

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
Cisco'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 Cisco'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

Cisco 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, Cisco 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 Cisco 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.

Cisco's Response





Cisco Systems Recommendations for NeedToChange: Modernizing the WAN for Mobility, Cloud, and IoT

Introduction

NeedToChange network administrators, like many organization administrators, face unprecedented change in their network environment. The traditional WAN was once a well-controlled perimeter of static point-to-point connections to the data center. Most, if not all, applications were hosted inside the enterprise, and measures of success focused on network uptime.

Today, NeedToChange must adapt to a mobile-cloud world, where more and more applications are hosted in multiple places, including the public cloud and infrastructure-as-a-service (IaaS) cloud. Applications are also distributed across private data centers, requiring more data transfer over the WAN. Users expect access from any device from anywhere at any time. And the nature of applications is changing, becoming more immersive and bandwidth-intensive.

Cloud and mobility open a host of security concerns, which is amplified for businesses that are also considering direct Internet access for software as a service (SaaS) and mobile devices. The Internet of Things (IoT) will only compound this problem. And of course, Network IT budget and resources will likely remain flat at best.

To remain competitive and meet growing business demands, NeedToChange must modernize its WAN for the world of mobility and cloud. [Cisco Intelligent WAN](#) follows [structured approach to optimize application performance without compromising security or reliability](#):

1. **Migrate to hybrid WAN:** Build a transport-independent architecture that enables the business to connect multiple access networks (Multiprotocol Label Switching [MPLS], Internet, third- and fourth-generation [3G and 4G LTE, respectively]), and Carrier Ethernet) with a single overlay for operational simplicity.
2. **Protect and optimize application performance:** Move to an application policy-based model that maximizes usage and improves the application experience, through services that provide greater visibility, granular control, and maximum optimization.
3. **Enable a secure, scalable, and resilient infrastructure:** Redesign WAN architecture to elevate security at the branch-office edge for direct Internet access, provide infrastructure that can quickly expand with the business, and ensure 99.99-percent reliability across connections that vary in reliability.
4. **Promote greater automation and orchestration:** Overcome network complexity with a software-based controller model that abstracts the network elements and services and allow IT to direct policy based on business intent with dramatically fewer resources.

Steps to Modernizing the WAN

Step 1: Migrate to a hybrid WAN overlay:

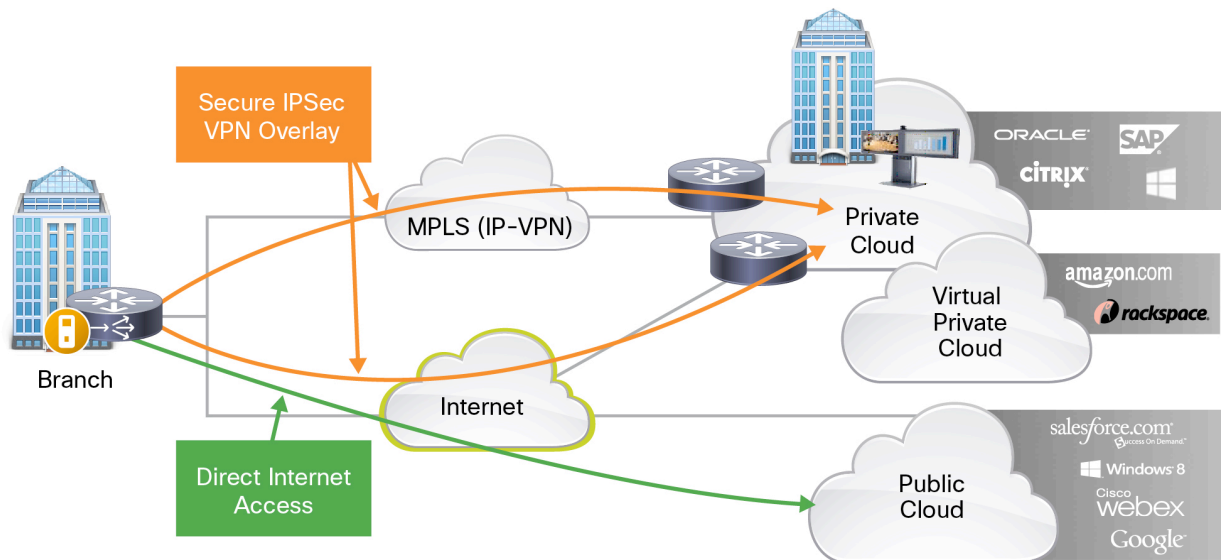
The WAN is a critical business resource that requires resilient design and architecture. NeedToChange will need path diversity and will likely have multiple service providers with different transport networks to support a multi-region WAN. The company must migrate to a hybrid access scheme to meet today's business needs. To increase WAN bandwidth and performance cost-effectively, NeedToChange should augment premium WAN connections with less-expensive transport such as Internet to meet growing traffic demands at lower costs. In addition, for fast branch-office deployment and disaster recovery backup (for example, mobile branch offices, construction, and disaster recovery), the company should also consider cellular 3G/4G LTE backup connectivity.

To accomplish these architectural changes to the WAN, NeedToChange should deploy a transport-independent WAN model that is a single, prescriptive overlay routing design that can be used over any type of WAN transport, with integrated security and the strongest cryptographic protection available to protect corporate data. NeedToChange will realize operational benefits from managing the same IP routing design across all transport networks and, by decoupling application path selection from routing, operations will be greatly simplified making it easier to roll-out new applications.

This architecture will enable NeedToChange to take advantage of hybrid access approaches with MPLS and Internet for private cloud transport as well as allowing future direct access to public cloud services.

- For branch-office access, NeedToChange should use the secure overlay for transport to the private cloud and Internet edge and take advantage of the cost and additional bandwidth afforded with a hybrid network design (MPLS + Internet).
- For future public cloud and Internet access, NeedToChange can build from the base architecture to move to a [direct Internet access](#) method when its organization feels ready.

Figure 1. WAN Design for Private and Public Clouds



- Secure WAN transport for private and virtual private cloud access
- Leverage local Internet path for public cloud and Internet access
- Increase WAN transport capacity and app performance cost effectively
- Improve application performance (right flows to right places)

Step 2: Protect and optimize application performance:

The hybrid WAN overlay design allows NeedToChange to have all connectivity in place with a “set it and forget it” approach, allowing the company to focus on optimizing and protecting application performance.

Intelligent path control: This layer is responsible for routing application traffic optimally, across multiple paths, and ensuring full use of all WAN resources. NeedToChange must move away from separate networks with static traffic mapping to a single dynamic WAN directed by application policy control. Path control assures that application traffic always follows the WAN path that is optimal for user experience. When a WAN path experiences performance impairment, it automatically moves priority traffic to the best-performing path available, protecting application performance and user experience.

To maximize use of expensive WAN resources, path control services automatically load balances traffic across all the WAN connections. There are no “hot spots” or underuse of available WAN circuits that result when static traffic mapping is used for path selection.

Path control and load balancing based on business-directed policies at the application level will greatly simplify the administration of application performance control for NeedToChange. For example, a path control policy may set the MPLS network as a preferred path for voice applications for guaranteed service-level agreements (SLAs) and high reliability provided by MPLS, and load balance other traffic across the network to maximize usage. However, if a brownout occurs, Intelligent Path Control (IPC) will dynamically reroute to the better path (now Internet) so the user experience is maintained, while alerting the network operator so the problem can be immediately addressed.

Application visibility: You can’t control what you can’t see. NeedToChange must have visibility into what applications are on the network and the performance of each application. This visibility is critical for capacity planning and to verify, tune, and troubleshoot problems that affect user experience.


Application-response-time measurement, instrumentation that is integrated as part of the system, should be deployed for mission-critical business applications to isolate where delays are occurring in the network (for example: client, LAN, WAN, or server response time).

Quality of service (QoS): Another important component of the architecture is QoS. After NeedToChange gains visibility into all applications running over the WAN, it can apply QoS policies to groups of key applications to help ensure the priority applications get properly scheduled into the WAN with the proper bandwidth allocation. In the case of Internet transport, with no service guarantees, QoS can be used to ensure proper priority scheduling at the egress interface, with advanced, adaptive QoS enabled to dynamically shape egress traffic to the actual, real-time offered rate as measured end-to-end between WAN routers. Thus QoS can help avoid excess drops during times of congestion, which can result in retransmission of packets, negatively affecting application performance.

Application optimization: Finally, NeedToChange will want to further accelerate application performance through application-optimization principles:

- **Traffic optimization:** TCP optimization, data compression, and data-redundancy elimination allow enterprises to squeeze more out of their existing pipe while maintaining the ability for applications to travel at normal speeds, even during usage spikes.
- **Application-specific optimizers:** These optimizers recognize chatty traffic such as Messaging Application Programming Interface (MAPI) or Microsoft Exchange, or Server Message Block (SMB) for Microsoft file sharing and Citrix ICA for desktop virtualization. They provide latency-mitigation techniques including prefetching data ahead of client requests, asynchronously acknowledging packets to allow the clients and servers to continue sending data, and providing server responses locally to certain client requests.
- **Intelligent caching:** Although bandwidth can relieve traffic congestion, web and cloud applications have introduced new levels of latency that only HTTP object caching can truly address (refer to Figure 2). In many cases, intelligent caching can offload 40 to 90 percent of network traffic, while giving users a near instant application experience.

Figure 2. Latency and Bandwidth Impact on Page Load Time



Step 3: Enable a secure, scalable, and resilient infrastructure:

NeedToChange must [rethink where security should be enforced](#) as its users become more distributed, applications are no longer hosted locally and more devices connect to the network.

Today NeedToChange is backhauling traffic to the data center to their core security devices, which reduces threats but increases bandwidth usage. Secondly, as NeedToChange adopts more SaaS applications and the demand for guest internet access the branch increases, they will likely adopt [direct Internet access \(DIA\)](#) to offload the WAN. Lastly, as more devices connect to the network protection against zero-day threats becomes critical. As a result of these changes NeedToChange will need to evolve its security architecture to address the following needs:

-
- Securing user traffic by moving the security policy enforcement from the data center edge to a centrally managed cloud model to enable businesses to split security services at the remote site between on-premises and the cloud with an HTTP proxy to complete requests and scan for malware, and allow, block, or warn based on the user, group, or business policy
 - Securing the perimeter of the corporate network from Internet threats with local firewalls and intrusion detection and prevention systems at the remote office location
 - Network isolation with routing separation and user-group segmentation for secure access control
 - Data confidentiality and integrity, by providing the strongest encryption possible, including a choice of advanced cryptographic algorithms such as 256-bit Advanced Encryption Standard Elliptical Curve Cryptography (AES-256-GCM or “Suite B”) coupled with Internet Key Exchange Version 2 (IKEv2)
 - Industry compliance; for example, Payment Card Industry (PCI), Network Equipment Building Standards (NEBS), etc.

In addition, NeedToChange requires infrastructure that can grow as their business does, adding new services or more performance through simple software updates. And, NeedToChange must design for resiliency, including instant failover of applications if one network is down, quick disaster recovery (for example, 4G LTE connections to data center), and immediate threat mitigation.

Step 4: Promote greater automation and orchestration:

To promote greater agility, NeedToChange will require [controlled-based architecture](#) with open interfaces, and a software-defined networking (SDN) services plane that can abstract the device layer. This solution must automate and orchestrate WAN deployments in minutes with an intuitive browser-based GUI. A branch-office platform can be provisioned in just minutes without any knowledge of how to configure the devices (i.e., command-line interface or CLI). The application business priorities are translated by the controller into network policies using best practices and validated designs. The controller dramatically reduces the time required for configuration of advanced network services such as VPN, application visibility, path control, and QoS through simple, predefined work flows to deliver these services that align to business policies. The controller-based application offers an easily deployed solution that allows NeedToChange IT to get out of the complexity of managing low-level semantics such as VPN, QoS, and access list policies. Instead, NeedToChange IT can focus on the bigger picture: aligning network resources with the business priorities and delivering outstanding user experiences that result in better business outcomes.

In addition, NeedToChange will need to look at services beyond the WAN that will need to be managed across the branch-office environment, including unified communications, wireless LAN configuration, and more. The company will need full branch-office service automation through virtualized network services. By deploying a branch-office customized standard x86-based appliance and virtualized network services, NeedToChange can deploy new services to the branch office, reducing complete equipment upgrades and eliminating branch-office visits, ultimately resulting in both capital expenditures (CapEx) and operating expenses (OpEx) savings. The solution must include lifecycle management for the virtual machines and service chaining automation between the services. In some cases local applications can also be virtualized on the same platform. The customized x86 appliance must also include physical elements to enhance operation and scalability of the virtual machine and also LAN and WAN interfaces such as 3G and 4G and embedded switch ports, to maintain a single branch-office platform for operational simplicity.

As NeedToChange makes infrastructure investments, the company must have flexibility as it moves from physical to virtual devices, which can be managed by a single management system with full investment protection. The management model must allow for out-of-the-box prescriptive deployments and more sophisticated customized deployments, and it also must work with third-party systems to meet unique business requirements.

Summary

Modernizing the WAN for NeedToChange and other organizations can be a daunting journey. It is essential that benefits from infrastructure investments can be realized today and still scale for tomorrow. The strategy outlined herein allows NeedToChange to lower costs with a hybrid WAN design; improve and protect the application experience; and elevate security from growing threats. As we move to greater automation and orchestration, IT will be able to free resources and accelerate time to market. And, with an open platform, NeedToChange is better prepared for new trends including virtualization of network 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**.

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.

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.