



Technical FAQ

All Flash FAS ONTAP 9.2

Scott Bell, NetApp
Version 1.0
May 2017

Abstract

This document answers frequently-asked technical questions about the All Flash FAS (AFF) A-Series series systems. The FAQ covers the requirements and architecture that define AFF systems, as well as configuration recommendations, performance expectations, sizing guidelines, and interoperability support. AFF A-Series systems are supported starting with the ONTAP 9.1 GA release.

Data Classification

NetApp Confidential—Limited Use

Version History

Version	Date	Document Version History
1.0	May 2017	Initial revision

TABLE OF CONTENTS

Version History	1
List of Tables	5
List of Figures	5
1 Overview	6
How are AFF systems different from FAS systems?	6
Can a FAS system configured with SSDs be converted to an AFF A-Series system?	6
Can an AFF A-Series system be converted into a FAS system?	6
What are the benefits of using an AFF A-Series system?	6
What applications and workloads are best-suited for AFF A-Series systems?	6
What break-fix support is available for the SSDs and other hardware used on AFF systems?	6
2 Requirements	7
What All Flash FAS (AFF) A-Series systems are available?	7
What is the minimum required ONTAP release for AFF A-Series 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	8
How does ONTAP deliver high performance and flash media endurance with AFF systems?	8
Does ONTAP include features that are specifically benefit AFF A-Series systems?	8
What key features are available with AFF A-Series systems running ONTAP 9.1?	8
What new features in ONTAP 9.2 releases benefit AFF A-Series systems?	8
What ONTAP features enable storage efficiency (space savings) on AFF A-Series systems?	9
What is Advanced Drive Partitioning?	9
How is partitioning of root aggregates on AFF A-Series systems enabled?	9
What does the RD2 configuration of an AFF system with RD2 partitioning look like with ONTAP 9?	9
What is the maximum number of SSDs using RD2 partitioning on an AFF A-Series systems?	10
What size are the root partitions in half shelf, single-shelf, 1.5 shelf, and 2 shelf configurations?	10
What type of flash media is used in the SSDs shipped with AFF A-Series systems?	10
How can SSD media usage be monitored on an AFF A-Series system?	11
4 Configuration	12
What are the best practice recommendations for configuring storage on AFF A-Series HA pairs?	12
What is the recommended configuration when the two nodes in an AFF HA pair have different storage capacity requirements?	13
What are the RAID group size recommendations for AFF A-series systems?	13
What is the RAID group configuration recommendation when data partitions and non-partitioned SSDs are provisioned in the same aggregate?	14
What are the hot spare policy recommendations for AFF A-series HA pairs?	14
What are the recommended data aggregate and RAID group configurations new AFF systems?	14

What are the recommended data aggregate and RAID group configurations when adding 12 SSDs to an AFF A-Series system running ONTAP 9.1?	15
What is the maximum number of shelves supported on an AFF A-Series HA pair?	16
Can different capacity SSDs be provisioned in the same aggregate?	16
Can different capacity SSDs be configured in the same shelf?	16
Can shelves with different capacity SSDs be configured in the same stack?	16
What is inline adaptive compression?	16
Is inline adaptive compression enabled on AFF A-Series systems?	17
Can data in a volume that has not been previously compressed inline be adaptively compressed?	17
What is inline volume deduplication?	17
Is inline volume deduplication enabled on AFF A-Series systems?	17
What is always-on deduplication?	17
What is inline zero block deduplication?	17
Is inline zero block deduplication enabled on AFF A-Series systems?	17
What is inline data compaction?	17
Is inline data compaction enabled on AFF A-Series systems?	18
Is thin provisioning of volumes recommended on AFF A-Series systems?	18
What is inline aggregate deduplication?	18
Is inline aggregate deduplication enabled on AFF A-Series systems?	18
What are the configuration recommendations for an AFF A-Series cluster used as a SAN-only storage?	18
5 Performance	19
What is the performance of AFF A-Series systems?	19
How do AFF A-Series systems perform with VMware virtual desktop deployments?	19
How does AFF performance compare to hybrid FAS systems configured with Flash Cache or Flash Pool? ...	19
Are ONTAP reallocation methods recommended with AFF A-Series systems?	19
6 Interoperability	20
Can AFF A-Series HA pairs be configured in a ONTAP cluster that contains FAS HA pairs?	20
Is an AFF HA pair supported in a ONTAP MetroCluster configurations?	20
Can an AFF A-Series system be configured as a NetApp Storage Encryption (NSE) system?	20
What happens with compressed and deduplicated blocks when volmove is used to migrate data non-disruptively from a volume on an AFF node to a volume on a FAS node?	20
Will blocks that are not compressed on a source FAS volume be adaptively compressed inline when volmove is used to transfer data to an AFF volume?	20
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?	20
Is 7MTT supported to move volumes on a 7-Mode FAS system to an AFF A-Series system?	21
Is FlexArray supported with AFF A-Series systems?	21
Is disk sanitization supported for SSDs in an AFF A-Series system?	21
Does Flash Cache work in an AFF A-Series system?	21
7 Additional Resources	22
8 Contact us	23

List of Tables

Table 1) SSDs supported with AFF A-Series systems.	7
Table 2) Major new features for AFF A-Series systems with ONTAP 9.2.....	8
Table 3) AFF A-Series storage efficiency features.	9
Table 4) SSDs supported with AFF A-Series systems.	10
Table 5) Data aggregate RAID group and spare recommendations for new systems running a ONTAP 9 releases ...	14
Table 6) Data aggregate RAID group configurations after adding 24 SSDs to a system running ONTAP 9.1	15
Table 7) Data aggregate RAID group configurations after adding a 24 SSDs to a system running ONTAP 9.2	16
Table 8) Maximum number of shelves supported per HA pair.....	16

List of Figures

Figure 1) AFF A-Series SSD assignment and RD2 partitioning of 12 SSDs	9
Figure 2) AFF A-Series SSD assignment and RD2 partitioning of 24 SSDs	10
Figure 3) AFF A-Series SSD assignment and RD2 partitioning of 48 SSDs with ONTAP 9.2.....	10
Figure 4) A sample output	11
Figure 5) Split shelf SSD assignment of a 96-drive AFF A300 system.....	12
Figure 6) SAS cabling for a two shelf AFF A300 or AFF A700	13
Figure 7) Illustration of space savings with inline data compaction	18

1 Overview

How are AFF systems different from FAS systems?

- Controllers: Must be configured as an HA pair or MetroCluster configuration
- Storage: Only SSDs
- Personality: Flash-optimized personality must be enabled. Enablement is done on new AFF systems at time of manufacturing. In-line storage efficiencies are enabled by default.

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

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

Can an AFF A-Series system be converted into 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 A-Series 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 loads. The key advantages AFF A-Series 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
- non-disruptive data mobility across storage tiers within a cluster that consists of AFF and FAS HA pairs
- data fabric and public cloud integration
- the 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 A-Series 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 standard NetApp 3-year support with 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 All Flash FAS (AFF) A-Series systems are available?

Answer: AFF A200, AFF A300, AFF A700s, and AFF A700.

What is the minimum required ONTAP release for AFF A-Series systems?

Answer: ONTAP 9.1 GA is the minimum, recommended release for all AFF A-Series platforms.

Which AFF systems provide the best performance?

Answer: The AFF A700s and AFF A700 are the best performing models due to the number of CPU cores and system memory.

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

Answer: DS224C with 960GB, 3.8TB, 3.8TB NSE are offered with all four models.

15.3TB are standard offerings on AFF A300, AFF A700s and AFF A700, and also available with PVR approval with the AFF A200. 800GB NSE are configuration on AFF A200 and AFF A300 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) and 1.6TB NSE (X365A-R6).

Table 1) SSDs supported with AFF A-Series 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	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	1.6 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	
X670A-R6	15.3 TB	14,650,368 MiB	DS224C

Note: (1) Supported but not offered when configuring a new AFF A-Series 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 A-Series systems?

Answer: Yes. ONTAP 9.1 releases include performance, storage efficiency, and other features specifically designed for AFF A-Series systems. Some features are also available on FAS systems configured with SSD aggregates, Flash Pool aggregates, or HDD-only aggregates, however AFF A-Series systems get the greatest benefit from these enhancements. Other features are available only with AFF A-Series systems.

What key features are available with AFF A-Series systems running ONTAP 9.1?

Answer:

- Storage efficiency: inline adaptive compression, inline volume deduplication, inline data compaction, root-data-data (RD2) partitioning
- Security: NetApp Volume Encryption (NVE), NetApp Storage Encryption (NSE), Onboard Key Manager (OKM)
- FlexGroup, Quality of Service (bully workload management)

What new features in ONTAP 9.2 releases benefit AFF A-Series systems?

Answer: The major new features in ONTAP 9 releases, and the associated benefits, are list below.

Table 2) Major new features for AFF A-Series systems with ONTAP 9.2

Release	New Feature	Benefit
9.2	Inline aggregate deduplication	Overall storage efficiency
9.2	QoS Minimum Throughput Policy	Performance Management: AFF SAN volumes only
9.2	Fabric Pool	Auto-tiering to lower cost storage
9.2	Shared SSD for data partition spares up to 48 drives	Storage efficiency: ~2.5% more usable capacity

What ONTAP features enable storage efficiency (space savings) on AFF A-Series systems?

Answer: Many ONTAP storage efficiency features help customers increase the effective capacity of AFF A-Series systems. Table 3 summarizes these features.

Table 3) AFF A-Series storage efficiency features.

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

What is Advanced Drive Partitioning?

Answer: ONTAP 9 includes a new drive partitioning scheme called Root-Data-Data (RD2) partitioning. RD2 partitioning builds on the existing Advanced Drive Partitioning feature available in ONTAP 8.3.

The root aggregate of each node of an AFF system uses an equally-sized small area, or thin partition, from each of multiple SSDs. The remaining space is partitioned into 2 equally-sized larger areas, or thick partitions, for provisioning a data aggregate on each node. The primary benefit of RD2 is ≈16% additional storage capacity due to less space consumed by the root aggregates compared to using dedicated SSDs.

How is partitioning of root aggregates on AFF A-Series systems enabled?

Answer: Root-Data-Data (RD2) partitioning is the default method of provisioning root aggregates on AFF A-Series systems. New AFF systems are shipped with root aggregates already provisioned with partitions, and data partitions are assigned to each node as spares. If an AFF system is re-provisioned after destroying root and data aggregates, ADP will be used when the new root aggregates are created.

What does the RD2 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 A-Series systems running ONTAP 9 releases. Figure 1 depicts a single-shelf configuration, figure 2 shows a 2-shelf system and figure 3 shows a half-shelf configuration.

Note: Data (R), Parity (DP and P), and Spare (S) root partitions 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 within 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 A-Series SSD assignment and RD2 partitioning of 12 SSDs with ONTAP 9




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





Figure 3) AFF A-Series 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 A-Series 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: Approximate root partition sizes, per drive quantity, can be found in Table 4.

Table 4) SSDs supported with AFF A-Series systems.

SSDs	Partition Size
12	144 GiB
24	54 GiB
36	31 GiB
48	22 GiB

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

Answer: AFF A-Series 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 run 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 rather 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 A-Series 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 the better indicator of how much of the 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
```

Field Descriptions:

- **Power-on Hours:** This is the number of hours the SSD has been powered on.
- **Glist count:** Also known as the Growth list, which 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 A-Series HA pairs?


Answer: AFF A-Series 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 A-Series 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 less per shelf. This is accomplished by following two important recommendations:

- **12xSSD Shelf Configuration:** Systems shipped from the factory will assign SSDs in shelf bays 0-5 to Node-A, while SSDs in shelf bays 6-11 will be assigned to Node-B.
- **24xSSD Shelf Configuration:** Systems shipped from the factory will assign 12xSSDs in each full shelf to each controller. The SSDs in bays 0-11 will be assigned to Node-A), while SSDs in bays 12-23 will be assigned to Node-B.

This is referred to as “split-shelf” assignment, and it is the default automatic assignment policy for full shelves add to AFF systems beginning with ONTAP 8.3.2. Figure 5 illustrates an example split-shelf configuration of an AFF A300 with 4 external shelves of SSDs.

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




It is recommended that a maximum of 2 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 configuring two shelves per stack with 12 SSDs per SAS channel.

An additional 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. As of 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 recommendation SAS port usage for a 2-shelf configuration.

Figure 6) SAS cabling for a two shelf AFF A300 or AFF A700



What is the recommended configuration when the two nodes in an AFF HA pair have different storage capacity requirements?

Answer: If one node requires more capacity than its partner, the SSD-to-node assignment will be asymmetric in one or more shelves. It is recommended to minimize 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 1 or 2 shelves only; this keeps the number of SSDs per channel at or below 24.

Scenario: An AFF A300 HA pair is configured with 72xSSDs. Node-A requires twice the capacity of Node-B. The recommendation is to split ownership of 2 shelves equally and assign the remaining 24 SSDs to Node-A. The third shelf should be configured alone on a SAS stack. This configuration results in 12xSSDs per SAS channel.

What are the RAID group size recommendations for AFF A-series systems?

Answer: RAID DP groups of up to 23 or 24 data partitions or un-partitioned SSDs are recommended to enable high usable capacity and storage efficiency. The high reliability and fast performance of SSDs enables configuring large RAID DP groups with high confidence. When a new RAID group is added to an existing aggregate, it is recommended 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 non-partitioned SSDs are provisioned in the same aggregate?

Answer: In most cases, provisioning data partitions and non-partitioned SSDs in separate RAID groups is recommended as this 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 non-partitioned SSDs (which are then automatically partitioned) is preferred. This approach provides more usable capacity than creating a new RAID group with non-partitioned SSDs.

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

Answer: For HA pairs configured with 48 or fewer SSDs that has data aggregates provisioned with partitions, one spare data partition per node is recommended. 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, retaining one spare, non-partitioned SSD per node is recommended. A non-partitioned SSD can replace either a data partition on a failed SSD or a failed non-partitioned 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. personnel scheduling, support agreement or shipment time from the nearest spare parts depot), then consider reserving two SSDs as hot spares per node.

If a node is configured with more than 100 total data partitions plus non-partitioned SSDs, then a minimum 2 spares is recommended.

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

Answer: Table 5 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 non-partitioned SSD and *p* signifies a parity partition or non-partitioned SSD. Most of the recommendations assume equal numbers of SSDs are assigned to each node. The 1½ and 2½ shelf examples illustrate unbalanced assignments of SSDs.

Table 5) Data aggregate RAID group and spare recommendations for new systems running a ONTAP 9 releases

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: 21 <i>d</i> partitions + 2 <i>p</i> partitions rg1: 10 <i>d</i> partitions + 2 <i>p</i> partitions	1 data partition
48	rg0: 21 <i>d</i> partitions + 2 <i>p</i> partitions rg1: 22 <i>d</i> partitions + 2 <i>p</i> partitions	1 data partition
60	Node A and B: rg0: 21 <i>d</i> partitions + 2 <i>p</i> partitions rg1: 22 <i>d</i> partitions + 2 <i>p</i> partitions Node A only: rg2: 9 <i>d</i> + 2 <i>p</i> non-partitioned SSDs	Node A: 1 data partition 1 non-partitioned SSD Node B: 1 data partition

SSDs	Data Aggregate RAID Groups per Node	Hot Spare per Node
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> non-partitioned SSDs	1 non-partitioned SSD

What are the recommended data aggregate and RAID group configurations when adding 12 SSDs to an AFF A-Series system running ONTAP 9.1?

Answer: When adding a full shelf to an AFF A-Series system, RAID group recommendations are similar to those for provisioning RAID groups on new systems. However, there are additional considerations.

A general guideline is to provision data partitions and non-partitioned SSDs in different RAID groups to avoid creating root partitions on non-partitioned SSDs that are only usable in root aggregates. There can be situations where doing exactly that is recommended, however, because it provides more usable capacity than creating a new RAID group of non-partitioned SSDs.

Additionally, a new RAID group should have a minimum of 11 SSDs. If expanding an existing RAID group, at least 6 SSDs should be added at once.

Table 6) 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 a 12 SSDs
12	rg0: 9 <i>d</i> partitions + 2 <i>p</i> partitions	rg0: 21 <i>d</i> partitions + 2 <i>p</i> partitions ¹
24	rg0: 21 <i>d</i> partitions + 2 <i>p</i> partitions	Node A and B : rg0: 21 <i>d</i> partitions + 2 <i>p</i> partitions Node A only : rg1: 9 <i>d</i> + 2 <i>p</i> non-partitioned SSDs
36	rg0: 21 <i>d</i> partitions + 2 <i>p</i> partitions rg1: 10 <i>d</i> partitions + 2 <i>p</i> partitions	rg0: 21 <i>d</i> partitions + 2 <i>p</i> partitions rg1: 22 <i>d</i> partitions + 2 <i>p</i> partitions ¹
48	rg0: 21 <i>d</i> partitions + 2 <i>p</i> partitions rg1: 22 <i>d</i> partitions + 2 <i>p</i> partitions	Node A and B: rg0: 21 <i>d</i> partitions + 2 <i>p</i> partitions rg1: 22 <i>d</i> partitions + 2 <i>p</i> partitions Node A only: rg2: 9 <i>d</i> + 2 <i>p</i> non-partitioned SSDs

Note: (1) The 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 only usable by the root aggregates.

Table 7) 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: 21 <i>d</i> partitions + 2 <i>p</i> partitions	rg0: 21 <i>d</i> partitions + 2 <i>p</i> partitions rg1: 9 <i>d</i> + 2 <i>p</i> non-partitioned SSDs
48	rg0: 21 <i>d</i> partitions + 2 <i>p</i> partitions rg1: 22 <i>d</i> partitions + 2 <i>p</i> partitions	rg0: 21 <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> non-partitioned SSDs

Note: If a node has one or more spare data partitions, and a new RAID group of non-partitioned SSDs will be added to an existing data aggregate, select only non-partitioned SSD spares to create the new RAID group. If a spare data partition is included when creating the RAID group, non-partitioned SSDs will be automatically partitioned to 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 based on system model, shelf type, and configuration. Table 8 provides the details.

Table 8) Maximum number of shelves supported per HA pair.

AFF Model/Configuration	Maximum External Shelves
AFF A200	5
AFF A300	16
AFF A700s	8
AFF A700	20

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

Answer: Yes. To keep shelf configurations simple, it is recommended to configure no more than two different SSD capacity models per shelf 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 (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 A-Series systems.

Is inline adaptive compression enabled on AFF A-Series systems?

Answer: Yes. All AFF A-Series systems enable inline adaptive compression by default on all volumes as they are created. Inline adaptive compression is not enabled automatically on volumes created with the volume move process. Inline adaptive compression can be enabled/disabled by modifying the volume efficiency policy.

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

Answer: Yes. Starting with ONTAP 8.3.2, a background adaptive compression scanner 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 snapshots, `-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 via 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 VMs with very similar content that reside in the same volume.

Is inline volume deduplication enabled on AFF A-Series systems?

Answer: Yes. All AFF A-Series systems enable inline deduplication by default on all volumes as they are created. Inline deduplication can be enabled/disabled by modifying the volume efficiency policy.

What is always-on deduplication?

Answer: Always-on deduplication is normal background deduplication with volume dedupe 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 contains all zeroes. Rather than writing a block of zeroes to storage, metadata is updated to reflect that a block contains only zeroes. 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 A-Series systems?

Answer: Yes. All AFF A-Series systems have inline zero block deduplication enabled by default on all volumes created.

What is inline data compaction?

Answer: Inline data compaction detects incoming I/O sizes less than 4KB blocks and attempts to group multiple small writes together within 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) Illustration of space savings with inline data compaction



Is inline data compaction enabled on AFF A-Series systems?

Answer: Yes. All AFF A-Series systems have inline data compaction enabled by default on all thin-provisioned volumes created. Inline data compaction can be enabled/disabled by modifying the volume efficiency policy.

Is thin provisioning of volumes recommended on AFF A-Series systems?

Answer: Yes. The recommendation is 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?

Answer: Inline aggregate deduplication, or cross-volume inline 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 already exist within the aggregate. Fingerprint and byte-for-byte comparison are done in memory before incoming data is committed to storage via a consistency point (CP). Inline aggregate deduplication was introduced with ONTAP 9.2. The primary benefit is lower capacity requirements when patching or moving multiple VMs with very similar content that are provisioned in volumes contained within the same aggregate.

Is inline aggregate deduplication enabled on AFF A-Series systems?

Answer: Yes. All AFF A-Series systems running ONTAP 9.2 have inline data compaction enabled by default on all thin-provisioned volumes created. Inline aggregate deduplication can be enabled/disabled per volume by modifying the volume efficiency policy.

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

Answer: The [FC SAN Optimized AFF Setup Guide](#) explains how to complete configuration of factory-configured SAN optimized AFF systems. FC SAN optimized systems are an orderable configuration.

Additional information can be found in [TR-4480: All Flash FAS SAN-Optimized Configuration Details](#).

5 Performance

What is the performance of AFF A-Series systems?

Answer: The best way to determine the performance capability of an AFF A-Series HA pair or multi-node cluster is to use the System Performance Modeler, or SPM. SPM v2.5.5 and later releases, model random workloads on AFF A-Series systems at low latencies. Latencies from 0.5 ms to 3 ms, done in 0.5 ms increments, are supported. Inline compression, inline deduplication, and advanced drive partitioning (ADP) for root aggregates are also supported.

How do AFF A-Series systems perform with VMware virtual desktop deployments?

Answer: Technical report 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 A-Series systems. The report includes performance and storage efficiency results, and the link to the report is in the Resources section.

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 will be similarly fast, but reads from HDDs will be slower as 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 will be 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 will generally be less for a workload on an AFF system as compared to when it is provisioned on FAS with cache acceleration.

Are ONTAP reallocation methods recommended with AFF A-Series 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 due to NAND flash technology. In addition, AFF systems are expected to be very busy and aggregates will often be relatively small and quite full. Thus, enabling read reallocate on AFF volumes or LUNs, and free space reallocation on SSD aggregates will provide minimal, if any, benefits.

If a volume's data is randomly written and updated, and subsequently read sequentially, there may be benefit in running the reallocate command periodically during periods of relatively low workload on the system. The command to is:

```
reallocate start -vserver <svm name> -force true -space-optimized -path <path>
```

6 Interoperability

Can AFF A-Series HA pairs be configured in a 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 a ONTAP MetroCluster configurations?

Answer: Yes. The AFF A300 and AFF A700 platforms support up to an 8-Node Fabric Attached MetroCluster configuration. A200 and A700s are not supported in MetroCluster systems.

Note: Root aggregates on AFF nodes configured in a MetroCluster are provisioned with dedicated SSDs. Root aggregate partitioning is not supported with MetroCluster.

Can an AFF A-Series system be configured as a NetApp Storage Encryption (NSE) system?

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

Note: NSE drives are not supported with clustered ONTAP MetroCluster (MCC) systems. FAS controllers with FlexArray enabled, configured with EF560 and self-encrypting SSDs is the recommended encryption solution for MCC all-flash systems.

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

Answer: Volmove 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 will be similarly compressed on the destination volume. In addition, any deduplicated blocks will remain deduped on the destination volume.

Will blocks that are not compressed on a source FAS volume be adaptively compressed inline when volmove is used to transfer data to an AFF volume?

Answer: No. Volmove uses physical block replication to move data. Blocks from the source volume are copied as-is, without any change to format, to the destination AFF volume. Inline adaptive compression is not applied during volmove, even when that is the policy on the destination volume.

Starting with ONTAP 8.3.2, volume data that is transferred from a FAS system to an AFF system using volmove can be adaptively compressed after blocks have been written to the AFF volume, by running background adaptive compression one time per volume. It is recommended to run background adaptive compression on one volume at a time when an AFF system is serving application or end-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 both during transfer and on the destination. When the policies are different, compression and deduplication savings 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 7MTT supported to move volumes on a 7-Mode FAS system to an AFF A-Series system?

Answer: Yes. However, inline compression and deduplication policies that are enabled on a destination volume are not applied to data that are transferred using 7MTT. 7MTT uses physical block replication technology to move data. Compressed and deduplicated blocks on the source 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. It is recommended to run background adaptive compression on only one volume at a time while an AFF system is serving host/client applications. Background deduplication can be also be run to dedupe transferred data.

Is FlexArray supported with AFF A-Series systems?

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

Is disk sanitization supported for SSDs in an AFF A-Series system?

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

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

Does Flash Cache work in an AFF A-Series system?

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

The AFF A-Series, by default, does not come with Flash Cache cards.

7 Additional Resources

All Flash FAS 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-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?assetComponentId=324953>
- Storage Subsystem Technical FAQ <https://fieldportal.netapp.com/Core/DownloadDoc.aspx?documentID=75058&contentID=81763>

Presentations

- AFF Technical Content Collection <https://fieldportal.netapp.com/collections/500339?ia=true>

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.