



TECHNICAL FAQ

Data ONTAP 8.1 Operating in Cluster-Mode Networking

NetApp
January 2012

TABLE OF CONTENTS

- 1 GENERAL QUESTIONS..... 3
- 2 HARDWARE REQUIREMENTS 4
 - CLUSTER-INTERCONNECT 4
 - MANAGEMENT NETWORK 5
 - NETWORK INTERFACES 6
- 3 LOGICAL INTERFACES 6

1 GENERAL QUESTIONS

FOR WHAT IS THE CLUSTER-INTERCONNECT USED?

The cluster interconnect network is used for a variety of purposes, including:

- Allowing access to data on any node from any other node in the cluster
- Synchronizing configurations
- Intracluster mirroring
- Volume movement

CAN ANYTHING OTHER THAN 10 GIGABIT ETHERNET BE USED FOR THE CLUSTER-INTERCONNECT?

No. The only exception is the FAS2040. See additional information in the next section.

ARE SPECIFIC SWITCHES REQUIRED FOR THE CLUSTER-INTERCONNECT?

Yes. NetApp has qualified the Cisco Nexus® 5010 and 5020 for use as the cluster-interconnect switches.

CAN I DIRECTLY CONNECT NODES INSTEAD OF USING THE CLUSTER-INTERCONNECT SWITCH?

No. Direct connection of nodes is not supported. Even if the cluster is only a single HA pair, the cluster-interconnect switches are still required.

DOES DATA ONTAP OPERATING IN CLUSTER-MODE SUPPORT IPV6?

No. Currently IPv6 is only supported in Data ONTAP® 8.1 7-Mode and Data ONTAP 7.3.6.

CAN MY CLUSTER BE DISTRIBUTED GEOGRAPHICALLY?

No. Geographically distributed Cluster-Mode systems are not allowed.

Connection Source	Connection Destination	Maximum Distance
Cluster node	Cluster-interconnect switch	300 meters
Cluster-interconnect switch	Cluster-interconnect switch	300 meters

The distance between nodes within an HA pair is limited by the NVRAM interconnect. The distance between HA pairs is limited by the cabling standard used to connect the node to the interconnect switch.

2 HARDWARE REQUIREMENTS

CLUSTER-INTERCONNECT

WHAT IS THE CLUSTER-INTERCONNECT?

The cluster-interconnect is a dedicated and redundant 10 Gigabit Ethernet network. Only the approved switches that have been qualified for use with Cluster-Mode systems can be used and must not be shared with other networks.

WHY DO I HAVE TO USE THE CISCO NEXUS 5010 OR 5020 10GBE SWITCHES FOR THE CLUSTER-INTERCONNECT?

These switches are an integral part of a Data ONTAP Cluster-Mode system. They provide a low-latency and high-bandwidth communication transport for the cluster interconnect. NetApp has tested specific Cisco® NX-OS software releases and NetApp® reference configuration files (RCFs) on these switches to make sure of interoperability with a given release of Data ONTAP operating in Cluster-Mode. These are the only supported switches at this time.

IF I ALREADY HAVE A CISCO NEXUS 5010/5020 OR CISCO NEXUS 7000 SWITCH, CAN I JUST UPGRADE THE NX-OS SOFTWARE AND APPLY THE REFERENCE CONFIGURATION FILES?

Only the Cisco Nexus 5010/5020 can be used for the Cluster-Interconnect. The use of customer owned switches may be possible through an approved PVR. The customer owned switches must meet the following criteria to be considered for approved use.

- Customer owned switches must be covered under an active Cisco SmartNet support contract.
- Customer owned switches must be completely dedicated to the Cluster-Interconnect function.
- Customer owned switches must run the version of Cisco NX-OS and NetApp Reference Configuration File (RCF) that is prescribed for the version of Data ONTAP Operating in Cluster-Mode being used.

DOES THE CISCO NX-OS SOFTWARE HAVE TO BE UPGRADED WHEN I UPGRADE THE VERSION OF DATA ONTAP RUNNING IN CLUSTER-MODE?

Yes. If an upgrade is specified on the cluster-interconnect or management switch software page located on the NetApp NOW® support site. The version of Cisco NX-OS running on the cluster-interconnect switches has been tested with specific versions of Data ONTAP operating in Cluster-Mode to make sure of interoperability.

WHERE CAN I FIND THE CORRECT VERSION OF CISCO NX-OS TO INSTALL ON THE SWITCH?

You can find the latest version of Cisco NX-OS software under the cluster and management switch section of the software download area of the NOW support site.

https://now.netapp.com/NOW/download/software/cm_switches/

DOES THE REFERENCE CONFIGURATION FILE NEED FOR THE CLUSTER INTERCONNECT NEED TO BE UPDATED WHEN I UPGRADE THE VERSION OF DATA ONTAP RUNNING IN CLUSTER-MODE?

You can find the latest version of reference configuration files (RCFs) under the cluster and management switch section of the software download area of the NOW support site.

https://now.netapp.com/NOW/download/software/cm_switches/

ARE THE CLUSTER-INTERCONNECT SWITCH UPGRADES NONDISRUPTIVE?

Yes. Redundant cluster-interconnect switches allow nondisruptive upgrades when the prescribed upgrade procedures are followed.

CAN THE CLUSTER-INTERCONNECT SWITCHES BE SHARED?

No. The cluster-interconnect switches are dedicated to the Cluster-Mode interconnect only and are not to be used for any other function.

THE FAS2040 DOESN'T HAVE 10 GIGABIT ETHERNET. HOW DO I CONNECT IT TO A CLUSTER?

The FAS2040 can connect to the cluster switches using 1GbE SFPs. Only the first 8 ports on the Cisco Nexus 5010 or the first 16 ports on the Cisco Nexus 5020 can operate with 1GbE SFPs. Mixing of 1GbE and 10GbE nodes in a cluster is not recommended.

WHY DO I NEED TO HAVE TWO CLUSTER INTERCONNECT SWITCHES FOR DATA ONTAP OPERATING IN CLUSTER-MODE?

Two cluster-interconnect switches are required to provide system resilience and redundancy for nondisruptive operations.

WHY ARE THERE EIGHT INTER SWITCH LINK (ISL) CONNECTIONS BETWEEN THE TWO CLUSTER-INTERCONNECT SWITCHES?

Eight ISL connections are required between the two cluster-interconnect switches to provide system resilience and redundancy for nondisruptive operations.

CAN I CONFIGURE MORE THAN TWO CLUSTER LIFS PER NODE?

No. Currently only two cluster LIFs are currently supported in Data ONTAP operating in Cluster-Mode.

MANAGEMENT NETWORK**ARE SPECIFIC SWITCHES REQUIRED FOR THE MANAGEMENT NETWORK?**

No. The requirement is to have redundant management switches. NetApp can provide a set of Cisco Catalyst® 2960-24TT switches as part of a new Cluster-Mode hardware order. Refer to the Data ONTAP 8.1 cluster and management switch product bulletin located on the Field Portal for any specific questions regarding the purchase.

<https://fieldportal.netapp.com/viewcontent.asp?qv=1&docid=36697>

WHY DO I HAVE TO USE THE CISCO CATALYST 2960 FOR THE MANAGEMENT NETWORK?

These switches are an integral part of a Data ONTAP Cluster-Mode system. They serve as a common interconnect between a customer's network and the management interfaces on a Data ONTAP cluster operating in Cluster-Mode. NetApp has tested specific Cisco Catalyst software releases and NetApp reference configuration files (RCFs) on these switches to make sure of interoperability with each release of Data ONTAP operating in Cluster-Mode.

WHERE CAN I FIND THE CORRECT VERSION OF CISCO IOS SOFTWARE TO INSTALL ON THE SWITCH?

You can find the latest version of Cisco IOS® Software under the cluster/management switch section of the software download area of the NOW support site.

https://now.netapp.com/NOW/download/software/cm_switches/

WHERE CAN I FIND THE CORRECT VERSION OF THE REFERENCE CONFIGURATION FILE FOR THE MANAGEMENT SWITCHES?

You can find the latest version of reference configuration files (RCFs) under the cluster/management switch section of the software download area of the NOW support site.

https://now.netapp.com/NOW/download/software/cm_switches/

ARE THE MANAGEMENT SWITCHES DEDICATED?

Yes. The switches used to provide connectivity for cluster management are dedicated and should be considered part of the entire system.

CAN MY CUSTOMERS USE THEIR EXISTING ETHERNET SWITCHES FOR THE MANAGEMENT NETWORK?

It might be possible for a customer to use existing Ethernet switches. You will need to submit a product variance request (PVR) to get approval for using nonstandard switches.

NETWORK INTERFACES

CAN I USE LONG-RANGE (LR) OPTICS FOR CONNECTIVITY?

The use of LR optics is not currently supported in Data ONTAP operating in Cluster-Mode. If you have a business case for supporting LR optics, file a PVR.

HOW DO I KNOW WHICH PORTS CAN BE USED FOR DATA, CLUSTER, AND MANAGEMENT?

More information about supported platforms and configurations can be found in the system configuration guide location on the NetApp Support site at <http://now.netapp.com/NOW/knowledge/docs/hardware/NetApp/syscfg>.

CAN I USE CISCO 7-METER OR 10-METER ACTIVE SFP+ TWINAX CABLES WITH DATA ONTAP OPERATING IN CLUSTER-MODE?

No. The use of active SFP+ twinax cables is not supported. They might be supported in future versions of Data ONTAP operating in Cluster-Mode.

IS IT CONSIDERED A BEST PRACTICE TO USE FLOW CONTROL BETWEEN CLUSTER NODES AND THE SWITCHES TO WHICH THEY ARE CONNECTED?

No. It is not a recommended best practice to configure flow control between a switch and cluster node that it is connected to. Note that all Ethernet interfaces in Data ONTAP default to full flow control. This should be set to off as a best practice.

3 LOGICAL INTERFACES

WHAT IS A LIF?

A LIF is a logical interface. LIFs are logical resources tied to a Vserver and have an IP address that is mapped to physical network ports on the controller. This logical component allows the IP address to remain the same even if the LIF is migrated to another physical port on any node within the cluster.

HOW MANY PORTS CAN BE IN A LIF?

A LIF points to a single network port, a single VLAN, or a single interface group.

CAN I CONFIGURE A INTERFACE GROUP ACROSS CONTROLLERS

No. Interface groups are per controller. If you want redundancy across switches with interface groups from a node, you can configure multi-chassis LACP if it is supported.

HOW MANY LIFS CAN I CREATE?

Each cluster node can support a maximum of 262 LIFs, of which six are reserved:

- One management LIF
- Two cluster LIFs
- Three intercluster LIFs

However, when using SFO for the HA configuration, no more than 128 LIFs per node should be used to account for failover. It's important to account for failover scenarios when nodes might become responsible for another node's LIFs.

CAN I CONFIGURE MORE THAN THREE INTERCLUSTER LIFS ON A CLUSTER NODE?

Yes. If you configure less than 128 data LIFs, you can use the remaining LIFs as intercluster.

HOW MANY LIFS CAN I ASSIGN TO A PORT/INTERFACE?

There is no limit to the number of LIFs that can be placed on a single port; however, the per-node limit of 256 LIFs cannot be exceeded.

WHAT TYPES OF LIFS ARE THERE?

There are five different types of LIFs:

- Cluster management: Used to manage the cluster. There is one cluster-management LIF in a cluster.
- Node management: This LIF is used to manage the node directly in the cluster for system maintenance.
- Cluster: Used for the cluster-interconnect. There are two cluster-interconnect LIFs on each node in the cluster.
- Data: Used to serve data to clients/hosts. These LIFs are designated either SAN or NAS. At least one is required for each Vserver. A best practice is to have one LIF per node per Vserver for NAS access, two LIFs per node per Vserver for iSCSI access, and two LIFs per fabric per node per Vserver for FC access.
- Intercluster: Used for mirroring between clusters. These are node specific. At least one intercluster LIF is required per node for replication between clusters. Data LIFs can also be used to host intercluster replication.

WHAT IS THE RANGE OF VLAN TAGS SUPPORTED?

You can configure VLAN tags from 1 to 4,094.

HOW MANY VLAN TAGS ARE SUPPORTED PER CLUSTER NODE?

The number of VLAN tags supported is dependent on the hardware platform.

Platform	Maximum Number of VLANs
FAS6xx0	512
FAS3xx0	256
FAS2xx0	128

© 2011 NetApp, Inc. All rights reserved. No portions of this document may be reproduced without prior written consent of NetApp, Inc. Specifications are subject to change without notice. NetApp, the NetApp logo, Go further, faster, Data ONTAP, and NOW are trademarks or registered trademarks of NetApp, Inc. in the United States and/or other countries. Catalyst, Cisco, Cisco IOS, and Cisco Nexus are registered trademarks of Cisco Systems, Inc. All other brands or products are trademarks or registered trademarks of their respective holders and should be treated as such.