

Cisco MDS 9000 I/O Accelerator

Q. What is the Cisco® MDS 9000 I/O Accelerator (IOA)?

A. The Cisco MDS 9000 IOA feature is a solution for accelerating disk and tape I/O operations, including disk and tape write and tape read I/O operations, to mitigate the effects of latency.

Q. What other acceleration solutions were available prior to IOA?

A. Other acceleration solutions include:

- Fibre Channel over IP (FCIP) Write Acceleration (FCIP-WA) and Tape Acceleration (FCIP-TA) on the Cisco MDS 9000 18/4-Port Multiservice Module (MSM) and 9222i Multiservice Modular Switch (MMS)
- Fibre Channel Write Acceleration (FC-WA) on Cisco MDS 9000 Storage Services Modules (SSMs)

Q. Why use IOA?

A. IOA provides a unified and comprehensive acceleration service that is transport and speed agnostic. IOA delivers disk write acceleration and tape acceleration in a highly scalable and available service. It has the following unique characteristics:

- **Transport agnostic:** It works independent of the type of data center interconnect (DCI). It supports Fibre Channel and Fibre Channel over IP (FCIP) and has been designed to support Fibre Channel over Ethernet (FCoE).
- **Speed agnostic:** It works independent of DCI link speed. It supports 1/2/4/8/10-Gbps Fibre Channel and Gigabit Ethernet and has been designed to support 16-Gbps Fibre Channel.
- **Highly available and scalable:**
 - It supports multipath connectivity for DCI over PortChannels and Equal-Cost Multipath (ECMP), to enhance link availability of the DCI.
 - Lightweight Reliability Transport Protocol (LRTP) developed by Cisco provides resiliency after Inter-Switch Link (ISL) failures, thereby protecting the application against planned and unplanned metropolitan-area network (MAN) and WAN ISL failures.
 - Engine clustering protects against acceleration and compression engine failures by reassigning flows on the failed devices to other engines.
 - Multipath and clustering support makes IOA scalable, providing the flexibility to dynamically increase DCI bandwidth and acceleration capacity.

Q. What is LRTP and how does IOA provide resiliency after ISL failure?

A. Lightweight Reliability Transport Protocol, or LRTP, is a proprietary Cisco protocol that provides reliability and in-order delivery for IOA flows. It provides TCP-like reliability and retransmission capabilities with complex semantics and congestion-control mechanisms. LRTP recovers from packet loss due to ISL failure and from reordered packets due to natural delay variances over the various ISLs in PortChannels or over ECMP. End applications such as tape backup and remote replication are unaware of transient packet loss due to ISL failures.

Note that the objective of LRTP is not to solve persistent drops in Fibre Channel networks as TCP does for IP networks, but instead to handle out-of-order frames and lost frame retransmissions.

Note also that TCP retransmission will still occur for the FCIP links in the IP network. However, the LRTP retransmission interval is greater than that for TCP.

Q. What are the top three business benefits (value propositions) of IOA?

A. In addition to the traditional business benefits of increasing replication distances and reducing backup windows, IOA offers the following unique value propositions:

- Consolidation of all acceleration services
 - IOA provides a unified and comprehensive acceleration service that works independent of media and transport. Hence, disk and tape acceleration can be unified and managed within a single system, simplifying management.
 - IOA clustering allows IOA engines to be gathered into a single group, enabling acceleration flows to be managed as one entity.
 - IOA delivers acceleration as a service with no topology restrictions and hence can be deployed anywhere in the fabric.
- Highly resilient infrastructure for business continuity and data recovery (BC/DR):
 - Support for Port Channels and multiple paths greatly reduce disruption of backup jobs and replication by enhancing link availability between data centers.
 - LRTP protects against application failures and aborts due to lost packets on failed physical links.
 - IOA clustering guards against engine failures, thereby reducing downtime.
- Highly scalable and available infrastructure, increasing cost savings
 - Backup jobs are completed faster due to fewer failures and backup job aborts.
 - The complexity of tape backup infrastructure management is reduced.
 - Support for PortChannels and multiple paths enables better MAN and WAN bandwidth utilization.

Q. What hardware modules are supported for IOA?

A. IOA is supported on the Cisco MDS 9000 18/4-Port MSM, 9222i MMS fixed module, and 16-Port Storage Service Node (SSN) module. The Cisco MDS 9000 18/4-Port MSM and 9222i MMS fixed module can host one IOA engine, and the Cisco MDS 9000 16-Port SSN can host up to four engines.

Q. What do I need to know about IOA before provisioning?

A. From a provisioning perspective, for each initiator-target pair that needs acceleration, an IOA flow needs to be created to perform either disk or tape acceleration. Flows are part of IOA flow groups. Flow groups are hosted by IOA engines on IOA switches. IOA engines are part of IOA clusters. The following list provides information about the concepts and features needed to provision IOA.

- **IOA site:** A local set of switches within a fabric; a switch can exist at only one site
- **IOA engine (interface):** Represents an application engine that performs acceleration on the supported hardware modules
 - From a configuration perspective, this engine is also called an interface.
 - An engine can be associated with only one IOA site.
 - Each engine needs a separate IOA license.
 - An engine can be part of only one IOA cluster.
- **IOA switch:** A switch containing one or more IOA engines
- **IOA cluster:** A set of IOA engines in a pair of IOA sites that operate as one group to provide IOA service
 - A cluster can span only two sites.
 - A site can share multiple clusters (for example, a bunker site with all the tapes).
 - A switch can belong to multiple clusters.

- An engine is bound to only one cluster.
- **IOA flow:** A tuple of initiator and target that is accelerated within an IOA cluster
 - Each flow can perform write acceleration or tape acceleration, or both.
 - Each flow is identified by the port world wide name (PWWN) and VSAN ID: (initiator PWWN, target PWWN, VSAN ID); for example, (I1, T1, V1) identifies a flow.
- **IOA flow group:** A set of IOA flows classified for a given purpose, such as a tape backup flow group, EMC Symmetrix Replication Data Facility (SRDF) flow group, and Tivoli Storage Manager (TSM) tape flow group

Q. Can the same IOA engine handle both tape backup and disk replication flows?

A. Yes. The same engine can perform both disk and tape write acceleration. In addition, a flow can perform tape backup or disk replication write acceleration.

Q. What does IOA's Service-Oriented Architecture mean?

A. IOA can be deployed as a service anywhere in the SAN. This is accomplished by using a Cisco developed technology called FC-Redirect. After IOA is configured, the IOA flows are automatically redirected to an IOA engine by FC-Redirect no matter where the IOA engine is located. As long as FC-Redirect is available on the switch connected to the backup server or storage array replication port, the IOA engine can accelerate the flows.

Q. What is FC-Redirect?

A. FC-Redirect is a distributed flow redirection mechanism that can enable redirection of a set of traffic flows to a special engine that performs services such as IOA and storage media encryption (SME). FC-Redirect is supported on all Cisco MDS 9500 Series Multilayer Directors and 9200 Series Multilayer Switches since Cisco MDS 9000 SAN-OS Software Release 3.0.

Q. Do IOA clusters span fabrics? Do I need separate clusters for fabric A and fabric B?

A. IOA clusters do not span fabrics. Separate clusters are needed for fabric A and fabric B.

Q. What are the deployment considerations for IOA?

A. You need to consider the following criteria when deploying IOA:

- Location of the IOA engine
- Number of IOA engines required
- Hardware options
- ISL capacity of the links connecting the IOA engines
- Number of IOA clusters needed

Q. Although IOA can be placed anywhere in the SAN, what is the best place for the IOA engine?

A. The best place for the IOA engine is the core, which has WAN connectivity, for the following reasons:

- **Easy planning and transition:** Capacity can be planned based on WAN ISL throughput on the same switch.
- **Optimal routing:** Flows always traverse core switches to reach the remote site.
- **Consolidation and scalability:** Consolidation reduces the number of engines (and hence the number of switches) required. Without consolidation, more IOA switches may be required. Only four IOA switches per cluster are supported.

Special considerations for IOA engines hosted on multiple core switches (in a single fabric) are:

- The core switches should be interconnected to provide the best possible redundancy.
- The connecting ISLs should have at least the same capacity as the aggregate MAN or WAN link connecting the two sites to help ensure that the ISLs can carry the traffic in case of an IOA engine failure.

Q. How many IOA engines (per fabric) are required for my setup?**A.** Consider these factors to determine the number of IOA engines required:

- The application throughput required, availability requirements, and the number of interconnected sites.
- Each IOA engine provides 10-Gbps application throughput.
- A pair of engines is required for any flow between two sites.
- An engine can be in only one cluster.
- More than one engine per site can be deployed to increase the availability of the acceleration service.
- The number of IOA engines per site depends on the throughput requirements at each site. If two sites need to service the same number of flows and applications, then the same number of IOA engines may be required.

Note that the same number of IOA engines is needed for the redundant or B fabric as well.

Q. How do I choose the right hardware from the available options?

A. The choice depends on the IOA engine throughput requirements, the number of Fibre Channel ports, and high-availability requirements. For example, if higher IOA engine density is needed, Cisco MDS 9000 16-Port SSN is the choice. Cisco MDS 9000 16-Port SSN is the preferred option for IOA sites that consolidate replication and backup applications. Depending on other Fibre Channel connectivity requirements, either Cisco MDS 9000 18/4-Port MSM or 9222i MMS engines can be used.

Q. What is the performance data for IOA?

A. IOA performance will be at least the same as FCIP-WA and FCIP-TA. However, two aspects of the IOA architecture help IOA achieve better overall performance than FCIP:

- IOA does not use TCP and has no TCP overhead.
- IOA supports multiple paths and PortChannels and hence can support higher aggregate application throughput.

Q. Are IOA and FCIP supported on the same engine?

A. No. For performance reasons, IOA and FCIP are not supported on the same engine.

Q. Who should migrate to IOA?

A. The current acceleration solutions (including those of competitors) have limitations and restrictions that increase deployment and management complexity. IOA offers a superior option.

Note that for migration from FC-WA, either Cisco MDS 9000 18/4-Port MSM or 16-Port SSN is required since Cisco MDS 9000 SSMs do not support IOA.

Q. I currently have FCIP-WA and FCIP-TA? What else do I need to migrate to IOA?

A. IOA and FCIP cannot be hosted on the same engine. Depending on the number of the Cisco MDS 9000 18/4-Port MSM or 9222i MMS modules, you may need a different engine for IOA. Alternatively, you can consolidate both FCIP and IOA on the same module using a Cisco MDS 9000 16-Port SSN.

Q. What is the licensing model for IOA?

A. Each IOA engine needs a separate IOA license. For Cisco MDS 9000 18/4-Port MSM and 9222i MMS, one license is required. To run four IOA engines, four IOA licenses are required.

Q. How do I manage IOA?

A. You manage IOA just like any other Cisco MDS 9000 Family features, using both the command-line interface (CLI) and the Cisco Fabric Manager. The Cisco Fabric Manager has an IOA wizard that simplifies the provisioning into just a few steps. It also provides a snapshot of the IOA configuration throughout the fabric. In addition, the CLI has an IOA flow setup wizard to help with automatic batch provisioning of flows.

Q. Both IOA and FCIP can perform compression. In a deployment with both FCIP and IOA engines, where should compression be enabled?

A. FCIP and IOA can provide similar compression services and throughput per engine, so it is important to understand which engines should be used for compression. IOA provides an overall superior compression solution for the following reasons:

- **Better high availability:** Although FCIP uses TCP retransmissions, link failures are not transparent to the applications and hence may cause disruption. With IOA, LRTP protects against link failure and reduces their effects on applications.
- **Flexible scalability:** IOA compression engine capacity can be scaled up by adding more engines without the need for fabric reconfiguration or rewiring. FCIP engines need Gigabit Ethernet port connectivity and wiring to the fabric. In addition, if FCIP-WA is being used, DCI bandwidth is limited to 16 Gbps (16 Gigabit Ethernet ports in a Port Channel) because multiple paths are not supported for FCIP-WA. No such limit exists for IOA.

The decision also depends on the total number of FCIP and IOA engines (the Cisco MDS 9000 18/4-Port MSM can support one FCIP or IOA engine, and the Cisco MDS 9000 SSN-16 can support four FCIP or IOA engines) and hence the aggregate compression bandwidth. A good approach is to enable compression on the engine type with higher aggregate compression throughput. For example, if you have two FCIP engines and one IOA engine, FCIP provides higher compression throughput.

Q. FCIP-WA is bundled with an FCIP license. IOA needs a separate license. Are there any advantages in moving to IOA from FCIP-WA?

A. The main advantages of IOA over FCIP-WA are better scalability and high availability.

FCIP-WA uses the TCP retransmission mechanism to protect against lost packets at the IP level. However, link failures affect applications. IOA uses LRTP to protect against link failures at the Small Computer System Interface (SCSI) level, thereby reducing the effect on the applications.

IOA engine capacity can be scaled up by adding more engines without the need for fabric reconfiguration or rewiring. FCIP engines need Gigabit Ethernet port connectivity and wiring to the fabric. In addition, if FCIP-WA is being used, DCI bandwidth is limited to 16 Gbps (16 Gigabit Ethernet ports in a PortChannel) because multiple paths are not supported for FCIP-WA. No such limit exists for IOA.

IOA also provides flexibility, enabling the same infrastructure to be shared to deploy both replication and backup applications independent of transport in the future.



Americas Headquarters
Cisco Systems, Inc.
San Jose, CA

Asia Pacific Headquarters
Cisco Systems (USA) Pte. Ltd.
Singapore

Europe Headquarters
Cisco Systems International BV
Amsterdam, The Netherlands

Cisco has more than 200 offices worldwide. Addresses, phone numbers, and fax numbers are listed on the Cisco Website at www.cisco.com/go/offices.

CCDE, CCENT, CCSI, Cisco Eos, Cisco HealthPresence, Cisco IronPort, the Cisco logo, Cisco Nurse Connect, Cisco Pulse, Cisco SensorBase, Cisco StackPower, Cisco StadiumVision, Cisco TelePresence, Cisco Unified Computing System, Cisco WebEx, DCE, Flip Channels, Flip for Good, Flip Mino, Flipshare (Design), Flip Ultra, Flip Video, Flip Video (Design), Instant Broadband, and Welcome to the Human Network are trademarks; Changing the Way We Work, Live, Play, and Learn, Cisco Capital, Cisco Capital (Design), Cisco.Financed (Stylized), Cisco Store, Flip Gift Card, and One Million Acts of Green are service marks; and Access Registrar, Aironet, AllTouch, AsyncOS, Bringing the Meeting To You, Catalyst, CCDA, CCDP, CCIE, CCIP, CCNA, CCNP, CCSP, CCVP, Cisco, the Cisco Certified Internetwork Expert logo, Cisco IOS, Cisco Lumin, Cisco Nexus, Cisco Press, Cisco Systems, Cisco Systems Capital, the Cisco Systems logo, Cisco Unity, Collaboration Without Limitation, Continuum, EtherFast, EtherSwitch, Event Center, Explorer, Follow Me Browsing, GainMaker, iLYNX, IOS, iPhone, IronPort, the IronPort logo, Laser Link, LightStream, Linksys, MeetingPlace, MeetingPlace Chime Sound, MGX, Networkers, Networking Academy, PCNow, PIX, PowerKEY, PowerPanels, PowerTV, PowerTV (Design), PowerVu, Prisma, ProConnect, ROSA, SenderBase, SMARTnet, Spectrum Expert, StackWise, WebEx, and the WebEx logo are registered trademarks of Cisco Systems, Inc. and/or its affiliates in the United States and certain other countries.

All other trademarks mentioned in this document or website are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (0910R)