

The 2017 Guide to WAN Architecture & Design

Part 3: Planning for a Successful Transition to a New WAN

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

Platinum Sponsors:



Table of Contents

Executive Summary	1
Call to Action	2
Introduction	2
Identify the Focus of the Project and the WAN Challenges	3
Agree on the Extent of the Analysis	3
Create an Effective Project Team	4
Choose Vendors	4
Rate Alternative Solutions	5
Manage existing contracts	5
Build a business case	6
Key WAN Architecture and Design Considerations	7
The Role of Cellular	7
Location of Key WAN Functionality	7
The Use of Dynamic Multi-Pathing	8
The Use of Policy	8
Network Topologies	9
Support for Real-Time Applications	9
Optimization	10
Security	11
Automation	11
Visibility	11
Customer Premise Equipment	11

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 for two reasons. One reason is that for the first time in well over a decade the wired WAN is the focus of considerable innovation which is leading to the deployment of a wide range of new WAN-related products and services. The second reason is that on a going forward basis, the WAN needs to support a new set of requirements such as providing connectivity to a growing number of mobile workers and public cloud providers as well as to the Internet of Things (IoT).

The primary goals of the [2017 Guide to WAN Architecture and Design](#) (The Guide) are to make enterprise network organizations aware of the emerging alternatives to the traditional approaches to WAN architecture, management and security and to help them understand the key differences in those alternatives.

The Guide will be published both in its entirety and in a serial fashion. This document is the third of the serial publications and it has two primary sub-sections. One sub-section is a detailed call to action and the other sub-section summarizes the key WAN architecture, management and security considerations that were brought out in Part 2 of The Guide.

The other sections of The Guide are:

- [Part 1](#)
This section focused on providing insight into the current state of the WAN and it contained the results of a survey that was distributed in May of 2016. Throughout The Guide the network professionals who completed the survey will be referred to as The Survey Respondents.
- [Part 2](#)
This section contained the description of a hypothetical company called NeedsToChange and it also contained how each of the sponsors suggested that NeedsToChange should evolve its WAN.
- [Complete copy](#)
The final publication will consist of an executive summary and Parts 1 – 3 as described above.

Call to Action

Introduction

In the novel *Alice and Wonderland*¹, Lewis Carroll used the following dialogue between Alice and the Cheshire Cat to explain the need for planning.

Alice: “Would you tell me, please, which way I ought to go from here?”

The Cheshire Cat: “That depends a good deal on where you want to get to.”

Alice: “I don't much care where.”

The Cheshire Cat: “Then it doesn't much matter which way you go.”

Alice: “...So long as I get somewhere.”

The Cheshire Cat: “Oh, you're sure to do that, if only you walk long enough.”

The relevance of the preceding dialogue to the process of a company's migration from their current to their next WAN is that without a plan that includes a clear sense of what the company is trying to accomplish, then the only way that the company is guaranteed of success is if it implements all possible WAN solutions.

The creation of a business case to justify adopting a new WAN solution is the last topic discussed in this sub-section of The Guide. However, network organizations should create an outline of the business case at the very beginning of the project and use that outline to drive the creation of the project plan. The reason for doing this is to ensure that the project is set up in such a way that it gathers all of the information necessary to create a compelling business case.

At the same time that the network organization creates the outline of the business case they should also begin a dialogue with anyone who is a key stakeholder in the process. In this context, the *key stakeholders* are whoever signs to authorize paying for the new solution as well as anyone who has a significant influence over the decision process, particularly those people who can either cause the project to be delayed or cancelled. A key component of this dialogue is to identify the stakeholder's primary business and technology concerns as well as to get their input on the overall direction of the project. The reason to start the dialogue early in the process is because at various times during the project, whether that is getting permission to do a trial or requesting financial authorization to acquire a solution, the project team is going to need management's buy-in. It's a lot easier and faster to get that buy-in if the team identifies up front the issues that are most important to the key stakeholders and works to address those issues throughout the project.

The following sub-sections outline some of the key components of a project plan for evaluating WAN solutions. The intention is that network organizations will modify this outline to suit their environment.

¹ <http://www.goodreads.com/quotes/225938-would-you-tell-me-please-which-way-i-ought-to>

Identify the Focus of the Project and the WAN Challenges

The term *WAN* refers to a wide range of types of connectivity. The primary uses of the term *WAN* refer to connecting a:

- Data center to either another data center or a public cloud facility;
- Branch office to either a data center, a public cloud facility or a web site;
- Home office to either a data center, a public cloud facility or a web site;
- Remote user to either a data center, a public cloud facility or a web site;
- Thing, such as a car or a school bus, to either a data center, a public cloud facility or a web site.

As part of creating the project plan, the network organization needs to decide on the focus of the project because the type of solutions that are appropriate for some classes of WAN challenges, such as providing connectivity between and amongst a company's data centers, may not be appropriate for a different class of WAN challenges, such as providing connectivity to remote users or to things. The network organization should also decide the type of solution or solutions that it wants to evaluate; e.g., Do-It-Yourself (DIY), managed service or Network-as-a-Service (NaaS). Those decisions should be reviewed with the key stakeholders.

Once the focus has been determined, the project team should identify the WAN challenges that they are currently facing or expect to face and use these challenges to structure their analysis of alternative WAN solutions. 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. The challenges and the weights that are assigned to them should be reviewed with the key stakeholders.

Agree on the Extent of the Analysis

In conjunction with the key stakeholders, the project team needs to determine how broad and how deep of an analysis it will do. A broad and deep analysis can yield more insight than would be produced by a more cursory analysis. However, the broader and deeper the analysis the more it costs and the longer it takes.

Network organizations who want to do a broad and deep analysis often create a Request for Information (RFI) to be sent to numerous possible providers. However, a large and increasing number of organizations are avoiding issuing formal RFIs and instead are engaging in somewhat brief conversations with a small number of WAN providers. They hold these conversations prior to moving forward with a production test by either piloting a WAN solution or conducting a POC of one.

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.

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.

One of the primary challenges of this approach is being able to understand vendors' strategies well enough to choose a feasible set of vendors while having minimum, if any, direct vendor interaction. One way to respond to this challenge is to subscribe to expensive third party services that analyze vendor offerings. As an alternative or as a supplement to relying on information from expensive third party services, this e-book provides detailed insight into the WAN vision and strategy of several key vendors.

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

Rate Alternative Solutions

Assume that the project team has come up with the challenges and weights shown in the first two columns of **Table 1**. Also 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 application 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 1**, 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 application 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 application 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 application 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

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".

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.

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 potential 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 cancelling an MPLS service and replacing it with a less expensive Internet circuit. In almost all 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 the 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 replacing MPLS bandwidth with Internet bandwidth. 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.

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.

The Role of Cellular

Cellular services have long been used as a back-up to wireline WAN services. One of the reasons for this is that the types of issues, such as a backhoe cutting the wired access lines, that would cause a wireline access service to fail would have no impact on a cellular service.

Increasingly cellular services are being used as either the primary WAN link or are used in conjunction with a wireline service in an active-active configuration. In the latter case, traffic is typically load-balanced over the cellular and wirelines services using the type of policy capability that is described below.

Some of the other key use cases for cellular services in an enterprise WAN include:

- **Temporary networks**
The time that it takes to get a wireline service such as MPLS installed is typically a month or longer. In the vast majority of cases that means that wireline services are not a feasible solution for the types of temporary networks that are needed to support locations such as construction trailers or pop-up stores.
- **In-vehicle networks**
While it may or may not be desirable to use an MPLS or DSL-based Internet service to provide connectivity to a fixed site such as a branch office, it isn't possible to use these services to provide connectivity to vehicles such as cars, trucks and school buses.
- **Internet of Things (IoT)**
IoT is a phrase that refers to the internetworking of a wide range of physical devices, buildings and other things that are embedded with electronics and/or sensors. For example, a *thing* may be a sensor inside of a traffic light. In situations like this, similar to in-vehicle networks, cellular services are the only feasible option.

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 [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. Alternatively, functionality such as optimization could be provided in a CO and other functionality, such as security, could be provided onsite at the customer's facility.

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.

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.

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. One goal of this approach is to reduce the capacity, and hence the cost, of the MPLS links and to replace 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.

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 a 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 consider the business criticality and the delay sensitivity of that application.

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, or possibly a fully meshed design is more appropriate.

Support for Real-Time Applications

The 2016 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 likely have the most impact on their WAN over the next twelve months. One of the top factors mentioned by the respondents was supporting real-time applications such as voice and/or video.

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.

Optimization

Improving application performance is a key issue facing network organizations. **Table 2** lists some of WAN characteristics that impact application delivery and identifies WAN optimization techniques that can mitigate the impact of those characteristics.

Table 2: 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)

Security

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, including:

- Encryption;
- Device authentication;
- URL filtering;
- Network access control;
- IDS/IPS;
- Micro-segmentation;
- Anti-malware.

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.

Visibility

There are many tools in the 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 2016 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.

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.

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

Division Cofounders:

[Jim Metzler](#)
[Steven Taylor](#)

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.

Copyright © 2017 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.