



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, page 1](#)
- [Deployment Guidelines, page 6](#)
- [Guidelines and Limitations, page 7](#)
- [Configuration Limits, page 10](#)

Supported Topologies

This section includes the following topics:


Core-Edge Topology

[Figure 1: Core-Edge Topology, on page 2](#) illustrates the core-edge topology where you are recommended to place the IOA interfaces (MSM-18/4 or SSN-16) 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 1: Core-Edge Topology



Edge-Core-Edge Topology

Figure 2: Edge-Core-Edge Topology, on page 2 illustrates the edge-core-edge topology where you are recommended to place the MSM-18/4 Module or SSN-16 Module at the core switches that interconnect the two sites.


Figure 2: Edge-Core-Edge Topology



Collapsed Core Topology

Figure 3: Collapsed Core Topology, on page 3 illustrates the collapsed core topology where you are recommended to place the MSM-18/4 Module or SSN-16 Module (IOA interfaces) in the core switches that interconnect the two sites.

Figure 3: Collapsed Core Topology




Extended Core-Edge Topology

Figure 4: Extended Core-Edge Topology, on page 4 illustrates the extended core-edge topology where you are recommended to place the IOA interfaces (MSM-18/4 Module or SSN-16 Module) 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 4: Extended Core-Edge Topology




Extending Across Multiple Sites

Figure 5: [Extended Across Multiple Sites](#), on page 5 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 5: 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 MSM-18/4 or SSN-16 module 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.

**Note**

IOA is supported for IVR flows starting from Cisco MDS NX-OS Release 5.0(1a).

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 and drastically reduces the scalability requirements on IOA.
- Deploy IOA interfaces (MSM-18/4 or SSN-16) 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 Generation 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 a 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](#).
 - 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 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.
- You can provision only one intelligent application on a single service engine. In SSN-16 there are 4 service engines and each service engine can host a single intelligent application.
- In Cisco NX-OS Release 4.2(1), only IOA and FCIP can run on the same SSN-16 as in the following example:
 - If one of the service engines runs SME on an SSN-16, you cannot configure another application to run the remaining service engines on this SSN-16.
 - If one of the service engines runs IOA or FCIP, you can configure other service engines to run either FCIP or IOA.
- IOA uses the image that is bundled as a part of the Cisco MDS NX-OS Release. In Cisco MDS NX-OS Release 4.2(1), SSI images are not supported for IOA.
- 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 from into the cluster before adding switches from the remote site.
- IOA clustering framework uses IP connectivity for its internal operation. In Cisco NX-OS Release 4.2(1) and later releases, 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.
- The targets must be connected to a FC-Redirect-capable switch running Cisco MDS NX-OS Release 4.2(1) or later. The hosts must be connected to a FC-Redirect-capable switch running Cisco MDS SAN-OS Release 3.3(1c) or later.
- In Cisco MDS NX-OS Release 4.2(1), the following features cannot coexist with IOA for a specific flow: SME, DMM, IVR, NPV and NPIV, F PortChannel or Trunk. In Cisco NX-OS Release 5.0(1), IVR is supported with IOA.
- To implement IOA on IVR flows, the host switches, target switches, border switches, and the IOA switches must all be running AAM-supported Cisco MDS NX-OS Release 5.0(1) or later. For more information, refer to the *Cisco MDS 9000 Family NX-OS Inter-VSAN Routing Configuration Guide*.
- If there are multiple Cisco 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 for these 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 `clear fc-redirect decommission-switch` 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 twinpeak line card 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 ELS entries on the twinpeak line card 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 using FCIP write acceleration instead of IOA.
- IOA flow takes a few seconds to become active upon certain triggers such as host or target port flaps. PLOGI from the hosts are buffered until the IOA flow becomes active. Once the IOA flow becomes active, it sends an 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 would be 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.

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



Note When using IOA:

- SSI images should not be loaded and installed on Cisco MDS 9250i switches and 18+4 or SSN-16 cards. The boot variables also should not be set to load these images.
 - Do not configure the IPS ports on the Cisco MDS 9250i Switch to 1 Gbps speed.
-

When you configure IOA, consider the following limitations for the Cisco MDS 9250i Switch:

- If the MDS 9250i Switch is part of the cluster as an IOA node, then maximum number of flows supported is 203 in one VSAN. If multiple VSAN are used, the maximum number of flows is 256.

- If the MDS9250i 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 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, and the Cisco MDS 9513 or the Cisco MDS 9509 or the Cisco MDS 9222i 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.

Configuration Limits

[Table 1: Cisco I/O Accelerator Configuration Limits](#) lists the IOA configurations and the corresponding limits.

Table 1: Cisco I/O Accelerator Configuration Limits

Parameter	MSM-18/4 or SSN-16 Module on MDS 9222i and MDS 9500 modular Chassis and MDS 9250i Fabric Switch
Number of switches in a cluster	4
Number of clusters per switch	16
Number of switches in a SAN fabric for FC-Redirect	34
Number of hosts per target	128
Number of concurrent flows per IOA service engine	128

Parameter	MSM-18/4 or SSN-16 Module on MDS 9222i and MDS 9500 modular Chassis and MDS 9250i Fabric Switch
Number of flows per IOA service engine (hard limit)	128 - Release 4.2(1) on MDS 9222i/MDS 9500 512 - Release 4.2(7) or later on MDS 9222i/MDS 9500 512 - Release 6.2(5) or later on MDS 9250i
Number of flows per IOA service engine (soft limit)	64 - Release 4.2(1) on MDS 9222i/MDS 9500 256 ¹ - Release 4.2(7) or later on MDS 9222i/MDS 9500 256 ¹ - Release 6.2(5) or later on MDS 9250i
Number of flows in a cluster	1024 - Release 4.2(7d) 1248 - Release 5.2(6b)

¹ If initiators or targets participating in IOA are present on MDS 9250i or Cisco MDS 9148S switches, then the limit is 203 for tape and 160 for disk. If initiators or targets participating in IOA are present on MDS 9250i switches, then the limit is 203 for tape and 160 for disk.

**Note**

When the new flows are load balanced again to the functional IOA interface, the soft limit is enforced to account for IOA interface failures. If the number of switches in the SAN exceeds the scalability limit, consider using CFS regions as described in [Using FC-Redirect with CFS Regions](#).
