Optimizing Dell Compellent Remote Instant Replay with Silver Peak Replication Acceleration

A Dell Technical White Paper

Abstract

This technical report details the benefits that Silver Peak's replication acceleration software provides for Dell Compellent Remote Instant Replay deployments. This report includes performance results from the Compellent test lab.

March 2013

Contents

Tables

Figures

Overview

Replication is fundamental to data protection and disaster recovery. Moving important data to a different geographic region helps protect business from natural and man-made disasters. In this paper we will review the problems that exist when replicating data over distance, how these problems impact Recovery Point Objectives (RPO) and availability, and how Silver Peak's replication acceleration software solves the problems.

Dell Compellent provides an efficient replication solution known as Remote Instant Replay, however, like all solutions that replicate data over a Wide Area Network (WAN), Remote Instant Replay is subject to the problems that exist when moving data over distance, specifically network quality and the amount of bandwidth.

By solving the problems that exist with replicating data over distance, Silver Peak enables Dell Remote Instant Replay to move more data, over a longer distance, and in less time. This paper also includes use cases, lab validation and best practices for deploying Silver Peak with Dell Compellent Remote Instant Replay.

Replication challenges

Data replication helps to address the core requirement of disaster recovery, safeguarding data against natural and man-made disasters. The key to being prepared for an event that requires the use of data at a remote site is maintaining a consistent RPO. RPO is easily defined as the amount of data that will be lost in the event of a disaster recovery event. For example, if the RPO is defined as 3 hours and a disaster is declared at 12pm, any data created between 9am and 12pm will be lost. Based on business or compliance requirements the RPO can be anything from days, to hours, minutes, or even seconds.

RPO is directly affected by the throughput of the data that is being replicated. If data is replicated over high-speed, dedicated links within a campus or a metro area, RPO is easily met. However, as replication is stretched over longer distance, three factors impact throughput and thus RPO:

- Distance (latency) can significantly affect replication performance because the longer the distance, the longer the remote DR site requires to respond to the primary site during the replication process. The impact of latency alone can be a factor of 10-20X slower replication.
- WAN quality (amount of packet loss) can also affect replication performance – even 0.1% loss (typical of MPLS WAN links) can decrease replication performance by 50-99%.
- Bandwidth is also a limiting factor for replication throughput. Simply put, if there is more data to replicate in a given period of time and bandwidth isn't sufficient, the RPO will be missed.

Silver Peak's replication acceleration software, deployed as either a virtual instance or physical appliance, quickly and easily repairs WAN quality problems, removes the impact of latency, and eliminates many bandwidth bottlenecks. The result is dramatically increased replication throughput for maintaining RPO's

Silver Peak Solution

The Silver Peak solution is built on the company's Virtual Acceleration Open Architecture (VXOA), which uses real-time optimization techniques to overcome bandwidth, latency, and packet loss issues that are common to most WANs. Silver Peak's optimization techniques are all performed in real-time and ensure maximum performance across the widest range of application and WAN environments.

Solution components

A joint Silver Peak and Dell Compellent solution includes a combination of hardware and software that has been verified to work together. Dell Compellent hardware is required to store and provide access to data. The Dell Compellent Storage Center software includes Remote Instant Replay to provide replication to a remote site. Silver Peak can be deployed as a physical or virtual appliance to meet the requirements of any data center. The solution can be deployed using any Dell Compellent hardware and any Silver Peak physical or virtual appliance.

Dell Compellent

Remote replication has come to play a major role in data protection over the past several years. However, many replication solutions still require duplicate arrays and expensive, high-speed data links. Never mind that the replication process itself can be slow and unreliable, especially when it requires the transfer of full-volume clones. Dell Compellent helps to cut the cost and complexity of remote replication, making disaster recovery a viable solution for organizations of all sizes.

Remote Instant Replay technology leverages space-efficient snapshots to provide a cost-effective solution for multi-site data protection. Following initial site synchronization, only incremental changes in data need to be replicated. This approach, also known as thin replication, not only reduces hardware costs, but minimizes bandwidth requirements and speeds recovery in the event of a disaster. Users can choose between synchronous and asynchronous replication options.

Silver Peak VXOA

Silver Peak virtual instances and physical appliances leverage the following Virtual Acceleration Open Architecture (VXOA) technology components to accelerate all enterprises applications in a secure and reliable way.

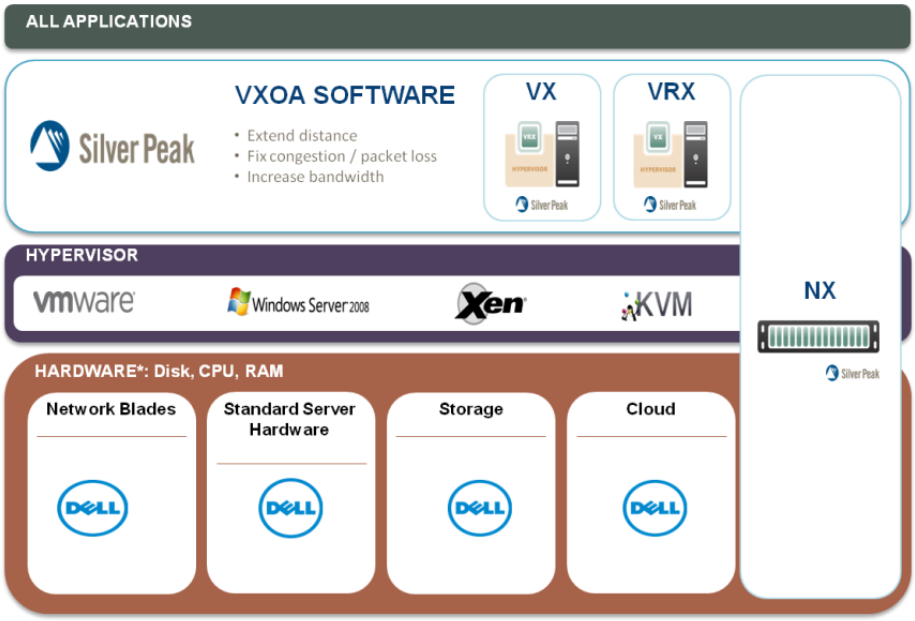

Figure 1. Silver Peak VXOA

Network Memory: Silver Peak's patented 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 devices. 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, Internet VPNs, and cloud. 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 due to distance by using various TCP acceleration techniques, like adjustable window sizing and selective acknowledgements, as well as CIFS acceleration techniques, such as read-aheads and write-behinds. These tools help to overcome inherent chattiness that can otherwise hamper application performance across a WAN

Secure Content Architecture™: Silver Peak keeps enterprise data secure with its Secure Content Architecture. All devices are equipped with hardware based AES encryption to protect local data stores. Optional IPsec keeps content safe when it traverses the WAN. Intelligent security policies can be established for granular control of WAN traffic.

Global Management System: Global Management System (GMS) is a powerful platform for the deployment, management, and monitoring of a Silver Peak-enabled Wide Area Network (WAN). GMS gives IT manager's detailed visibility into all aspects of application delivery across a distributed enterprise, including WAN performance statistics, application analysis, and tools for the configuration and management of Silver Peak appliances.

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

Silver Peak offers two separate appliance product lines:

- VX-Series are for general purpose WAN optimization and are sized by WAN bandwidth
- VRX-Series are replication-specific and are sized by throughput (GB or TB/hr) to meet a RPO

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

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

Table 1. Virtual solution options

Solution validation

Silver Peak and Dell worked collaboratively to test and validate the impact of replication acceleration on Remote Instant Replay traffic. During the testing, two Compellent arrays were used to simulate a primary and disaster recovery site. The Compellent arrays were connected to a 10 Gbps Ethernet switch via iSCSI cards in the arrays. A WAN emulator was used to simulate various bandwidth constraints, latency, and packet loss between the two sites. Two Silver Peak VRX-8 virtual appliances were used to optimize the simulated WAN between the primary and disaster recovery site.

Tool/Device	*Name*	*Versions*
Primary Compellent Array	SC8000 Controller	SCOS 6.2.1.107
Secondary Compellent Array	Series 40 Controller	SCOS 6.2.1.107
Replication Software	Remote Instant Replay	
Silver Peak Appliance	VRX-8	
Silver Peak Software Version	VXOA v.5.2.2	VXOA 5.2.2_42396
WAN Emulator	KWANEM	FC6

Table 2. Solution validation equipment

Test methodology

During testing, an environment was configured that allows multiple WAN sizes and types to be simulated so that the performance benefits of VXOA with Remote Instant Replay can be measured against a non-optimized baseline. A traffic generator is used to write a dataset to the Compellent array that is a representative mix of small, medium, and large files. A certain percentage of the files were then edited to measure the difference between an initial sync and an incremental sync. After setting a baseline of non-optimized performance, Silver Peak VXOA was enabled to set an optimized throughput.

The following WAN bandwidths, latencies, and loss variables were used during the testing:

- WAN bandwidths of 45 Mbps, 155 Mbps, and 1 Gbps
- Round trip latencies of 0, 40, and 80 ms
- Packet loss of 0%, 0.1%, and 1%

The latency and packet loss values were chosen to help simulate a variety of networks with an increasing geographical separation between data centers. A latency of 0ms is used to set a best case scenario for each bandwidth while also measuring the benefits of Silver Peak's Network Memory technology. 40 ms of latency is used to simulate a WAN connecting Washington DC to Phoenix, Arizona, while the 80 ms of latency is used to simulate a WAN connection San Francisco to New York City.

The following table lists common WAN types and their average packet loss.

WAN Type	Average Packet Loss
Private Line	0%
MPLS	.1%
Internet VPN	1%
Satellite	5%

Table 3. Common WAN average packet loss

The following table displays average latencies between geographic regions. Latency will vary based on the actual distance of the telecom infrastructure and can be significantly higher.

Distance	Average Latency
Metro	5ms
Regional US	25ms
Coast to Coast US	80ms
US to Europe	150ms
US to Asia	300ms

Table 4. Geographic average latency

Test results

During testing, an improvement in Remote Instant Replay throughput of 5X was seen with as little as 20ms of latency. As latency was increased the throughput benefit became more pronounced, yielding a 10X or greater benefit at 300ms of latency.

As the quality of the WAN link was degraded (increased loss), the disparity in throughput became even greater. Both the 40ms and 80ms WAN with .1% packet loss had an average 12X performance increase. Finally, when 1% packet loss was added to a 40ms or 80ms WAN the average throughput increases by 10X.

As Silver Peak fixed the errors on the WAN, Remote Instant Replay was able to use all of the bandwidth available. At the same time, Silver Peak Network Memory increased the size of the virtual bandwidth across the WAN. For most Remote Instant Replay installations Network Memory will be able to reduce the data sent across the WAN by 4-6X. This reduction results in more data being sent across the same amount of bandwidth, creating a virtual bandwidth increase.

Use cases

The following use cases are based off of data that has been collected during joint Silver Peak and Dell testing as well as an average of data from installed customers.

Regional replication across an Internet VPN

Replicating across an Internet VPN connection is becoming more common as WAN bandwidth is shared between multiple applications and end user traffic. Having a single Internet connection that is also used for Remote Instant Replay traffic can be a cost effective solution for small and mid-sized businesses.

Without Silver Peak optimizing the connection, throughput would be severely limited due to the amount of lost and out-of-order packets that are present across Internet connections.

For this use case we will use the following assumptions:

- WAN
  - Bandwidth limit of 45 Mbps
  - Latency 40ms – the equivalent of Atlanta to Austin
  - 1% packet loss – the average loss across an Internet VPN connection
- Data Set
  - 5TB – mix of user data, databases, email, and other applications
  - 10% daily change rate between the hours of 8am and 5pm
  - 4 hour RPO

For most Internet VPN replication deployments it can be very difficult, or impossible, to maintain a connection between the replication partner arrays. For this scenario we assume that the arrays are constantly connected during the time that replication is happening.

Our 5TB dataset with a 10% change rate will have an average of 512 GB per day of data that needs to be replicated. In most scenarios the change rate will vary during the day creating periods when more data must be replicated to maintain the specified RPO. To make the math simple we will assume a uniform change rate across 8 hours resulting in 64 GB of data per hour that needs to be replicated.

- Without Silver Peak the average throughput will be 8.2GB per hour
  - The 64GB per hour of changed data will take approximately 7.5 hours to replicate
  - The 4 hour RPO will be **missed by 3.5 hours** during the period when data is changed

- With Silver Peak VXOA the average throughput will be 71.1GB per hour – an 8.6X increase
  - The 64 GB per hour of changed data will take approximately 54 minutes to replicate
  - The 4 hour RPO will be **maintained and reduced to .9 hours**

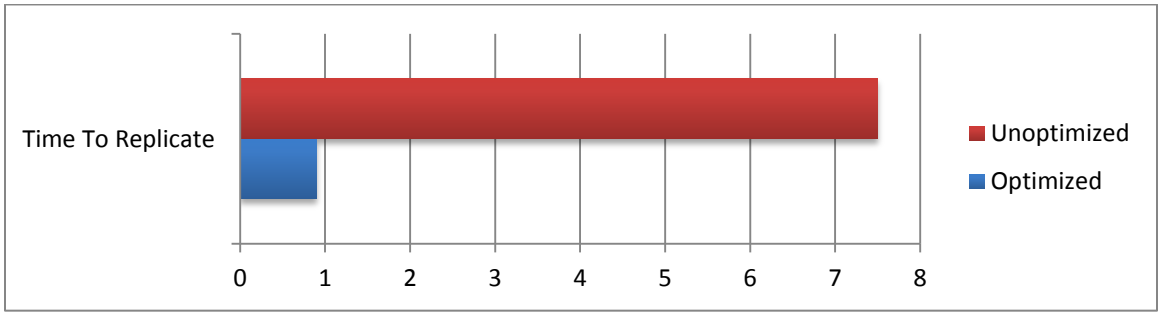

Figure 2. Internet VPN replication time with and without Silver Peak VXOA

For this use case Silver Peak provides immediate value by repairing the quality issue of replication across the Internet. Loss and out-of-order packets are repaired in real-time, resulting in a connection that is stable and reliable, while Network Memory provides virtual bandwidth that makes the 45 Mbps WAN behave like a 150 Mbps WAN.

Coast to coast replication in the USA across an MPLS network

MPLS connections provide good value for the money compared to private or leased lines. However, MPLS WANs have and inherent packet loss of approximately .1% and can have high out-of-order ratios.

When replication is dropped, or delivered consistently out-of-order, it must be retransmitted across the WAN. Retransmits increase the latency, reduce the available bandwidth, and decrease the throughput of replication. Even with the potential for loss and out-of-order MPLS still provides a good value for all businesses. Without Silver Peak, throughput will be limited due to a combination of latency and packet loss.

For this use case we will use the following assumptions:

- WAN
    - Bandwidth limit of 155 Mbps
    - Latency 80 ms – the equivalent of replicating from New York to San Francisco
    - .1% packet loss – the average loss on an MPLS connection
- Data Set
    - 20 TB – mix of user data, several databases, application data, email data
    - 10% daily change rate between the hours of 8am and 5pm
    - 8 hour RPO

Our 20 TB dataset will have an average of 2 TB per day that needs to be replicated based on a 10% change rate. In most scenarios the change rate will vary during the day creating periods when more data must be replicated to maintain the specified RPO. To make the math simple we will assume a uniform change rate across 8 hours resulting in 250 GB of data per hour that needs to be replicated.

- Without Silver Peak VXOA the average throughput will be 9.61GB per hour
    - The 250 GB of changed data will take approximately **26 hours** to replicate
    -
- With Silver Peak VXOA the average throughput will be closer to 300GB per hour
    - The 250 GB of changed data will take approximately **1 hour** to replicate

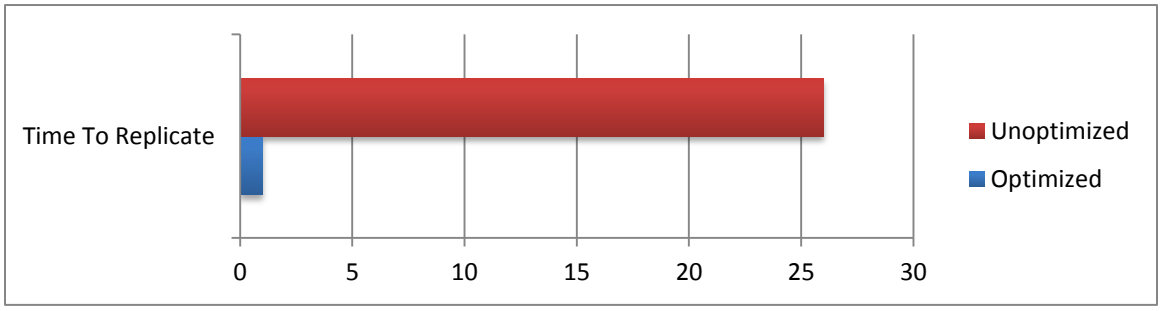

Figure 3. MPLS time to replicate with and without Silver Peak VXOA

For this use case the performance increase can be attributed to Silver Peak's ability to repair all loss and out-of-order on the WAN, making the low **cost MPLS connection behave like a high cost private or leased line.**

Simply repairing the quality of the WAN will allow Remote Instant Replay to take advantage of the entire 155 Mbps of available bandwidth. Silver Peak Network Memory is able to provide a virtual bandwidth increase of 4 – 6X, making the 155 Mbps WAN appear to be a 622 – 930 Mbps WAN.

International replication using a private line

While private lines have a higher cost compared to MPLS or Internet VPNs, they also provide a stable and reliable connection that is typically not subject to packet loss or out-of-order. In this scenario throughput will be limited due to the high amount of latency, throughput will also be limited to the speed of the WAN. For this use case the primary benefits that Silver Peak VXOA will provide are mitigating the impact of latency on the Remote Instant Replay traffic and providing higher virtual bandwidth when compared to the actual WAN speed.

For this use case we will use the following assumptions:

- WAN
  - Bandwidth limit of 1 Gbps
  - Latency 100ms – the equivalent of replicating between New York and London
  - 0% packet loss
- Data
  - 45TB
  - Primarily database and application data, including email, low amounts of user data
  - 2 hour RPO
  - 1 iSCSI port in each array used for replication

The 45 TB data set will have an average change rate of 10% resulting in 4.5TB of changed data to be replicated daily. We will assume that the change rate is spread across a 12 hour period due to the international nature of the business. During the 12 hour period 375GB of data will be changed per hour and require replication in 2 hours to the remote site.

- Without Silver Peak the average throughput will be 112.4GB per hour
  - The 375 GB will take approximately **3.4 hours** to replicate
  -
- With Silver Peak the average throughput will be 295GB per hour
  - The 375 GB will take approximately **1.9 hours** to replicate

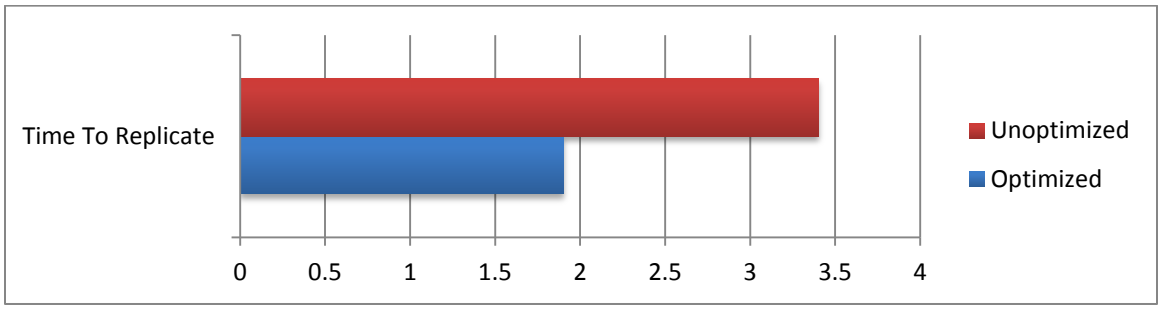

Figure 4. Private line time to replicate with and without Silver Peak VXOA

In this case, latency is the primary factor limiting the throughput. Without Silver Peak the 2 hour RPO is impossible to maintain without adding additional bandwidth, resulting in a recurring monthly cost.

Silver Peak is able to reduce the impact of latency on the Remote Instant Replay traffic and increase throughput by almost 3X resulting in an RPO of 1.9 hours. Network Memory provides additional benefit in this scenario by increased the amount of virtual bandwidth available to Remote Instant Replay. Because this is a high-bandwidth environment throughput will be limited based on the number of

volumes that are being replicated at any given time and the number of iSCSI ports that are being used for replication. For this environment a 10 GB LAN infrastructure is recommended to allow the Dell Compellent array to transmit more data to the Silver Peak appliance, resulting in a throughput higher than the 1 GB WAN limit.

When to deploy Silver Peak with Remote Instant Replay

Silver Peak's replication acceleration software should be deployed with Dell Compellent Remote Instant Replay in the following scenarios:

- Latency across the WAN is greater than 20 milliseconds
- The WAN has intermittent, or consistent, packet loss of 0.1% or greater
- The WAN is shared with multiple applications creating the potential for congestion
- Replication is being constrained by the size of the WAN and the available bandwidth
- Any combination of the above scenarios

Deployment architectures

Silver Peak NX appliances and VX software can be deployed to optimize all IPtraffic. Or, VRX software can be deployed solely for replication acceleration. The deployment method used does not impact the performance improvement that VXOA provides for Remote Instant Replay. For more information on Silver Peak deployment options see the Silver Peak Network Deployment Guide located at http://www.silver-peak.com/Support/user_docs.asp.

In-path

For in-path (also known as bridge mode) deployment, the Silver Peak device will sit in line between a network switch and the WAN router. An inline deployment does not require any changes to be made on the network switches or WAN router. In the event that an appliance fails, it will simply go into hardware bypass mode and all traffic will continue to flow as if the appliance was never there. Bridge mode is the simplest deployment option and requires the least amount of configuration.

In the diagram below the Silver Peak appliance is deployed between the network switch and WAN router. All traffic that is sent to the WAN router will travel through the Silver Peak appliance.

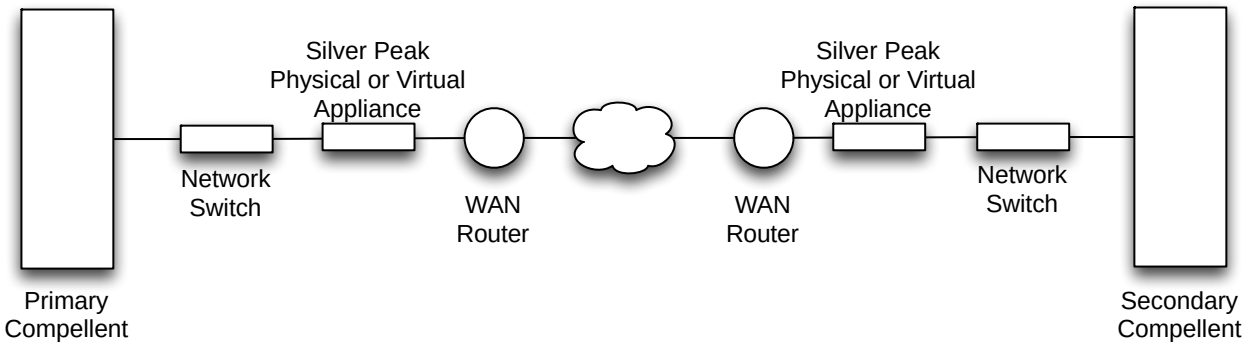

Figure 5. In-path deployment architecture

Routed mode

With routed mode the Silver Peak device is not in the direct path of network traffic, and traffic must be redirected to the device by a router or switch. Routed mode deployments are more robust and have

the benefit of supporting a more varied set of failure recovery mechanisms. A routed mode deployment does require configuration of the WAN router to redirect traffic to the Silver Peak device.

In the diagram below network traffic is redirected to the Silver Peak device by the WAN router. After the traffic is optimized, the Silver Peak device will forward the traffic to the remote site, for replication the remote site would be the DR data center.

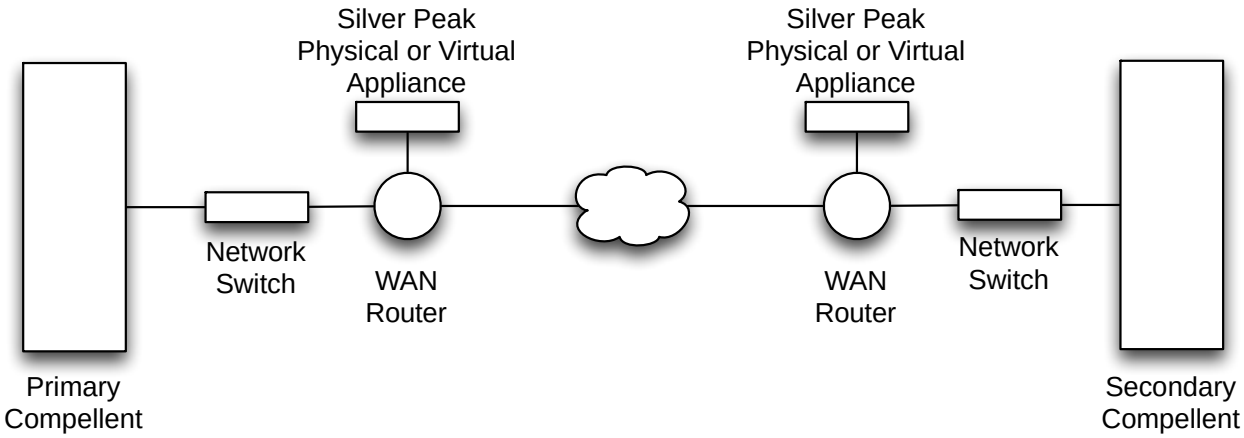

Figure 6. Routed mode deployment architecture

Storage forwarding

Silver Peak software can also be deployed by redirecting replication traffic from the storage array. In the Compellent Remote Instant Replay environment, this is done by changing the default gateway setting on the iSCSI ports that are used for replication. The IP Address of the Silver Peak device will be used as the default gateway for the iSCSI ports that are performing replication. For more information, a deployment guide is available at : http://www.silver-peak.com/products-solutions/velocity-dell-compellent

Using this mode to optimize Remote Instant Replay requires no changes to the network and therefore has no deployment impact.

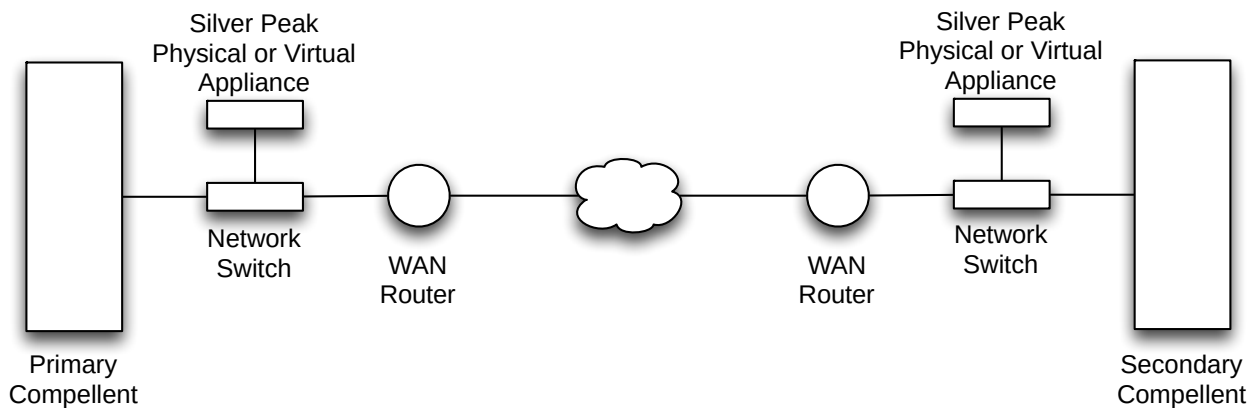

Figure 7. Velocity mode deployment architecture

Recommended configuration options

When Silver Peak is deployed to optimize Dell Compellent Remote Instant Replay replication there are several items that can be tuned to improve performance and provide the maximum throughput. While these changes are optional they are highly recommended.

Minimum Versions:

- Compellent SCOS 6.2.1.107 or higher
- Silver Peak VXOA 5.0 or higher

Dell Compellent Configuration Recommendations:

1. Remove any QoS definitions for Remote Instant Replay replications optimized by Silver Peak. If a QoS definition exists the full benefit of Silver Peak's Network Memory will not be able to be utilized and throughput will be limited by the Compellent QoS setting. If throughput needs to be limited across the WAN this can be configured in the Silver Peak appliance.
2. Increase the Window Size on the iSCSI adapter that is being used for replication. During testing a significant performance increase was seen when the Window Size was increased to the maximum setting available on the iSCSI adapter.
3. Disable de-duplication on the Compellent array. The de-duplication feature on the Compellent array tracks the changed blocks in each replay and only send these changes instead of the entire replay. Disabling this feature allows the Silver Peak appliance to de-duplicate all data across the WAN and also reduces the CPU workload on the Compellent system.

Silver Peak Configuration Recommendations:

While Silver Peak will work well with default settings, there are certain situations where one might wish to adjust the default configuration

1. Configure Network Memory to use Minimize Latency. Minimize Latency will provide the highest throughput and response time for Remote Instant Replay traffic. If a higher level of data reduction is needed Balanced Mode can be used.
2. Configure FEC to 1:5 when packet loss is consistent. Using a static setting for FEC will use some WAN bandwidth for parity but the impact of packet loss will be negligible. This will also provide a consistent throughput that can be used for reliable RPO planning.

Summary

Silver Peak offers a proven replication acceleration solution that is easily deployed with Dell Compellent Remote Instant Replay. The combined Dell Compellent and Silver Peak solution helps provide better utilization of WAN bandwidth, while dissolving distance and quality limitations. This results in lower RPOs and better disaster preparedness.

In summary, optimizing Dell Compellent Remote Instant Replay with Silver Peak provides the following benefits for customers who are replicating over longer distances, with limited bandwidth, or lower quality WAN links (or all of the above)

- Replication throughput can be increased up 20X or greater
- WAN bandwidth requirements can be reduced by 5-7X
- Utilize lower cost, lower quality bandwidth while maintaining high performance
- Replicate over longer distances with consistent throughput to maintain RPO