

Integration Guide

EMC Data Domain and Silver Peak VXOA

4.4.10 Integration Guide

August 2013

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.

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

Part Number: 302-999-817 Rev. 01

Table of Contents

About this document	4
Audience.....	4
Related documents	4
Overview	4
Replication Challenges	5
Solution Components	7
EMC Data Domain.....	7
Silver Peak VXOA.....	7
Network Memory	7
Network Integrity	8
Network Acceleration	8
Secure Content Architecture™	8
Global Management System.....	8
Customer Example: Global Telecommunications Company	8
Replication across the United States using a private line.....	8
WAN.....	9
Throughput	9
Deploying Silver Peak with Data Domain	10
Recommended Versions.....	10
Silver Peak Configuration	10
Silver Peak Deployments Options.....	10
Sizing Silver Peak Appliances.....	11
Data Domain System Configuration	11
Replication Settings	11
Bandwidth Throttle.....	12
Data Domain Static Route.....	13
Monitoring Performance	15
Silver Peak Performance Monitoring.....	15
Data Domain Performance Monitoring.....	17

About this document

Note: Prior to performing this integration, be sure to consult the *EMC Data Domain Product Compatibility Matrix* on Powerlink or the EMC Data Domain support portal (<https://support.datadomain.com>) for the versions of Silver Peak and DD OS certified for use with the respective products.

This document explains how to deploy EMC Data Domain® deduplication storage systems with Silver Peak VXOA WAN optimization appliances.

The instructions provided in this document by EMC are for customer convenience and are not warranted or supported by EMC. EMC expects users to customize installation of third-party software for use at a particular site, but EMC is not responsible for the usability of third-party software after installation.

This document is applicable for the following versions of Silver Peak:

- Silver Peak VXOA 4.4.10 and higher

This document is applicable for the following versions of the Data Domain operating system (DD OS):

- EMC Data Domain systems running DD OS versions 5.1.0.9 and higher

Audience

This paper is intended for customers, technical field consultants, and customer engineers who have familiarity with either Silver Peak and/or Data Domain operations.

Related documents

The Data Domain and Silver Peak documents referred to in this guide provide additional information for configuring and using Data Domain and Silver Peak.

The Silver Peak documents are available at: http://www.silver-peak.com/Support/user_docs.asp.

Documentation for the Data Domain system is available at: <https://my.datadomain.com/documentation> and on Powerlink.

Overview

EMC Data Domain deduplication storage systems optimize storage efficiency through the use of deduplication and compression of data at rest. When the deduplicated and compressed data is replicated using Data Domain, it is a common belief that WAN optimization is not required because a smaller amount of data is being sent across the WAN. However, in addition to limited bandwidth, at least two factors can affect the ability for replication to meet customer RPOs over a Wide Area Network (WAN): replication distance (latency) and quality of the WAN itself (packet loss and out of order packets). By solving many of the underlying network problems that plague

replicating over distance, Silver Peak VXOA, when combined with EMC Data Domain Replicator, enables customers to move more data, over a longer distance, and in less time. In this paper we will describe the benefits that Silver Peak's Virtual Acceleration Open Architecture, VXOA, provides for the Data Domain system; provide deployment and configuration examples; detail best practices; and illustrate expected performance improvements based on a customer deployment.

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 Recovery Point Objective (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 12:00 P.M., any data created between 9:00 A.M. and 12:00 P.M. 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 a 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 a 0.1% loss (typical of MPLS WAN links) can decrease replication performance significantly
- 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.

For many storage replication solutions, all three of the above issues plague replication performance and the ability for customers to meet RPOs. For Data Domain, one of the key issues (amount of bandwidth) is already addressed since the data is already deduplicated and therefore does not need to be sent over the WAN. In situations where Data Domain replication shares the WAN pipe with other replication (or non-replication) applications, however, bandwidth can still be a constraint (for the other applications).

Silver Peak VXOA, deployed as either a virtual or physical appliance, quickly and easily repairs the above WAN quality problems, removes the impact of latency, and

eliminates many bandwidth bottlenecks. The result is dramatically increased replication throughput for maintaining RPOs.

For example, a 1 Gbps WAN with 0.1% of packet loss and 80 ms of latency will be limited to approximately 7.5 Mbps per TCP session. The limitation comes from the combination of latency and packet loss. When packets are dropped TCP will limit the throughput (congestion control) while recovering from the error and resending the data that was lost. It is the combination of congestion control, retransmission, and latency that cause the throughput limitations seen when replicating across distance over a WAN with congestion or loss. [Figure 1](#) illustrates the maximum throughput across a WAN with different packet loss and latency combinations.

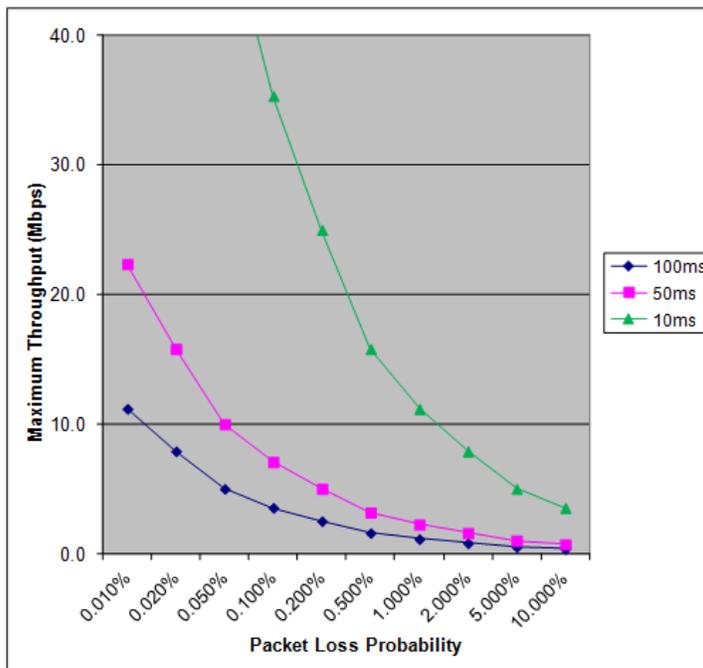


Figure 1 Packet Loss Probability

Silver Peak VXOA will complement Data Domain replication in the following scenarios:

- Replication occurs over longer distances, resulting in increased latency (and decreased replication throughput).
- The WAN suffers from poor quality (packet loss and/or out of order packets).
- The WAN is shared with multiple applications with the potential for congestion.
- Any combination of the above.

In summary, using Silver Peak with Data Domain allows customers to:

- Replicate at greater distances, while maintaining replication throughput.
- Repair network problems, loss, and out-of-order, allowing the full use of available bandwidth for replication, even across the Internet.

Solution Components

This solution can be deployed using any EMC Data Domain hardware and any Silver Peak VXOA physical or virtual software appliance.

EMC Data Domain

EMC Data Domain deduplication storage systems reduce the amount of disk capacity required to store and protect enterprise data with inline deduplication and compression. Redundant data is identified, as it is stored, reducing the amount of capacity required by 10-30X on average. After data is stored on the Data Domain system, it can be replicated for disaster recovery by using DD Replicator software.

DD Replicator software provides network-efficient and encrypted replication for disaster recovery (DR), remote office data protection, and multi-site tape consolidation. DD Replicator software vaults or asynchronously replicates only the compressed and deduplicated data over the wide area network (WAN), making network-based replication economically and operationally feasible.

Silver Peak VXOA

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.

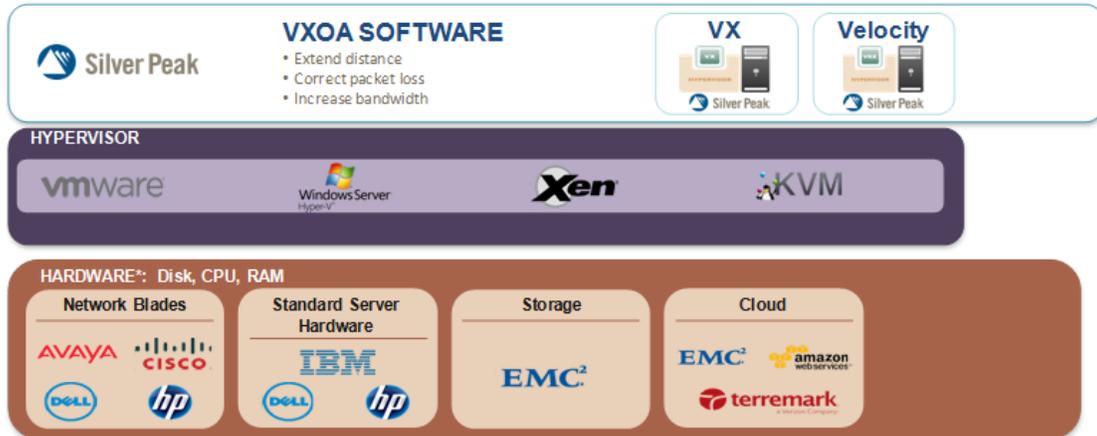


Figure 2 Silver Peak VXOA

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

Network Memory

Silver Peak has a 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 products are built on a Secure Content Architecture™ that enables enterprises to deploy WAN acceleration with complete confidence. The Silver Peak solution incorporates the latest in encryption technology to protect data at all times: at rest and in transit across the WAN. In addition, Silver Peak makes it easy to configure, enforce, and monitor security policies from a central location through the Silver Peak Global Management System (GMS), and employs mechanisms to ensure that security does not come at the expense of network performance or scalability. The result is end-to-end secure WAN acceleration.

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 managers 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. With Silver Peak, enterprises can deliver business critical applications centrally in a secure, scalable, and cost effective manner.

Customer Example: Global Telecommunications Company

Replication across the United States using a private line

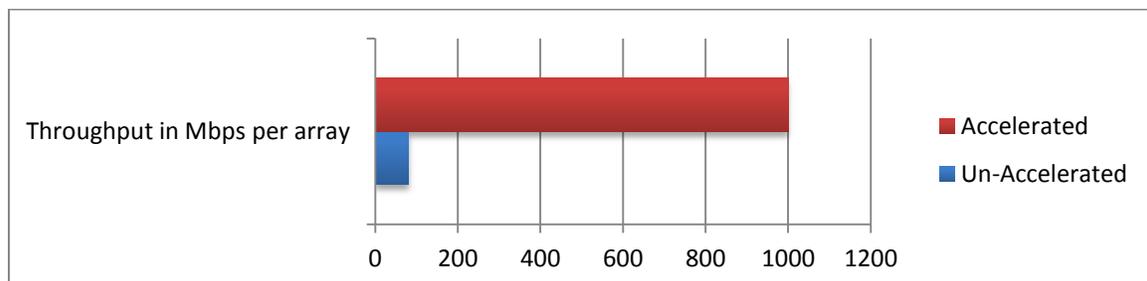
In this use case the customer is using Data Domain Replicator to replicate several Data Domain systems simultaneously across the United States using a private line.

Typically a private line offers the highest available throughput due to an absence of out-of-order and dropped packets. This customer faced the problem of high latency, limiting throughput and causing replication jobs to take an excessively long amount of time, resulting in missed RPOs.

Deploying Silver Peak VXOA allowed the customer to use all of the available bandwidth on the WAN by overcoming the effects of latency. Before Silver Peak was deployed, each Data Domain system was able to use approximately 80 Mbps of bandwidth. After Silver Peak VXOA was deployed, the amount of bandwidth utilized by each array increased to 1 Gbps, allowing the entire 10 Gbps private line to be fully utilized, which in turn allowed the customer to meet their RPO.

Note that the Data Domain system allows optimization of the networking connection by changing the network setting for Delay and Bandwidth, resulting in higher throughput when using the Data Domain Replicator. (Refer to the *EMC Data Domain Operating System Administration Guide* for instructions.) The networking settings for this customer had been optimized, but still suffered from high amounts of latency, resulting in the Transmission Control Protocol (TCP) reducing the throughput. Deploying Silver Peak VXOA accelerates Data Domain Replicator without any additional tuning on the Data Domain arrays.

Silver Peak VXOA is also able to provide additional benefits by deduplicating replication data across the multiple Data Domain systems as it is replicating. Note that this deduplication occurs only while the data is being replicated and does not replace the deduplication that each Data Domain system provides for data at rest.



WAN

- 10 Gbps private line stretching across the United States
- 90 ms of latency
- 0% packet loss and out-of-order

Throughput

- Without Silver Peak VXOA – 80 Mbps per array
- With Silver Peak VXOA – 1 Gbps per array

Deploying Silver Peak with Data Domain

Recommended Versions

- Silver Peak VXOA 4.4.10 and higher
- Data Domain OS 5.1.0.9 and higher

Silver Peak Configuration

Use the default VXOA settings shown in [Table 1](#) unless the WAN connection has consistent levels of packet loss. When packet loss is variable, Forward Error Correction (FEC) should be set to Auto. This lets VXOA dynamically adjust the number of parity packets that are used to recover from the loss. When packet loss is consistent, a static Forward Error Correction can be used to help maintain throughput.

Note: If you specify a fixed value for Forward Error Correction, some bandwidth is always used for parity packets. For example, if you specify 1:10 for FEC, 10% of the available bandwidth is used for parity.

VXOA includes an application definition for Data Domain systems. By default, all traffic is optimized with the default settings. Except for adjusting Forward Error Correction as described above, do not change these default settings.

Table 1 Default Silver Peak Settings

Setting	Default for Data Domain Systems
Network Memory Mode	Balanced
Payload Compression	Enabled
TCP Acceleration	Enabled
Forward Error Correction (FEC)	Disabled

Silver Peak Deployments Options

You can deploy Silver Peak VXOA to optimize all network traffic or just specific applications. When you deploy Silver Peak to optimize the network, you use Silver Peak in bridge or router mode; this requires changes to the network. When you deploy Silver Peak for specific applications, no network changes are required. An application-specific Silver Peak deployment with a Data Domain system uses a static route that is configured on the Data Domain system. See [“Data Domain Static Route” on page 13](#). The choice of deployment method does not impact the performance of the Data Domain system.

For more information on Silver Peak deployment options, see the *Silver Peak Network Deployment Guide*.

Sizing Silver Peak Appliances

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 virtual 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/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

Data Domain System Configuration

Replication Settings

Configure replication by following the instructions in the *EMC Data Domain Operating System Administration Guide*. When the Data Domain system is used with Silver Peak VXOA, the network settings for Delay and Bandwidth, shown in **Error! Reference source not found.**, should be set to the default values (no settings). The Silver Peak appliances will provide optimization across the WAN and negate the need to change these settings. If these settings are configured the throughput increase that is provided by Silver Peak VXOA will be reduced.

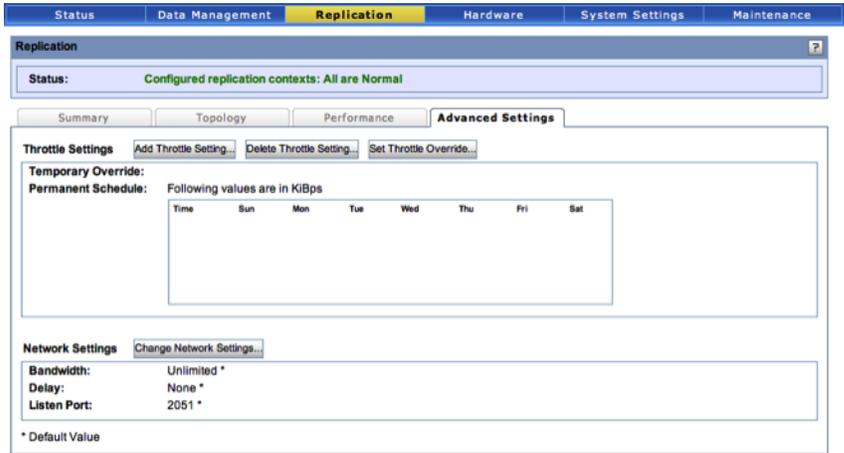


Figure 3 Data Domain Replication Settings

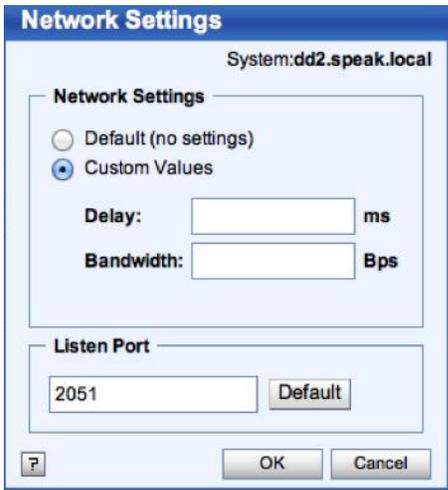


Figure 4 Data Domain Network Settings

Bandwidth Throttle

The Data Domain system provides a feature to throttle bandwidth based on time and day of the week (see Figure 5). When the bandwidth throttle is used, Silver Peak VXOA will not be able to accelerate replication traffic beyond the throttle. For example, if a 100 Mbps WAN is used and the throttle is set to 50 Mbps the maximum replication speed will be 50 Mbps.

For deployment with Silver Peak VXOA, the bandwidth throttle should be disabled initially. Disabling the throttle will allow the maximum throughput to be measured. Once the maximum throughput is known, the throttle can be set to reduce the amount of bandwidth that the Data Domain system uses. A throttle override option exists for times when the replication bandwidth needs to be changed for a limited amount of time (see Figure 6).

An additional option to control bandwidth usage is available with the Silver Peak QoS feature.



Figure 5 Data Domain Bandwidth Throttle

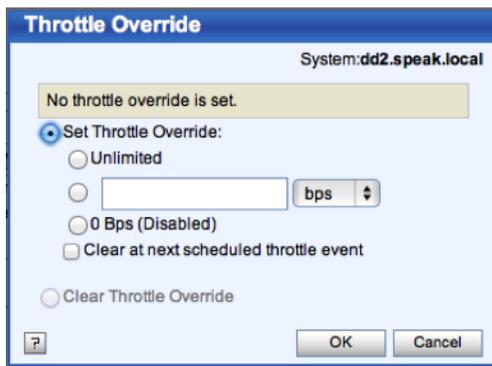


Figure 6 Data Domain Throttle Override

Data Domain Static Route

Silver Peak VXOA can be deployed with the Data Domain system using a static route. When a static route is used, no changes are required on the network or the WAN router. A route will be added to each Data Domain appliance that is replicating to use the local Silver Peak appliance, physical or virtual, as the next hop for replication traffic.

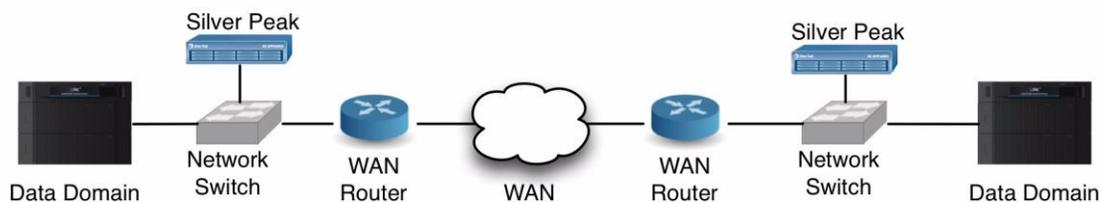


Figure 7 Silver Peak and Data Domain Static Route Network Diagram

To deploy Silver Peak with the Data Domain system using a static route:

1. Choose Hardware > Network > Routes in the EMC Data Domain System Manager.
2. Click Create under static routes (see [Figure 8](#)).
3. Select the interfaces that are used for replication and click Next (see [Figure 9](#)).

- For the destination, select Host and enter the IP Address for the target Data Domain system (see Figure 10).
- For the gateway, select Specify different gateway for this node and enter the IP Address for the local Silver Peak appliance (see Figure 10).

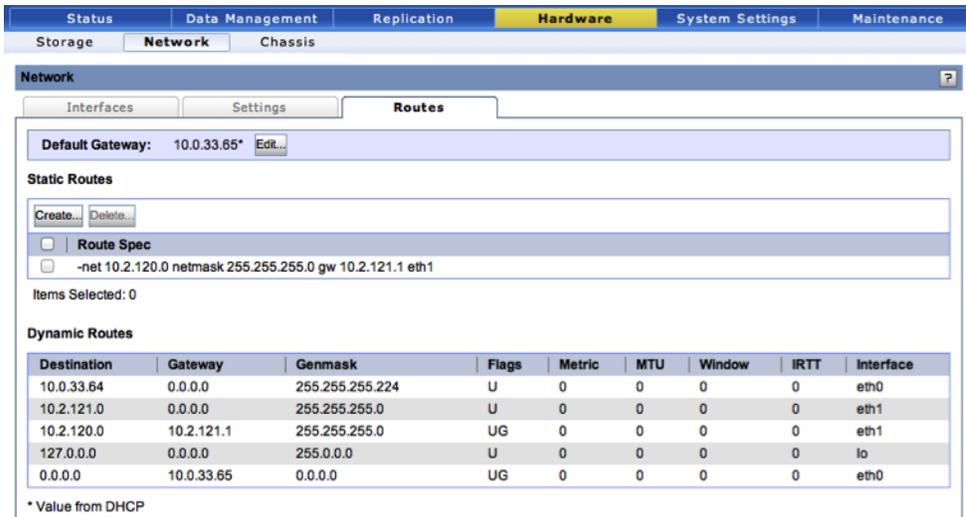


Figure 8 Data Domain Static Route

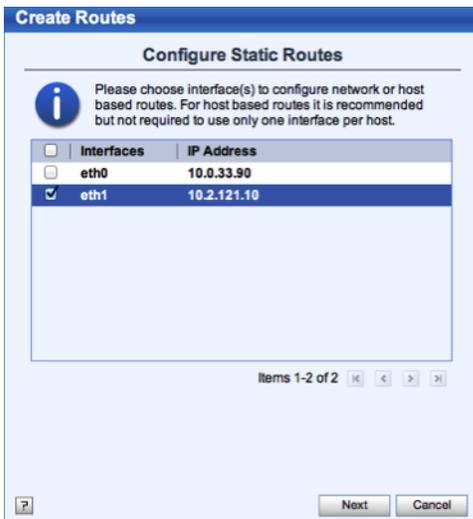


Figure 9 Data Domain Create Route

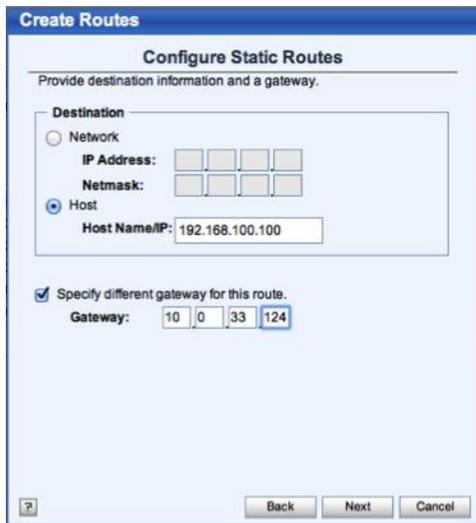


Figure 10 Data Domain Create Route

Monitoring Performance

Silver Peak Performance Monitoring

Silver Peak VXOA will automatically identify and categorize the Data Domain system traffic. By identifying Data Domain traffic, the Silver Peak reports will contain detailed information for all the Data Domain system traffic. The VXOA home screen, shown in [Figure 11](#), shows real time information for bandwidth utilization, latency, loss, top applications, and the top ten flows. A flow is identified as a connection between two end points. For example, a Data Domain replication session between two systems will be identified as a flow. Additional data is available from the Monitoring menu.

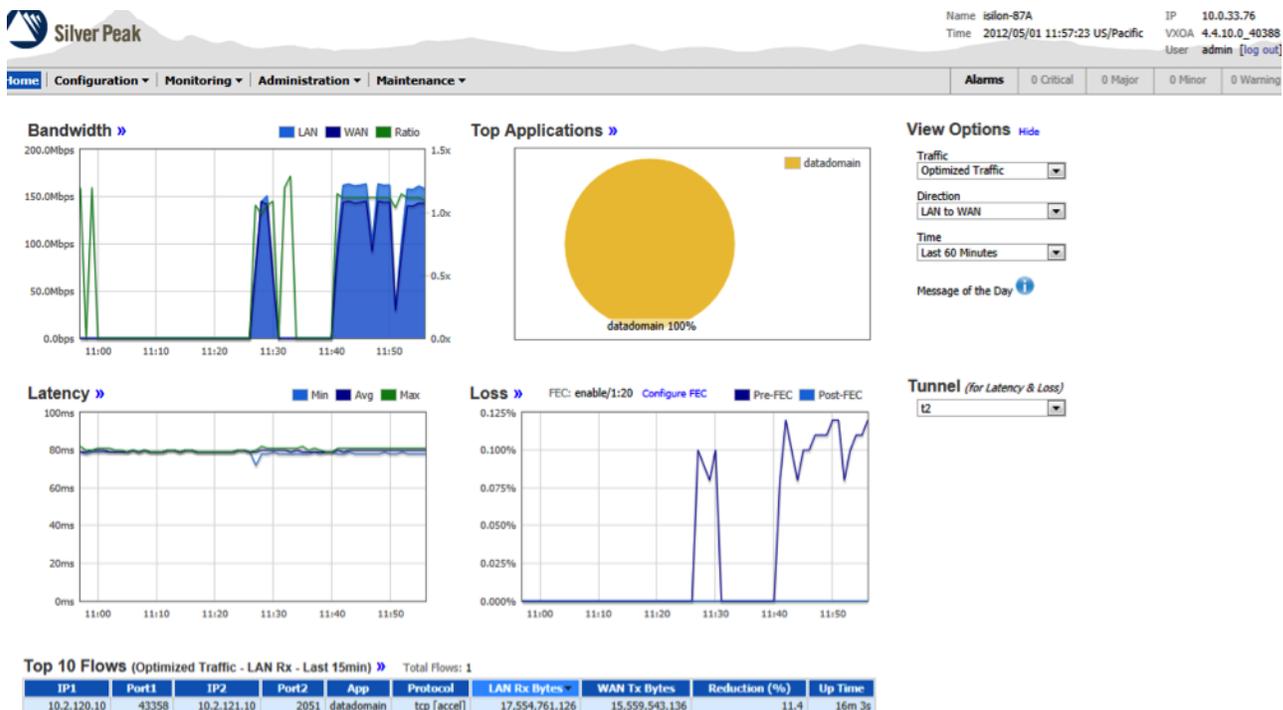


Figure 11 Silver Peak Home Screen

The Silver Peak Monitoring – Bandwidth chart shows real-time and historical information for LAN throughput, WAN throughput, and the bandwidth reduction ratio. LAN throughput is the amount of bandwidth that the Data Domain system is sending directly from the array. WAN throughput is the amount of data that the Silver Peak appliance is sending across the WAN. Reduction ratio is the amount of bandwidth reduction that VXOA is able to provide and measures the difference between the LAN and WAN throughput.

Note: When loss is present on the WAN, the reduction ratio (%) will be reduced. In rare circumstances it is possible for VXOA to report that more traffic is sent over the WAN when compared to the LAN. This traffic looks like a negative reduction percentage. This report is due to parity packets being sent to repair loss or when Forward Error Correction is disabled and VXOA has to retransmit packets that are lost. In this case, VXOA is not increasing the size of the Data Domain traffic; the reporting is accurately reflecting the amount of the data that has been sent across the WAN, including the retransmitted and parity packets.

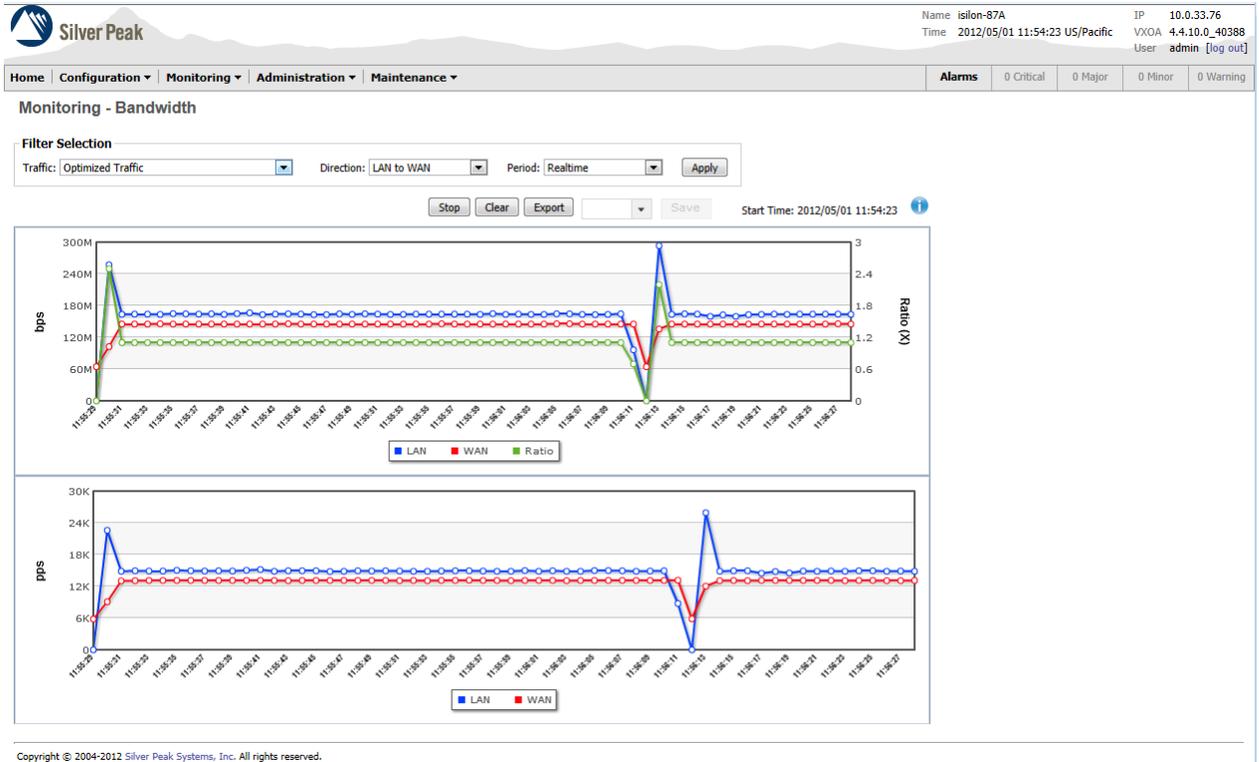


Figure 12 Silver Peak Bandwidth Monitoring

Data Domain Performance Monitoring

The Data Domain System Manager GUI provides real-time information on the Network usage for all Ethernet interfaces and the MB/s in and out for replication (see Figure 13). Replication specific information is available by selecting Replication > Performance.

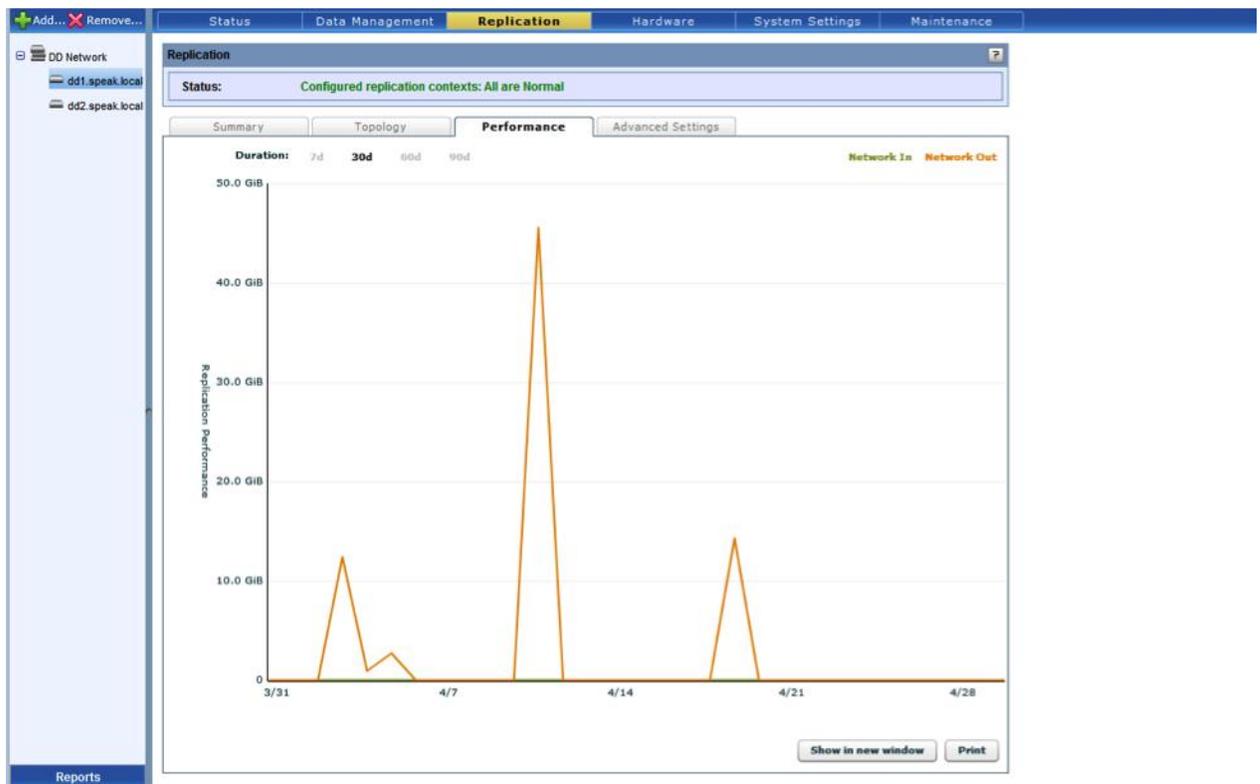


Figure 13 Data Domain Replication Performance