

## Cisco MDS 9000 Family Acceleration Services: Enhance Synchronous Replication Performance

### What You Will Learn

A comprehensive business continuity and disaster recovery (BC/DR) plan mandates the deployment of multiple data centers located far enough apart to protect against regional power failures or disasters, yet close enough together for synchronous data replication without affecting application performance. Synchronous replication solutions such as EMC Symmetrix Remote Data Facility/Synchronous (SRDF/S) and Hitachi TrueCopy Synchronous provide a continuous, nondisruptive, host-independent remote data replication solution for disaster recovery or data migration over distance. However, application performance is affected by the distance and latency between the data centers, which restricts the location of the data centers.

The biggest challenge for synchronous replication is its distance limitation. Fibre Channel, the primary enterprise storage transport protocol, is limited only by its physical-layer flow-control mechanism. However, latency becomes a problem because propagation delays lengthen with increased distance. Propagation delays can significantly slow down a system by forcing it to wait for confirmation of the completion of each storage operation at both the local and remote sites. This means that the practical distance for synchronous replication is about 100 to 200 km, or 60 to 120 miles, depending on the application response time tolerance and other factors.

Cisco® MDS 9000 Family switches offer traffic acceleration services capable of dramatically reducing I/O latency between data centers and thereby increasing the distance for synchronous replication solutions. Cisco MDS 9000 I/O Accelerator (IOA) offers industry-leading, highly resilient, clustered, and transport-independent architecture for both Fibre Channel and Fibre Channel over IP (FCIP) transports, providing the best I/O acceleration service in the industry.

This document describes tests that demonstrate that Cisco MDS 9000 IOA enhances synchronous replication performance and extends replication distance between primary and secondary data centers.

### Challenge

Today's businesses operate more than one data center at different locations as part of a BC/DR plan. Such a model helps optimize the cost of running data centers while complying with government regulations. An effective BC/DR plan mandates the deployment of multiple data centers located far enough apart to protect against regional power failures or disasters, yet close enough together for synchronous data replication without affecting application performance. Achieving this balance for successful business continuance poses a significant challenge for businesses.

Synchronous replication solutions provide remote replication for BC/DR. Replication can be either storage-array based or host based. EMC SRDF/S and Hitachi TrueCopy Synchronous are examples of array-based replication solutions.

The recovery point objective (RPO)—the amount of data loss during a failure—and recovery time objective (RTO)—the amount of time that the business needs to recover after a disaster—determines the type of replication deployed. Synchronous data replication provides the best RPO but requires the data centers to be located close enough to maintain consistent application performance.

In the case of synchronous replication, the initiator (primary storage array) waits until the target (secondary storage array) successfully stores the data and acknowledges the current outstanding I/O. This process helps ensure that both data centers have the same data at the same time, enabling zero data loss and the best RPO. However, the distance between data centers may not meet the requirements of the BC/DR plan.

Asynchronous or semisynchronous replication supports greater distance by not waiting for the secondary storage array to complete the I/O operation before completing the I/O operation at the application level. Therefore, the secondary data center lags behind the primary data center by a fixed amount of time, and the data is not the same at both data centers at any given point in time. This approach implies that data can be lost if a failure occurs.

Therefore, a method to reduce latency or increase distance is required for comprehensive BC/DR plans that require synchronous replication to meet RPO and RTO goals.

## Business Benefits

Figure 1 shows the effect of latency in the communication channel on the time taken to complete the I/O operation during a Small Computer System Interface (SCSI) write operation. The SCSI write operation is a four-step process consisting of the write command, a transfer ready (XferReady) command from the target, the transmission of the data, and the final acknowledgment (status) from the target. Therefore, the latency between the data centers is added to each step of the process and is equal to twice the round-trip time (RTT) between data centers.

Cisco MDS 9000 Family traffic acceleration services mitigate the effects of latency by proxying the XferReady command and allowing the data to be sent to the remote data center without waiting for the target XferReady message to be transferred all the way back to the host, as shown in the Figure 2. To preserve data integrity, the status message is not proxied. The net effect is that the application latency can be reduced by up to 50 percent at the same distance, or the distance between data centers can be increased by up to 100 percent with the same latency.

Figure 1: Traditional Write Command Operation Without Acceleration

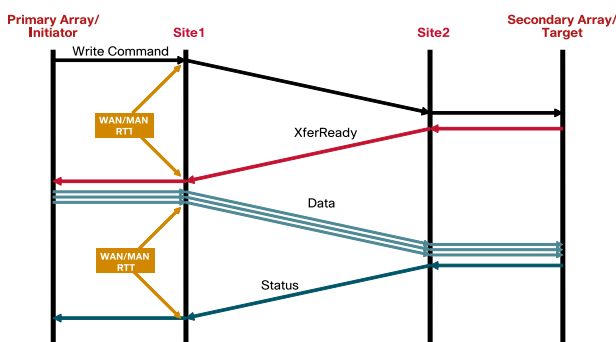
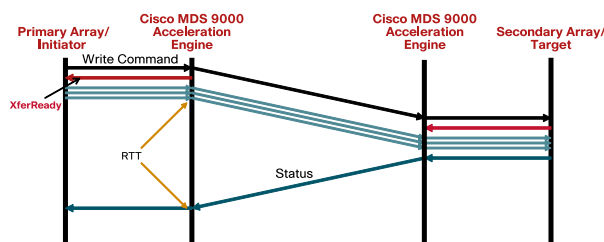


Figure 2: Cisco MDS 9000 Switches Proxying the XferReady Command



## Cisco Solution

Cisco MDS 9000 Family switches offer traffic acceleration services capable of reducing I/O latency and increasing the distance between data centers for synchronous replication solutions.

FCIP Write Acceleration (FCIP-WA) offers a cost-effective solution for replication over FCIP and is included in the FCIP SAN Extension license.

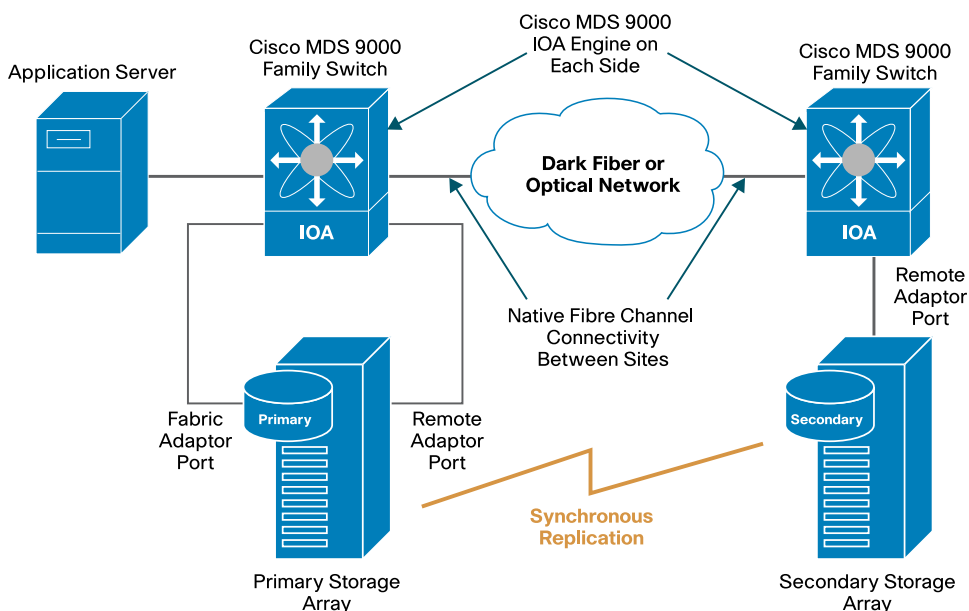
Cisco MDS 9000 IOA provides increased scalability and is a comprehensive, next-generation, transport-independent acceleration solution for both Fibre Channel and FCIP. The Cisco MDS 9000 IOA industry-leading, highly resilient, cluster technology, with integrated compression and an error-recovery mechanism, provides the best I/O acceleration service in the industry.

## Solution Details

Figure 3 shows a typical collapsed-core deployment of synchronous replication with Cisco MDS 9000 IOA. Cisco MDS 9000 IOA is deployed on the same switch as both the application hosts and storage. This setup was used for the test examples described in this document. Both EMC SRDF/S and Hitachi TrueCopy Synchronous were used as the replication methods for the tests.

Alternatively, Cisco MDS 9000 IOA can be deployed in a core-edge or edge-core-edge topology because of the Cisco MDS 9000 Fibre Channel Redirect (FCR) function that transparently reroutes traffic in the switch fabric to the Cisco MDS 9000 IOA engine.

**Figure 3.** Storage Array Replication Deployment with Cisco MDS 9000 IOA (Only SAN A Is Shown)



Although Cisco MDS 9000 IOA eliminates one RTT and improves application throughput, the net application performance improvement depends on the following factors:

- Distance between data centers (sites)
- Write I/O size
- Number of outstanding I/O operations

Distance between the data centers adds additional latency, and any RTT savings has a positive effect on the throughput. The effect of distance is somewhat masked by larger I/O sizes. Larger I/O sizes can use more of the physical link and may not have as much of an effect on application performance. However, I/O sizes are generally dictated by the application and its performance requirements.

The effect of distance is more severe if the number of outstanding I/O operations between the initiator and the target is small. In the case of synchronous replication, there is only one outstanding I/O operation, which means that the initiator does not send the next write command before it gets a response from the previous write command. In asynchronous replication, there can be many outstanding I/O operations. This means that multiple write operations may be performed at the same time, better utilizing the physical link. Hence, the performance improvement for synchronous replication is greater than for asynchronous replication.

To demonstrate the performance characteristics of Cisco MDS 9000 IOA, tests using an I/O meter were performed using both EMC SRDF/S and Hitachi TrueCopy Synchronous with SCSI write-only traffic. The performance increase was measured in terms of relative increase in application throughput and decrease in latency over different distances and I/O sizes. These tests were performed accelerating the traffic flow between two storage arrays running synchronous replication.

Figure 4 shows the effect of distance on application throughput. The results show that Cisco MDS 9000 IOA eliminates the negative performance impact of increased distance on replication throughput and helps reduce application latency. Note that similar improvements are seen with synchronous replication when the solution is deployed with FCIP-WA.

**Figure 4.** Cisco MDS 9000 Family Acceleration Services Mitigate Effect of Distance on Application Throughput and Latency (I/O Size = 256 KB)

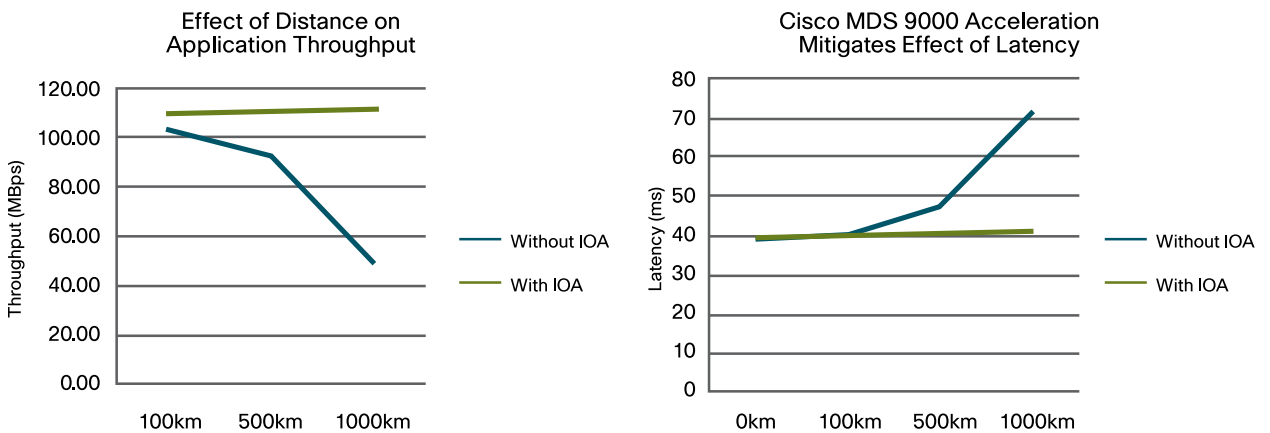
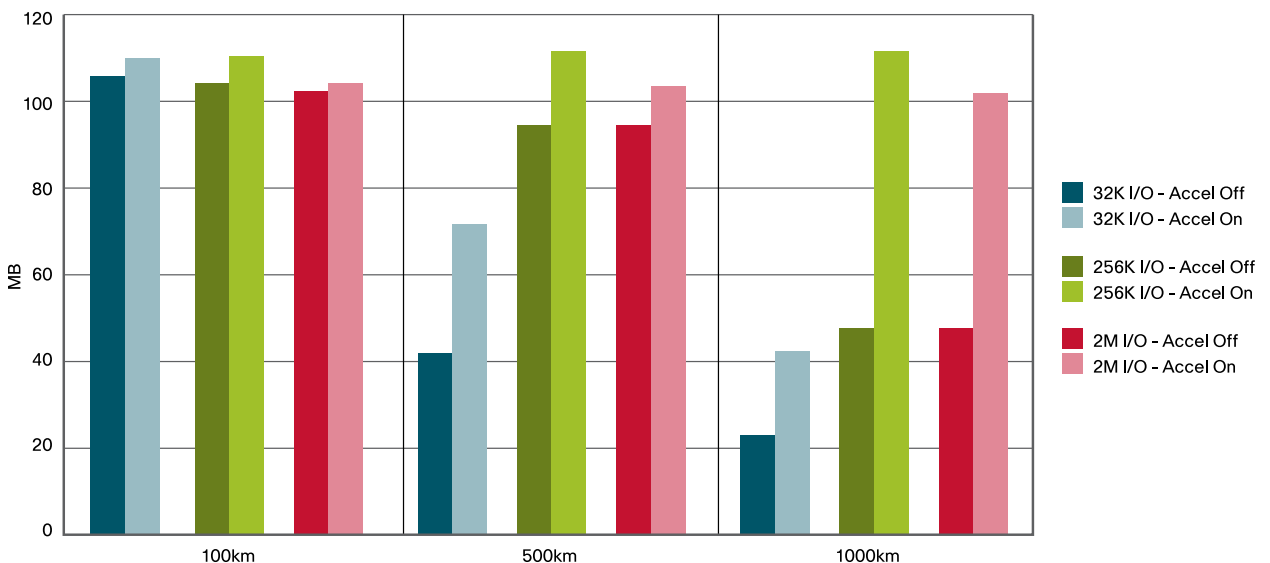


Figure 5 shows the relationship between distance and I/O size and the effect of Cisco MDS 9000 IOA on performance. With Cisco MDS 9000 IOA, the performance improvements are greater with larger I/O sizes. However, the improvement is significant even at the smaller I/O sizes (50 to 70 percent improvement in application throughput compared to the case without Cisco MDS 9000 IOA).

**Figure 5.** Performance Improvement Is Relatively Greater for Larger I/O Sizes



## Deployment Guidelines

- One Cisco MDS 9000 IOA engine is required per site per fabric. However, the use of two engines per site per fabric is recommended for increased redundancy and capacity. The number of Cisco MDS 9000 IOA engines required also depends on application throughput. Each engine can provide 10 Gbps of raw application throughput.
- Native Fibre Channel is typically the transport of choice for the data center interconnect (DCI) when running synchronous replication between storage arrays. Depending on the application latency requirements, FCIP may be used.
- The use of multiple paths for DCI provides additional resiliency after physical link failures. Cisco MDS 9000 Family PortChannels—logical links composed of more than one physical link—are recommended for best link availability.

## Conclusion

As demonstrated in the test results, Cisco MDS 9000 IOA enables increased distance between data centers for synchronous replication by reducing latency and provides increased flexibility in the choice of primary and secondary data center locations.

## For More Information

Cisco MDS I/O Accelerator: <http://www.cisco.com/en/US/products/ps10531/index.html>

Cisco MDS 9000 SAN Extension over IP Package:

[http://www.cisco.com/en/US/prod/collateral/ps4159/ps6409/ps4358/product\\_data\\_sheet09186a00801cc917.html](http://www.cisco.com/en/US/prod/collateral/ps4159/ps6409/ps4358/product_data_sheet09186a00801cc917.html)



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](http://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, iLync, 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)