Technical Report

Silver Peak WAN Optimization and Dell EqualLogic PS Series

Abstract

This technical report details the benefits that Silver Peak's WAN optimization appliances provide for Dell EqualLogic Auto-Replication deployments. This report includes performance results from the EqualLogic test lab.

July 2012

Preface

Dell EqualLogic PS Series arrays optimize resources by automating performance and network load balancing. Additionally, PS Series arrays offer all-inclusive array management software, host software, and free firmware updates.

Audience

The information in this guide is intended for storage administrators that are involved in the planning and deployment of Dell EqualLogic storage arrays.

Related Documentation

For detailed information about PS Series arrays, groups, volumes, array software, and host software, log in to the Documentation page at the customer support site.

Dell Online Services

You can learn about Dell products and services using this procedure:

1. Visit http://www.dell.com or the URL specified in any Dell product information.

2. Use the locale menu or click on the link that specifies your country or region.

Dell EqualLogic Storage Solutions

To learn more about Dell EqualLogic products and new releases being planned, visit the Dell EqualLogic TechCenter site: http://delltechcenter.com/page/EqualLogic. Here you can also find articles, demos, online discussions, technical documentation, and more details about the benefits of our product family.

Table of Contents

Revision Information v
Introduction 1
Replication Challenges 1
Solution Components 4
    EqualLogic PS Series Storage Array 4
    EqualLogic Auto-Replication 4
    Silver Peak VXOA 4
Deployment Options 5
Testing Overview 8
Test Environment 9
Test Process 10
Test Results 10
Replication Challenges 11
    Packet Loss 12
    WAN Latency 13
    Limited Bandwidth 14
Conclusion 16

Revision Information

The following table describes the release history of this Technical Report.

Report	Date	Document Revision
1.0	July 2012	Initial Release

The following table shows the software and firmware used for the preparation of this Technical Report.

Vendor	Model	Software Revision
Dell	EqualLogic PS6000XV	5.1.1h2
Dell	EqualLogicPS6510E	5.1.1h2
Silver Peak	VX-5000	4.4.10

The following table lists the documents referred to in this Technical Report. All PS Series Technical Reports are available on the Customer Support site at: *support.dell.com*

Vendor	Document Title
Dell	Dell EqualLogic Auto-Replication: Best Practices and Sizing Guide
Silver Peak	Network Deployment Guide
Silver Peak	Appliance Manager Operator's Guide
Silver Peak	Deploying Silver Peak VXOA with a Dell EqualLogic Isolated IP SAN

INTRODUCTION

How much data can your business afford to lose?

The answer to that question should drive the design of any disaster recovery solution, particularly one that replicates data over a WAN. Dell EqualLogic Auto-Replication periodically replicates data from primary sites to designated disaster recovery sites, and Silver Peak appliances help optimize that replication across the WAN. Should any outage occur, IT managers can leverage these secondary sites as interim data centers, utilizing this remote data until primary sites recover, to help maintain high business continuity with minimal loss. When Silver Peak VXOA is deployed with Dell EqualLogic Auto-Replication, you can move more data, over a longer distance, and in less time, than trying to replicate without WAN optimization.

This paper will illustrate the challenges of replicating data across a WAN, and how Silver Peak can help businesses overcome the challenges and meet their requirements for data protection and disaster recovery.

REPLICATION CHALLENGES

Disaster recovery strategies address the need to safeguard data against natural disasters, human error and malice. Enterprises can manage against these risks by replicating primary data to secondary sites. However, replication across a WAN can be adversely impacted by three things — packet loss, latency, and bandwidth. In most cases, it is a combination of these three factors that degrades replication performance. Many times, bandwidth is added to increase replication throughput, when the real issue is the quality of the WAN, or simply the distance between data centers, that is causing the performance problem.

A WAN that has even a small amount of packet loss (0.1 percent), can dramatically reduce replication performance. Similarly, as latency across the WAN increases, throughput will also decrease. When latency is combined with packet loss, the reduction in performance is even higher. The Dell EqualLogic Auto-Replication Best Practices and Sizing Guide describes in more detail how these factors influence EqualLogic replication.

Bandwidth can also play a large part in replication performance. When latency is low and there is no packet loss, bandwidth will be the limiting factor for replication throughput. A 10 Mbps WAN will only transmit a maximum of 10 Mbps, far less than the EqualLogic arrays are capable of sending.

PACKET LOSS AND LATENCY

Packet loss is typically seen on congested WANs, MPLS networks, satellite connections, and Internet VPN connections. When data is replicated across a WAN and there is packet loss, throughput is reduced due to the retransmission of the dropped packets and the congestion avoidance mechanism of the Transmission Control Protocol (TCP). TCP slows the amount of data that is being transmitted to reduce congestion and help manage the packet loss. The reduced throughput typically results in missed Recovery Point Objectives (RPOs) and puts business data in danger of not being available in the event of a disaster.

WAN Type	*Average Packet Loss*
Private Line	0%
MPLS	0.1%
Internet VPN	1%
Satellite	5%

Table 1: Common WAN Loss

For a business that is trying to save money by using low cost MPLS networks, 0.1 percent packet loss can be the difference between meeting a 4-hour RPO and being forced to accept a 24 or 36-hour RPO. Packet loss and congestion on the WAN reduce the throughput enough that the available bandwidth in unusable.

The impact of packet loss on throughput can easily be calculated using the Mathis formula, which calculates maximum throughput by dividing the Maximum TCP Segment Size (MSS) by the round trip time (RTT), multiplied by the square root of the probability of packet loss.

$$\text{Throughput} = MSS/(RTT\sqrt{P_{loss}})$$

In *Figure 1* it is easy to see that 0.1 percent packet loss might not sound like a lot, but it will have a dramatic impact on throughput, especially if there is latency on the WAN. A 100 Mbps WAN with 0.1 percent loss and 50 ms of latency will result in a throughput of 7.4 Mbps, far less than the 100 Mbps of available bandwidth.

When latency is present on a WAN, throughput is reduced due to the default behavior of TCP. TCP throughput is governed by the size of the receive window (RWIN). The RWIN controls the amount of data that can be sent over the WAN before the sender receives an acknowledgement. As the distance between data centers increases, more data is required to be sent across the WAN at any moment to "fill the pipe." If the RWIN is too small the sender will only transmit a small amount of data and then wait for an acknowledgement from the receiver.

This is the reason that it is difficult to fill a WAN connection that has latency. Typically, performance starts to suffer with as little as 20 ms of latency, equivalent to roughly 1600 kilometers.

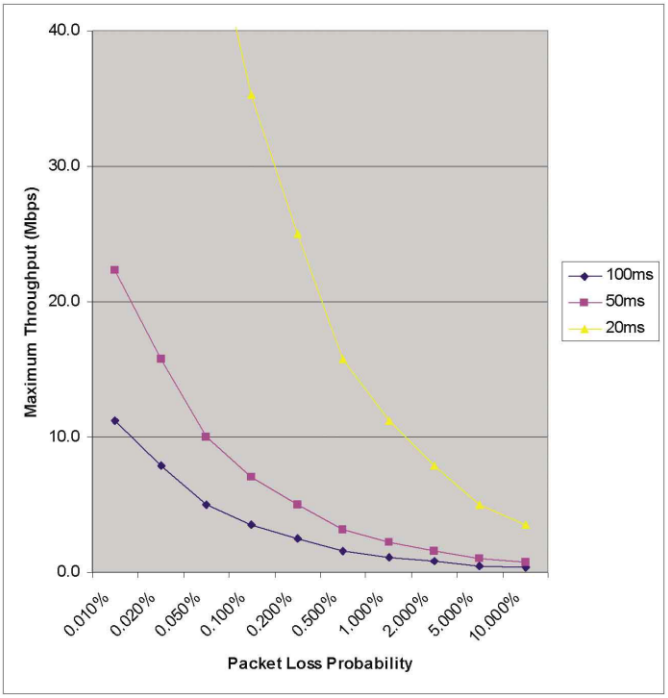

Figure 1: Impact of Packet Loss and Latency on Throughput

SILVER PEAK SOLUTION OVERVIEW

Silver Peak VXOA provides a variety of real-time optimization techniques to "clean up" WAN links for a higher effective throughput. *Forward Error Correction (FEC)* rebuilds lost packets on the Silver Peak appliance without causing data to be retransmitted across the WAN, and Packet Order Correction ensures that all packets are delivered in the order in which they were sent. Because packets are not retransmitted, the congestion control mechanism of TCP doesn't start and more data can be sent across the WAN. Forward error correction and packet order correction are both enabled by default on all Silver Peak appliances.

Silver Peak appliances can also perform advanced Quality of Service (QoS) techniques to prioritize traffic and ensure that necessary bandwidth requirements are met.

SOLUTION COMPONENTS

EqualLogic PS Series Storage Array

The EqualLogic PS Series storage solution provides consolidated network storage in a self-managing, iSCSI storage area network (SAN). The peer storage architecture used in the PS Series delivers high performance, availability, and scalability. The architecture allows all of the arrays to share resources, distribute workloads, and provide the highest levels of performance while reducing complexity.

EqualLogic Auto-Replication

Every EqualLogic array includes storage system-based replication software. This software, Auto-Replication, provides the ability to replicate data volumes to peer EqualLogic storage groups in remote locations. The remote EqualLogic storage group can then be used for disaster recovery. EqualLogic Auto-Replication provides point-in-time incremental data synchronization between primary and remote sites. Replication events are scheduled and include all changes that have occurred since the last synchronization. EqualLogic replication is typically performed across a wide area network (WAN).

Silver Peak VXOA

Silver Peak's VX and NX appliances complement the EqualLogic SAN solution by overcoming the common WAN challenges that adversely impact replication performance between source and target locations. More specifically, the Silver Peak solution mitigates bandwidth, distance, and WAN quality issues by optimizing Auto-Replication traffic transmitted across the WAN.

Network Memory: Silver Peak's patent-pending solution for disk based WAN deduplication. 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, 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 NX appliances 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.

DEPLOYMENT OPTIONS

Silver Peak appliances are easily deployed with EqualLogic Auto-Replication. Silver Peak VXOA is a symmetric solution that requires appliances to be deployed at all locations requiring optimization. For EqualLogic Auto-Replication, a Silver Peak appliance must be deployed in the primary data center, where the primary array is located, and at the DR data center, where the DR array is located.

For Silver Peak appliances to be able to optimize EqualLogic Auto-Replication traffic, the traffic must flow through, or be re-directed to, the appliances. There are two ways to accomplish this -- routed mode or bridge mode.

Note that Silver Peak physical and virtual appliances support the same deployment modes.

Bridge Mode

Bridge mode is the simplest deployment option and requires the least amount of configuration.

For a bridge mode deployment, the Silver Peak appliance is deployed inline between a network switch and the WAN router. An inline deployment does not require any changes to be made to the network switches or WAN router. If 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. When the appliance is operating in hardware bypass mode, it will behave differently, depending on whether copper or fiber interfaces are used. When copper interfaces are used, the appliance will behave as a crossover cable between the LAN switch and WAN router. When the fiber interfaces are used, the appliance will behave as a straight-through cable between the LAN switch and WAN router. Hardware bypass mode is only possible when using physical Silver Peak appliances.

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.

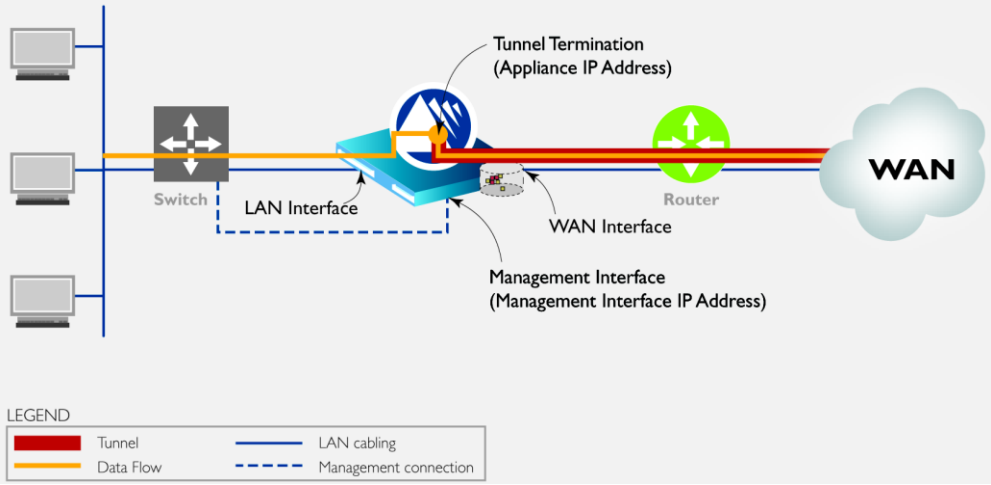

Figure 2: Bridge Mode Deployment

Routed Mode

With routed mode the Silver Peak appliance is not in the direct path of network traffic, and traffic must be redirected to the appliance 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 requires configuration of the WAN router to redirect traffic to the Silver Peak appliance. In the diagram below network traffic is redirected to the local Silver Peak appliance by the WAN router.

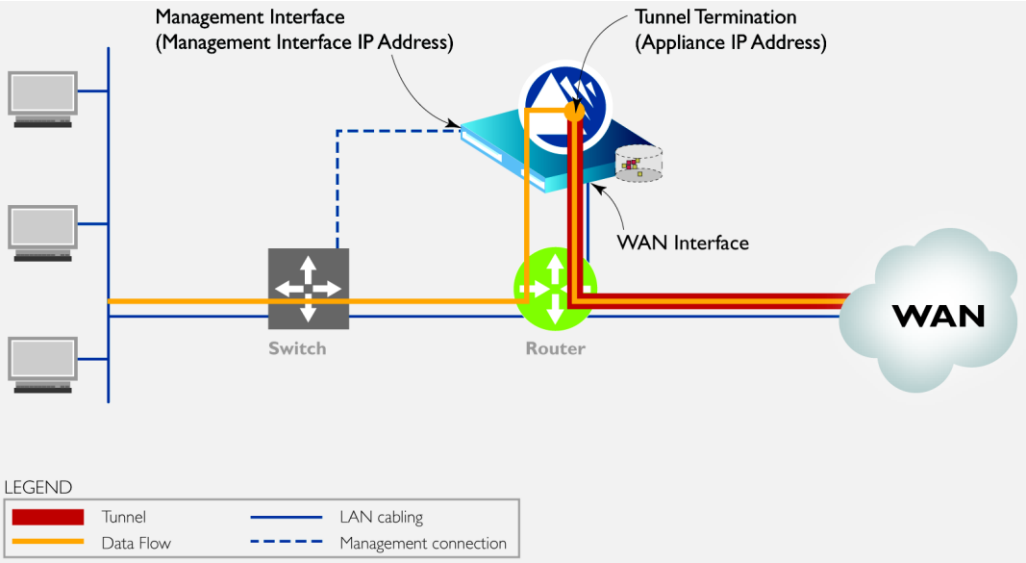

Figure 3: Routed Mode Deployment

For more information on network deployment options, refer to the Silver Peak Network Deployment Guide at http://www.silver-peak.com/Support/user_docs.asp.

Isolated IP SAN

A common deployment model for EqualLogic customers that desire a high level of security and performance is to isolate the IP SAN from the rest of the network. Best practices for designing an IP SAN include the only connectivity into the SAN being the hosts that are accessing iSCSI LUNs. Because the IP SAN is isolated from the rest of the network, replication is difficult to deploy. Solving this problem typically results in network connections being made between the LAN and IP SAN. An additional option is to connect ports on the WAN router to the IP SAN. Adding ports to the WAN router is not always an option due to port restrictions on the router.

To solve the isolated IP SAN problem Silver Peak uses a hybrid routed/inline mode and is connected to the IP SAN and the LAN. For this deployment the Silver Peak appliance is deployed in routed mode, and the EqualLogic groups are configured to use the Silver Peak appliance as the default gateway. The Silver Peak appliance then forwards the replication traffic to the WAN router.

When Silver Peak is deployed this way, the only traffic passing through the appliance is from the remote Silver Peak appliance. This means the IP SAN remains isolated from the production network, with the exception of replication traffic traversing the Silver Peak tunnel, keeping the IP SAN secure.

For more information on deploying Silver Peak with an isolated IP SAN see the following document: *Deploying Silver Peak VXOA with an EqualLogic Isolated IP SAN*.

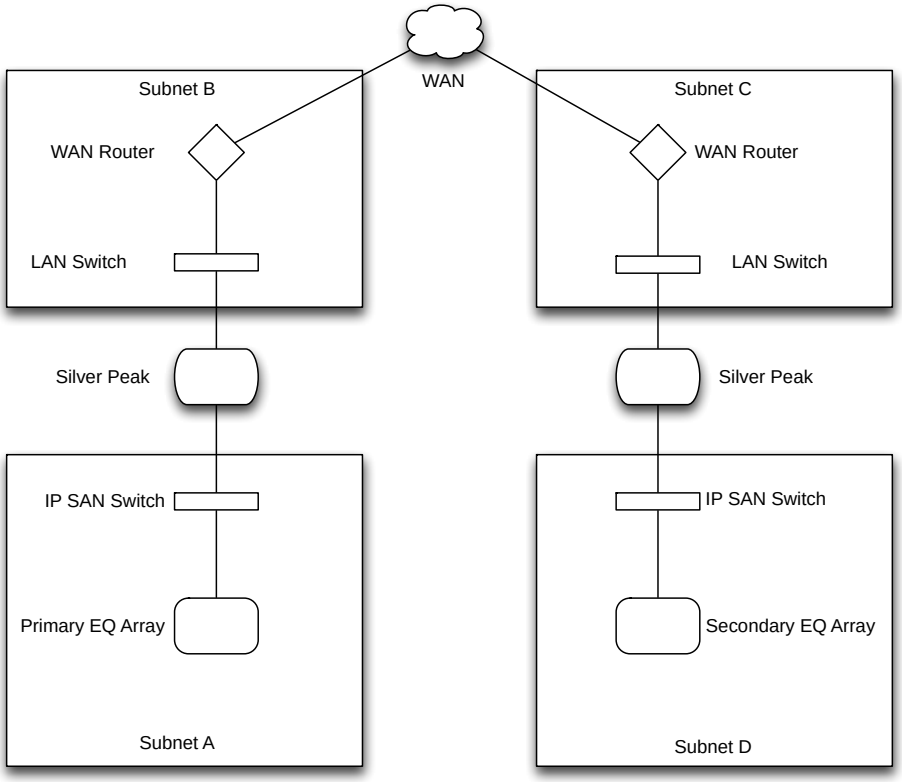

Figure 4: Isolated IP SAN Architecture

Deployment Considerations

Deploying Silver Peak VXOA with EqualLogic is very straightforward and does not require major changes to the existing infrastructure. Silver Peak and Dell have determined that the best performance is obtained by using the following versions:

Recommended Versions

- EqualLogic – Version 5.1.1h2 or higher
- Silver Peak VXOA – 4.4.10 or higher

TESTING OVERVIEW

Dell and Silver Peak conducted joint testing to qualify EqualLogic replication with Silver Peak VXOA. During this testing, the impact of various bandwidth limitations, loss, and latency settings on replication performance was studied. Based on these observations, best practices for WAN-optimized replication were developed.

To establish a reference baseline, EqualLogic replication performance was measured with and without Silver Peak VXOA using a simulated 45 Mbps WAN and packet loss values of 0 percent, 0.1 percent and 1 percent -- losses seen in

production WANs. A 0 percent loss is equivalent to a private line, 0.1 percent loss is commonly seen on MPLS WANs, and 1 percent and higher is commonly seen across the Internet and on satellite connections.

Latency of 20 and 80 milliseconds was added to the testing to determine performance degradation due to latency.

NOTE: *20 milliseconds of latency is equivalent to a WAN link between Austin, TX and Santa Clara, CA, while 80 milliseconds is equivalent to a WAN connection between Santa Clara, CA and New York City, NY. The throughput was measured in the EqualLogic and Silver Peak management interfaces and was converted from megabytes per second (MBps) to megabits per second (Mbps) for easier comparison to standard WAN speeds.*

TEST ENVIRONMENT

The joint test environment was set up as follows:

- Two Silver Peak VX-5000 virtual appliances were used with VMware™ vSphere 5 servers.
- The Silver Peak VX-5000 appliances were deployed in routed mode and connected to virtual switches.
- The EqualLogic arrays used the Silver Peak appliances as their default gateways.
- The vSphere server for the primary site also had a WAN emulator running to add latency and loss, as well as to simulate various WAN bandwidths.

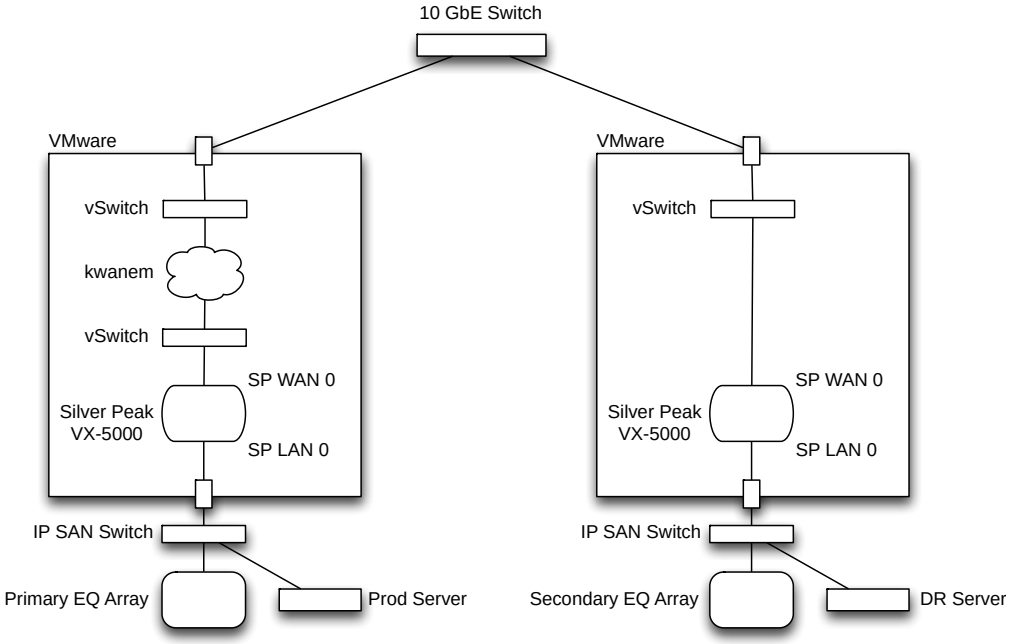

Figure 5: Test Lab Diagram

TEST PROCESS

- Data was written to the primary EqualLogic volumes.
- After the data was written to the primary volume, it was replicated to the secondary group across the simulated WAN using EqualLogic Auto-Replication.
- The WAN emulator limited the speed across the simulated WAN to 45 Mbps and also added 20 or 80 milliseconds (ms) of latency and, packet loss of 0.1 percent or 1 percent.
- Each test was performed multiple times to set baseline throughput without Silver Peak optimization and then again with Silver Peak optimization.

TEST RESULTS

In lab testing, Silver Peak VXOA provided a consistent performance increase when packet loss was added to a simulated WAN with EqualLogic Auto-Replication. To isolate the impact of packet loss on Auto-Replication, the latency of the WAN was set to 0 ms and the loss was set to 0 percent, 0.1 percent, and 1 percent.

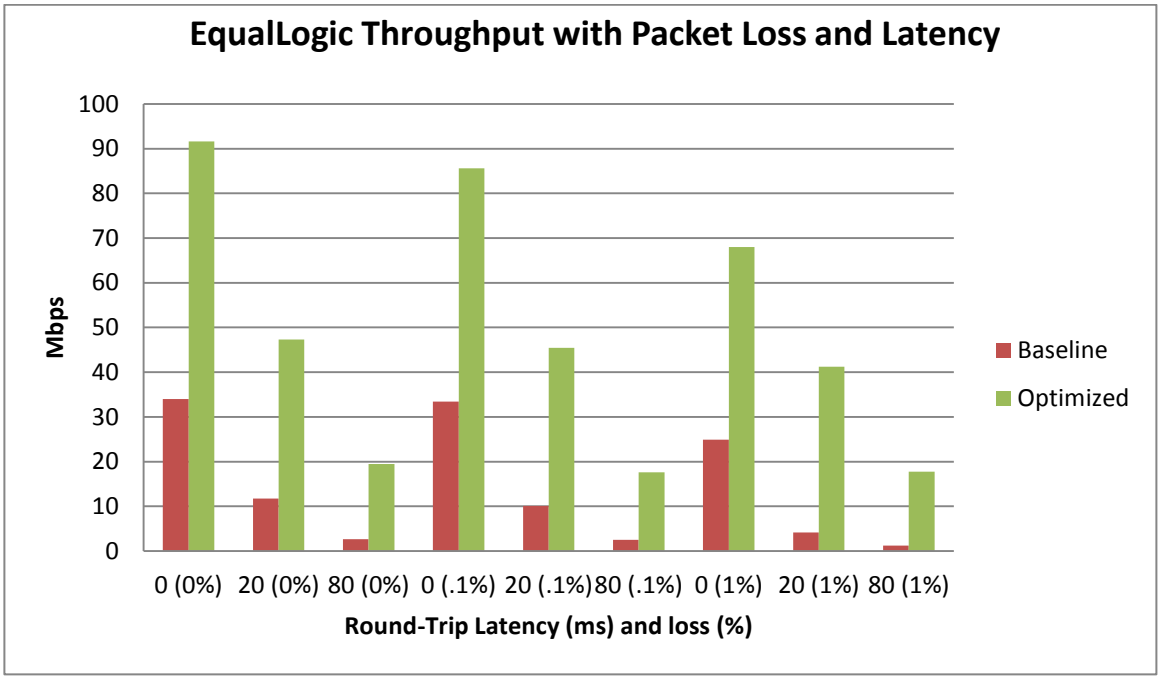

Figure 6: EqualLogic throughput with loss and latency

- The baseline 0 percent loss test showed Silver Peak VXOA providing a 3X performance increase. This increase can be attributed to the bandwidth reduction capabilities on Silver Peak VXOA. These results are reflected in the first three columns of Figure 6 -- 0 (0%), 20 (0%), and 80 (0%).
- When 0.1 percent loss is added without latency, Auto-Replication was able to maintain throughput; however, when latency was added throughput dropped by almost 70 percent. See Figure 6 columns for 0 (.1%), 20 (.1%), and 80 (.1%). This additional drop in throughput is due to TCP retransmitting data across the WAN. Whenever data has to be retransmitted, the impact of latency and loss on the throughput is increased.
- When packet loss was increased to 1 percent, the throughput dropped 5 percent without latency, and when latency was added throughput was reduced by nearly 90 percent. See Figure 6 columns for 0 (1%), 20(1%), 80(1%).

REPLICATION CHALLENGES: ADDRESSED

The following sections discuss some of the aforementioned replication challenges — packet loss, latency, bandwidth — and specifically how the joint Silver Peak/EqualLogic solution addresses them.

Packet Loss

VXOA is able repair the dropped packets using *Forward Error Correction (FEC)*. FEC takes the concept of parity used to protect data with RAID 5 in storage and applies it at the IP packet level. When loss is detected on the WAN, VXOA adds parity packets to the WAN traffic, similar to RAID 5 parity. When packets are dropped they can be rebuilt from the remaining packets and the additional parity information. *Figure 7* illustrates the impact of packet loss on throughput and the subsequent benefits of deploying VXOA to repair loss.

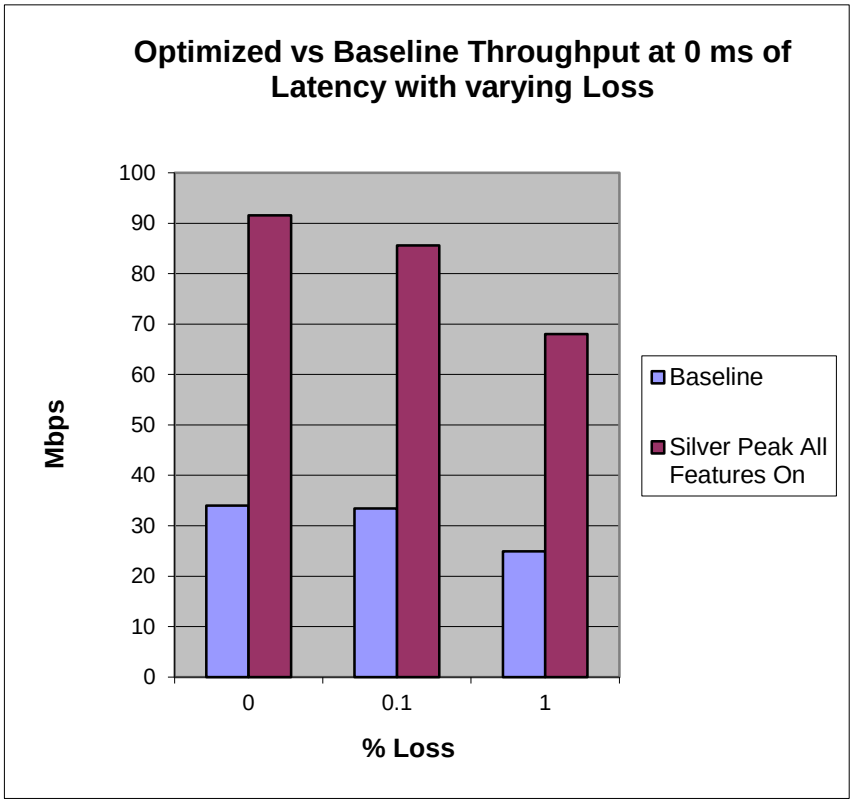

Figure 7: Throughput with packet loss

FEC can operate in static or auto mode. In static mode, a FEC packet is always sent for every *n* number of data packets. This setting is configurable by the administrator and offers flexibility from 1:2, 1:5, 1:10, and 1:20. A setting of 1:20 would generate one parity packet for every twenty data packets. In Auto mode parity packets are added automatically as needed.

In the graph above, FEC is disabled when the loss is set to 0 percent, resulting in no parity packets being sent. FEC is enabled with a static setting of 1:20 when loss is set to 0.1 percent, and 1:5 when loss is set to 1 percent.

A minor throughput reduction is seen in the testing at 0.1 percent and 1 percent loss because a fixed amount of parity packets are being sent. 1/5th and 1/20th of

the bandwidth will be used to send the parity packets. This is due to FEC being set to a static value of 1:5 and 1:20.

Using auto mode allows VXOA to proactively manage loss on the WAN and also obtain the highest throughput possible at all times. Typically, deployments use auto mode, resulting in parity packets being adjusted in real-time based on the level of packet loss present at any given time.

WAN Latency

The impact of latency on throughput can be seen in *Figure 8*. As latency increases throughput is reduced. As little as 20 ms of latency is enough to reduce EqualLogic Auto-Replication throughput to 12 Mbps across a 45 Mbps WAN.

VXOA uses a technique called Network Acceleration to overcome the impact of latency. When VXOA is deployed with EqualLogic Auto-Replication, VXOA will manage the RWIN. The default RWIN for EqualLogic arrays is 72KB; however VXOA will provide an RWIN of up to 1 GB. VXOA will also perform TCP acknowledgements local to the sender in order to keep data streaming across the WAN, further reducing the impact of latency.

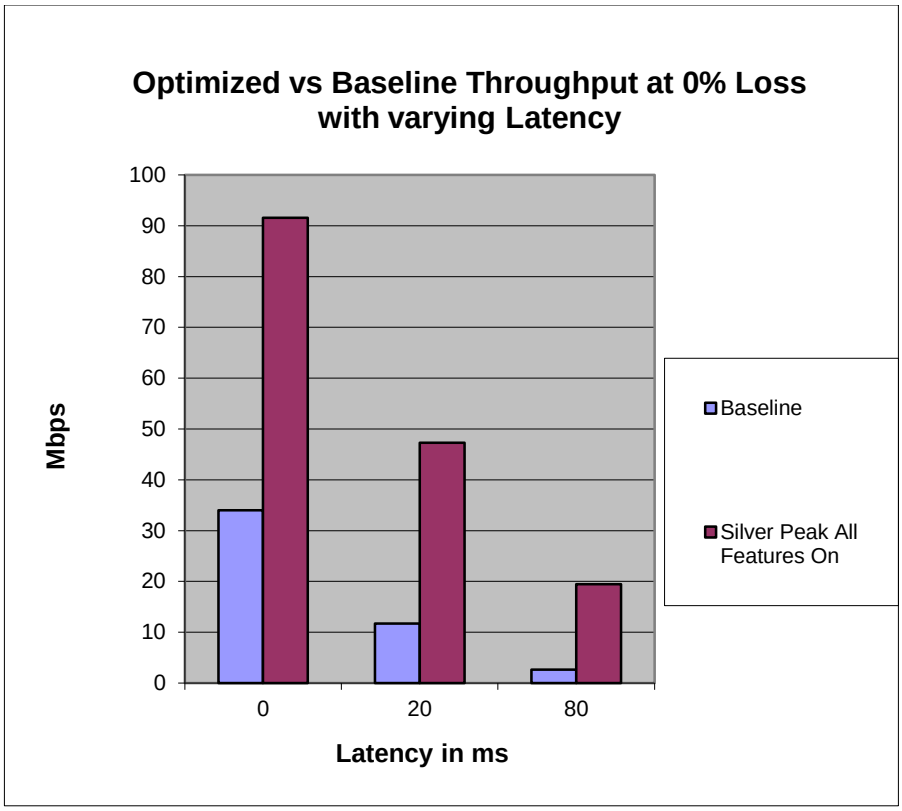

Figure 8: Throughput with latency

Limited Bandwidth

Bandwidth limitations can exist on the WAN itself or caused by other traffic using the WAN. For example, traffic can be limited to 10 Mbps because that is the size of the WAN, or because the WAN is 20 Mbps but 10 Mbps is reserved for application traffic, VoIP and video conferencing. Typically a larger WAN connection is installed to provide more bandwidth, increasing monthly telecom costs.

Silver Peak VXOA utilizes a feature called Network Memory to provide bandwidth reduction for data sent across the WAN. Network Memory uses disk-based deduplication to eliminate the transfer of duplicate data across the WAN. After being deduplicated, traffic is compressed to provide additional reduction. Network Memory works at the byte level and does not have a block size, fixed or variable. Because Network Memory is free to find redundant data across block and packet boundaries, typical reduction rates are 60-90 percent for replication data.

Figure 9 shows a chart from the VXOA management interface with a peak of 40 Mbps of Auto-Replication traffic sent across the WAN while using only 7 Mbps of bandwidth.

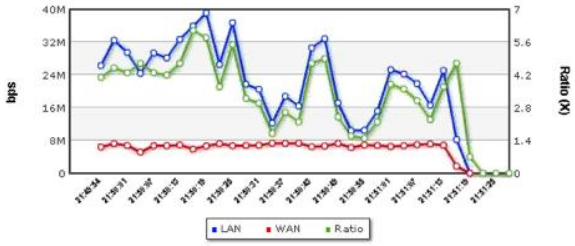

Figure 9: throughput with limited bandwidth

To test Network Memory with Auto-Replication, a data set was created that had a known reducible size of 10x. By using a known dataset, we were able to determine that Auto-Replication doesn't make any changes to the data as it is replicated and Network Memory can be utilized to provide the highest level of reduction possible.

Packet Loss and Latency on the Same WAN

Most replication environments have a combination of loss, latency and limited bandwidth. When these three things are combined the problems are exacerbated: Loss requires retransmissions, which are impacted by latency; the already limited bandwidth is used by retransmitted data; and finally, if the loss is caused by congestion, the problem can be made worse by the dropped data being retransmitted.

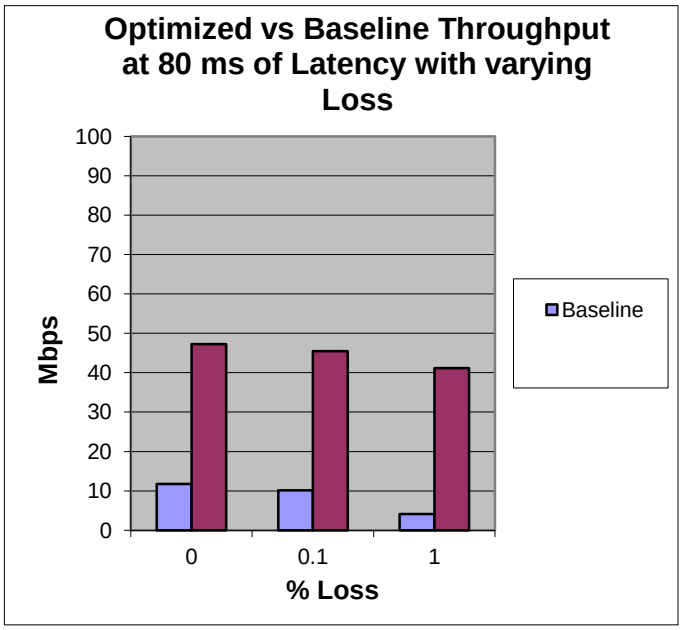

Figure 10: Throughput with loss and latency

In lab testing, latencies of 0, 20 and 80 ms were used with loss rates of 0 percent, 0.1 percent, and 1 percent. The combination of these latencies and loss rates provides guidelines on how Auto-Replication behaves across typical WAN connections, and to what level Silver Peak VXOA is able to optimize throughput.

We have already presented the results of 0 percent loss with 80 ms of latency — 11 Mbps without Silver Peak and 47 Mbps with Silver Peak. When 0.1 percent loss is added, Auto-Replication is able to maintain 11 Mbps of throughput and Silver Peak VXOA is able to increase throughput to 45 Mbps. With 1 percent loss on the same WAN, throughput drops to 4 Mbps, and Silver Peak VXOA is able to increase this to 41 Mbps.

CONCLUSION

EqualLogic Auto-Replication provides a reliable solution for disaster recovery and data protection. When Silver Peak VXOA WAN optimization is combined with Auto-Replication, data can be replicated across longer distances, in a shorter time, and using less bandwidth. The combined EqualLogic and Silver Peak solution helps save money by optimizing the use of WAN bandwidth and can even improve recovery point objectives. The joint Dell and Silver Peak solution has been tested extensively and is proven to provide the highest throughput and flexibility using Auto-Replication.

In addition to optimizing EqualLogic Auto-Replication, Silver Peak VXOA can optimize any application that is using the WAN. Among these applications are VoIP, VDI, video, and file protocols like CIFS and NFS. For additional information, refer to the Silver Peak Info Center at http://www.silver-peak.com/InfoCenter/.

Deploying Silver Peak VXOA with Dell EqualLogic Auto-Replication provides the following benefits:

- Increases replication throughput and reduce bandwidth usage.
- Replicates over longer distances.
- Lowers bandwidth costs by sharing the WAN with other applications.
- Meets or beats challenging RPOs.
- Replicates over low-cost WAN links, like MPLS or Internet VPN.
- Enables replication for isolated iSCSI SANs.