Silver Peak

AN OVERVIEW OF SILVER PEAK'S WAN ACCELERATION TECHNOLOGY

TABLE OF CONTENTS

Understanding WAN Challenges

Every Wide Area Network (WAN) connection has unique challenges. In some instances, there is simply not enough bandwidth. In other instances, the bandwidth that does exist cannot be utilized effectively because packets are being dropped and/or delivered out-of-order (a problem common in MPLS and IP VPN deployments). In still other scenarios, the biggest challenge is WAN latency, as long distances and "chatty" protocols make it difficult to communicate in real-time.

Bandwidth, latency, and loss all have an impact on "effective throughput" across the WAN. As a result, large enterprises must deploy a WAN acceleration solution that is capable of handling all three of these challenges in a secure and scalable manner. Silver Peak has the technology capable of achieving that objective.

Silver Peak Technology

Silver Peak provides a robust WAN acceleration 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.

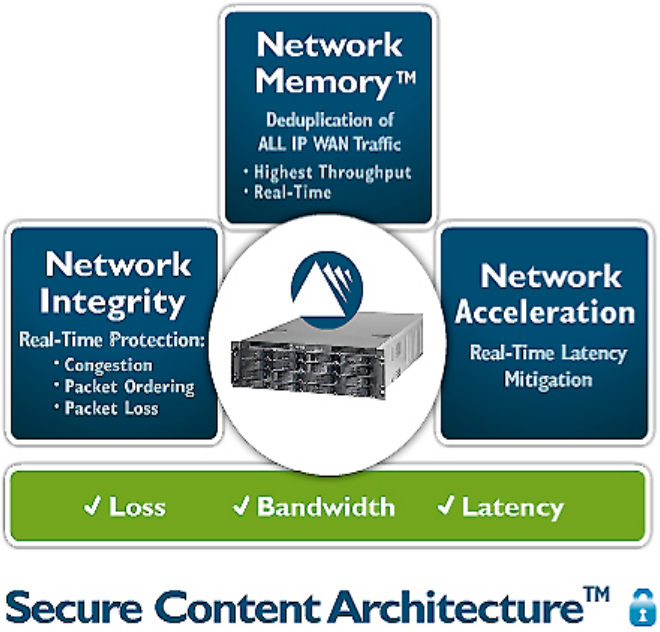

Silver Peak's network based solution addresses bandwidth, latency, and loss across the WAN.

Silver Peak NX Series appliances leverage the following technology components to accelerate all enterprises applications in a secure and reliable fashion.

Network Memory™ — Maximize Bandwidth Efficiency

Network Memory™ is Silver Peak's solution for overcoming WAN bandwidth limitations. It leverages patent-pending deduplication technology to eliminate the transfer of repetitive data across the WAN, resulting in maximum WAN utilization.

As packets flow through the NX appliances, Network Memory technology looks at the information and stores it as a local copy, or instance, in disk drives located on the appliances. As new packets arrive at an NX appliance, Network Memory computes fingerprints of the data contained within the packets, and checks to see whether these fingerprints match data that is stored in the local instance of the NX appliance at the destination location. (This is all done at the byte level). If the remote NX appliance contains the information, there is no need to resend it over the WAN. Instead, specific start-stop instructions are sent to deliver the data locally.

In addition to deduplication, Network Memory incorporates a variety of standard compression techniques for optimum "first pass" gains. This includes cross-flow payload and header compression using LZ and other technologies.

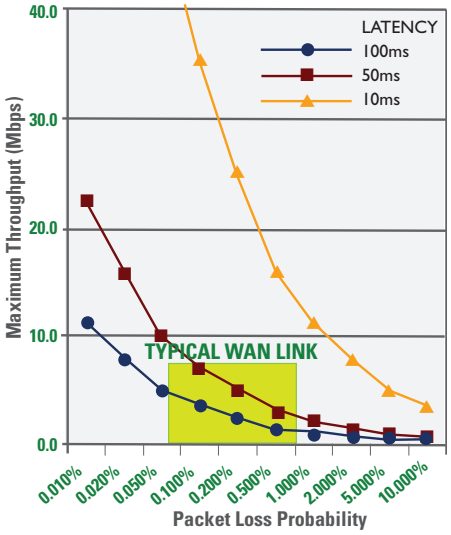

Bandwidth, latency, and loss have a significant impact on effective throughput across the WAN.

How is Silver Peak Different?

Silver Peak's Network Memory is different from other WAN deduplication solutions in the following ways:

- **Works on all IP Applications** – Silver Peak offers the only WAN deduplication solution that works at the IP layer of the OSI stack. This enables Silver Peak's deduplication to work on all IP traffic, regardless of transport protocol – from traditional TCP applications to those that run over UDP, proprietary or encapsulated protocols.

- **Byte Level Granularity** – Silver Peak uses an instruction-based approach to data referencing that leverages specific start-stop instructions to indicate where duplicate data can be found and retrieved. This is much more granular than alternative token based approaches to deduplication, whereby tokens are created by a hash function to represents "segments" of data. The Silver Peak architecture provides true byte-level granularity, the benefit of which is most apparent in applications where a change to a dataset is spread randomly throughout the original dataset, as is the case with AutoCAD, Microsoft Excel, and video streaming.

- **Scalable** – Silver Peak's Network Memory does not rely on the underlying operating system to perform TCP re-assembly and re-segmentation operations. This enables the Silver Peak solution to support hundreds of thousands of simultaneous user sessions (ie TCP flows), which is significantly higher than competitive offerings.

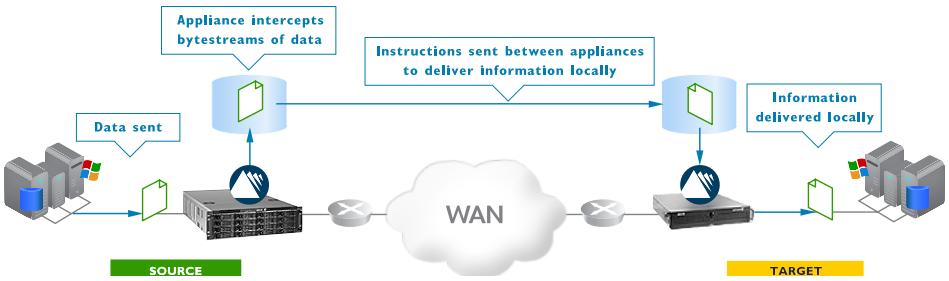

Network Memory uses real-time byte-level deduplication to eliminate the transfer of repetitive data across the WAN.

"By Maximizing bandwidth utilization, Silver Peak improved our backup capabilities without requiring costly upgardes. Of all the vendors we tested, only Silver Peak could handle the volume of traffic and number of flows in our environment."
— Transplace

- **Low Latency** – Silver Peak performs pattern matching in real-time on each individual IP packet, enabling deduplication to take place with very little latency (< 1 ms on average). This is in contrast to other deduplication solutions that add significant latency by buffering traffic prior to matching patterns. The result is that Silver Peak is the only vendor capable of deduping real-time and interactive applications, like Citrix, video streaming, and Remote Desktop Protocol (RDP).

- **Maximum Data Store Efficiency** – Silver Peak stores a single instance of information for all offices connected to an appliance, which maximizes the data store in each NX appliances. Some other solutions store a separate instance of information for each WAN link, resulting in a large disparity between stated disk capacity and effective capacity.

Network Integrity — Overcome Congestion and Packet Loss

Even when the physical-layer of a WAN is 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 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%, with averages in the .1% to 1% range. When packet loss is .5% or higher, it is very difficult to get more than 10 Mbps of effective throughput per flow across the WAN (regardless of the size of the WAN link). This can cause significant problems for many applications, like data replication, where high sustained data throughput is required on a per flow basis.

Silver Peak employs a variety of real-time techniques to address packet delivery issues, which include:

- **Forward Error Correction (FEC)** is a technology that is well known for its ability to correct bit errors at the physical-layer. Silver Peak adapted this technology to operate on packets at the network-layer to improve application performance across WANs that have high loss characteristics.

  Packet-Level FEC works by adding an error-recovery packet for every "N" packets that are sent. This FEC packet contains information that can be used to reconstruct any single packet within the group of N. If one of these N packets happens to be lost during transfer across the WAN, the FEC packet is used on the far end of the WAN link to reconstitute the lost packet. This eliminates the need to retransmit the lost packet across the WAN, which dramatically reduces application response time and improves WAN efficiency.

  The Silver Peak solution dynamically adjusts the FEC overhead in response to changing link conditions for maximum effectiveness in environments with high packet loss.

- **Packet Order Correction (POC)** is used to re-sequence packets on the far end of a WAN link "on the fly" to avoid re-transmissions that occur when packets arrive out of order. By performing the functionality in a dedicated WAN optimization device (as opposed to an end station or router), enterprises have the scalability needed to handle high volume, high throughput data streams with minimal added latency. POC is performed in real-time and across all IP flows (regardless of transport protocol).

"Silver Peak's Network Integrity features are indispensable for real-time applications like VoIP and Citrix."

— Linklaters

The figure below shows a detailed report pulled from the Silver Peak Global Management System (GMS), showing actual packet loss before and after Silver Peak FEC and POC. In this enterprise, packet loss is spiking to over 6% without Silver Peak. With Silver Peak in the equation, packet loss is kept consistently between 0% and .1%. The Silver Peak solution has enabled this enterprise to deliver real-time applications across their WAN, including VoIP, streaming video, and Citrix XenApp.

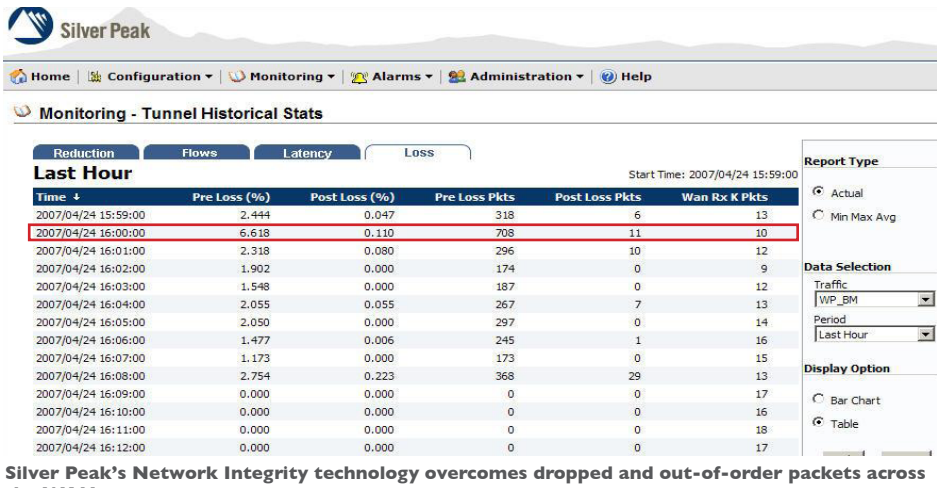

Silver Peak's Network Integrity technology overcomes dropped and out-of-order packets across the WAN.

- **Quality of Service (QoS)** techniques are used to maximize WAN utilization when WAN links are oversubscribed. Silver Peak's QoS implementation consists of an extensive deep packet inspection based classification engine and advanced queuing and service engine.

The classification engine can distinguish traffic based on any combination of source IP, destination IP, source port, destination port, incoming DSCP settings, and application type. The application

"With Silver Peak, we can better allocate our existing bandwidth resources and prioritize key applications, such as VoIP."

— Triquint Semiconductor

type is determined not based just on the port, but on stateful deep inspection of packets. This enables Silver Peak to track ephemeral ports for applications like FTP, Cisco Skinny, H323, SIP, Microsoft Port Mapper, Sun RPC, and more by monitoring the control channel for these protocols and identifying new dynamic connections as they are established. Classification decisions based on multiple combinations of these factors can be implemented with application-aware ACLs. Policies are used to assign different traffic types to different traffic classes.

Each traffic class can be assigned by the following parameters:

- Absolute priority
- Maximum and minimum bandwidth
- Excess bandwidth allocation weight
- Maximum number of packets/bytes per class
- Maximum number of packets/bytes per flow
- Maximum weight time

By supporting arbitrary adds and rewrites of the DSCP field, Silver Peak solutions can be configured to work seamlessly with any particular service providers' MPLS QoS class assignments. This includes the ability to translate between an interior (enterprise) an exterior (service provider) QoS maps.

How is Silver Peak Different?

Silver Peak is the only vendor to address Network Integrity issues in real-time with its FEC and POC technologies. Other WAN optimization technologies rely on aggressive TCP retransmission methods to address packet loss, which do not work with time-sensitive applications like voice and video because they add additional latency. Furthermore, these methods only work on TCP traffic, rendering them useless for applications that leverage UDP and other protocols. Lastly, aggressive retransmissions are often quite "unfriendly" to other applications trying to use the WAN at the same time as it consumes all available WAN resources.

Silver Peak provides a wide array of management tools that make it easy to setup its Network Integrity features and monitor the ongoing effect that they have on WAN performance. This includes a variety of advanced templates for QoS policy configuration and management, as well as detailed reporting tools to monitor dropped and out-of-order packets across the WAN.

Network Acceleration — Mitigate WAN Latency

The time it takes for information to go from sender to receiver and back is called network latency. Since 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

CONFIGURING QOS POLICY

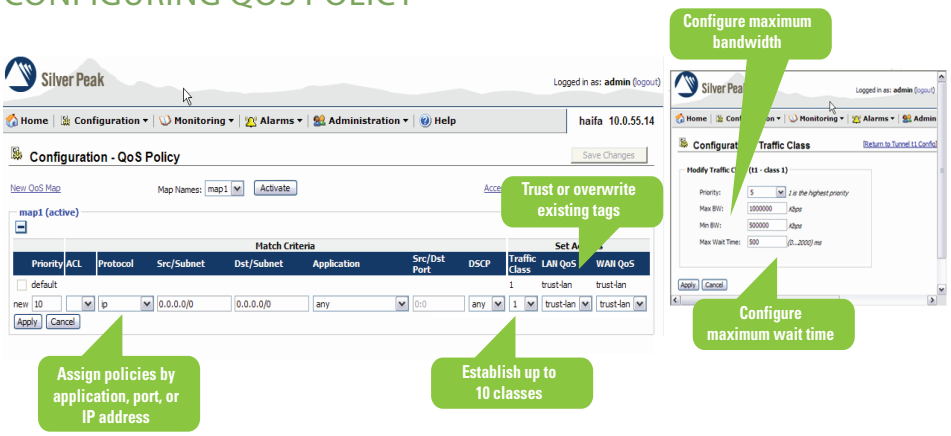

Silver Peak provides easy to use templates for configuring advanced QoS policies.

functions, additional latency is added to the equation. Below is what typical enterprises will experience by way of network latency:

LATENCY	
LAN	5 ms
Western US to Eastern US	80-100 ms
International	100-300 ms
Satellite	500+ ms

Network latency becomes even more problematic when higher level protocols are "chatty", as is the case with TCP and Microsoft CIFS. These protocols involve numerous acknowledgements before sending data, which can severely impact transfer times.

Silver Peak offers a variety of "TCP acceleration" techniques to mitigate WAN latency, which include:

- **Window Scaling** – Silver Peak utilizes the TCP Window Scale option to deliver window sizes as large as 1GB thus overcoming the throughput limitation imposed by the standard 64KB TCP window size in higher latency networks.
- **Selective Acknowledgement** – Selective acknowledgements (SACK) provide a mechanism for handling multiple packet loss in a WAN environment. Unlike the normal process whereby a cumulative acknowledgment is provided across all TCP packets, selective acknowledgments give the sender a complete picture of which segments are queued at the receiver and which have not yet arrived, so the sender only needs to retransmit the missing data segments.
- **Round-Trip Measurement** – Silver Peak utilizes a proprietary round trip measurement scheme that enables RTTs to be calculated more efficiently. This leads to more accurate RTO (retransmission timeout) measurements which, in turn, improves throughput.
- **HighSpeed TCP** – Silver Peak has implemented HighSpeed TCP, which is a modification to TCP's congestion control mechanism for use with TCP connections with large congestion windows. It alters how the window is opened on each round trip and closed on congestion events as a function of the absolute size of the window. When the window is small, HighSpeed TCP behaves exactly like ordinary TCP. But when the window is large, it increases the window by a larger amount and decreases it by a smaller amount, where these amounts are chosen based on the precise value of the window in operation. The effect of these changes is that TCP achieves high throughput with more realistic requirements for packet drop rates. Equivalently, HighSpeed TCP has more realistic requirements for the number of round-trip times between loss events, enabling TCP to perform better in high-bandwidth, high-latency environments. When HighSpeed TCP is used in conjunction with Silver Peak's Network Integrity features, enterprises have a complete arsenal of tools for improving performance in high loss WAN environments.

"By overcoming latency challenges across our WAN, Silver Peak reduced transfer times from hours to several minutes between the US, Canada and Europe."
— Embarcadero Technologies

Silver Peak also offers various techniques to overcome latency caused by the Microsoft CIFS protocol. At a high level, these techniques include:

- **CIFS Read-Ahead** – When a user is working with a file, Silver Peak appliances generate read ahead requests within the file in order to pipeline operations to the server, thus eliminating the round-trip delay associated with waiting for acknowledgement. This minimizes the latency associated with read operations.
- **CIFS Write-Behind** – Silver Peak NX appliances pipeline write operations on behalf of a client, thus eliminating round-trip delays associated with waiting for acknowledgements. This minimizes the amount of round trips required to perform write operations, improving performance without risking data integrity.
- **CIFS Meta Data Optimizations** – Standard Microsoft clients make many unneccessary, independent, and duplicative requests for file and directory meta-data. Performance is enhanced by locally consolidating and responding to these requests.

How is Silver Peak Different?

Silver Peak

Silver Peak is able to perform its Network Acceleration functions at very high data rates. With appliances that scale up to 1 Gbps on the WAN side and 10 Gbps on the LAN side, Silver Peak can overcome latency on even the largest of WAN links.

Silver Peak also combines standard Network Acceleration features with its Network Integrity features to address the combined impact of loss and latency. The two of these issues together have a "saw tooth" effect on WAN performance, which can result in erratic effective throughput. By addressing both of these issues in tandem, Silver Peak offers a robust solution for WAN acceleration.

Secure Content Architecture™ - Real-Time Data Protection

Silver Peak keeps enterprise data secure with its Secure Content Architecture. This is achieved via the following capabilities:

- **Disk Encryption** – Silver Peak employs 128 bit AES encryption to protect all data stored on NX appliances. This is the only fail-safe way to ensure the privacy of information in the event that an entire device (or individual hard drive) is stolen, compromised, or changes hands.
- **Secure Transmission (IPsec)** – Silver Peak supports 128 bit IPsec (using AES for encryption) to ensure that data is secure as it is transferred over the WAN.
- **Secure Socket Layer (SSL) Acceleration** – Silver Peak accelerates SSL traffic using a variety of techniques, including Quality of Service (QoS) and TCP acceleration. By offering 128 bit encryption for data on the disk and on the wire, with hardware acceleration, additional SSL features can be added via a software upgrade.
- **Centralized Control** – Silver Peak's Global Management System (GMS) enables advanced authentication policies to be centrally configured and enforced. This includes "peer authentication", whereby only valid Silver Peak appliances are allowed on the network, and connectivity can only be established between trusted Silver Peak devices. These features protect an enterprise from session hijacking or man-in-the-middle (MiM) types of attacks that can compromise WAN acceleration solutions.
- **Secure Access** – Access to all Silver Peak devices is tightly controlled using TACACS+ and RADIUS. This ensures complete AAA protection, including user tracking and auditing per-command authorization, and group based authentication privileges. Enterprises can use existing databases to facilitate administration and avoid potential security holes.

LATENCY AND LOSS

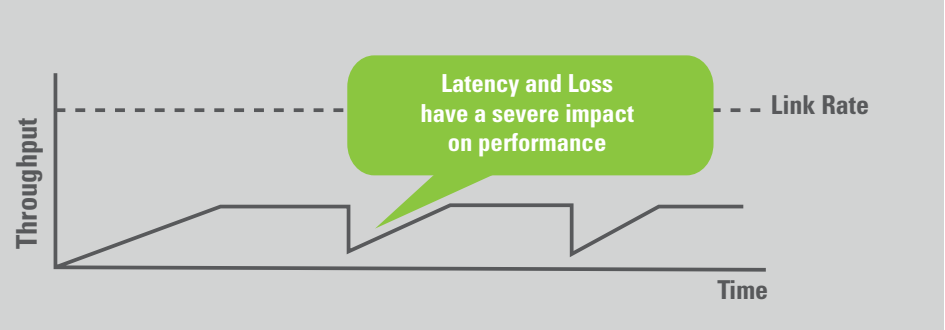

Periods of packet loss can have a "saw tooth" effect on WAN throughput, exacerbating the ill affects of latency.

Future-Proof Your WAN

Bandwidth, latency, and loss can all impact WAN performance. Silver Peak has a robust WAN acceleration solution that lets enterprises address each of these WAN challenges in a secure and scalable manner.

By offering a predominantly network-based solution to WAN acceleration, Silver Peak addresses WAN performance challenges right at the source. This not only ensures the best performance across the widest range of enterprise applications, but it provides a future-proof architecture designed to handle the unique requirements of every WAN link in every enterprise location.

"We can pump as much secure data as we want through the Silver Peak boxes without impacting scalability or performance."
— Global Professional Services Company