

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	1
Introduction and Background	2
Definition of WAN	2
WAN Evolution	2
WAN Services	2
Traditional WAN Design	3
Hypothetical Company: NeedToChange	4
Viptela's Response	7
Key WAN Architecture and Design Considerations	12
Call to Action	16

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 Viptela'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 is 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

Viptela 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, Viptela 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 Viptela 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 network¹ 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.

Viptela's Response



Modified Enterprise Requirements

- Number of branches could range from 100 – 10,000
- 10Mbps – 20Mbps bandwidth required for Telepresence and video collaboration
- Need to converge multiple WAN infrastructures to a single overlay infrastructure
- Infrastructure should be policy controlled and centrally managed
- WAN capacity needs to be augmented on-demand and in a cost-effective manner with option of MPLS, Internet or LTE bandwidth at any site
- Operationally, the overlay WAN should either be managed by in-house teams or outsourced to a SP
- SaaS/IaaS/PaaS applications need to have efficient routes to the cloud to achieve requisite application latencies
- Health and visibility information of the entire WAN must be available to the admins in real-time, even if managed by the SP
- Guest Wi-Fi and Business Partner traffic must be isolated from the rest of the enterprise
- No delays in change control or site bring-up. All change requests should be implemented between 1 – 7 days, including integration of new acquisitions

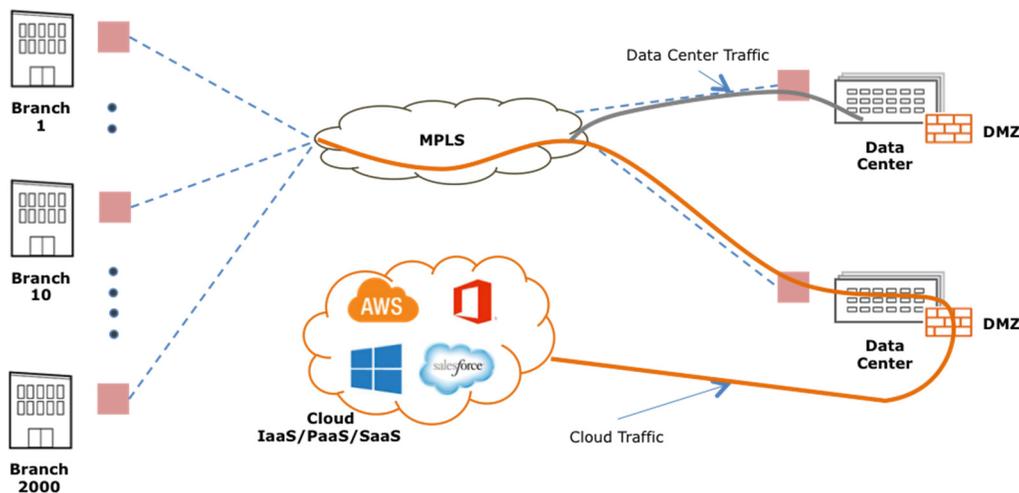


Figure 1: Present MPLS-based WAN

Technology Requirements for Software-Defined WANs

- Must integrate all transport links of MPLS, Broadband and LTE onto a single overlay infrastructure
- Zero-trust network security: device authentication and traffic encryption (full-mesh)
- Should enable flexible, service-based topologies as per application needs (Full-mesh for Telepresence, or hub-and-spoke for ERP implemented on the same overlay infrastructure)
- Centralized provisioning, monitoring and management of the WAN. Dashboard for network health and visibility including detailed application performance stats
- Non-disruptive integration into existing networks with full interoperability with existing routing hardware and routing protocols
- Support centralized App-route policies to honor network-wide SLA for critical applications like Voice and ERP even during failures of MPLS links

- Support end-to-end segmentation to securely isolate Guest WiFi traffic and Business Partner traffic
- Must support efficient traffic paths for cloud applications to prevent hair-pinning of IaaS/PaaS/SaaS traffic through a centralized DMZ
- Scales to tens of thousands of sites globally

Network Transformation Steps

Viptela recommends a phased, non-disruptive WAN transformation approach as detailed below.

Phase 1: Seamless SD-WAN insertion on a sample number of sites (say 10)

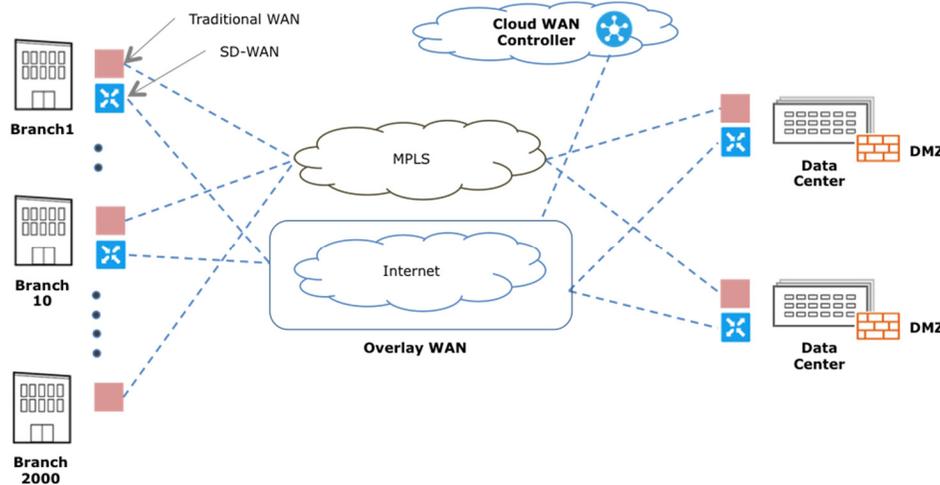


Figure 2: Insert SD-WAN at 10 sites over broadband or LTE

Goal: Insert SD-WAN into 10 sites on the network. This could either be 10 new sites on broadband, or 10 existing sites that need bandwidth augmented.

- Implement the one-time setup for installing the centralized controller and orchestrator. These are virtual machines (VMs) that are either hosted in the data center or cloud.
- Install an SD-WAN router in the data-center and at each of the 10 sites. The transport link is broadband or LTE.
- Peer the SD-WAN routers with the existing branch routers/switches. This enables route learning, so different traffic types can be split between the traditional WAN router and SD-WAN router. SD-WAN policy management is implemented on the centralized controller.
- Traffic routing between SD-WAN sites and MPLS sites happens automatically due to the routing relationships established between these sites. One or more sites are designated as hub sites (connected to both broadband and MPLS) that facilitate the traffic interchange.
- The SD-WAN dashboard provides information on application and link stats of all 10 SD-WAN sites in real time

Note that, in this mode the enterprise inserts SD-WAN without disturbing the existing network by supporting standardized routing protocols like BGP, OSPF and VRRP.

Phase 2: Expanding the overlay on all the WAN transports for the 10 sites

Goal: On the 10 SD-WAN sites, expand the overlays to include all transports (MPLS, LTE, Metro-E, and Broadband). This enables global visibility & policy definition.

- a. The SD-WAN overlay can now be expanded from just broadband to include all other underlay transports

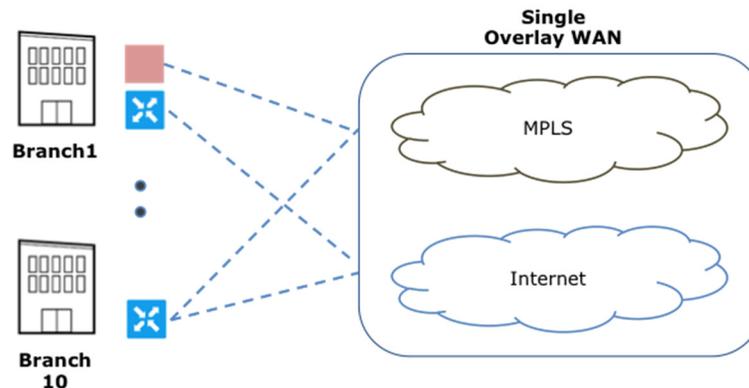


Figure 2: Expand WAN Overlay to All Transport Links

- b. This can be achieved one of two ways.
 - i. Replacing the traditional branch router and connecting all transport links directly to the SD-WAN router (see Branch 10), OR,
 - ii. Keeping the traditional router and extending the overlay through the traditional router. Essentially the overlay tunnels are cutting through the existing router. (see Branch 1)
- c. All the 10 sites are now on a single SD-WAN overlay. The health of all the WAN links and the performance of the overlay WAN can now be monitored on the dashboard

Phase 3: Network-wide, App-route policies to achieve SLA of critical applications

Goal: Setting centralized app-route policies to ensure the SLA requirements of critical applications (voice, ERP) and less critical Telepresence on the 10 sites.

- a. The topology stays the same as the previous step. On the controller, set application policies that meet the following criteria
 - i. Voice traffic is branch-branch, and meets 50ms latency and 50ms jitter
 - ii. ERP traffic is hub-spoke, high priority and must use a no-loss link
 - iii. Video traffic should not traverse MPLS links
 - iv. Employee Internet traffic (facebook, youtube) is prioritized lowest and always uses least expensive, broadband links
- b. These policies ensure that voice and ERP application-SLAs are always met even during failures. The SD-WAN solution monitors all links in real-time and steers traffic based on the centralized policies and link quality
- c. Centralized dashboard provides real-time information and historical information of all application stats and link quality stats.

Phase 4: Expand SD-WAN solution to all sites

Goal: Achieve an enterprise-wide, single overlay WAN

- a. The SD-WAN solution can be expanded to as many sites or all sites if needed. The principles are similar to Phase1, Phase2 and Phase 3 above.

Phase 5: End-to-end segmentation on the SD-WAN network

Goal: Use WAN segmentation to achieve Guest WiFi offload, expeditious integration of M&A acquisitions, or a protected business partner network

- a. Segmentation provides secure logical isolation on the SD-WAN overlay and thus can provide end-to-end segmentation.
 - i. Acquisitions can be integrated on the parent network and yet kept separate. Policies control what applications the acquired company can access.
 - ii. Guest WiFi can be maintained on a separate, low-priority segment and offloaded onto the Internet at closest exit points
 - iii. Business partners can be each defined on a separate segment, or on a collective business-partner network segment. Policies control the access of business partners to data-center applications
- b. Segments are defined as separate VPN instances and controlled centrally by access-control policies

Phase 6: Regional DMZs to optimize latencies of Cloud applications

Goal: Centralized DMZ architectures introduce inefficient paths for cloud applications like SaaS/PaaS/IaaS (as shown in Figure 1). This can be corrected by introducing multiple regional DMZs that are cost-effective options for optimizing latencies for cloud applications.

- a. Instead of one central DMZ, define 3-5 Regional DMZs that are geographically distributed at colocation facilities. Install SD-WAN routers at these locations. This automatically extends the secure footprint of the enterprise to these Regional DMZ locations.
- b. Centralized policies can achieve the following performance and compliance requirements:
 - i. IaaS/PaaS/SaaS application use the closest DMZ exit
 - ii. Sensitive cloud applications like EMR/EHR or Mortgage Transactions or PCI use centralized DMZ

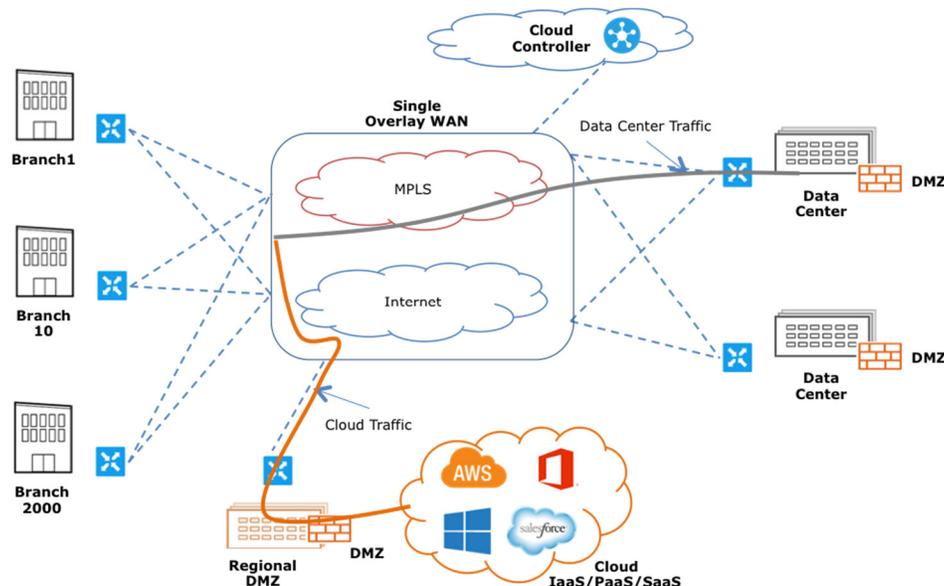


Figure 3: End state of a fully transformed SD-WAN network

Summary

With Software-Defined WANs enable a step-wise transformation to an Overlay WAN that works over any transport, is centrally managed, and with the use of flexible topologies and app-route policies it can meet the SLA goals of all kinds of applications on the network.

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: <ul style="list-style-type: none"> • Data Compression • Differencing (a.k.a., de-duplication) • Intelligent Caching Complementary bandwidth <ul style="list-style-type: none"> • 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: <ul style="list-style-type: none"> • MAPI • SMB Protocol Acceleration: <ul style="list-style-type: none"> • TCP • HTTP • CIFS • NFS Mitigate Round-trip Time <ul style="list-style-type: none"> • 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 components². 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**.

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.

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 challenges³. 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](#)

Copyright © 2015 Webtorials

For editorial and sponsorship information, contact Jim Metzler or Steven Taylor. The Webtorials Editorial/Analyst Division is an analyst and consulting joint venture of Steven Taylor and Jim Metzler.