silver peak™

White Paper

Optimize Every SaaS Application with an Intelligent WAN

Optimize Every SaaS Application with an Intelligent WAN

Widespread adoption of SaaS has vastly expanded enterprise IT. Silver Peak's Unity WAN fabric gives IT the visibility and control that has been missing, while giving users the performance they need.

It's no secret that Software as a Service (SaaS), and the cloud in general, offers the enterprise a tremendous opportunity. SaaS delivers the breadth, simplicity and affordability of packaged applications in a fraction of time and for a fraction of the cost of traditional models.

But SaaS also changes the dynamics of IT. As anyone can sign up for a SaaS application, purchasing authority is often spread throughout the organization. IT can no longer guarantee suppliers meet corporate standards or use the company's size and power to negotiate better terms and conditions.

IT also can do little to ensure the performance of SaaS, which relies on the Internet – the unpredictable, vortex of networks. Congested one day and tolerable the next, Internet "weather" undermines the consistency of SaaS delivery, even in regions with highly developed Internet infrastructure. Predictability and consistency, critical factors for the success of any business, are the very antithesis of the new cloud-driven Internet (see Figure 1).

CIOs must rethink how to complement SaaS with the IT organization's core values – strategic alignment with business objectives, risk management, resource management, and more. This is particularly challenging given that many CIOs may not even know the extent of SaaS penetration within their organization. More than 60 percent of business units reported circumventing their own IT shops and going directly to the cloud. According to a study conducted for McAfee by Stratecast, a unit of Frost & Sullivan, 81 percent of line-of-business workers and 83 percent of IT staff admit to using non-approved SaaS apps. Detecting those SaaS apps is often impossible as they are hidden from legacy monitor and control systems.

Silver Peak Unity is an intelligent wide area network (WAN) fabric that enables IT to achieve something that's never been possible: the ability to monitor and control connectivity to the cloud while ensuring consistent SaaS performance (as well as any enterprise application). Unity dynamically compensates for Internet weather. SaaS applications perform far better as latency issues are solved, packet loss and other effects of network congestion are minimized or eliminated, and bandwidth requirements are reduced by as much as 40x. With Unity, the enterprise becomes SaaS-optimized.

SaaS Changes IT Governance and Policies

SaaS forces IT to rethink how it governs and delivers enterprise applications. Traditionally, IT policies and processes assumed full control over the applications, data, and compute environment. Regulatory compliance, SLA adherence, creation of new IT products, and integration with existing ones were all possible because IT could allocate resources as necessary to meet business objectives. IT had full control.

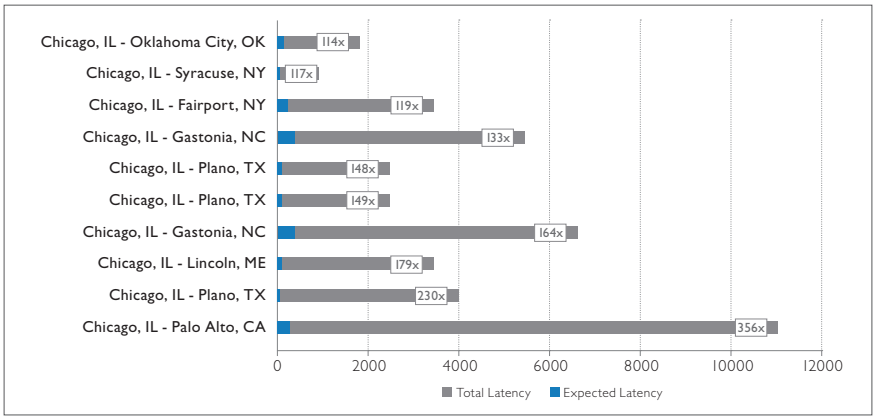

Figure 1: Although Internet infrastructure is highly developed within the US, latencies between cities can vary widely on any given day (Source: Internetweathmap.com)

[1] "2013 Cloud Adoption Rate Survey," 2nd Watch, Inc.

With SaaS, applications and data shift outside of the enterprise data center to the SaaS provider. IT loses control of the application environment, which raises questions as to how CIOs can fulfill their responsibilities around strategic alignment, risk management, resource management, accounting, budgeting, and value delivery:

Strategic Alignment – Users may be drawn to a given SaaS offering because of the interface and its features, but IT must consider the longer term. Questions around vendor stability must be addressed:

- Can the SaaS infrastructure scale to meet customer requirements?
- What operational processes need to change to accommodate SaaS or cloud adoption?
- What are the actual tradeoffs, in detail, and how would these impact the business overall?
- Invariably, as business units adopt SaaS, they will want to integrate that service with enterprise data and applications. Who ensures that such an integration will be possible?

Risk Management – Enterprises need to comply with industry and government regulations, and they need ways to ensure any SaaS offering adopted by their users meets those regulations. At the most basic level, IT needs a way to ensure users adopt compliant services and that those services offer guarantees around privacy, ownership and use (or lack thereof) of the enterprise data stored in that respective service. IT will want to understand:

- Where does enterprise data reside within the SaaS offering and can IT secure and delete that data if necessary?
- Does the SaaS provider implement the required standards and procedures to comply with industry and government regulations? How does IT maintain compliance?
- What guarantees, processes or tools are in place to facilitate the removal of enterprise data in the event a SaaS provider fails or discontinues its service?

Resource Management – Aside from the costs of personnel and computing systems, IT also manages vendor relationships, negotiating terms and contracts, and leveraging its size to extract more favorable terms from those providers. As users subscribe to SaaS services on their own, enterprises lose that leverage and one single point of contact.

- How does the IT organization maintain a more strategic relationship with the service provider?
- Does it make sense to add the SaaS provider to the company's list of sanctioned vendors?
- What kind of controls can be put into place to extract the most favorable terms for the enterprise and monitor contractual compliance?

Accounting and Budgeting – Departmental and user adoption of SaaS poses significant challenges from an accounting and budgeting perspective. The actual spend on cloud services is often unclear. A recent Executive Guidance report published by CEB found that non-IT executives are spending an additional $0.40 on technology for every $1 managed in the corporate IT budget, also known as "Shadow IT."

- How do organizations ensure the proper accounting and planning so those dollars are in alignment with business strategies and goals?
- How does the procurement process and guidelines change?
- What other hidden costs, such as network consumption, are being driven by SaaS?

Value Delivery – By overseeing application deployment and acquisition, CIOs are in the position to advise business units on options and best practices. This knowledge exchange can only be done through engagement and ongoing communication across the business. But if SaaS acquisition and deployment occurs beyond its domain, how can IT continue to gain the expertise to fulfill that role of the "trusted IT advisor?"

Then there is the tactical challenge of SaaS performance. If the IT organization is to include SaaS in its package of sanctioned services – and it should – performance expectations must be met. Even if users circumvent IT and subscribe to cloud services directly, IT will need a way to help users achieve a faster, more consistent SaaS experience. When users grow frustrated waiting for screen refreshes, or when files take too long to transfer, the fingers will start pointing to IT. And regardless of responsibility or fault, the entire business suffers.

Internet Weather Undermines SaaS Performance

Ensuring SaaS performance over the Internet is far more complicated than conventional applications that run over your MPLS or private network. If a conventional application has a performance problem, the IT organization typically has the right levels of visibility and control to respond to trouble tickets. Compute or storage resources can be added to improve performance, and if the network is the issue, organizations can reposition servers closer to the office or add bandwidth.

SaaS is more complicated. For one, IT managers may not even be aware that SaaS applications are running on the network. Many conventional monitoring systems cannot differentiate between the various Internet applications and services sharing Transmission Control Protocol (TCP) port 443, the port used by many SaaS applications.

Even if they can identify the SaaS application, IT managers may be unable to improve its performance. With compute and storage resources outside of the IT manager's control, adjusting them is impossible.

Good news: SaaS providers typically design their offerings for availability and scale. Bad news: neither the SaaS provider nor IT is able to appreciably improve the network.

And it is the network that is so critical to SaaS. Significant bandwidth is often consumed by SaaS applications, in part because they rely heavily on excessive use of Cascading Style Sheets (CSS) and downloading large JavaScript (~5 Megabyte) files to provide a desktop-like experience. Third-party plugins exacerbate the issue, gathering and pushing data across the connection for marketing and sales purposes.

Across relatively short connections with small amounts of delay, these issues may not be noticeable. But as SaaS traverses longer distances with larger delays, the throughput of the underlying TCP session decreases dramatically. The issue, commonly called TCP's "long fat pipe problem," has to do with how data is sent across TCP. TCP hosts gradually increase their transmission rate until reaching the maximum number of packets that can be sent without dropping a packet, also known as the TCP window. Once the packets of the TCP window are sent, the host waits for an acknowledgement of receipt from the destination before sending the next window of packets.

Over long distances, the time spent waiting to receive the TCP acknowledgement limits the throughput of a session. In fact, regardless of the amount of bandwidth, TCP throughput will continue to be limited by the latency of the session. A perfect 10 gigabit-per-second (Gbps) coast-to-coast session (about 75 milliseconds), for example, will still perform like a 20 Mbps connection simply because of the session's delay (see Figure 2).

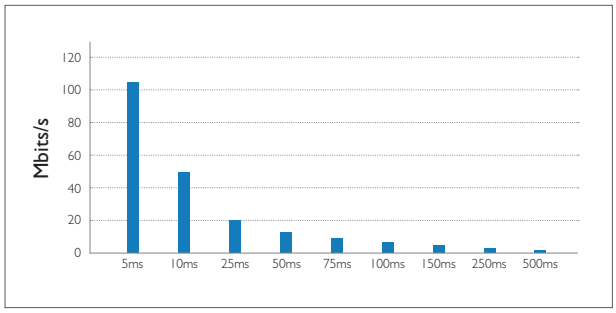

Figure 2: The Impact of Latency

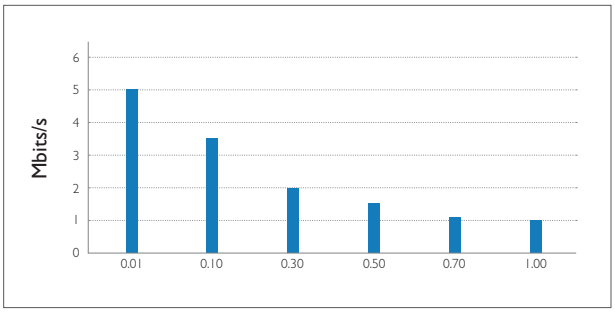

Figure 3: The Impact of Loss

Delay is only exacerbated by the nature of Internet routing and peering. Today's routing protocols do not consider end-to-end delay when delivering an application's packets. Route calculations are designed to move packets forward and not necessarily along the most direct path. This is particularly true as routing decisions are being made by service providers whose interests are not necessarily aligned with the enterprise. A service provider may hold traffic on its own network, even when there is a better path available via an alternative provider. There may also be financial incentives for a service provider to dump traffic off its network as early as possible even though it can provide a better path.

Delay is not the only challenge. Internet connections are always congested, dropping packets or delivering them out-of-order. As TCP encounters packet loss, the transmission throughput exponentially decreases, further undermining session performance. A TCP session across the same 10 Gbps coast-to-coast connection could be limited to less than 2 Mbps (see Figure 3). When combined together, latency and packet loss wreak havoc on SaaS connections.

Even in regions with well developed, dense Internet infrastructures, such as the US, latency and loss can change significantly from day-to-day. Research by Keynote, an Internet testing authority, revealed that over a two-week period, the average packet loss across hundreds of routes was just .34 percent. More than half of those routes exhibited average packet loss rates in excess of one percent for at least one 24-hour period. However, packet loss fluctuations we not confined to inter-backbone traffic. Some backbones showed packet loss rates peaking at 13 percent.

Latency was no better. Over the same two-week interval, data from the Internetweathermap.com, another Internet research firm, found that latency on some routes can fluctuate widely. The most extreme case involved two well-traversed locations – Chicago, IL to Palo Alto, CA – with maximum delay exceeding 11,000 milliseconds or 350x more than the minimum measured delay. (Download the Internet Weather Report for more information)

Enterprise Networks Are Ill-Suited for SAAS

Enterprise backbones do not help the SaaS problem and often make the issues worse. Many organizations centralize Internet access for security and management reasons. Branch office Internet traffic is backhauled across an MPLS network to the data center and then delivered through the central Internet connection (see Figure 4).

This is an inefficient architecture for SaaS for numerous reasons. Additional latency is often added traveling across the MPLS backbone when delivery directly over the Internet may be quicker. The economics of such an approach are also dubious as the organization must use expensive bandwidth to accommodate a significant amount of low-priority data. Any burst in Internet traffic will also increase MPLS WAN costs.

Today's routers also lack the necessary application-awareness to choose the right network for a given application. They cannot dynamically determine, for example, whether it is shorter to route SaaS traffic through a native Internet connection or across the MPLS network (see Figure 5).

The old static ways of determining the best path no longer work well in an environment where the enterprise's "effective WAN" includes large portions of the Internet. Without a new way of addressing this, a lot of traffic is routed sub-optimally with no knowledge of the end-to-end service or application being carried.

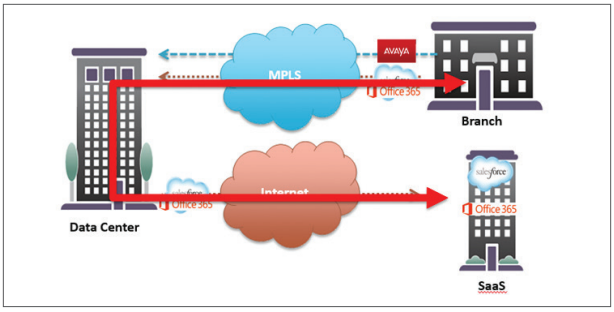

Figure 4: The Enterprise and SaaS

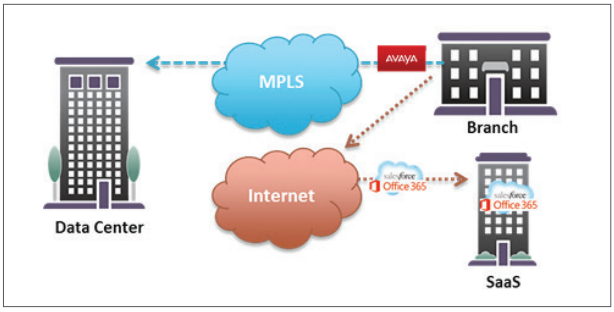

Figure 5: The Hybrid WAN

[2] "Netskope Cloud Report" July, 2014

Current Approaches are Insufficient

SaaS providers, particularly the larger ones, have tried to address the issue by storing a customer's data in the closest data center. For example, a Paris-based company signing up for Microsoft Office 365 will likely have its data stored in Microsoft's Ireland data center, per Microsoft's terms of service.

But problems still exist. Many elements of a requested page will still be served from other locations in the service provider's global network. And having data hosted within the region of enterprises headquarters is of no help when users are located outside the region.

Conventional WAN optimization architectures fail to address the challenges of optimizing SaaS on several accounts:

- Many approaches only optimize one or two SaaS applications, when the average enterprise will utilize 508 SaaS applications.
- They require situating an end-point in or near the SaaS data center, but that is typically impossible.
- Even then they fail to detect changes in loss or latency. As such, can they cannot correct for the effects of those conditions, nor can they choose an alternative path better suited for the SaaS traffic.

WAN optimization end-points are also ill-suited to the dynamic nature of SaaS. Part of the attraction of SaaS is its low cost and ease of deployment. WAN optimization end-points traditionally require proprietary hardware, particularly when encryption is required. Hardware only interferes with the SaaS model. It's expensive and made more expensive when requiring a subscription to a content delivery network (CDN). It also delays overall deployment times as enterprises must test, configure, ship, and deploy the hardware appliance.

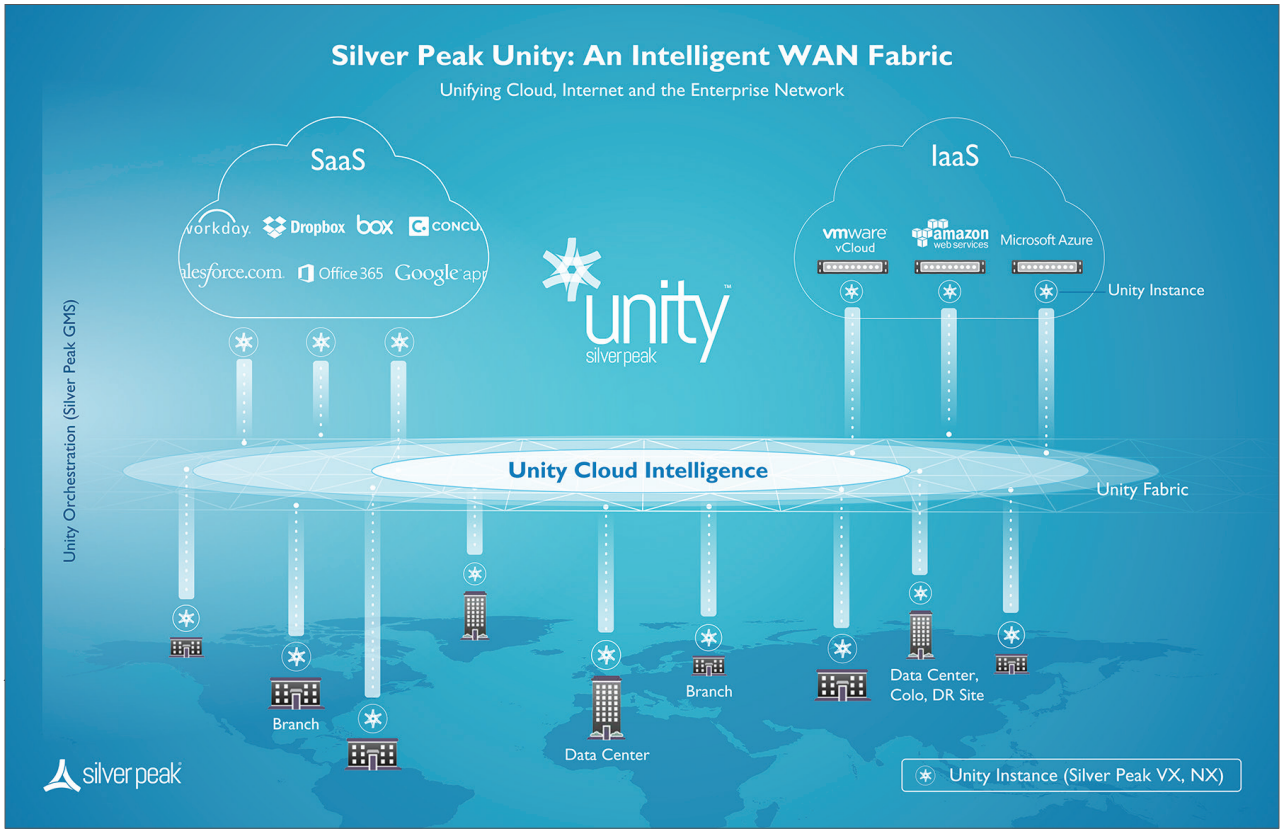

Figure 6: The Unity WAN fabric is formed from Silver Peak instances in the cloud and at the customer premises.

Session scalability becomes an enormous challenge for WAN optimization when supporting a multitude of SaaS applications, many of which rely on parallel sessions between the client and server for performance reasons. Conventional WAN optimization hardware, though, is limited in the number of simultaneous sessions. As a result, many SaaS sessions will be left unoptimized even when bandwidth is available, forcing IT to upgrade to a more expensive end-point.

Newer software-defined WAN offerings address some of these issues, but they lack the acceleration and control capabilities of WAN optimization. Not only do many SDWAN solutions fail to fix the effects of network congestion, but they also lack deduplication and compression techniques used by WAN optimization to overcome bandwidth problems. If encryption is offered, it's often at the expense performance.

Silver Peak Unity Optimizes Enterprise SaaS

Silver Peak Unity is the first solution to solve the performance challenges every SaaS application, while also giving IT managers visibility and control of how the network is being used, including for SaaS.

Silver Peak software installed in data centers, branch offices and cloud interconnection hubs generates the Unity fabric, a network overlay that controls and accelerates connectivity to any combination of enterprise services and SaaS applications. Each Silver Peak instance on the Unity fabric communicates with Silver Peak's Cloud Intelligence service, which aggregates constantly changing information about cloud providers and Internet weather. Unity uses this information, along with calculations from each software instance, to dynamically route traffic to the cloud over the optimal path.

Orchestration is handled by Silver Peak's Global Management System (GMS), which provides IT managers with complete visibility and control over the deployment and use of cloud services (see Figure 6).

Silver Peak Unity combines unique technology for optimizing SaaS:

- **Advanced Exterior Routing:** Advanced Exterior WAN Routing identifies the closest egress to SaaS providers and directs traffic to the SaaS application over an optimal path
- **Advanced Interior Routing:** Advanced Interior WAN Routing dynamically selects the fastest, least-congested, or most available path for traffic by monitoring packet loss, latency, and bandwidth in real-time
- **Cloud Intelligence:** This subscription service delivers information about cloud services to Silver Peak software instances, creating an Internet weather map for intelligent routing decisions
- **Accelerated Encryption:** The Unity fabric is built on IPsec VPN security without any performance degradation, using AES-256 encryption with SHA-1 authentication and a simplified deployment model
- **Data Reduction:** WAN compression and deduplication inspects all traffic in real-time, eliminating repetitive transmission of duplicate data
- **Path Conditioning:** Adaptive Forward Error Correction reconstitutes dropped packets in real-time while Packet Order Correction re-sequences packets that may traverse multiple paths across the network
- **Traffic Shaping:** SaaS applications are classified to prioritize critical traffic classes while constraining recreational or personal-use SaaS traffic classes.
- **Global Visibility:** Web-based Global Management System provides centralized orchestration of WAN deployments, advanced application classification, and detailed SaaS performance metrics.

Silver Peak Unity adds awareness of the subnets and IP addresses used by the SaaS providers. Every Silver Peak instance can then measure its loss, latency and other metrics to these subnets and distribute that information to the Silver Peak instances across the Unity fabric. Those instances can then determine the optimal end-to-end path for connecting to the SaaS application.

Silver Peak Unity is built to optimize every SaaS application, including popular applications like Microsoft Office 365 and Salesforce.com. And through Silver Peak's Cloud Intelligence service, Silver Peak is constantly adding support for new SaaS applications.

IT managers also gain an accurate picture of the extent of SaaS penetration. Silver Peak software-based instances detect and identify SaaS traffic on the enterprise network, and specifically what SaaS applications are being used. With that information, IT can address SaaS usage on a number of fronts.

At the highest level, CIOs and the executive team can develop the proper policies and procedures for addressing new SaaS implementations. Tactically, IT teams can determine how best to treat that traffic. Is this a hosted-VoIP or video conferencing service? If so, delay and packet loss must be kept to a minimum. Is it a CRM or hosted e-mail service? If so, packet loss constraints might be more relaxed. These and many more factors are used to align SaaS requirements against existing network demands, and ultimately become the Quality of Service (QoS) and traffic shaping policies enforced by the Silver Peak instances. This enables IT to prioritize business-critical SaaS applications while constraining YouTube and other recreational services.

Information about the service is also used by Silver Peak's interior routing technology to determine the optimal route across the Unity fabric. The path with the least loss may be selected for hosted-VoIP or video traffic, while CRM or hosted e-mail might be sent along a path with higher loss, but lower latency.

For SaaS traffic, Silver Peak's exterior routing comes into play, and the destination of that path becomes the Silver Peak instance closest to the SaaS data center as determined by Silver Peak's Cloud Intelligence service. Instead of SaaS traffic bouncing around the Internet or enterprise networks, Silver Peak sends traffic to the cloud over the shortest or least-congested path.

Good for SaaS, Smart for IT

Not only does Silver Peak Unity and the SaaS optimization feature address the challenges faced by wide spread SaaS adoption, but it does so in a way that is in-line with IT requirements. No major upgrades or network changes are required to adopt SaaS optimization. Even new hardware is not necessary. Silver Peak uses software instances that can be deployed across every major platform and within cloud hubs or IaaS providers These instances are optimized for cloud services that require many simultaneous sessions. Unity instances can scale up to 256,000 simultaneous sessions, ensuring that all cloud sessions can be optimized.

The use of scalable software means that deployment of Unity instances is easy, and activation can occur in minutes. Silver Peak software can be downloaded from the Silver Peak software marketplace and deployed in as little as 20 minutes. The use of subscription-based software and a flat, low-cost subscription to cloud intelligence makes building a Silver Peak Unity fabric affordable for small businesses, yet powerful enough for the largest enterprises. And with comprehensive support for SaaS applications, Silver Peak can optimize SaaS as fast as the enterprise can spin it up.

The Silver Peak Unity fabric works across every corporate network, allowing IT to extend its SaaS intelligence across MPLS, private lines, Internet VPNs, and more. However the enterprise connects to remote offices, Silver Peak Unity secures and accelerates the connectivity while giving IT the monitoring and control it requires.