The 2015
Guide to WAN Architecture & Design

By ***Dr. Jim Metzler, Ashton Metzler & Associates***
Distinguished Research Fellow and Co-Founder
Webtorials Analyst Division

Sponsored in part by:

Produced by:

Table of Contents

Executive Summary

The wide area network (WAN) is a critically important topic for number of reasons. Those reasons include:

- The latency, jitter and packet loss that is associated with the WAN often cause the performance of applications to degrade;
- The WAN can be a major source of security vulnerabilities;
- Unlike most of the components of IT, the price/performance of WAN services doesn't obey Moore's Law;
- The outage of a WAN link often causes one or more sites to be offline;
- The lead time to either install a new WAN link or to increase the capacity of an existing WAN link can be quite lengthy.

A discussion of wide area networking is extremely timely because after a long period with little if any fundamental innovation, the WAN is now the focus of considerable innovation. As a result, for the first time in a decade network organizations have an opportunity to make a significant upgrade to their WAN architecture.

This e-book describes a hypothetical company, referred to as NeedToChange, which has a traditional approach to WAN design. It then presents Silver Peak's response to how NeedToChange should evolve its WAN. This e-book includes a summary of the key components of some of the emerging approaches to WAN architecture and design and concludes with a call to action that outlines a project plan that network organizations can use to evolve their WAN.

Introduction and Background

Definition of WAN

To many network professionals the term *WAN* doesn't refer to the Internet but refers exclusively to enterprise WAN services such as Frame Relay, ATM or MPLS. The distinction is that enterprise WAN services were designed primarily to connect a given enterprise's branch offices and data centers while the Internet provides connectivity to a huge range of resources with myriad owners. That is an arbitrary distinction that is quickly losing relevance and as a result throughout this e-book the term WAN refers to any combination of the Internet and enterprise WAN services.

WAN Evolution

The modern WAN got its start in 1969 with the deployment of the ARPANET which was the precursor to today's Internet. The technology used to build the Internet began to be commercialized in the early 1970s with the development of X.25 based packet switched networks.

In addition to the continued evolution of the Internet, the twenty-year period that began around 1984 saw the deployment of four distinct generations of enterprise WAN technologies. For example, in the mid to late 1980s, it became common for enterprise IT organizations to deploy integrated TDM-based WANs to carry both voice and data traffic. In the early 1990s, IT organizations began to deploy Frame Relay-based WANs. In the mid to late 1990s, some IT organizations replaced their Frame Relay-based WANs with WANs based on ATM (Asynchronous Transfer Mode) technology. In the 2000s, many IT organizations replaced their Frame Relay or ATM-based WANs with WANs based on MPLS. Cost savings was the primary factor that drove the adoption of each of the four generations of WAN technologies.

WAN Services

As discussed in The 2014 State of the WAN Report, network organizations currently make relatively little use of WAN services other than MPLS and the Internet and the use they do make of those other services is decreasing somewhat rapidly. That report also identified the concerns that network organizations have with those two services. Those concerns are shown in **Table 1** in descending order of importance.

Table 1: Concerns with WAN Services	
Concerns with MPLS	**Concerns with the Internet**
Cost	Security
Uptime	Uptime
Latency	Latency
Lead time to implement new circuits	Cost
Security	Packet loss
Lead time to increase capacity on existing circuits	Lead time to increase capacity on existing circuits
Packet loss	Lead time to implement new circuits
Jitter	Jitter

Traditional WAN Design

The traditional approach to designing a branch office WAN is to have T1 access to a service provider's MPLS network at each branch office and to have one or more higher speed links at each data center. In this design, it is common to have all or some of a company's Internet traffic be backhauled to a data center before being handed off to the Internet. One of the limitations of this design is that since the Internet traffic transits the MPLS link this adds both cost and delay.

One alternative to the traditional approach to designing a branch office WAN is to supplement the T1 access link in a branch office with direct Internet access and to also leverage technology such as Policy Based Routing (PBR). PBR allows network administrators to create routing policies to allow or deny paths based on factors such as the identity of a particular end system, the protocol or the application.

One advantage of this alternative design is that it enables network administrators to take Internet traffic off the relatively expensive MPLS link and put it on the relatively inexpensive Internet link. One disadvantage of this approach is that configuring PBR is complex, time consuming and error prone. Another limitation of this approach it that it creates a static allocation of traffic to multiple links which means that it isn't possible to reallocate the traffic when the quality of one of the links degrades.

Hypothetical Company: NeedToChange

Silver Peak was given the description of a hypothetical company, referred to as NeedToChange, that has a traditional WAN and they were asked to provide their insight into how the company should evolve its WAN.

Within the context of a traditional WAN there is a wide breadth of options relative to a company's WAN topology, services, applications and goals. As a result of this breadth, it wasn't feasible to cover all possible options in a reasonably sized description of NeedToChange's WAN. In order to limit the size of the description of NeedToChange's WAN and yet still bring out some important WAN options, Silver Peak was allowed to embellish the description of NeedToChange's WAN. They could, for example, add additional data centers or key applications; vary the amount of traffic that was backhauled; prioritize the factors impacting NeedToChange's WAN or identify business drivers such as the need to support mergers and acquisitions.

Below is the description of NeedToChange's WAN that Silver Peak received.

1. Data Centers

   NeedToChange has a class A data center in Salt Lake City, Utah. The site has two diversely routed T3 links into an MPLS network[1] and a 100 Mbps link to the Internet.

2. Traffic Prioritization

   In the current environment, traffic is prioritized in a static manner; e.g., voice traffic always gets top priority and it receives a set amount of bandwidth.

3. Business Critical Data Applications

   Two of NeedToChange's business critical applications are SAP and Product Data Management (PDM). PDM is NeedToChange's most bandwidth intensive application, however it is widely understood that NeedToChange runs its business on SAP. In addition to the applications that NeedToChange uses to run its business, the company uses an Infrastructure as a Service (IaaS) provider for disaster recovery (DR).

4. Public Cloud Computing Services

   Other than its use of an IaaS site for DR, NeedToChange currently makes relatively modest use of public cloud computing services. However, the decision has been made that on a going forward basis, unless there is a compelling reason not to do it, any new application that the company needs will be acquired from a Software as a Service (SaaS) provider.

5. Voice and Video

   NeedToChange supports a modest but rapidly growing amount of real time IP traffic, including voice, traditional video and telepresence.

[1] Throughout the description of NeedToChange, the MPLS network the company uses is provided by a carrier.

6. Internet Access

   NeedToChange currently backhauls over half of its Internet traffic to its data center in Salt Lake City. The company is looking to enable direct Internet access from their branch offices but they are concerned about security. NeedToChange is also concerned that it is supporting non-business related Internet traffic that is negatively impacting business traffic.

7. Remote Workers

   Roughly half of NeedToChange's employees regularly works either from home or from some remote site.

8. Guest Workers

   NeedToChange's network organization is considering offering guest WiFi access from at least some of its facilities.

9. Branch Offices

   NeedToChange categorizes its branch offices into three categories: small, medium and large.

   - A small office/site has between 5 and 25 employees. These sites are connected by an MPLS network with each site having either a single T1 link or multiple T1 links that are bonded. All of its Internet traffic is backhauled.
   - A medium office/site has between 25 and 100 employees. These sites are connected by an MPLS network with each site having capacity between a single T1 link and a link running at 10 Mbps. All of its Internet traffic is backhauled.
   - A large office/site has more than 100 employees. These sites are connected to an MPLS network either by using bonded T1 links or by a T3 link. They also have direct Internet connectivity which in most cases runs at 10 Mbps over DSL.

10. Visibility

    In the majority of instances in which the performance of one of NeedToChange's business critical applications begins to degrade, the degradation is noticed first by the end users.

11. Regulations

    NeedToChange is subject to PCI compliance. As such, NeedToChange needs a network infrastructure that provides robust security.

12. Factors Driving Change

    While not in priority order, the following factors are driving NeedToChange to seek alternative WAN designs:

    - Improve application performance;
    - Reduce cost;
    - Increase uptime;
    - Reduce complexity;
    - Provide access to public cloud computing services;

- Provide better support for real time applications;
- Reduce the time it takes to implement new network services;
- Increased agility both in terms of supporting new facilities and in supporting growth within existing facilities

Balancing off the factors driving NeedToChange to seek alternative WAN designs is the fact that NeedToChange will not be allowed to increase the size of its network organization.

Silver Peak's Response

NeedToChange Redesigns its WAN with Broadband using Silver Peak Unity

Illustrates Proven SD-WAN Model for Reducing MPLS Dependency and Costs

NeedToChange is a thriving enterprise headquartered in the United States. The company is currently using a traditional wide area network (WAN) with MPLS for remote office connectivity, and a mix of centralized and direct Internet access. The NeedToChange sites include:

- 13 small sites located across the United States
- 8 medium sites located across the United States
- 3 large sites located in the United States
- 1 large site co-located with the primary data center

NeedToChange is facing many of the problems as other enterprises today, mainly to reduce cost and improve performance. The WAN needs to be flexible enough to support new applications (local or in the cloud), secure against attacks, and be able to meet regulatory requirements. The network problems faced by NeedToChange are the same as other companies with multiple locations, a disaster recovery initiative, and a move to cloud based applications.

To meet these requirements and provide flexibility for the future, Silver Peak recommends a software-defined WAN (SD-WAN) fabric using Silver Peak's software and appliance solutions. The initial deployment involves applying a virtual WAN overlay using the current WAN infrastructure, which will improve connectivity, performance, and security. Starting with the current infrastructure will allow NeedToChange to complete the initial deployment quickly while also providing flexibility to transform their network from MPLS to Internet connectivity over time.

A Silver Peak proposed solution includes the following:

- Deploying Silver Peak Unity EdgeConnect virtual or physical appliances (based on the requirements and existing infrastructure) at each physical location
- Using the existing MPLS investment, but over time transitioning to an Internet WAN infrastructure to reduce the monthly cost for bandwidth
- Regional hubs will be utilized as consolidation points for Internet traffic, creating a more secure environment with centralized firewalls
- Security for corporate data in-flight with AES-256 bit accelerated IPsec encryption, done as part of the Silver Peak Unity (SD-WAN) fabric
- Detailed WAN monitoring and usage reporting for each remote site and all Internet traffic
- Dynamic traffic prioritization and control using Silver Peak Quality-of-Service (QoS)
- Full usage of redundant WAN connections at key location

Solution Overview

The Silver Peak solution consists of the Silver Peak Unity EdgeConnect software deployed as virtual machines, or as dedicated hardware appliances, at each NeedToChange location. The Unity EdgeConnect hardware and virtual machines are used to create an SD-WAN. The SD-WAN provides connectivity between all corporate locations, and cloud providers, using any bandwidth. The essential site-to-site connectivity that is provided by MPLS and point-to-point connections today will instead be provided by software. By building the WAN in software, the bandwidth becomes nothing more than a simple transport mechanism that can be changed as needed. This is similar to changing server vendors when running a hypervisor - the applications remain online regardless of the server they are running on.

A regional hub approach will be used to consolidate connectivity to cloud applications and the Internet. The data center in Salt Lake City will serve as the hub for the western United States, Chicago will service the central region, and New York City will service the east. Branch-to-branch and branch-to-data center connectivity will happen directly over the SD-WAN fabric. For the remote sites that do not have virtualized infrastructure, Silver Peak Unity EdgeConnect appliances will be deployed. For sites with a virtualized infrastructure, including the Salt Lake City data center and the two regional hubs, Silver Peak Unity EdgeConnect software appliances will be used. Silver Peak Unity EdgeConnect software can run on nearly any server and supports all common hypervisors.

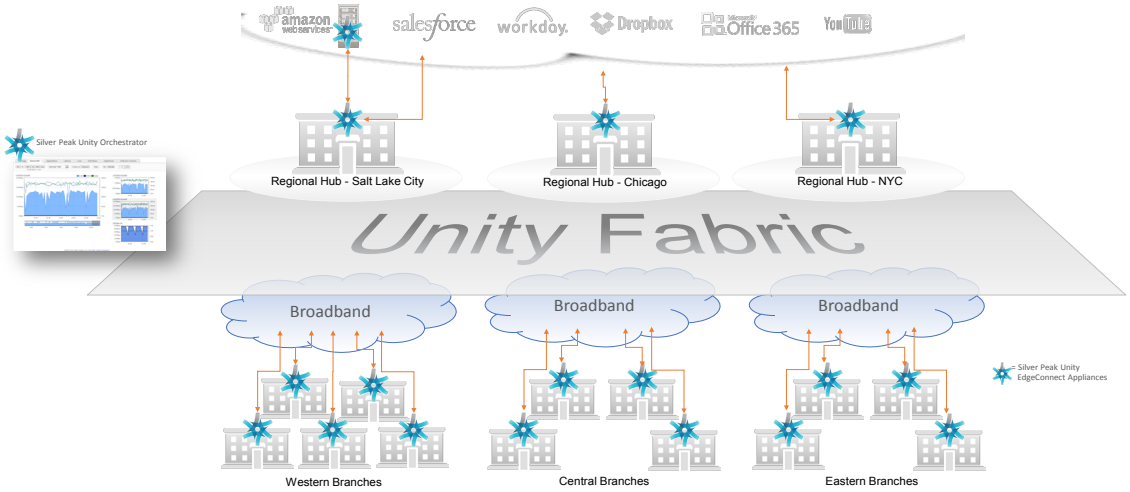

The Silver Peak Unity Fabric will be implemented using a phased approach. Initially, all sites will maintain their current bandwidth (MPLS and Internet) while the WAN fabric is deployed. For sites with multiple connections, Silver Peak's Dynamic Path Control (DPC) feature will be used to load-balance traffic across all connections while also providing redundancy in the event of a connection failure. Any new locations that are opened will only be connected to the Silver Peak fabric over the Internet.

Once Silver Peak is deployed, the IT organization will be able to gain increased visibility into the traffic traversing the WAN. Using Silver Peak's Unity Orchestrator, NeedToChange will apply business intent policies to secure and control all WAN traffic. Powerful features such as a policy builder, traffic shaper, parallel tunnels, and an overlay builder will be used to protect the performance of critical applications, including NeedToChange's SAP, PDM, and VoIP deployments. The added visibility and control will also allow the IT department to find rogue applications and control Internet usage of non-business critical applications, such as Facebook.

The second phase will include an upgrade to the Internet bandwidth at the regional hubs. Each hub will have diversely routed Internet connections of 1 Gbps. The regional hub upgrade will be used to migrate remote offices to an Internet-based WAN, replacing MPLS everywhere. A Silver Peak Unity EdgeConnect software appliance will also be deployed at the company's Infrastructure as a Service (IaaS) cloud provider. The higher bandwidth and optimized path in the Salt Lake City data center will also allow NeedToChange to reduce the recovery point objective (RPO) for data and applications that are protected in the IaaS DR site.

The third phase begins the migration from MPLS to an all-broadband-based WAN. While the Silver Peak solution can use any network for connectivity, MPLS will be replaced with broadband for remote sites to save on bandwidth costs. As sites are migrated to larger or new Internet connections, MPLS connections are decommissioned. When all sites have been converted to broadband access, the MPLS connection at the Salt Lake City data center will be removed. Based on current bandwidth averages, NeedToChange can expect to save ~$91,000 per year for all small sites, ~$140,000 per year for all medium sites, and $84,000 per year for all large sites. The Salt Lake City data center can expect to reduce bandwidth costs from $90,000 to $48,000 annually. The total annual savings across all NeedToChange locations will be $358,346.

WAN Bandwidth Costs for all sites - Current (MPLS)

Site Size	# of Sites	Annual Cost	Monthly Cost
Small	13	$110,136	$9,178
Medium	8	$160,000	$13,333
Large	2	$180,000	$15,000
Data Center	1	$90,000	$7,500
Total		$540,136	$45,011

WAN Bandwidth Costs for all sites - with Silver Peak

Site Size	# of Sites	Annual Cost	Monthly Cost
Small	13	$18,590	$1,549
Medium	8	$19,200	$1,600
Large	2	$96,000	$8,000
Data Center	1	$48,000	$4,000
Total		$181,790	$15,149

Silver Peak Details

Deployment and Management

All Silver Peak Unity EdgeConnect appliances will be configured using zero-touch deployment, with all policies delivered via a cloud management portal. Additional sites can quickly be added with no expertise in the remote office. Anyone can un-box and plug-in the appliance. When the appliance is active on the network, it will connect to the cloud and download its configuration. Additional policies can be pushed to all appliances, or a subset, simply by making a change in the Silver Peak Unity Orchestrator. The Silver Peak Unity Orchestrator provides a single screen from which to easily implement network-wide business intent policies, eliminating the need to make complex and error-prone policy changes at every branch. These simplified deployment capabilities can significantly reduce deployment time and ongoing management costs.

Connectivity

Silver Peak provides the foundation for building a secure and optimized corporate network over the Internet. The DPC feature within the Silver Peak software allows NeedToChange to use multiple WAN connections wherever high availability is a concern. DPC can intelligently control the flow of data over multiple WAN links, including the ability to quickly and seamlessly fail-over an active connection from one path to another. By using WAN links actively, instead of a typical active/passive deployment, NeedToChange benefits from more continuous connectivity. With DPC, network administrators can also "pin" applications and traffic types to a specific WAN. The pin will remain active as long as the WAN connection is active. Should the connection fail, the application traffic will fail over to a remaining active connection.

Security

The Silver Peak Unity deployment will utilize the Salt Lake City data center, a large office in Chicago, and another large office in New York City as regional hubs for connectivity. Each regional hub will have redundant firewalls for security and infiltration protection. NeedToChange initially wanted to deploy broadband access at each remote office, however the problems associated with managing many distributed firewalls are difficult to overcome. Centralized security presence allows NeedToChange to have much tighter control over policies. All Silver Peak Unity EdgeConnect appliances apply WAN hardening, where the only traffic allowed into the device is from other Silver Peak appliances. The Silver Peak Unity EdgeConnect appliances also act as routers, replacing another piece of equipment in the branch. Connections between sites are secured using AES-256 IPsec encryption. The IPsec connection is essential for securing data in-flight over the Internet while also meeting regulatory requirements, like those outlined by the Payment Card Industry (PCI) for retailers.

Visibility and Monitoring

The Silver Peak Unity Orchestrator will give NeedToChange detailed information on everything happening on the WAN. Bandwidth usage, bandwidth savings (if WAN optimization is enabled), top talkers, packet loss, latency, and even a detailed list of Software as a Service (SaaS) traffic is all available from a single management console. The Unity Orchestrator will be instrumental when NeedToChange's network administrators start to build more detailed traffic prioritization policies, and when they decide that a SaaS application or website needs to be blocked. A daily email report can be created that provides the information that is important to NeedToChange's executive team as they work through this project.

WAN Optimization

Silver Peak brings more than 10 years of WAN optimization heritage and expertise to the table for NeedToChange. While WAN optimization capabilities are not required for each of the NeedToChange links, the Silver Peak Unity solution has an optional component called Unity Boost, which allows NeedToChange to add WAN optimization to higher-volume

or more congested segments of the network. Silver Peak WAN optimization features can further reduce the effects of latency, and apply leading deduplication and compression technology to all data sent across the WAN. This could be applied to NeedToChange's backup and replication traffic, where it is essential to meet a stringent (RPO), even as data volumes increase. With Silver Peak, NeedToChange is able to benefit from the flexibility, visibility, and control that an SD-WAN provides today, and easily add WAN performance tomorrow without needing to rip and replace hardware or software. A simple license key will enable WAN optimization wherever it is needed.

Summary

NeedToChange is currently facing a problem that many businesses are also trying to solve, growing their network while controlling costs and maintaining security. The Silver Peak Unity WAN fabric lets NeedToChange grow their network simply, with zero-touch deployment and cloud-based policy management.

- **Silver Peak Unity EdgeConnect** – Silver Peak's virtual or physical products that deploy in data centers, regional hubs and branch offices to build an SD-WAN overlay
- **Silver Peak Unity Orchestrator** – Centralized global management provides a single screen to see all traffic and implement network-wide business intent policies
- **Silver Peak Unity Boost** – WAN optimization that can be added to any Silver Peak Unity EdgeConnect deployment by simply adding a license key to the deployment

With Silver Peak Unity, the time and cost of adding a new site is significantly reduced. Furthermore, the hardware footprint in remote offices is reduced to a single device that controls access, security, routing, and prioritization across multiple WAN links. The regulatory requirements that NeedToChange must follow under the PCI compliance standard are satisfied by Silver Peak's use of IPsec encryption for all traffic sent over the Internet between sites. Silver Peak's detailed reporting and rich policy engine give NeedToChange the ability to prioritize and monitor traffic across the WAN. Performance can be guaranteed for mission critical and real-time applications, while rogue applications and websites are denied from using crucial bandwidth. With Silver Peak, NeedToChange can easily build a broadband-based WAN that provides flexibility, visibility and control, performance, and dramatic cost savings.

Silver Peak Benefit Highlights

- **Flexibility** – NeedToChange can rapidly and non-disruptively augment or replace their existing MPLS networks with any form of Internet connectivity
- **Visibility & Control** – NeedToChange benefits from unprecedented levels of visibility into both legacy and cloud applications, and the unique ability to centrally assign business intent policies to secure and control all WAN traffic
- **Performance** – End-user satisfaction is significantly improved due to consistent and dramatically enhanced application performance
- **Savings** – NeedToChange reduces bandwidth costs and the dependency on MPLS, and has also reduced its hardware footprint and equipment costs in the remote sites.

About Silver Peak

Silver Peak helps enterprises and service providers flexibly and securely connect users to applications via the most cost-effective source of connectivity available. With Silver Peak's WAN solutions, customers can augment or replace MPLS networks with secure Internet connectivity, (often referred to as an SD-WAN) while dramatically reducing WAN costs and complexity. Customers benefit from unprecedented levels of visibility, control and security over all traffic traversing the WAN, while improving application and network performance. With Silver Peak, sites can be rapidly and non-disruptively extended, moved, or changed as business demands evolve. Learn more at www.silver-peak.com.

Key WAN Architecture and Design Considerations

Below is a description of some of the considerations that network organizations need to include in their evaluation of alternative WAN architectures and designs.

1. Location of key WAN functionality

   In a traditional WAN, functionality such as optimization is typically provided onsite. That's still a viable option. However, there are a number of other viable options. Below are some examples of where key functionality may be provided. In many instances network organizations will find that the best solution is for WAN functionality to be located in multiple types of sites.

   Service Provider's Central Office (CO)
   As described in a recent blog, one of the Network Functions Virtualization (NFV) use cases that the European Telecommunications Standards Institute (ETSI) defined is referred to as Virtual Network Functions (VNF) as a Service (VNFaaS). This is more commonly referred to as virtual CPE (vCPE). As part of a vCPE offering a service provider would enable customers to access functionality, such as optimization, that is provided on servers in one or more of the service provider's COs.

   A Service Provider's Central Facility
   Some network organizations have historically outsourced the management of their WAN to a service provider. If that is of interest, network organizations need to ensure that approach to management remains an option as they evaluate alternative WAN solutions.

   A Software-as-a-Service (SaaS) Site
   The initial SaaS offerings focused on business applications such as supply chain management. However, in the current environment most if not all L4 – L7 functionality can be acquired from a SaaS provider. For example, branch office traffic can be tunneled to a SaaS provider's site where the traffic is inspected for malware.

   An Infrastructure-as-a-Service (Iaas) Site or at a Colocation site
   One example of the use of an IaaS/Colocation site is that instead of having firewall functionality at each branch office, traffic from branch offices is tunneled to a nearby IaaS/Colocation site which provides the firewall functionality. This minimizes the overhead that is associated with the management of firewalls and potentially presents some cost savings due to the economy of scale that is associated with providing this functionality in a centralized manner. This approach also enables the company to optimize the performance of the Internet traffic as it flows from a branch office to a central site.

   A Company's Central Facilities
   Instead of using an IaaS or SaaS provider for the type of functionality described in the preceding two paragraphs, a network organization can implement that functionality in one or more of their own facilities, such as a data center or a regional headquarters building.

2. The Use of Dynamic Multi-Pathing

   Being able to load balance traffic over multiple WAN links isn't a new capability. However, in a traditional WAN this capability was difficult to configure and the assignment of traffic to a given WAN link was usually done in a static fashion.

   Functionality currently exists that enables load balancing over WAN links to be done based on a combination of policy and the characteristics of the WAN links. One approach to leveraging this functionality is to dynamically load balance traffic over both MPLS and Internet links with the goal of reducing the capacity, and hence the cost, of the MPLS links and replacing the reduced MPLS bandwidth with relatively inexpensive Internet bandwidth. An alternative approach is to use this functionality to load balance traffic over multiple Internet links.

3. The Use of Policy

   There is a broad movement to implement a policy based approach to all aspects of IT, including networking. Policies can be based on hierarchical system of rules designed to deal with the complexities of the environment, and to manage the relationships among users, services, SLAs, and device level performance metrics. One way that policy can be implemented is at the application level. For example, if the performance of an application begins to degrade because the CPU utilization of a physical server hosting a virtualized network function (VNF) that is used by that application becomes excessive, the VNF may be moved to a server with lower utilization, if that is in line with the policy that exists for that application. As was alluded to in the discussion of dynamic multi-pathing, another way to implement policy-based networking is to control which WAN link application traffic transits based in part on centralized policies that indicate among other things, the business criticality of that application.

4. Network Topologies

   A traditional branch office WAN is often based on a hub and spoke design. That topology is efficient in an environment in which the bulk of the traffic flows from a branch office to a data center. That topology becomes notably less efficient if the bulk of the traffic flows between branch offices. In that type of a network, a highly meshed design, or possibly a fully meshed design is more appropriate.

5. Support for Real-Time Applications

   The 2015 State of the WAN Report contained the results of a survey in which the survey respondents were given a set of a dozen factors and were asked to indicate which factors would like have the most impact on their WAN over the next twelve months. The three factors that were indicated the most were:

   - Support real-time applications such as voice and/or video;
   - Increase security;
   - Improve application performance.

   There are a number of ways that a WAN can provide support for real-time applications. One way was already mentioned – the use of a policy engine that can steer certain traffic to the most appropriate WAN link. In some cases, the optimization techniques that are mentioned below can make it easier to support real-time applications.

6. Optimization

As noted above, improving application performance is a key issue facing network organizations. **Table 1** lists some of WAN characteristics that impact application delivery and identifies WAN optimization techniques that can mitigate the impact of those characteristics.

Table 1: Techniques to Improve Application Performance

WAN Characteristics	WAN Optimization Techniques
Insufficient Bandwidth	Data Reduction: • Data Compression • Differencing (a.k.a., de-duplication) • Intelligent Caching Complementary bandwidth • Utilize low cost alternative circuits (Internet) to offload non-critical business traffic. • Use policy based networking to assign security processes (encryption)
High Latency	Application Acceleration: • MAPI • SMB Protocol Acceleration: • TCP • HTTP • CIFS • NFS Mitigate Round-trip Time • Request Prediction • Response Spoofing
Packet Loss	Congestion Control Forward Error Correction (FEC) Packet Reordering
Network Contention	Quality of Service (QoS)

7. Security

   As noted above, increasing security is a key issue facing network organizations. As they examine new WAN solutions, network organizations need to look at functionality such as firewalls and determine whether that functionality should be in a branch office or in a central site. They also need to evaluate whether or not to implement other security functionality such as encryption and device authentication.

8. Automation

   The use of policy for managing application performance was already discussed. Another use of policy is for device configuration and security policy management. Some WAN solutions make it possible to create device configurations and security policies in a centralized location and push them out to branch offices in a way that requires no manual intervention at the branch offices.

9. Visibility

   There are many tools in marketplace that are positioned as being able to provide network organizations with all of the visibility into their WAN that they need for troubleshooting problems related to network and/or application performance degradation. However, whether it is the deficiencies of those tools or the troubleshooting processes used by network organizations, survey data contained in the 2015 State of the WAN Report showed that less than one out of five network organizations has all of the visibility that they need to effectively troubleshoot problems. In addition, roughly half of network organizations report having visibility into their WAN that either has frequent gaps or that is barely adequate.

   Evaluating new WAN solutions creates an opportunity and a challenge for network organizations. The opportunity is that by implementing a new WAN design, network organizations might be able to increase their visibility into the WAN. The challenge is that network organizations need to ensure that as they explore new WAN alternatives that they evaluate the visibility provided by each of those alternatives.

10. Customer Premise Equipment

    There are alternatives for the customer premise equipment (CPE) that is available both at the branch office and at the data center. One key option is whether the network organization wants to continue to use their existing routers or to replace them with a new device. Another consideration is the ability of the CPE to support the dynamic insertion of L4 – L7 services.

Call to Action

For the first time in a decade, the WAN is the focus of considerable innovation. As a result of this innovation, network organizations have the opportunity to make a significant upgrade to their current WAN architecture and design. Below is the outline of a project plan that network organizations can use to evaluate how to best make that upgrade.

Create an Effective Project Team

As part of evaluating alternative WAN designs, there are a number of components of each design that need to be analyzed. For the sake of example, let's assume there are four primary components of each design which need to be analyzed and those components are the:

- Underlying technologies;
- Ability to manage the technologies;
- Security implications associated with the new technologies and design;
- Financial implications of each design.

One viable option is to have a four person team where each team member is a subject matter expert (SME) on one of the above components[2]. For example, the team could include a SME from the organization's Network Operations Center (NOC). The role of that team member is to ensure that the NOC will be able to manage whatever technologies are eventually implemented.

Establish an Ongoing Dialogue with Senior Management

A key component of this dialogue is to identify management's key business and technology concerns. The reason to do that is because at various times in the project, whether that is getting permission to do a trial or requesting money to buy new equipment, the project team is going to need management's buy-in. It's a lot easier to get that buy-in if the team identifies up front the issues that are most important to management and works to address those issues throughout the project.

Identify the WAN Challenges

For most companies the key WAN challenges include improving application performance, increasing availability, reducing cost and increasing security. However, since every company is somewhat unique, just identifying these challenges isn't enough. The team should also assign a weight to each challenge.

One technique that can be used to assign those weights is to give each project team member 100 points and ask them to assign weights to each challenge. To exemplify how this works assume that there are just two team members, team member A and team member B, and just the four WAN challenges mentioned above. As shown in Table 1, team member A thinks that all challenges are equally important while team member B thinks that improving application performance is much more important than the other challenges. One way to deal with the fact

[2] Other team members could include additional technologists, an application architect, a systems analyst or a business systems analyst.

that there is often a wide variation in how the team members weight the challenges is to come up with an average weighting as shown in the right hand column of **Table 2**.

Table 2: Sample Weighting			
Challenge	**Team Member A**	**Team Member B**	**Average Weight**
Improving app performance	25	55	40
Increase availability	25	25	25
Reduce cost	25	15	20
Increase security	25	5	15

As part of the ongoing dialogue with senior management, the project team should review and possibly revise both the WAN challenges and their weighting.

Agree on the Extent of the Analysis

In conjunction with senior management, the project team needs to determine how broad and how deep of an analysis it will do. For example, consider the four person project team described above and assume that as part of analyzing the choices they have for redesigning their WAN that they identified two alternative approaches:

1. Do a moderately detailed analysis of the solutions provided by their two incumbent vendors and by two other vendors to be chosen by the team.
2. Do a very detailed analysis of the solutions provided by all of the eight vendors that seem viable.

Assume that a very detailed analysis takes twice as much effort as a moderately detailed analysis. That fact combined with the fact that approach #2 involves twice as many vendors as approach #1 means that approach #2 will take roughly four times as much effort as approach #1. To complete this analysis further assume that:

1. The loaded compensation (salary plus benefits) of each of the four project team members is $130,000 or roughly $2,500 per week.
2. Approach #1 will consume 10 weeks of work from each team member.

In the hypothetical situation described above, approach #1 would cost $100,000 and approach #2 would cost $400,000. Approach #2 would definitely provide more insight, but senior management needs to decide if that additional insight worth dedicating an extra $300,000 worth of internal resources.

Choose Vendors

As described above, the decisions that are made relative to the breadth and depth of the analysis of alternative solutions can have a dramatic impact on the amount of time and resources consumed by the process. That is just one of the reasons why the project team needs to choose potential vendors carefully. A reasonable strategy is to enter into a high level conversation with what the team determines to be a feasible set of vendors. If the content of those conversations impresses the team, they can do a deeper analysis with a short list of vendors who they believe can best meet their needs. This approach balances off the desire to do a broad analysis of emerging solutions with the need to conserve IT resources.

Rate Alternative Solutions

Once the team has come up with a set of weights for the key WAN challenges, it should use those weights to rate alternative solutions. For the sake of example, assume there are two viable alternative WAN designs, one from Vendor A and the other from Vendor B.

Table 3: Evaluating Vendors

Challenge	Weighting	Vendor A Scores	Vendor A Total	Vendor B Scores	Vendor B Total
Improving app performance	40	9	360	7	280
Increase availability	25	8	200	8	200
Reduce cost	20	7	140	8	160
Increase security	15	7	105	6	90
Grand Total			**805**		**730**

As shown in Table 2, the team used a 10 point scale to evaluate how the two solutions responded to each of the WAN challenges[3]. The fourth column from the left demonstrates how the total score for vendor A was determined. The team gave Vendor A a 9 for improving app performance. That 9 was multiplied by the weight of that challenge (40) to arrive at a score of 360. That process was repeated for each challenge and the sum of the four scores (805) was determined. That process was also applied to Vendor B, whose total score of 730 is significantly lower than Vendor A's total score. If the scores were closer, it might be valuable to do a "what-if" analysis. For example, what-if reducing cost was weighted higher than 20? What-if Vendor B got an 8 for improving app performance?

When the team presents their vendor evaluation to management there should be little if any discussion of either the set of WAN challenges or the weights that were used in the evaluation as those items should already have been reviewed with management and adjusted based on their feedback. This limits the discussion with management to a small set of well-defined, well-confined questions such as why vendor A got a 9 for improving app performance and vendor B got a 7. In most cases, management, particularly senior management, won't spend much time on questions like that.

Manage existing contracts

One possible decision that a network organization could make after evaluating alternative WAN designs is to decide to significantly reduce their use of MPLS. The implementation of that decision might not be possible in the short term based on the contract that they have with their WAN service provider. That follows because most contracts for WAN services include a Minimum Revenue Commitment (MRC) on the part of the company acquiring the services. If the company significantly reduces their use of MPLS, the company's spend with the service provider could fall below their MRC which would result in some form of penalty or other action, such as extending the life of the contract.

[3] The team needs to agree on the meaning of the 10 point scale. For example, the team may decide that a "6" means "meets most requirements" and that a "10" means "far exceeds all expectations".

The fact that a company isn't able to significantly reduce their use of MPLS in the short terms isn't necessarily a major problem as few companies would want to do a flash cut of a new WAN architecture. An approach that incorporates the need to minimize the risk of implementing a new WAN architecture, with the need to honor existing contracts, and the typical requirement to work within the current manpower limits of the network organization is to phase in the new WAN architecture over time. While this approach makes a lot of sense, it will reduce the savings that results from the WAN upgrade and this needs to be reflected in the business case.

Build a business case

The easiest and most compelling way to build a business case for a WAN upgrade is to base the business case on hard savings. Hard savings refers to a verifiable reduction in spending such as the reduction that results from either canceling an MPLS circuit or cancelling an MPLS service and replacing it with a less expensive Internet circuit. In some cases the network organization will want to pilot the proposed products and/or services to verify the potential savings prior to building the business case.

Soft savings, while important, can be both harder to measure and more difficult to use as justification for upgrading the WAN. There are many types of soft savings associated with a WAN upgrade including:

- Improving the quality of VoIP;
- Protecting the company's revenue stream by increasing availability of key applications;
- Improving employee productivity;
- Responding to compliance requirements;
- Enabling one or more of the company's key business initiatives such as pursuing mergers and acquisitions;
- Improving the performance of one or more applications;
- Supporting mobile workers;
- Enabling one or more of the IT organizations key initiatives such as implementing virtual desktops or making additional use of public cloud services.

Depending on your company, cost avoidance may be considered a hard saving or it may be considered a soft savings. As mentioned, one example of cost reduction is the savings that results from decommissioning an MPLS circuit. An example of cost avoidance is the savings that occurs from not having to increase the capacity, and hence the cost, of an MPLS circuit.

About the Webtorials® Editorial/Analyst Division

The Webtorials® Editorial/Analyst Division, a joint venture of industry veterans Steven Taylor and Jim Metzler, is devoted to performing in-depth analysis and research in focused areas such as Metro Ethernet and MPLS, as well as in areas that cross the traditional functional boundaries of IT, such as Unified Communications and Application Delivery. The Editorial/Analyst Division's focus is on providing actionable insight through custom research with a forward looking viewpoint. Through reports that examine industry dynamics from both a demand and a supply perspective, the firm educates the marketplace both on emerging trends and the role that IT products, services and processes play in responding to those trends.

Jim Metzler has a broad background in the IT industry. This includes being a software engineer, an engineering manager for high-speed data services for a major network service provider, a product manager for network hardware, a network manager at two Fortune 500 companies, and the principal of a consulting organization. In addition, he has created software tools for designing customer networks for a major network service provider and directed and performed market research at a major industry analyst firm. Jim's current interests include cloud networking and application delivery.

For more information and for additional Webtorials® Editorial/Analyst Division products, please contact Jim Metzler or Steven Taylor.

Published by Webtorials Editorial/Analyst Division
www.Webtorials.com

Professional Opinions Disclaimer

All information presented and opinions expressed in this publication represent the current opinions of the author(s) based on professional judgment and best available information at the time of the presentation. Consequently, the information is subject to change, and no liability for advice presented is assumed. Ultimate responsibility for choice of appropriate solutions remains with the reader.

Division Cofounders:
Jim Metzler
Steven Taylor