

Cisco MDS 9000 Series I/O Accelerator Configuration Guide, Release 9.x

First Published: 2022-09-02

Last Modified: 2026-03-04

Americas Headquarters
Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA
http://www.cisco.com
Tel: 408 526-4000
800 553-NETS (6387)
Fax: 408 527-0883

© 2021–2026 Cisco Systems, Inc. All rights reserved.



CONTENTS

PREFACE

Preface xi

Audience xi

Document Conventions xi

Documentation Feedback xiii

Related Documentation xiii

Communications, Services, and Additional Information xiii

CHAPTER 1

New and Changed Information 1

Change Summary 1

CHAPTER 2

Overview 3

About Cisco I/O Accelerator 3

IOA Supported Switches 3

IOA Unsupported Switches 3

Unified Acceleration Service 4

Topology Independent 4

Transport Agnostic 4

High Availability and Resiliency 4

Improved Tape Acceleration Performance 4

Load Balancing 4

Example IOA Topology 5

Terminology 5

Clustering 7

Hardware Requirements 7

Software Requirements 7

License Requirements 8

CHAPTER 3

Getting Started 11
Enabling SSH 11
Enabling CFS 12
IP Access Lists 12
Zone Default Policy 12
FC-Redirect 12
FC-Redirect Supported Switches 12
FC-Redirect Unsupported Switches 13
FC-Redirect Requirements 13
Configuring FC-Redirect v2 Mode 14
Using FC-Redirect with CFS Regions 15
Guidelines for Designing CFS Regions For FC-Redirect 15
Configuring CFS Regions For FC-Redirect 16
Using IOA Cluster with IPFC Interface 16
Task Flow for Configuring IOA Cluster To Use the IPFC Interface 17
Configuring an IOA Cluster To Use the IPFC Interface 17
Creating a VSAN Interface and Configuring IPv4 Addresses 17
Enabling IPv4 Routing 18
Verifying Connectivity 18
Creating an IOA Cluster and IOA Interface in the Local Node 19
Verifying Cluster Configuration 19
Adding a Remote Node and IOA Interface to the Remote Node 20
Verifying the Cluster Configuration 20
Configuration Example 21
Creating an Interface VSAN 21
Verifying the Configuration 21
Verifying the Connectivity 22
Configuring IOA Site on Switch sw-231-14 22
Configuring IOA Site on Switch sw-231-19 22
Configuring IOA Cluster cltr1 on Switch sw-231-14 22
Changing the Node to Use IPFC Interface Address 22
Adding a Remote Node to the IOA Cluster 23
Adding an IOA Interface to the Switch sw-231-14 23

Adding an IOA Interface to the Switch sw-231-19 23
Verifying the Cluster Configuration 23
Verifying the IP Address 23
Verifying the IOA Interface 24
Task Flow for Converting an Existing IOA Cluster to Use the IPFC Interface 24
Configuration Example for Converting IOA Cluster to Use the IPFC Interface 24
Verifying the IOA Cluster Configuration 24
Verifying the IP Address 24
Verifying the Flow Status 25
Shutting Down IOA Cluster on a Local Node 25
Shutting Down the IOA Cluster on the Remote Node 25
Removing the IOA Cluster from the Remote Node 25
Verifying the IOA Cluster in the Remote Node 26
Removing the Remote Node from the Cluster in the Local Switch 26
Changing the Local Node Configuration To Use IPFC Address 26
Activating the Single Node Cluster 26
Adding Remote Node with IPFC Address 26
Adding IOA Interfaces to the Remote Node 27
Verifying the Cluster Nodes 27
Verifying the Flow Status 27

CHAPTER 4 **Deployment Considerations** 29
Supported Topologies 29
Core-Edge Topology 29
Edge-Core-Edge Topology 30
Collapsed Core Topology 30
Extended Core-Edge Topology 31
Extending Across Multiple Sites 32
IVR Topologies 32
Other Topologies 33
Deployment Guidelines 33
General Guidelines 33
Scalability and Optimal Performance Considerations 33
Resiliency Considerations 34

Guidelines and Limitations 34
Configuration Limits 37

CHAPTER 5 Configuring IOA Using Cisco DCNM-SAN 39
IOA Manager 39
Toolbar 40
Launching IOA Manager 41
Configuring Sites 41
Adding a New Site 42
Removing a Site 43
Viewing a Site 44
Adding Switches to a Site 45
Removing Switches from a Site 45
Configuring Clusters 46
Adding a New Cluster 46
Removing a Cluster 48
Viewing Clusters 49
Configuring Interfaces 50
Assigning Interfaces to a Cluster 50
Removing Interfaces from a Cluster 51
Configuring Flows 52
Adding a Flow 52
Removing a Flow 54
Viewing Interface Statistics 55

CHAPTER 6 Configuring IOA Using the CLI 57
Configuring IOA 57
Enabling Clustering 58
Enabling the IOA Service 58
Classifying the Switch to IOA Site 59
Configuring IOA Interfaces 59
Displaying IOA Interface Status 60
Configuring an IOA Cluster 61
Displaying IOA Cluster Status 61

Adding Nodes to an IOA Cluster 62
Adding Interfaces to an IOA Cluster 63
Adding N Ports to an IOA Cluster 65
Configuring the IOA Flows 66
Verifying the IOA Flow Connectivity 67
IOA Flow Setup Wizard 67
Using the IOA Flow Setup Wizard 68
Creating Multiple IOA Clusters on a Single Switch 71
Configuring IOA with NPV 73
Guidelines for Configuring IOA with NPV 73
Configuring NPIV on an NPV Core Switch, NPV on an NPV Device, and Activating NP Link 74
Configuring NPIV on the NPV Core Switch 74
Configuring NPV on the NPV Device, Bringing Up the NP Port and NP Uplink 75
Verifying the NPV Configuration 76
Creating and Activating an IOA Cluster 77
Configuring NPV on IOA 77
Enabling NPV 77
Enabling NPIV on the NPV Core Switches 78
Verifying the Configured NP Uplinks 78
Enabling IOA on the IOA Nodes 80
Classifying the Switches into IOA Sites 80
Configuring IOA Interfaces 81
Configuring IOA Interfaces 82
Configuring IOA Cluster 82
Configuring Nodes to the IOA Cluster 83
Verifying the IOA Cluster Configuration 83
Configuring Interfaces in the IOA Cluster 84
Verifying the Cluster Interface Configuration 84
Adding N Ports to the IOA cluster 85
Verifying the Configured N Ports in the IOA Cluster 85
Configuring IOA Flows in the Cluster 86
Verifying the Configured IOA Flow 86
Displaying Interface Statistics 87
Additional Configurations for the Features Supported by NPV on IOA 87

NP Link Trunking 87
Configuring an NP Uplink Port 87
Verifying the Configured Trunking NP Uplink Port on the NPV Core Switch 88
Verifying the Configured Trunking NP Uplink Port on NPV Device Switch 89
Configuring F-PortChannel 90
Configuring F-PortChannel on the NPV Core Switch 90
Configuring NP-PortChannel on an NPV Device Switch 91
Turning on the Administrative State of PortChannel Member Interfaces in the NPV Core Switch 92
Turning on the Administrative State of PortChannel member interfaces in NPV Device Switch 92
Verifying the Configured PortChannel of NP Links 92
Example for Configuring TF-TNP PortChannel Links 93
Configuring the PortChannel on the NPV Core Switch 94
Configuring PortChannel on the NPV Device Switch 94
Verifying the Configured PortChannel of TF-TNP Links 96
Configuring FlexAttach Virtual pWWN on an NPV Switch 97
Automatically Enabling FlexAttach Virtual pWWN 97
Manually Enabling FlexAttach Virtual pWWN 98
Verifying the Configured FlexAttach Virtual pWWN 98
Verifying the Configured FlexAttach Virtual pWWN 99
Configuring NPV Traffic Management on NPV Switches with IOA 99
Configuring a List of External Interfaces per Server Interface 99
Enabling or Disabling the Global Policy for Disruptive Load Balancing 100
Verifying the NPV Traffic Management on an NPV Switch 100
Example for Implementing IOA with NPV 101
Verifying the IOA Configuration 101
Applying ASCII Configuration back to IOA Switches 101
Additional Configurations 103
Shutting Down a Cluster 103
Load Balancing the Flows 103
Setting the Tunable Parameters 104
Changing the Node Description and IP Address of an IOA Cluster 105
Configuration Example for Changing the Node Description and Node IP Address of an IOA Cluster 106
Shut Down the IOA Cluster on switch1 106

Shut Down the IOA Cluster on switch2 107
Remove the IOA Cluster on switch2 107
Remove the Node of switch2 in switch1 107
Change the Management Interface IP Address on Switches 107
Change the Node Description and IP Address on switch1 107
Bring Up IOA Cluster on switch1 107
Add switch2 Node with New Description and the IP Address 108
Add IOA Interfaces on switch1 108
Verify the Node Description and IP Address and Flows 108
Displaying Interface Statistics 109

CHAPTER 7 SCSI Write Acceleration and Tape Acceleration 113
SCSI Write Acceleration 113
SCSI Tape Acceleration 115

CHAPTER 8 Cluster Management and Recovery Scenarios 117
Cluster Quorum and Master Switch Election 117
Cluster Quorum 118
Master Switch Election 118
Two-Switch Cluster Scenarios 118
Three-Switch Cluster Scenarios 119
Four-Switch Cluster Scenarios 120
In-Service Software Upgrade (ISSU) in a Two-Node Cluster 120
Supported Topologies 121
Single-Fabric Topology 121
Cluster Recovery Scenarios 121
Deleting an Offline Switch from a Cisco IOA Cluster 121
Deleting a Cisco IOA Cluster with One or More Offline Switches while the Master Switch is Online 122
Deleting a Cisco IOA Cluster when All Switches Are Offline 123
Reviving a Cisco IOA Cluster 124



Preface

This preface describes the audience, organization of, and conventions used in the Cisco MDS 9000 Series Configuration Guides. It also provides information on how to obtain related documentation, and contains the following sections:

- Audience, on page xi
- Document Conventions, on page xi
- Documentation Feedback, on page xiii
- Related Documentation, on page xiii
- Communications, Services, and Additional Information, on page xiii

Audience

This publication is for network administrators who install, configure, and maintain Cisco MDS 9000 Series Switches.

Document Conventions

Command descriptions use these conventions:

Convention	Description
bold	Bold text indicates the commands and keywords that you enter literally as shown.
Italic	Italic text indicates arguments for which the user supplies the values.
[x]	Square brackets enclose an optional element (keyword or argument).
[x | y]	Square brackets enclosing keywords or arguments separated by a vertical bar indicate an optional choice.
{x | y}	Braces enclosing keywords or arguments separated by a vertical bar indicate a required choice.

Convention	Description
[x {y | z}]	Nested set of square brackets or braces indicate optional or required choices within optional or required elements. Braces and a vertical bar within square brackets indicate a required choice within an optional element.
`variable`	Indicates a variable for which you supply values, in context where italics cannot be used.
string	A nonquoted set of characters. Do not use quotation marks around the string or the string will include the quotation marks.

Examples use these conventions:

Convention	Description
`screen font`	Terminal sessions and information the switch displays are in screen font.
`boldface screen font`	Information you must enter is in boldface screen font.
italic screen font	Arguments for which you supply values are in italic screen font.
< >	Nonprinting characters, such as passwords, are in angle brackets.
[ ]	Default responses to system prompts are in square brackets.
!, #	An exclamation point (!) or a pound sign (#) at the beginning of a line of code indicates a comment line.

This document uses the following conventions:



Note Means *reader take note*. Notes contain helpful suggestions or references to material not covered in the manual.



Caution Means *reader be careful*. In this situation, you might do something that could result in equipment damage or loss of data.



Warning IMPORTANT SAFETY INSTRUCTIONS

This warning symbol means danger. You are in a situation that could cause bodily injury. Before you work on any equipment, be aware of the hazards involved with electrical circuitry and be familiar with standard practices for preventing accidents. Use the statement number provided at the end of each warning to locate its translation in the translated safety warnings that accompanied this device.

SAVE THESE INSTRUCTIONS

Documentation Feedback

To provide technical feedback on this document, or to report an error or omission, please send your comments to mds-docfeedback@cisco.com. We appreciate your feedback.

Related Documentation

The entire Cisco MDS 9000 Series switches documentation set is available at the following URL:

https://www.cisco.com/c/en/us/support/storage-networking/mds-9000-nx-os-san-os-software/series.html

Documentation Roadmap

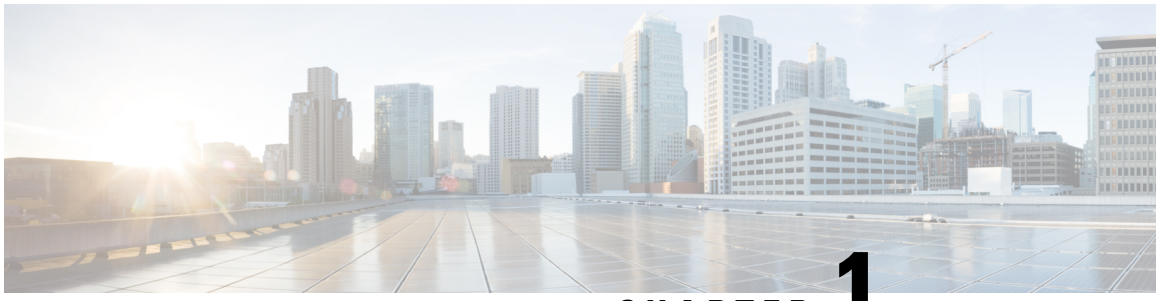
https://www.cisco.com/c/en/us/td/docs/storage/san_switches/mds9000/roadmaps/rel90.html

Communications, Services, and Additional Information

- To receive timely, relevant information from Cisco, sign up at Cisco Profile Manager.
- To get the business results you're looking for with the technologies that matter, visit Cisco Services.
- To submit a service request, visit Cisco Support.
- To discover and browse secure, validated enterprise-class apps, products, solutions and services, visit Cisco DevNet.
- To obtain general networking, training, and certification titles, visit Cisco Press.
- To find warranty information for a specific product or product family, access Cisco Warranty Finder.

Cisco Bug Search Tool

Cisco Bug Search Tool (BST) is a web-based tool that acts as a gateway to the Cisco bug tracking system that maintains a comprehensive list of defects and vulnerabilities in Cisco products and software. BST provides you with detailed defect information about your products and software.



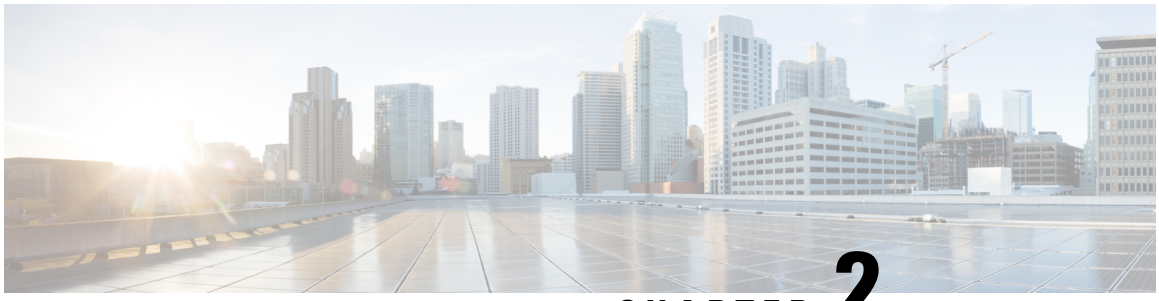
CHAPTER 1

New and Changed Information

- Change Summary, on page 1

Change Summary

Feature	New or Changed Topics	Changed in Release	Where Documented
IOA on Cisco MDS 24/10 Port SAN Extension Module	New guidelines related to IOA on the Cisco MDS 24/10 Port SAN Extension Module have been added.	8.2(1)	Guidelines and Limitations, on page 34



CHAPTER 2

Overview

This chapter provides an overview of the Cisco I/O Accelerator feature and includes the following sections:

- About Cisco I/O Accelerator, on page 3
- Example IOA Topology, on page 5
- Terminology, on page 5
- Clustering, on page 7
- Hardware Requirements, on page 7
- Software Requirements, on page 7
- License Requirements, on page 8

About Cisco I/O Accelerator

The Cisco MDS 9000 Family I/O Accelerator (IOA) feature provides Small Computer System Interface (SCSI) acceleration in a storage area network (SAN) where the sites are interconnected over long distances using Fibre Channel or Fibre Channel over IP (FCIP) Inter-Switch Links (ISLs).

IOA provides these features, which are described in the following sections:

IOA Supported Switches

IOA is supported on the following switches:

- Cisco MDS 24/10-Port SAN Extension Module on Cisco MDS 9700 Series switches
- Cisco MDS 9250i Switch

IOA Unsupported Switches

IOA is not supported on the following switches:

- Cisco MDS 9220i Switch
- Cisco MDS 9132T Switch
- Cisco MDS 9148T Switch
- Cisco MDS 9148S Switch

- Cisco MDS 9396T Switch
- Cisco MDS 9396S Switch

Unified Acceleration Service

IOA provides both SCSI write acceleration and tape acceleration features as a unified fabric service. These services were provided in previous releases in the form of Fibre Channel write acceleration for remote replication over Fibre Channel links and FCIP write acceleration and tape acceleration over FCIP links. Fibre Channel write acceleration was offered on the Storage Services Module (SSM) and FCIP write acceleration and tape acceleration were offered on the IP storage services modules. IOA offers both the write acceleration and tape acceleration services on the Cisco MDS MSM-18/4 module, SSN-16 module, and 9222i switch as a fabric service. This eliminates the need to buy separate hardware to obtain Fibre Channel write acceleration and FCIP write acceleration and tape acceleration.

Topology Independent

IOA can be deployed anywhere in the fabric without rewiring the hardware or reconfiguring the fabric. There are no restrictions on where the hosts and targets are connected to. Both the Fibre Channel and FCIP write acceleration is supported only on PortChannels but do not support multiple equal-cost links. FCIP tape acceleration is not supported on PortChannels. IOA eliminates these topological restrictions.

Transport Agnostic

IOA is completely transport-agnostic and is supported on both Fibre Channel and FCIP ISLs between two sites.

High Availability and Resiliency

IOA equally supports both PortChannels and equal-cost multiple path (ECMP) links across two data centers. This allows you to seamlessly add ISLs across the two data centers for capacity building or redundancy. IOA is completely resilient against ISL failures. IOA uses a Lightweight Reliable Transport Protocol (LRTP) to guard against any ISL failures as long as there is an alternate path available across the two data centers. Remote replication and tape backup applications are completely unaffected by these failures.

Improved Tape Acceleration Performance

IOA tape acceleration provides higher throughput numbers than the FCIP tape acceleration, which is limited by a single Gigabit Ethernet throughput.

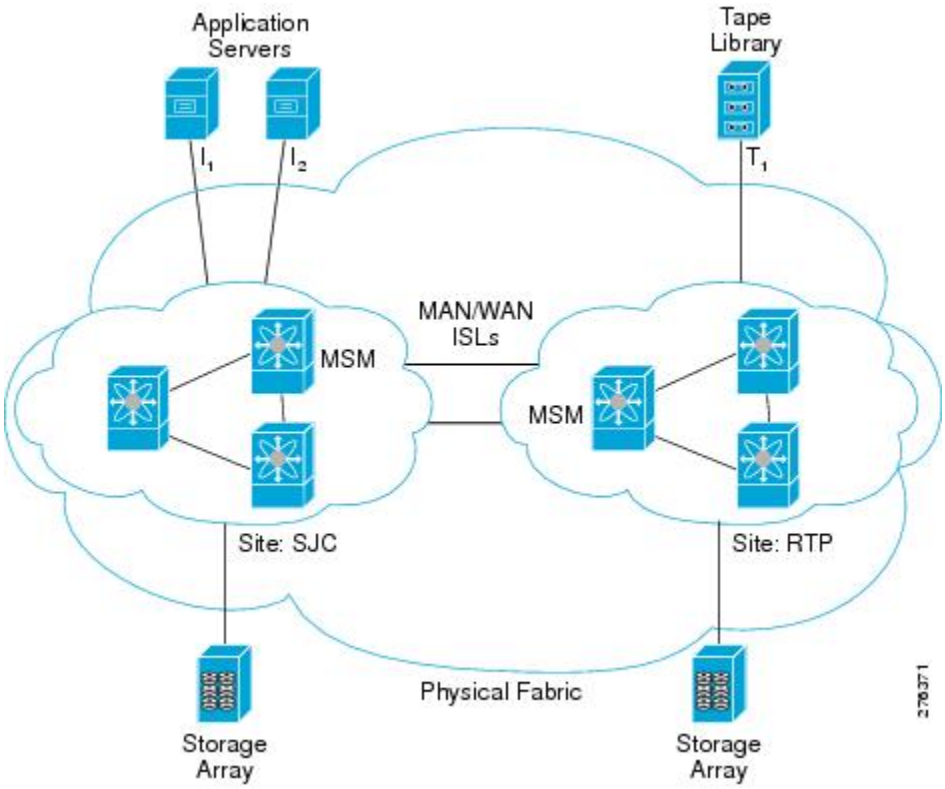
Load Balancing

IOA uses clustering technology to provide automatic load balancing and redundancy for traffic flows across multiple IOA service engines that can be configured for the IOA service. When an IOA service engine fails, the affected traffic flows are automatically redirected to the available IOA service engines to resume acceleration.

Example IOA Topology

Figure 1: Fabric with Two Sites, on page 5 illustrates a physical fabric that consists of two sites in different locations interconnected across the MAN or WAN using Fibre Channel or FCIP links. Remote replication and remote tape backup services run across these two data centers.

Figure 1: Fabric with Two Sites



Note This topology illustrates a single fabric only. In a dual fabric, the second fabric is an exact replica of this topology, and the concepts that are described in this document are applicable to the second fabric as well.

Terminology

The following Cisco IOA-related terms are used in this book:

- **Fabric**—A physical topology of switches interconnected by Fibre Channel or FCIP ISLs.
- **IOA Site**—Represents a set of switches within the physical fabric that is in a specific physical location. Multiple IOA sites within the physical fabric are typically interconnected over a MAN or WAN using Fibre Channel or FCIP links. IOA provides the acceleration service for flows traversing across sites. As a part of the IOA configuration, the switches must be classified into appropriate IOA sites. Acceleration

is provided for flows traversing the MAN or WAN across sites. The main reason to classify the sites is to select the intersite flows for acceleration. No intrasite flows will be allowed to participate in acceleration.



Note When using the CLI, only the switches where IOA is deployed need to be classified into a site. When using the Cisco DCNM-SAN, all the switches in a physical location need to be classified into a site. The site classification is used internally by the Cisco DCNM-SAN to automate the classification of the flows that traverse across sites.

- **IOA Interface**—Represents a single service engine in the MSM-18/4 Module or the SSN-16 Module. An IOA interface must be provisioined to enable IOA service on the service engine. The MSM-18/4 Module has one service engine and the SSN-16 Module has four service engines, which directly represents the number of IOA interfaces that can be created on these modules. In the CLI, an IOA interface is represented as **interface ioa** *x*/*y* where *x* represents the slot and *y* represents the service engine ID. With the SSN-16, the service engine ID can be 1 to 4. Each IOA interface requires a IOA license to be checked out.

  An IOA interface must be brought up administratively to enable the IOA service on the service engine.

- **IOA Switch**—Represents a switch that has one or more IOA interfaces configured for the IOA service. The terms IOA switch and IOA node are used interchangeably in this configuration guide.
- **IOA Cluster**—A set of IOA switches that can operate in a coordinated manner to provide the IOA service. An IOA cluster can only span two IOA sites. If there is a consolidation site that has connectivity to various other sites, each site pair must be represented by a unique IOA cluster. A switch may participate in multiple IOA clusters due to this reason, but each IOA interface is bound only to one IOA cluster. This architecture allows for cluster scalability and limiting the scope of configuration distribution as appropriate.
- **IOA N Port**—Represents a Fibre Channel N port represented by a port world-wide name. IOA requires that the site to which the N port belongs and the VSAN ID be configured. The site classification is required to identify how to redirect the traffic flow for acceleration.
- **FC-Redirect** —Fibre Channel Redirect (FC-Redirect) infrastructure provides the ability to redirect a flow to a specific service engine in the fabric to provide certain intelligent services such as Storage Media Encryption and Data Mobility Manager. This infrastructure has been extended for IOA to redirect the flow to two service engines in the fabric that can then work together to provide the acceleration intelligence.

Both the host and the target or tape must be directly attached to a FC-Redirect-capable switch.

- **IOA Flow**—A flow that is accelerated across the MAN or WAN by the IOA cluster. Each IOA flow is identified by initiator PWWN and target PWWN.

  IOA provide bidirectional acceleration for each configured flow. A separate reverse flow configuration is not required.

- **IOA Flow Group**—A set of IOA flows classified for a specific purpose. For example, if the same IOA cluster is being used for remote replication and backup, you can have all the replication flows classified into the replication flow group and all the backup flows classified into the backup flow group.

Note You can have more than one IOA service engine in the same site in the IOA cluster. This is the preferred configuration when an IOA service engine fails, then all the flows bound to it can be automatically moved to another available IOA service engine in the same site. This function is taken care of by the IOA cluster based load balancer.

Clustering

IOA is offered as a clustered service that consists of a set of switches that operates in coordination with each other. Clustering provides the following advantages:

- **Single point management**— IOA can be managed as a fabric service from a single switch. You do not need not configure multiple switches individually to provide IOA as a fabric service.
- **Automatic load-balancing**— You can provision all of the flows that need to be accelerated through IOA. Clustering allows these flows to load balance automatically across all the available IOA service engines within the cluster. It also makes it easy to plan for capacity as you only need to add an additional IOA service engine when you need to add more throughput within IOA.
- **Resiliency**— Allows automatic failover of the IOA flows whenever an IOA service engine fails on any of the switches. If a switch fails, an alternate switch in the cluster takes over the failed flows to maintain the continuity of the IOA service.

IOA clustering uses standard algorithms to provide consistency and reliability of the configuration metadata required for the service to be operational. A master switch is internally elected by the clustering infrastructure to perform certain tasks such as load-balancing and failover. To keep the process simple, we recommend that you provision the IOA from the master switch. If the network fails, which partitions the switches in a cluster, a standard majority node-based quorum algorithm is used to decide which partition should be operational to guarantee the consistency.

An internal node ID that is allocated as a part of adding the switches to the cluster is used in the master election algorithm. If you intend to manage IOA from a specific switch or a site, we recommend that you use this switch as a seed switch when a IOA cluster is configured, and also add all the nodes in this site before you add the nodes from the remote site into the IOA cluster.

Hardware Requirements

IOA is supported on the Cisco MDS 9000 Family 18/4-port Multiservice Module MSM-18/4, the Cisco MDS 9222i Switch, the Cisco MDS 9250i Switch, and the 16-Port Storage Services Node (SSN-16) module. Each MSM-18/4 Module, 9222i Switch, and 9250i Switch has one service engine that can be configured for the Cisco IOA service. The SSN-16 module has four service engines that can be used for the IOA service.

IOA scaling with NX-OS release 6.2(3) is supported only on Supervisor 2a module and not supported on Supervisor 2 module.

Software Requirements

To enable the IOA feature on the MSM-18/4 Module or SSN-16 Module, the MDS 9000 Family switch must run Cisco NX-OS Release 4.2(1) or later. To enable IOA on the MDS 9250i Switch it should run Cisco NX-OS

Release 6.2(5). You must also use Cisco DCNM-SAN 5.2(1) to manage the switches. Hosts must be connected to a switch running Cisco SAN-OS 3.3(1c) or later. Targets must be connected to a switch running Cisco NX-OS Release 4.2(1) or later.

The Cisco MDS 9710 Switch can have Cisco NX-OS Release 6.2(1) or 6.2(3) or above, if the hosts are directly connected to the Cisco MDS 9710 Switch that is participating in the IOA. But if the targets that are participating in IOA are directly connected to it (Cisco MDS 9710 Switch), then we must have Cisco NX-OS release 6.2(3) and above running on that Cisco MDS 9710 Switch.

After the target is connected to the Cisco MDS 9710 Switch with Cisco NX-OS Release 6.2(3) or above, IOA (master/slave) switch should be able to communicate with the Cisco MDS 9710 Switch (targets are connected) with at least one path where all FCR-capable switches are connected in between the IOA (master/slave) and Cisco MDS 9710 Switch.

License Requirements

The Cisco MDS 9000 Family IOA package is licensed per service engine and is tied to the chassis. The number of licenses required is equal to the number of service engines on which the intelligent fabric application is used.

IOA runs on the MDS 9222i Switch (native), MDS 9250i Switch (native) and on the MSM-18/4 Module and SSN-16 Module. The modules are supported in the MDS 9500 Directors and the MDS 9222i Switch.

On the SSN-16 Module, a separate license is required for each engine that will run IOA. Each SSN-16 engine configured for IOA checks out a license from the pool managed at the chassis level. For convenience, SSN-16 Module licenses can be purchased singly (the usual model) or in a package of four. Once they are installed into an MDS 9000 chassis, there is no difference between the IOA package of four and four single IOA licenses.

On the SSN-16 Module, because each engine is licensed independently, different licensed features can be configured on the four engines based on the following requirements for NX-OS Release 4.2(1):

- As with the MDS 9222i Switch and the MSM-18/4 Module, only one licensed feature can run on an engine at a time.
- On the SSN-16 Module, mix and match is supported for IOA and SAN Extension over IP in any combination (4+0, 1+3, 2+2, 3+1, or 0+4).
- Storage Media Encryption (SME) is not supported for mix and match in NX-OS Release 4.2(1).

To use the IOA features, Cisco MDS NX-OS Release 4.2(1) or later must be installed on a Cisco MDS 9000 Family switch and Cisco MDS9250i Switch must be running Cisco NX-OS Release 6.2(5).

Table 1: Cisco I/O Accelerator Licenses, on page 8 lists the available Cisco IOA licenses.

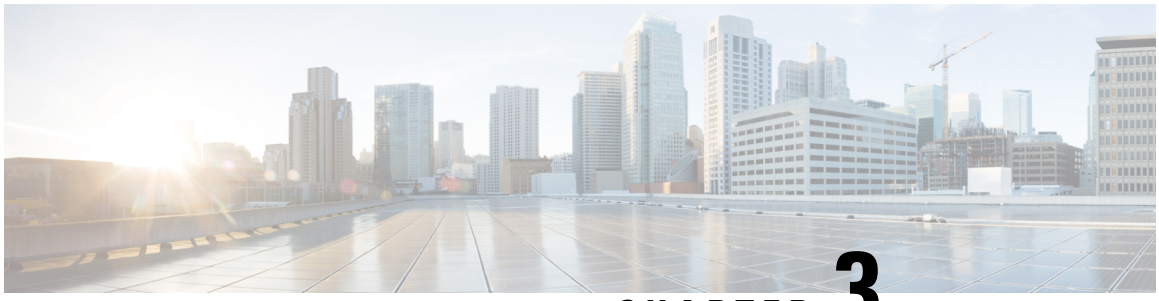
Table 1: Cisco I/O Accelerator Licenses

Part Number	Description	Applicable Product
M9250IIOA	Cisco I/O Acceleration License for MDS 9250i.	MDS 9250i
M92IOA184	Cisco I/O Acceleration License for MSM-18/4 on MDS 9200, spare.	MSM-18/4 on MDS 9200
M95IOA184	Cisco I/O Acceleration License for MSM-18/4 on MDS 9500, spare.	MSM-18/4 on MDS 9500

Part Number	Description	Applicable Product
M95IOASSN	Cisco IOA License (1 engine) for SSN-16 on MDS 9500, spare.	SSN-16 on MDS 9500
M92IOASSN	Cisco IOA License (1 engine) for SSN-16 on MDS 9200, spare.	SSN-16 on MDS 9200
M95IOASSN4X	Cisco IOA License (4 engines) for SSN-16 on MDS 9500, spare.	SSN-16 on MDS 9500
M92IOASSN4X	Cisco IOA License (4 engines) for SSN-16 on MDS 9200, spare.	SSN-16 on MDS 9200
M9222IIOA	Cisco I/O Accelerator License for MDS 9222i, spare.	MDS 9222i Switch



Note A device is either a switch or a module. When you enter the serial number for the device, make sure that you enter the serial number for the correct device; either the switch or the module for which you want to get the license. You can use the **show license** *host-id* command to find out which serial number to lock the license against.



CHAPTER 3

Getting Started

This chapter provides an overview of the basic configurations that need to be completed before getting started with IOA-specific configurations:

- Enabling SSH, on page 11
- Enabling CFS, on page 12
- IP Access Lists, on page 12
- Zone Default Policy, on page 12
- FC-Redirect, on page 12
- Configuring FC-Redirect v2 Mode, on page 14
- Using FC-Redirect with CFS Regions , on page 15
- Using IOA Cluster with IPFC Interface, on page 16

Enabling SSH

SSH needs to be enabled on all the IOA switches for Cisco DCNM-SAN to provision IOA. By default, the SSH service is enabled with the RSA key.

Procedure

Step 1 switch# **configure terminal**

Enters configuration mode.

Step 2 switch(config)# **feature ssh** updated

Enables the use of the SSH service.

For more information about the SSH service, refer to the *Cisco MDS 9000 Family NX-OS Security Configuration Guide*.

Enabling CFS

CFS must be enabled on the IOA switches as well as those switches of which the hosts and targets are directly connected to. FC-Redirect internally uses CFS to configure the rules for any given flow in the fabric.

Procedure

Step 1 switch# **configure terminal**

Enters configuration mode.

Step 2 switch(config)# **cfs distribute**

Enables (default) CFS distribution on the switch.

For more information about CFS, refer to the *Cisco MDS 9000 Family NX-OS System Management Configuration Guide* .

IP Access Lists

Cluster communication requires the use of the Management interface. IP ACL configurations must allow UDP and TCP traffic on ports 9333, 9334, 9335, and 9336.

Zone Default Policy

For FC-Redirect to work correctly, the default zone policy on all the switches in the IOA environment must be configured to deny and the initiator-target pairs must be configured in user-defined zones.

FC-Redirect

FC-Redirect Supported Switches

FC-Redirect (FCR) is supported on the following switches and is enabled by default:

- Cisco MDS 9250i Switch
- Cisco MDS 9148S Switch
- Cisco MDS 9396S Switch
- Cisco MDS 9706 Switch
- Cisco MDS 9710 Switch

- Cisco MDS 9718 Switch

Note Starting from Cisco MDS NX-OS Release 7.3(0)DY(1), FCR is supported on Cisco MDS 9148S and Cisco MDS 9396S switches.

FC-Redirect Unsupported Switches

IOA does not support any FCoE connected devices including devices connected through any of the MDS FCoE modules.

Starting from Cisco MDS NX-OS Release 5.2(x), you cannot install an FCoE module in a switch that is running DMM, SME, or IOA.

FC-Redirect is not supported on the following switches:

- Cisco MDS 9220i Switch
- Cisco MDS 9132T Switch
- Cisco MDS 9148T Switch
- Cisco MDS 9396T Switch

FC-Redirect Requirements

FC-Redirect requirements for IOA include the following:

- The MDS switch with the MSM-18/4 Module installed or the 9222i Switch needs to be running Cisco MDS NX-OS Release 4.2(1) or later or the 9250i Switch needs to be running Cisco MDS NX-OS Release 6.2(5).
- The targets must be connected to a FC-Redirect-capable switch running Cisco MDS NX-OS Release 4.2(1) or later.
- 32 targets per MSM-18/4 Module can be FC-Redirected.
- In FC-Redirect v2 mode, up to 128 hosts per target are supported. If you do not enable FC-Redirect v2, this is limited to 16 hosts per target.
- CFS is enabled by default. Ensure that the CFS is enabled on the switches that have the host and the target connected. Also ensure that the CFS is not disabled on switches that are part of the IOA cluster.
- Advanced zoning capabilities such as quality of service (QoS), logical unit number (LUN) zoning, and read-only LUNs must not be used for FC-Redirect hosts and targets.

Note FCR support for IVR is disabled by default. If FCR support for IVR is enabled on one switch in a fabric, it gets propagated through the entire fabric and is enabled on all the switches. We recommend that you disable FCR support for IVR using the **no fc-redirect ivr-support enable command** in global configuration mode.

Configuring FC-Redirect v2 Mode

To enable the v2 mode in FC-Redirect, use the **fc-redirect version2 enable** command in configuration mode. To disable the v2 mode in FC-Redirect, use the **no** form of the command.

This command is used to increase scalability of FC-Redirect. Disabling v2 mode after it is enabled in the fabric is not recommended. However, if you want to disable v2 mode, you cannot disable it until all FC-Redirect configurations are deleted. FC-Redirect configurations can be deleted only by deleting all corresponding application configurations.

The MDS switches not running Cisco MDS NX-OS Release 4.2(1) and later cannot be added to the fabric after the v2 mode is enabled. If the switches are added, all further FC-Redirect configuration changes will fail across the fabric. This could lead to traffic disruption for applications such as IOA, SME, and DMM.

Use the **show fc-redirect configs** command to see the list of applications that create FC-Redirect configurations.

If v2 mode is enabled in the fabric and you want to move a switch to a different fabric, use the **clear fc-redirect decommission-switch** command before moving the switch to a different fabric. If the mode is not enabled, all switches in the new fabric will be converted to v2 mode automatically.



Note Ensure that there are no fabric changes or upgrades in progress. For more information see Software Requirements, on page 7. Use the show **fc-redirect peer-switches** command (UP state) to see all the switches in the fabric.

Procedure

Step 1 Enter the following commands:

switch# **configure terminal**

switch(config)# **fc-redirect version2 enable**

Step 2 Enter **yes**.

```
Please make sure to read and understand the following implications
before proceeding further:
1) This is a Fabric wide configuration. All the switches in the
fabric will be configured in Version2 mode.Any new switches
added to the fabric will automatically be configured in version2
mode.
2) SanOS 3.2.x switches CANNOT be added to the Fabric after Version2
mode is enabled. If any 3.2.x switch is added when Version2 mode
is enabled, all further FC-Redirect Configuration changes will Fail
across the fabric. This could lead to traffic disruption for
applications like SME.
3) If enabled, Version2 mode CANNOT be disabled till all FC-Redirect
configurations are deleted. FC-Redirect configurations can be
deleted ONLY after all the relevant application configurations
are deleted. Please use the command 'show fc-redirect configs'
to see the list of applications that created FC-Redirect
configurations.
4) 'write erase' will NOT disable this command. After 'write erase'
```

```
on ANY switch in the fabric, the user needs to do:
'clear fc-redirect decommission-switch'
on that that switch. Without that, if the user moves the switch
to a different fabric it will try to convert all the switches
in the fabric to Version2 mode automatically. This might lead
to Error conditions and hence Traffic disruption.
Do you want to continue? (Yes/No) [No]Yes
```

Step 3 Enter **yes**.

```
Before proceeding further, please check the following:
1) All the switches in the fabric are seen in the output of
'show fc-redirect peer-switches' command and are in 'UP' state.
2) All switches in the fabric are running SanOS version 3.3.x or
higher.
3) Please make sure the Fabric is stable ie.,
No fabric changes/upgrades in progress
Do you want to continue? (Yes/No) [No] Yes
```

Using FC-Redirect with CFS Regions

The FC-Redirect feature uses Cisco Fabric Services (CFS) regions to distribute the FC-Redirect configuration. By default, the configuration is propagated to all FC-Redirect-capable switches in the fabric. CFS regions can be used to restrict the distribution of the FC-Redirect configuration.

Note Using FC Redirect with CFS regions is an optional configuration only if the number of switches in the SAN exceeds the scalability limit supported by IOA. As of MDS NX-OS Release 4.2(1), the number of switches supported in a fabric is 34.

To learn more about CFS regions, refer to the *Cisco MDS 9000 Family NX-OS System Management Configuration Guide*.

Guidelines for Designing CFS Regions For FC-Redirect

To design CFS regions for FC-Redirect, follow these guidelines:

- Ensure that the CFS region configuration for FC-Redirect can be applied to all FC-Redirect-based applications. The applications include Cisco SME, Cisco DMM, Cisco IOA, and any future applications.
- All switches in the region must have a common VSAN.
- For existing IOA installations, refer to Configuring CFS Regions For FC-Redirect, on page 16 for steps on migrating to CFS regions.
- Remove all instances of the previous configurations when a switch is moved to a region or moved out of a region.

Configuring CFS Regions For FC-Redirect

To configure the CFS regions for FC-Redirect, do the following tasks:

Procedure

Step 1 Configure a switch in the CFS region as shown in the following example:

```
switch# config t
switch(config)# cfs region 2
switch(config-cfs-region)# fc-redirect
switch(config)# end
```

Repeat this step for all the switches that are included in the specified region.

Step 2 Confirm that all the required switches are available in the CFS region by entering the **show fc-redirect peer-switches** command.

Step 3 To migrate existing Cisco IOA installations to CFS regions for FC-Redirect, delete all the existing FC-Redirect configurations created by the switches in other regions from each switch. To remove the configurations, perform the following steps:

a) Obtain a list of all FC-Redirect configurations by entering the **show fc-redirect configs** command.

b) Remove all configurations created by the switches in other regions by using the **clear fc-redirect configs** command. The configurations are removed from the switches but the switches remain active in the region in which they are created.

Using IOA Cluster with IPFC Interface

Internet protocol over Fibre Channel (IPFC) provides IP forwarding or in-band switch management over a Fibre Channel interface (instead of management using the Gigabit Ethernet mgmt 0 interface). You can use IPFC to specify that IP frames be transported over Fibre Channel using encapsulation techniques. IP frames are encapsulated into Fibre Channel frames so that cluster management information can transmit across the Fibre Channel network without using an overlay Ethernet network.

When an IOA cluster communicates via the IPFC interface, the cluster management messages can be sent and received on Fibre Channel ISLs by encapsulating cluster management messages in Fibre Channel frames instead of using the management interface.



Note

- Configuring IOA cluster with the IPFC interface is optional and is supported in Cisco MDS NX-OS Release 5.0(4c) or later. Support for GUI for configuring IOA cluster with the IPFC interface might be added in the future releases.
- You must configure the nodes in an IOA cluster either to use an IPFC interface or a management interface. We do not recommend using the combination of two interface configurations.

Task Flow for Configuring IOA Cluster To Use the IPFC Interface

To configure IOA cluster using the IPFC interface, follow these steps:

Procedure

Step 1 Create an IPFC interface.
a) Create a VSAN to use for in-band management.
b) Configure an IPv4 address and subnet mask for the VSAN interface.
c) Enable IPv4 routing.
d) Verify connectivity.

Step 2 Create an IOA cluster.

Step 3 Change the local node to use IPFC interface's IPv4 address.

Step 4 Add the IOA interfaces to the cluster.

Step 5 Add the remote node with IPFC interface IPv4 address.

Step 6 Add the IOA interface of the remote cluster.

Configuring an IOA Cluster To Use the IPFC Interface

The process of configuring an IOA cluster to use the IPFC interface involves a number of configuration tasks that should be completed in the following order:

Creating a VSAN Interface and Configuring IPv4 Addresses

The first step in the process of configuring an IOA cluster to use the IPFC interface is to create a VSAN interface and configure IPv4 addresses.

To create an interface VSAN, perform this task:

Procedure

Step 1 Switch# **configure terminal**

Enters configuration mode

Step 2 Switch(config)# **interface vsan 1**

Configures the interface for the specified VSAN (1).

Step 3 Switch (config-if)# **ip address 10.1.1.1 255.255.255.0**

Configures the IPV4 address and netmask for the selected interface.

Step 4 Switch (config-if)# **no shutdown**

Enables the interface.

Example

After creating the VSAN and configuring the IPv4 address, use the **show interface vsan** command to verify the configuration:

```
sw-231-14# show interface vsan 1
vsan1 is up, line protocol is up
    WWPN is 10:00:00:0d:ec:18:a1:05, FCID is 0xec03c0
    Internet address is 10.1.1.1/24
    MTU 1500 bytes, BW 1000000 Kbit
    0 packets input, 0 bytes, 0 errors, 0 multicast
    6 packets output, 384 bytes, 0 errors, 0 dropped
sw-231-14#
```

Enabling IPv4 Routing

To enable IPv4 routing, perform this task:

Procedure

Step 1 Switch# **configure terminal**

Enters configuration mode.

Step 2 **Switch(config)# ip routing**

Enables IPV4 routing.

Step 3 **Switch(config)# no ip routing**

Disables IPV4 routing.

Example

After enabling IPv4 routing, use the show ip routing to verify the configuration.

```
sw-231-14(config)# show ip routing
ip routing is enabled
```

Verifying Connectivity

To verify the connectivity, use the show ip route and ping commands.

```
sw-231-14# show ip route
Codes: C - connected, S - static
C 10.1.1.0/24 is directly connected, vsan1
```

```
sw-231-14# ping 10.1.1.2
PING 10.1.1.2 (10.1.1.2) 56(84) bytes of data.
64 bytes from 10.1.1.2: icmp_seq=1 ttl=64 time=0.875 ms
64 bytes from 10.1.1.2: icmp_seq=2 ttl=64 time=0.866 ms
64 bytes from 10.1.1.2: icmp_seq=3 ttl=64 time=0.884 ms
64 bytes from 10.1.1.2: icmp_seq=4 ttl=64 time=0.875 ms
--- 10.1.1.2 ping statistics ---
4 packets transmitted, 4 received, 0% packet loss, time 3023ms
rtt min/avg/max/mdev = 0.866/0.875/0.884/0.006 ms
```

Creating an IOA Cluster and IOA Interface in the Local Node

To create an IOA cluster and IOA interface in the local node, perform this task:



Note To configure an IOA cluster, you can use the name of the switch if the network supports DNS service. The IOA cluster requires switch name to IP address resolution.

Procedure

Step 1 Switch# **configure terminal**

Enters configuration mode

Step 2 Switch(config)# **ioa cluster cluster name**

Creates IOA cluster with specific name.

Step 3 Switch(config-ioa-cl)# **node switchname/ip address ip-address 10.1.1.1**

Adds or changes the node address from the mgmt0 address to the IPFC interface address.

Step 4 Switch(config-ioa-cl-node)# **int ioa 1/1**

Adds IOA interfaces to the cluster.

Verifying Cluster Configuration

To verify the cluster configuration, use the **show ioa cluster** *name* **node summary** command.

```
sw-231-14# sh ioa cluster  cltr1 node sum
--------------------------------------------------------------------------------
Switch             Site            Status          Master     Node ID
--------------------------------------------------------------------------------
sw-231-14(L)       site2           online          yes        1
```

To verify the IP address of the node, use the show ioa cluster name node command.

```
sw-231-14# show ioa cluster  cltr1 node
Node sw-231-14 is local switch
Node ID is 1
IP address is 10.1.1.1
Status is online
```

```
Belongs to Site site2
Node is the master switch
```

Adding a Remote Node and IOA Interface to the Remote Node

To add a remote node, perform this task:

Procedure

Step 1 Switch# **configure terminal**

Enters configuration mode.

Step 2 Swtich(config)# **ioa cluster** *cluster name*

Enter IOA cluster.

Step 3 Switch(config-ioa-cl)# **node <switchname/ip address> ip-address 10.1.1.2**

Adds remote node to the cluster with the IPFC interface address.

Step 4 Switch(config-ioa-cl-node)# **int ioa 4/1**

Adds IOA interfaces to the cluster.

Verifying the Cluster Configuration

To verify the node configuration, use the **show ioa cluster** *name* **node summary** command:

```
sw-231-14# show ioa cluster  cltr1 node summary
------------------------------------------------------------------------------
Switch             Site            Status          Master     Node ID
------------------------------------------------------------------------------
sw-231-14(L)       site2           online          yes        1
sw-231-19          site1           online          no         2
```

To verify the IP address of the node, use the **show ioa cluster** *name* **node** command:

```
Node sw-231-14 is local switch
  Node ID is 1
  IP address is 10.1.1.1
  Status is online
  Belongs to Site site2
  Node is the master switch
Node sw-231-19 is remote switch
  Node ID is 2
  IP address is 10.1.1.2
  Status is online
  Belongs to Site site1
  Node is not master switch
sw-231-14#
```

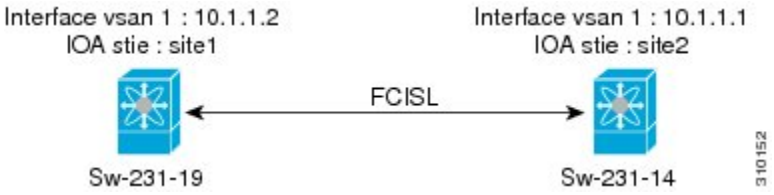
To see all of the configured interfaces in the IOA cluster, use the **show ioa cluster** *name* **interface summary** command:

```
sw-231-14# show ioa cluster cltr1 interface summary
--------------------------------------------------------------------------------
Switch             Interface         Status       Flows
--------------------------------------------------------------------------------
sw-231-14(L)       ioa1/1            up           0
sw-231-14(L)       ioa1/2            up           0
sw-231-19          ioa4/1            up           0
sw-231-19          ioa4/2            up           0
sw-231-14#
```

Configuration Example

This section includes an example for creating an IOA cluster using IPFC interface. Figure 2: Configuration Example, on page 21 illustrates the IOA cluster configuration used in this example. The sample topology shows the FC ISL between sw-231-14 and sw-231-19 switches.

Figure 2: Configuration Example



Creating an Interface VSAN

The following example creates an interface VSAN and configures IP address on sw-231-14 and enables IP routing:

```
sw-231-14(config)# int vsan 1
sw-231-14(config-if)# ip address 10.1.1.1 255.255.255.0
sw-231-14(config-if)# no shut
sw-231-14(config-if)# exit
sw-231-14(config)# ip routing
sw-231-14(config)#
```

The following example creates an interface VSAN and configure IP address on sw-231-19 and enables IP routing.

```
sw-231-19(config)# int vsan 1
sw-231-19(config-if)# ip address 10.1.1.12 255.255.255.0
sw-231-19(config-if)# no shut
sw-231-19(config-if)# exit
sw-231-19(config)# ip routing
```

Verifying the Configuration

The following example verifies the configuration of sw-231-14 using the **show interface** command.

```
sw-231-14# show interface vsan 1
vsan1 is up, line protocol is up
WWPN is 10:00:00:0d:ec:18:a1:05, FCID is 0xec03c0
Internet address is 10.1.1.1/24
MTU 1500 bytes, BW 1000000 Kbit
```

```
758 packets input, 110841 bytes, 0 errors, 42 multicast
651 packets output, 122577 bytes, 0 errors, 0 dropped
sw-231-14#
```

The following example verifies the configuration of sw-231-19 using the **show interface** command:

```
sw-231-19# show interface vsan 1
vsan1 is up, line protocol is up
WWPN is 10:00:00:05:30:01:9f:09, FCID is 0xc60000
Internet address is 10.1.1.2/24
MTU 1500 bytes, BW 1000000 Kbit
675 packets input, 124613 bytes, 0 errors, 36 multicast
755 packets output, 111785 bytes, 0 errors, 0 dropped
sw-231-19#
```

Verifying the Connectivity

The following example verifies the connectivity using the **ping** command:

```
sw-231-14# ping 10.1.1.2
PING 10.1.1.2 (10.1.1.2) 56(84) bytes of data.
64 bytes from 10.1.1.2: icmp_seq=1 ttl=64 time=0.868 ms
64 bytes from 10.1.1.2: icmp_seq=2 ttl=64 time=0.898 ms
64 bytes from 10.1.1.2: icmp_seq=3 ttl=64 time=0.906 ms
--- 10.1.1.2 ping statistics ---
3 packets transmitted, 3 received, 0% packet loss, time 2017ms
rtt min/avg/max/mdev = 0.868/0.890/0.906/0.038 ms
sw-231-14#
```

Configuring IOA Site on Switch sw-231-14

The following example configures IOA site on switch sw-231-14:

```
sw-231-14(config)# ioa site-local site2
sw-231-14(config)#
```

Configuring IOA Site on Switch sw-231-19

The following example configures IOA site on switch sw-231-19:

```
sw-231-19(config)# ioa site-local site1
sw-231-19(config)#
```

Configuring IOA Cluster cltr1 on Switch sw-231-14

The following example configures IOA cluster ctrl1 on switch sw-231-14:

```
sw-231-14(config)# ioa cluster  cltr1
2011 Apr  8 05:00:46 sw-231-14 %CLUSTER-2-CLUSTER_LEADER_ANNOUNCE: Node 0x1 is the new
Master of cluster 0x2e05000dec18a133 of 1 nodes
2011 Apr  8 05:00:46 sw-231-14 %CLUSTER-2-CLUSTER_QUORUM_GAIN: Cluster 0x2e05000dec18a133
now has quorum with 1 nodes
```

Changing the Node to Use IPFC Interface Address

The following example forces the node to use IPFC interface addresss:

```
sw-231-14(config-ioa-cl)# node sw-231-14 ip-address 10.1.1.1
sw-231-14(config-ioa-cl-node)# ex
```

Adding a Remote Node to the IOA Cluster

The following example adds a remote node to an IOA cluster:

```
sw-231-14(config-ioa-cl)# node sw-231-19 ip-address 10.1.1.2
2011 Apr  8 05:02:47 sw-231-14 %CLUSTER-2-CLUSTER_QUORUM_GAIN: Cluster 0x2e05000dec18a133
now has quorum with 1 nodes
2011 Apr  8 05:02:52 sw-231-14 %CLUSTER-2-CLUSTER_QUORUM_GAIN: Cluster 0x2e05000dec18a133
now has quorum with 2 nodes
sw-231-14(config-ioa-cl-node)# ex
```

Adding an IOA Interface to the Switch sw-231-14

The following example adds an IOA interfaces on the switch sw-231-14:

```
sw-231-14(config-ioa-cl)# node sw-231-14
sw-231-14(config-ioa-cl-node)# int ioa 1/1
sw-231-14(config-ioa-cl-node)# ex
```

Adding an IOA Interface to the Switch sw-231-19

The following example adds an IOA interface on the switch sw-231-19:

```
sw-231-14(config-ioa-cl)# node sw-231-19
sw-231-14(config-ioa-cl-node)# int ioa 4/1
sw-231-14(config-ioa-cl-node)# exit
```

Verifying the Cluster Configuration

The following example verifies the cluster configuration using the **show** *cluster name* **node summary** command:

```
sw-231-14# show ioa cluster cltr1 node summary
-------------------------------------------------------------------------------
Switch             Site            Status          Master     Node ID
-------------------------------------------------------------------------------
sw-231-14(L)       site2           online          yes        1
sw-231-19          site1           online          no         2
```

Verifying the IP Address

The following example verifies the IP address that is configured on the switch using the **show ioa cluster** *cluster name* **node** command:

```
sw-231-14# show ioa cluster  cltr1 node
Node sw-231-14 is local switch
  Node ID is 1
  IP address is 10.1.1.1
  Status is online
  Belongs to Site site2
  Node is the master switch
Node sw-231-19 is remote switch
```

```
Node ID is 2
IP address is 10.1.1.2
Status is online
Belongs to Site site1
Node is not master switch
```

Verifying the IOA Interface

The following example verifies the IOA interface that is configured on the switch using the **show ioa cluster** *cluster name* **interface summary** command:

```
sw-231-14# show ioa cluster cltr1 int summary
--------------------------------------------------------------------------------
Switch             Interface         Status       Flows
--------------------------------------------------------------------------------
sw-231-14(L)       ioa1/1            up           0
sw-231-14(L)       ioa1/2            up           0
sw-231-19          ioa4/1            up           0
sw-231-19          ioa4/2            up           0
sw-231-14#
```

Task Flow for Converting an Existing IOA Cluster to Use the IPFC Interface

To convert an existing IOA cluster to use the IPFC Interface, follow these steps:

1. Shut down IOA cluster on both the nodes.
2. Remove the IOA cluster that is configured on remote node.
3. Remove the remote node from the cluster on local switch and convert it as a single node cluster.
4. Change the local node to use IPFC by entering the commands **node id** *id nodename* and **ip-address** *IPFC address*.
5. Bring the single node cluster by **no-shut** on local switch.
6. Add the remote node and its interface on local switch.
7. Verify using **show** commands.

Configuration Example for Converting IOA Cluster to Use the IPFC Interface

Verifying the IOA Cluster Configuration

The following example verifies the IOA cluster configuration that is configured on the switch using **show ioa cluster** *cluster name* **node summary** command:

```
sw-231-14(config)# show ioa cluster  cltnew node summary
--------------------------------------------------------------------------------
Switch             Site              Status           Master     Node ID
--------------------------------------------------------------------------------
sw-231-14(L)       site2             online           yes        1
sw-231-19          site1             online           no         2
```

Verifying the IP Address

The following example verifies the IP address that is configured on the switch using the **show ioa cluster** *cluster name* **node** command:

```
sw-231-14(config)# show ioa cluster  cltnew node
Node sw-231-14 is local switch
  Node ID is 1
  IP address is 172.25.231.14
  Status is online
  Belongs to Site site2
  Node is the master switch
Node sw-231-19 is remote switch
  Node ID is 2
  IP address is 172.25.231.19
  Status is online
  Belongs to Site site1
  Node is not master switch
```

Verifying the Flow Status

The following example verifies the status of the flows using the **show ioa cluster** *cluster name* **flows** command. The nodes in this example are using mgmt0 interface address.

```
sw-231-14(config)# show ioa cluster  cltnew flows
--------------------------------------------------------------------------------
Host WWN,                VSAN     WA  TA  Comp  Status    Switch,Interface
Target WWN                                                Pair
--------------------------------------------------------------------------------
21:01:00:1b:32:22:55:df, 1        Y    Y    N   online    sw-231-14, ioa1/1
21:01:00:0d:77:dd:f8:9d, 1                                sw-231-19, ioa4/1
```

Shutting Down IOA Cluster on a Local Node

The following example shuts down the IOA cluster on a local node using **shut down** command:

```
sw-231-14(config)# ioa cluster cltnew
sw-231-14(config-ioa-cl)# shut down
This change can be disruptive. Please ensure you have read the IOA Cluster Recovery Procedure"
 in the configuration guide. -- Are you sure you want to continue? (y/n) [n] y
2011 Apr  8 05:36:41 sw-231-14 %CLUSTER-2-CLUSTER_LOCAL_NODE_EXIT: Local Node 0x1 has left
 the Cluster 0x2e06000dec18a133
```

Shutting Down the IOA Cluster on the Remote Node

The following example shuts down the IOA cluster on the remote node using the **shut down** command:

```
sw-231-19(config)# ioa cluster cltnew
sw-231-19(config-ioa-cl)# shut
This change can be disruptive. Please ensure you have read the IOA Cluster Recovery Procedure"
 in the configuration guide. -- Are you sure you want to continue? (y/n) [n] y
2011 Apr  8 05:37:03 sw-231-19 %CLUSTER-2-CLUSTER_LOCAL_NODE_EXIT: Local Node 0x2 has left
 the Cluster 0x2e06000dec18a133
sw-231-19(config-ioa-cl)# exit
```

Removing the IOA Cluster from the Remote Node

The following example removes the IOA cluster from the remote node using the **no ioa cluster** *cluster name* command:

```
sw-231-19(config)# no ioa cluster  cltnew
```

Note In case of removal of IOA cluster from the master node, clear all the FC-redirect configurations from the fabric, refer to "Configuring CFS Regions For FC-Redirect, on page 16".

Verifying the IOA Cluster in the Remote Node

The following example verifies the absence of the IOA cluster on the remote node using the **show ioa cluster** *cluster name* command:

```
sw-231-19(config)# show ioa cluster
sw-231-19(config)#
```

Removing the Remote Node from the Cluster in the Local Switch

The following example removes the remote node from the cluster in the local switch:

```
sw-231-14(config-ioa-cl)# no node sw-231-19
sw-231-14(config-ioa-cl)# show ioa cluster  cltnew node summary
------------------------------------------------------------------------------
Switch             Site            Status           Master     Node ID
------------------------------------------------------------------------------
sw-231-14(L)       --              unknown (cluster is offline)            1
```

Changing the Local Node Configuration To Use IPFC Address

The following example changes the local node to use the IPFC address:

```
sw-231-14(config-ioa-cl)# node id 1 sw-231-14 ip-address 10.1.1.1
sw-231-14(config-ioa-cl-node)# exit
```

Activating the Single Node Cluster

The following example activates the single node cluster:

```
sw-231-14(config-ioa-cl)# no shut
This change can be disruptive. Please ensure you have read the IOA Cluster Recovery Procedure"
 in the configuration guide. -- Are you sure you want to continue? (y/n) [n] y
sw-231-14(config-ioa-cl)# 2011 Apr  8 05:39:17 sw-231-14 %CLUSTER-2-CLUSTER_LEADER_ANNOUNCE:
 Node 0x1 is the new Master of cluster 0x2e06000dec18a133 of 1 nodes
2011 Apr  8 05:39:17 sw-231-14 %CLUSTER-2-CLUSTER_QUORUM_GAIN: Cluster 0x2e06000dec18a133
now has quorum with 1 nodes
```

Adding Remote Node with IPFC Address

The following example adds a remote node with IPFC address:

```
sw-231-14(config-ioa-cl)# node sw-231-19 ip-address 10.1.1.2
2011 Apr  8 05:39:36 sw-231-14 %CLUSTER-2-CLUSTER_QUORUM_GAIN: Cluster 0x2e06000dec18a133
now has quorum with 1 nodes
2011 Apr  8 05:39:41 sw-231-14 %CLUSTER-2-CLUSTER_QUORUM_GAIN: Cluster 0x2e06000dec18a133
now has quorum with 2 nodes
```

Adding IOA Interfaces to the Remote Node

The following example adds the IOA interfaces to the remote node:

```
sw-231-14(config-ioa-cl-node)# int ioa 4/1
sw-231-14(config-ioa-cl-node)# end
sw-231-14#
```

Verifying the Cluster Nodes

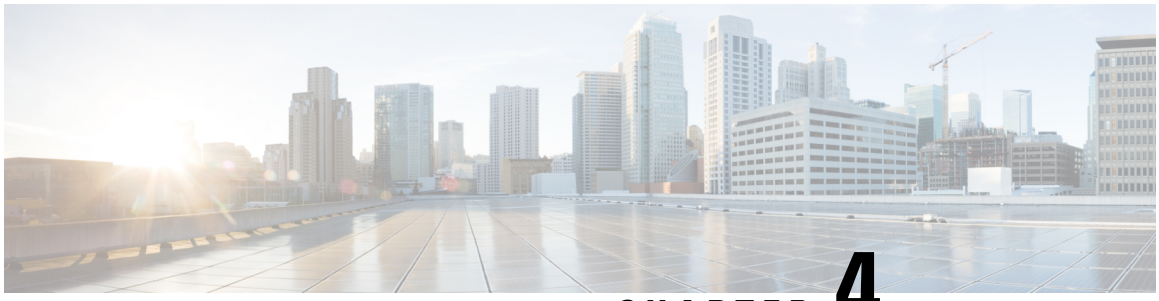
The following example verifies the status of the IOA clusters using the **show ioa cluster** *cluster name* **node summary** command:

```
sw-231-14# show ioa cluster  cltnew node summary
------------------------------------------------------------------------------
Switch             Site            Status           Master     Node ID
------------------------------------------------------------------------------
sw-231-14(L)       site2           online           yes        1
sw-231-19          site1           online           no         2
```

Verifying the Flow Status

The following example verifies the status of the IOA clusters using the **show ioa cluster** *cluster name* **flows** command:

```
sw-231-14# show ioa cluster  cltnew flows
------------------------------------------------------------------------------
Host WWN,                VSAN     WA  TA  Comp  Status    Switch,Interface
Target WWN                                                Pair
------------------------------------------------------------------------------
21:01:00:1b:32:22:55:df, 1        Y    Y    N   online    sw-231-14, ioa1/1
21:01:00:0d:77:dd:f8:9d, 1                                sw-231-19, ioa4/1
sw-231-14#
```



CHAPTER 4

Deployment Considerations

This chapter describes the requirements and guidelines that are necessary to successfully deploy your Cisco I/O Accelerator SAN. Read this chapter before installing or configuring Cisco I/O Accelerator (IOA).

This chapter includes the following sections:

- Supported Topologies, on page 29
- Deployment Guidelines, on page 33
- Guidelines and Limitations, on page 34
- Configuration Limits, on page 37

Supported Topologies

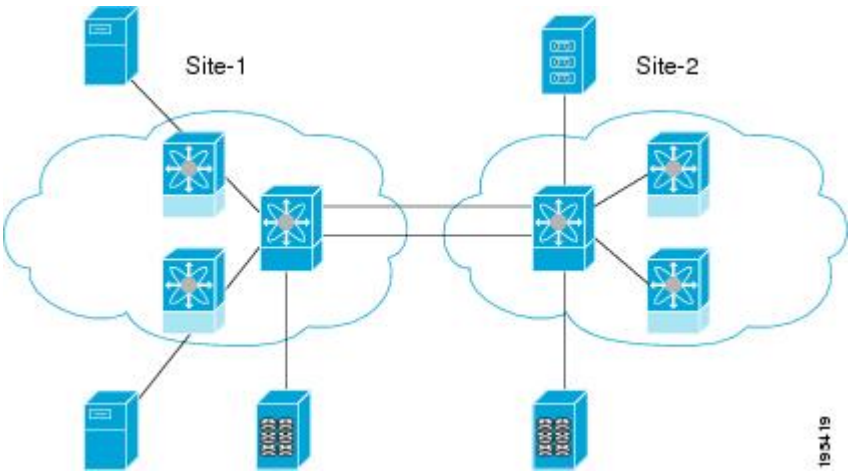
This section includes the following topics:

Core-Edge Topology

Figure 3: Core-Edge Topology, on page 30 illustrates the core-edge topology where you are recommended to place the IOA interfaces (24/10 port SAN Extension Module and 9250i Switch) in the core switches that interconnect the two sites. The ISLs interconnecting the two sites over a MAN or WAN are typically on the core switches as well, so this becomes a natural place to deploy the IOA service. This deployment provides the following benefits:

- Provides consolidation of IOA service at the core.
- Allows easy scalability of the IOA service engines based on the desired throughput.
- Allows you to plan and transition from FC or FCIP acceleration solutions to IOA. This is because these acceleration solutions will likely be deployed at the core switches already and will allow for a smooth transition to IOA.
- Facilitates planning the capacity based on WAN ISL throughput on the core switches themselves.
- Provides optimal routing as the flows have to traverse these core switches to reach the remote sites.

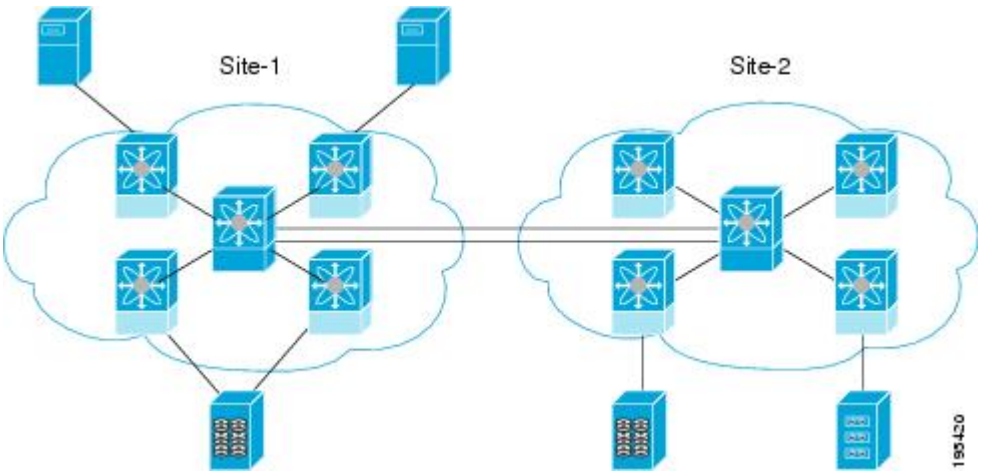
Figure 3: Core-Edge Topology



Edge-Core-Edge Topology

Figure 4: Edge-Core-Edge Topology, on page 30 illustrates the edge-core-edge topology where you are recommended to place the Cisco MDS 24/10 port SAN Extension Module and Cisco MDS 9250i Switch at the core switches that interconnect the two sites.

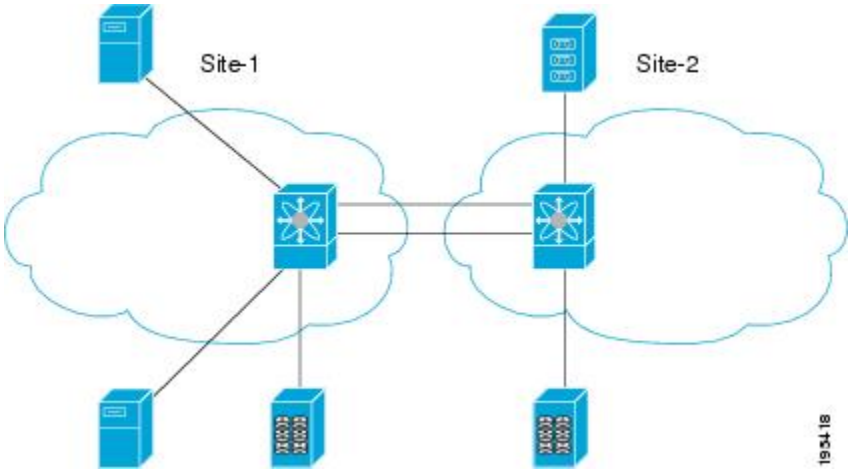
Figure 4: Edge-Core-Edge Topology



Collapsed Core Topology

Figure 5: Collapsed Core Topology, on page 31 illustrates the collapsed core toplogy where you are recommended to place the Cisco MDS 24/10 port SAN Extension Module or Cisco MDS 9250i Switch (IOA interfaces) in the core switches that interconnect the two sites.

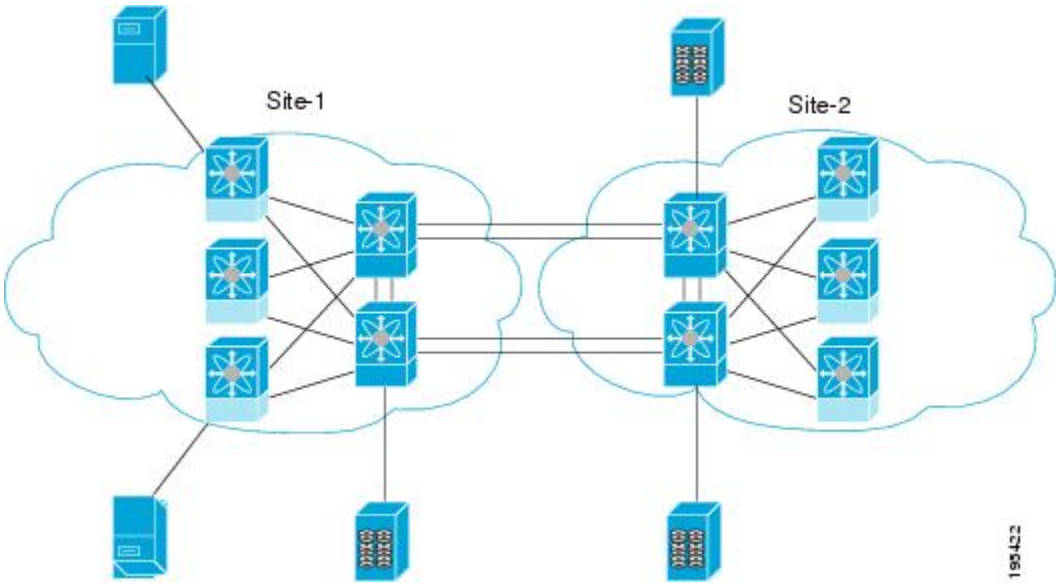
Figure 5: Collapsed Core Topology



Extended Core-Edge Topology

Figure 6: Extended Core-Edge Topology, on page 31 illustrates the extended core-edge topology where you are recommended to place the IOA interfaces Cisco MDS 24/10 port SAN Extension Module or Cisco MDS 9250i Switch) in all the core switches. As the IOA service load balances the traffic by selecting any IOA interface from each site and forms the IOA interface pair for a given flow, certain failures may result in suboptimal routing. The recommendation is to interconnect the core switches within each site for maximum availability of the IOA service. The ISLs between the core switches in the specific site has as much throughput as the WAN ISLs between the sites.

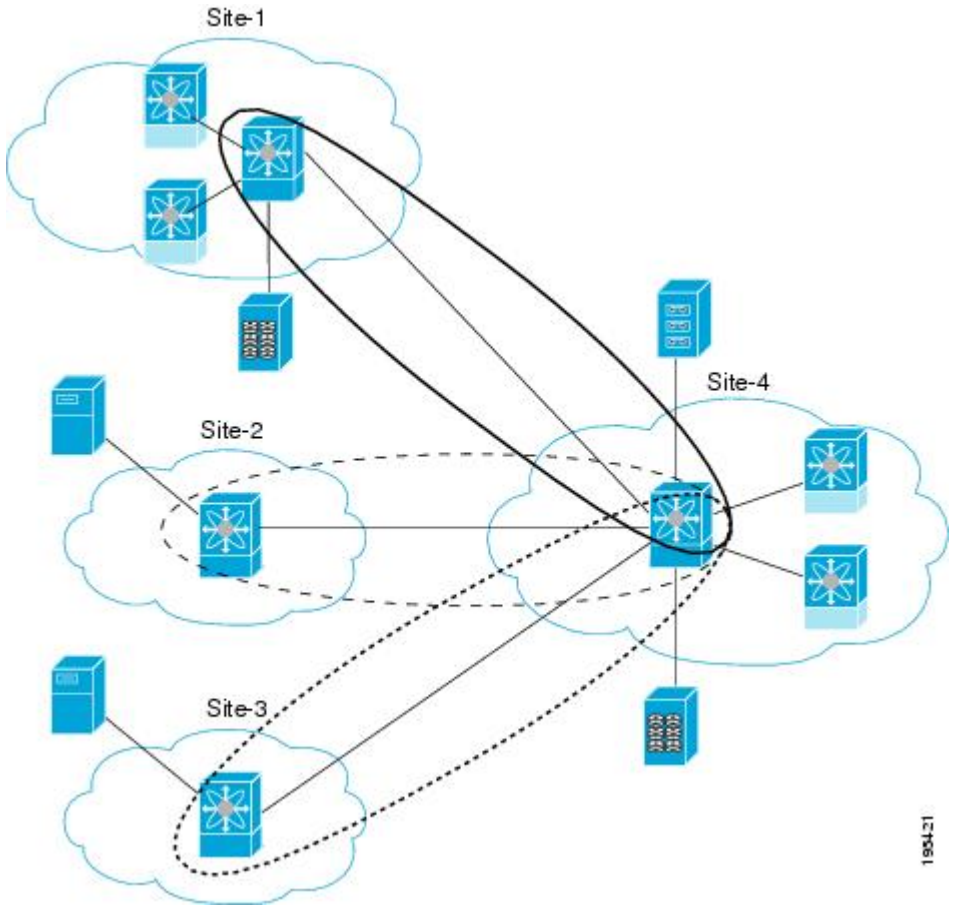
Figure 6: Extended Core-Edge Topology



Extending Across Multiple Sites

Figure 7: Extended Across Multiple Sites, on page 32 illustrates the IOA implementation where the IOA service is extended across multiple sites. In this example, Site-4 consolidates the tape backup from Site-1, Site-2, and Site-3. Each IOA cluster represents a site pair, which means that there are three unique clusters. This topology provides segregation and scalability of the IOA service across multiple sites. In Site-4, a single switch participates in multiple IOA clusters.

Figure 7: Extended Across Multiple Sites



IVR Topologies

Note Starting from Cisco MDS NX-OS Release 6.2(1), IOA with IVR is not supported.

For IOA to support IVR flows, we recommend that you place the IOA interfaces on the Cisco MDS 24/10 port SAN Extension Module or Cisco MDS 9250i Switch in the IVR border switches for optimum routing. IOA must always be deployed on the host and target VSANs. Packets from the host get redirected to the IOA interface in the host VSAN, traverses the IVR transit VSANs for routing, and again gets redirected to the IOA

interface in the target VSAN before it reaches the target and vice-versa. IVR transit VSANs are used only for FC routing. IOA is not supported or deployed on transit VSANs.

For more information, refer to the *Cisco MDS 9000 Family NX-OS Inter-VSAN Routing Configuration Guide*.

Other Topologies

In certain other topologies, the edge switches are connected across the WAN. With these topologies, we recommend that you do the following:

- Transition the WAN links from the edge to core switches to provide consolidation and optimal routing services.
- Deploy the IOA service in the core switches.

Deployment Guidelines

General Guidelines

When you deploy IOA, consider these general configuration guidelines:

- The IOA flows bound to the IOA interfaces on the module undergoing an upgrade will be affected.
- Clustering infrastructure uses the management IP network connectivity to communicate with the other switches. In the case of a switchover, the management IP network connectivity should be restored quickly to preserve the cluster communication. If the management port is connected to a Layer 2 switch, spanning tree must be disabled on these ports. In a Cisco Catalyst 6500 Series switch, you can implement this disabling by configuring the **spanning-tree portfast** command on these ports which considers these ports as access or host ports.

Scalability and Optimal Performance Considerations

For maximum scalability and optimal performance, follow these IOA configuration guidelines:

When you configure IOA, consider the following Zoning requirements:

- In certain tape backup environments, a common practice is to zone every backup server with every tape drive available to allow sharing of tape drives across all the backup servers. For small and medium tape backup environments, this may be retained when deploying IOA. For large backup environments, the scalability limit of number of flows in IOA must be considered to check if the zoning configuration can be retained. Best practice for such an environment is to create multiple tape drive pools, each with a set of tape drives and zones of only a set of backup servers to a particular tape drive pool. This practice sharing of tape drives drastically reduces the scalability requirements on IOA.
- Deploy IOA interfaces (Cisco MDS 24/10 port SAN Extension Module or Cisco MDS 9250i Switch) in the core switches in both core-edge and edge-core-edge topologies.
- When multiple core switches are interconnected across the MAN or WAN, do the following:
  - Deploy the IOA interfaces equally among the core switches for high availability.
  - Interconnect core switches in each site for optimal routing.
- Plan for Geneneration 2 and above line cards to avoid any FC-Redirect limitations.

- There is a limit of only 32 targets per switch if Generation 1 modules are used to link the ISLs connecting the IOA switch and target switches or if the host is directly connected to a Generation 1 module.
- Depending on the WAN transport used, you may have to tune the Fibre Channel extended B2B credits for the round-trip delay between the sites.

Resiliency Considerations

When you configure IOA, consider the following resiliency guidelines:

- Plan to have a minimum of one additional IOA service engine for each site for handling IOA service engine failures.
- Plan for E_D_TOV: Fibre Channel Error Detect Timeout Value (E_D_TOV) is used by Fibre Channel drives to detect errors if any data packet in a sequence takes longer than the specified timeout value. The default timeout value for E_D_TOV is 2 seconds. IOA has an built-in reliability protocol (LRTP) to detect and recover from ISL failures by doing the necessary retransmissions. However, you need to ensure that it recovers before the expiry of E_D_TOV. LRTP is not required if the FCP-2 sequence level error-recovery procedures are enabled end-to-end (primarily in the tape drivers) because this helps to recover from timeout issues.
- When the FCP-2 sequence level error-recovery procedure is not enabled, you must tune certain timers in order to protect the site from ISL failures.
  - Reduce the LRTP retransmit value from the default value of 2.5 seconds to 1.5 seconds. For more information, see the Setting the Tunable Parameters, on page 104 .
  - If the ISLs are FCIP links, the FCIP links must be tuned in order to detect link flaps quickly. By default, FCIP links detect a link failure in 6 seconds based on TCP maximum retransmissions. To reduce the time taken to detect failures, you need to set the maximum retransmission attempts in the FCIP profile from the default value of 4 to 1.



Caution Modifying the default setting to a lower value results in quick link failure detections. You must make sure that this is appropriate for your deployment. We recommend that you modify the default setting only for those applications which are sensitive to E_D_TOV values. For other applications, the default configuration is sufficient.

Guidelines and Limitations

When you configure IOA, consider the following guidelines and limitations:

- Starting from Cisco MDS NX-OS Release 6.2(1), IOA with IVR is not supported. Before configuring IOA, disable IVR support for Fibre Channel Redirect (FCR) using the **no fc-redirect ivr-support enable** command in global configuration mode.
- Only 512 flows are supported when IOA and IVR co-exists.
- IOA decides the master based on a master election algorithm. If you have multiple switches in the IOA cluster, you must add all the switches in the site that you manage through the cluster before adding switches from the remote site.
- IOA clustering framework uses IP connectivity for its internal operation. If an IOA cluster becomes nonoperational due to IP connectivity, IOA flows are brought down to offline state. In this state, the hosts

may not be able to see the targets. To accelerate the IOA flows, the IOA cluster must be operational and there must be at least one IOA switch in each site that is online within this IOA cluster.

- If there are multiple IOA clusters in a region, a target can be part of the IOA configuration in only one cluster. To change the target to a different cluster, the configuration in the first cluster must be deleted before creating the configuration in the second cluster.
- IOA licenses are not tied to a specific IOA service engine. IOA licenses are checked out when any of the following event occurs:
  - An IOA interface is configured.
  - A line card that contains the IOA interface comes online. There are no links between an IOA license and a IOA service engine. If a line card goes offline, another IOA interface can be brought up using the same IOA license. In such cases, when the line card comes back online, the IOA interface is automatically brought down with status displaying No License. You need to install licenses corresponding to the number of IOA interfaces configured regardless of the status of the line card.
- If IOA flows are configured and a copy running to startup is not performed, FCR rules are removed automatically during a reload for all flows in all VSANs except VSAN 1. VSAN 1 is a default VSAN that is always persistent even without a copy running to startup and so FCR rules are preserved for this VSAN. To recover from this, you can enter the **clear fc-redirect decommission-switch** command before rebooting the switch to purge the FCR configurations in VSAN 1. Alternately, you can clean up the entire IOA flow configuration before rebooting the switch.
- If an MDS switch is connected through an ISL using a DS-X9248-96K9 line card module and the targets are connected to the MDS switch, then this MDS switch can connect to a maximum of 160 targets. This is because the maximum number of Extended Link Service (ELS) entries on the DS-X9248-96K9 line card module is 320 entries. For example, in an IOA configuration that has 5 flows (1 host: 1 target) you can have 10 ELS entries on a module with ISL and in a IOA configuration that has 10 flows (2 hosts: 1 target), you only can have 10 ELS entries. This is because ELS entries depends on the number of targets. The workaround for this situation is to implement an allowed VSAN on ISL. For example, if ISL-1 is connected to module 9 and is limited to VSAN 2000, then all the ELS entries specific to VSAN 2000 will be on module 9. If ISL-2 is connected to module 2 and is limited to VSAN-3000, then all the ELS entries specific to targets of VSAN-3000 will be on module 2.
- When IOA is used to accelerate the EMC SRDF family of products, switching between SRDF adaptive Copy and SRDF/A can cause your RDF pairs to go into TransIdle state. If your SRDF deployment requires switching between these two modes, we recommend that you use the FCIP write acceleration feature instead of IOA.
- IOA flow takes a few seconds to become active upon certain triggers such as host or target port flaps. Port login (PLOGI) from the hosts are buffered until the IOA flow becomes active. Once the IOA flow becomes active, it sends a Registered State Change Notification (RSCN) to request the host to PLOGI again. Certain target arrays perform a few back-to-back PLOGIs before the flow becomes active and when determining that a failure requires a manual corrective action. To prevent this, IOA flows that have been configured for write acceleration are set up with a default timeout of 10 seconds after which the flow becomes unaccelerated. This is useful specifically in cases where IOA is unable to take over the flow before the timeout. For example, the line card reloads where no other IOA interface is available to handle the flow. In certain target arrays, the 10-second timeout is not sufficient and these arrays may require manual recovery using the storage management interfaces. One example of this target array is HDS AMS.

The workaround for this situation is to set the timeout to 5 seconds using the CLI command **tune wa-fcr-rule-timeout 5** under the IOA cluster configuration submode. This configuration is cluster-wide persistent across reboots.

- If the MDS 9250i Switch is part of the cluster as an IOA node, then the maximum number of flows supported is 203 in one VSAN. If multiple VSAN are used, the maximum number of flows is 256.
- If the MDS 9250i Switch is connected through an ISL and the targets are connected to that ISL, then the MDS switch can connect to a maximum of 203 targets. This is because the maximum number of ELS entries on a MDS9250i switch is 406 entries. The 203 targets involved in IOA includes all VSANs. The 203 target limit exists if no IVR entries are programmed. In case of IVR, the corresponding number of targets will reduce accordingly depending upon the availability of the ELS region.
- ISSU on the Cisco MDS9250i with IOA disk flows greater than 180 flows is not supported.
- In a 4-node IOA cluster of the Cisco MDS 9250i Switch having 1020 flows with a 3:1 ratio and hosts or targets being distributed equally, you may get the following message:

```
%ACLTCAM-2-ACL_TCAM_NO_TCAM_LEFT: ACLTCAM resource exhausted for interface on fcx/y.
```

The above message indicates that the ACLTCAM usage for Region2 Security on the Cisco MDS 9250i Switch or the Cisco MDS 9148S Switch is full. Due to this, a few IOA flows may be offline. This is the expected behavior. In such a case, ensure that the number of flows that get bound to the IOA node on the Cisco MDS 9250i Switch is not more than 203.

If either the hosts or the targets that are participating in IOA are connected to Cisco MDS 9148S Switch, then the maximum number of hosts or targets you can have is 203.

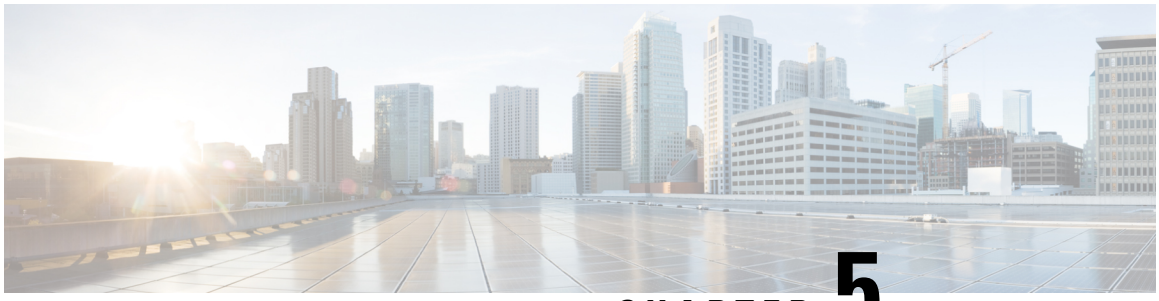


Note To view the ACLTCAM usage on the Cisco MDS 9250i Switch or the Cisco MDS 9148S Switch, use the **show system internal acltcam-soc tcam-usage** command.

- If FCIP is configured on a Cisco MDS 9700 Series switch and IOA VSAN is a part of that FCIP tunnel, then the targets participating in IOA cannot be locally present on that switch.
- When an IOA engine is reset, traffic in flows accelerated by it will revert to unaccelerated exchanges. This occurs during normal switch operations, namely—ISSU, ISSD, or reload of a module or switch hosting the IOA engine.
- Compression between Cisco MDS 24/10 port SAN Extension Module and Cisco MDS 9250i Switch works only when the switch is running in Cisco MDS NX-OS Release 8.2(1).
- From Cisco MDS NX-OS Release 8.2(1), there is no interoperability between Cisco MDS 24/10 port SAN Extension Module and Cisco MDS 18/4-Port Multiservice Module (MSM).
- From Cisco MDS NX-OS Release 8.2(1), there are only two IOA interfaces on a single Cisco MDS 24/10 port SAN Extension Module.
- From Cisco MDS NX-OS Release 8.2(1), FCIP and IOA together are not supported on a single Cisco MDS 24/10 port SAN Extension Module.

Configuration Limits

See the Cisco MDS NX-OS Configuration Limits document for information on the configuration limits for IOA.



CHAPTER 5

Configuring IOA Using Cisco DCNM-SAN

This chapter describes how to configure I/O Accelerator (IOA) using Cisco DCNM-SAN.

This chapter contains the following sections:

- IOA Manager , on page 39
- Launching IOA Manager, on page 41
- Configuring Sites, on page 41
- Configuring Clusters, on page 46
- Configuring Interfaces, on page 50
- Configuring Flows, on page 52

IOA Manager

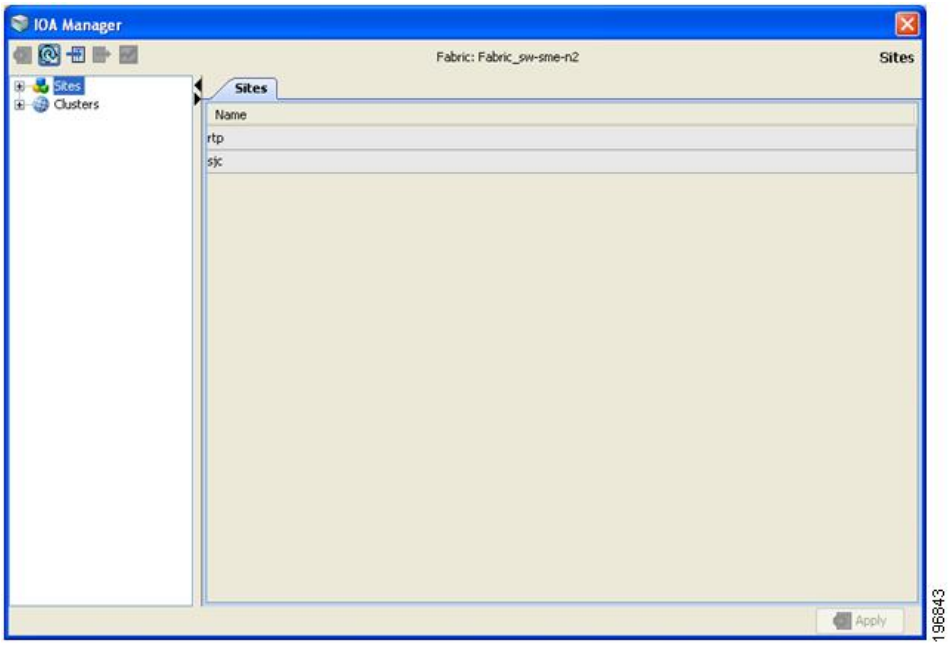
The IOA Manager is a graphical user interface (GUI) for configuring and managing IOA. The IOA Manager user interface consists of a navigation pane on the left that displays a hierarchy and an information pane on the right that displays the contents of the item that you click in the navigation pane. The hierarchy is a tree structure that contains elements that you can configure with IOA Manager. It also consists of a toolbar for quick access to the most commonly used options and a Fabric drop-down list box. The Fabric drop-down list box allows you to directly access the fabrics managed by Cisco DCNM-SAN. The Fabric drop-down list box will be available only if more than one fabric is open.

Note

- Cisco DCNM-SAN Client standalone supports IOA Manager from Release 5.0(1a).
- When you perform some of the time-consuming configuration activities using IOA Manager, the progress bar indicates that the configuration actions are in progress. You need to wait until the action is complete. You can click Stop to cancel the action. However, stopping the action may not roll back the transactions that were executed.

Figure 8: IOA Manager Main Window, on page 40 shows the IOA Manager interface.

Figure 8: IOA Manager Main Window



Toolbar

The IOA Manager main toolbar provides icons for accessing the most commonly used operations as shown in Table 2: I/O Accelerator Toolbar, on page 40.

Table 2: I/O Accelerator Toolbar

Icon	Description
	Applies the changes.
	Refreshes the window.
	Adds a cluster or interface.
	Deletes an existing entry.
	Displays a real-time chart of the selected switch.

Launching IOA Manager

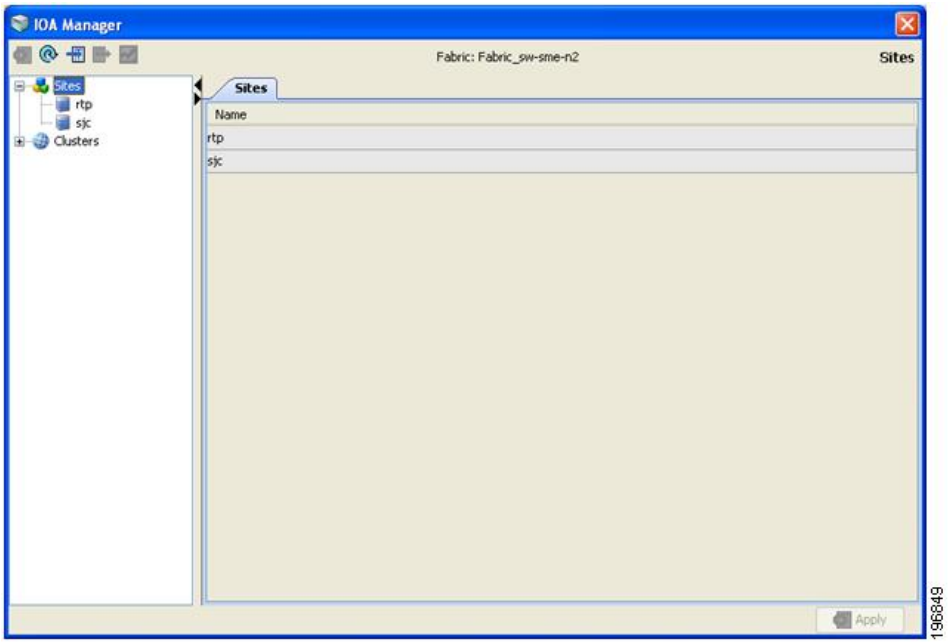
To launch IOA Manager, follow the below step:

Procedure

Choose **Tools** > **I/O Acceleration**.

You see the Cisco DCNM-SAN main window as shown in Figure 9: Cisco DCNM-SAN Window, on page 41.

Figure 9: Cisco DCNM-SAN Window



Note
When you select IOA Manager, it opens the tree for the fabric that is selected. If there is no active fabric, IOA Manager launches with the first fabric in the tree.

Configuring Sites

A site is described as a named set of switches. You can click the sites node to view the list of defined sites. There are two tables in the information pane: one for the assigned switches on the top and the another one for unassigned switches below the assigned switches table. You can click the name of the site to display the details in the information pane. Only active sites can be used for creating a clusters.

Adding a New Site

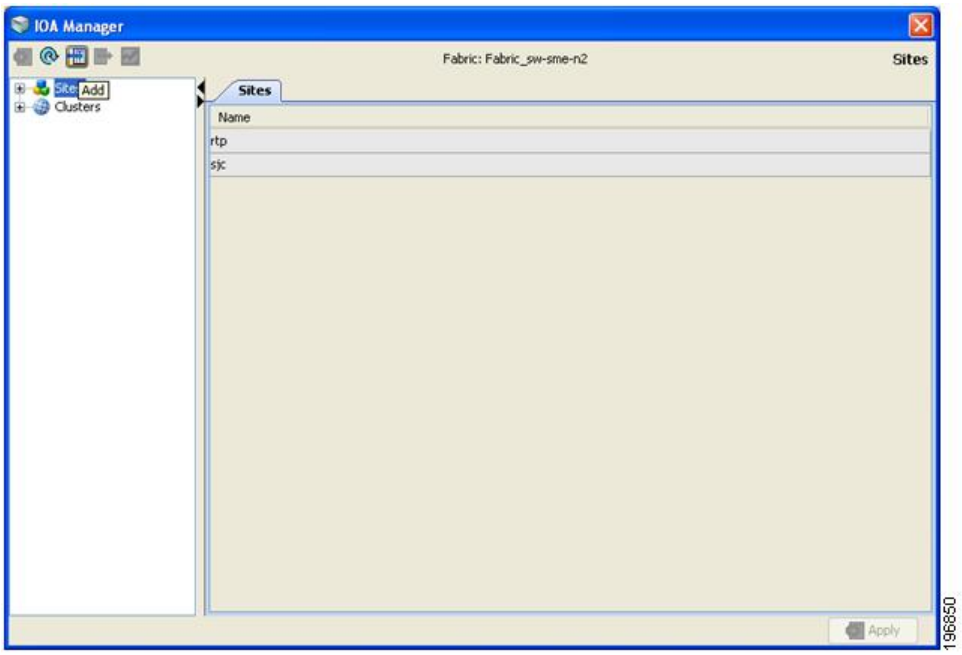
To create a new site using IOA Manager, follow these steps:

Procedure

Step 1 Select Sites in the navigation pane.

You see the IOA Manager window as shown in Figure 10: IOA Manager, on page 42.

Figure 10: IOA Manager



Step 2 Click the Add icon on the toolbar.

You see the site name dialog box as shown in Figure 11: Site Name Dialog Box, on page 42.

Figure 11: Site Name Dialog Box



Step 3 Enter the site name and then click OK.

You see the select switch dialog box as shown in Figure 12: Select Switch Dialog Box, on page 42.

Figure 12: Select Switch Dialog Box



Step 4 Select a switch from the drop-down list box and then click OK.

Step 5 Click OK in the dialog box to confirm that you have successfuly created the site.

Removing a Site

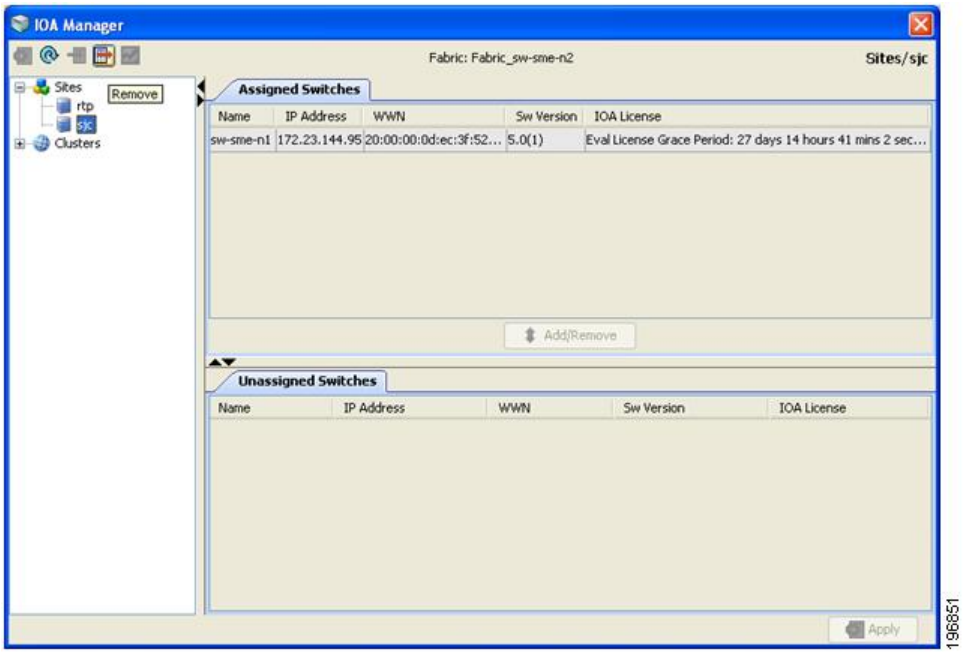
To remove a site using IOA Manager, follow these steps:

Procedure

Step 1 In the navigation pane, click the name of the site you want to delete.

You see the IOA Manager window as shown in Figure 13: IOA Manager Window, on page 43.

Figure 13: IOA Manager Window



Step 2 Click the **Remove** icon on the toolbar.

You see the confirmation dialog box in as shown in Figure 14: Delete Confirmation Dialog Box, on page 43.

Figure 14: Delete Confirmation Dialog Box



Step 3 Click **Yes** to confirm that you want to remove the site.

Viewing a Site

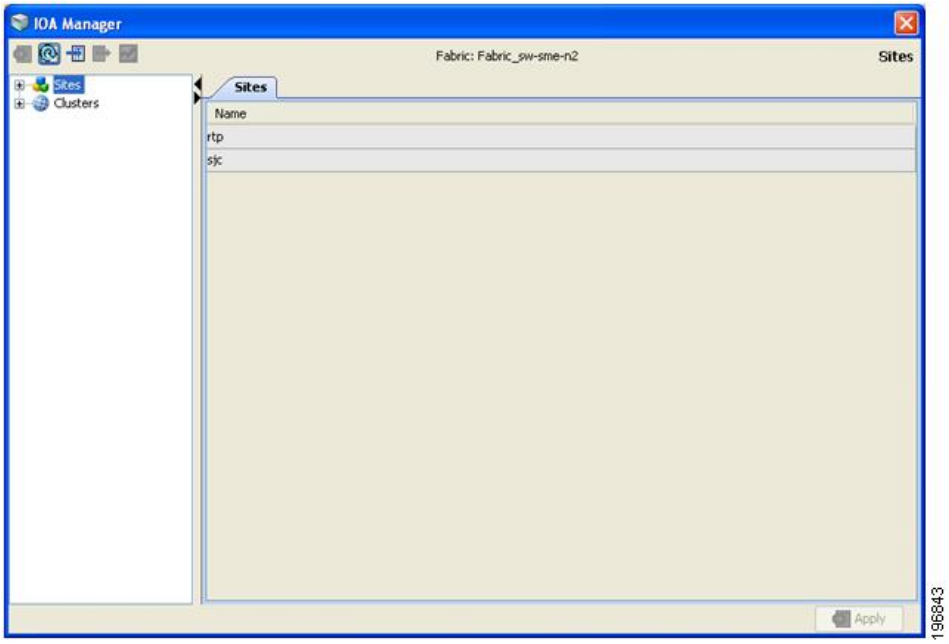
To view a site using IOA Manager, follow these steps:

Procedure

Step 1 In the navigation pane, click Sites.

You see the IOA Manager window as shown in Figure 15: Viewing Sites Using IOA Manager, on page 44.

Figure 15: Viewing Sites Using IOA Manager

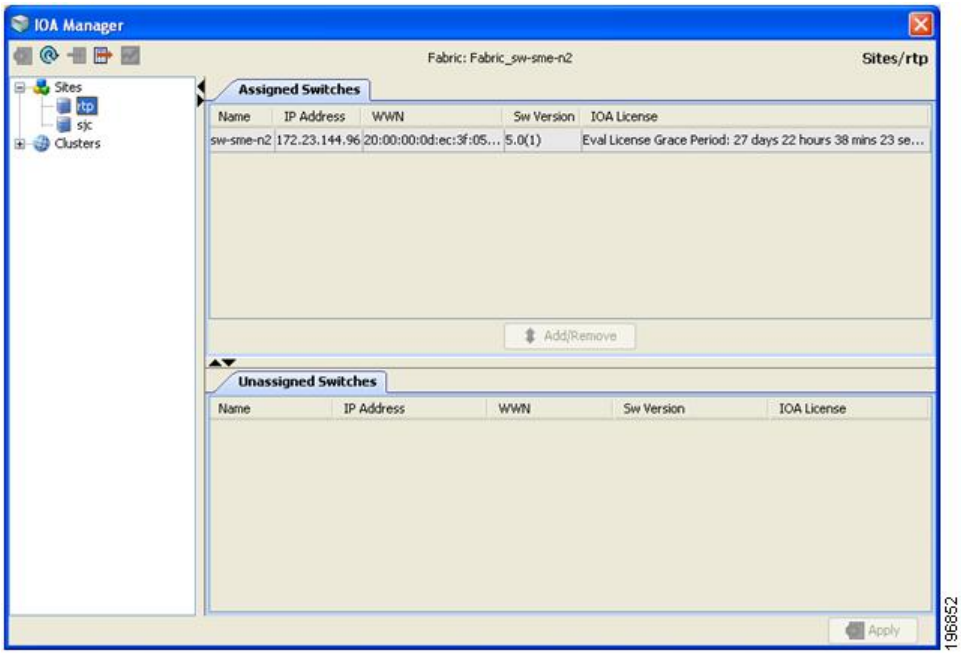


Step 2 Expand the sites in the hierarchy.

Step 3 Click the name of the site to view the details in the information pane.

You see the site details as shown in Figure 16: Viewing Site Details Using IOA Manager, on page 45.

Figure 16: Viewing Site Details Using IOA Manager



Adding Switches to a Site

To add a switch to a site, follow these steps:

Procedure

Step 1 In the navigation pane, click **Sites**.

Step 2 Select the switches that you want to add from the Unassigned Switches table.

Step 3 Click **Add**, and then click **Apply**.

Removing Switches from a Site

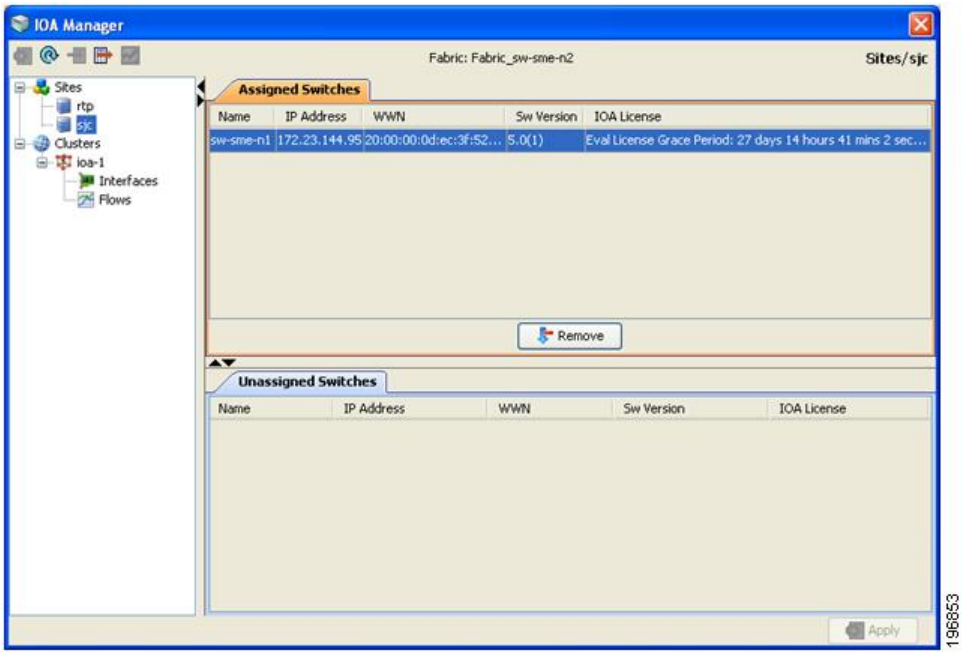
To remove a switch from a site, follow these steps:

Procedure

Step 1 In the navigation pane, click Sites.

You see the IOA Manager window as shown in Figure 17: Removing Switches from a Site Using IOA Manager, on page 46.

Figure 17: Removing Switches from a Site Using IOA Manager



Step 2 Click to select the switches you want to remove from Assigned Switches table.

Step 3 Click Remove, and then click Apply.

Configuring Clusters

You can select a cluster to see the details in the information pane. The upper table in the information pane displays the members of a named cluster, and the table below displays the statistical information about the cluster's active IOA interfaces.

Adding a New Cluster

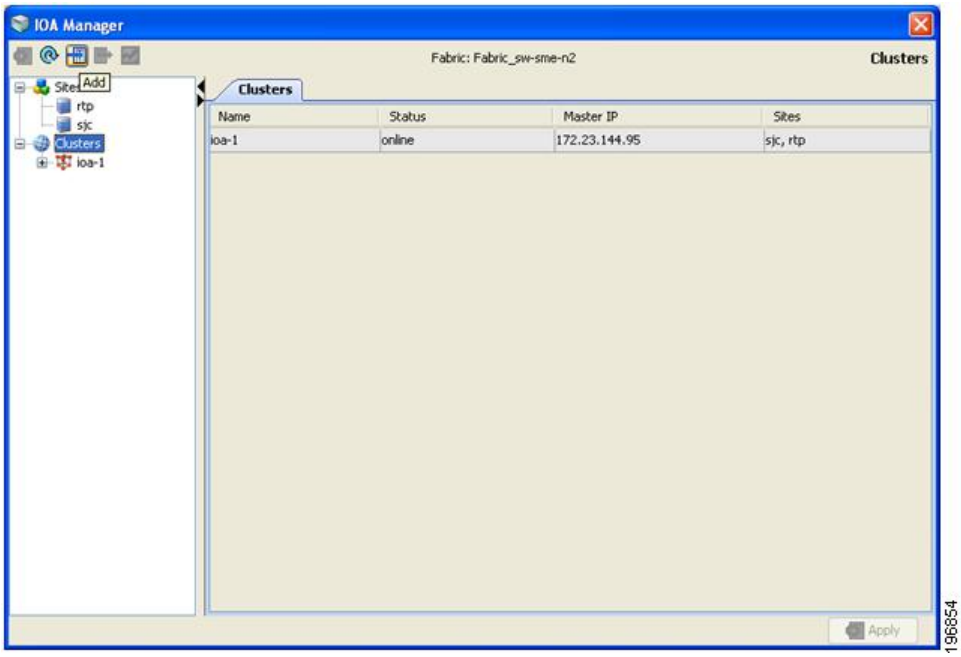
To create a new cluster using IOA Manager, follow these steps:

Procedure

Step 1 Select Cluster in the navigation pane and then click the **Add** icon on the toolbar.

You see the IOA Manager window as shown in Figure 18: IOA Manager - Add Clusters, on page 47.

Figure 18: IOA Manager - Add Clusters



Step 2 Enter the Cluster name, and then click **OK**.

You see the add Cluster name dialog box as shown in Figure 19: Add Cluster Name Dialog box, on page 47.

Figure 19: Add Cluster Name Dialog box

Step 3 Enter the Cluster name and then click OK.

You see the select switch dialog box as shown in Figure 20: Select Switch Dialog Box, on page 47.

Figure 20: Select Switch Dialog Box

Step 4 Select a switch from the drop-down list, and then click **OK**.

Note
You need to select a switch that you would like it to be the master switch as the seed switch when you create the IOA cluster. If you have multiple switches in a site, you may add all the switches in a site that you would like to manage from to the cluster before adding the switches from the remote site.

You see a message box upon successfully creating a cluster as shown in Figure 21: Message Box, on page 47.

Figure 21: Message Box

Step 5 Click **OK**.

Note
If the master switch that you selected is not a member of the site, you may either need to add the switch to an existing site or to create a new site.

Removing a Cluster

To remove a site using IOA Manager, follow these steps:

SUMMARY STEPS

1. In the navigation pane, click the name of the cluster that you want to delete.
2. Click the **Remove** icon on the toolbar.
3. Click **Yes** to remove the cluster.

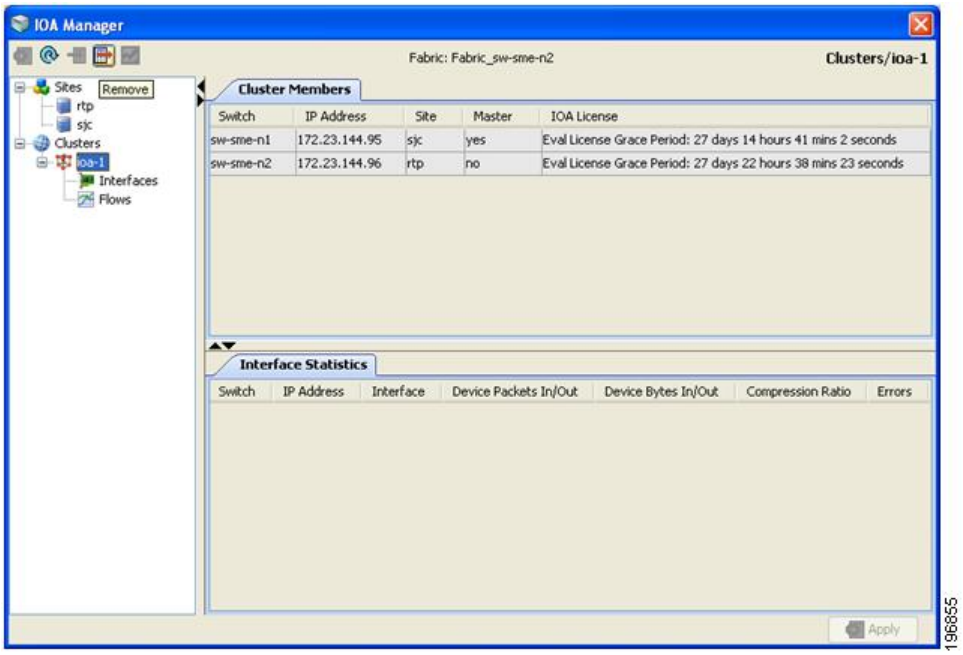
DETAILED STEPS

Procedure

Step 1 In the navigation pane, click the name of the cluster that you want to delete.

You see the IOA Manager window as shown in Figure 22: Removing a Cluster Using IOA Manager, on page 48.

Figure 22: Removing a Cluster Using IOA Manager



Step 2 Click the **Remove** icon on the toolbar.

You see the Delete confirmation dialog box as shown in Figure 23: Remove Confirmation Dialog Box, on page 49.

Figure 23: Remove Confirmation Dialog Box



Step 3 Click **Yes** to remove the cluster.

Viewing Clusters

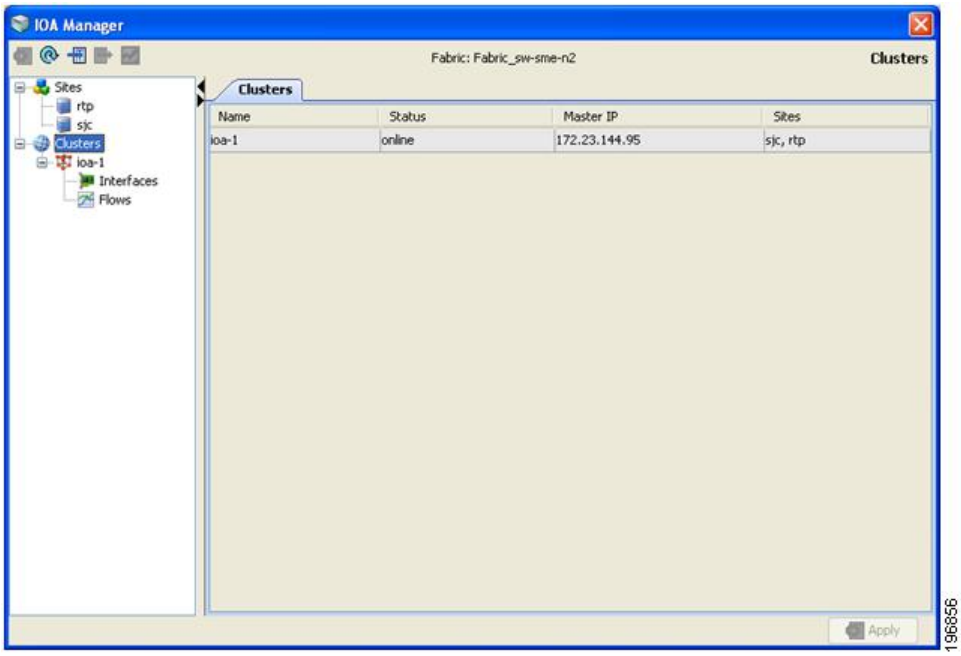
To view a cluster using IOA Manager, follow these steps:

Procedure

Step 1 Click clusters in the navigation pane.

You see the IOA Manager window with clusters selected as shown in Figure 24: Viewing Clusters Using IOA Manager, on page 49.

Figure 24: Viewing Clusters Using IOA Manager

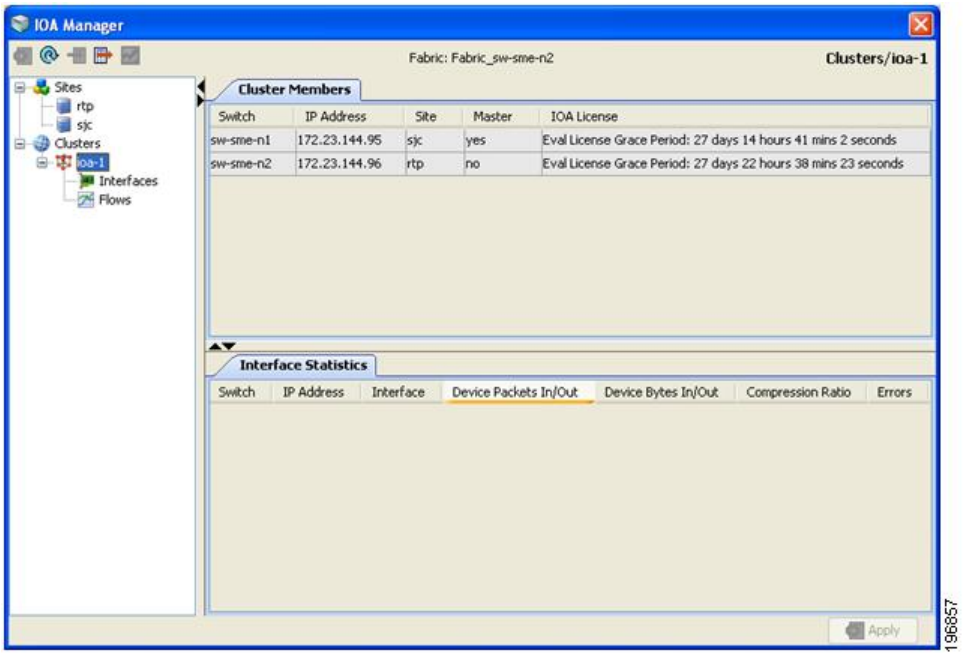


Step 2 Expand the cluster in the hierarchy.

Step 3 Click the name of the cluster to view the details in the information pane.

You see the IOA Manager window with the cluster details as shown in Figure 25: Viewing Cluster Details, on page 50.

Figure 25: Viewing Cluster Details



Note
DCNM does not support configuration of multiple clusters with the same name in a fabric. This action is supported through CLI only.

Configuring Interfaces

You can select the interfaces in a named cluster to see the details in the information pane. The upper table in the information pane displays information about active and configured IOA interface pairs associated with the cluster. The lower table in the information pane displays information about IOA interface candidates that are ready for use in the cluster.

Assigning Interfaces to a Cluster

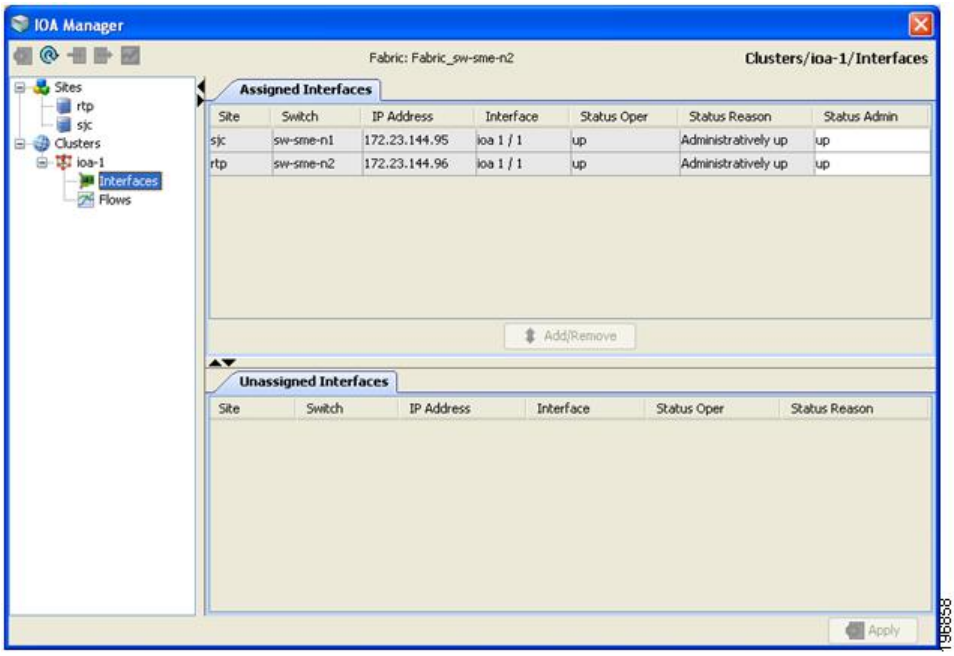
To add a new interface to a cluster using IOA Manager, follow these steps:

Procedure

Step 1 Expand the cluster node in the navigation pane and click **Interfaces**.

You see the IOA Manager window as shown in Figure 26: Adding Interfaces Using IOA Manager, on page 51.

Figure 26: Adding Interfaces Using IOA Manager



The information pane displays the Assigned Interfaces and Unassigned Interfaces tabs.

Step 2 Select one or more interfaces from the Unassigned Interfaces table in the information pane and then click **Add**.

Step 3 Click **Apply** to apply changes.

Note

- You can change the administrative status of an assigned interface by selecting up or down from the admin status drop-down list box and then click **Apply**.
- Cisco DCNM-SAN denotes all the candidate service engines that are not currently provisioned for any service as unconfigured in the unassigned interfaces table. When you select these interfaces, it will automatically provision these service engines for IOA, and configure them as a part of this IOA cluster.

Removing Interfaces from a Cluster

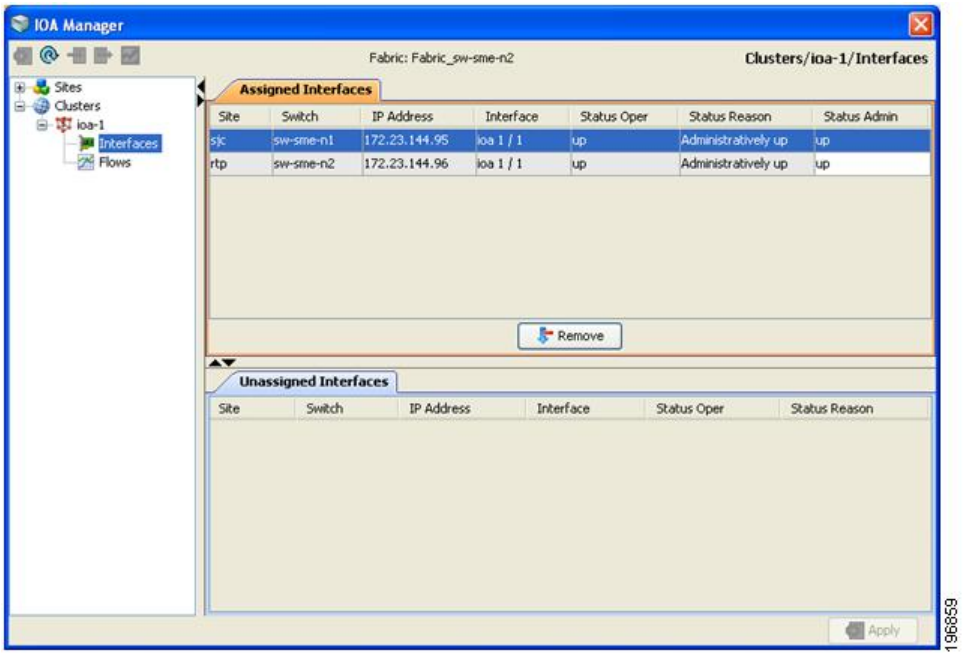
To remove an interface from a cluster, follow these steps:

Procedure

Step 1 Expand the cluster node in the navigation pane and click **Interfaces**.

You see the IOA Manager window as shown in Figure 27: Removing Interfaces Using IOA Manager, on page 52.

Figure 27: Removing Interfaces Using IOA Manager



Step 2 Select the switches from the Assigned Interfaces table that you want to remove.

Step 3 Click **Remove** to move the switches to Unassigned Interfaces table.

Step 4 Click **Apply**.

Configuring Flows

You can select the flows in a named cluster to see the details in the information pane. The upper table in the information pane displays information on active IOA flows. The lower table in the information pane displays information on candidate IOA flows.

Adding a Flow

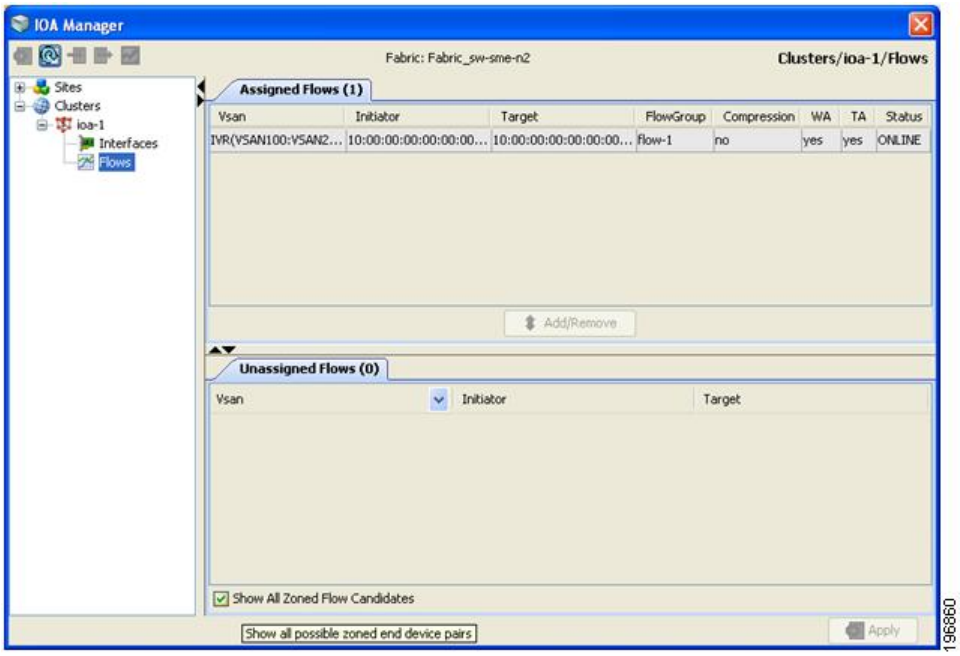
To add a flow in the cluster using IOA Manager, follow these steps:

Procedure

Step 1 Expand the Cluster node in the navigation pane and then click **Flows**.

You see the IOA Manager window displaying the Assigned Flows and Unassigned Flows as shown in Figure 28: Adding Flows Using IOA Manager, on page 53.

Figure 28: Adding Flows Using IOA Manager



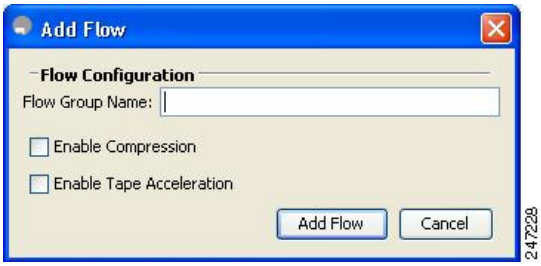
Note
If IVR zoneset is activated, Cisco DCNM-SAN will automatically consider the IVR zoneset and list the candidate IVR flows in the Unassigned flows section.

Step 2 Check the **Click Show All Zoned Flow Candidates** check box to display all the zoned members.

Step 3 Select one or more switches from the Unassigned Flows in the information pane and then click Add.

You see the Add Flows dialog box as shown in Figure 29: Flow Configuration Dialog Box, on page 53.

Figure 29: Flow Configuration Dialog Box



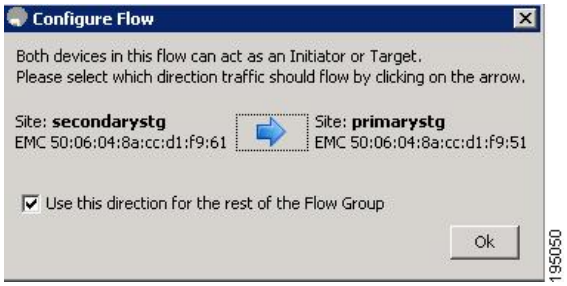
Step 4 Enter a flow group name.

Step 5 Check the **Enable Compression** check box to enable compression.

Step 6 Check the **Enable Tape Acceleration** check box to enable tape acceleration.

Note
Write accleration is enabled by default.

Figure 30: Configure Flow Dialog Box



Step 7 Click the arrow icon to configure the flow in this direction.

Step 8 (Optional) Check the **Use this directon for the rest of the Flow group** check box to apply the same direction to rest of the flow group.

Note
You may use this step only if some of the N ports are registered as both initiators and targets, especially in cases of remote replication flow.

Step 9 Click **Add** and then click **Apply**.

Removing a Flow

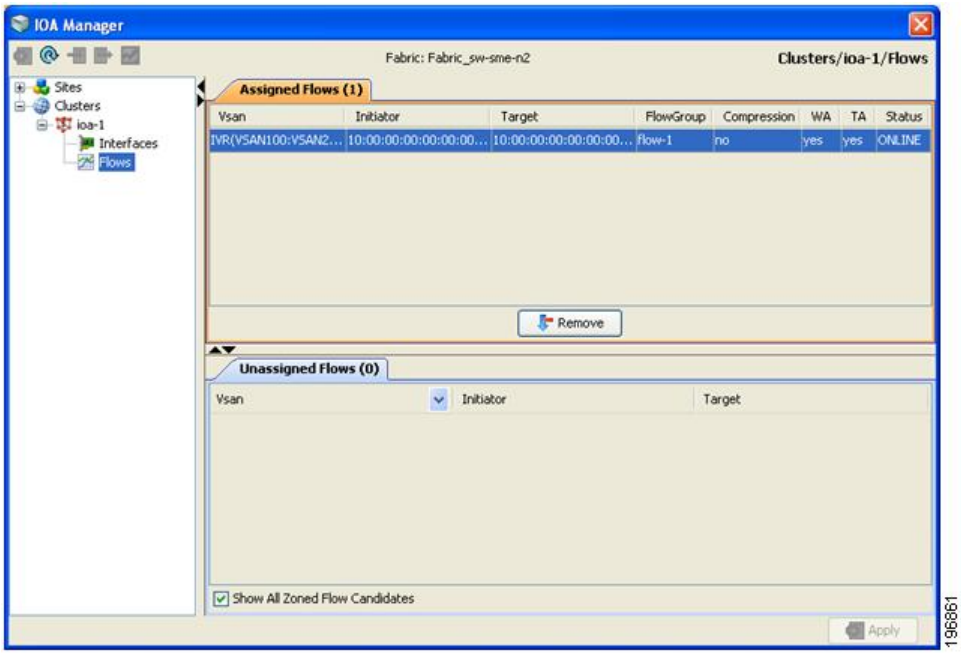
To add a flow in the cluster using IOA Manager, follow these steps:

Procedure

Step 1 Expand the Cluster node in the navigation pane and then click **Flows**.

You see the IOA Manager window displaying the Assigned Flows and Unassigned Flows as shown in Figure 31: Removing Flows Using IOA Manager, on page 55.

Figure 31: Removing Flows Using IOA Manager



Step 2 Select one or more switches from the Assigned Flows in the information pane and then click **Remove**.

Step 3 Click **Apply**.

Viewing Interface Statistics

To view real-time charts using IOA Manager, follow these steps:

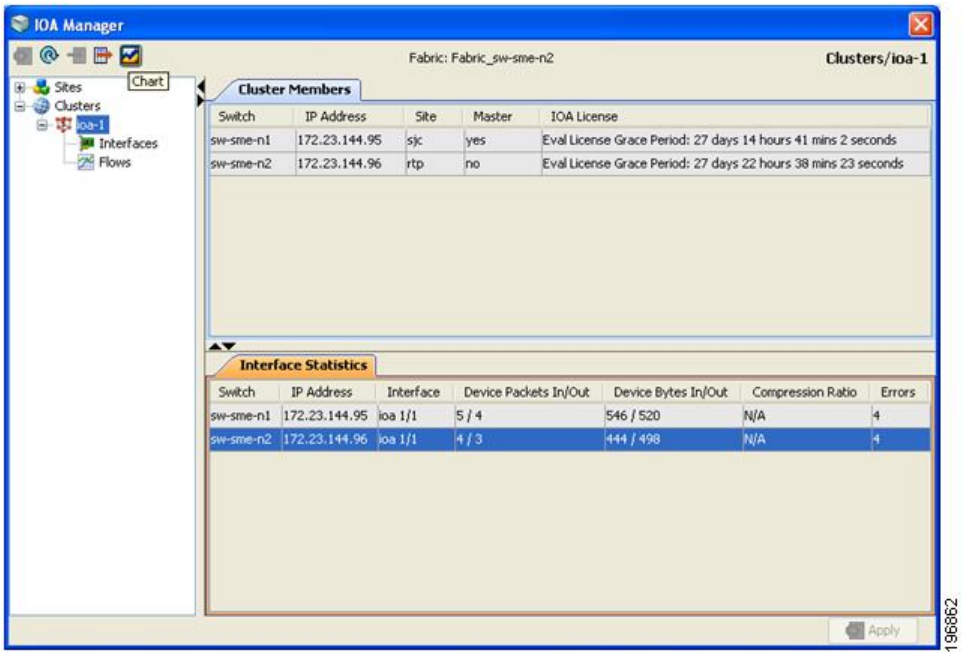
Procedure

Step 1 Expand the Cluster node in the navigation pane and click the name of the cluster.

Step 2 Select a switch from the Interfaces Statistics table in the information pane.

You will see the IOA Manager window as shown in Figure 32: Select IOA Manager Real-time Chart, on page 56.

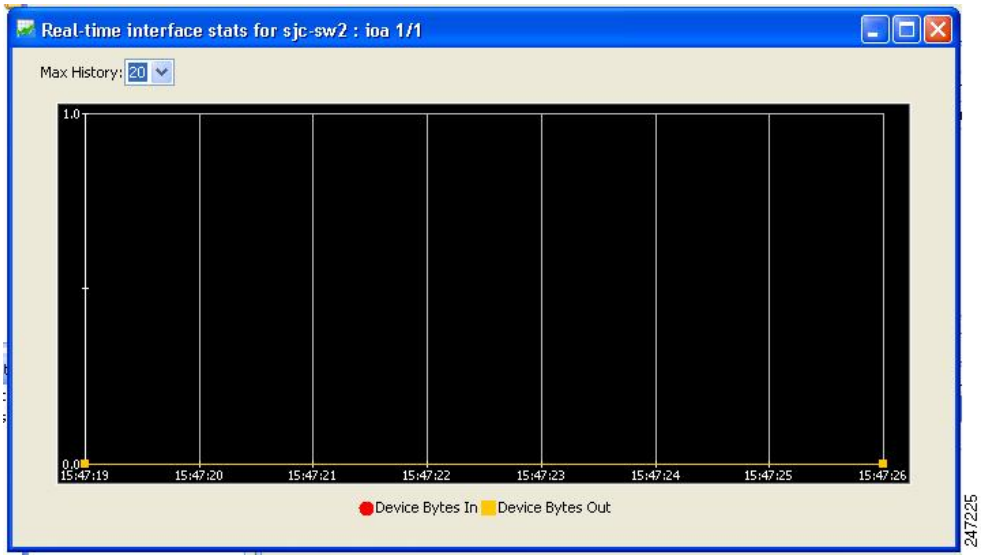
Figure 32: Select IOA Manager Real-time Chart



Step 3 Click the chart icon on the toolbar to monitor real-time charts.

You see the chart as shown in Figure 33: IOA Manager Real-time Chart, on page 56.

Figure 33: IOA Manager Real-time Chart





CHAPTER 6

Configuring IOA Using the CLI

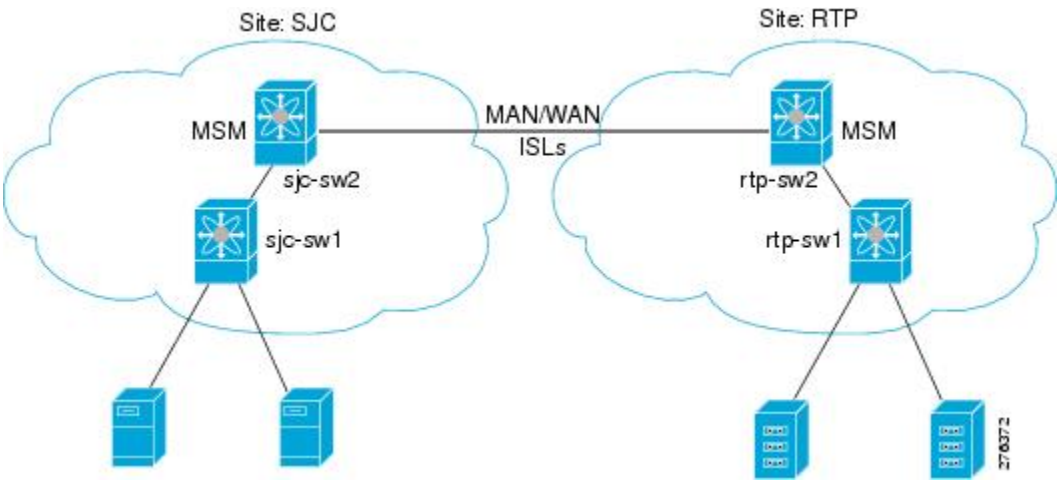
This chapter describes how to configure IOA using the command line interface (CLI).

- Configuring IOA, on page 57
- Configuring an IOA Cluster , on page 61
- Verifying the IOA Flow Connectivity, on page 67
- IOA Flow Setup Wizard, on page 67
- Creating Multiple IOA Clusters on a Single Switch, on page 71
- Configuring IOA with NPV, on page 73
- Additional Configurations for the Features Supported by NPV on IOA, on page 87
- Additional Configurations, on page 103

Configuring IOA

In this chapter, all configuration steps relate to a reference topology shown in Figure 34: IOA CLI Reference Topology, on page 57 where SJC and RTP represent two sites connected through the WAN or MAN ISLs. In this example, sjc-sw2 and rtp-sw2 represent the core switches where IOA is deployed. sjc-sw1 and rtp-sw1 are edge switches that has the hosts or targets connected to them.

Figure 34: IOA CLI Reference Topology



The process of configuring IOA involves a number of configuration tasks that should be completed in order.

On each IOA switch, complete the following configurations:

Enabling Clustering, on page 58
Enabling the IOA Service, on page 58
Classifying the Switch to IOA Site, on page 59
Configuring IOA Interfaces, on page 59

On the master IOA switch, complete the following configurations:

Configuring an IOA Cluster , on page 61
Adding Nodes to an IOA Cluster, on page 62
Adding Interfaces to an IOA Cluster, on page 63
Adding N Ports to an IOA Cluster, on page 65
Configuring the IOA Flows, on page 66

Enabling Clustering

The first step in the process of configuring IOA is to enable clustering in all of the IOA switches.

To enable or disable the IOA cluster on sjc-sw2, perform this task:

Procedure

Step 1 sjc-sw2# **configure terminal**

Enters configuration mode.

Step 2 sjc-sw2(config)# **feature cluster**

Enables clustering.

Step 3 sjc-sw2(config)# **no feature cluster**

Disables clustering.

To complete the configuration for the reference topology, enable clustering in rtp-sw2.

Enabling the IOA Service

After enabling the IOA cluster, the second step in the process of configuring IOA is to enable the IOA service on each of the IOA switches.

To enable the IOA service on sjc-sw2, perform this task:

Procedure

Step 1 sjc-sw2# **configure terminal**

Enters configuration mode.

Step 2 sjc-sw2(config)# **feature ioa**

Enables IOA feature.

Step 3 sjc-sw2(config)# **no feature ioa**

Disables IOA feature.

To complete the configuration for the reference topology, enable the IOA service in rtp-sw2.

Classifying the Switch to IOA Site

Each of the IOA switches need to be classified into a site. Make sure that you classify only the IOA switches within the physical site into an IOA site.

To classify an IOA switch into the SJC site, perform this task:

Procedure

Step 1 sjc-sw2# **configure terminal**

Enters configuration mode.

Step 2 sjc-sw2(config)# **ioa site-local SJC**

Configures the site to which the switch belongs. The maximum name length is restricted to 31 alphabetical characters.

Note
This command configures the site to which the switch belongs across all the IOA clusters that the switch participates in.

To complete the configuration for the reference topology, classify rtp-sw2 into the RTP site.

Configuring IOA Interfaces

After enabling the cluster and enabling IOA, configure the IOA interfaces on the switch.

Procedure

Step 1 sjc-sw2# **configure terminal**

Enters configuration mode.

Step 2 sjc-sw2(config)# **interface ioa 2/1**

Configures IOA on service engine 1 in slot 2.

Step 3 sjc-sw2(config)# **interface ioa 2/1**

Configures IOA on service engine 2 in slot 2.

Note
Service engines 2, 3, and 4 are available only on the SSN-16 module. The appropriate IOA license is checked out as a part of the creation of the interface.

A standard MDS notation is used to denote the IOA interfaces: ioa *slot* /*service engine* . For example, ioa2/1 refers to Slot 2, Service Engine 1. In the case of the MSM-18/4 Module and 9222i Switch, only one service engine exists and so only ioa2/1 is valid. In the case of the SSN-16 Module, four service engines exist and so ioa2/1, ioa2/2, ioa2/3, and ioa2/4 are valid interfaces.

Step 4 sjc-sw2(config)# **no interface ioa 2/2**

Deletes the IOA interface.

Note
Before deleting an IOA interface, you must remove the IOA interface from the cluster.

Step 5 sjc-sw2(config-if)# **no shutdown**

Enables the IOA interface.

Step 6 sjc-sw2(config-if)# **shutdown**

Disables the IOA interface.

Note
FCIP and IOA are not supported on the same engine.

To complete the configuration for the reference topology, configure the interfaces in rtp-sw2.

Displaying IOA Interface Status

After configuring the IOA interface, use the **show int** command to show whether the IOA interface is down. The interface is down until the interface is added to a cluster.

```
sjc-sw2# show interface ioa 2/1
ioa2/1 is down (Not in any Cluster)
    0 device packets in, 0 device packets out
    0 device bytes in, 0 device bytes out
    0 peer packets in, 0 peer packets out
    0 peer bytes in, 0 peer bytes out
    0 i-t create request, 0 i-t create destroy
    0 i-t activate request, 0 i-t deactivate request
```

Possible reasons for the interface being down are as follows:

- Administratively down—The interface is shut down.
- Not in any cluster—The interface is not part of any IOA cluster.

- Port software failure—A software failure has occured causing a reset of the IOA service engine.
- No license—The interface does not have a valid IOA license. The license is either not installed or all the available licenses are in use.

Configuring an IOA Cluster

Procedure

Step 1 sjc-sw2# **configure terminal**

Enters configuration mode.

Step 2 sjc-sw2(config)# **ioa cluster tape_vault**

sjc-sw2(config-ioa-cl)#

Assigns a user-specified name (tape_vault) to the IOA cluster. The maximum length of the name is 31 alphabetical characters. Enters the cluster configuration submode. The local switch is implicitly added to the cluster as part of this command.

Step 3 sjc-sw2(config)# **no ioa cluster tape_vault**

Deletes the specified IOA cluster.

Note
You need to select a switch that you want to be the master switch as the seed switch when you create the IOA cluster. If you have multiple switches in a site, you may add all the switches in a site that you want to manage the cluster before adding the switches from the remote site.

This section includes the following topics:

Displaying IOA Cluster Status

The following examples display the cluster information:



Note You must configure at least one IOA interface on each site for the cluster to be online.

```
sjc-sw2# show ioa cluster

IOA Cluster is tape_vault
  Cluster ID is 0x213a000dec3ee782
  Cluster status is online
  Is between sites SJC and RTP
  Total Nodes are 2
  Cluster Infra Status : Operational
  Cluster is Administratively Up
  Cluster Config Version : 26
  SSL for ICN : Not Configured
```

```
sjc-sw2# show ioa cluster tape_vault
IOA Cluster is tape_vault
  Cluster ID is 0x213a000dec3ee782
  Cluster status is online
  Is between sites SJC and RTP
  Total Nodes are 2
  Cluster Infra Status : Operational
  Cluster is Administratively Up
  Cluster Config Version : 26
  SSL for ICN : Not Configured
```

A cluster can have the following statuses:

- Pending—An IOA interface needs to be added to the cluster.
- Online—The cluster is online. IOA services can be run on the cluster.
- Offline—The cluster is offline. Check the infrastructure status for more information.

The infrastructure status has the following values:

- Operational—The cluster infrastructure is operational on this switch. The IOA service will be able to use the cluster on this switch.
- Not Operational—The cluster infrastructure is not operational on this node. The IOA service will not run on this cluster on this switch.

The administrative status has the following values:

- Administratively Up—If the cluster is not online, check this status to make sure that the cluster is administratively up.
- Administratively Shutdown—The cluster was shut down.

Adding Nodes to an IOA Cluster

Procedure

Step 1 sjc-sw2# **configure terminal**

Enters configuration mode.

Step 2 sjc-sw2(config)# **ioa cluster tape_vault**

sjc-sw2(config-ioa-cl)#

Enters the cluster configuration submode and adds the local switch where this command is executed into the IOA cluster.

Step 3 sjc-sw2(config-ioa-cl)# **node local**

Enters the node configuration submode for the local switch. The **local** keyword denotes the switch where the CLI command is executed.

Note
You may also specify the node name of the local switch to enter submode. The node name could be either the IP address or the DNS name of the local switch.

Step 4 sjc-sw2(config-ioa-cl)# **node sjc-sw2**

sjc-sw2(config-ioa-cl-node)# **end**

Includes the switch as part of the cluster. Enters the node configuration submode.

Step 5 sjc-sw2(config-ioa-cl)# **node rtp-sw2**

sjc-sw2(config-ioa-cl-node)# **end**

Includes the remote switch as part of the cluster. Alternatively, use an IPv4 or IPv6 address. Enters the node configuration submode.

Step 6 sjc-sw2(config-ioa-cl)# **no node rtp-sw2**

Removes the local or the remote node from the cluster.

The following examples display the nodes information:

```
sjc-sw2# show ioa cluster summary

------------------------------------------------------------------------------
Cluster          Sites              Status    Master Switch
------------------------------------------------------------------------------
tape_vault       SJC,               online    172.23.144.97
                 RTP
sjc-sw2# show ioa cluster tape_vault node summary
------------------------------------------------------------------------------
Switch             Site            Status           Master
------------------------------------------------------------------------------
172.23.144.97(L)   SJC             online           yes
172.23.144.98      RTP             online           no
sjc-sw2# show ioa cluster tape_vault node
Node 172.23.144.97 is local switch
  Node ID is 1
  Status is online
  Belongs to Site SJC
  Node is the master switch
Node 172.23.144.98 is remote switch
  Node ID is 2
  Status is online
  Belongs to Site RTP
  Node is not master switch
```

Adding Interfaces to an IOA Cluster

To add IOA interfaces to an IOA cluster, perform this task:

Procedure

Step 1 sjc-sw2# **configure terminal**

Enters configuration mode.

Step 2 sjc-sw2(config)# **ioa cluster tape_vault**

sjc-sw2(config-ioa-cl)#

Enters the cluster configuration submode.

Step 3 sjc-sw2(config-ioa-cl)# **node local**

Includes the local switch as part of the cluster. Enters the node configuration submode for the local switch. The **local** keyword denotes the switch where the CLI command is executed.

Note
You may also specify the node name of the local switch to enter submode. The node name could be either the IP address or the DNS name of the local switch.

Step 4 sjc-sw2(config-ioa-cl-node)# **interface ioa 2/1**

sjc-sw2(config-ioa-cl-node)# **interface ioa 2/2**

Adds the interfaces to the IOA cluster.

Step 5 sjc-sw2(config-ioa-cl-node)# **no interface ioa 2/2**

(Optional) Removes the interface from the IOA cluster.

Step 6 sjc-sw2(config-ioa-cl)# **node rtp-sw2**

Includes the remote switch as part of the cluster. Alternatively, use a IPv4 or IPv6 address. Enters the node configuration submode.

Step 7 sjc-sw2(config-ioa-cl-node)# **interface ioa 2/1**

sjc-sw2(config-ioa-cl-node)# **interface ioa 2/2**

Adds the interfaces to the IOA cluster.

Step 8 sjc-sw2(config-ioa-cl-node)# **no interface ioa 2/2**

(Optional) Removes the interface from the IOA cluster.

The following examples display IOA interfaces information:

```
The following examples display IOA interfaces information:
sjc-sw2# show interface ioa2/1
ioa2/1 is up
    Member of cluster tape_vault
    0 device packets in, 0 device packets out
    0 device bytes in, 0 device bytes out
    0 peer packets in, 0 peer packets out
    0 peer bytes in, 0 peer bytes out
    303 i-t create request, 300 i-t create destroy
    300 i-t activate request, 0 i-t deactivate request
sjc-sw2# show ioa cluster tape_vault interface summary

--------------------------------------------------------------------------------
Switch             Interface         Status       Flows
--------------------------------------------------------------------------------
172.23.144.97(L)   ioa2/1            up           --
172.23.144.97(L)   ioa2/2            up           --
172.23.144.98      ioa2/1            up           --
172.23.144.98      ioa2/2            up           --
sjc-sw2# show ioa cluster tape_vault interface

Interface ioa2/1 belongs to 172.23.144.97(L)(M)
  Status is up
Interface ioa2/2 belongs to 172.23.144.97(L)(M)
  Status is up
Interface ioa2/1 belongs to 172.23.144.98
  Status is up
```

```
Interface ioa2/2 belongs to 172.23.144.98
  Status is up
```

Note
(L) indicates the Local switch.(M) indicates the Master switch.

Adding N Ports to an IOA Cluster

Procedure

Step 1 sjc-sw2# **configure terminal**

Enters configuration mode.

Step 2 switch(config)# **ioa cluster tape_vault**

Enters the cluster configuration submode.

Step 3 sjc-sw2(config-ioa-cl)# **nport pwwn 10:0:0:0:0:0:0:1 site SJC vsan 100**

sjc-sw2(config-ioa-cl)# **nport pwwn 11:0:0:0:0:0:0:1 site RTP vsan 100**

sjc-sw2(config-ioa-cl)# **nport pwwn 10:0:0:0:0:0:0:2 site SJC vsan 100**

sjc-sw2(config-ioa-cl)# **nport pwwn 11:0:0:0:0:0:0:2 site RTP vsan 100**

sjc-sw2(config-ioa-cl)# **end**

Configures the site and VSAN ID of the N ports that will be a part of accelerated flows.

Step 4 sjc-sw2(config-ioa-cl)# **no nport pwwn 10:0:0:0:0:0:0:1**

Removes the N port from the IOA cluster.

This example shows how to display N ports configuration:

```
sjc-sw2# show ioa cluster tape_vault nports
--------------------------------------------------------------------------------
P-WWN                          Site                           Vsan
--------------------------------------------------------------------------------
10:00:00:00:00:00:00:01        SJC                            100
11:00:00:00:00:00:00:01        RTP                            100
10:00:00:00:00:00:00:02        SJC                            100
11:00:00:00:00:00:00:02        RTP                            100
```

Configuring the IOA Flows

Procedure

Step 1 sjc-sw2# **configure terminal**

Enters configuration mode.

Step 2 switch(config)# **ioa cluster tape_vault**

Enters the cluster configuration submode.

Step 3 switch(config-ioa-cl)# **flowgroup tsm**

Creates an IOA flow group.

Step 4 switch(config-ioa-cl)# **no flowgroup tsm**

Deletes an IOA flow group.

Step 5 sjc-sw2(config-ioa-cl-flgrp)# **host 10:0:0:0:0:0:0:1 target 11:0:0:0:0:0:0:1**

Creates a flow with write acceleration.

Step 6 sjc-sw2(config-ioa-cl-flgrp)# **host 10:0:0:0:0:0:0:2 target 11:0:0:0:0:0:0:2 tape**

Creates a flow with tape acceleration.

Step 7 sjc-sw2(config-ioa-cl-flgrp)# **host 10:0:0:0:0:0:0:3 target 11:0:0:0:0:0:0:3 compression**

Creates a flow with write acceleration and compression.

Step 8 sjc-sw2(config-ioa-cl-flgrp)# host 10:0:0:0:0:0:0:4 target 11:0:0:0:0:0:0:4 tape compression

Creates a flow with tape acceleration, and compression.

Step 9 sjc-sw2(config-ioa-cl-flgrp)# **no host 10:0:0:0:0:0:0:1 target 11:0:0:0:0:0:0:1**

Removes the configured flow.

Note
We recommend that you suspend the traffic while enabling IOA for a given flow.

The following examples display the configured flow information without device alias:

```
sjc-sw2# show ioa cluster tape_vault flows
--------------------------------------------------------------------------------
Host WWN,                VSAN     WA  TA  Comp  Status    Switch,Interface
Target WWN                                                Pair
--------------------------------------------------------------------------------
10:00:00:00:00:00:00:01, 100      Y    Y    N   online    172.23.144.97, ioa2/1
11:00:00:00:00:00:00:01, 100                              172.23.144.98, ioa2/1
10:00:00:00:00:00:00:02, 100      Y    Y    Y   online    172.23.144.97, ioa2/2
11:00:00:00:00:00:00:02, 100                              172.23.144.98, ioa2/2
```

The following examples display the configured flow information with device alias:

```
sjc-sw2# show ioa cluster tape_vault flows
```

```
--------------------------------------------------------------------------------
Host WWN,                VSAN     WA  TA  Comp  Status    Switch,Interface
Target WWN                                               Pair
--------------------------------------------------------------------------------
host-1                  , 100      Y    Y    N   online    172.23.144.97, ioa2/1
target-1                , 100                              172.23.144.98, ioa2/1
host-2                  , 100      Y    Y    Y   online    172.23.144.97, ioa2/2
target-2                , 100                              172.23.144.98, ioa2/2
sjc-sw2# show ioa cluster tape_vault flows detail

Host 10:00:00:00:00:00:00:01, Target 11:00:00:00:00:00:00:01, VSAN 100
  Is online
  Belongs to flowgroup tsm
  Is enabled for WA, TA
  Is assigned to
    Switch 172.23.144.97    Interface ioa2/1 (Host Site)
    Switch 172.23.144.98    Interface ioa2/1 (Target Site)
Host 10:00:00:00:00:00:00:02, Target 11:00:00:00:00:00:00:02, VSAN 100
  Is online
  Belongs to flowgroup tsm
  Is enabled for WA, TA, Compression
  Is assigned to
    Switch 172.23.144.97    Interface ioa2/2 (Host Site)
    Switch 172.23.144.98    Interface ioa2/2 (Target Site)
```

Verifying the IOA Flow Connectivity

To validate the connectivity between the master switch and the specified target device (for a specific flow), perform this task:

Procedure

switch# **ioa-ping host hpwwn target tpwwn vsan vid interface if0**

switch# **ioa-ping host 10:00:00:00:11:a1:01:0a target 50:0a:09:80:11:4b:01:0a vsan 11 interface ioa 1/1**

Validates the connectivity between the master switch and the specified target device. The host and target pWWNs should be part of the same IOA flow.

Note: The **ioa-ping** command functions from 6.2(5) Release and later releases and the command has to be executed from the exec mode of IOA master switch only.

IOA Flow Setup Wizard

You can use the IOA Flow Setup Wizard to simplify the provisioning of flows especially when there are many flows to provision, and when you add, remove, or replace host HBAs, tape drives or storage controllers.

This section includes the following topics:

Using the IOA Flow Setup Wizard

To configure flows using the Flow Setup Wizard, follow these steps:

Before you begin

The following prerequisites must be met before you can invoke the IOA Flow Setup Wizard:

- All of the N ports of both initiators and targets that need to be accelerated must be online.
- The zoning configuration must already be in place to permit the flows that need to communicate with each other. If you are replacing a host HBA, you must update the zoning configuration to remove the faulty HBA and to add the new HBA before you invoke the IOA Flow Setup Wizard.

Procedure

Step 1 sjc-sw1# **ioa flow-setup cluster tape_vault flowgroup repln-fg vsan 100**

Invoke the Flow Setup Wizard on a specific VSAN. In the case of an IVR deployment, you can enter the following CLI command on an IVR border switch where IOA is deployed:

a) sjc-sw1# **ioa ivr flow-setup cluster tape_vault flowgroup repln-fg**

The wizard processes the active zone set for the VSAN and creates a set of candidate flows. When you use the **ivr flow-setup** command, the active IVR zone set is considered. The zone set may have local flows as well as flows that traverse across sites. The IOA Flow Setup Wizard runs through a series of steps as listed in this procedure to prune the list to capture only the flows that traverse across the sites that need to be accelerated.

Step 2 Classify the switches in the candidate switch list into appropriate sites.

This step is only for those switches where none of the hosts or targets have been configured yet for acceleration. From the flows in the active zone set, a candidate switch list is prepared based on where the hosts and targets are logged into.

The following switches need to be classified into appropriate sites:

Example:

```
--------------------------------------------------------------------------------
Do you want to classify sjc-sw1 into site sjc or rtp [sjc]
Do you want to classify 172.23.144.96 into site sjc or rtp [sjc] rtp
```

The candidate flow list is now pruned to contain only the inter-site flows that need to be accelerated.

Step 3 The wizard displays all of the N ports that need to be classified into sites. Enter **yes** to classify the N ports into sites.

The following nport to site mapping needs to be configured:

Example:

```
--------------------------------------------------------------------------------
N-Port PWWN: 10:00:00:00:00:00:00:00 Site: sjc
N-Port PWWN: 10:00:00:00:00:00:01:00 Site: sjc
N-Port PWWN: 10:00:00:00:00:00:02:00 Site: sjc
N-Port PWWN: 10:00:00:00:00:00:03:00 Site: sjc
N-Port PWWN: 10:00:00:00:00:00:04:00 Site: sjc
N-Port PWWN: 11:00:00:00:00:00:00:00 Site: rtp
N-Port PWWN: 11:00:00:00:00:00:01:00 Site: rtp
N-Port PWWN: 11:00:00:00:00:00:02:00 Site: rtp
```

```
N-Port PWWN: 11:00:00:00:00:00:03:00 Site: rtp
N-Port PWWN: 11:00:00:00:00:00:04:00 Site: rtp
Do you want to configure the n-port to site mappings? (yes/no) [yes] yes
```

Step 4 (Optional) Use this step only when some of the N ports such as those used in remote replication are represented as **scsi-fcp(both)** in the FCNS database. Enter the primary direction of the traffic that will be used by IOA to decide on what should be configured as host and target in IOA.

Replication traffic can flow in either direction.

Example:

```
Certain N-ports in this VSAN can act as both initiator and targets
```

Is the traffic flow primarily from sjc to rtp? (yes/no) [yes] **yes**

Step 5 The wizard configures the list of flows that are not already configured in IOA and attempts to delete the IOA flows that are not part of the zone set. This operation specifically handles removing HBAs or storage controllers. Enter **yes** to accept the flows that need to be accelerated. New flows that need to be accelerated are displayed.

The following flows will be configured:

Example:

```
------------------------------------------------------------------------------
Host: 10:00:00:00:00:00:00:00 VSAN: 100 Target: 11:00:00:00:00:00:00:00 VSAN:100
Host: 10:00:00:00:00:00:00:00 VSAN: 100 Target: 11:00:00:00:00:00:01:00 VSAN:100
Host: 10:00:00:00:00:00:00:00 VSAN: 100 Target: 11:00:00:00:00:00:02:00 VSAN:100
Host: 10:00:00:00:00:00:00:00 VSAN: 100 Target: 11:00:00:00:00:00:03:00 VSAN:100
Host: 10:00:00:00:00:00:01:00 VSAN: 100 Target: 11:00:00:00:00:00:00:00 VSAN:100
Host: 10:00:00:00:00:00:01:00 VSAN: 100 Target: 11:00:00:00:00:00:01:00 VSAN:100
Host: 10:00:00:00:00:00:01:00 VSAN: 100 Target: 11:00:00:00:00:00:02:00 VSAN:100
Host: 10:00:00:00:00:00:01:00 VSAN: 100 Target: 11:00:00:00:00:00:03:00 VSAN:100
Host: 10:00:00:00:00:00:02:00 VSAN: 100 Target: 11:00:00:00:00:00:00:00 VSAN:100
Host: 10:00:00:00:00:00:02:00 VSAN: 100 Target: 11:00:00:00:00:00:01:00 VSAN:100
Host: 10:00:00:00:00:00:02:00 VSAN: 100 Target: 11:00:00:00:00:00:02:00 VSAN:100
Host: 10:00:00:00:00:00:02:00 VSAN: 100 Target: 11:00:00:00:00:00:03:00 VSAN:100
Host: 10:00:00:00:00:00:03:00 VSAN: 100 Target: 11:00:00:00:00:00:00:00 VSAN:100
Host: 10:00:00:00:00:00:03:00 VSAN: 100 Target: 11:00:00:00:00:00:01:00 VSAN:100
Host: 10:00:00:00:00:00:03:00 VSAN: 100 Target: 11:00:00:00:00:00:02:00 VSAN:100
Host: 10:00:00:00:00:00:03:00 VSAN: 100 Target: 11:00:00:00:00:00:03:00 VSAN:100
Host: 10:00:00:00:00:00:04:00 VSAN: 100 Target: 11:00:00:00:00:00:04:00 VSAN:100
Do you want to configure these flows? (yes/no) [yes] yes
```

Step 6 You can display the configured flow information by using the following commands:

Example:

```
sjc-sw1# show ioa cluster tape_vault nports

------------------------------------------------------------------------------
P-WWN                          Site                           Vsan
------------------------------------------------------------------------------
10:00:00:00:00:00:00:00        sjc                            100
10:00:00:00:00:00:01:00        sjc                            100
10:00:00:00:00:00:02:00        sjc                            100
10:00:00:00:00:00:03:00        sjc                            100
10:00:00:00:00:00:04:00        sjc                            100
11:00:00:00:00:00:00:00        rtp                            100
11:00:00:00:00:00:01:00        rtp                            100
11:00:00:00:00:00:02:00        rtp                            100
11:00:00:00:00:00:03:00        rtp                            100
```

```
11:00:00:00:00:00:04:00        rtp                          100
sjc-sw1# show ioa cluster tape_vault flows

--------------------------------------------------------------------------------
Host WWN,                VSAN     WA  TA  Comp  Status    Switch,Interface
Target WWN                                                Pair
--------------------------------------------------------------------------------

10:00:00:00:00:00:00:00, 100      Y    N    N   offline   --, --
11:00:00:00:00:00:00:00                                   --, --
10:00:00:00:00:00:01:00, 100      Y    N    N   offline   --, --
11:00:00:00:00:00:00:00                                   --, --
10:00:00:00:00:00:02:00, 100      Y    N    N   offline   --, --
11:00:00:00:00:00:00:00                                   --, --
10:00:00:00:00:00:03:00, 100      Y    N    N   offline   --, --
11:00:00:00:00:00:00:00                                   --, --
10:00:00:00:00:00:00:00, 100      Y    N    N   offline   --, --
11:00:00:00:00:00:01:00                                   --, --
10:00:00:00:00:00:01:00, 100      Y    N    N   offline   --, --
11:00:00:00:00:00:01:00                                   --, --
10:00:00:00:00:00:02:00, 100      Y    N    N   offline   --, --
11:00:00:00:00:00:01:00                                   --, --
10:00:00:00:00:00:03:00, 100      Y    N    N   offline   --, --
11:00:00:00:00:00:01:00                                   --, --
10:00:00:00:00:00:00:00, 100      Y    N    N   offline   --, --
11:00:00:00:00:00:02:00                                   --, --
10:00:00:00:00:00:01:00, 100      Y    N    N   offline   --, --
11:00:00:00:00:00:02:00                                   --, --
10:00:00:00:00:00:02:00, 100      Y    N    N   offline   --, --
11:00:00:00:00:00:02:00                                   --, --
10:00:00:00:00:00:03:00, 100      Y    N    N   offline   --, --
11:00:00:00:00:00:02:00                                   --, --
10:00:00:00:00:00:00:00, 100      Y    N    N   offline   --, --
11:00:00:00:00:00:03:00                                   --, --
10:00:00:00:00:00:01:00, 100      Y    N    N   offline   --, --
11:00:00:00:00:00:03:00                                   --, --
10:00:00:00:00:00:02:00, 100      Y    N    N   offline   --, --
11:00:00:00:00:00:03:00                                   --, --
10:00:00:00:00:00:03:00, 100      Y    N    N   offline   --, --
11:00:00:00:00:00:03:00                                   --, --
10:00:00:00:00:00:04:00, 100      Y    N    N   offline   --, --
11:00:00:00:00:00:04:00                                   --, --
```

Step 7 If data is currently being transmitted through the flow, it is considered to be online and active. A throughput number in megabytes per second is shown for each flow that is online and active. Use the following commands to display all flows assigned to a single interface, or to display all flows assigned to all interfaces:

Example:

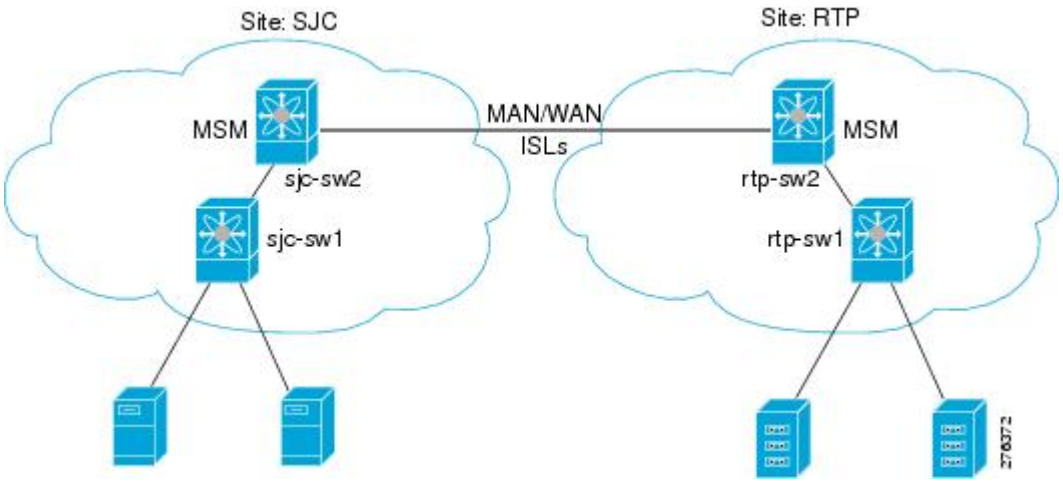
```
switch# show ioa online flows interface ioa2/1
--------------------------------------------------------------------------------
                                                                    A O
                                                                    c n
                                                                    t l
                                                                    i i
                                                                    v n
 FLOW ID     FLOW HOST               FLOW TARGET             VSAN   e e    MBps
--------------------------------------------------------------------------------
       0     10:00:00:00:00:00:00:10 11:00:00:00:00:00:00:10    1   N Y   0.00
      17     42:00:00:00:00:00:00:11 41:00:00:00:00:00:00:11    1   N Y   0.00
      18     42:00:00:00:00:00:00:12 41:00:00:00:00:00:00:12    1   N Y   0.00
--More--
switch# show ioa online flows interface all
```

```
--------------------------------------------------------------------------------
                                                                    A O
                                                                    c n
                                                                    t l
                                                                    i i
                                                                    v n
 FLOW ID     FLOW HOST               FLOW TARGET             VSAN   e e    MBps
--------------------------------------------------------------------------------
       0     10:00:00:00:00:00:00:10 11:00:00:00:00:00:00:10    1   N Y   0.00
      17     42:00:00:00:00:00:00:11 41:00:00:00:00:00:00:11    1   N Y   0.00
      18     42:00:00:00:00:00:00:12 41:00:00:00:00:00:00:12    1   N Y   0.00
      19     42:00:00:00:00:00:00:13 41:00:00:00:00:00:00:13    1   N Y   0.00
      20     42:00:00:00:00:00:00:14 41:00:00:00:00:00:00:14    1   N Y   0.00
      21     42:00:00:00:00:00:00:15 41:00:00:00:00:00:00:15    1   N Y   0.00
      22     42:00:00:00:00:00:00:16 41:00:00:00:00:00:00:16    1   N Y   0.00
      23     42:00:00:00:00:00:00:17 41:00:00:00:00:00:00:17    1   N Y   0.00
      24     42:00:00:00:00:00:00:18 41:00:00:00:00:00:00:18    1   N Y   0.00
      25     42:00:00:00:00:00:00:19 41:00:00:00:00:00:00:19    1   N Y   0.00
      26     42:00:00:00:00:00:00:1a 41:00:00:00:00:00:00:1a    1   N Y   0.00
      27     42:00:00:00:00:00:00:1b 41:00:00:00:00:00:00:1b    1   N Y   0.00
      28     42:00:00:00:00:00:00:1c 41:00:00:00:00:00:00:1c    1   N Y   0.00
      29     42:00:00:00:00:00:00:1d 41:00:00:00:00:00:00:1d    1   N Y   0.00
      30     42:00:00:00:00:00:00:1e 41:00:00:00:00:00:00:1e    1   N Y   0.00
      31     42:00:00:00:00:00:00:1f 41:00:00:00:00:00:00:1f    1   N Y   0.00
      32     42:00:00:00:00:00:00:20 41:00:00:00:00:00:00:20    1   N Y   0.00
      33     42:00:00:00:00:00:00:21 41:00:00:00:00:00:00:21    1   N Y   0.00
      34     42:00:00:00:00:00:00:22 41:00:00:00:00:00:00:22    1   N Y   0.00
      35     42:00:00:00:00:00:00:23 41:00:00:00:00:00:00:23    1   N Y   0.00
      36     42:00:00:00:00:00:00:24 41:00:00:00:00:00:00:24    1   N Y   0.00
      37     42:00:00:00:00:00:00:25 41:00:00:00:00:00:00:25    1   N Y   0.00
      38     42:00:00:00:00:00:00:26 41:00:00:00:00:00:00:26    1   N Y   0.00
      39     42:00:00:00:00:00:00:27 41:00:00:00:00:00:00:27    1   N Y   0.00
      40     42:00:00:00:00:00:00:28 41:00:00:00:00:00:00:28    1   N Y   0.00
      41     42:00:00:00:00:00:00:29 41:00:00:00:00:00:00:29    1   N Y   0.00
      42     42:00:00:00:00:00:00:2a 41:00:00:00:00:00:00:2a    1   N Y   0.00
      43     42:00:00:00:00:00:00:2b 41:00:00:00:00:00:00:2b    1   N Y   0.00
      44     42:00:00:00:00:00:00:2c 41:00:00:00:00:00:00:2c    1   N Y   0.00
      45     42:00:00:00:00:00:00:2d 41:00:00:00:00:00:00:2d    1   N Y   0.00
      46     42:00:00:00:00:00:00:2e 41:00:00:00:00:00:00:2e    1   N Y   0.00
      47     42:00:00:00:00:00:00:2f 41:00:00:00:00:00:00:2f    1   N Y   0.00
      48     42:00:00:00:00:00:00:30 41:00:00:00:00:00:00:30    1   N Y   0.00
      49     42:00:00:00:00:00:00:31 41:00:00:00:00:00:00:31    1   N Y   0.00
switch#
```

Creating Multiple IOA Clusters on a Single Switch

Figure 35: Extended Across Multiple Sites, on page 72 illustrates the IOA implementation where the IOA service is extended across multiple sites. In the illustration, Site-SJC consolidates the tape backup from Site-RTP and Site-SAC. Each IOA cluster represents a site pair, which means there are two unique clusters. This topology provides segregation and scalability of the IOA service across multiple sites. In the Site-SJC, a single switch can participate in multiple IOA clusters.

Figure 35: Extended Across Multiple Sites



Note Before creating another cluster on sjc-sw2, create a third site SAC with the sac-sw2 switch. Clustering and IOA service must be enabled, and IOA interfaces must have been provisioned on the sac-sw2 switch.

To create another IOA cluster on sjc-sw2 with SAC, follow these steps:

Procedure

Step 1 sjc-sw2# **configure terminal**

Enters configuration mode.

Step 2 sjc-sw2(config)# **ioa cluster tape_vault_site2**

Specifies the cluster name and enters IOA cluster configuration submode. A cluster name can include a maximum of 31 alphabetical characters.

Step 3 sjc-sw2(config-ioa-cl)# **node local**

Adds the local switch to the cluster. Enters the node configuration mode.

Step 4 sjc-sw2(config-ioa-cl-node)**# interface ioa2/3**

Adds the IOA interface to the cluster.

Step 5 sjc-sw2(config-ioa-cl)# **node sac-sw2**

Adds the remote node to the cluster and enters the node configuration mode.

Step 6 sjc-sw2(config-ioa-cl-node)#**interface ioa2/3**

Adds the IOA interface to the cluster.

Example

The following example displays the multiple clusters created using the SJC site:

```
sjc-sw2# show ioa cluster summary

--------------------------------------------------------------------------------
Cluster          Sites              Status    Master Switch
--------------------------------------------------------------------------------
tape_vault       SJC,               online    172.25.231.19
                 RTP
tape_vault_site2 SAC,               online    172.25.231.19
                 SJC
```



Note

- You need to select a switch that you want to be the master switch as the seed switch when you create the IOA cluster. If you have multiple switches in a site, you add all the switches in a site that you want to manage the cluster before adding the switches from the remote site.
- In this example, the SJC site may be a natural consolidation point for management, and you may choose a switch from this site as the preferred master switch

Configuring IOA with NPV

You can use the Cisco MDS 9000 Family I/O Accelerator (IOA) with N port virtualization (NPV) to reduce the number of Fibre Channel domain IDs in SANs. Switches operating in NPV mode does not join a fabric or exchange traffic between NPV core switch links and end devices. You can deploy multiple edge switches without any shortage of domain IDs. NPV is not available in switch mode. To make NPV available on a switch, you must turn on NPV mode.

You can use the Cisco MDS 9000 Family I/O Accelerator (IOA) with N port ID virtualization (NPIV). NPIV efficiently utilizes the HBA ports on the blade servers in a data center and reduces the number of FCIDs assigned to the HBA ports.

The switches are not in NPV mode by default. NPV is supported in the following Cisco MDS 9000 switches:

- Cisco MDS 9124 Multilayer Fabric Switch
- Cisco MDS 9134 Fabric Switch
- Cisco Fabric Switch for HP c-Class BladeSystem
- Cisco Fabric Switch for IBM BladeCenter



Note Configuring IOA with NPV is supported in Cisco NX-OS Release 5.2(2) and later.

Guidelines for Configuring IOA with NPV

Follow these guidelines to configure IOA with NPV:

- Enable NPV on Cisco MDS 9124 switch or Cisco MDS 9134 switch.
- Enable NPIV on the NPV core switch.

Note To enable NPV on the NPV device switch, follow the guidelines specified in *Cisco MDS 9000 Family NX-OS Interfaces Configuration Guide, Release 5.0(1a)*.

- You must make sure that the NP link is active.
- You must configure NPV traffic map, F-port PortChannel and multiple NP links.
- Enable IOA and configure IOA cluster on the NPV core switch and on another node in the SAN. The IOA node can reside on any other Cisco MDS switches in the SAN other than the NPV core switch.
- Add IOA interfaces to the cluster.
- Add remote node and interface of the remote node.
- Activate the IOA flows with WA, TA, compression as per the requirement.
- You can activate multiple IOA flows and multiple IOA clusters.
- You must enable NPIV on NPV devices if you are using VMware hosts or servers for multiple FDISC, fabric discovery configuration over the same NP link.
- You can use up to 100 hosts with IOA active flows over a single NP link.
- You can use up to 100 hosts with IOA active flows over a FPC (F-port PortChannel)
- Beginning with Cisco MDS 9000 NX-OS Release 5.2(2), features such as FPC, TFPC, and FlexAttach virtual pWWN are supported.
- You can have an IOA node on the NPV core switch and also on any other switches.

Configuring NPIV on an NPV Core Switch, NPV on an NPV Device, and Activating NP Link

The following procedures are used to enable NPV and NPIV:

- Enabling NPIV on the NPV core switch
- Enabling NPV on the NPV device
- Configuring the interfaces connected to the NPV core switch as NP ports
- Configuring the port VSAN for the NP ports
- Configuring NPV link as an F port on the NPV core switch
- Configuring the port VSAN for the F ports
- Configuring the other server and target ports on the NPV device as F ports

Configuring NPIV on the NPV Core Switch

To enable NPIV and NPV, perform this task:

Procedure

Step 1 switch# **configure terminal**

Enters configuration mode.

Step 2 switch(config)# **feature npiv**

Enables NPIV mode on the NPV core switch.

Step 3 switch(config)# **interface fc 2/1**

switch(config-if)# **switchport mode F**

switch(config-if)# **no shutdown**

Configures the NPV core switch port as a F port and enables the interface.

Step 4 switch(config)# **vsan database**

switch(config-vsan-db)# **vsan 500 interface fc2/1**

Configures the port VSANs for the F port on the NPV core switch.

Configuring NPV on the NPV Device, Bringing Up the NP Port and NP Uplink

To configure NPV on an NPV device, perform this task:

Procedure

Step 1 switch# **configure terminal**

Enters configuration mode.

Step 2 switch(config)# **feature npv**

Enables NPV mode on an NPV device.

The switch is rebooted, and when it comes back up it is in NPV mode.

Note
A write-erase is performed during the reboot.

Step 3 switch(config)# **interface fc 2/1**

switch(config-if)# **switchport mode NP**

switch(config-if)# **no shutdown**

On the NPV device, selects the interfaces that will be connected to the aggregator switch, configures them as NP port and enables the interface.

Step 4 switch(config)# **exit**

Exits interface mode for the port.

Step 5 switch(config)# **vsan database**

switch(config-vsan-db)# **vsan 500 interface fc 1/1**

Configures the port VSANs for the NP port on the NPV device.

Step 6 switch(config)# **exit**

Exits interface mode for the port.

Step 7 switch(config)# **interface fc 1/2 - 6**

switch(config-if)# **switchport mode F**

switch(config-if)# **no shutdown**

Selects the remaining interfaces (2 through 6) which might be connected to end device such as hosts or targets on the NPV-enabled device, configures them as F ports and enables the interface.

Verifying the NPV Configuration

To view all the NPV devices in all the VSANs on the NPV core switch, enter the show fcns database command.

```
switch# show fcns database
VSAN 1:
--------------------------------------------------------------------------
FCID TYPE PWWN (VENDOR) FC4-TYPE:FEATURE
--------------------------------------------------------------------------
0x010000 N 20:01:00:0d:ec:2f:c1:40 (Cisco) npv
0x010001 N 20:02:00:0d:ec:2f:c1:40 (Cisco) npv
0x010200 N 21:00:00:e0:8b:83:01:a1 (Qlogic) scsi-fcp:init
0x010300 N 21:01:00:e0:8b:32:1a:8b (Qlogic) scsi-fcp:init
Total number of entries = 4
```

To display a list of the NPV devices that are logged in, along with VSANs, source information, pWWNs, and FCIDs, on the NPV device, enter the show npv flogi-table command.

```
switch# show npv flogi-table
--------------------------------------------------------------------------------
SERVER EXTERNAL
INTERFACE VSAN FCID PORT NAME NODE NAME INTERFACE
--------------------------------------------------------------------------------
fc1/19 1 0xee0008 10:00:00:00:c9:60:e4:9a 20:00:00:00:c9:60:e4:9a fc1/9
fc1/19 1 0xee0009 20:00:00:00:0a:00:00:01 20:00:00:00:c9:60:e4:9a fc1/1
fc1/19 1 0xee000a 20:00:00:00:0a:00:00:02 20:00:00:00:c9:60:e4:9a fc1/9
fc1/19 1 0xee000b 33:33:33:33:33:33:33:33 20:00:00:00:c9:60:e4:9a fc1/1
Total number of flogi = 4.
```

To display the status of the different servers and external interfaces, on the NPV device, enter the show npv status command.

```
switch# show npv status
npiv is enabled
External Interfaces:
====================
Interface: fc1/1, VSAN: 2, FCID: 0x1c0000, State: Up
Interface: fc1/2, VSAN: 3, FCID: 0x040000, State: Up
Number of External Interfaces: 2
Server Interfaces:
==================
Interface: fc1/7, VSAN: 2, NPIV: No, State: Up
Interface: fc1/8, VSAN: 3, NPIV: No, State: Up
Number of Server Interfaces: 2
```

Creating and Activating an IOA Cluster

To configure IOA flows, follow the guidelines specified in *Cisco MDS 9000 Family I/O Accelerator Configuration Guide, Release 4.2(1)*.

To verify the IOA configuration, follow the procedures specified in *Cisco MDS 9000 Family I/O Accelerator Configuration Guide, Release 4.2(1)*.

Configuring NPV on IOA

This section describes the following configuration procedures used to configure NPV on IOA:

Enabling NPV

To enable NPV, perform this task:

Procedure

Step 1 switch# **configure terminal**

Enters configuration mode.

Step 2 switch(config)# **feature npv**

Enables NPV mode on a NPV device.

The switch is rebooted, and when it comes back up, it is in NPV mode.

Note
A write-erase is performed during the reboot.

Step 3 switch(config)# **switchname sjc-sw1**

Configures the switch name.

Step 4 sjc-sw1(config)# **interface fc 2/1**

sjc-sw1(config-if)# **switchport mode NP**

sjc-sw1(config-if)# **no shutdown**

On the NPV device, selects the interfaces that will be connected to the aggregator switch, configures them as NP port and enables the interface.

Step 5 sjc-sw1(config)# **vsan database**

sjc-sw1(config-vsan-db)# **vsan 500 interface fc 1/6**

Configures the port VSANs for the NP port on the NPV device.

Step 6 sjc-sw1(config)# **exit**

Exits VSAN database mode for the port.

Step 7 sjc-sw1(config)#**interface fc 1/7 - 9**

sjc-sw1(config-if)#**switchport mode F**

sjc-sw1(config-if)# **no shutdown**

Configures the remaining interfaces (7 through 9) which might be connected to hosts as F ports and enables the interfaces.

Enabling NPIV on the NPV Core Switches

To enable NPIV on the NPV core switches, perform this task:

Procedure

Step 1 switch# **configure terminal**

Enters configuration mode.

Step 2 switch(config)# **feature npiv**

Enables NPIV mode on a NPV core switch.

Step 3 sjc-sw2(config)# **vsan database**

sjc-sw2(config-vsan-db)# **vsan 500 interface fc 1/6**

Configures the port VSANs for the NP port on the NPV device.

Step 4 sjc-sw2(config)# **exit**

Exits VSAN dtabase mode for the port.

Step 5 sjc-sw2(config)# **interface fc 1/6**

sjc-sw2(config-if)# **switchport mode F**

sjc-sw2(config-if)# **no shutdown**

Configures the interfaces as F mode and enables the interface.

Verifying the Configured NP Uplinks

Use the following **show** commands to confirm the functioning of the configured NP uplinks on the NPV device sjc-sw1:

```
sjc-sw1# show npv status
npiv is enabled
External Interfaces:
====================
Interface: fc1/6, VSAN: 500, FCID: 0xaf0000, State: Up
Number of External Interfaces: 1
Server Interfaces:
==================
Interface: fc1/7, VSAN: 500, State: Up
Interface: fc1/8, VSAN: 500, State: Up
Number of Server Interfaces: 2
sjc-sw1# show interface fc 1/6
```

```
fc1/6 is up
    Hardware is Fibre Channel, SFP is short wave laser w/o OFC (SN)
    Port WWN is 20:06:00:0d:ec:3d:92:00
    Admin port mode is NP, trunk mode is off
    snmp link state traps are enabled
    Port mode is NP
    Port vsan is 500
    Speed is 2 Gbps
    Rate mode is dedicated
    Transmit B2B Credit is 16
    Receive B2B Credit is 16
    Receive data field Size is 2112
    Beacon is turned off
    5 minutes input rate 1956320 bits/sec, 244540 bytes/sec, 3617 frames/sec
    5 minutes output rate 132841568 bits/sec, 16605196 bytes/sec, 11309 frames/sec
      6219674043 frames input, 349356203708 bytes
        0 discards, 0 errors
        0 CRC,  0 unknown class
        0 too long, 0 too short
      36666335463 frames output, 64666483082476 bytes
        512 discards, 0 errors
      36 input OLS, 23 LRR, 2 NOS, 0 loop inits
      29 output OLS, 17 LRR, 14 NOS, 0 loop inits
      0 receive B2B credit remaining
      16 transmit B2B credit remaining
      14 low priority transmit B2B credit remaining
    Interface last changed at Mon Oct 10 10:07:54 2011
sjc-sw1# sh npv flogi-table
--------------------------------------------------------------------------------
SERVER                                                                  EXTERNAL
INTERFACE VSAN FCID             PORT NAME               NODE NAME       INTERFACE
--------------------------------------------------------------------------------
fc1/7     500  0xbe005a 10:00:02:c8:01:cc:01:21 10:00:00:00:11:86:00:00 fc1/6
fc1/8     500  0xbe0214 10:00:02:c8:01:cc:01:81 10:00:00:00:11:86:00:00 fc1/6
Total number of flogi = 1
```

Use the following **show** commands to confirm the functioning of the configured NP uplinks on the NPV device sjc-sw2:

```
sjc-sw2# show npiv status
NPIV is enabled
sjc-sw2# show  int fc 1/5
fc1/9 is up
    Hardware is Fibre Channel, SFP is short wave laser w/o OFC (SN)
    Port WWN is 20:09:00:0d:ec:3d:92:00
    Admin port mode is F, trunk mode is off
    snmp link state traps are enabled
    Port mode is F, FCID is 0xbe0044
    Port vsan is 500
    Speed is 2 Gbps
    Rate mode is dedicated
    Transmit B2B Credit is 16
    Receive B2B Credit is 16
    Receive data field Size is 2112
    Beacon is turned off
    5 minutes input rate 0 bits/sec, 0 bytes/sec, 0 frames/sec
    5 minutes output rate 8 bits/sec, 1 bytes/sec, 0 frames/sec
      4283 frames input, 231280 bytes
        0 discards, 0 errors
        0 CRC,  0 unknown class
        0 too long, 0 too short
      4348 frames output, 2295004 bytes
        0 discards, 0 errors
```

```
  1 input OLS, 1 LRR, 2 NOS, 0 loop inits
  1 output OLS, 1 LRR, 1 NOS, 0 loop inits
  16 receive B2B credit remaining
  16 transmit B2B credit remaining
  16 low priority transmit B2B credit remaining
Interface last changed at Fri Sep 30 09:24:40 2011
```

Enabling IOA on the IOA Nodes

To enable IOA on the first IOA node sjc-sw2 in site SJC, perform this task:

Procedure

Step 1 sjc-sw2# **configure terminal**

Enters configuration mode.

Step 2 sjc-sw2(config)# **feature cluster**

Enables the feature cluster on IOA node.

Step 3 sjc-sw2(config)# **feature ioa**

Enables the feature IOA on IOA node.

To enable IOA on the first IOA node rtp-sw2 in Site RTP, perform this task:

Procedure

Step 1 sjc-sw2# config t

Enters configuration mode.

Step 2 sjc-sw2(config)# feature cluster

Enables the feature cluster on IOA node.

Step 3 sjc-sw2(config)# feature ioa

Enables the feature IOA on IOA node.

Classifying the Switches into IOA Sites

To configure the IOA site on sjc-sw2, perform this task:

Procedure

Step 1 sjc-sw2# **configure terminal**

Enters configuration mode.

Step 2 sjc-sw2(config)# **ioa site-local SJC**

Classifies the switches into IOA site.

To configure the IOA site on rtp-sw2, perform this task:

Procedure

Step 1 rtp-sw2# **configure terminal**

Enters configuration mode.

Step 2 rtp-sw2(config)# **ioa site-local RTP**

Classifies the switches into IOA site.

Configuring IOA Interfaces

To configure IOA interface on sjc-sw2, perform this task:

Procedure

Step 1 sjc-sw2# **configure terminal**

Enters configuration mode.

Step 2 sjc-sw2(config)# **interface ioa 1/1**

sjc-sw2(config-if)# **no shutdown**

Configures IOA on service engine 1 in slot 1 and enables the interface.

To configure IOA interface on rtp-sw2, perform this task:

Procedure

Step 1 rtp-sw2# **configure terminal**

Enters configuration mode.

Step 2 rtp-sw2(config)# **interface ioa 1/1**

rtp-sw2(config-if)# **no shutdown**

Configures IOA on service engine 1 in slot 1 and enables the interface.

Configuring IOA Interfaces

To configure IOA interface on rtp-sw2, perform this task:

SUMMARY STEPS

1. rtp-sw2# config t
2. rtp-sw2(config)# interface ioa 1/1

DETAILED STEPS

Procedure

	Command or Action	Purpose
Step 1	rtp-sw2# config t	Enters configuration mode.
Step 2	rtp-sw2(config)# interface ioa 1/1 **Example:** `rtp-sw2(config-if)# no shutdown`	Configures IOA on service engine 1 in slot 1 and enables the interface.

What to do next

Configuring IOA Cluster

To configure IOA cluster on sjc-sw2, perform this task:

Procedure

Step 1 sjc-sw2# **configure terminal**

Enters configuration mode.

Step 2 sjc-sw2(config)# **ioa cluster DC1**

Configures IOA cluster. Cluster names are case sensitive.

Configuring Nodes to the IOA Cluster

To add an IOA cluster on sjc-sw2 , perform this task:

Procedure

Step 1 sjc-sw2# **configure terminal**

Enters configuration mode.

Step 2 sjc-sw2(config)# **ioa cluster DC1**

Enters the IOA cluster sub mode

Step 3 sjc-sw2(config-ioa-cl)# **node local**

Adds the switch sjc-sw2 to the cluster.

Step 4 sjc-sw2(config-ioa-cl-node)# **exit**

Exits the IOA cluster node sub mode.

Step 5 sjc-sw2(config-ioa-cl)# **node rtp-sw2**

Adds the remote IOA node into the same cluster. The remote node can be added into the cluster by using its switch name or IPv4/IPv6 management interface address.

Verifying the IOA Cluster Configuration

Use the following **show** commands to confirm the functioning of the IOA cluster on sjc-sw2:

```
sjc-sw2# show ioa cluster
IOA Cluster is DC1
  Cluster ID is 0x2003000573cbe602
  Cluster status is online
  Is between sites SJC and RTP
  Total Nodes are 2
  Cluster Infra Status : Operational
  Cluster is Administratively Up
  Cluster Config Version : 707
  SSL for ICN : Not Configured
sjc-sw2# show ioa cluster DC1 summary
--------------------------------------------------------------------------------
Cluster          Sites              Status    Master Switch
--------------------------------------------------------------------------------
DC1              SJC,                online    10.65.217.48
                      RTP
sjc-sw2# show ioa cluster DC1 node
Node 10.65.217.48 is local switch
  Node ID is 1
  IP address is 10.65.217.48
```

```
  Status is online
  Belongs to Site SJC
  Node is the master switch
Node 10.65.217.56 is remote switch
  Node ID is 2
  IP address is 10.65.217.56
  Status is online
  Belongs to Site RTP
  Node is not master switch
```



Note You can use the same **show** command to verify the IOA configuration on rtp-sw2.

Configuring Interfaces in the IOA Cluster

To add IOA interfaces to the IOA cluster on the Master switch sjc-sw2, perform this task:

Procedure

Step 1 sjc-sw2# **configure terminal**

Enters configuration mode.

Step 2 sjc-sw2(config)# **ioa cluster DC1**

Enters the IOA cluster submode.

Step 3 sjc-sw2(config-ioa-cl)# **node local**

Adds the switch sjc-sw2 to the cluster.

Step 4 sjc-sw2(config-ioa-cl-node)# **interface ioa 1/1**

Adds the interface of the local IOA node into the cluster.

Step 5 sjc-sw2(config-ioa-cl-node)# **exit**

Exits the IOA cluster submode.

Step 6 sjc-sw2(config-ioa-cl)# **node rtp-sw2**

Adds the remote IOA node into the same cluster. The remote node can be added into the cluster by using its switch name or IPv4/IPv6 management interface address.

Step 7 sjc-sw2(config-ioa-cl-node)# **interface ioa 1/1**

Adds the interface of the remote IOA node into the cluster.

Verifying the Cluster Interface Configuration

Use the following **show** commands to confirm the functioning configured cluster interface:

```
sjc-sw2# show interface ioa 1/1
```

```
ioa1/1 is up
    Member of cluster DC1
    21368133123 device packets in, 6851375618 device packets out
    31397026863066 device bytes in, 476831158620 device bytes out
    914301804 peer packets in, 8706253930 peer packets out
    56107433228 peer bytes in, 17877494274392 peer bytes out
    0 i-t create request, 0 i-t create destroy
    0 i-t activate request, 0 i-t deactivate request
sjc-sw2# show ioa cluster DC1 interface summary
--------------------------------------------------------------------------------
Switch             Interface         Status       Flows
--------------------------------------------------------------------------------
10.65.217.48(L)    ioa1/1            up           --
10.65.217.56       ioa1/1            up           --
```



Note You can use the same **show** command to verify the IOA cluster and interface configuration on rtp-sw2.

Adding N Ports to the IOA cluster

To add N ports (hosts and targets) to the IOA cluster on the master switch sjc-sw2, perform this task:

Procedure

Step 1 sjc-sw2# **configure terminal**

Enters configuration mode.

Step 2 sjc-sw2(config)# **ioa cluster DC1**

Enters the IOA cluster submode.

Step 3 sjc-sw2(config-ioa-cl)# **nport pwwn 10:00:02:c8:01:cc:01:01 site SJC vsan 500**

Adds the N port in VSAN 500 to the cluster.

Step 4 sjc-sw2(config-ioa-cl)# **nport pwwn 10:00:02:c8:01:cc:02:01 site RTP vsan 500**

Adds another N port in remote IOA site in the same VSAN to the cluster.

Step 5 sjc-sw2(config-ioa-cl-node)# **exit**

Exits the IOA cluster submode.

Verifying the Configured N Ports in the IOA Cluster

Use the following **show** command to confirm the functioning of the configured N ports in the IOA cluster:

```
sjc-sw2# show ioa cluster DC1 nports
--------------------------------------------------------------------------------
P-WWN                          Site                           Vsan
--------------------------------------------------------------------------------
```

```
10:00:02:c8:01:cc:01:01                SITE sjc                    500
10:00:02:c8:01:cc:02:01        SITE rtp                                500
```

Note You can use the same command to verify the IOA cluster and interface configuration on rtp-sw2.

Configuring IOA Flows in the Cluster

To configure IOA flows in the IOA cluster on the master switch sjc-sw2, perform this task:

Procedure

Step 1 sjc-sw2# **configure terminal**

Enters configuration mode.

Step 2 sjc-sw2(config)# **ioa cluster DC1**

Enters the IOA cluster submode. Cluster names are case sensitive.

Step 3 sjc-sw2(config-ioa-cl)# **flowgroup Dep1**

Configures an IOA flow group.

Step 4 sjc-sw2(config-ioa-cl-flgrp)# **host 10:00:02:c8:01:cc:01:01 target 10:00:02:c8:01:cc:02:01**

Creates an IOA flow with write acceleration.

Step 5 sjc-sw2(config-ioa-cl-flgrp)# **exit**

Exits IOA cluster flow group submode.

Verifying the Configured IOA Flow

Use the following **show** commands to confirm the functioning of the IOA flow configuration and to verify status of the flow on the master switch sjc-sw2:

```
sjc-sw2# show ioa cluster DC1 flows flowgroup Dep1
--------------------------------------------------------------------------------
Host WWN,                VSAN     WA  TA  Comp  Status    Switch,Interface
Target WWN                                               Pair
--------------------------------------------------------------------------------
10:00:02:c8:01:cc:01:01, 500      Y    N    N   online    10.65.217.48, ioa1/1
10:00:02:c8:01:cc:02:01  500                              10.65.217.56, ioa1/1
sjc-sw2# show ioa cluster DC1 flows flowgroup Dep1 detail
Host 10:00:02:c8:01:cc:01:01, VSAN 500, Target 10:00:02:c8:01:cc:02:01, VSAN 500
  Is online
  Belongs to flowgroup Dep1
  Is enabled for WA,
  Is assigned to
    Switch 10.65.217.48     Interface ioa1/1 (Host Site)
    Switch 10.65.217.56     Interface ioa1/1 (Target Site)
```

Displaying Interface Statistics

Use the following **show** commands to verify the IOA interface counters when live packets are ran over the IOA flow:

```
sjc-sw2# show interface ioa 1/1 counters
ioa1/1
    21523240117 device packets in, 6901040984 device packets out
    31625069090806 device bytes in, 480287657508 device bytes out
    920937376 peer packets in, 8769431691 peer packets out
    56514685912 peer bytes in, 18007222544310 peer bytes out
    1 i-t create request, 0 i-t create destroy
    1 i-t activate request, 0 i-t deactivate request
```

The following example shows the average for 5minutes, 12 hours and 24 hours respectively:

```
sjc-sw2(config-if)# show interface ioa 1/1 counters brief
-------------------------------------------------------------------------------
Interface          Rate          Rate          Rate          Total
                        MB/s          MB/s          MB/s          Bytes
                       (5min)        (12hr)        (24hr)        (MB)
-------------------------------------------------------------------------------
                       To  Device (Average)
ioa1/1              0.00        0.00             0.00             0.02
                       To  Peer   (Average)
                        0.00             0.00         0.00             0.05
sjc-sw2# show ioa internal interface ioa 1/1 summary
---- ---------------------- ---- ------------- ---- ---
FLOW HOST                   VSAN STATUS        COMP ACC
     TARGET
---- ---------------------- ---- ------------- ---- ---
1    10:00:02:c8:01:cc:01:01 500  ACTIVE        NO  TA
     10:00:02:c8:01:cc:02:01
```

Additional Configurations for the Features Supported by NPV on IOA

NP Link Trunking

Configuring an NP Uplink Port

To configure an NP link, you must bring up the TF-TNP link between an F port in the NPIV core switch and then configure a NP port in the NPV switch.

To configure an NPV core switch, perform this task:

Procedure

Step 1 sjc-sw2# **configure terminal**

Enters configuration mode.

Step 2 sjc-sw2#(config)# **feature fport-channel-trunk**

Enables the F port trunking and channeling protocol on the NPV core switch.

Step 3 sjc-sw2#(config)# **feature npiv**

Enables NPIV on the NPV core switch.

Step 4 sjc-sw2#(config)# **interface fc1/2**

sjc-sw2#(config-if)# **switchport mode F**

sjc-sw2#(config-if)# **switchport trunk mode on**

Configures the port mode to auto, F, or Fx on the NPV core switch and enables Trunk mode on.

Step 5 sjc-sw2#(config)# **interface fc1/2**

sjc-sw2#(config-if)# **no shut**

Turns on the port administrative state on NPV core switch.

To configure an NPV device switch, perform this task:

Procedure

Step 1 sjc-sw1# **configure terminal**

Enters configuration mode.

Step 2 sjc-sw1(config)# **interface fc 1/2**

sjc-sw1(config)# **switchport mode NP**

sjc-sw1(config-if)# **switchport trunk mode on**

Configures the port mode to NP on the NPV switch and enabled Trunk mode on.

Step 3 sjc-sw1(config)# **interface fc1/2**

sjc-sw1(config-if)# **no shut**

Turns on the port administrative state on NPV core switch.

Verifying the Configured Trunking NP Uplink Port on the NPV Core Switch

Use the following **show** command to confirm the functioning configured NPV core switch:

```
sjc-sw2(config-if)# show int fc 1/2
fc1/2 is trunking
    Hardware is Fibre Channel, SFP is short wave laser w/o OFC (SN)
    Port WWN is 20:04:00:05:73:cb:e6:00
    Admin port mode is auto, trunk mode is on
    snmp link state traps are enabled
```

```
    Port mode is TF
    Port vsan is 9
    Speed is 4 Gbps
    Rate mode is dedicated
    Transmit B2B Credit is 16
    Receive B2B Credit is 16
    Receive data field Size is 2112
    Beacon is turned off
    Belongs to port-channel 21
    Trunk vsans (admin allowed and active) (9-13)
    Trunk vsans (up)                       (9,10)
    Trunk vsans (isolated)                 ()
    Trunk vsans (initializing)             (11-13)
    5 minutes input rate 0 bits/sec, 0 bytes/sec, 0 frames/sec
    5 minutes output rate 8 bits/sec, 1 bytes/sec, 0 frames/sec
      231 frames input, 16680 bytes
        0 discards, 0 errors
        0 CRC,  0 unknown class
        0 too long, 0 too short
      248 frames output, 114660 bytes
        0 discards, 0 errors
      1 input OLS, 1 LRR, 1 NOS, 0 loop inits
      2 output OLS, 3 LRR, 0 NOS, 1 loop inits
      16 receive B2B credit remaining
      16 transmit B2B credit remaining
      14 low priority transmit B2B credit remaining
```

Verifying the Configured Trunking NP Uplink Port on NPV Device Switch

Use the following **show** commands to confirm the functioning configured NPV device switch:

```
sjc-sw1(config-if)# show int fc 1/2
fc1/2 is trunking
    Hardware is Fibre Channel, SFP is short wave laser w/o OFC (SN)
    Port WWN is 20:06:00:0d:ec:3d:92:00
    Admin port mode is NP, trunk mode is on
    snmp link state traps are enabled
    Port mode is TNP
    Port vsan is 9
    Speed is 4 Gbps
    Rate mode is dedicated
    Transmit B2B Credit is 16
    Receive B2B Credit is 16
    Receive data field Size is 2112
    Beacon is turned off
    Belongs to port-channel 21
    Trunk vsans (admin allowed and active) (9-13)
    Trunk vsans (up)                       (9,10)
    Trunk vsans (isolated)                 ()
    Trunk vsans (initializing)             (11-13)
    5 minutes input rate 0 bits/sec, 0 bytes/sec, 0 frames/sec
    5 minutes output rate 0 bits/sec, 0 bytes/sec, 0 frames/sec
      2837806124 frames input, 147817029296 bytes
        0 discards, 0 errors
        0 CRC,  0 unknown class
        0 too long, 0 too short
      26077437111 frames output, 49186719497132 bytes
        512 discards, 0 errors
      36 input OLS, 23 LRR, 2 NOS, 0 loop inits
      29 output OLS, 17 LRR, 14 NOS, 0 loop inits
      16 receive B2B credit remaining
      16 transmit B2B credit remaining
```

```
  14 low priority transmit B2B credit remaining
Interface last changed at Mon Oct 10 10:07:54 2011
```

Note In the case of ports, after the handshake, one of the allowed VSANs is moved to Up state. All other VSANs will be in initial state even though the handshake with the peer is completed successfuly. Each VSAN is moved from initial state to Up state when a server or target logs in through the trunked F or NP ports in the corresponding VSAN. For more information about configuring ports and TF-TNP ports, refer to the *Cisco MDS 9000 Family NX-OS Interfaces Configuration Guide, Release 5.0(1a)*.

Configuring F-PortChannel

To configure F-PortChannel (FPC) in shared mode and bring up the link between F ports on the NPIV core switches and NP ports on the NPV use the procedure in this section.

Note Configuring FPC is not supported on the MDS 91x4 switches.

Configuring F-PortChannel on the NPV Core Switch

To configure the F-PortChannel on an NPV core switch, perform this task:

Procedure

Step 1 sjc-sw2# **configure terminal**

Enters configuration mode.

Step 2 sjc-sw2(config)# **feature fport-channel-trunk**

Enables the F port trunking and channeling protocol on the NPV core switch.

Step 3 sjc-sw2(config)# **feature npiv**

Enables NPIV on the NPV core switch.

Step 4 sjc-sw2(config)# **interface port-channel 1**

sjc-sw2(config-if)# **switchport mode F**

sjc-sw2(config-if)# **channel mode active**

sjc-sw2(config-if)# **switchport trunk mode off**

sjc-sw2(config-if)# **switchport rate-mode shared**

sjc-sw2(config-if)# **exit**

Creates the PortChannel on the NPV core switch.

Step 5 sjc-sw2(config)# **interface fc2/1-3**

sjc-sw2(config-if)# **shut**

sjc-sw2(config-if)# **switchport mode F**

sjc-sw2(config-if)# **switchport trunk mode off**

sjc-sw2(config-if)# **switchport speed 4000**

sjc-sw2(config-if)# **switchport rate-mode shared**

sjc-sw2(config-if)# **channel-group 1**

sjc-sw2(config-if)# **exit**

Creates the PortChannel member interfaces on the NPV core switch.

Configuring NP-PortChannel on an NPV Device Switch

To configure NP-PortChannel on an NPV device switch, perform this task:

Procedure

Step 1 sjc-sw2# **configure terminal**

Enters configuration mode.

Step 2 switch(config)# **interface port-channel 1**

switch(config-if)# switchport mode NP

switch(config-if)# **switchport mode NP**

switch(config-if)# **switchport rate-mode shared**

switch(config-if)# **exit**

Creates the PortChannel on the NPV switch.

Step 3 switch(config)# **interface fc1/1-3**

switch(config-if)# **shut**

switch(config-if)# **switchport mode NP**

switch(config-if)# **switchport speed 4000**

switch(config-if)# **switchport rate-mode shared**

switch(config-if)# **switchport trunk mode off**

switch(config-if)# **channel-group 1**

switch(config-if)# **no shut**

switch(config-if)# **exit**

Creates the PortChannel member interfaces on the NPV switch.

Turning on the Administrative State of PortChannel Member Interfaces in the NPV Core Switch

To turn on the administrative state of all the PortChannel member interfaces in the NPV core switch, perform this task:

Procedure

Step 1 sjc-sw2# **configure terminal**

Enters configuration mode.

Step 2 switch(config)# **interface fc1/1-3**

switch(config-if)# **shut**

switch(config-if)# **no shut**

switch(config-if)# **exit**

Turns on the administrative state of the PortChannel members.

Turning on the Administrative State of PortChannel member interfaces in NPV Device Switch

To turn on the administrative state of all the PortChannel member interfaces in the NPV device switch, perform this task:

Procedure

Step 1 sjc-sw2# **configure terminal**

Enters configuration mode.

Step 2 switch(config)# **interface fc2/1-3**

switch(config-if)# **shut**

switch(config-if)# **no shut**

switch(config-if)# **exit**

Turns on the administrative state of the PortChannel members.

Verifying the Configured PortChannel of NP Links

Use the following **show** command to verify the configured PortChannel on the NPV core switch side:

```
sjc-sw2(config-if)# show interface port-channel 1
port-channel 1 is up
    Hardware is Fibre Channel
    Port WWN is 24:15:00:05:73:cb:e6:00
```

```
    Admin port mode is NP, trunk mode is off
    snmp link state traps are enabled
    Port mode is NP
    Port vsan is 500
    Speed is 8 Gbps
    Trunk vsans (admin allowed and active) (500-512)
    Trunk vsans (up)                       (500,512)
    Trunk vsans (isolated)                 ()
    Trunk vsans (initializing)             (501-511)
    5 minutes input rate 0 bits/sec, 0 bytes/sec, 0 frames/sec
    5 minutes output rate 0 bits/sec, 0 bytes/sec, 0 frames/sec
      792 frames input, 51848 bytes
        0 discards, 0 errors
        0 CRC,  0 unknown class
        0 too long, 0 too short
      811 frames output, 417880 bytes
      0 discards, 0 errors
      2 input OLS, 2 LRR, 1 NOS, 0 loop inits
      3 output OLS, 4 LRR, 0 NOS, 2 loop inits
    Member[1] : fc2/1
    Member[2] : fc2/2
    Member[3] : fc2/3
    Interface last changed at Wed Oct 12 08:12:36 2011
```

Use the following **show** commands to verify that the configured PortChannel on the NPV device switch side:

```
switch# show interface port-channel 1
port-channel 1 is trunking
    Hardware is Fibre Channel
    Port WWN is 24:15:00:05:73:cb:e6:00
    Admin port mode is auto, trunk mode is off
    snmp link state traps are enabled
    Port mode is NP
    Port vsan is 500
    Speed is 8 Gbps
    Trunk vsans (admin allowed and active) (500-512)
    Trunk vsans (up)                       (500,512)
    Trunk vsans (isolated)                 ()
    Trunk vsans (initializing)             (501-511)
    5 minutes input rate 0 bits/sec, 0 bytes/sec, 0 frames/sec
    5 minutes output rate 0 bits/sec, 0 bytes/sec, 0 frames/sec
      792 frames input, 51848 bytes
        0 discards, 0 errors
        0 CRC,  0 unknown class
        0 too long, 0 too short
      811 frames output, 417880 bytes
      0 discards, 0 errors
      2 input OLS, 2 LRR, 1 NOS, 0 loop inits
      3 output OLS, 4 LRR, 0 NOS, 2 loop inits
    Member[1] : fc1/1
    Member[2] : fc1/2
    Member[3] : fc1/3
    Interface last changed at Wed Oct 12 08:12:36 2011
```

Example for Configuring TF-TNP PortChannel Links

This example shows the following configuration procedures used to change the PortChannels in dedicated mode to bring up the TF-TNP PortChannel link between TF ports in the NPIV core switch, and TNP ports in the NPV switch.

Configuring the PortChannel on the NPV Core Switch

To configure the PortChannel on an NPV core switch, perform this task:

Procedure

Step 1 sjc-sw2# **configure terminal**

Enters configuration mode.

Step 2 switch(config)# **feature fport-channel-trunk**

Enables the F port trunking and channeling protocol on the NPV core switch.

Step 3 switch(config)# **feature npiv**

Enables NPIV on the NPV core switch.

Step 4 switch(config)# **interface port-channel 1**

switch(config-if)# **switchport mode F**

switch(config-if)# **switchport rate-mode dedicated**

switch(config-if)# **channel mode active**

switch(config-if)# **exit**

Creates the PortChannel on the NPV switch.

Step 5 switch(config)# **interface fc2/1-3**

switch(config-if)# **shut**

switch(config-if)# **switchport mode F**

switch(config-if)# **switchport speed 4000**

switch(config-if)# **switchport rate-mode dedicated**

switch(config-if)# **switchport trunk mode on**

switch(config-if)# **channel-group 1**

switch(config-if)# **exit**

Creates the PortChannel member interfaces on the NPV core switch.

Configuring PortChannel on the NPV Device Switch

To configure the PortChannel on an NPV device switch, perform this task:

Procedure

Step 1 sjc-sw1# **configure terminal**

Enters configuration mode.

Step 2 switch(config)# **interface port-channel 1**

switch(config-if)# **switchport rate-mode dedicated**

switch(config-if)# **switchport mode NP**

switch(config-if)# **no shutdown**

switch(config-if)# **exit**

Creates the PortChannel on the NPV device switch.

Step 3 switch(config)# **interface fc2/1-3**

switch(config-if)# **shut**

switch(config-if)# **switchport mode NP**

switch(config-if)# **switchport speed 4000**

switch(config-if)# **switchport rate-mode dedicated**

switch(config-if)# **switchport trunk mode on**

switch(config-if)# **channel-group 1**

switch(config-if)# **exit**

Creates the PortChannel member interfaces on the NPV device switch.

Turning on the Administrative State of ProvinceChannel Member Interfaces in the NPV Core Switch

To turn on the administrative state of all the PortChannel member interfaces in NPV core switch, perform this task:

Procedure

Step 1 sjc-sw2# **configure terminal**

Enters configuration mode.

Step 2 switch(config)# **interface fc2/1-3**

switch(config-if)# **shut**

switch(config-if)# **no shut**

switch(config-if)# **exit**

Turn on the administrative state of the PortChannel members.

Turning on the Administrative State of ProvinceChannel Member Interfaces in the NPV Device Switch

To turn on the administrative state of all the PortChannel member interfaces in NPV device switch, perform this task:

Procedure

Step 1 sjc-sw2# **configure terminal**

Enters configuration mode.

Step 2 switch(config)# **interface fc1/1-3**

switch(config-if)# **shut**

switch(config-if)# **no shut**

switch(config-if)# **exit**

Turn on the administrative state of the PortChannel members.

Note
The speed configuration must be the same for all member interfaces in a PortChannel. You must ensure that the required bandwidth is available to all the ports while configuring the channel in dedicated mode.

Verifying the Configured PortChannel of TF-TNP Links

Use the following **show** commands to verify the configured PortChannel on the NPV core switch side:

```
sjc-sw2# show interface port-channel 1
port-channel 1 is trunking
    Hardware is Fibre Channel
    Port WWN is 24:15:00:05:73:cb:e6:00
    Admin port mode is auto, trunk mode is on
    snmp link state traps are enabled
    Port mode is TF
    Port vsan is 500
    Speed is 8 Gbps
    Trunk vsans (admin allowed and active) (500-512)
    Trunk vsans (up)                       (500,512)
    Trunk vsans (isolated)                 ()
    Trunk vsans (initializing)             (501-511)
    5 minutes input rate 0 bits/sec, 0 bytes/sec, 0 frames/sec
    5 minutes output rate 0 bits/sec, 0 bytes/sec, 0 frames/sec
      792 frames input, 51848 bytes
        0 discards, 0 errors
        0 CRC,  0 unknown class
        0 too long, 0 too short
      811 frames output, 417880 bytes
      0 discards, 0 errors
      2 input OLS, 2 LRR, 1 NOS, 0 loop inits
      3 output OLS, 4 LRR, 0 NOS, 2 loop inits
    Member[1] : fc2/1
    Member[2] : fc2/2
    Member[3] : fc2/3
    Interface last changed at Wed Oct 12 08:22:36 2011
```

Use the following **show** commands to verify the configured PortChannel on the NPV device switch side:

```
sjc-sw2# show interface port-channel 1
port-channel 1 is trunking
    Hardware is Fibre Channel
    Port WWN is 24:15:00:05:73:cb:e6:00
    Admin port mode is auto, trunk mode is on
    snmp link state traps are enabled
    Port mode is TNP
    Port vsan is 500
    Speed is 8 Gbps
    Trunk vsans (admin allowed and active) (500-512)
    Trunk vsans (up)                       (500,512)
    Trunk vsans (isolated)                 ()
    Trunk vsans (initializing)             (501-511)
    5 minutes input rate 0 bits/sec, 0 bytes/sec, 0 frames/sec
    5 minutes output rate 0 bits/sec, 0 bytes/sec, 0 frames/sec
      792 frames input, 51848 bytes
        0 discards, 0 errors
        0 CRC,  0 unknown class
        0 too long, 0 too short
      811 frames output, 417880 bytes
      0 discards, 0 errors
      2 input OLS, 2 LRR, 1 NOS, 0 loop inits
      3 output OLS, 4 LRR, 0 NOS, 2 loop inits
    Member[1] : fc1/1
    Member[2] : fc1/2
    Member[3] : fc1/3
    Interface last changed at Wed Oct 12 08:22:36 2011
```

Configuring FlexAttach Virtual pWWN on an NPV Switch

The FlexAttach virtual pWWN feature facilitates server and configuration management. In a SAN environment, the server installation or replacement requires interaction and coordination among the SAN and server administrators. It is important that the SAN configuration does not change when a new server is installed, or when an existing server is replaced.

FlexAttach virtual pWWN minimizes the interaction between the server administrator and the SAN administrator by abstracting the real pWWN using virtual pWWNs. When FlexAttach virtual pWWN is enabled on an interface, a virtual pWWN is assigned to the server interface. The real pWWN is replaced by a virtual pWWN, which is used for a SAN configuration such as zoning.

With pWWNs configured on NPV switch in various forms as described in the next section, IOA works seamlessly with pWWNs. The pWWNs feature is enabled automatically, manually, or by mapping pWWN to virtual pWWN.

Automatically Enabling FlexAttach Virtual pWWN

The virtual pWWN is enabled automatically on all of the NPV switches or per port on the NPV device. When enabled automatically, a virtual WWN is generated from the device switch WWN. This WWN is used as the virtual pWWN. Virtual pWWNs are generated using the local switch WWNs.



Note The port must be in a shut state when the virtual pWWN is enabled.

To enable virtual pWWN automatically, perform this task:

Procedure

Step 1 sjc-sw2# **configure terminal**

Enters configuration mode.

Step 2 swtich(config)# **flex-attach virtual-pwwn auto [interface interface-list]**

Enables FlexAttach virtual pWWN automatically for the interfaces.

Step 3 switch(config)# **flex-attach commit**

Commits the configuration.

Manually Enabling FlexAttach Virtual pWWN

You can manually assign a WWN to the interface, without generating it through the switch. Several checks are done by the NPV core to ensure the uniqueness of virtual pWWNs in the switch. When duplicate virtual pWWNs are configured, the subsequent logins are rejected by the NPV core switch.



Note
- Some ports may be in automode, some in manual mode, and the virtual pWWNs need not be assigned.
- The port must be in a shut state when a virtual pWWN is enabled.
- The interface mentioned in the interface value must be in a shut state.

To enable virtual pWWN manually, perform this task:

Procedure

Step 1 switch# **configure terminal**

Enters configuration mode.

Step 2 swtich(config)# **flex-attach virtual-pwwn vpwwn interface interface**

Enables FlexAttach virtual pWWN manually for the interfaces.

Step 3 switch(config)# **flex-attach commit**

Commits the configuration.

Verifying the Configured FlexAttach Virtual pWWN

Use the following **show** command to verify the type and value of virtual pWWNs are correct:

```
sjc-sw1# show flex-attach virtual-wwn
```

```
VIRTUAL PORT WWNS ASSIGNED TO INTERFACES
------------------------------------------------------------------------
VSAN INTERFACE VIRTUAL-PWWN AUTO LAST-CHANGE
------------------------------------------------------------------------
1 fc1/1 00:00:00:00:00:00:00:00
1 fc1/2 22:73:00:05:30:01:6e:1e TRUE Thu Jan 31 01:58:52 2008
1 fc1/3 22:5e:00:05:30:01:6e:1e TRUE Thu Jan 31 01:58:52 2008
1 fc1/4 22:5f:00:05:30:01:6e:1e TRUE Thu Jan 31 01:58:52 2008
1 fc1/5 22:74:00:05:30:01:6e:1e TRUE Thu Jan 31 01:26:24 2008
1 fc1/6 22:60:00:05:30:01:6e:1e TRUE Thu Jan 31 01:58:52 2008
1 fc1/7 22:61:00:05:30:01:6e:1e TRUE Thu Jan 31 01:58:52 2008
1 fc1/8 22:62:00:05:30:01:6e:1e TRUE Thu Jan 31 01:58:52 2008
1 fc1/9 22:63:00:05:30:01:6e:1e TRUE Thu Jan 31 01:58:52 2008
1 fc1/10 22:64:00:05:30:01:6e:1e TRUE Thu Jan 31 01:58:52 2008
1 fc1/11 22:65:00:05:30:01:6e:1e TRUE Thu Jan 31 01:58:52 2008
1 fc1/12 22:66:00:05:30:01:6e:1e TRUE Thu Jan 31 01:58:52 2008
```

Verifying the Configured FlexAttach Virtual pWWN

Use the following **show** commands to verify that the end device is logged with the correct virtual WWNs:

```
switch# show fcns database
VSAN 1:
------------------------------------------------------------------------
FCID TYPE PWWN (VENDOR) FC4-TYPE:FEATURE
------------------------------------------------------------------------
0x010000 N 20:01:00:0d:ec:2f:c1:40 (Cisco) npv
0x010001 N 20:02:00:0d:ec:2f:c1:40 (Cisco) npv
0x010200 N 21:00:00:e0:8b:83:01:a1 (Qlogic) scsi-fcp:init
0x010300 N 21:01:00:e0:8b:32:1a:8b (Qlogic) scsi-fcp:init
Total number of entries = 4
```

Configuring NPV Traffic Management on NPV Switches with IOA

Configuring NPV traffic management involves configuring a list of external interfaces to the servers, and enabling or disabling disruptive load balancing. The NPV traffic management feature is enabled after configuring NPV.

Configuring a List of External Interfaces per Server Interface

A list of external interfaces is linked to the server interfaces when the server interface is down, or if the specified external interface list includes the external interface already in use.

To configure the list of external interfaces per server interface, perform this task:

Procedure

Step 1 sjc-sw2# **configure terminal**

Enters configuration mode.

Step 2 switch(config)# **npv traffic-map server-interface fc 1/1-3 external-interface fc 1/8-10**

Allows you to configure a list of external FC interfaces to a range of server interface.

Step 3 switch(config)# **npv traffic-map server-interface fc1/1-3 external-interface port-channel 10-12**

Allows you to configure a list of external PortChannel interfaces per server interface.

Step 4 switch(config)# **no npv traffic-map server-interface fc 1/1-3 external-interface fc 1/8-10**

Disables the NPV traffic management feature on the NPV device switch.

Note
You must map the non-PortChannel interfaces and PortChannel interfaces to the server interfaces, in two steps.

Enabling or Disabling the Global Policy for Disruptive Load Balancing

Disruptive load balancing allows you to review the load on all the external interfaces and balance the load disruptively. Disruptive load balancing is done by moving the servers using heavily loaded external interfaces, to the external interfaces running with fewer loads.

To enable or disable the global policy for disruptive load balancing, perform this task:

Procedure

Step 1 sjc-sw2# **configure terminal**

Enters configuration mode.

Step 2 switch(config)# **npv auto-load-balance disruptive**

Enables disruptive load balancing on the NPV device.

Step 3 switch(config)# **no npv auto-load-balance disruptive**

Disables disruptive load balancing on the NPV device.

Verifying the NPV Traffic Management on an NPV Switch

Use the following **show** command to display the NPV traffic map on an NPV switch:

```
switch# show npv traffic-map
NPV Traffic Map Information:
----------------------------------------
Server-If External-If(s)
----------------------------------------
fc1/3 fc1/10,fc1/11
fc1/5 fc1/1,fc1/2
----------------------------------------
```

Use the following **show** command to display the NPV internal traffic details on an NPV switch:

```
switch# show npv internal info traffic-map
NPV Traffic Map Information:
----------------------------------------
Server-If External-If(s)
----------------------------------------
fc1/3 fc1/10,fc1/11
```

```
fc1/5 fc1/1,fc1/2
----------------------------------------
```

Example for Implementing IOA with NPV

In this implementation example, an NPIV-capable server is the host directly connected to the NPV core (NPIV-enabled) switch which also acts as an IOA node. The host sends data to the target over IOA flows.

To enable NPIV on NPV core switch, perform this task:

Procedure

Step 1 sjc-sw2# **configure terminal**

Enters configuration mode.

Step 2 sjc-sw1(config)# **feature npiv**

Enables NPIV mode on a NPV core switch.

To enable IOA on all the IOA nodes and to bring up the IOA flows, perform this task:

Procedure

Step 1 sjc-sw2# **configure terminal**

Enters configuration mode.

Step 2 sjc-sw1(config)# **feature npiv**

Enables NPIV mode on a NPV nodes switches.

Verifying the IOA Configuration

Use the following **show** command to verify the IOA configuration:

```
sjc-sw1# show npiv status
NPIV is enabled
```

Applying ASCII Configuration back to IOA Switches

To write erase, reload and then restore two switches that have an ioa cluster configured, perform this task:

Procedure

Step 1 switch# **configure terminal**

Enters configuration mode.

Step 2 switch(config)# **copy running-config bootflash:** *ioa_master*

Saves the running-config locally on the switch.

Step 3 Switch(config)# **boot system bootflash:** *m9250-s5ek9-mz.6.2.5b2.bin*

Sets the boot variable to a different kickstart image.

Step 4 Switch(config)# **boot system bootflash:** *m9250-s5ek9-mz.6.2.5b2.bin*

Sets the boot variable to a different system image.

Step 5 Switch(config)# **copy running-config startup-config**

Saves the boot variables that were configured in step 3 and step 4 to the startup configuration.

Step 6 Switch(config)# **write erase**

Will erase the startup configuration, but the boot variables and management IP address will not be erased.

Step 7 Switch(config)# **reload**

(Do steps 1-7 on slave IOA switches as well)

Reload the switches with the new kickstart and system image.

Warning
Once the switch has come up, you will need to configure the admin account password. Once you have done that, you will be asked if you would like to enter the basic configuration dialog. You can choose no and log onto the switch with the admin account password that you just configured. This must be done from the console port.

Step 8
a) **configure terminal**
b) **feature cluster**
c) **feature ioa**
d) **ioa site-local** *site name*
e) **interface ioa** *x/y*
f) **no shut**

Enables the cluster and ioa features on the switch. Configure the ioa site name and provision an ioa interface. Perform these commands on both master and slave switches.

Step 9
a) **configure terminal**
b) **ioa cluster** *name*
c) **node local**
d) **interface ioa** *x/y*
e) **node** *remote IP*
f) **interface ioa** *x/y*

Create a cluster and add the local and remote node to the cluster by executing the following commands on the master swich.

Note
These commands only need to be executed on the master switch. Do not run these commands on the slave switch. The interface specified in step d must exist on the master and the interface specified in step f must exist on the slave.

Step 10 **Switch# Copy bootflash: ioa_master running-config [echo-commands]**

Copy the ASCII configuration to the running configuration on the master switch.

echo-commands will echo the command before applying the config.

Step 11 **Switch# copy bootflash: ioa_slave running-config [echo-commands]**

Copy the ASCII configuration to the running configuration on the slave switch.

echo-commands will echo the command before applying the config.

Note
After copying the ASCII config, the boot variables will not be pointing to the images that were specified in steps 3 and 4. The boot variables will be pointing to whatever the saved boot variables are in the saved startup-config file. User should repeat steps 3 and 4 and do a **copy running-config startup-config**.

Applying the ascii configuration back to the switches will result in some errors on the master and slave switch but they can be ignored since the IOA cluster has already been set up on both switches.

Additional Configurations

Shutting Down a Cluster

To shut down a cluster, perform this task:

Procedure

Step 1 sjc-sw2# **configure terminal**

Enters configuration mode.

Step 2 sjc-sw2(config)# **ioa cluster tape_vault**

Specifies the cluster name and enters IOA cluster configuration submode. A cluster name can include a maximum of 31 alphabetical characters.

Step 3 sjc-sw2(config-ioa-cl)# **shut**

Shuts down the cluster. This command must be used to recover a cluster when it is partitioned. The change can be disruptive. For more information, see Cluster Recovery Scenarios, on page 121.

Load Balancing the Flows

To load balance the flows, perform this task:

Procedure

Step 1 sjc-sw2# **configure terminal**

Enters configuration mode.

Step 2 sjc-sw2(config)# **ioa cluster tape_vault**

Enters the cluster configuration mode.

Step 3 sjc-sw2(config-ioa-cl)# **load-balancing**

Load balances all the IOA flows. This process is disruptive and causes the hosts to relogin to targets. This command will take some time to execute depending on the number of flows. You should not terminate the command in the middle of its execution.

Step 4 sjc-sw2(config-ioa-cl)# **load-balancing enable**

This command turns on the load-balancing attribute for the new flows. You may enter this command only when you terminate the load-balancing command process.

Step 5 sjc-sw2(config-ioa-cl)# **load-balancing 11:22:33:44:55:66:77:88**

Load balances specified targets in the IOA flows. This process is disruptive and causes the hosts to relogin to targets. The **load-balancing** command will take some time to execute depending on the number of flows. You should not terminate the command in the middle of its execution.

Setting the Tunable Parameters

To set the following tunable parameters based on your deployment requirements, perform this task:

Command	Purpose
`sjc-sw2(config-ioa-cl)# tune round-trip-time ms`	Specifies the round-trip time in milliseconds. It is the time taken by the IOA data packet to traverse between two sites. The value can vary from 1 to 100 ms. 15 ms is the default.
`sjc-sw2(config-ioa-cl)# tune lrtp-retx-timeout msec`	Specifies the LRTP retransmit timeout in milliseconds. It is the time to wait before LRTP starts retransmitting packets. The value can vary from 500 to 5000 msec. 2500 msec is the default.

Caution The following are advanced tunable parameters, and you must consult the Cisco Services and Support team before tuning these parameters.

To set the the following advanced tunable parameters based on your deployment requirements, perform this task:

Procedure

Step 1 sjc-sw2# **configure terminal**

Enters configuration mode.

Step 2 sjc-sw2(config)# **ioa cluster tape_vault**

Enters the cluster configuration mode.

Step 3 sjc-sw2(config-ioa-cl)# **tune timer rscn-suppression** *seconds*

Specifies the IOA RSCN suppression timer value. It is the amount of time the IOA process waits before it queries FCNS (name server) after learning about changes in the network. This helps to alleviate the number of duplicate or repeating queries in case of rapid network changes. The value can vary from 1 to 10 seconds. 5 seconds is the default.

Step 4 sjc-sw2(config-ioa-cl)# **tune timer load-balance target** *seconds*

Specifies a IOA target load-balance timer value. It is the amount of time the IOA process waits before it attempts to load balance all IT Nexuses of a certain target port after a change in connectivity has been detected. The value can vary from 2 to 30 seconds. 2 seconds is the default.

Step 5 sjc-sw2(config-ioa-cl)# **tune timer load-balance global** *seconds*

Specifies a global IOA load-balance timer value. It is the amount of time the IOA process waits before it attempts to load balance all IT Nexuses configured in a cluster after a change in connectivity has been detected. The value can vary from 5 to 30 seconds. 5 seconds is the default.

Step 6 sjc-sw2(config-ioa-cl)# **tune ta-buffer-size** *KB*

Specifies the tape acceleration buffer size in KB. It is the amount of buffering allowed for flow control during tape acceleration. The value can vary from 64 to 12288 KB or Auto. Auto is the default. Auto option takes WAN latencies and speed of the tape device into account to provide optimum performance.

Step 7 sjc-sw2(config-ioa-cl)# **tune wa-buffer-size** *MB*

Specifies the write acceleration buffer size in MB. It is the amount of buffering allowed for flow control during write acceleration. The value can vary from 50 to 100 MB. 70 MB is the default.

Step 8 sjc-sw2(config-ioa-cl)# **tune wa-max-table-size** *KB*

Specifies the Write Max Table size in KB. It is the maximum number of active exchanges supported on an IOA flow. The value can vary from 4 to 64 KB. 4 KB is the default.

Changing the Node Description and IP Address of an IOA Cluster

To perform any of the following tasks, follow the steps defined in the Changing the Node Description and IP Address of an IOA Cluster, on page 105:

- Change the node-description (IP address) and node IP-address of a cluster.
- Change node-description(DNS name) of a cluster.
- Change the node-description from IP address to DNS name and vice versa.

To change the node description and IP address of an IOA node in the existing IOA cluster:

Procedure

Step 1 Shut down the IOA cluster on the switch1.

Step 2 Shut down the IOA cluster on the switch2.

Step 3 Remove the IOA cluster on the switch2.

Step 4 Remove the node of switch2 in the switch1.

Step 5 Do one of the following tasks based on what you want to perfom on the switch:

- Change the management interface IP address.
- Change the IP address and the switch name.
- Enable or disable the DNS configuration.

Step 6 Change the node description using the node id id node-description ip-address ip address command on switch1.

This step may vary depending on when the node description (DNS name) needs to be changed or node description and node IP address to be changed.

Step 7 Shut down the IOA cluster on the switch1.

Step 8 Add switch2 node with new description and the IP address.

Step 9 Add IOA interfaces on switch2.

Configuration Example for Changing the Node Description and Node IP Address of an IOA Cluster

This example shows the following configuration procedures used to change the description and IP address:

Shut Down the IOA Cluster on switch1

To shut down the IOA cluster on switch1 enter this command:

```
sw-231-19(config)# show ioa cluster  c1 node summary
--------------------------------------------------------------------------------
Switch             Site            Status           Master     Node ID
--------------------------------------------------------------------------------
172.25.231.14      site3           online           no         2
172.25.231.19(L)   site2           online           yes        1
sw-231-19(config)# ioa cluster  c1
sw-231-19(config-ioa-cl)# sh
This change can be disruptive. Please ensure you have read the "IOA Cluster Recovery
Procedure" in the configuration guide. -- Are you sure you want to continue? (y/n) [n] y
2011 Apr 12 07:02:21 sw-231-19 %CLUSTER-2-CLUSTER_LOCAL_NODE_EXIT: Local Node 0x1 has left
 the Cluster 0x5000530019f08076
```

Shut Down the IOA Cluster on switch2

To shut down the IOA cluster on switch2 enter this command:

```
sw-231-14(config)# ioa cluster  c1
sw-231-14(config-ioa-cl)# sh
This change can be disruptive. Please ensure you have read the "IOA Cluster Recovery
Procedure" in the configuration guide. -- Are you sure you want to continue? (y/n) [n] y
2011 Apr 12 07:02:30 sw-231-14 %CLUSTER-2-CLUSTER_LOCAL_NODE_EXIT: Local Node 0x2 has left
 the Cluster 0x5000530019f08076
sw-231-14(config-ioa-cl)# sh ioa cluster  c1 node sum
--------------------------------------------------------------------------------
Switch             Site            Status           Master     Node ID
--------------------------------------------------------------------------------
192.125.231.14(L)   --              unknown (cluster is offline)            2
192.125.231.19      --              unknown (cluster is offline)            1
```

Remove the IOA Cluster on switch2

To remove the IOA cluster on switch2, enter this command:

```
sw-231-14(config-ioa-cl)# no ioa cluster  c1
sw-231-14(config)#
```

Remove the Node of switch2 in switch1

To remove the node of switch2 in switch1, enter this command:

```
sw-231-19(config-ioa-cl)# no node 192.125.231.14
sw-231-19(config-ioa-cl)# sh ioa cluster  c1 node sum
--------------------------------------------------------------------------------
Switch             Site            Status           Master     Node ID
--------------------------------------------------------------------------------
192.125.231.19(L)   --              unknown (cluster is offline)            1
sw-231-19(config-ioa-cl)#
```

Change the Management Interface IP Address on Switches

```
sw-231-19(config)# int mgmt0
sw-231-19(config-if)# ip address 192.125.231.72 255.255.255.0
```

Change the Node Description and IP Address on switch1

To change the node description and IP address on switch1, enter this command

node id *id new-description* **ip-address** *new-ip address*

```
sw-231-19(config-ioa-cl)# node id 1 192.125.231.72 ip-address 192.125.231.72
```

Bring Up IOA Cluster on switch1

To bring up the IOA cluster on a switch, enter this command:

```
sw-231-19(config-ioa-cl-node)# no sh
This change can be disruptive. Please ensure you have read the "IOA Cluster Recovery
Procedure" in the configuration guide. -- Are you sure you want to continue? (y/n) [n] y
```

```
sw-231-19(config-ioa-cl)# 2011 Apr 12 07:04:54 sw-231-19 %CLUSTER-2-CLUSTER_LEADER_ANNOUNCE:
 Node 0x1 is the new Master of cluster 0x5000530019f08076 of 1 nodes
2011 Apr 12 07:04:54 sw-231-19 %CLUSTER-2-CLUSTER_QUORUM_GAIN: Cluster 0x5000530019f08076
now has quorum with 1 nodes
sw-231-19(config-ioa-cl)# show ioa cluster  c1 node summary
--------------------------------------------------------------------------------
Switch             Site            Status           Master     Node ID
--------------------------------------------------------------------------------
192.125.231.72(L)   site2           online           yes        1
```

Add switch2 Node with New Description and the IP Address

To add switch2 node with a new description and IP address, enter this command:

```
sw-231-19(config-ioa-cl)# node 172.25.231.25
2011 Apr 12 07:05:30 sw-231-19 %CLUSTER-2-CLUSTER_QUORUM_GAIN: Cluster 0x5000530019f08076
now has quorum with 1 nodes
2011 Apr 12 07:05:30 sw-231-19 %CLUSTER-2-CLUSTER_QUORUM_GAIN: Cluster 0x5000530019f08076
now has quorum with 2 nodes
```

Add IOA Interfaces on switch1

To add IOA interfaces on the switch, enter this command:

```
sw-231-19(config-ioa-cl-node)# int ioa 1/1
sw-231-19(config-ioa-cl-node)# int ioa 1/2
sw-231-19(config-ioa-cl-node)#
```

Verify the Node Description and IP Address and Flows

To confirm the functioning of the cluster with the new IP address, use the following show commands:

```
sw-231-19(config)# show ioa cluster  c1 node summary
--------------------------------------------------------------------------------
Switch             Site            Status           Master     Node ID
--------------------------------------------------------------------------------
172.25.231.25      site3           online           no         2
192.125.231.72(L)   site2           online           yes        1
```

```
sw-231-19(config)# show ioa cluster  c1 int summary
--------------------------------------------------------------------------------
Switch             Interface         Status       Flows
--------------------------------------------------------------------------------
172.25.231.25      ioa1/1            up           20
172.25.231.25      ioa1/2            up           16
192.125.231.72(L)   ioa4/1            up           20
192.125.231.72(L)   ioa4/2            up           16
```

```
sw-231-19(config)# show ioa cluster  c1 node
Node 172.25.231.25 is remote switch
  Node ID is 2
  IP address is 172.25.231.25
  Status is online
  Belongs to Site site3
  Node is not master switch
Node 192.125.231.72 is local switch
  Node ID is 1
  IP address is 192.125.231.72
```

```
  Status is online
  Belongs to Site site2
  Node is the master switch
sw-231-19(config)#
```

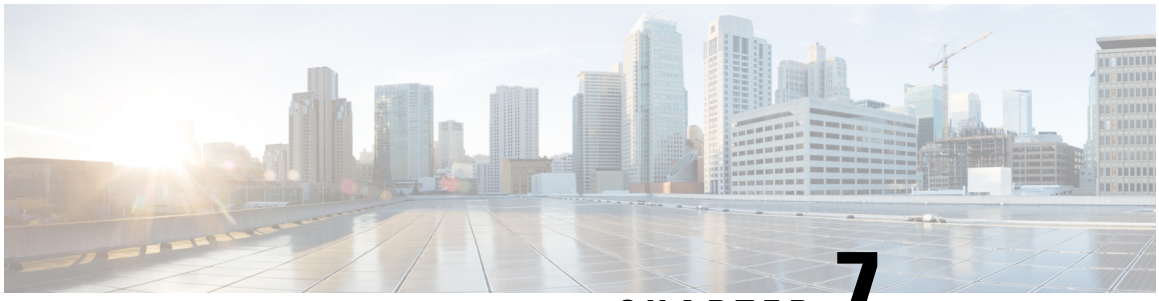
Displaying Interface Statistics

The following examples display interface statistics:

```
sw231-19# show int ioa 1/1 counters
ioa1/1
    4454232796 device packets in, 375748229 device packets out
    8948409208760 device bytes in, 24047886946 device bytes out
    526563297 peer packets in, 2471396408 peer packets out
    45198770258 peer bytes in, 4697995629324 peer bytes out
    8 i-t create request, 4 i-t create destroy
    8 i-t activate request, 0 i-t deactivate request
sw231-19# show int ioa 1/1 counters  brief
-------------------------------------------------------------------------------
Interface          To Device  (rate is 5 min avg) To Peer  (rate is 5 min avg)
                   -------------------------  -------------------------
                   Rate     Total                Rate     Total
                   MB/s     Bytes                MB/s     Bytes
-------------------------------------------------------------------------------
ioa1/1             0.56     24049257618          109.66   4698262901274
sw231-19# show ioa int int ioa 1/1 summary
---- ---------------------- ---- ------------- ---- ---
FLOW HOST                    VSAN STATUS        COMP ACC
     TARGET
---- ---------------------- ---- ------------- ---- ---
1    10:00:00:00:00:00:03:00 200  ACTIVE        YES WA
     11:00:00:00:00:00:03:00
2    10:00:00:00:00:00:02:00 200  ACTIVE        NO  WA
     11:00:00:00:00:00:02:00
3    10:00:00:00:00:00:01:00 100  ACTIVE        YES TA
     11:00:00:00:00:00:01:00
4    10:00:00:00:00:00:00:00 100  ACTIVE        NO  TA
     11:00:00:00:00:00:00:00
sw231-19# show ioa int int ioa 1/1 stats
  Adapter Layer Stats
    4457312829 device packets in, 376008035 device packets out
    8954596919462 device bytes in, 24064514554 device bytes out
    526927441 peer packets in, 2473105321 peer packets out
    45230025550 peer bytes in, 4701244024682 peer bytes out
    8 i-t create request, 4 i-t create destroy
    8 i-t activate request, 0 i-t deactivate request
    0 i-t create error, 0 i-t destroy error
    0 i-t activate error, 0 i-t deactivate error
    48 i-t-n  not found, 0 i-t-n stale logo timer expiry
    4 logo sent, 8 logo timer started
    4 logo timer fired, 4 logo timer cancelled
    4 plogi 4 plogi-acc 4 logo-acc 4 prli 4 prli-acc 0 els-q-err
    to-device 214279940 orig pkts 12743547488 orig bytes
    to-peer 8748538 orig pkts 682386268 orig bytes
    0 queued 0 flushed 0 discarded
  LRTP Stats
    0 retransmitted pkts, 0 flow control
    2464072014 app sent 2464072014 frags sent 0 tx wait
    0 rexmt bulk attempts  0 rexmt bulk pkts  2 delayed acks
    376008013 in-order 0 reass-order 0 reass-wait  0 dup-drop
    376008013 app deliver 376008013 frags rcvd
    150919428 pure acks rx  376008013 data pkts rx 0 old data pkts
    0 remove reass node, 0 cleanup reass table
```

```
    Tape Accelerator statistics
       2 Host Tape Sessions
       0 Target Tape Sessions
      Host End statistics
       Received 26275926 writes, 26275920 good status, 2 bad status
       Sent 26275914 proxy status, 10 not proxied
       Estimated Write buffer 4 writes 524288 bytes
       Received 0 reads, 0 status
       Sent 0 cached reads
       Read buffer 0 reads, 0 bytes
      Host End error recovery statistics
       Sent REC 0, received 0 ACCs, 0 Rejects
       Sent ABTS 0, received 0 ACCs
       Received 0 RECs, sent 0 ACCs, 0 Rejects
       Received 0 SRRs, sent 0 ACCs, 0 Rejects
       Received 0 TMF commands
      Target End statistics
       Received 0 writes, 0 good status, 0 bad status
       Write Buffer 0 writes, 0 bytes
       Received 0 reads, 0 good status, 0 bad status
       Sent 0 reads, received 0 good status, 0 bad status
       Sent 0 rewinds, received 0 good status, 0 bad status
       Estimated Read buffer 0 reads, 0 bytes
       Target End error recovery statistics
       Sent REC 0, received 0 ACCs, 0 Rejects
       Sent SRR 0, received 0 ACCs
       Sent ABTS 0, received 0 ACCs
       Received 0 TMF commands
  Write Accelerator statistics
   Received 726357548 frames, Sent 529605035 frames
   0 frames dropped, 0 CRC errors
   0 rejected due to table full, 0 scsi busy
   0 ABTS sent, 0 ABTS received
   0 tunnel synchronization errors
   Host End statistics
     Received 188004026 writes, 188004000 XFER_RDY
     Sent 188004026 proxy XFER_RDY, 0 not proxied
     Estimated Write buffer 1146880 bytes
     Timed out 0 exchanges, 0 writes
   Target End statistics
     Received 0 writes, 0 XFER_RDY
     Write buffer 0 bytes
     TCP flow control 0 times, 0 bytes current
     Timed out 0 exchanges, 0 writes
  Compression Statistics
    Pre Comp Batch size 131072
    Post Comp Batch size 2048
    4375494911078 input bytes, 50140348947 output compressed bytes
    0 non-compressed bytes, 0 incompressible bytes
    0 compression errors
    0 Compression Ratio
  De-Compression Statistics
    0 input bytes, 0 output decompressed bytes
    11883488326 non-compressed bytes
    0 de-compression errors
sw231-19# show ioa int int ioa 1/1 init-pwwn 10:00:00:00:00:00:03:00 targ-pwwn
11:00:00:00:00:00:03:00 vsan 200 counters
  Adapter Layer Stats
    1366529601 device packets in, 160768174 device packets out
    2699458644986 device bytes in, 10289163140 device bytes out
    160844041 peer packets in, 165188790 peer packets out
    18652597246 peer bytes in, 47736122724 peer bytes out
    0 i-t create request, 0 i-t create destroy
    0 i-t activate request, 0 i-t deactivate request
```

```
    0 i-t create error, 0 i-t destroy error
    0 i-t activate error, 0 i-t deactivate error
    0 i-t-n  not found, 0 i-t-n stale logo timer expiry
    1 logo sent, 2 logo timer started
    1 logo timer fired, 1 logo timer cancelled
    1 plogi 1 plogi-acc 1 logo-acc 1 prli 1 prli-acc 0 els-q-err
    to-device 80384094 orig pkts 4662277452 orig bytes
    to-peer 0 orig pkts 0 orig bytes
    0 queued 0 flushed 0 discarded
  LRTP Stats
    0 retransmitted pkts, 0 flow control
    160768190 app sent 160768190 frags sent 0 tx wait
    0 rexmt bulk attempts  0 rexmt bulk pkts  1 delayed acks
    160768162 in-order 0 reass-order 0 reass-wait  0 dup-drop
    160768162 app deliver 160768162 frags rcvd
    75879 pure acks rx  160768162 data pkts rx 0 old data pkts
    0 remove reass node, 0 cleanup reass table
  Write Accelerator statistics
   Received 1607681842 frames, Sent 1527297774 frames
   0 frames dropped, 0 CRC errors
   0 rejected due to table full, 0 scsi busy
   0 ABTS sent, 0 ABTS received
   0 tunnel synchronization errors
   Host End statistics
     Received 80384094 writes, 80384082 XFER_RDY
     Sent 80384094 proxy XFER_RDY, 0 not proxied
     Estimated Write buffer 524288 bytes
     Timed out 0 exchanges, 0 writes
   Target End statistics
     Received 0 writes, 0 XFER_RDY
     Write buffer 0 bytes
     TCP flow control 0 times, 0 bytes current
     Timed out 0 exchanges, 0 writes
sw231-19# show ioa int int ioa 1/1 init-pwwn 10:00:00:00:00:00:03:00 targ-pwwn
11:00:00:00:00:00:03:00 vsan 200 counters brief
--------------------------------------------------------------------------------
Interface          Input (rate is 5 min avg)      Output (rate is 5 min avg)
                   ------------------------------  ------------------------------
                   Rate     Total                 Rate     Total
                   MB/s     Frames                MB/s     Frames
--------------------------------------------------------------------------------
ioa1/1
Device             60       9573683               0        1126308
Peer               0        1126833               1        1157161
sjc-sw2#
```



CHAPTER 7

SCSI Write Acceleration and Tape Acceleration

This appendix describes the concepts of SCSI write acceleration, tape acceleration, and compression.

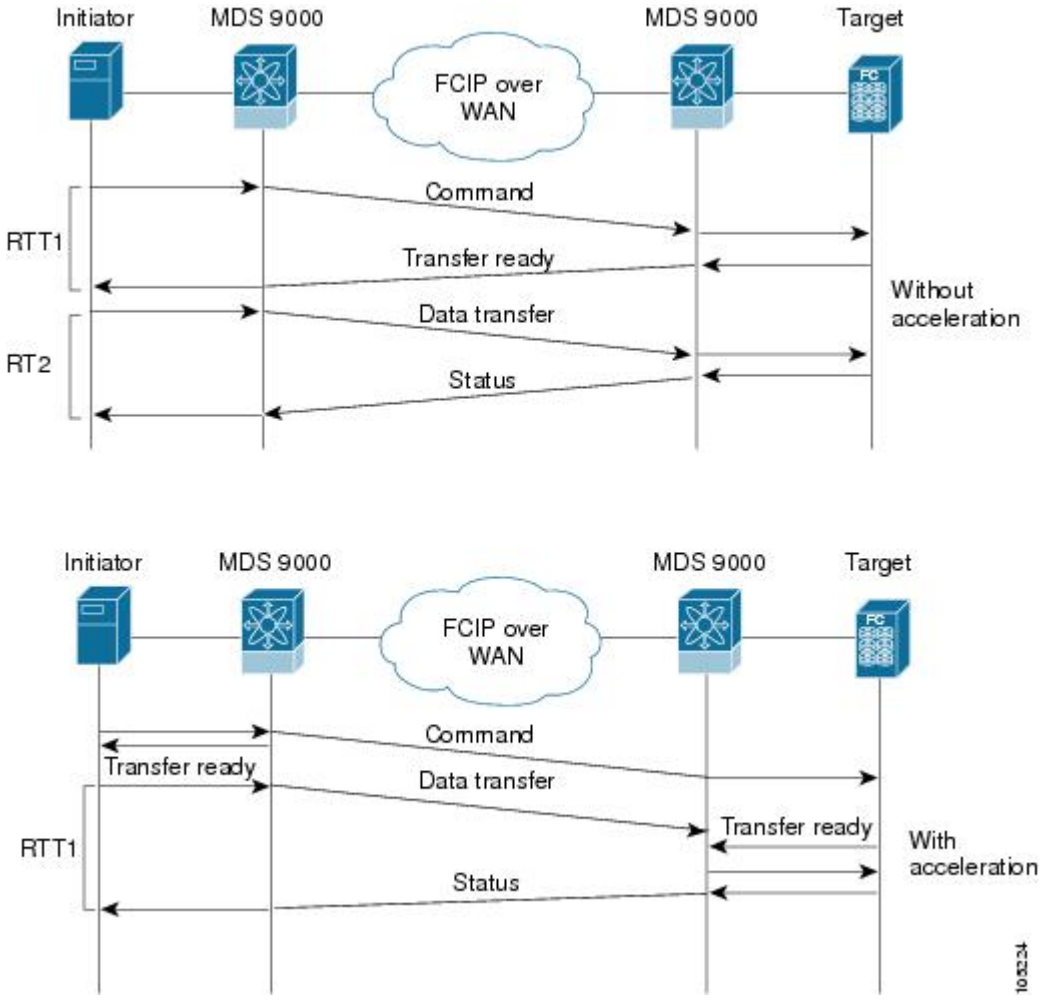
- SCSI Write Acceleration, on page 113
- SCSI Tape Acceleration, on page 115

SCSI Write Acceleration

The SCSI write acceleration feature enables you to significantly improve application write performance when storage traffic is routed over wide area networks using FCIP or Fibre Channel. When write acceleration is enabled, WAN throughput is maximized by minimizing the impact of WAN latency for write operations.

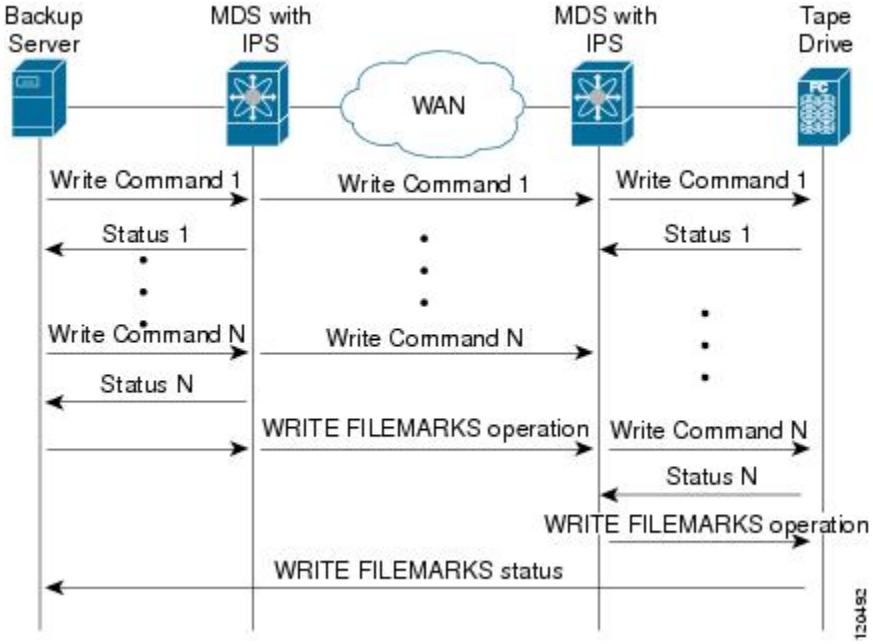
In Figure 36: SCSI Write Acceleration, on page 114, the WRITE command without write acceleration requires two round-trip transfers (RTT), while the WRITE command with write acceleration only requires one RTT. The maximum sized Transfer Ready is sent from the host side of the FCIP or Fibre Channel link back to the host before the WRITE command reaches the target. This enables the host to start sending the write data without waiting for the long latency over the FCIP or Fibre Channel link of the WRITE command and Transfer Ready. It also eliminates the delay caused by multiple Transfer Readys needed for the exchange going over the FCIP or Fibre Channel link.

Figure 36: SCSI Write Acceleration



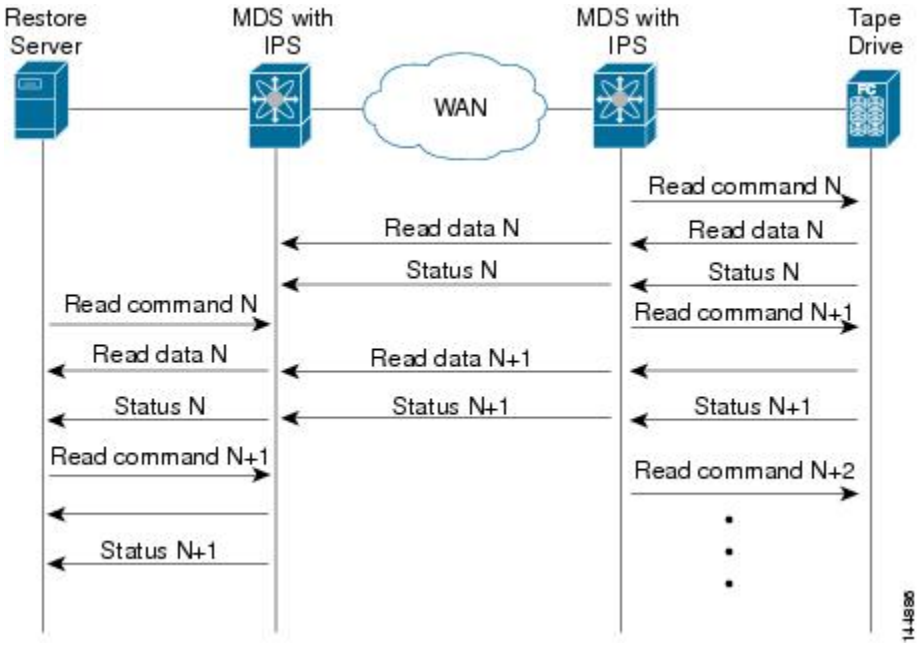
SCSI Tape Acceleration

Figure 37: SCSI Tape Acceleration for Write Operations

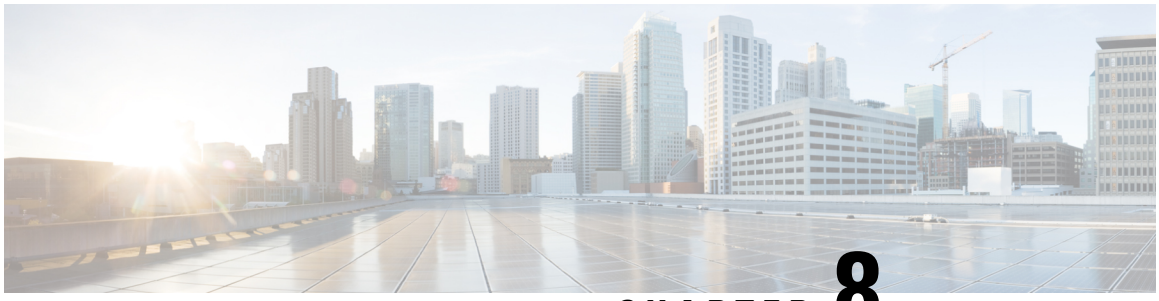


At the tape end of the FCIP or Fibre Channel tunnel, another Cisco MDS switch buffers the command and data it has received. It then acts as a backup server to the tape drive by listening to a transfer ready from the tape drive before forwarding the data.

Figure 38: SCSI Tape Acceleration for Read Operations



The Cisco NX-OS provides reliable data delivery across the WAN. While tape media errors during the read operation are returned to the restore server for error handling, the Cisco NX-OS software recovers from any other errors.



CHAPTER 8

Cluster Management and Recovery Scenarios

This appendix includes information on cluster management and recovery procedures that are used when one or more switches in a Cisco IOA cluster is offline or when you want to change the master switch assignment from one switch to another switch.

Note

- The procedures in this appendix describe troubleshooting solutions that use the CLI.
- The Cisco IOA cluster configuration for an offline switch must be done using the CLI. Cisco IOA cluster configuration for an online switch can be done using Cisco DCNM-SAN or the CLI.

This appendix includes the following sections:

- Cluster Quorum and Master Switch Election, on page 117
- In-Service Software Upgrade (ISSU) in a Two-Node Cluster, on page 120
- Supported Topologies, on page 121
- Cluster Recovery Scenarios, on page 121

Cluster Quorum and Master Switch Election

This section describes the Cisco IOA cluster quorum and the process for electing the master switch in a cluster.

Node ID

Every switch in a cluster has a node ID. Cisco IOA assigns a node ID to every new switch as it is added to the cluster. The switch where the cluster is created is assigned the node ID of 1. This is the master switch. When a new switch is added to the cluster, it is assigned the next available higher node ID. For example, when a second switch is added to the cluster it gets the node ID of 2 and the third switch gets the node ID of 3, and so on.

Cluster View

The cluster view is the set of switches that are part of the operational cluster.

Cluster Quorum

For a cluster to be operational, it must include more than half the number of configured switches in the cluster view. In an N-node cluster, N/2 + 1 nodes form a cluster quorum.

If N is even, the cluster quorum requires N/2 nodes and also, the presence of the switch with the lowest node ID.

The quorum logic ensures that in the event of cluster partitions at least one partition can be operational. All other switches are nonoperational. This guarantees the consistency of the cluster.

Master Switch Election

When a cluster is created, the switch on which the cluster is created becomes the cluster master switch. When the master switch fails or is rebooted, another switch takes over as the master switch. The master election logic uses the node ID and the latest cluster configuration to determine which switch in the cluster will become the master switch. The master election logic is described as follows:

- If the master switch fails in an operational cluster, the switch with the next lowest node ID takes over as the master switch. Note that in an operational cluster, all the switches run the same cluster configuration.
    - When the previous master switch comes back online and joins the cluster, it does not immediately become the master.
- When all the switches of a cluster are coming up, the switch that has the latest cluster configuration becomes the master switch. If there are multiple switches with the same configuration, the switch with the lowest node ID is chosen to be the master switch.
    - Once a master switch is chosen and the cluster is operational (there is a quorum), even if a switch with a lower node ID joins the cluster at a later time, the master switch does not change.

For example, there are three switches S1, S2, and S3 with node IDs 1, 2, and 3, respectively. If switches S2 and S3 form a quorum then switch S2 becomes the master switch. Even if switch S1 with the node ID of 1 comes up and joins the cluster at a later time, switch S2 continues to be the master. However, if switch S2 goes down for any reason, switch S1 will become the master switch.

Two-Switch Cluster Scenarios

According to the cluster quorum logic, a cluster with two configured switches can be operational if both switches are operational or the switch with the lowest node ID is operational.

In the latter case, the switch with the lowest node ID is the master of the one-switch cluster. The other switch could have failed or simply lost connectivity to the operational switch. In either case, the switch with the higher node ID would become nonoperational. If the node with the lower node ID failed, the other switch cannot form an operational cluster.

The examples that follow describe these scenarios. The first three examples consider single switch failures.

1. Assume that in a two-switch cluster with switches S1 (node ID 1) and S2 (node ID 2), S1 is the master (the master has the lower node ID).

   When the switches lose connectivity between them, the master switch S1 continues to be operational since it has the lower node ID and can form an (N/2) switch cluster. Switch S2 becomes nonoperational.

2. Assume that in a two-switch cluster with switches S1 (node ID 1) and S2 (node ID 2), S2 is the master (note that the master has the higher node ID because it has the latest configuration when both the switches came online).

   When the switches lose connectivity between them, switch S2 becomes nonoperational and S1 takes over as the master to form a 1-switch cluster. This is consistent with the quorum logic in a two-switch cluster (N/2 with lowest node ID).

3. Assume that in a two-switch cluster with switches S1 (node ID 1) and S2 (node ID 2). If S1 fails (regardless of which switch was the master), S2 will also become non-operational as long as S1 is down.

   When S1 comes up, S1 and S2 will form a two-switch cluster.

The next set of examples describe reboots of both switches (S1 with node ID 1 and S2 with node ID 2):



Caution If you perform any configuration change on a cluster, you must save the running configuration to the startup configuration by entering the **copy running-config startup-config** CLI command on all switches before rebooting them. Otherwise, the cluster may not form correctly after the reboot (see example Example 3).

1. After a reboot, if both switches S1 and S2 come up about the same time, a two-switch cluster will be formed.
   a. If the cluster configurations are the same, S1 (with the lower node ID) will become the master.
   b. If the cluster configurations are different, the switch with the latest cluster configuration will become the master.
2. After a reboot, if switch S2 comes up first, it will not be able to form a cluster until S1 also comes up. After that, the algorithm explained in the previous case will be used.
3. After a reboot, if switch S1 comes up first, it will form a one-switch cluster (N/2 with lowest node ID). When S2 comes up, it will join the cluster to form a two-switch cluster.

When S2 comes up and if it happens to have the latest cluster configuration in the startup configuration (this can happen if you did not save the running configuration to the startup configuration on S1 but did so on S2), it will not be able to join the cluster formed by S1.



Caution It is critical that you save the running configuration on all switches before a reboot.

Three-Switch Cluster Scenarios

In a three-switch cluster, the quorum requires two switches to be in the cluster view (N/2 + 1). The examples below explain three scenarios in a three-switch cluster with switches S1 (node ID 1), S2 (node ID 2) and S3 (node ID 3). S1 is the master switch.

1. In a three-switch operational cluster, if switch S3 fails or loses connectivity with the other two switches, then S3 becomes nonoperational. Switches S1 and S2 will form an operational cluster. When S3 comes up again, it will rejoin the cluster.
2. In a three-switch operational cluster, if the master switch S1 fails or loses connectivity with the other two switches, then S1 becomes nonoperational. Switches S2 and S3 will form an operational cluster and S2 will be the master. When S1 comes up again, it will rejoin the cluster. Note that S2 will continue to be the master.
3. If two switches fail, the cluster will become nonoperational.

These examples describe reboots on all switches in the cluster:

Caution If you perform any configuration change on a cluster, you must save the running configuration to the startup configuration by entering the **copy running-config startup-config** command on all switches before rebooting them. Otherwise, the cluster may not form correctly after the reboot.

1. After a reboot, if all switches come up at about the same time, first a 2-switch cluster will be formed and later the third switch will be added.
   a. If the cluster configurations are the same, S1 (with the lower node ID) will become the master switch and form the 2-switch cluster first; and then add the third switch.
   b. If the cluster configurations are different, the switch that is running the latest configuration will become the master switch and then form a 2-switch cluster; and then add the third switch.
2. After a reboot, if the switches come up one at a time, a 2-switch cluster will be formed after the first two switches are up. Later, when the third switch comes online, it will join the cluster.

   If the third switch happens to be running the latest cluster configuration in the startup configuration (this can happen if you save the running configuration only on this switch but not on the other two), the third switch will not be able to join the cluster.

Caution It is critical that you save the running configuration on all switches before a reboot.

Four-Switch Cluster Scenarios

The four-switch cluster scenario is very similar to the examples above. The cluster will be operational if the cluster view has at least three switches (N/2 + 1), or if the cluster view has two switches including the switch with the lowest node ID (N/2 with lowest node ID).

In-Service Software Upgrade (ISSU) in a Two-Node Cluster

In-Service Software Upgrade (ISSU) is a comprehensive, transparent software upgrade application that allows you to deploy bug fixes and add new features and services without any disruption to the traffic.

In a cluster comprising of the MDS 9222i Switches as nodes, if the nodes are not able to communicate, then the node having the lowest node identifier (node ID) remains in the cluster while the other node leaves the cluster. However, when an ISSU is performed on a node having the lowest node identifier, a complete loss of the cluster results because both the nodes leave the cluster.

This undesirable situation is addressed in a two-node cluster as follows:

- The upgrading node sends a message to the other node of the intent to leave the cluster. The upgrading node can either be a master node or a slave node.
- The remaining node remains in the cluster and performs the role of the master node if it was a slave node. This node continues to remain in the cluster with the quorum intact.
- After the ISSU is completed and the switches boot up, the upgraded node rejoins the cluster as a slave node.

Note This feature is tied to ISSU logic and no additional commands need to be executed.

Supported Topologies

Cisco IOA supports a single-fabric topology. Multiple modules can be deployed in a Fibre Channel fabric to easily scale-up performance, to enable simplified load balancing, and to increase availability. In a typical configuration, one IOA engine per site is required in each IOA cluster.

IOA clusters include designated backup servers, tape libraries, and one or more MDS switches running Cisco SAN-OS Release 3.2(2c) or alter. One cluster switch must include an IOA engine per site. With easy-to-use provisioning, traffic between any host and tape on the fabric can utilize the IOA services.

Required Cisco IOA engines are included in the following Cisco products:

- Cisco MDS 9000 Family 18/4-port Multiservice Module (MSM-18/4)
- Cisco SSN-16 Module Switch

Single-Fabric Topology

The MSM-18/4 Module can be anywhere in the fabric. Cisco IOA does a one-to-one mapping of the information from the host to the target and forwards the encrypted data to the dedicated HR tape. Cisco IOA also tracks the barcodes on each encrypted tape and associates the barcodes with the host servers.

Encryption and compression services are transparent to the hosts and storage devices. These services are available for devices in any virtual SANs (VSANs) in a physical fabric and can be used without rezoning.

In certain topologies, edge switches are interconnected across the WAN. Plan for deployment at the core and transition of WAN links to core switches for optimal routing.

Cluster Recovery Scenarios

Refer to this section for information on recovery procedures that are used when one or more switches in a Cisco IOA cluster is offline or when you want to change the master switch assignment from one switch to another switch.

This section includes the following topics:

Deleting an Offline Switch from a Cisco IOA Cluster

To delete an offline switch when one or more switches are offline and the master switch is online, use these procedures.

Procedure

Step 1 switch# **configure terminal**

Enters configuration mode.

Step 2 switch(config)# **ioa cluster** *ABC*

Enter the IOA Cluster mode.

Step 3 switch(config-ioa-cl)# **shutdown**

Shuts down the ABC cluster on the offline switch.

Note
Repeat the procedure for every offline switch.

Step 4 switch(config-ioa-cl)# **no node switch 2**

Deletes switch2 from the ABC cluster configuration.

Note
Repeat this step for every offline switch that was shut down in Step 3.

Step 5 switch(config-ioa-cl)# **exit**

Exits the IOA Cluster mode and enters the Global configuration mode.

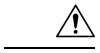
Step 6 switch(config)# **no ioa cluster** *ABC*

Deletes the ABC cluster configuration.

Note
Delete the cluster on every offline switch that was shut down in the first procedure.

Deleting a Cisco IOA Cluster with One or More Offline Switches while the Master Switch is Online

To delete a Cisco IOA cluster that includes one or more offline switches and online master switch, use these procedures.



Caution Do not remove a cluster master switch from a cluster and then try to revive the cluster on an offline switch. Since the offline switch was not part of the operational cluster, the cluster master may have progressed beyond what is in the offline switch's state. Deleting the cluster master and reviving the cluster on an offline switch can result in stale configuration.

Procedure

Step 1 switch# **configure terminal**

Enters configuration mode.

Step 2 switch(config)# **ioa cluster** *ABC*

Enter the IOA Cluster mode.

Step 3 switch(config-ioa-cl)# **shutdown**

Shuts down the ABC cluster on the offline switch.

Note
Repeat the procedure for every offline switch.

Step 4 switch(config-ioa-cl)# **no node switch2**

Deletes switch2 from the ABC cluster configuration.

Note
Repeat this step for every offline switch that was shut down in Step 3.

Step 5 switch(config-ioa-cl)# **exit**

Exits the IOA Cluster mode and enters the Global configuration mode.

Step 6 switch(config)# **no ioa cluster** *ABC*

Deletes the ABC cluster configuration.

Note
Delete the cluster on every offline switch that was shut down in the first procedure.

Deleting a Cisco IOA Cluster when All Switches Are Offline

To delete a Cisco IOA cluster when the master switch and all other switches are offline, use these procedures.



Note When all switches are offline, the cluster is offline.

On the offline switch (for example, switch2), shut down the cluster by performing this task:

Procedure

Step 1 switch# **configure terminal**

Enters configuration mode.

Step 2 switch(config)# **ioa cluster** *ABC*

Enter the IOA Cluster mode.

Step 3 switch(config-ioa-cl)# **shutdown**

Shuts down the ABC cluster on the offline switch.

Note
Repeat the procedure for every offline switch.

Step 4 switch(config-ioa-cl)# **exit**

Exits the IOA Cluster mode and enters the Global configuration mode.

Step 5 switch(config)# **no ioa cluster** *ABC*

Deletes the ABC cluster configuration.

Note
Delete the cluster on every offline switch that was shut down in the first procedure.

Reviving a Cisco IOA Cluster

To revive a cluster on the switch that has the latest Cisco IOA configuration version, use these procedures.

This procedure is used to revive a cluster when one or more switches are offline and the cluster is nonoperational (for example, due to a quorum loss). The recovery procedure includes deleting one or more offline switches and then reviving the cluster on the remaining switches.

Caution A Cisco IOA cluster must only be revived on the switch with the latest IOA configuration version as displayed by the **show IOA cluster detail** command. Reviving the cluster on a switch that does not have the highest configuration version can result in stale configuration.

Note The following procedure assumes that switch1 has the latest IOA configuration version. The steps shown for switch2 should be carried out for every switch that needs to be removed before reviving the cluster.

Procedure

Step 1 switch# **configure terminal**

Enters configuration mode.

Step 2 switch(config)# **ioa cluster** *ABC*

Enter the IOA Cluster mode.

Step 3 switch(config-ioa-cl)# **shutdown**

Shuts down the ABC cluster on the offline switch.

Step 4 switch(config-ioa-cl)# **exit**

Exits the IOA Cluster mode and enters the Global configuration mode.

Step 5 switch(config)# **no ioa cluster** *ABC*

Deletes the ABC cluster configuration.

Note

Delete the cluster on every offline switch that was shut down in the first procedure.

Step 6 switch(config)# **ioa cluster** *ABC*

Enter the IOA Cluster mode.

Step 7 switch(config-ioa-cl)# **no node switch 2**

Deletes switch2 from the ABC cluster configuration.

Note
Repeat for every switch that needs to be deleted.

Step 8 switch(config-ioa-cl)# **no shutdown**

Restarts the cluster on the switch.