



Technical FAQ

All Flash FAS8000

Data ONTAP 8.3 GA and later releases

Skip Shapiro, NetApp
Version 1.1
April 2016

Abstract

This document answers frequently-asked technical questions about All Flash FAS (AFF) 8000 series systems. The FAQ covers the requirements and architecture that define AFF8000 systems, as well as configuration recommendations, performance expectations and sizing guidelines, and interoperability support. AFF8000 systems are supported starting with the clustered Data ONTAP 8.3 GA release. AFF8000 systems are not supported with Data ONTAP operating in 7-Mode.

Data Classification

NetApp Confidential—Limited Use

Version History

Version	Date	Document Version History
1.1	April 2016	Configuration figures updated and added. Data aggregate RAID group recommendation tables updated and added. Additional questions added.
1.0.2	Jan/Feb 2016	Initial AFF8000 FAQ release plus minor corrections.

TABLE OF CONTENTS

Version History	1
List of Tables	5
List of Figures	5
1 Overview.....	6
What is an All Flash FAS (AFF) system?.....	6
Is a FAS8000 system that is configured with SSDs only and running Data ONTAP 8.2.x considered to be an AFF system?	6
When a FAS8000 system configured with SSDs only is upgraded to Data ONTAP 8.3.x does it automatically become an AFF8000 system?	6
Can a FAS8000 system configured with SSDs only be converted to an AFF8000 system?.....	6
Can an AFF system be deployed as part of a Data ONTAP cluster that also has FAS HA pairs?	6
Can an AFF8000 system be converted into a FAS8000 system?	6
What are the benefits of using an All Flash FAS (AFF) system?	6
What applications and workloads are best-suited for AFF8000 systems?	7
What break-fix support is available for the SSDs and other hardware used on AFF systems?	7
2 Requirements.....	8
What is the minimum required Data ONTAP release for All Flash FAS (AFF) systems?.....	8
Which controller models are supported on All Flash FAS (AFF) systems?.....	8
Which AFF8000 systems provide the best performance?	8
Are special All Flash FAS (AFF) 8000 system configurations offered for sale?	8
What SSDs and storage enclosures are supported with AFF8000 systems?.....	8
3 Architecture.....	9
How does Data ONTAP deliver high performance and flash media endurance with AFF systems?	9
Does Data ONTAP include features that are specifically designed to benefit AFF8000 systems?	9
What new features in Data ONTAP 8.3 releases benefit AFF8000 systems?	9
What Data ONTAP features enable storage efficiency (space savings) on AFF8000 systems?	10
What is advanced drive partitioning (ADP) of root aggregates?	10
How is partitioning of root aggregates on AFF8000 systems enabled?	10
What does the storage configuration of an AFF8000 system with root aggregate partitioning look like?	10
.....	11
What is the maximum number of SSDs partitioned for root aggregates on an AFF8000 system?	11
What size are the root partitions in the single-shelf, 1½-shelf and 2-shelf AFF8000 configurations?	11
What type of flash media is used in the SSDs shipped with AFF8000 systems?	11
How can SSD media usage be monitored on an AFF8000 system?	12
4 Configuration	13
What are the best practice recommendations for configuring storage on AFF8000 HA pairs?.....	13
What is the recommended configuration when the two nodes in an AFF HA pair have different storage capacity requirements?.....	14
What are the RAID group size recommendations for AFF8000 systems?	14

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 AFF8000 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 a full shelf (24 SSDs) to AFF systems?	15
What are the recommended data aggregate and RAID group configurations when adding a half- shelf (12 SSDs) to AFF systems?	16
What is the maximum number of shelves supported on an AFF8000 HA pair?	17
Can different capacity SSDs be provisioned in the same aggregate?	17
Can different capacity SSDs be configured in the same shelf?	17
Can shelves with different capacity SSDs be configured in the same stack?	17
What is the default RAID parity scrub schedule for AFF8000 systems?	17
What is adaptive compression?	18
How is inline adaptive compression enabled on AFF8000 systems?	18
Can data in a volume that has not been previously compressed inline be adaptively compressed?	18
What is inline deduplication?	18
How is inline deduplication enabled on AFF8000 systems?	18
What is always-on deduplication?	18
How is always-on deduplication implemented on AFF8000 systems?	18
What is inline zero block deduplication?	19
How is inline zero block deduplication enabled on AFF8000 systems?	19
Is thin provisioning of volumes recommended on AFF8000 systems?	19
What are the configuration recommendations for an AFF8000 cluster used as a SAN-only storage?	19
5 Performance	20
What is the performance of AFF8000 systems?	20
How do AFF8000 systems perform with Oracle databases?	20
How do AFF8000 systems perform with Microsoft SQL Server databases?	20
How do AFF8000 systems perform with VMware virtual desktop deployments?	20
Do AFF8000 systems perform better with Data ONTAP 8.3.1 or 8.3.2 than with Data ONTAP 8.3?	20
How does the performance of AFF compare to hybrid FAS systems configured with Flash Cache or Flash Pool aggregates?	20
Are Data ONTAP reallocation methods recommended with AFF8000 systems?	20
6 Interoperability	22
Can AFF8000 HA pairs be configured in a Data ONTAP cluster that contains FAS HA pairs?	22
Is an AFF HA pair supported in a clustered Data ONTAP MetroCluster configurations?	22
Can an AFF8000 HA pair be configured as a NetApp Storage Encryption (NSE) system?	22
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?	22
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?	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?	22

Is 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 AFF8000 systems.	8
Table 2) Major new features for AFF8000 systems in Data ONTAP 8.3 releases.....	9
Table 3) AFF8000 storage efficiency features.	10
Table 4) Per node data aggregate RAID group and hot spare recommendations for new systems.	15
Table 5) Data aggregate RAID group configurations after adding one full shelf.	16
Table 6) Data aggregate RAID group configurations after adding one half-shelf.	16
Table 7) Maximum number of shelves supported per AFF8000 HA pair.....	17

List of Figures

Figure 1) Single-shelf AFF8000 SSD assignment and partitioning.....	10
Figure 2) Two-shelf SSD assignment and root-data partitioning.	11
Figure 3) One-and-a-half shelf SSD assignment and root-data partitioning.	11
Figure 4) A sample storage disk show -a output	12
Figure 5) SAS cabling for a two shelf AFF8040, AFF8060 or AFF8080 EX system	13

1 Overview

What is an All Flash FAS (AFF) system?

Answer: Four characteristics must be true for a system to be an All Flash FAS system:

- Controllers: Must be 8000 series (8080 EX, 8060, 8040 or 8020) and configured as a HA pair
- Storage: SSDs only
- Data ONTAP: 8.3 GA or a later release
- Flash optimized: the flash-optimized personality must be enabled, which is done on new AFF systems as they are shipped.

If one of the four criteria is not satisfied, then the system is not an AFF system.

Is a FAS8000 system that is configured with SSDs only and running Data ONTAP 8.2.x considered to be an AFF system?

Answer: No. The proper term for this is a FAS8000 with SSD storage system; it is not an AFF system.

When a FAS8000 system configured with SSDs only is upgraded to Data ONTAP 8.3.x does it automatically become an AFF8000 system?

Answer: No. Upgrading a FAS8000 with SSD system to Data ONTAP 8.3.x does not enable the flash-optimized personality so the system remains a FAS with SSD system.

Can a FAS8000 system configured with SSDs only be converted to an AFF8000 system?

Answer: No. Conversion of a FAS8000 system to an AFF8000 system is not generally supported. A Policy Variation Request (PVR) for one-time conversion of a FAS8000 system configured with SSDs only—that was sold as “all-flash FAS” configuration and shipped between June 2014 and June 2015—can be submitted for approval. Approval is required before converting the system and for it to be supported as an AFF8000 system.

Can an AFF system be deployed as part of a Data ONTAP cluster that also has 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.

Can an AFF8000 system be converted into a FAS8000 system?

Answer: No. Converting an AFF8000 system into a FAS8000 system is not supported. Customers who need the flexibility of using both SSDs and HDDs in the same HA pair should purchase a FAS system.

What are the benefits of using an All Flash FAS (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 loads. The key advantages AFF8000 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
- the rich enterprise data management features of the NetApp® Data 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 AFF8000 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 AFF8000 hardware is covered by standard NetApp 3-year support that is included with new system purchases. Customers have the option of purchasing extended support for up to an additional 4 years – for a total of 7 years of supported – that covers all hardware, including SSDs. There are no 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.

There is also a promotion for that provides customers the option to upgrade AFF controller hardware at no cost when extended support for years 4 through 6 is purchased. More information about the promotion is available in the [AFF Controller Upgrade Promo Brief](#).

2 Requirements

What is the minimum required Data ONTAP release for All Flash FAS (AFF) systems?

Answer: Data ONTAP 8.3 GA is the minimum required release. AFF8000 systems are only supported with clustered Data ONTAP.

Which controller models are supported on All Flash FAS (AFF) systems?

Answer: AFF8080 EX, AFF8060, AFF8040 and AFF8020 only. No other controllers are supported.

Which AFF8000 systems provide the best performance?

Answer: AFF8080 EX and AFF8060 are the best performing models because of the number and clock-speed of the CPU cores. These are the models to position and sell to customers when competing for most AFF opportunities.

Are special All Flash FAS (AFF) 8000 system configurations offered for sale?

Answer: Yes. Starting with the June 2015 quote tool, two AFF8080 EX pre-configured systems (bundles) are offered. See the AFF product bulletin and sales FAQ on the AFF Field Portal landing page for more information.

What SSDs and storage enclosures are supported with AFF8000 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. DS2246 is the only supported storage enclosure for SSDs used with AFF8000 systems.

Table 1) SSDs supported with AFF8000 systems.

SSD Marketing Part Number	Marketed Capacity	Right-sized Used Capacity	Supported Storage Enclosure
X446A-R6 ¹ X446B-R6 ¹	200GB	190,532 MiB	DS2246
X438A-R6	400GB	381,304 MiB	
X447A-R6	800GB	762,847 MiB	
X440A-R6	800GB NSE	762,847 MiB	
X439A-R6	1.6TB	1,525,935 MiB	
X365A-R6	1.6TB NSE	1,525,935 MiB	
X356A-R6	3.8TB	3,662,580 MiB	

Note: (1) Supported but not offered when configuring a new AFF8000 system.

3 Architecture

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

Answer: Data 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 Data ONTAP 8.3 and 8.3.1 dramatically increased random read performance (both IOPS and latency), while adding inline storage efficiency capabilities.

Streamlined read paths, and read-ahead algorithms, in Data ONTAP keep read latencies low. Write operations and SSD firmware are tightly couple and data layouts align are well aligned, delivering 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 Data 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 Data ONTAP include features that are specifically designed to benefit AFF8000 systems?

Answer: Yes. Data ONTAP 8.3.x releases include performance, storage efficiency, and other features specifically designed for AFF8000 systems. Some features are also available on FAS systems configured with SSD aggregates, Flash Pool aggregates, or HDD-only aggregates, however AFF8000 systems get the greatest benefit from these enhancements. Other features are available only with AFF8000 systems.

What new features in Data ONTAP 8.3 releases benefit AFF8000 systems?

Answer: The major new features in Data ONTAP 8.3 releases, and the associated benefits, are list below.

Table 2) Major new features for AFF8000 systems in Data ONTAP 8.3 releases

Release	New Feature	Benefit
8.3	Read fast path	Performance: Random read IOPS and latency
8.3	Advanced drive partitioning (ADP)	Storage efficiency: root-data aggregate provisioning
8.3	Inline zero block deduplication (Enabled by default with 8.3.1)	Faster virtual machine (VM) and LUN provisioning
8.3.1	Read fast path to storage	Performance: Random read IOPS and latency
8.3.1	Inline adaptive compression	Storage efficiency
8.3.1	Storage driver and write optimizations	Performance: Random read and write latencies
8.3.2	Inline deduplication	Storage efficiency with software patching and storage vmotion of many similar VMs in a volume
8.3.2	Auto assignment policy for AFF	Assigns SSDs in bays 0–11 to one node and SSDs in bays 12–23 to the partner node when shelves are added to an AFF system.
8.3.2	One-time background adaptive compression scan	Enables adaptive compression of volumes moved to AFF via 7MTT or non-disruptive volmove, and AFF volumes created with Data ONTAP 8.3

What Data ONTAP features enable storage efficiency (space savings) on AFF8000 systems?

Answer: Many Data ONTAP storage efficiency features help customers increase the effective capacity of AFF8000 systems. Table 3 summarizes these features.

Table 3) AFF8000 storage efficiency features.

Standard Data ONTAP Features	All Flash FAS Default Features
RAID-DP	Advanced drive partitioning of root aggregates
Thin provisioning	Inline adaptive compression
Snapshot® copies	Inline and always-on deduplication
FlexClone® technology	Inline zero block deduplication

What is advanced drive partitioning (ADP) of root aggregates?

Answer: Advanced drive partitioning (ADP) provisions the root aggregate of each node of an AFF system using an equal-sized small area from each of multiple SSDs. Each partitioned SSD has a larger partition that is used to provision a data aggregate on each node. The primary benefit of ADP is more storage capacity because less space is consumed for the root aggregates compared to using dedicated SSDs.

How is partitioning of root aggregates on AFF8000 systems enabled?


Answer: Advanced drive partitioning (ADP) is the default method of provisioning root aggregates on AFF8000 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 is used when the new root aggregates are created.

What does the storage configuration of an AFF8000 system with root aggregate partitioning look like?

Answer: Figures 1, 2 and 3 are pictorial representations of SSD partitioning on as-shipped AFF8000 systems. Figure 1 depicts a single-shelf configuration, figure 2 shows a 2-shelf system and figure 3 shows a 1½ shelf configuration. One half the SSDs in each shelf are assigned to each controller (node) applying the recommendation for best performance.

Each numbered rectangle represents a physical SSD in a DS2246 shelf bay. The small square within each SSD is a root partition, and the larger rectangle is the data partition. 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). All data partitions are spares on as-shipped and initialized systems.


Figure 1) Single-shelf AFF8000 SSD assignment and partitioning.



Note: (1) Assignment of data (R), parity (DP and P) and spare (S) root partitions to SSD numbers may differ. A typical example is shown.

Note: (2) Root partition size is approximately 54 GiB (gibibytes).


Figure 2) Two-shelf SSD assignment and root-data partitioning.



Note: (1) Assignment of data (R), parity (DP and P) and spare (S) root partitions to SSD numbers may differ. A typical example is shown.

Note: (2) Root partition size is approximately 22 GiB (gibibytes).

Figure 3) One-and-a-half shelf SSD assignment and root-data partitioning.



Note: (1) Assignment of data (R), parity (DP and P) and spare (S) root partitions to SSD numbers may differ. A typical example is shown.

Note: (2) Root partition size is approximately 31 GiB (gibibytes).

What is the maximum number of SSDs partitioned for root aggregates on an AFF8000 system?

Answer: Up to 48 SSDs can be partitioned on an HA pair—24 SSDs per node.

What size are the root partitions in the single-shelf, 1½-shelf and 2-shelf AFF8000 configurations?

Answer: Approximate root partition sizes for the three system configurations are:

- Single-shelf: 54 GiB
- 1½-shelf: 31 GiB
- 2-shelf: 22 GiB

What type of flash media is used in the SSDs shipped with AFF8000 systems?

Answer: AFF8000 systems use eMLC and cMLC SSDs, and with the introduction of 3.8TB SSDs in November 2015, 3D NAND TLC drives are supported. 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 useful way of determining SSD fitness for enterprise storage system use than whether the SSDs used are eMLC, cMLC, TLC or some other NAND flash technology. In addition, 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 AFF8000 system?

Answer: SSD media can be monitored using the node-level command `storage disk show -a` command 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 1 followed by descriptions of some of the fields.

Figure 4) A sample storage disk show -a 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.
- **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 Data 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 Data ONTAP currently shows the spare-blocks-consumed limit as N/A.

4 Configuration

What are the best practice recommendations for configuring storage on AFF8000 HA pairs?

Answer: AFF8000 systems are deployed to deliver low latency for random read-write workloads over a wide range of IOPS. Consequently, best practice AFF8000 configuration recommendations and support limits differ from those that apply to hybrid or HDD-only FAS systems.


To achieve best performance from an AFF8000 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. This is accomplished by following two important recommendations:

- Assign one-half of the 24 SSDs in each **full** shelf to each controller. The SSDs in bays 0 through 11 are assigned to one node (call it Node A) and the SSDs in bays 12 through 23 are assigned to the partner node (Node B). This is called “split-shelf” SSD assignment, and it is the automatic assignment policy for full shelves add to AFF8000 systems starting with Data ONTAP 8.3.2.
- Configure each SAS stack with 1 or 2 shelves whenever possible. Up to 4 shelves per stack are supported, however no more than 2 shelves per stack are recommended 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. Four shelves per stack results in 24 SSDs per SAS channel; this can result in higher latencies 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.

One 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. 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. The figure below illustrates the recommendation SAS port usage for a 2-shelf configuration.

Figure 5) SAS cabling for a two shelf AFF8040, AFF8060 or AFF8080 EX system



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 to each 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.

For example, if an AFF8000 HA pair is configured with three shelves (72 SSDs total), and one node requires twice the capacity of its partner, the recommendation is to split ownership equally on two of the three shelves and assign all 24 SSDs in the third shelf to one node. The third shelf should be configured alone on a SAS stack; that results in 12 SSDs per SAS channel.

What are the RAID group size recommendations for AFF8000 systems?

Answer: Large RAID-DP groups—typically 23 or 24 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 (in terms of 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 AFF8000 HA pairs?

Answer: For a HA pair configured with 48 or fewer SSDs and that has data aggregates provisioned with partitions only, one spare data partition per node that has a data aggregate is recommended. (In an active-passive HA pair, only one node is configured with a data aggregate and only that node needs a data partition spare.)

For a HA pair configured with three or more shelves, one spare non-partitioned SSD per node is recommended. A non-partitioned SSD can replace 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 (for example due to personnel scheduling, support agreement or shipment time from the nearest spare parts depot), then consider reserving two SSDs as hot spares.

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

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

Answer: The table below shows the recommended RAID group configurations for the data aggregate on each node, including the minimum hot spare recommendations, for configurations with up to eight shelves. 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 show unequal assignments of SSDs.

Data aggregate RAID groups for systems configured with more than eight shelves follow the same pattern as shown in the table for 7- and 8-shelf configurations when a RAID group size of 24 is used.

Table 4) Per node data aggregate RAID group and hot spare recommendations for new systems.

Shelves	Data Aggregate RAID Groups per Node	Hot Spare per Node
1	rg0: 9d partitions + 2p partitions	1 data partition
1½	rg0: 15d partitions + 2p partitions or Node A rg0: 21d partitions + 2p partitions Node B rg0: 9d partitions + 2p partitions	1 data partition 1 data partition 1 data partition
2	rg0: 21d partitions + 2p partitions	1 data partition
2½	Node A rg0: 22d partitions + 2p partitions rg1: 9d + 2p non-partitioned SSDs Node B rg0: 21d partitions + 2p partitions	1 non-partitioned SSD 1 data partition
3	rg0: 22d partitions + 2p partitions rg1: 9d + 2p non-partitioned SSDs	1 non-partitioned SSD
3½	rg0: 22d partitions + 2p partitions rg1: 15d + 2p non-partitioned SSDs	1 non-partitioned SSD
4	rg0: 22d partitions + 2p partitions rg1: 21d + 2p non-partitioned SSDs	1 non-partitioned SSD
5	rg0: 22d partitions + 2p partitions rg1: 22d + 2p non-partitioned SSDs rg2: 9d + 2p non-partitioned SSDs	1 non-partitioned SSD
6	rg0: 22d partitions + 2p partitions rg1: 22d + 2p non-partitioned SSDs rg2: 21d + 2p non-partitioned SSDs	1 non-partitioned SSD
7	rg0: 22d partitions + 2p partitions rg1: 22d + 2p non-partitioned SSDs rg2: 22d + 2p non-partitioned SSDs rg3: 9d + 2p non-partitioned SSDs	1 non-partitioned SSD
8	rg0: 22d partitions + 2p partitions rg1: 22d + 2p non-partitioned SSDs rg2: 22d + 2p non-partitioned SSDs rg3: 21d + 2p non-partitioned SSDs	1 non-partitioned SSD

What are the recommended data aggregate and RAID group configurations when adding a full shelf (24 SSDs) to AFF systems?

Answer: Data aggregate RAID group recommendations when adding a full shelf to an installed AFF8000 system are similar to those for provisioning RAID groups on new systems. That said, 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. However, there are situations where doing exactly that is recommended because it provides more usable capacity than creating a new RAID group of non-partitioned SSDs.

In addition, a new RAID group should have a minimum of 11 SSDs and if an existing RAID group is expanded at least 6 SSDs should be added at once.

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

Table 5) Data aggregate RAID group configurations after adding one full shelf.

Starting Shelf Quantity	Original Data Aggregate RAID Group Configuration per Node	RAID Group Configuration per Node After Adding a Full Shelf
1	rg0: 9d partitions + 2p partitions	rg0: 21d partitions + 2p partitions ¹
2	rg0: 21d partitions + 2p partitions	rg0: 21d partitions + 2p partitions rg1: 9d + 2p non-partitioned SSDs
3	rg0: 22d partitions + 2p partitions rg1: 9d + 2p non-partitioned SSDs	rg0: 22d partitions + 2p partitions rg1: 21d + 2p 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.

What are the recommended data aggregate and RAID group configurations when adding a half-shelf (12 SSDs) to AFF systems?

Answer: The existing RAID group configurations influence whether it is recommended to add six SSDs to each node or to add all 12 SSDs to one node. If assigning six SSDs to each node will result in a RAID group that has only 5 or 6 SSDs, the recommendation is to add all 12 SSDs to one node. Half shelves typically ship with SSDs installed in bays 0–11; if six SSDs are added to each node, it is recommended to move the SSDs that will be assigned to Node B to drive bays on the right side of the shelf (bays 12–23). This maintains the split-shelf (left-right) assignment policy for AFF systems and allows the desired assignment of the next 12 SSDs that are added to be done automatically.

Table 6) Data aggregate RAID group configurations after adding one half-shelf.

Starting Shelf Quantity	Original Data Aggregate RAID Group Configuration per Node	RAID Group Configuration per Node After Adding a Half Shelf
1	rg0: 9d partitions + 2p partitions	rg0: 15d partitions + 2p partitions ¹
1½	rg0: 15d partitions + 2p partitions	rg0: 21d partitions + 2p partitions ¹
2	rg0: 21d partitions + 2p partitions	rg0: 21d partitions + 2p partitions rg1: 9d + 2p SSDs (Node A only) ²
2½	Node A rg0: 22d + 2p partitions rg1: 9d + 2p SSDs Node B rg0: 21d + 2p partitions	Node A rg0: 22d + 2p partitions rg1: 9d + 2p SSDs Node B rg0: 21d + 2p partitions rg1: 9d + 2p SSDs
3	rg0: 22d partitions + 2p partitions rg1: 9d + 2p non-partitioned SSDs	rg0: 22d + 2p partitions rg1: 15d + 2p non-partitioned SSDs

Note: (1) The 6 SSDs added to RAID group rg0 in each node's data aggregate are partitioned automatically, and the resulting root partitions become spares that are not usable by data aggregates.

Note: (2) Adding all 12 SSDs to Node A is recommended, rather than assigning six SSDs to each node which would result in a $3d + 2p$ RAID group on each node. One spare non-partitioned SSD is needed for each node.

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

Answer: The maximum number of shelves supported per AFF8000 HA pair varies based on system model and configuration. Table 5 provides the details.

Table 7) Maximum number of shelves supported per AFF8000 HA pair.

AFF Model/Configuration	Maximum Number of Shelves
AFF8020	10
AFF8040	16
AFF8060 A (2 controllers per chassis)	16
AFF8060 AE (1 controller + IOXM per chassis)	20
AFF8080 A (2 controllers per chassis)	16
AFF8080 AE (1 controller + IOXM per chassis)	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 in a shelf and to group like capacity SSDs together.

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

Answer: Yes.

What is the default RAID parity scrub schedule for AFF8000 systems?

With Data ONTAP 8.3.1, the default weekly RAID parity scrub schedule for AFF8000 systems is six (6) hours per week starting at 1 am on Sunday.

AFF systems running 8.3 GA releases have a default RAID parity scrub schedule with daily 4-hour scrubs. It is recommended that this be changed to the same schedule as the 8.3.1 default. You can do that using the following command:

```
storage raid-options modify -node <nodename> -name raid.scrub.schedule  
-value 6@sun@1
```

The schedule change can be confirmed by running the command:

```
storage raid-options show raid.scrub.schedule
```

What is 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 Data ONTAP 8.3.1, and it is supported as an inline process only with AFF8000 systems.

How is inline adaptive compression enabled on AFF8000 systems?

Answer: AFF8000 systems running Data ONTAP 8.3.1 (and later releases) enable inline adaptive compression by default on all volumes as they are created. However, inline adaptive compression is not enabled automatically on volumes created with the volume move (volmove) process. Inline adaptive compression can be enabled on volumes that were created prior to Data ONTAP 8.3.1 by changing the volume efficiency policy.

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

Answer: Yes. Starting with Data 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 deduplication?

Answer: 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 were recently written. Fingerprint and byte-for-byte comparison are done in memory before incoming data is committed to storage via a consistency point (CP). Inline deduplication is supported starting with Data ONTAP 8.3.2, and the main benefit is lower capacity requirements when patching or moving many VMs with very similar content that are provisioned in the same volume.

How is inline deduplication enabled on AFF8000 systems?

Answer: AFF8000 systems running Data ONTAP 8.3.2 (and later releases) enable inline deduplication by default on all volumes as they are created. Inline deduplication can be enabled on volumes that were created prior to Data ONTAP 8.3.2 by changing 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.

How is always-on deduplication implemented on AFF8000 systems?

Answer: Always-on deduplication is implemented by enabling deduplication on a volume and scheduling deduplication jobs to run as often as every 1 to 5 minutes. Always-on dedupe has been tested with, and is intended for use on, volumes that contain virtual machines (VMs) for end-user computing (virtual desktop) operating system and application data. The number of concurrent deduplication jobs to allow should be considered because running multiple jobs simultaneously can impact performance.

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 only is updated to reflect which that a block contains only zeroes. The principal benefit of inline zero block deduplication is faster provisioning of virtual machines and LUNs that are created with all zero blocks. Reduced write workload is a secondary benefit.

How is inline zero block deduplication enabled on AFF8000 systems?

Answer: AFF8000 systems running Data ONTAP 8.3.1 have inline zero block deduplication enabled by default on all volumes created. For systems running Data ONTAP 8.3 and for volumes created with 8.3, inline zero block deduplication is enabled by enabling deduplication; if running background deduplication jobs is not planned, then no schedule is set and only zero block deduplication will occur.

Is thin provisioning of volumes recommended on AFF8000 systems?

Answer: Yes. The recommendation is to thin provision all volumes, that is volumes that contain LUNs and also volumes that contain files. Starting with the Data ONTAP 8.3.2 release, this is the default policy for newly-created volumes.

What are the configuration recommendations for an AFF8000 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 AFF8000 systems?

Answer: The best way to determine the performance capability of an AFF8000 HA pair or multi-node cluster is to use the System Performance Modeler, more commonly known as SPM. SPM v2.2 and later releases model random workloads on AFF8000 systems at low latencies. Latencies from 0.5 ms to 3 ms – in 0.5 ms increments – are supported. Inline compression, inline deduplication, and advanced drive partitioning (ADP) for root aggregates are also supported.

How do AFF8000 systems perform with Oracle databases?

Answer: Technical report TR-4415: NetApp AFF8080 EX Performance and Server Consolidation with Oracle Database covers performance, storage efficiency and server consolidation with an OLTP workload on Oracle 12c. The link to the report is in the Resources section.

How do AFF8000 systems perform with Microsoft SQL Server databases?

Answer: Technical report TR-4403: NetApp AFF8080 EX Performance and Server Consolidation with Microsoft SQL Server 2014 covers performance, storage efficiency and server consolidation with an OLTP workload. The link to the report is in the Resources section.

How do AFF8000 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 AFF8000 systems. The report includes performance and storage efficiency results, and the link to the report is in the Resources section.

Do AFF8000 systems perform better with Data ONTAP 8.3.1 or 8.3.2 than with Data ONTAP 8.3?

Answer: Yes. Data ONTAP 8.3.1 provides comparable or better performance for 8KB random workloads, with inline adaptive compression enabled, as compared to Data ONTAP 8.3 without inline compression. Data ONTAP 8.3.1 includes read and write path optimizations for SSDs and AFF systems, beyond what is implemented in Data ONTAP 8.3. Data ONTAP 8.3.2 extends comparable random performance to 16KB and 32KB I/O sizes with an enhancement to the way inline compression processing is done.

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

Answer: All reads are served from flash media with an AFF system. Reads of blocks that are already in cache with a hybrid FAS system will be similarly fast, but reads from HDDs will be slower and 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 Data ONTAP reallocation methods recommended with AFF8000 systems?

Answer: No. Data ONTAP reallocation methods are specifically designed to either improve sequential read performance, or write performance, for volumes and LUNs provisioned on HDD-only or Flash Pool aggregates. SSDs don't have mechanical or rotational seek latencies, and writes to SSDs are always to contiguous blocks because that's how NAND flash works. In addition, AFF systems are expected to be very busy and often aggregates will be relatively small and quite full. Thus, enabling read reallocate on AFF volumes or LUNs, and free space reallocation on SSD aggregates provides no benefit.

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 use is:

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

6 Interoperability

Can AFF8000 HA pairs be configured in a Data ONTAP cluster that contains FAS HA pairs?

Answer: Yes. AFF and FAS HA pairs can be configured in the same cluster. In fact, this is an important use case that enables non-disruptive data movement between all flash and hybrid or all-HDD storage tiers within a storage virtual machine (SVM). This capability differentiates AFF and clustered Data ONTAP from competitive all flash array systems.

Is an AFF HA pair supported in a clustered Data ONTAP MetroCluster configurations?

Answer: Yes. Starting with Data ONTAP 8.3.1, two-node and four-node MetroCluster clusters configured with AFF8000 HA pairs are supported.

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 AFF8000 HA pair be configured as a NetApp Storage Encryption (NSE) system?

Answer: Yes. 800GB encrypting SSDs are offered to configure a NSE AFF8000 system. All SSDs on the HA pair must be self-encrypting SSDs, and the SSDs are FIPS 140-2 compliant.

Note: NSE drives are not supported with clustered Data ONTAP MetroCluster (MCC) systems. FAS controllers with FlexArray enabled and 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; this means 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, which means blocks from the source volume are copied as-is (that is with no 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 Data 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 using, 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 also 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 on both the source and destination volumes. When the source and destination volumes have the same compression or deduplication policy, then the associated savings are preserved during transfer and on the destination. However, 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 system?

Answer: Yes. However, inline compression and deduplication policies that are enabled on a destination volume—in this case on an AFF system—are not applied to data that are transferred using 7MTT. 7MTT uses physical block replication technology which means compressed and de-duplicated 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 (available starting with Data ONTAP 8.3.2 RC1) 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 systems?

Answer: No. FlexArray uses array LUNs provisioned on HDDs or SSDs for storage. Data ONTAP cannot determine what type of physical drive a LUN is provisioned 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 has to be enabled on a node and all SSDs must un-partitioned spares to be sanitized. Refer to the [Clustered Data ONTAP 8.3 Physical Storage Management Guide](#) for more information.

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

Does Flash Cache work in an AFF system?

Answer: No. Data 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.

Furthermore, Flash Cache cards should be removed from any controller that is configured with SSDs only. The presence of Flash Cache cards in the controller results in memory being allocated for cache metadata during controller boot. This reduces the amount memory available for use as read buffers, and more read buffers improves performance.

7 Additional Resources

All Flash FAS Product Page – Field Portal

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

Technical Reports

- TR-4403: NetApp AFF8080 EX Performance and Server Consolidation with Microsoft SQL Server 2014 <http://www.netapp.com/us/media/tr-4403.pdf>
- TR-4415: NetApp AFF8080 EX Performance and Server Consolidation with Oracle Database <http://www.netapp.com/us/media/tr-4415.pdf>
- 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 Data 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-4322: NetApp All Flash FAS for Databases Performance Testing Summary for FAS8000 Series www.netapp.com/us/media/tr-4322.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

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

Presentations

- AFF SE Technical Presentation <https://fieldportal.netapp.com/content/202265?assetComponentId=202363>

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.