

EMC Celerra Replicator V2 with Silver Peak WAN Optimization

Applied Technology

Abstract

This white paper discusses the interoperability and performance of EMC[®] Celerra Replicator[™] V2 with Silver Peak's WAN optimization. Where applicable, it also discusses deployment and configuration best practices.

March 2010

Copyright © 2010 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

All other trademarks used herein are the property of their respective owners.

Part Number H6988

Table of Contents

Executive summary	4
Introduction	4
Audience	4
Terminology	4
Data replication with Celerra Replicator	4
Technical challenges of data replication over distance	5
WAN optimization with Silver Peak	6
Performance comparison with Silver Peak optimization.....	6
Sizing Silver Peak with Celerra Replicator	8
Conclusion	8
References	9

Executive summary

Point-in-time asynchronous replication is a common disaster recovery solution for data centers that have remote secondary sites located at varying distances from each other. EMC® Celerra Replicator™ provides the ability for file systems to replicate to remote file systems and ensure business continuance. However, replicating over large distances can slow data transfer due to network limitations such as latency and bandwidth congestion at one or more points across the connection. This is particularly crucial during the initial synchronization between the primary and secondary file systems, where large amounts of data are transferred over the network connection.

Silver Peak provides a variety of wide area network (WAN) optimization techniques to address these common network issues. With Silver Peak optimization, Celerra Replicator experiences a dramatic increase in data transfer performance.

Introduction

This white paper introduces the benefits of Silver Peak and the techniques used to optimize Celerra® replication across a WAN environment. The white paper compares the performance of a WAN replication with Silver Peak to that of a WAN replication without Silver Peak. The paper does not provide specific quantitative results of the performance increase. Rather, it provides the reader with a guideline to what the general expectations should be with Silver Peak optimization.

This white paper discusses the common issues that may appear in replication over long distances. Then, the paper moves onto a brief discussion about Silver Peak and the techniques that it leverages to provide significantly improved transfer speeds. Finally, the paper makes a general comparison between having Silver Peak in the data path and not having it in the data path. The comparisons are made across different distances: LAN, cross-country, and intercontinental.

Audience

This white paper is intended to be used by EMC customers, EMC employees, partners, IT planners, storage architects, administrators, and any others involved in evaluating, acquiring, managing, operating, or designing an EMC networked storage environment.

Terminology

Network Data Management Protocol (NDMP) — Open standard network protocol designed for enterprise-wide backup and recovery of heterogeneous network-attached storage.

recovery point objective (RPO) — Describes the acceptable amount of data loss, measured in units of time, for example, 12 minutes, 2 hours, or 1 day. This represents a target that is derived from conditions specified in an SLA, RTO, and relevant analyses. The RPO in conjunction with the recovery time objective (RTO) is the basis on which a data protection strategy is developed.

recovery time objective (RTO) — How long a business process can be down before consequences are unacceptable. This represents a target that is derived from conditions specified in an SLA and business impact analysis.

replication — Service that produces a read-only, point-in-time copy of a source file system. The service periodically updates the copy, making it consistent with the source file system.

Data replication with Celerra Replicator

Celerra Replicator is a leading solution for point-in-time asynchronous replication of file systems and iSCSI LUNs. It provides a variety of flexible options for both local and remote replication (across multiple sites, if desired), with rapid asynchronous data recovery. Celerra Replicator offers advanced replication

features such as one-to-many and cascading replication which, when leveraged together, can support up to 16 remote copies. Celerra Replicator also provides automatic replication management with adaptive technology, which synchronizes the source and destination based on a single configurable update interval value.

In enterprise environments, where large file systems are common, a process called tape silvering is used to create the initial copy through NDMP tape backup and then the copy is shipped to the destination. This process is frequently referred to as “tape and truck.” Initial data synchronization can also take place through an IP network. Either method can involve the transfer of large amounts of data, including file systems, Virtual Data Movers (CIFS server data, audit logs, local groups), or iSCSI LUNs. After the initial synchronization, Celerra Replicator uses differential snapshots to send only incremental block-level changes between hosts.

Technical challenges of data replication over distance

When Celerra replication is configured across a WAN, a variety of challenges with the WAN can affect the replication and recovery process. These challenges include:

- Network latency
- Packet loss
- Bandwidth congestion

Network latency is the time it takes for information to go from sender to receiver and back. Because the speed of light is constant, WAN latency is directly proportional to the distance traveled between the two network endpoints. When routers and other network elements perform queuing and processing functions, additional latency is added to the equation. Network latency becomes problematic when higher-level protocols are “chatty.” Many protocols involve numerous acknowledgements before sending data, which can severely affect transfer times. The table below details some typical latency values for various network topologies.

Table 1. Latency values for network topologies

Network description	Average latency
LAN	5 ms
Eastern U.S. to western U.S.	80 ms–100 ms
International WAN connection	100 ms–300 ms
Satellite connection	Over 500 ms

Even though lower layer technologies are error-free, packet loss can still occur at the network layer due to congestion in routers, link failures, network re-routes, and other equipment problems. Packet loss is especially prevalent on Multiprotocol Label Switching (MPLS) and IP VPN WANs, where oversubscribed network resources can experience congestion during periods of peak usage. In these environments, it is common to see packet loss rates as high as 5 percent, with averages in the .1 percent to 1 percent range. This can cause significant problems for data replication, where sustained data throughput is required on a per-flow basis.

Bandwidth congestion may occur at several points across a WAN connection. Within the enterprise, many applications and services may be contending for the same bandwidth. Network providers may oversubscribe bandwidth within the core of their network. Network hardware could be undersized or improperly configured. All of these situations can restrict replication throughput. A typical solution to these situations would be to purchase dedicated bandwidth to minimize the impact of bandwidth congestion on data replication. However, dedicated bandwidth is expensive and may not always provide the desired results.

Poor WAN conditions, as described above, can slow data replication to a crawl or even prevent replication from happening altogether. This jeopardizes RPOs and RTOs, which are critical to the enterprise. Traditionally, solving performance issues of data replication is expensive and difficult to implement.

WAN optimization with Silver Peak

Silver Peak provides a variety of WAN optimization techniques that address poor WAN conditions and enable all applications to run at near LAN performance levels while providing services over global distances. The testing reviewed in this paper confirmed that Celerra Replicator performance dramatically increased throughput when run over WAN links that have been Silver Peak optimized. These optimization techniques break down into the following three categories:

- **Network integrity:** Silver Peak provides a variety of tools that address packet delivery issues common to shared WAN technologies, such as MPLS and IP VPN. For example, Forward Error Correction (FEC) is used to rebuild lost packets in real time and Packet Order Correction (POC) is used to ensure that all packets are delivered in the order that they were sent. These techniques eliminate the need to retransmit dropped or out-of-order packets, which maximizes the effective throughput of traffic sent during the replication process. In addition, Silver Peak employs advanced quality of service (QoS) techniques to prioritize replication traffic and guarantee that necessary bandwidth requirements are met.
- **Network acceleration:** Silver Peak uses a variety of techniques to address latency across the WAN. These include TCP acceleration techniques such as 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 network traffic that can otherwise hamper application performance across a WAN.
- **Network memory:** One of the main tools in Silver Peak's arsenal is Network Memory technology, a premier platform for WAN deduplication. Network Memory uses disk-based data reduction to eliminate the transfer of duplicate information across the WAN, resulting in three to 10 times more virtual bandwidth between Celerra devices. Network Memory works on all IP WAN traffic, in real time, on a packet-by-packet basis. For this reason, it is complementary to other deduplication processes that may exist on the host or within the application layer. As a result, significant data reduction is achieved during the initial synchronization process and subsequent differential copy processes between Celerra Data Movers.

Performance comparison with Silver Peak optimization

A common baseline configuration to evaluate the performance of WAN optimization with Silver Peak is the 155 Mb/s WAN link. This speed represents a generally achievable rate at which replication sessions operate in a WAN environment. This rate will serve as a comparison for the baseline and optimized tests with loss, latency, and bandwidth impairments included. Figure 1 on page 7, Figure 2 on page 7, and Figure 3 on page 8 compare the performances with and without Silver Peak in the network data path in LAN environments, cross-country environments (eastern U.S. to western U.S.), and international WAN environments. The results are based on the use of 16 replication sessions in parallel. This performance comparison chart is intended as a guideline and not as a real-world predictor for actual network or replication performance.

The three figures set the baseline configuration (throughput without Silver Peak enabled) at 100 percent. Figure 1 on page 7 shows that in a LAN environment an optimized throughput can be expected to provide over three times the performance improvement that the baseline throughput provides.

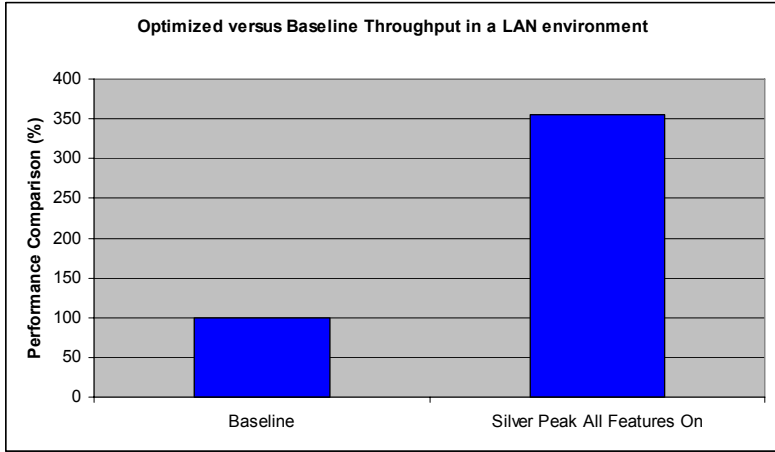


Figure 1. Performance comparison in a LAN environment

Figure 2 shows the performance comparison in a cross-country environment. With Silver Peak enabled, there is an average of 700 percent improvement compared to the baseline throughput.

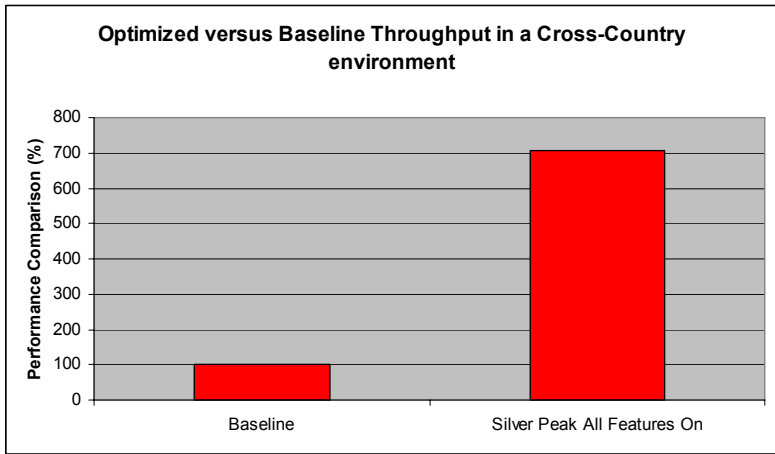


Figure 2. Performance comparison in a cross-country environment

Figure 3 shows the performance comparison in an environment replicated internationally. With Silver Peak enabled, the performance improves an average of 16 times compared to the baseline throughput. Silver Peak uses WAN optimization techniques to address poor WAN conditions. As the trend from these three figures show, Silver Peak dramatically improves performance over WANs, where the network path is more susceptible to loss, bandwidth impairments, and in particular latency.

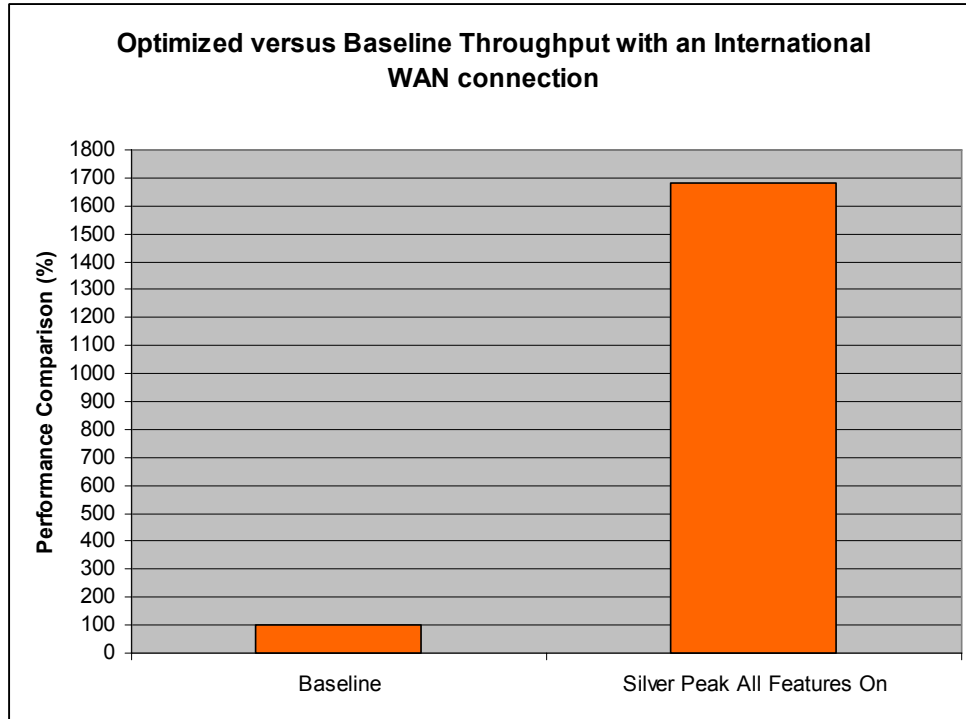


Figure 3. Performance comparison in an environment replicated internationally through a WAN connection

Sizing Silver Peak with Celerra Replicator

Generally, the Silver Peak NX appliance is sized according to the expected amount of overall WAN throughput required. Silver Peak appliances can be deployed on any size WAN links, from a low-end, fractional T1 WAN link (less than 1.5 Mb/s) to a high-end full 1 Gb/s WAN link. Most often, the Silver Peak appliance is sized to match the line speed in a given environment. However, it is acceptable to size the Silver Peak appliance based on only the expected optimized throughput. This is done in situations where only a subset of the total traffic is to be optimized.

When sizing the Silver Peak deployment, the number of flows (for example, in this context, the number of replication sessions) traversing the WAN connection must also be considered. However, it will be a rare occurrence that Celerra Replicator flow counts alone exceed the capacity of even the smallest Silver Peak appliance. This is because Silver Peak appliances can handle a minimum of 64,000 optimized flows and a maximum of 256,000 optimized flows. If the WAN connection is shared by applications and services other than Celerra Replicator, it will be necessary to assess the total number of flows, including optimized and unoptimized flows.

For additional information about sizing Silver Peak appropriately in Celerra Replicator environments, please contact your EMC technical account representative. You can also involve your local Silver Peak account team.

Conclusion

The comparisons in this white paper show that significant benefits can be achieved by deploying Celerra Replicator in a Silver Peak-optimized WAN environment. Silver Peak's network integrity, network acceleration, and optimization techniques for network memory enable Celerra Replicator deployments to perform at LAN speeds across long distances. Without Silver Peak, enterprises are limited in their choices when designing data replication networks.

For instance, to minimize network latency and to curb costs, enterprises have traditionally sought to keep the replication distance relatively short. Moreover, to avoid transmission errors and packet loss, companies have traditionally been forced to pay a premium for clean bandwidth with no packet loss. Silver Peak addresses both of these obstacles by allowing infrastructure architects to deliver LAN-like throughput across any distance by using any type of bandwidth. With Silver Peak, enterprises can use lower-cost shared links such as MPLS or Internet VPN connectivity, for data replication.

Silver Peak can also assist with the deployment of Celerra Replicator. Typically, methods such as “tape silvering,” also called “tape and truck,” are utilized during initial deployment to synchronize arrays. With Silver Peak deployed, enterprises can opt to do the initial sync through the Silver Peak-optimized WAN connection. This saves significant time, money, and resources.

Finally, Silver Peak enables applications like Celerra Replicator to achieve maximum application throughput across a WAN. Enterprises can decrease RPOs and RTOs while obtaining strong data reduction and maintaining high throughput. Doing so also reduces RPOs by enabling faster replication and saves costs

References

Name: *EMC Celerra Version 5.6 Technical Primer: SLA-Driven Replication with Celerra Replicator (V2) - Technology Concepts and Business Considerations*

Type: White paper

URL: <http://powerlink.emc.com>

Audience: Customer, Employees, Partners

Technical Depth: High