, LUN, and Virtual Volume (VVol) levels
FC-NVMe support for mid-range and high-range AFF A-Series arrays	Increased performance
Fabric Pool enhancements	Auto-tiering of primary file system data, in addition to Snapshot data, to lower cost storage

What ONTAP features enable storage efficiency on AFF systems?

Answer: Multiple ONTAP storage efficiencies help customers increase the effective capacity of AFF systems.

Table 6) AFF storage efficiency features

Standard ONTAP Features	AFF Core ONTAP Features
NetApp RAID DP®	Advanced drive partitioning (RD2)
NetApp RAID-TEC®	Inline adaptive compression
Thin provisioning	Inline zero block deduplication
Snapshot copies	Inline volume and aggregate deduplication
NetApp FlexClone technology	Inline data compaction

What is Advanced Drive Partitioning?

Answer: Advanced Drive Partitioning (ADP) is a storage efficiency feature available in ONTAP 8.3 and later releases. ADP enables capacity sharing between aggregates and controllers in a HA pair.

What is the benefit of ADP?

Answer: ADP provides more usable capacity when compared to provisioning with whole, unpartitioned drives. Less capacity is consumed to provision root aggregates, which results in more capacity available for provisioning data aggregates. Because parity and spare drives are shared between both controllers, even more capacity is available in the HA pair.

What are the available types of ADP?

Answer: ADP can be implemented in one of two ways on AFF systems:

- Root-Data (RD) partitioning: Available in ONTAP 8.3 GA and later. Each SSD is divided into 2 partitions: A smaller, or thin, root partition and a larger, or thick, data partition. Actual partition sizes will be dependent upon the number of drives present at the time of system initialization.
- Root-Data-Data (RD2) partitioning: Available in ONTAP 9.0 GA and later. Each SSD is divided into 3 partitions: A smaller root partition and 2 larger data partitions. Actual partition sizes will depend on the number of attached SSDs at the time of system initialization.

For more information, refer to the [Advanced Drive Partitioning SE Presentation](#) on Field Portal.

How is partitioning of root aggregates on AFF systems enabled?

Answer: RD2 partitioning is the default method of provisioning root aggregates on AFF systems running ONTAP 9.0 or later. New AFF systems are shipped with root aggregates already provisioned with partitions, and data partitions are assigned to each node as spares.

ADP can be enabled easily through Option 9 of the Special Boot Menu in ONTAP 9.2 or later.

What does the configuration of an AFF system with RD2 partitioning look like with ONTAP 9?

Answer: Figures 1, 2, and 3 illustrate examples of RD2 partitioning on AFF systems running ONTAP 9. Figure 1 shows a 12-drive system, Figure 2 a 24-drive system, and Figure 3 shows a 48-drive system.

Note: Data (R), parity (DP and P), and spare (S) root partition assignment to bay numbers may differ. Typical examples are illustrated

Each numbered rectangle represents a physical SSD in a DS2246/DS224C shelf bay. The small square in each SSD is a root partition, and the two larger rectangles are the data partitions. Each partition is configured as: a root that stores root volume data (R), a parity or diagonal parity (P or DP), or a spare (S).

Figure 1) AFF SSD assignment and RD2 partitioning of 12 SSDs with ONTAP 9




Figure 2) AFF SSD assignment and RD2 partitioning of 24 SSDs with ONTAP 9





Figure 3) AFF SSD assignment and RD2 partitioning of 48 SSDs with ONTAP 9.2



What is the maximum number of SSDs using RD2 partitioning on an AFF systems?

Answer: Up to 48 SSDs can be partitioned per HA pair, or 24 SSDs per node.

What size are the root partitions in half shelf, single-shelf, 1.5 shelf, and 2 shelf configurations?

Answer: Root partition sizes vary based on controller model, ONTAP release, and the number of SSDs attached during system initialization. Refer to NetApp [Hardware Universe](#) for additional information.

What type of flash media is used in the SSDs shipped with AFF systems?

Answer: AFF systems use MLC and TLC SSDs. All supported SSDs are tested to meet the same requirements for performance, quality, reliability, and longevity of use (that is, the ability to perform write-intensive workloads for 7 years).

Media endurance measured in drive writes per day (DWPD) over the expected life of a SSD is a more reliable way of determining SSD fitness for enterprise storage system use than whether the SSDs used are MLC and TLC or some other NAND flash technology. As the capacity of a drive increases, lower DWPD SSDs meet endurance requirements when handling enterprise workloads.

How can SSD media usage be monitored on an AFF system?

Answer: SSD media can be monitored using the node-level command **storage disk show -a** and specifying a SSD ID. The output includes a percentage of the rated media life that has been used, and a percentage of the spare blocks that have been consumed. The spare-blocks-consumed percentage is a better indicator of how much SSD media has been used.

A sample **storage disk show -a** output is shown in Figure 4 followed by some field descriptions.

Figure 4) A sample output

```
stl6280DLEcmodel1:> node run -node stl6280DLEcmodel1-01
stl6280DLEcmodel1-01> storage disk show -a 6a.00.0
Disk: 6a.00.0
Shelf: 0
Bay: 0
Serial: XXVE0P4A
Vendor: NETAPP
Model: X446_RALS200MCHT
Rev: NA01
RPM: N/A
WWN: 5:000:cca013:17ad88
UID: 5000CCA0:1317AD88:00000000:00000000:00000000:00000000:00000000:00000000:00000000:00000000
Downrev: no
Pri Port: A
Sec Name: 6b.00.0
Sec Port: B
Power-on Hours: 8156
Blocks read:      8827073939
Blocks written:   11862011056
Time interval: 196:10:38
Glist count: 0
Scrub last done: 00:15:37
Scrub count: 161
LIP count: 0
Dynamically qualified: No
Power cycle count: 0
Power cycle on error: 0
Current owner: 1873768414
Home owner: 1873768414
Reservation owner: 1873768414
Rated life used: 0 %
Spare blocks consumed: 0 %
Spare blocks consumed limit: N/A
```

- **Power-on Hours:** This is the number of hours the SSD has been powered on.
- **Glist count:** Also known as the Growth list, the Glist count records the number of bad sectors on the SSD.
- **Rated life used:** This parameter is a vendor-specific estimate of the percentage of device life used based on the actual device usage and the manufacturer's prediction of device life. A value of 100 (or greater) indicates that the estimated endurance has been consumed, but it might not indicate a device failure.
- **Spare blocks consumed:** Reported by the SSD as a percentage, this is the current worst-case state of the internal resource pool of the SSD. This parameter is monitored by ONTAP to trigger an Event Management System (EMS) and AutoSupport™ event when the spare-blocks-consumed limit has been reached.
- **Spare blocks consumed limit:** This is the threshold value used internally by the SSD to trigger a predictive failure event based on spare blocks consumed. The limit for currently shipping SSDs is 90% even though ONTAP currently shows the spare-blocks-consumed limit as N/A.

4 Configuration


What are the best practice recommendations for configuring storage on AFF HA pairs?

Answer: AFF systems are deployed to deliver low latency for random read-write workloads over a wide range of IOPS. Consequently, best practice configuration recommendations and support limits differ from those that apply to hybrid or HDD-only FAS systems.

To achieve best performance from an AFF HA pair, it is important to distribute I/O traffic to and from the SSDs across all available SAS channels (ports) connected to each storage shelf, and to minimize the number of shelves per SAS stack. The recommendation is to keep the number of SSDs assigned to each SAS channel to 12 or fewer per shelf. To accomplish this, there are two important recommendations:

- **12x SSD shelf configuration:** Systems shipped from the factory assign SSDs in shelf bays 0-5 to Node-A, and SSDs in shelf bays 6-11 are assigned to Node-B.
- **24x SSD shelf configuration:** Systems shipped from the factory assign 12x SSDs in each full shelf to each controller. The SSDs in bays 0-11 are assigned to Node-A, while SSDs in bays 12-23 are assigned to Node-B.

Figure 5) Split-shelf SSD assignment of a 96-drive AFF A300 system




This is referred to as “split-shelf” assignment, and it is the default automatic assignment policy for full shelves added to AFF systems beginning with ONTAP 8.3.2. Figure 5 illustrates split-shelf configuration of an AFF A300 with four external shelves of SSDs.

NetApp recommends that a maximum of two shelves per stack be implemented to enable best performance under very high IOPS workloads. Two shelves per stack, when combined with split-shelf SSD assignment, results in 12 SSDs per SAS channel. Implementing four shelves per stack results in 24 SSDs per SAS channel. This can result in higher latency when SSDs on the stack are very busy.

Splitting the SSD assignments in each shelf between both nodes of the HA pair enables all four SAS channels connected to the stack, two from each controller, to actively carry I/O. This enables the administrator to configure two shelves per stack, with 12 SSDs per SAS channel.

Another recommendation is to use one onboard SAS port and one SAS adapter card port per controller to dual-path connect each stack of shelves, when possible. This maintains access from the controller to SSDs should one SAS port fail. Since September 2015, all AFF8000 systems ship with a 4-port SAS adapter card in each controller in addition to the onboard SAS ports. Figure 6 illustrates the recommended SAS port usage for a two-shelf configuration.

Figure 6) SAS cabling for an AFF A300 system with 2 shelves attached



What configuration is recommended when a HA pair's nodes require different capacity?

Answer: If one node requires more capacity than its partner, the SSD-to-node assignment will be asymmetric in one or more shelves. NetApp recommends minimizing the number of shelves that do not use the default split-shelf assignment (12 SSDs per node). Shelves with asymmetric drive assignment should be configured on SAS stacks that have one or two shelves only. This keeps the number of SSDs, per channel, at or below 24.

Scenario: An AFF A300 HA pair is configured with 72x SSDs. Node-A requires twice the capacity of Node-B.

Recommendation: Split ownership of two shelves equally. Assign the 24 remaining SSDs to Node-A. The third shelf should be configured as a stand-alone stack. This results in 12x SSDs per SAS channel.

What are the RAID group size recommendations for AFF systems?

Answer: NetApp recommends using RAID DP groups of up to 23 or 24 data partitions or unpartitioned SSDs to enable high usable capacity and storage efficiency. The high reliability and fast performance of SSDs enables the administrator to configure large RAID DP groups with high confidence. When a new RAID group is added to an existing aggregate, NetApp recommends that the RAID group be at least half the size (capacity and number of SSDs) as the largest existing RAID group in that aggregate.

What is the RAID group configuration recommendation when data partitions and unpartitioned SSDs are provisioned in the same aggregate?

Answer: In most cases, NetApp recommends provisioning data partitions and unpartitioned SSDs in separate RAID groups, because it yields the most usable capacity. However, when a shelf is added to an AFF system that was first deployed with only one shelf, expanding the existing RAID group in each node's data aggregate with unpartitioned SSDs is preferred. This approach provides more usable capacity than creating a new RAID group with unpartitioned SSDs.

What are the hot spare policy recommendations for AFF HA pairs?

Answer: For HA pairs configured with 48 or fewer SSDs that have data aggregates provisioned with partitions, NetApp recommends retaining one spare data partition per node. In an active-passive HA pair configuration, only one node is configured with a data aggregate. Only the active node needs a data partition spare.

For a HA pair configured with three or more shelves, NetApp recommends retaining one spare, unpartitioned SSD per node. An unpartitioned SSD can replace either a data partition on a failed SSD or a failed unpartitioned SSD. A separate data partition spare is not required or recommended.

If replacing a failed SSD is likely to be deferred or delayed for more than two days (e.g. due to personnel scheduling, support agreement or shipment time from the nearest spare parts depot), consider retaining two SSDs as hot spares per node.

If a node is configured with more than 100 total data partitions plus unpartitioned SSDs, NetApp recommends retaining a minimum of two spare unpartitioned SSDs.

What are the recommended data aggregate and RAID group configurations new AFF systems?

Answer: Table 8 shows the recommended RAID group configurations for the data aggregate on each node, including the minimum hot spare recommendations, for configurations with up to 72 SSDs. In the table, *d* signifies a data partition or unpartitioned SSD and *p* signifies a parity partition or unpartitioned SSD. Most of the recommendations assume that equal numbers of SSDs are assigned to each node. The 1.5 and 2.5 shelf examples illustrate unbalanced assignments of SSDs.

Table 7) Data aggregate RAID group and spare recommendations for new systems running ONTAP 9

SSDs	Data Aggregate RAID Groups per Node	Hot Spare per Node
12	rg0: 9 <i>d</i> partitions + 2 <i>p</i> partitions	1 data partition
24	rg0: 21 <i>d</i> partitions + 2 <i>p</i> partitions	1 data partition
36	rg0: 22 <i>d</i> partitions + 2 <i>p</i> partitions rg1: 9 <i>d</i> partitions + 2 <i>p</i> partitions	1 data partition
48	rg0: 22 <i>d</i> partitions + 2 <i>p</i> partitions rg1: 21 <i>d</i> partitions + 2 <i>p</i> partitions	1 data partition
60	Node-A - rg0: 22<i>d</i> partitions + 2<i>p</i> partitions - rg1: 22<i>d</i> partitions + 2<i>p</i> partitions - rg2: 9<i>d</i> + 2<i>p</i> unpartitioned SSDs Node-B - rg0: 22<i>d</i> partitions + 2<i>p</i> partitions - rg1: 22<i>d</i> partitions + 2<i>p</i> partitions	Node A: 1 unpartitioned SSD Node B: 1 data partition
72	rg0: 22 <i>d</i> partitions + 2 <i>p</i> partitions rg1: 22 <i>d</i> partitions + 2 <i>p</i> partitions rg2: 9 <i>d</i> + 2 <i>p</i> unpartitioned SSDs	1 unpartitioned SSD

What are the recommended RAID group configurations when adding 12 SSDs to an AFF system?

Answer: When a full shelf is added to an AFF system, RAID group recommendations are similar to those for provisioning RAID groups on new systems. However, there are other considerations.

A general guideline is to provision data partitions and unpartitioned SSDs in different RAID groups to avoid creating root partitions on unpartitioned SSDs that are only usable in root aggregates. In some situations, however, NetApp recommends doing exactly that because it provides more usable capacity than creating a new RAID group of unpartitioned SSDs.

Additionally, a new RAID group should have a minimum of 11 SSDs. If an existing RAID group is to be expanded, at least 6 SSDs should be added at a time.

Table 8) Data aggregate RAID group configurations after adding 24 SSDs to a system running ONTAP 9.1

Starting SSD Quantity	Original Data Aggregate RAID Group Configuration per Node	RAID Group Configuration per Node After Adding 12 SSDs
12	rg0: 9d partitions + 2p partitions	rg0: 21d partitions + 2p partitions ¹
24	rg0: 21d partitions + 2p partitions	Node-A and Node-B : rg0: 21d partitions + 2p partitions Node A only: rg1: 9d + 2p unpartitioned SSDs
36	rg0: 21d partitions + 2p partitions rg1: 10d partitions + 2p partitions	rg0: 21d partitions + 2p partitions rg1: 22d partitions + 2p partitions ¹
48	rg0: 21d partitions + 2p partitions rg1: 22d partitions + 2p partitions	Node A and B: rg0: 21d partitions + 2p partitions rg1: 22d partitions + 2p partitions Node A only: rg2: 9d + 2p unpartitioned SSDs

Note: ¹The 12 SSDs added to RAID group rg0 in each node's data aggregate are partitioned automatically, and the resulting root partitions become spares that are usable only by the root aggregates.

Table 9) Data aggregate RAID group configurations after adding a 24 SSDs to a system running ONTAP 9.2

Starting SSD Quantity	Original Data Aggregate RAID Group Configuration per Node	RAID Group Configuration per Node After Adding 24-drive Shelf
24	rg0: 21d partitions + 2p partitions	rg0: 21d partitions + 2p partitions rg1: 9d + 2p unpartitioned SSDs
48	rg0: 21d partitions + 2p partitions rg1: 22d partitions + 2p partitions	rg0: 21d partitions + 2p partitions rg1: 22d partitions + 2p partitions rg2: 9d + 2p unpartitioned SSDs

Note: If a node has one or more spare data partitions, and a new RAID group of unpartitioned SSDs will be added to an existing data aggregate, select only unpartitioned SSD spares to create the new RAID group. If a spare data partition is included when creating the RAID group, unpartitioned SSDs will be automatically partitioned to match the size of the data partition.

What is the maximum number of shelves supported on an AFF A-Series HA pair?

Answer: The maximum number of shelves supported per AFF A-Series HA pair varies according to system model, shelf type, and configuration. Table 11 provides the details.

Table 10) Maximum number of shelves supported per HA pair.

Model	Maximum External Shelves	Maximum Supported Drives
AFF A200	5	144
AFF A220	5	144
AFF A300	16	384
AFF A700s	8	216
AFF A700	20	480
AFF A800	8	240

What is the maximum number of shelves supported on an AFF8000 Series HA pair?

Answer: The maximum number of shelves supported per AFF8000 Series HA pair varies according to system model, shelf type, and configuration. Table 12 provides the details.

Table 11) Maximum number of shelves supported per HA pair.

AFF Model	Maximum External Shelves	Maximum Supported Drives
AFF8020	10	240
AFF8040	16	384
AFF8060	20	480
AFF8080 EX	20	480

Can SSDs of different capacity be provisioned in the same aggregate?

Answer: Yes. Different capacity SSDs should be provisioned in separate RAID groups so that the full capacity of all SSDs can be used.

Can SSDs of different capacity be configured in the same shelf?

Answer: Yes. NetApp recommends keeping shelf configurations as simple as possible. In a single shelf, NetApp recommends that no more than two SSDs of different capacity and to physically group like-capacity SSDs together.

Can shelves with different capacity SSDs be configured in the same stack?

Answer: Yes.

What is inline adaptive compression?

Answer: Adaptive compression evaluates an 8KB group of data, or two consecutive 4KB WAFL® blocks and compresses the data, if possible, into a single 4KB WAFL block. Inline adaptive compression performs this before data is written to storage media. Adaptive compression was introduced with ONTAP 8.3.1 and is supported as an inline process with AFF systems.

Is inline adaptive compression enabled on AFF systems?

Answer: Yes. Inline adaptive compression is enabled by default on all newly created thin-provisioned volumes. It is not enabled automatically on volumes created with volume move. Inline adaptive compression can be enabled or disabled by modifying the volume efficiency policy.

Can data in a volume that has not been previously compressed inline be adaptively compressed?

Answer: Yes. ONTAP 8.3.2 introduced a background adaptive compression scanner that can be run on one volume at a time to compress blocks.

The command to run the background adaptive compression scanner is:

```
volume efficiency start -volume vol_name -compression true-scan-old-data true
```

There is an advanced privilege option to compress shared blocks (that is deduplicated or cloned blocks), **-shared-blocks** (or **-a**) **true**, and a second advanced option to compress blocks in Snapshot copies, **-snapshot-blocks** (or **-b**) **true**.

What is inline volume deduplication?

Answer: Inline volume deduplication reduces the amount of data that is written to storage when new blocks written to memory from a host or client match the content of blocks that were recently written.

Fingerprint and byte-for-byte comparisons are done in memory before incoming data is committed to storage by means of a consistency point (CP). Inline deduplication was introduced with ONTAP 8.3.2. The principal benefit is lower capacity requirements when patching or moving multiple virtual machines (VMs) with very similar content that reside in the same volume.

Is inline volume deduplication enabled on AFF systems?

Answer: Yes. Inline volume deduplication is enabled by default on all newly created thin-provisioned volumes. Inline volume deduplication can be enabled or disabled by modifying the volume efficiency policy.

What is always-on deduplication?

Answer: Always-on deduplication is normal background deduplication with volume deduplication jobs scheduled to run as frequently as every minute.

What is inline zero block deduplication?

Answer: Inline zero block deduplication detects incoming 4KB blocks that contain all zeroes. Rather than writing a block of zeros to storage, metadata is updated to reflect that a block contains only zeros. The principal benefit of inline zero block deduplication is faster provisioning of virtual machines and LUNs that were created with all zero blocks. Reduced write workload is a secondary benefit.


Is inline zero block deduplication enabled on AFF systems?

Answer: Yes. Inline zero block deduplication is enabled by default on all newly created thin-provisioned volumes.

What is inline data compaction?

Answer: Inline data compaction detects incoming I/O sizes less than 4KB and attempts to group multiple small writes together in a single block. Inline data compaction was introduced with ONTAP 9. The principal benefit is greater storage efficiency. Inline data compaction occurs after inline adaptive compression and inline deduplication have completed.

Figure 7) Space savings with inline data compaction



Is inline data compaction enabled on AFF systems?

Answer: Yes. Inline data compaction is enabled by default on all newly created thin-provisioned volumes. Inline data compaction can be enabled and disabled via the volume efficiency policy.

Is thin provisioning of volumes recommended on AFF systems?

Answer: Yes. NetApp recommends to thin provision all volumes whether they contain files or LUNs. Beginning with ONTAP 8.3.2, this is the default policy for newly created volumes.

What is inline aggregate deduplication?

Inline aggregate deduplication, or cross-volume inline deduplication, was introduced with ONTAP 9.2. It aims to reduce the amount of data that is written to storage when new blocks written to memory from a host or client match the content of existing blocks in the aggregate. The principal benefit is lower capacity requirements when patching or moving multiple virtual machines (VMs) with very similar content that reside in volumes contained within the same aggregate.

Is inline aggregate deduplication supported on NVE-protected volumes?

Answer: No. Inline aggregate deduplication is not supported on NVE protected volumes.

What is aggregate background deduplication?

Answer: Aggregate background deduplication, or cross-volume background deduplication, was introduced with ONTAP 9.3. Its scope is limited to deduplicating newly written blocks that were not addressed by inline processing. Aggregate background deduplication can be enabled only on volumes with volume background deduplication enabled. Ensure that the volume efficiency policy is set to 'auto.'

ONTAP 9.4 introduced aggregate background deduplication of "cold" data, or data that has already been committed to disk. By deduplicating existing data, a greater storage efficiency gain can be realized.

Is aggregate background deduplication supported on NVE protected volumes?

Answer: No. Cross-volume background deduplication is not supported on NVE protected volumes.

What are the configuration recommendations for an AFF cluster used as a SAN-only storage?

Answer: The [FC SAN Optimized AFF Setup Guide](#) provides details regarding factory-configured SAN optimized AFF systems. More information can be found in [TR-4480: AFF SAN-Optimized Configuration](#).

5 Performance

What is the performance of AFF systems?

Answer: The best way to determine the performance capability of an AFF HA pair or multi-node cluster is to use NetApp System Performance Modeler (SPM). SPM performs modeling of random workloads on AFF systems at low latencies. Latencies from 0.5 ms to 3 ms, done in 0.5 ms increments, are supported. Inline storage efficiencies and ADP for root aggregates are also supported.

How do AFF systems perform with VMware virtual desktop deployments?

Answer: NetApp [TR-4428: NetApp All Flash FAS Solution for VMware Horizon 6 and vSphere Virtual Volumes: A Technical Preview](#) provides a reference architecture for deploying AFF systems. The report includes performance and storage efficiency results.

How does AFF performance compare to hybrid FAS systems configured with Flash Cache™ or Flash Pool?

Answer: All reads are served from flash media with an AFF system. Block reads that are already in cache with a hybrid FAS system are similarly fast, but reads from HDDs are slower because most application workloads will have some reads from HDDs. Consequently, for most random read-write workloads, average read latency from SSD aggregates on an AFF system are lower than when the same workloads are provisioned on a Flash Pool aggregate, or HDD aggregate fronted by Flash Cache. Furthermore, the variation in latency is generally less for a workload on an AFF system when compared to a FAS with cache acceleration.

How full can an aggregate get and still sustain excellent performance?

Answer: According to performance testing, aggregates that have more than 90% capacity consumed, continue to serve data with excellent response times.

Are ONTAP reallocation methods recommended with AFF systems?

Answer: No. ONTAP reallocation methods are specifically designed to improve sequential read-write performance for volumes and LUNs provisioned on HDD-only or Flash Pool aggregates.

SSDs do not have rotational seek latencies, and writes to SSDs are always to contiguous blocks because of NAND flash technology. In addition, AFF systems are expected to be very busy and aggregates are often relatively small and quite full. Therefore, enabling read reallocation on AFF volumes or LUNs, and free space reallocation on SSD aggregates, will provide minimal, if any, benefits.

6 Interoperability

Can AFF HA pairs be configured in an ONTAP cluster that contains FAS HA pairs?

Answer: Yes. Deploying a cluster that has both AFF and FAS HA pairs enables customers to move volumes, LUNs, and VVols non-disruptively between storage tiers. This is a key benefit of AFF systems that competitive all-flash arrays cannot match.

Is an AFF HA pair supported in ONTAP MetroCluster™ configurations?

Answer: Yes. ONTAP MetroCluster (MCC) support depends on the AFF system.

- AFF A-Series: AFF A300 and AFF A700 support up to an 8-Node Fabric Attached MetroCluster configuration. AFF A200, AFF A220, and AFF A700s are not supported in MetroCluster systems.
- AFF8000 Series: AFF8000 platforms running ONTAP 8.3 support up to a 4-Node Fabric Attached MetroCluster configuration. ONTAP 9.2 provides support for up to an 8-Node Fabric MetroCluster.

Note: Root aggregates on AFF nodes configured in a MCC are provisioned with dedicated SSDs. ADP is not supported with MCC, but is supported with MCCIP.

Is an AFF HA pair supported in ONTAP MetroCluster IP configurations?

Answer: Yes. ONTAP MetroCluster IP (MCC) support depends on the AFF system.

- AFF A-Series: AFF A700 and AFF A800 systems support MetroCluster IP (MCCIP). For more information about MetroCluster IP, refer to the [MetroCluster in ONTAP SE Technical Presentation](#).

Note: AFF A700 and AFF A800 systems also support ADP with MCCIP on systems running ONTAP 9.4 or later releases.

Can an AFF system be configured as a NSE system?

Answer: Yes. 800GB and 3.8TB NSE SSDs are offered for purchase with an AFF system. All SSDs attached to the HA pair must be NSE SSDs, and cannot be used concurrently with non-NSE SSDs. NSE SSDs offered by NetApp are FIPS 140-2 Level 2 validated.

Note: NSE drives are not supported with MCC systems. The recommended encryption solution for MCC consists of FAS controllers with NetApp FlexArray® enabled, configured with an EF570 array and self-encrypting SSDs.

What happens with compressed and deduplicated blocks when volume move is used to migrate data non-disruptively from a volume on an AFF node to a volume on a FAS node?

Answer: The `volume move` command uses physical block replication to move data. Meaning that the format of blocks on the source volume is copied to the destination volume. Adaptively compressed blocks on a source AFF volume are similarly compressed on the destination volume. In addition, any deduplicated blocks remain deduped on the destination volume.

Will uncompressed blocks on a source FAS volume be adaptively compressed inline when volume move is used to transfer data to an AFF volume?

Answer: No. The `volume move` command uses physical block replication to transfer data. Source volume data blocks are copied, without any format change, to the destination AFF volume. Inline storage efficiencies are not applied during volume move, even when that is the policy on the destination volume.

Starting with ONTAP 8.3.2, volume data transferred from a FAS system to an AFF system using `volume move` can be adaptively compressed after blocks have been transferred to the AFF volume by running background adaptive compression once per volume. NetApp recommends running background adaptive compression on one volume at a time when an AFF system is serving application or user workloads.

Is storage efficiency affected when a volume on an AFF system is the source of SnapMirror or SnapVault transfer and the destination is a volume on a FAS system?

Answer: Whether compression or deduplication space savings on a source AFF volume are preserved or not during transfer and when data is written to a destination FAS volume depends upon the storage efficiency policies of both the source and destination volumes.

When the source and destination volumes have the same compression or deduplication policy, the associated savings are preserved during transfer and on the destination volume. When compression policies are different, all storage efficiency savings (compression, deduplication, and compaction) are lost during transfer. The destination volumes policies are then applied to the data, either inline as data is being written, or post-write in a background job, depending on the destination volume policy.

Is 7-Mode Transition tool (7MTT) supported to move volumes on a 7-Mode FAS system to an AFF system?

Answer: Yes. However, inline compression and deduplication policies that are enabled on a destination volume are not applied to data that is transferred using 7MTT. 7MTT uses physical block replication technology to move data. Compressed and deduplicated blocks on the source volume are transmitted with the same format and storage efficiencies to the destination volume.

If data on the source volume is uncompressed, background adaptive compression can be applied one time on an AFF volume to compress blocks. NetApp recommends to run background adaptive compression on only one volume at a time while an AFF system is serving host/client applications. Background deduplication can also be run to dedupe transferred data.

Is FlexArray supported with AFF systems?

Answer: No. FlexArray uses array LUNs provisioned on HDDs or SSDs for storage. ONTAP cannot determine what type of physical drive a LUN resides on, so FlexArray is not supported.

Is disk sanitization supported for SSDs in an AFF system?

Answer: Yes. A single pass of random overwrites is more than enough to securely remove all user data from SSDs. Disk sanitization must be enabled on a node, and all spare SSDs must be unpartitioned to be sanitized.

The nodeshell command to sanitize SSDs is: `disk sanitize start -p -r <disk list>`

Does Flash Cache work in an AFF system?

Answer: No. ONTAP is coded not to insert blocks that are read from SSDs into Flash Cache. When data blocks are read from SSDs on AFF systems, the latency is as low when compared to reading data blocks stored in Flash Cache.

AFF systems, by default, do not come configured with Flash Cache cards.

7 Additional Resources

AFF Product Page – Field Portal

<https://fieldportal.netapp.com/content/208699?assetComponentId=208798>

Technical Reports

- TR-4428: NetApp All Flash FAS Solution for VMware Horizon 6 and vSphere Virtual Volumes: A Technical Preview <http://www.netapp.com/us/media/tr-4428.pdf>
- TR-4435: SAP HANA on NetApp All Flash FAS Systems with NFS – Configuration Guide <http://www.netapp.com/us/media/tr-4435.pdf>
- TR-4436: SAP HANA on NetApp All Flash FAS Systems with Fibre Channel Protocol – Configuration Guide <http://www.netapp.com/us/media/tr-4436.pdf>
- TR-4450: VMware Horizon Desktop as a Service on NetApp All Flash FAS www.netapp.com/us/media/tr-4450.pdf
- TR-4476: NetApp Data Compression, Deduplication, and Data Compaction www.netapp.com/us/media/tr-4476.pdf
- TR-4490: Oracle Database 12c Performance: Protocol Comparison Using Clustered ONTAP www.netapp.com/us/media/tr-4490.pdf
- TR-4492: MongoDB on the Data Fabric Enabled by NetApp Reference Architecture with All Flash FAS, Cloud ONTAP, and VMware vSphere www.netapp.com/us/media/tr-4492.pdf
- TR-4307: NetApp All Flash FAS Solution for Non-persistent Desktops with VMware Horizon View www.netapp.com/us/media/tr-4307.pdf
- TR-4324: Electronic Device Automation (EDA) Verification Workloads and All-Flash FAS (AFF) www.netapp.com/us/media/tr-4322.pdf
- TR-4335: NetApp All Flash FAS Solution for Persistent Desktops with VMware Horizon View www.netapp.com/us/media/tr-4335.pdf
- TR-4342: NetApp All-Flash FAS Solution for Persistent and Nonpersistent Desktops with Citrix XenDesktop and XenApp <http://www.netapp.com/us/media/tr-4342.pdf>

FAQs

- Foreign LUN Import (FLI) Technical FAQ <https://fieldportal.netapp.com/content/324335>
- Storage Subsystem Technical FAQ <https://fieldportal.netapp.com/content/195603>

Presentations

- AFF Technical Content Collection <https://fieldportal.netapp.com/collections/500339>
- ONTAP Storage Technical Content Collection <https://fieldportal.netapp.com/collections/593410>

8 Contact Us

Do you have a question that you think we should include in this FAQ?

Ask us at docfeedback@netapp.com.

Include AFF TECHNICAL FAQ in the subject line.

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.

Copyright Information

Copyright © 1994–2018 NetApp, Inc. All rights reserved. Printed in the U.S. No part of this document covered by copyright may be reproduced in any form or by any means—graphic, electronic, or mechanical, including photocopying, recording, taping, or storage in an electronic retrieval system—without prior written permission of the copyright owner.

Software derived from copyrighted NetApp material is subject to the following license and disclaimer:

THIS SOFTWARE IS PROVIDED BY NETAPP “AS IS” AND WITHOUT ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE, WHICH ARE HEREBY DISCLAIMED. IN NO EVENT SHALL NETAPP BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

NetApp reserves the right to change any products described herein at any time, and without notice.

NetApp assumes no responsibility or liability arising from the use of products described herein, except as expressly agreed to in writing by NetApp. The use or purchase of this product does not convey a license under any patent rights, trademark rights, or any other intellectual property rights of NetApp.

The product described in this manual may be protected by one or more U.S. patents, foreign patents, or pending applications.

RESTRICTED RIGHTS LEGEND: Use, duplication, or disclosure by the government is subject to restrictions as set forth in subparagraph (c)(1)(ii) of the Rights in Technical Data and Computer Software clause at DFARS 252.277-7103 (October 1988) and FAR 52-227-19 (June 1987).

Trademark Information

NETAPP, the NETAPP logo, and the marks listed at <http://www.netapp.com/TM> are trademarks of NetApp, Inc. Other company and product names may be trademarks of their respective owners.