



FC SAN Optimized AFF Setup Guide

FC SAN optimized All Flash FAS (AFF) clusters are preconfigured at the factory

Note: The procedures described in this document are intended for experienced NetApp SAN storage administrators. Contact your account representative if you need additional guidance configuring your system.

About FC SAN optimized clusters

FC SAN optimized clusters are preconfigured at the factory, so the setup process is different from the setup for other NetApp® clusters.

Most setup tasks have already been completed according to the latest All Flash FAS (AFF) FC SAN best practices. The cluster, licenses, aggregates, Storage Virtual Machine (SVM), and Logical Interfaces (LIFs) are preconfigured, and the optimum storage efficiency and performance parameters are set.

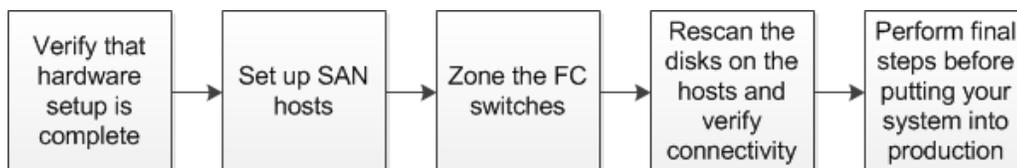
For the complete list of the factory-configured values, see [Appendix: Factory configuration details](#) on page 5.

Moving your system into production

The simplicity of the preconfigured SAN optimized cluster enables you to quickly complete a deployment that is ready for demonstration and testing.

This document provides a high-level summary of the steps you need to take to prepare your cluster to make it production-ready. The steps provided here assume experienced working knowledge of NetApp and SAN administration.

Setup overview



IMPORTANT: Before proceeding with this document, you must complete all steps in the printed *Installation and Setup Instructions* shipped with your AFF system, including cabling the system and logging into System Manager to complete the SAN setup wizard for AFF.

Do *not* attempt to run the System Setup program, the node setup wizard, or the cluster setup wizard. Do *not* change the storage efficiency or performance settings.

1. Verify that the hardware setup is complete

1. Confirm that all components are mounted and cabled.
2. Confirm that FC target ports are cabled to FC switches:
Use a minimum of 2 ports per node to fabric A, and 2 ports per node to fabric B.
3. Confirm that the system is powered on.
4. Download and install the Config Advisor tool:
http://mysupport.netapp.com/NOW/download/tools/config_advisor/
Config Advisor is a configuration validation and health check utility for NetApp systems that runs a series of commands on the hardware and then verifies cabling, configuration, availability, and conformance with NetApp best practice settings.
5. Confirm that your computer is physically connected to the management switch for the cluster and configured with an IP address in the 10.10.10.x/24 subnet.
6. Run Config Advisor, select the *Clustered Data ONTAP* execution profile, enter the cluster node and switch login information, and then click Collect Data.
7. Review the Config Advisor output and resolve any issues reported.

2. Set up SAN hosts

You need one or more host computers connected to the cluster. If desired, you can complete all of the host configuration steps in advance.

1. Using the NetApp Interoperability Matrix Tool (IMT), verify that the entire configuration is supported: <http://mysupport.netapp.com/matrix>
Note the Host Utilities version required for your configuration.
2. Install the Host Utilities.
3. Use 16-Gb FC initiators.
4. Enable and configure a supported multipathing (MPIO) solution.

5. Cable the FC initiator ports on the host to the FC switches.
6. Record the WWPN values of the FC initiators associated with each host.

For detailed host setup instructions, see the FC Configuration Express Guides:
<https://mysupport.netapp.com/documentation/docweb/index.html?productID=62092>

3. Zone the FC switches

Clustered Data ONTAP uses N_Port ID Virtualization (NPIV) for FC data. To connect hosts to the cluster, you must use 16-Gb FC switches and you must zone the FC switches by the WWPNs of the initiator and target ports.

1. Create one zone for each initiator port (single initiator zoning).
2. To each zone, add all FC Logical Interfaces (LIFs) on target ports connected to the same FC switch as the initiator (a minimum of two ports per node).
3. Save and activate the new zones.

For more information about configuring your SAN fabric, see the Clustered Data ONTAP *SAN Configuration Guide*:

https://library.netapp.com/ecm/ecm_download_file/ECMP12458217

4. Rescan disks and verify connectivity

On the connected hosts, discover the new LUNs and make them available to your application. The exact process depends on the host OS and the application.

1. Rescan the disks.
2. Configure the LUNs for use by the application.
3. Read and write data to verify access to the LUNs.

5. Final steps before putting your system into production

The SAN optimized factory configuration is intended for demonstration and testing purposes. For the complete list of the factory-configured values, see [Appendix: Factory configuration details](#) on page 5.

Before you put the cluster into production, you need to complete some additional configuration steps:

1. Register the system on the [NetApp Support Site](#).
2. Remove the test LUNs and volumes.
3. Update the password for the cluster admin account.
4. Update the password for the Storage Virtual Machine (SVM) admin account if you intend to keep the SVM.
5. Assign/change passwords for the cluster interconnect switches (if present).
6. Rename the cluster.
7. Change the management IP addresses to fit your network.
8. Configure AutoSupport™ messages.
9. Configure DNS for your network.
10. Configure NTP for your network.
11. Configure EMS messages.
12. Protect the root volumes.
13. Create production volumes and LUNs.
14. Optionally join additional AFF or FAS nodes to the cluster.

Note: You can join nodes to the FC SAN optimized cluster, but you cannot add the FC SAN optimized cluster to an existing cluster. You can only join nodes to an existing cluster, you cannot merge clusters.

For detailed instructions, see the Data ONTAP Express Guides and clustered Data ONTAP documentation:

<http://mysupport.netapp.com/documentation/docweb/index.html?productID=62092>

<http://mysupport.netapp.com/documentation/docweb/index.html?productID=62091>

For Oracle, see *Best Practices for Oracle Databases on NetApp Storage*:

<http://www.netapp.com/us/media/tr-3633.pdf>

Appendix: Factory configuration details

The factory configures the SAN optimized configuration as follows. These values have been tested and found to be optimal for AFF clusters. Do not modify the performance and space efficiency parameters unless directed by technical support.

Object	Default
Cluster management IP address	10.10.10.10/24
User name	admin
Password	netapp!123
Cluster name	aff
Node management (each node)	Port: e0m 10.10.10.1 (node 1) 10.10.10.2 (node 2)
Onboard UTA2 ports	Configured as FC target Fabric A: 0e, 0g Fabric B: 0f,0h
SVM	Name = AFF_SAN_DEFAULT_SVM -snapshot-policy none -ipSPACE default -rootvolume-security-style unix SVM management LIF: 10.10.10.100 SVM management user name: admin SVM management password: netapp!123
FCP service	Created
FC data LIFs on each node	aff_node-0n_0e_A aff_node-0n_0g_A aff_node-0n_0f_B aff_node-0n_0h_B -data-protocol fcp
Portsets	None
Licenses installed	Base cluster CIFS FCP FlexClone Insight_Balance iSCSI NFS SnapManagerSuite SnapMirror SnapProtectApps SnapRestore SnapVault
Disk ownership (all shelves)	Bays 0-11 to node 1 Bays 12-23 to node 2

Object	Default
Root aggregates (first two shelves only)	One root aggregate of 23 disks created on each node with root-data partitioning One hot spare partition per node name = node_0n_aggr0
Data aggregates (first two shelves only)	One data aggregate of 23 disks created on each node with root-data partitioning -space-nearly-full-threshold-percent 70 -space-full-threshold-percent 75 One hot spare partition per node name = node_0n_aggr1
Volume properties (all volumes same size as aggregate and thinly provisioned)	autosize-mode = off snapshot-policy = none fractional-reserve = 0 space-guarantee = none
Volume efficiency policy	-policy default -type scheduled -schedule daily -qos-policy best_effort inline compression on
SNMP community	-community-name public -type ro options snmp.enable on
VBN_ZERO	On
CSC	Off
Read_realloc	Off
Atime update	Off
Deduplication	sis.max_active_ops = 1
Automatic RAID-level scrubs	raid.scrub.schedule = 6h@sun@1