



NetApp®

Updated for 8.3.1

Clustered Data ONTAP® 8.3

Using All Flash FAS with Clustered Data ONTAP



NetApp, Inc.
495 East Java Drive
Sunnyvale, CA 94089
U.S.

Telephone: +1 (408) 822-6000
Fax: +1 (408) 822-4501
Support telephone: +1 (888) 463-8277
Web: www.netapp.com
Feedback: doccomments@netapp.com

Part number: 215-10179_B0
September 2015

Contents

What All Flash FAS is	4
How the AFF platform differs from the FAS platform	5
Setup checklist for AFF	6
Optimizing performance on AFF clusters	7
Assigning disk ownership on AFF systems	7
Creating a Storage Virtual Machine (SVM) on an AFF80xx system	7
Creating aggregates on AFF80xx systems	8
Creating volumes on an AFF80xx system	8
AFF80xx disk partition supported configurations	9
Readying the cluster for production	11
Where to find additional resources for AFF	12
Copyright information	13
Trademark information	14
How to send comments about documentation and receive update notifications	15
Index	16

What All Flash FAS is

All Flash FAS (AFF) systems optimize I/O and maximize application throughput while running flash-optimizing functions to accelerate a broad range of SAN and NAS workloads. In general, AFF80xx systems use the same Data ONTAP commands and features as do the FAS platforms, although there are a few differences.

AFF80xx systems can be ordered with a factory-configured FC SAN optimized configuration. If you have an FC SAN optimized AFF system and want more information about setting it up, see [Clustered Data ONTAP 8.3 FC SAN Optimized AFF Setup Guide](#).

How the AFF platform differs from the FAS platform

AFF platforms differ from FAS platforms in a few significant ways. You should understand these differences before using your AFF cluster.

The following table describes the differences between AFF and FAS systems:

Difference	Description
SSD drives	Only SSD drives are supported on AFF platforms. If non-SSD drives are attached to an AFF platform, the disks do not go online. If the root aggregate is HDD-based, the nodes do not boot. HDDs and FlexArray are not supported.
Performance	Using an AFF80xx provides maximum I/O performance. For example, AFF systems running Data ONTAP 8.3.1 have experienced 200 to 300 microsecond reductions in FCP latency for read-intensive workloads already at or around 1 microsecond in overall response times.
Volumes	All volumes created in aggregates owned by an AFF node have inline compression and inline deduplication of zeroes enabled by default.
SVMs	The storage efficiency policies for SVMs differ from those on a FAS system.
Root-data partitioning	The root-data partitioning capability is enabled by default on an AFF node.
Mixed platform support	You cannot pair an AFF node with a regular FAS node in an HA pair. You can mix AFF HA pairs with FAS HA pairs in a cluster. For more information, see the Clustered Data ONTAP Storage Platform Mixing Rules.
Zero block elimination	Instead of writing incoming 4 KB blocks that contain all zeroes to storage media, AFF80xx systems indicate the block's inode as all zeroes. With deduplication enabled, zero block elimination is enabled by volume by default, with no user configuration required.
FC SAN Optimization	An FC SAN optimized AFF configuration is available. The FC SAN optimized cluster enables you to quickly complete a configuration that is ready for demonstration and testing.

Related information

[Clustered Data ONTAP Storage Platform Mixing Rules](#)

Setup checklist for AFF

After you have cabled your system and performed the initial setup of your AFF80xx system, you can perform some of the AFF-specific configuration tasks.

- Verify that the hardware setup is complete and Config Advisor runs without errors.
- Prepare hosts that you want to connect to the cluster.
- Use System Setup to set up your cluster, or the Cluster Setup wizard to perform initial configuration.
- Zone switches (if applicable).

Optimizing performance on AFF clusters

While most configuration is the same for AFF clusters, you must configure the storage layout differently from FAS clusters to optimize performance.

Assigning disk ownership on AFF systems

Before using your AFF system, you should assign disk ownership to maximize performance.

About this task

AFF80xx systems use root-data partitioning to partition aggregates.

Twelve SSDs from each shelf are assigned to a node. One data aggregate is provisioned on each node.

Steps

1. Assign disks 0 through 11 to the first node in the HA pair.
2. Assign disks 12 through 23 to the second node in the HA pair.
3. If required, repeat Steps 1 and 2 for additional disk shelves.

You can partition up to 48 SSDs on an HA pair depending on your configuration.

Related references

[AFF80xx disk partition supported configurations](#) on page 9

Creating a Storage Virtual Machine (SVM) on an AFF80xx system

The process for creating SVMs is the same for either an AFF or FAS system, but the SVM on an AFF system uses different options.

Step

1. Create an SVM using the default IPspace, root volume security style of UNIX, and no snapshot policy:

If you want to use...	Then...
System Setup	System Setup creates your SVM when you first create the cluster.
System Manager	Create the SVM by using the SVM Create Wizard.
CLI	Enter the following command: <pre>vserver create -vserver <vserver> [-ipspace <IPspace>] -rootvolume <volume name> -aggregate <aggregate name> -rootvolume-security-style <security style> [-snapshot-policy <snapshot policy>]</pre>

Example

The following example creates an SVM named vs0.example.com in the default IPspace. The SVM's root volume is named root_vs0, is located on aggregate aggr0, and has UNIX security style. The snapshot policy is specified as None.

```
vserver create -vserver vs0.example.com -ipspace default -rootvolume
root_vs0 -aggregate aggr0 -rootvolume-security-style unix -snapshot-
policy none
```

Creating aggregates on AFF80xx systems

The process for creating aggregates on AFF80xx systems is the same as on a FAS system, but you should specify certain options that are specific to AFF.

Steps

1. Create one data aggregate on each node using 23 partitions and one hot spare partition per node:

If you want to use...	Then...
System Setup	System Setup partitions your storage automatically when you first create the cluster.
System Manager	Create the aggregates by using the Aggregate Create Wizard.
CLI	Enter the following command: <pre>storage aggregate create -aggregate <aggr_name> - diskcount <number_of_partitions> [-node <node_name>]</pre>

Example

The following example creates aggregate named aggr0 on a node named node0. The aggregate contains 23 disks.

```
storage aggregate create -aggregate aggr0 -node node0 -diskcount 23
```

2. Create additional aggregates if your system has more than 48 disks.

Creating volumes on an AFF80xx system

You should create volumes that are the same size as the aggregate on which they are located and ensure that they are thinly provisioned. Volumes on AFF systems also have different efficiency policies than volumes on FAS systems.

About this task

Volumes created in aggregates owned by an AFF node have inline compression enabled by default.

Steps

1. Create a volume that is the same size as the aggregate that contains it:

If you want to use...	Then do this...
System Manager	Create volumes using the UI.

If you want to use...	Then do this...
CLI	Enter the following command: <pre>volume create -vserver <vserver_name> -volume <volume_name> -aggregate <aggregate_name> [- size <volume_size>] [-autosize-mode {off grow grow_shrink}] [-space-guarantee {none volume}] [- snapshot-policy <snapshot policy>] [-fractional- reserve <percent>]</pre>

Example

The following example displays a volume being created with autosize mode off, no snapshot policy, no space guarantee, and a fractional reserve equaling zero:

```
volume create -vserver vs0.example.com -volume user_jdoe -aggregate
aggr0 - size <volume_size> -autosize-mode off -space-guarantee none -
snapshot-policy none -fractional-reserve 0
```

2. Create the volume efficiency policy using the default policy, daily schedule, and best-effort QoS policy:

If you want to use...	Then do this...
System Manager	Create the efficiency policy using the UI.
CLI	Enter the following command: <pre>volume efficiency policy create -vserver <vserver name> -policy <text> [-type <Efficiency policy type>] [-schedule <text>] [-qos-policy <Efficiency QoS policy>]</pre>

Example

The following example displays a volume efficiency policy being created with the default policy, daily schedule, and best-effort QoS policy:

```
volume efficiency policy create -vserver <vserver name> -policy default
-type scheduled -schedule daily -qos-policy best effort
```

AFF80xx disk partition supported configurations

AFF80xx platforms have different supported disk partitioning configurations based on the drive size of the platform.

AFF80xx platform	Partitions and drive size
AFF8020	48 x 200 GB
	36 x 400 GB
	36 x 800 GB
AFF8040	48 x 200 GB
	36 x 400 GB
	36 x 800 GB

10 | Using All Flash FAS with Clustered Data ONTAP

AFF80xx platform	Partitions and drive size
AFF8060	48 x 200 GB
	36 x 400 GB
	36 x 800 GB
AFF8080	36 x 400 GB
	36 x 800 GB

Readying the cluster for production

Before you put your cluster into production, you should complete some additional configuration steps.

Steps

1. Assign passwords for cluster interconnect switches (if applicable).
2. Configure EMS.
3. Configure the DNS network.
4. Protect the root volumes.
5. Create production volumes and LUNs.
6. Optionally join additional AFF or FAS nodes to the cluster.

You can add FAS nodes to an existing AFF cluster, but you cannot merge clusters.

Where to find additional resources for AFF

After performing AFF-specific configuration tasks, you can continue configuring the cluster using additional resources.

Additional AFF resources

[*Clustered Data ONTAP 8.3 FC SAN Optimized AFF Setup Guide*](#)

Additional clustered Data ONTAP resources

- [*NetApp Documentation*](#)
Links to product documentation, including the Data ONTAP product library.
- [*NetApp Documentation: Clustered Data ONTAP Express Guides*](#)
Specific express guides describe how to configure protocols and to verify setup.
- [*Clustered Data ONTAP 8.3 Network Management Guide*](#)
Describes how to configure and manage physical and virtual network ports (VLANs and interface groups), LIFs using IPv4 and IPv6, routing, and host-resolution services in clusters; optimize network traffic by load balancing; and monitor the cluster by using SNMP.
- [*Clustered Data ONTAP 8.3 Physical Storage Management Guide*](#)
Describes how to manage physical storage resources for FlexVol volumes and Infinite Volumes in clusters, including disks, RAID groups, and aggregates.
- [*Clustered Data ONTAP 8.3 System Administration Guide*](#)
Describes general system administration of a cluster, including the CLI interface, cluster access, node management, SVM setup, user account management, event monitoring, and performance evaluation.
- [*Clustered Data ONTAP 8.3 Cluster Management Using OnCommand System Manager*](#)
A printable version of System Manager online Help that describes how to configure, manage, and monitor storage objects and storage systems running clustered Data ONTAP.

Copyright information

Copyright © 1994–2015 NetApp, Inc. All rights reserved. Printed in the U.S.

No part of this document covered by copyright may be reproduced in any form or by any means—graphic, electronic, or mechanical, including photocopying, recording, taping, or storage in an electronic retrieval system—without prior written permission of the copyright owner.

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

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

NetApp reserves the right to change any products described herein at any time, and without notice. NetApp assumes no responsibility or liability arising from the use of products described herein, except as expressly agreed to in writing by NetApp. The use or purchase of this product does not convey a license under any patent rights, trademark rights, or any other intellectual property rights of NetApp.

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

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

Trademark information

NetApp, the NetApp logo, Go Further, Faster, AltaVault, ASUP, AutoSupport, Campaign Express, Cloud ONTAP, Clustered Data ONTAP, Customer Fitness, Data ONTAP, DataMotion, Fitness, Flash Accel, Flash Cache, Flash Pool, FlashRay, FlexArray, FlexCache, FlexClone, FlexPod, FlexScale, FlexShare, FlexVol, FPolicy, GetSuccessful, LockVault, Manage ONTAP, Mars, MetroCluster, MultiStore, NetApp Insight, OnCommand, ONTAP, ONTAPI, RAID DP, RAID-TEC, SANtricity, SecureShare, Simplicity, Simulate ONTAP, Snap Creator, SnapCenter, SnapCopy, SnapDrive, SnapIntegrator, SnapLock, SnapManager, SnapMirror, SnapMover, SnapProtect, SnapRestore, Snapshot, SnapValidator, SnapVault, StorageGRID, Tech OnTap, Unbound Cloud, and WAFL and other names are trademarks or registered trademarks of NetApp, Inc., in the United States, and/or other countries. All other brands or products are trademarks or registered trademarks of their respective holders and should be treated as such. A current list of NetApp trademarks is available on the web at <http://www.netapp.com/us/legal/netapptmlist.aspx>.

How to send comments about documentation and receive update notifications

You can help us to improve the quality of our documentation by sending us your feedback. You can receive automatic notification when production-level (GA/FCS) documentation is initially released or important changes are made to existing production-level documents.

If you have suggestions for improving this document, send us your comments by email to doccomments@netapp.com. To help us direct your comments to the correct division, include in the subject line the product name, version, and operating system.

If you want to be notified automatically when production-level documentation is released or important changes are made to existing production-level documents, follow Twitter account @NetAppDoc.

You can also contact us in the following ways:

- NetApp, Inc., 495 East Java Drive, Sunnyvale, CA 94089 U.S.
- Telephone: +1 (408) 822-6000
- Fax: +1 (408) 822-4501
- Support telephone: +1 (888) 463-8277

Index

A

- access
 - additional documentation [12](#)
- additional information
 - where to find for clustered Data ONTAP [12](#)
 - where to find for FC SAN optimized AFF clusters [12](#)
- AFF
 - See* All Flash FAS
- AFF platforms
 - how they differ from FAS platforms [5](#)
- AFF systems
 - assigning disk ownership [7](#)
 - creating aggregates on [8](#)
 - creating SVMs on [7](#)
 - creating volumes on [8](#)
 - readying for production [11](#)
- aggregates
 - creating on AFF80xx systems [8](#)
- All Flash FAS
 - what it is [4](#)

C

- checklist
 - for setting up AFF systems [6](#)
 - readying cluster for production [11](#)
- clustered Data ONTAP
 - where to find additional information [12](#)
- comments
 - how to send feedback about documentation [15](#)
- configuration
 - final steps to ready the cluster for production [11](#)

D

- differences
 - AFF and FAS [5](#)
- disk ownership
 - assigning on AFF systems [7](#)
- disk partitioning
 - supported configurations [9](#)
- disks
 - assigning ownership on AFF systems [7](#)
- documentation
 - additional information about protocol configuration after setup [12](#)
 - how to receive automatic notification of changes to [15](#)

- how to send feedback about [15](#)
- drive sizes
 - supported disk partitions [9](#)

F

- FAS platforms
 - how they differ from AFF platforms [5](#)
- FC SAN optimized AFF clusters
 - where to find additional information [12](#)
- feedback
 - how to send comments about documentation [15](#)

I

- information
 - how to send feedback about improving documentation [15](#)

P

- performance
 - optimizing [7](#)
- policies
 - setting on the SVM for an AFF80xx system [7](#)
- production
 - readying AFF systems for [11](#)

S

- setup checklist
 - AFF systems [6](#)
- storage layout
 - optimizing performance [7](#)
- suggestions
 - how to send feedback about documentation [15](#)
- SVMs
 - creating on an AFF80xx system [7](#)

T

- twitter
 - how to receive automatic notification of documentation changes [15](#)

V

- volumes
 - creating on an AFF80xx system [8](#)