

White Paper

SILVER PEAK REPLICATION ACCELERATION AND EMC SRDF/A

Copyright © 2013 EMC Corporation. All Rights Reserved.

EMC believes the information in this publication is accurate as of its publication date. The information is subject to change without notice.

The information in this publication is provided “as is.” EMC Corporation makes no representations or warranties of any kind with respect to the information in this publication, and specifically disclaims implied warranties of merchantability or fitness for a particular purpose.

Use, copying, and distribution of any EMC software described in this publication requires an applicable software license.

For the most up-to-date listing of EMC product names, see EMC Corporation Trademarks on EMC.com.

Silver Peak Replication Acceleration and EMC SRDF/A

Table of Contents

Silver Peak Replication Acceleration and EMC SRDF/A	1
Overview	3
Audience.....	3
The Challenge of Replicating Over Distance.....	4
Silver Peak VXOA Architecture.....	5
When to Deploy Silver Peak with SRDF	6
SRDF/A Testing and Validation With Silver Peak.....	7
Test Methodology.....	7
Testing Suite Overview	7
Test Results	8
Jumbo Frames	12
Deployment Architectures	13
Ethernet Fabric Speed	13
LAN Switch Performance Differences.....	14
WAN with multiple applications.....	15
Silver Peak Deployment Options	16
Silver Peak Out-of-Path (Router Mode)	16
Silver Peak In-Line Deployment (Bridge Mode)	17
Silver Peak Velocity Deployment.....	18
High Availability Silver Peak Deployments.....	18
Deployment Best Practices.....	19
Jumbo Frames	19
SRDF/A Settings	20
Silver Peak VXOA Settings	20
Tunnel configuration	21
Optimization Policy	22
Ethernet Switches	22
Monitoring Performance Using the Silver Peak GUI.....	23
Conclusion	24
Appendix A.....	25
Links to Additional Documentation	25

Appendix B	26
Environmental Variables Affecting Performance	26
IO size and type	26
Number of LUNs being replicated.....	26

Overview

The EMC Symmetrix Remote Data Facility / Asynchronous (SRDF/A) is a key component for disaster recovery with EMC Symmetrix storage arrays. SRDF/A provides the technology necessary to replicate large amounts of data between geographically separated data centers. SRDF/A throughput can be severely limited by problems with the WAN connection between the primary and secondary data centers. Limited bandwidth, latency, dropped packet, and out-of-order packets all reduce SRDF/A throughput, causing missed recovery point objectives (RPO) and failed replication.

Silver Peak develops data center class Replication Acceleration software that fixes the problems that exist on the WAN, enabling more data to be replicated, in less time, and over longer distances. Silver Peak's Virtual Acceleration Open Architecture (VXOA) can be deployed as virtual software (VX or VRX) on any hypervisor or as a hardware appliance (NX). VXOA solves WAN problems by deduplicating and compressing replication traffic, repairing dropped and out-of-order packets, and overcoming the effects of latency.

When the problems on the WAN have been repaired it is easy to meet, maintain, and even reduce RPOs.

This architecture guide details the benefits that Silver Peak Replication Acceleration provides for EMC SRDF/A environments using 1 GbE or 10 GbE directors. Specific use cases are included that detail the effect on SRDF/A throughput from environment variables like IO size and type on the SAN, LAN switch configuration, and WAN characteristics. Best practices for architecting and deploying EMC SRDF/A with Silver Peak's solution are included as well.

Silver Peak replication acceleration should be deployed to increase SRDF/A throughput when there is Latency of 40 ms or greater on the WAN (3991 km/2480 miles) and any of the following conditions:

- The network used for replication has packet loss or out-of-order (MPLS or Internet VPN)
- WAN bandwidth is limiting the throughput of SRDF/A
- The WAN is shared by multiple applications

Any of the conditions listed above will reduce SRDF/A throughput. Deploying Silver Peak replication acceleration will increase SRDF/A throughput up to 20X.

Audience

This white paper is intended for storage architects, IT engineers, storage administrators, professional service implementers, and SAN administrators who are responsible for architecting or implementing EMC SRDF/A with Silver Peak VXOA for disaster recovery purposes.

The Challenge of Replicating Over Distance

Due to the nature of TCP, SRDF/A throughput across a WAN is affected by several things that can cause disruption to replication, resulting in missed RPOs or SRDF/A going into suspend mode, stopping all replication traffic. Latency across the WAN caused by the geographic distance between data centers is a common problem that causes reduced throughput. The speed of light through fiber optic cable is approximately 124 miles per millisecond, or 199.5 Kilometers per millisecond. For every 124 miles of cable, there is 1 millisecond of latency. This doesn't take into account the additional latency from indirect network paths and network switching and routing infrastructure, which will often double the observed latency.

SRDF/A throughput is also affected by the amount and quality of the bandwidth used for replication. When the WAN suffers from dropped or out-of-order packets, it can have a severe impact on throughput. The problem of throughput with dropped and out-of-order packets is exacerbated when latency is present. Finally, the amount of bandwidth that is available for SRDF/A directly affects throughput. When replication is performed across a WAN that has 50 Mbps of available bandwidth, only 50 Mbps can be used for replication.

There is compression available on the GbE and 10 GbE directors for SRDF/A, but even at 3:1 compression, throughput will be limited to 150 Mbps. Of course, the real limitation on the amount of bandwidth that can be used for SRDF/A will be the latency and quality of the WAN. A WAN with 80 ms of latency (approximately 6,400 km or 3900 miles) and 0.1% packet loss (MPLS) will only be able to support 75 Mbps of SRDF/A replication traffic, regardless of the available bandwidth.

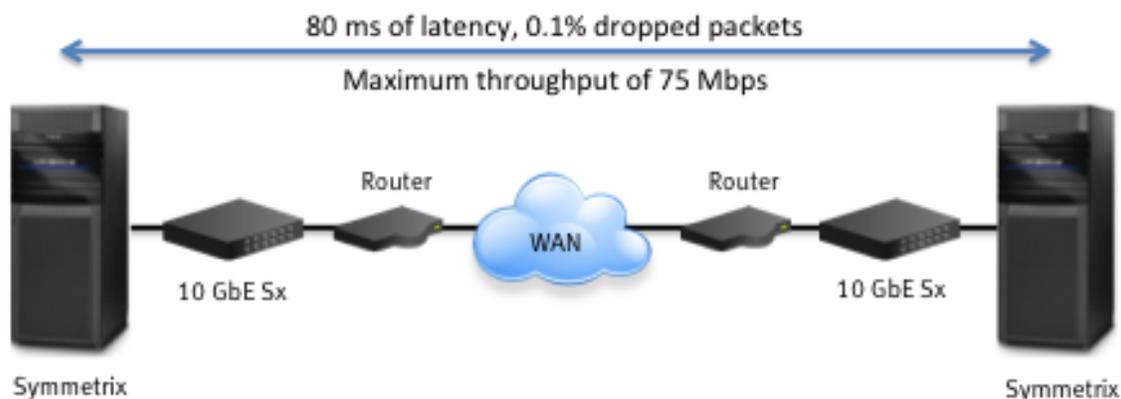


Figure 1: Example SRDF/A environment

Silver Peak's replication acceleration features can remove the impact of latency on throughput, while repairing dropped and out-of-order packets in real time, allowing all of the available bandwidth to be used. Silver Peak will also perform de-duplication and compression, resulting in a much higher effective SRDF/A throughput. In typical deployments Silver Peak is able to deduplicate and compress the replication traffic sent by SRDF/A by 6X, resulting in an effective 6X bandwidth increase. A 50 Mbps WAN will appear to be a 300 Mbps WAN to SRDF/A.

Silver Peak VXOA Architecture

Silver Peak provides a robust WAN optimization solution that addresses the bandwidth, latency, and packet loss issues that are common to most enterprise environments. Silver Peak's optimization techniques are all performed in real-time and primarily at the network (IP) layer to ensure maximum performance across the widest range of applications and WAN environments.



Figure 2: Silver Peak VXOA Architecture

Silver Peak appliances leverage the following Virtual Acceleration Open Architecture (VXOA) technology components to accelerate all enterprise applications in a secure and reliable fashion:

Network Memory: Silver Peak's patent-pending solution for disk-based WAN de-duplication. Network Memory inspects all traffic that is sent between clients and servers, storing information as a local instance in Silver Peak appliances. Repetitive information is delivered locally rather than sent across the WAN, improving application performance and WAN utilization. Cross-flow payload and header compression provide additional gains on first-time data transfers and non-repetitive traffic.

Network Integrity: Silver Peak employs a variety of real-time techniques to address packet delivery issues common to shared WAN technologies, such as MPLS and IP VPN. These include adaptive Forward Error Correction (FEC) and Packet Order Correction (POC) to overcome dropped and out-of-order packets, and advanced Quality of Service (QoS) techniques to prioritize traffic and guarantee network resources.

Network Acceleration: Silver Peak mitigates the impacts of latency across the WAN by using various TCP acceleration techniques, like adjustable window sizing and selective acknowledgements. These tools help to overcome inherent chattiness that can otherwise hamper application performance across a WAN.

Virtual Deployment Options: Silver Peak can be deployed as a physical appliance (NX) or a virtual instance (VX/VRX). Silver Peak’s virtual products can be deployed on any hypervisor (VMware, Microsoft Hyper-V, Xen, and KVM) and on any server hardware that meets the minimum requirements.

Silver Peak offers two virtual appliance product lines:

- VX-Series are for general purpose WAN optimization and are sized by bandwidth
- VRX-Series are for replication acceleration and are sized by throughput (GB/TB per hour)

VX Model	VX 500	VX 1000	VX 2000	VX 3000	VX 5000	VX 6000	VX 7000	VX 8000	VX 9000
WAN Size	2 Mbps	4 Mbps	10 Mbps	20 Mbps	50 Mbps	100 Mbps	200 Mbps	622 Mbps	1 Gbps

Table 1: VX Models (General WAN Optimization)

VRX Model	VRX-2	VRX-4	VRX-8
Throughput	Up to 60 GB/hr	Up to 300 GB/hr	Up to 1.5 TB/hr

Table 2: VRX Models (Replication Only)

When to Deploy Silver Peak with SRDF

Silver Peak Replication Acceleration provides benefits that can be classified into three categories: latency, quality, and bandwidth. Depending on the issues that are affecting SRDF/A replication throughput, Silver Peak can provide a solution that will accelerate replication, but there are scenarios where Silver Peak will provide limited value. For example, SRDF/S is used for synchronous replication and requires high bandwidth and low latency (typically less than 5ms) connections with no loss or quality problems. Because the connection does not have bandwidth constraints, high latency, or lost or out-of-order packets, Silver Peak will not be able to accelerate the SRDF throughput.

Silver Peak should be deployed with SRDF/A when more than 80 ms of latency is present on the WAN and 1 GbE RE ports are used. When 10 GbE RE ports are used, Silver Peak should be used when there is 40 ms or greater latency on the WAN. If there is packet loss or congestion on the WAN that is causing SRDF/A throughput problems, Silver Peak should be deployed with as little as 20 ms of latency. When SRDF/A throughput is limited because of available WAN bandwidth, and latency is 20 ms or higher, Silver Peak should be deployed to de-duplicate and compress the traffic creating an effective bandwidth increase of up to 6X.

Scenario	Deploy Silver Peak
SRDF/S replication	No
SRDF/A replication with no loss and less than 80ms latency and 1 GbE RE ports	No
SRDF/A replication with more than 80ms of latency and 1 GbE RE ports	Yes
SRDF/A replication with no loss and less than 40 ms of latency and 10 GbE RE ports	No
SRDF/A replication with more than 40 ms of latency and 10 GbE RE ports	Yes
SRDF/A replication with packet loss or congestion (MPLS or shared WAN) and more than 20 ms of latency	Yes
Any replication performed over an Internet VPN	Yes
SRDF/A throughput is limited because of WAN bandwidth	Yes

Table 3: Scenarios for deployment of VXOA with SRDF/A

SRDF/A Testing and Validation With Silver Peak

Test Methodology

A dedicated environment was set up to test deployment scenarios across multiple WAN settings. Silver Peak physical and virtual appliances were used for replication acceleration across multiple WAN bandwidths, latencies, and loss characteristics. Additional testing was performed to determine the impact of IO size and type, SRDF/A cycle time, and the number of simultaneous LUNs being replicated. All of the testing was performed across a 10 GbE LAN using 4 RE ports on each array with 16 total TCP connections so the maximum throughput between the arrays and the Silver Peak appliance could be obtained.

Testing Suite Overview

Device/Tool	Name	Version
Storage Arrays	EMC Symmetrix	VMAX – Single Engine
Software	Enginuity	5876.159.102
Replication Software	SRDF/A	5867.159.102
Silver Peak Appliances	NX, VRX, VX	NX-9700, VRX-8, VX-5000
Silver Peak Software	VXOA	5.2.3
Traffic Generator	IOrate Linux	2.12
WAN Emulator	KWANEM	52909

Table 4: Tested Devices and Software Versions

Bandwidth Tested	Latencies Tested	Loss Percentages Tested
155 Mbps	20, 40, 80, 120, 200 ms	0, 0.1, 1%
622Mbps	20, 40, 80, 120, 200 ms	0, 0.1, 1%
1 Gbps	20, 40, 80, 120, 200 ms	0, 0.1, 1%

Table 5: Test Parameters

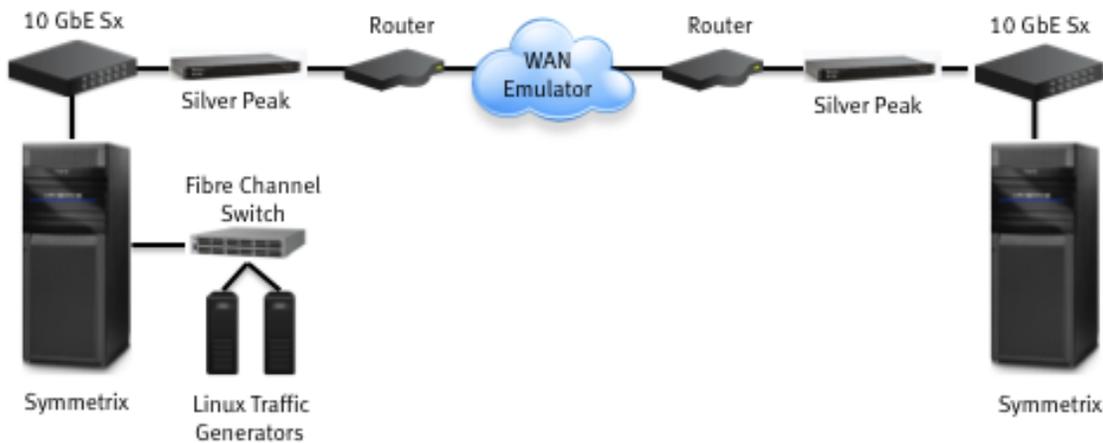


Figure 3: Silver Peak and EMC Symmetrix Test Diagram

Test Results

The average data reduction ratio with Silver Peak and SRDF/A in most customer environments is 6X, or 83.4%. During lab testing the average was 6.1X, or 83.6%. This number can change based on the data type being replicated, and whether the data is compressed, encrypted, or both when stored. Data that is compressed will yield a lower reduction, while data that is encrypted will have no decrease in size when replicated.

WAN Bandwidth	Packet Loss %	Average Reduction Ratio	Average Reduction %
155 Mbps	0.1	5.5X	81.8
155 Mbps	1	5.6X	82.4
622 Mbps	0.1	6.8X	85.3
622 Mbps	1	6.7X	85.2
1 Gbps	0.1	7X	85.8
1 Gbps	1	5.1X	80.5

Table 6: Average data reduction during testing with Silver Peak and SRDF/A

WAN Type	Typical Packet Loss
Private Line	0%
MPLS	0.1%
Internet VPN	1% or higher

Table 7: Typical WAN packet loss

During the un-optimized SRDF/A only testing, hardware compression on the RE ports was enabled. With hardware compression enabled the amount of data replicated by SRDF/A is higher than the WAN bandwidth used. For example, a 100 Mbps SRDF/A flow will be transmitting 200 – 300 Mbps of data.

When Silver Peak is deployed with SRDF/A the results show **6 – 20X** performance improvement across bandwidths, latencies, and loss values.

Figures 4, 5, and 6 illustrate SRDF/A throughput with and without Silver Peak acceleration and 0% packet loss (simulating a private line). These results show the benefit of Silver Peak's Network Memory de-duplication and compression technology. Throughput improvements for SRDF/A range from 1.5X to over 6X on a network with no quality issues.

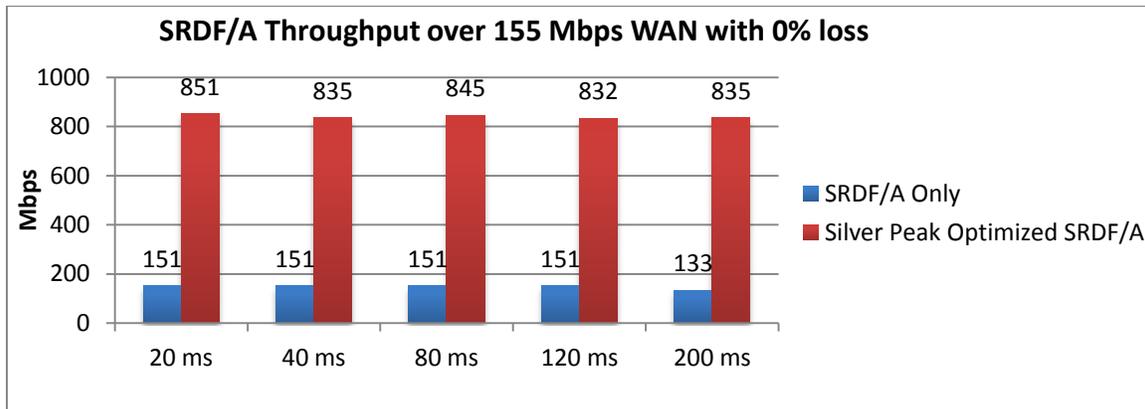


Figure 4: 155 Mbps WAN with 0% packet loss

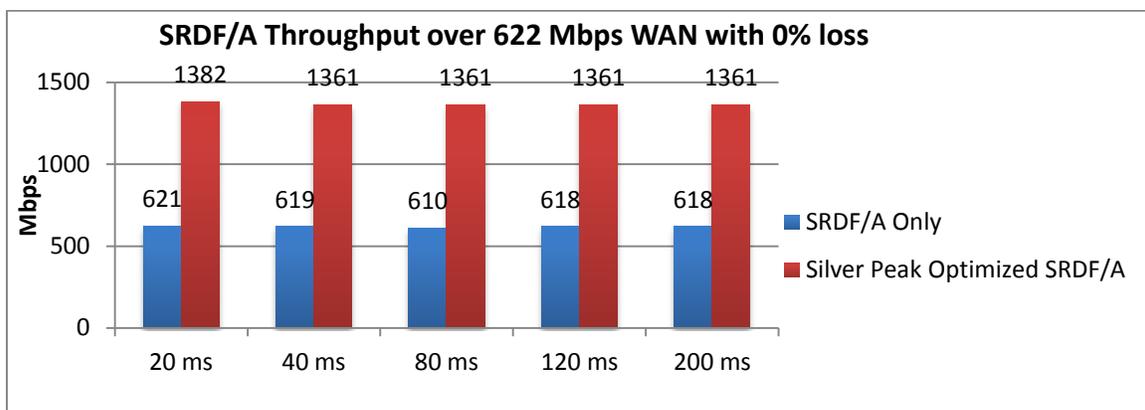


Figure 5: 622 Mbps WAN with 0% packet loss

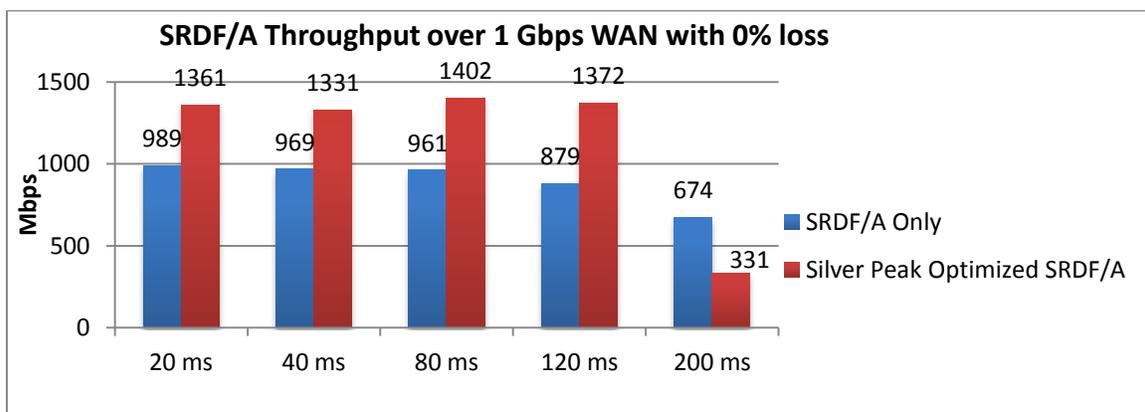


Figure 6: 1 Gbps WAN with 0% packet loss

In Figures 7 and 8, the un-optimized, SRDF/A-only, throughput across a 155 Mbps WAN starts at 151 Mbps and drops to 116 Mbps with latency when low packet loss, 0.1%, is present on the simulated WAN. However, when packet loss is increased to 1%, the throughput drops more dramatically as latency is increased. When Silver Peak is used to accelerate the SRDF/A replication, throughput is increased to 700 Mbps with .1%, and 650 Mbps with 1% packet loss.

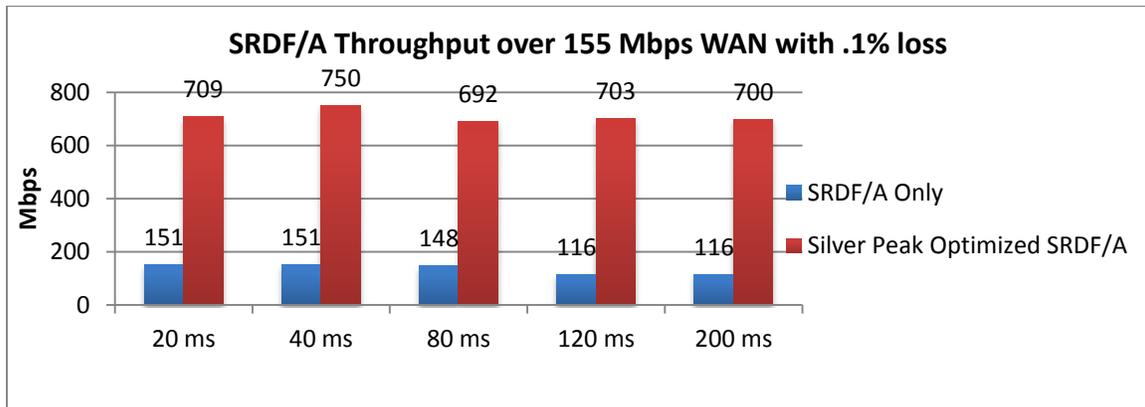


Figure 7: 155 Mbps WAN with .1% packet loss

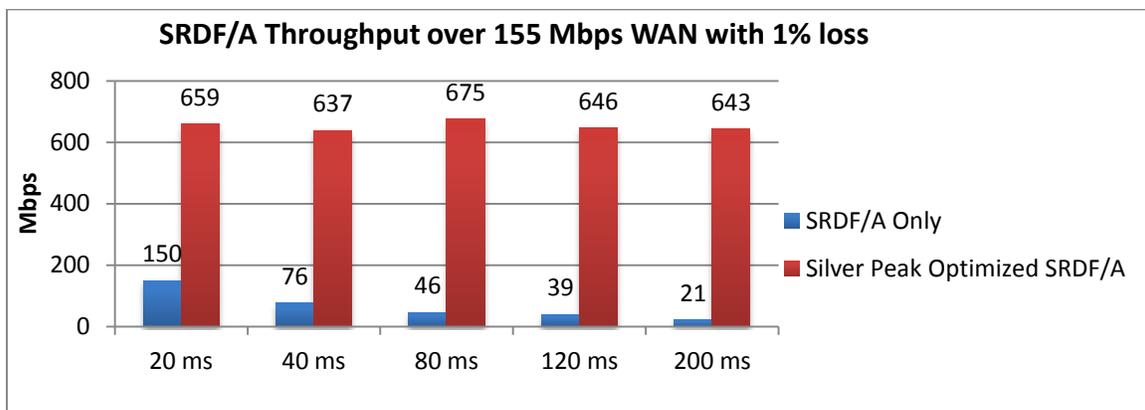


Figure 8: 155 Mbps WAN with 1% packet loss

When the WAN bandwidth is increased to 622 Mbps (figures 9 and 10) the un-optimized SRDF/A throughput starts at 592 Mbps with .1% packet loss and 20 ms of latency and then steadily decreases to 108 Mbps as latency is added. When the packet loss is increased to 1%, the unoptimized throughput drops to 276 Mbps with 20 ms of latency and continues to drop as latency is added, while the optimized throughput is maintained at an average of 1.3 Gbps on the same simulated WAN, regardless of packet loss.

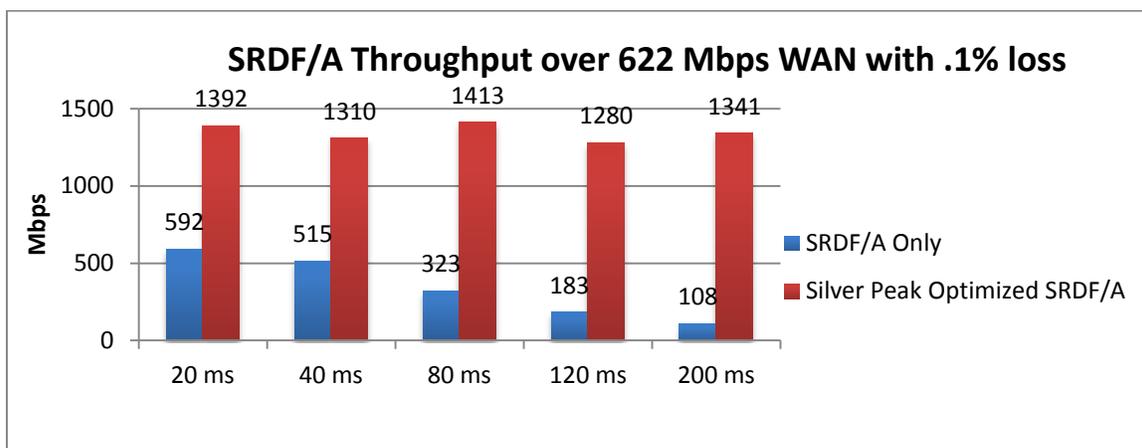


Figure 9: 622 Mbps WAN with .1% packet loss

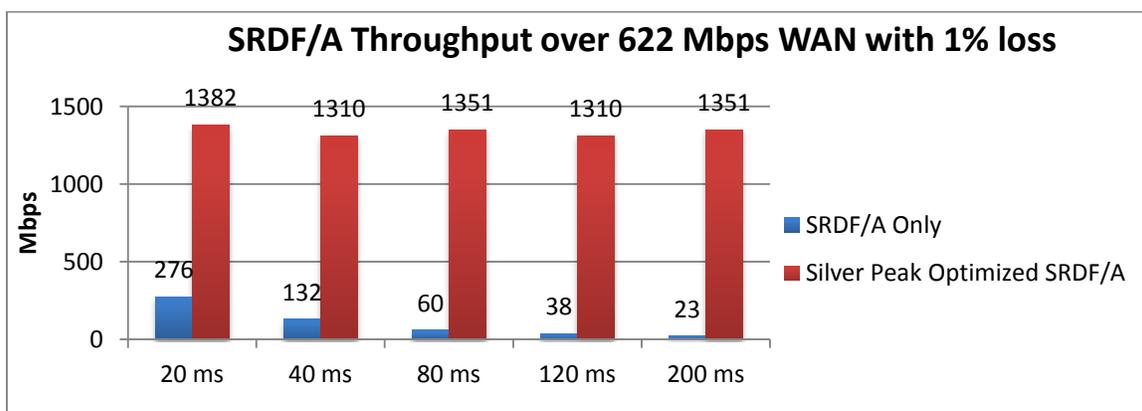


Figure 10: 622 Mbps with 1% packet loss

Finally, when bandwidth is increased to 1 Gbps (figures 11 and 12) SRDF/A throughput starts at 953 Mbps with 20 ms of latency and .1% packet loss and falls to 21 Mbps with 200 ms of latency and 1% packet loss, while the Silver Peak accelerated throughput is also consistent at 1.3 Gbps.

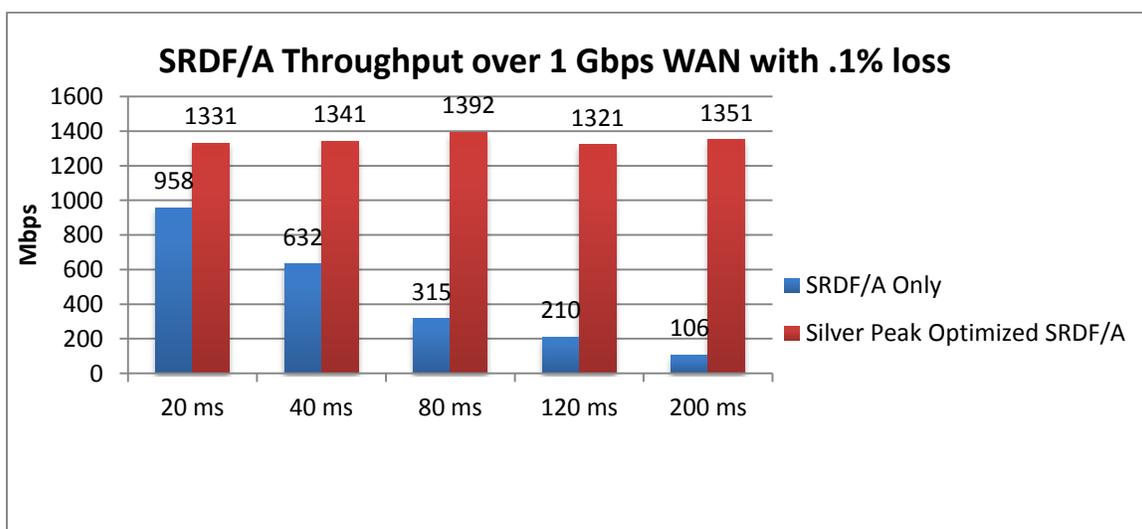


Figure 11: 1 Gbps WAN with .1% packet loss

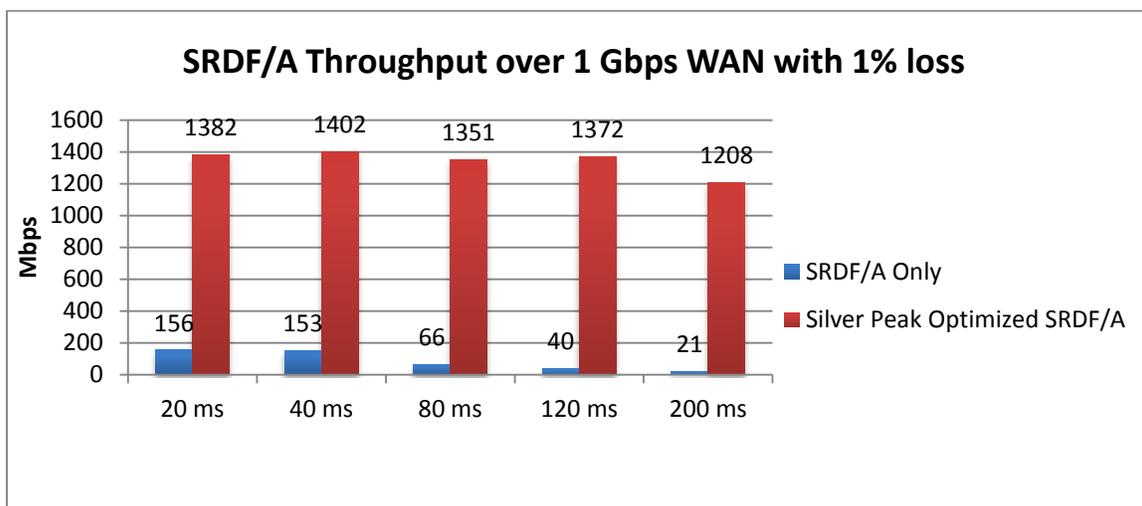


Figure 12: 1 Gbps WAN with 1% packet loss

These test results show a consistent increase in SRDF/A throughput when Silver Peak is used to optimize the WAN and accelerate the replication traffic. Actual throughput is going to be affected by traffic type, IO size, number of concurrent LUNs replicated, and the makeup of the data.

Jumbo Frames

Silver Peak VXOA has the ability to enable jumbo frame replication across typical WAN links. Every device that transmits data across a LAN or WAN has a limitation on the number of packets that it can process and send across the network. Jumbo frames give the device the ability to send more data inside every packet, increasing the overall throughput. In a typical WAN deployment, jumbo frames are not possible because every device in the network path must support the same size frames. When jumbo frames are enabled across a WAN that does not support them, the larger packets must be fragmented into a smaller size by routers or switches, severely reducing overall throughput.

Silver Peak Replication Acceleration enables jumbo frames across any WAN. This is accomplished by accepting jumbo frames from the LAN, where the Symmetrix systems would be located, and transmitting standard size frames across the WAN. The performance limitations resulting from fragmenting jumbo frames with normal network products, such as routers and switches, do not apply when the Silver Peak device is deployed. Silver Peak VXOA accepts jumbo frames and packages them into standard-size packets prior to transmission to the remote site. As with all other packets, the data is deduplicated and compressed.

The following chart (figure 13) illustrates the throughput of SRDF/A using jumbo frames and standard frames. A maximum throughput of 3 Gbps was achieved using standard 1500 byte frames. When the frame size was increased to 9000 bytes, throughput was increased to 5 Gbps. For both tests the Silver Peak appliances transmitted standard 1500-byte frames across the WAN.

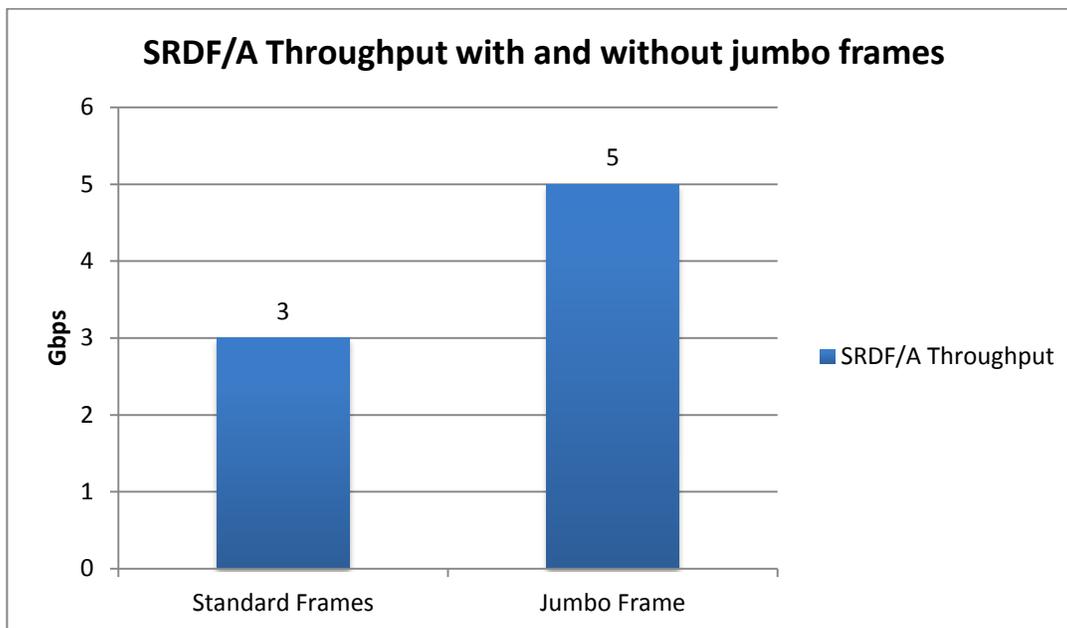


Figure 13: SRDF/A throughput with jumbo frames

Deployment Architectures

In all deployment scenarios the Ethernet fabric is key to performance. Replication traffic should be consolidated onto a single Ethernet switch where the Silver Peak appliance is also connected, Figure 14. Using a single switch eliminates the potential for packets to be dropped on the LAN over inter switch links between Ethernet switches. Ethernet speed also plays an important role in the effectiveness of replication acceleration.

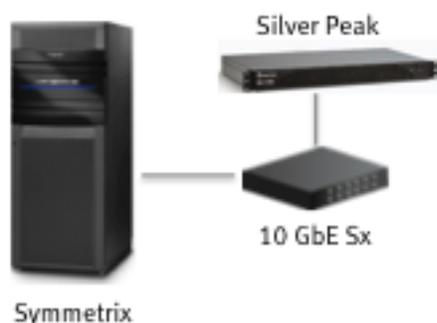


Figure 14: Same Ethernet Switch Deployment

Ethernet Fabric Speed

The speed of the Ethernet fabric and connections should be at least 6X the expected throughput of the WAN. The Silver Peak appliance must also have a connection to the Ethernet fabric that is 6X, or higher, the throughput of the WAN. For example, a WAN with 1 Gbps of bandwidth would require an Ethernet connection for the Silver Peak appliance of 10 Gbps. For this scenario multiple replication ports would be needed for a single Symmetrix when 1 GbE RE ports are used. When 10 GbE RE ports are used, a single port can be used due to the higher throughput.

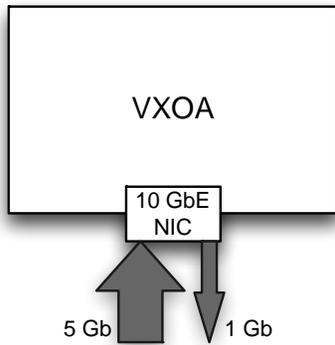


Figure 15: 1 Gbps WAN

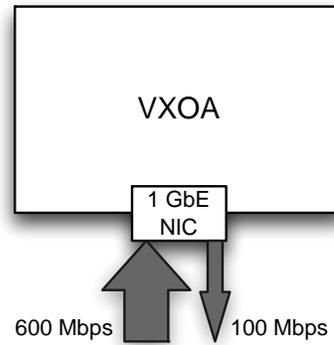


Figure 16: 100 Mbps WAN

When this requirement is not met, the throughput on the WAN will be lower because of the de-duplication and compression provided by Silver Peak VXOA. For example, if a 1 Gbps WAN is used and the Silver Peak appliance is connected to the Ethernet fabric with a 1 GbE port, SRDF/A throughput will be limited to approximately 867 Mbps. This is because the maximum throughput available to VXOA will be 1 Gbps. The 1 Gbps of available bandwidth will be shared between data sent to the VXOA and data transmitted by VXOA across the WAN. Because the bandwidth is shared, only 867 Mbps will be received by VXOA with the remaining bandwidth, 133 Mbps, being used to transmit the data across the WAN. When the Ethernet fabric speed is not high enough to support the requirements of VXOA, multiple Silver Peak appliances can be deployed with the SRDF/A traffic spread across them.

LAN Switch Performance Differences

The speed of the network switches at the primary and disaster recovery sites can have a major impact on the performance of accelerated SRDF/A replication traffic. Problems have been observed when the performance of one side of the DR network is faster than the other. For example, the primary data center has a 10 Gb Ethernet LAN connecting the Symmetrix replication ports to the Silver Peak appliance and the WAN router, while the disaster recovery site has a 1 Gb Ethernet LAN connecting the Symmetrix replication ports to the Silver Peak appliance and the WAN router.

In this type of environment, the speed of packets delivered at one site can cause the switch at the other site to have buffer size problems. This is due to the packets being optimized and delivered faster than the other switch can forward them to the Symmetrix array. When this happens a number of the packets will be dropped as the receiving buffer on the switch becomes full.

To remedy this problem, both the primary and disaster recovery site should have the same level of performance on the LAN.

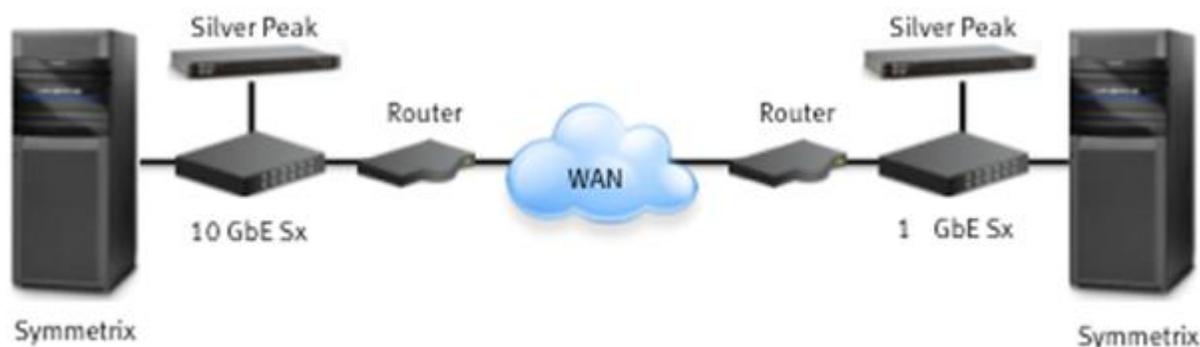


Figure 17: LAN Switch Performance Mismatch

WAN with multiple applications

In a typical enterprise environment, multiple applications use the WAN for replication and data transfer simultaneously. Sharing the WAN can put additional strain on the resources available for SRDF/A when optimization is not used. Adding Silver Peak to a shared replication environment allows control of bandwidth utilization through QoS. QoS policies set in VXOA only affect data flowing through the Silver Peak appliance. If VXOA is only being used to accelerate replication traffic and there is other traffic on the WAN, the other traffic will not be subject to any QoS policies in the Silver Peak appliance. For example, if 200 Mbps is dedicated to all shared replication traffic, the VXOA system throughput should be limited to 200 Mbps with all QoS policies managing this bandwidth. In this scenario QoS can be used to guarantee SRDF/A a minimum of 100 Mbps of bandwidth at all times. More information on configuring QoS can be found in the [Silver Peak Appliance Manager Operator's Guide](#).

Using a shared environment provides the additional benefit of being able to deduplicate replication data across multiple arrays, increasing the effectiveness. For example, data is stored on a Symmetrix array, replicated with SRDF/A, and backed up to a Data Domain array, replicated with Data Domain Replicator. Components of the data are also stored on an Isilon NAS system, replicated with SyncIQ. As the data is replicated across these multiple systems, the Silver Peak appliance will be able to deduplicate data across all of the replication traffic. As more replication applications are added to the WAN, like VPLEX Geo and RecoverPoint, the deduplication ratios will increase, resulting in higher replication throughput.

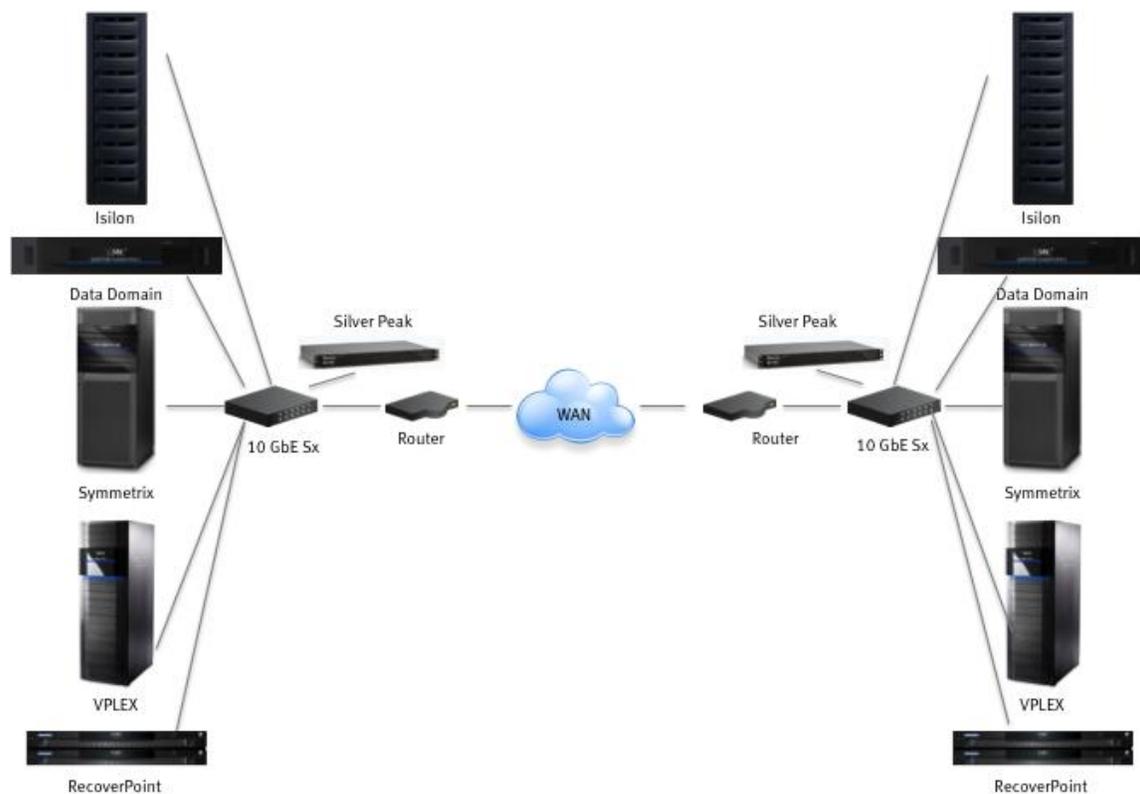


Figure 18: Consolidated Ethernet Fabric

Silver Peak Deployment Options

This section contains a high level overview of the various Silver Peak deployment options. Additional information is available in the [Silver Peak Network Deployment Guide](#).

Silver Peak Out-of-Path (Router Mode)

When an out-of-path deployment is used, the Silver Peak appliance is not in the direct path of network traffic. Because the appliance is not in the data path, a network-based redirection technique like Policy Based Routing (PBR), Web Cache Communication Protocol (WCCP), or Virtual Router Redundancy Protocol (VRRP), must be used to forward traffic to the appliance. The traffic forwarding is transparent to SRDF/A, and requires no changes to the Symmetrix array.



Figure 19: VXOA Out-of-path deployment

Silver Peak In-Line Deployment (Bridge Mode)

When an in-line deployment is used, the Silver Peak appliance is inserted in-line between the WAN router and the Ethernet switch. With this deployment method, the Silver Peak appliance will intercept all network traffic destined for the WAN. The appliance will accelerate traffic flows using the Silver Peak tunnel unless a rule has been created to pass through the traffic.

This deployment method does not require any changes to the network switch or WAN router. WAN downtime is required for the Silver Peak appliance to be inserted into the data path. In-line Silver Peak deployments are transparent to SRDF/A and require no changes on the Symmetrix array.



Figure 20: VXOA in-line deployment

An additional in-line deployment is possible where the Silver Peak appliances are connected directly to the Symmetrix RE ports. Connecting directly to the Symmetrix array is used in environments where there is high WAN bandwidth and the Ethernet fabric is limiting throughput. By deploying in-line, directly connected to the array, the SRDF/A traffic will be deduplicated and compressed before it is sent across the Ethernet fabric. By optimizing the SRDF/A traffic before the Ethernet fabric, the requirements on the fabric are reduced. For example, two Silver Peak appliances are deployed, each with a 1 Gbps connection directly to the Symmetrix array. SRDF/A will transmit 1 Gbps directly to the Silver Peak appliance, where it will be optimized. Using an average reduction of 6X, the 1 Gbps will be reduced to 166 Mbps. When both Silver Peak appliances are used, 2 Gbps of SRDF/A throughput will be reduced to 332 Mbps across the Ethernet fabric, well below the 1 Gbps limit of the switch ports.

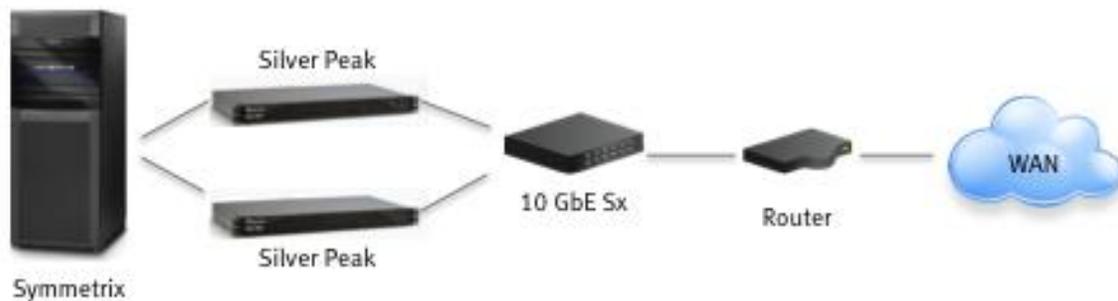


Figure 21: VXOA in-line deployment with multiple Silver Peak appliances

Silver Peak Velocity Deployment

Silver Peak Velocity is deployed as a virtual appliance on the network and uses a single IP Address. Velocity is designed as a Replication Accelerator and should be connected as close to the Symmetrix system on the network as possible, preferably on the same network switch. When the Silver Peak appliance is not connected to the same network switch, the SRDF/A traffic sent across the Ethernet fabric will be subject to any performance limitations or congestion on the fabric. No network changes are required for a Velocity deployment, as the default gateway of the RE will be changed to the Silver Peak appliance IP address. Changing the default gateway of the RE does require a change to the Symmetrix configuration file.

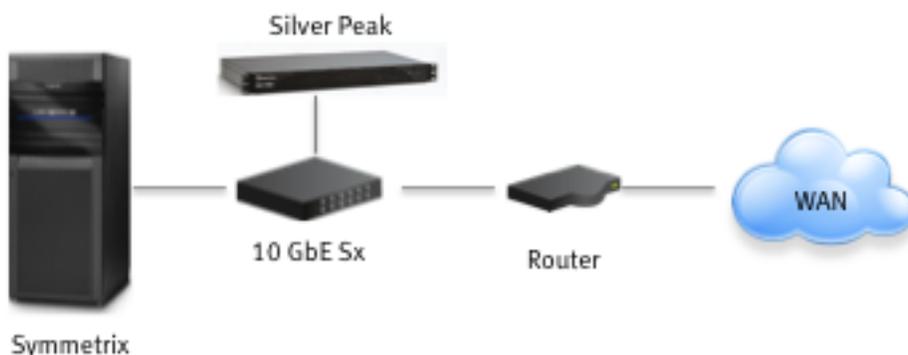


Figure 22: Velocity deployments

High Availability Silver Peak Deployments

Silver Peak appliances can be deployed in highly available configurations using either network redirection with out of path deployments, or Velocity deployments with multiple appliances. When an out of path deployment is used with WCCP, PBR, or VRRP, multiple Silver Peak appliances can be configured with the redirection method managing fail over.

A Velocity deployment uses a simple form to configure VRRP between two or more appliances. A virtual IP address is created that is shared among all of the appliances. In the event that an appliance fails, the virtual IP address will be moved to a surviving appliance.

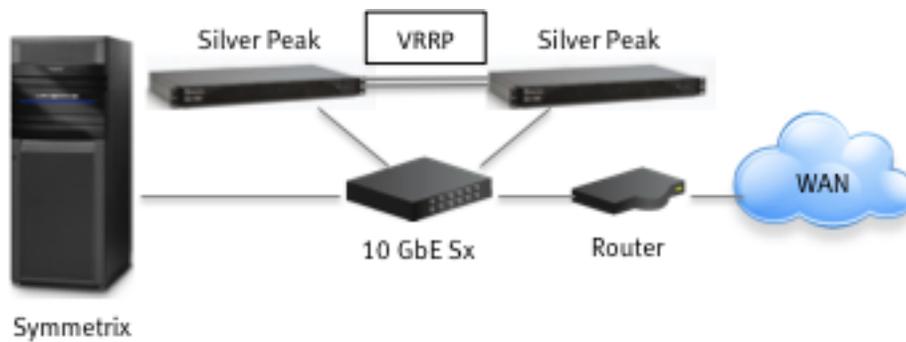


Figure 23: Silver Peak high availability deployment with VRRP

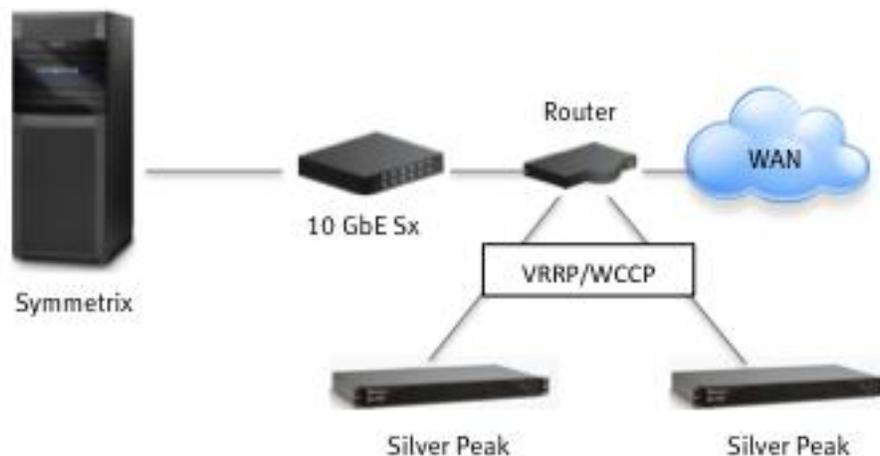


Figure 24: Silver Peak high availability deployment with WAN router

Deployment Best Practices

Jumbo Frames

Jumbo frames are best used when there is high WAN bandwidth, 1 Gb or greater, that cannot be fully utilized by SRDF/A. For environments where there is a lower amount of WAN bandwidth, less than 1 Gb, standard frames are recommended. When jumbo frames are used, all network devices connected to the Symmetrix RE ports and Silver Peak ports must support the same frame size. For example, if the RE port and Silver Peak appliance are connected to the same switch, the maximum frame size supported by the switch should be configured. If the switch supports 9000 byte frames, the RE ports and the Silver Peak appliance should be configured for 9000 byte frames.

When Silver Peak is used to enable jumbo frame support for SRDF/A across a WAN that only supports standard frames, a high availability deployment should be used for the Silver Peak appliances. If a high availability Silver Peak deployment is not used, SRDF/A replication traffic will stop if the Silver Peak appliance fails. This happens because jumbo frames are not supported across the WAN without the Silver Peak appliances optimizing the network and converting jumbo frames to standard frames.

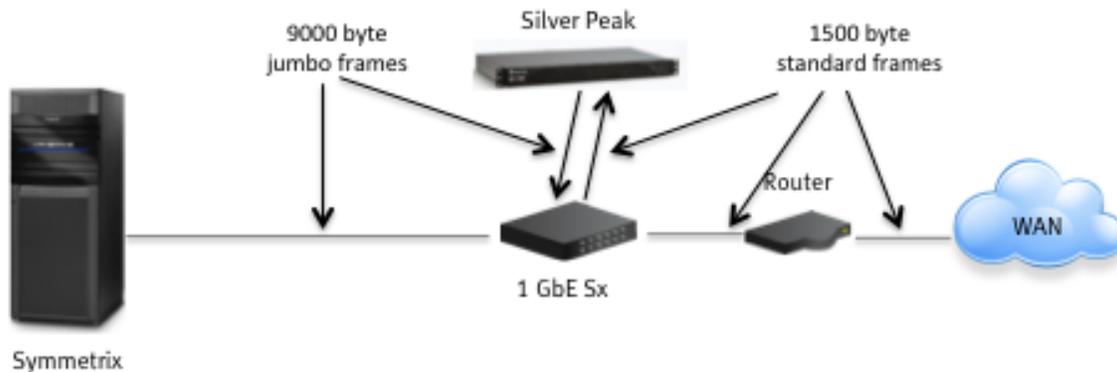


Figure 25: Jumbo frame deployments

SRDF/A Settings

The following items should be configured on all Symmetrix arrays that are being used for SRDF/A replication with Silver Peak VXOA.

Symmetrix 1 GbE and 10 GbE settings

- SRDF Flow Control = enabled
- SRDF Compression = enabled*

Any rate limits present on the RE ports should be disabled so that the maximum amount of data can be transmitted to the Silver Peak appliance for de-duplication and compression. If a rate limit is set on the RE port the maximum throughput of the Silver Peak appliance will be artificially limited. For example, if the network will support 100 Mbps of SRDF/A traffic, and the RE port is limited to 100 Mbps, the throughput will never exceed 100 Mbps even though the Silver Peak appliance is reducing the traffic by 6X. In this scenario, the data transmitted across the WAN by the Silver Peak appliance would be 16.66 Mbps (100/6). If the rate limit on the RE port is disabled, and set in the Silver Peak appliance, the data sent across the WAN will be 100 Mbps while the data sent from the Symmetrix will be 600 Mbps. The additional 500 Mbps will come from the deduplication and compression in the Silver Peak appliance. If the Silver Peak appliances are not deployed with high availability, HA, the WAN could be oversubscribed in the event of appliance failure. If the WAN becomes oversubscribed, SRDF/A replication can stop. HA Silver Peak deployments are recommended with SRDF/A.

For Silver Peak Velocity deployments, the default gateway for the RE ports should be set to the system IP address of the Silver Peak appliance.

*Note that the Silver Peak SRDF Protocol Acceleration feature will disable SRDF Compression. SRDF Compression should be enabled to provide some reduction on the replication data in the event that a Silver Peak appliance becomes unavailable.

Silver Peak VXOA Settings

The Silver Peak Quick Start Guide should be followed for the initial deployment. All Silver Peak documentation is located at <http://silver-peak.com/support/user-documentation>.

After deploying Silver Peak according to the Quick Start Guide, the following settings should be used for the tunnel configuration and the Optimization Policy.

Tunnel configuration

The tunnel that will be used to carry SRDF/A traffic should be configured with the following settings. The default values are noted to aid in troubleshooting and reconfiguration.

Forward Error Correction (FEC)

Forward Error Correction, FEC, is a tunnel option that is used to recover lost packets on the WAN in real time without requiring the data to be resent. Real-time packet loss can be measured in the Silver Peak appliance manager, figure 24.

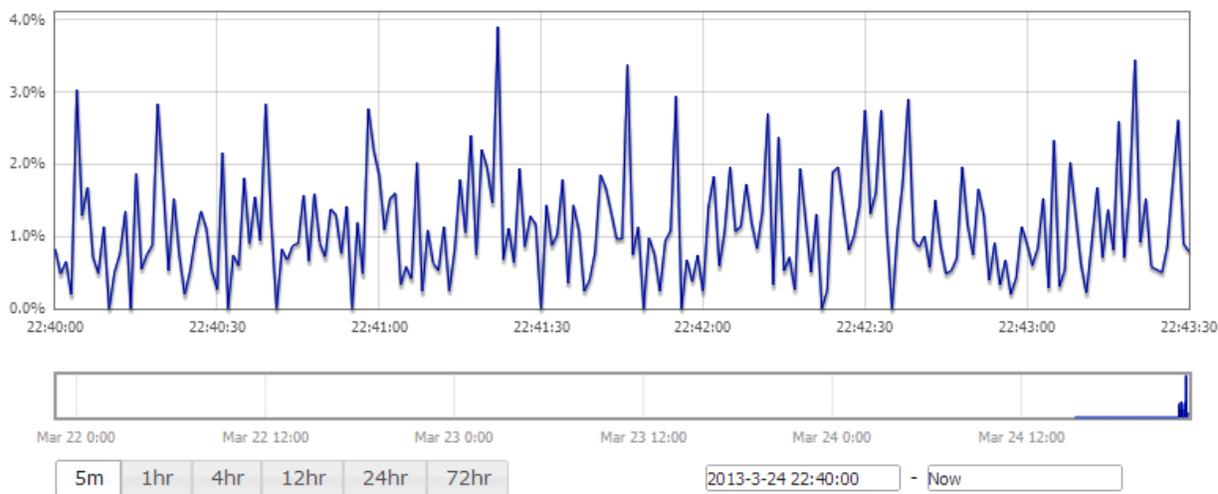


Figure 26: Monitoring loss with VXOA

When there is evidence of consistent packet loss on the WAN, a FEC setting of enable with a ratio of 1:5 is recommended. By using a FEC setting of enable with a ratio of 1:5, packet loss across the WAN will be corrected, increasing overall replication throughput. When packet loss is rare or intermittent, FEC should be set to auto with a ratio of 1:5. Auto will allow VXOA to dynamically control the amount of parity packets that are injected into the data stream, with the upper limit capped by the chosen Ratio value (1:5). For WANs that have zero packet loss, FEC can be set to auto or off. Off should only be used when there is no chance for bandwidth contention or dropped packets. When a setting of off is used, packet loss should be carefully monitored in the VXOA GUI.

Option	Value	Default Setting
Admin	Up	Yes
Auto Discover MTU	Checked	Yes
Auto Max BW	Checked	Yes
Mode	UDP	No
Reorder Wait	100	Yes
FEC with <.1% WAN packet loss	Auto 1:5	No
FEC with >.1% WAN packet loss	Enable 1:5	No

Table 8: VXOA tunnel configuration best practices

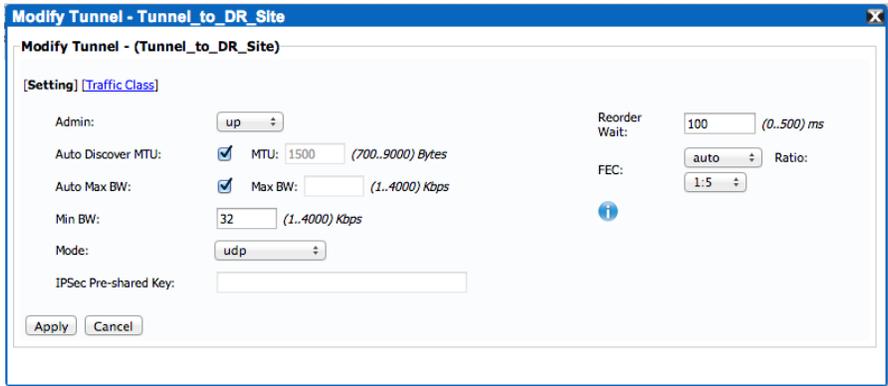


Figure 27: VXOA tunnel configuration screen

Optimization Policy

A new optimization policy must be added to enable the SRDF compression feature. After adding the new optimization policy, configure SRDF protocol acceleration and network memory according to the settings in table 9.

Note that the values for Advanced TCP Options should not be changed from the defaults without the assistance of a Silver Peak Technical Support Engineer.

SRDF Protocol Acceleration

SRDF protocol acceleration is used to disable the built-in compression on the GbE and 10 GbE directors. This prevents the need to change the configuration file to disable compression. If SRDF protocol acceleration is not enabled and SRDF compression is enabled on the director, the amount of deduplication and compression available from Silver Peak’s Network Memory will be reduced. When this happens the overall SRDF/A throughput will also be reduced.

SRDF protocol acceleration should be enabled.

Priority	ACL	Protocol	Src/Subnet	Dst/Subnet	Application	Src/Dst Port	DSCP	VLAN	Network Memory	Payload Compression	TCP Accel	Protocol Accel
<input type="checkbox"/> 10		ip	0.0.0.0/0	0.0.0.0/0	srdf		any	any.any	minimize latency	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	srdf

Figure 28: Optimization policy configuration screen

Option	Value	Default Setting
Protocol	ip	Yes
Application	srdf	No
Network Memory	Minimize latency	No
Payload Compression	Checked	Yes
TCP Acceleration	Checked	Yes
Protocol Acceleration	srdf	No

Table 9: Optimization policy best practices

Ethernet Switches

The LAN infrastructure at the source and target sites is very important to the performance of SRDF/A with Silver Peak VXOA. Enterprise class Ethernet switches should be used for all connectivity between the Symmetrix RE ports and the Silver Peak ports. Enterprise class switches are fully managed, have a redundant architecture, and have dedicated port buffers. When unmanaged

workgroup switches are used, there is a higher probability for packet loss on the LAN resulting in throughput problems.

Monitoring Performance Using the Silver Peak GUI

The Silver Peak user interface displays throughput based on 5 minute, 1 hour, 4 hour, 12 hour, 24 hour, or 72 hour intervals. SRDF/A uses cycles to transmit data; between cycles there is a gap where no data is sent. When the bandwidth chart period is set to a value higher than 5 minutes, a throughput average is used. The gaps between cycles will lower the average throughput for the time displayed, figure 26. To monitor peak throughput, use the 5 minute value for the chart period. The 5 minute value will display throughput peaks and the gaps between SRDF/A cycles, figure 29.

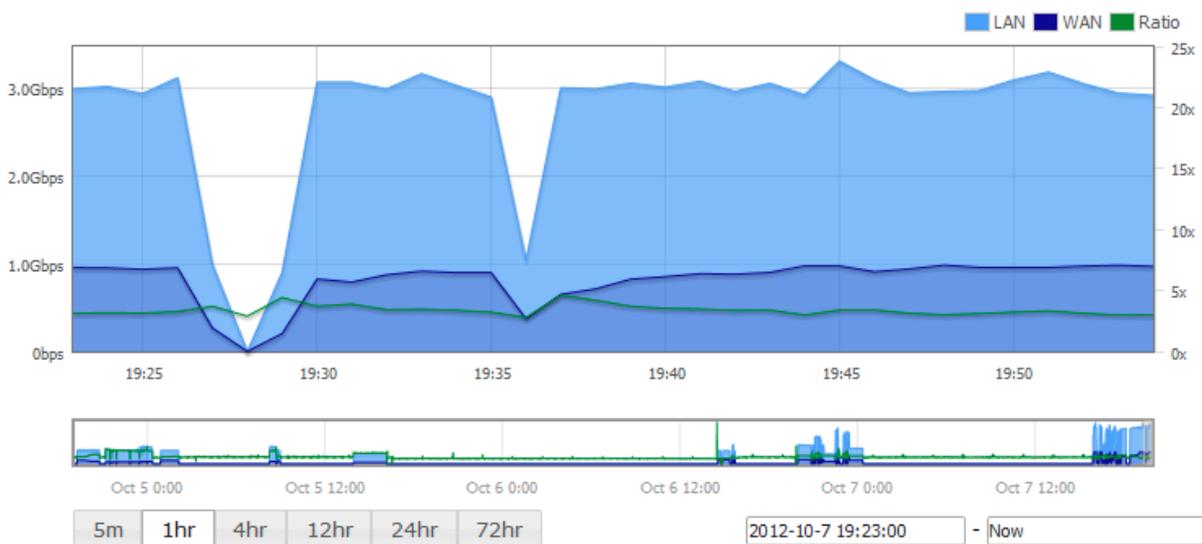


Figure 29: Silver Peak reporting 1 hour average

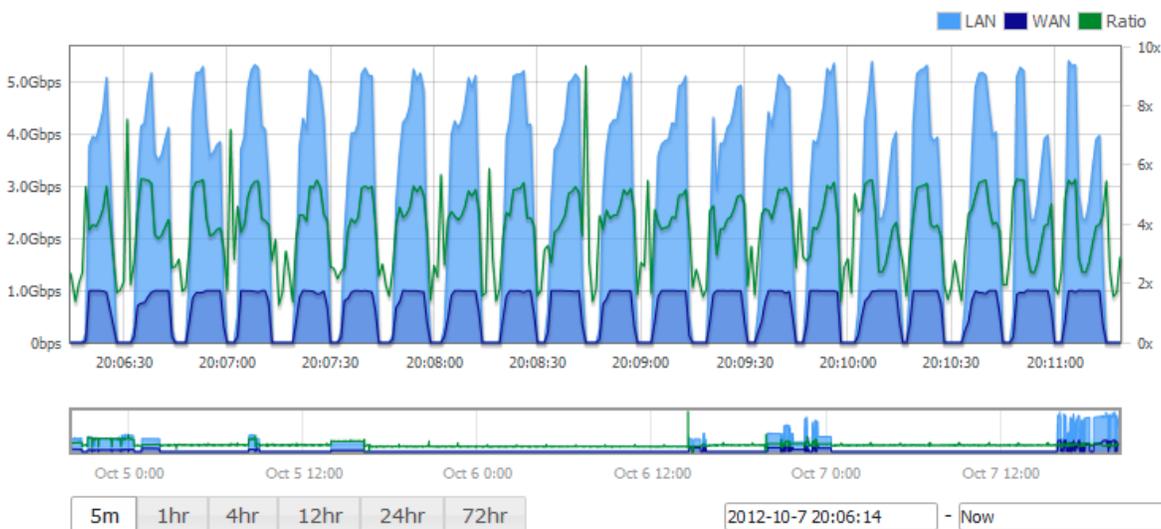


Figure 30: Silver Peak reporting 5 minute average

Conclusion

Silver Peak VXOA provides the features that enable SRDF/A replication across any distance and any WAN. Deployments that will fail without replication acceleration not only become possible, but also can be deployed successfully. Environments where the recovery point objective is not being met, or needs to be lowered, can quickly and easily be addressed with Silver Peak replication acceleration. Extensive testing has been performed by Silver Peak and EMC to verify compatibility and characterize the performance of VXOA and SRDF/A, and Silver Peak physical and virtual appliances are listed on the EMC Interoperability Support Matrix, (ESM). The ESM is available at elabnavigator.emc.com.

Appendix A

Links to Additional Documentation

EMC PowerLink – powerlink.emc.com

EMS ESM – elabnavigator.emc.com

Silver Peak - http://silver-peak.com/Support/user_docs.asp

Appendix B

Environmental Variables Affecting Performance

IO size and type

The size and type of IO greatly affects the throughput of SRDF/A across Silver Peak-optimized WAN connections. Most applications write data at a specific size, and these results illustrate the performance that can be expected based upon the IO size. As illustrated in Table 4, it is evident that as IO size increases, so does throughput.

To determine the impact of IO size and type on throughput, a workload was used that generated 10,000 IOPs across 56 LUNs. A 1 Gbps WAN was simulated between the VMAX arrays with Silver Peak NX-9700s optimizing the connection.

The throughput numbers included in this section are representative samples only and are used to illustrate the impact of IO size and type on SRDF/A throughput. These numbers are presented only to help determine the impact of the different traffic types on throughput in order to help with architecting and monitoring the environment.

IO Size and Type	SRDF/A Throughput
4KB random	593 Mbps
8KB random	670 Mbps
12KB random	1043 Mbps
16KB random	1015 Mbps
32KB random	1614 Mbps
64KB random	2328 Mbps
256KB random	2863 Mbps
512KB random	2913 Mbps
1024KB random	2954 Mbps
64KB sequential	2897 Mbps
128KB sequential	2932 Mbps
256B sequential	2976 Mbps
512KB sequential	2954 Mbps

Number of LUNs being replicated

As the number of LUNs replicated with SRDF/A increases, so does the throughput and reduction across the WAN. The configuration of the Symmetrix systems containing the R1 and R2 LUNs will also have an impact on performance. EMC best practices should be followed when configuring the systems that will be taking part in SRDF/A replication.

During testing an increase in throughput of 25% was observed when the number of LUNs was increased from 5 to 10. Throughput stayed consistent with up to 50 concurrent LUNs being replicated in the test environment. The amount of data reduction provided by Silver Peak Network

Memory increased as LUNs were added. The increased data reduction can be attributed to having a larger sample of data to use for deduplication.

Number of concurrent LUNs	SRDF/A Throughput	Reduction Ratio
5	2.5 Gbps	3.5
10	3.1 Gbps	4.4
15	3.1 Gbps	4.4
30	3.1 Gbps	5.5
40	3.0 Gbps	5.8
50	3.0 Gbps	5.9