The 2017
Guide to WAN Architecture & Design

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

Platinum Sponsors:

Table of Contents

Executive Summary

Introduction

2017 Guide to WAN Architecture and Design (The Guide) was published both in its entirety and in a serial fashion. The three serial publications were:

- Part 1: State of the WAN
  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.
- Part 2: WAN Evolution
  This section contained the description of a hypothetical company called NeedsToChange and it also contained how the sponsors of The Guide suggested that NeedsToChange should evolve its WAN.
- Part 3: Planning for a Successful Transition to a New WAN
  This section of The Guide contained a detailed call to action as well as a summary of the key WAN architecture, management and security considerations that were brought out in Part 2.

Below is a summary of The Guide.

State of the WAN

WAN Evolution

The modern WAN got its start in 1969 with the deployment of the ARPANET which was the precursor to today's Internet. 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 wired WAN technologies and services. This deployment started with Integrated TDM-based WANs in the early 1990s and ended in the early 2000s with MPLS.

The early to mid-1980s also saw the beginning of the deployment of four generations of cellular services. The next generation of cellular services, denoted 5G, should be in production in the 2018 to 2020 timeframe.

WAN Use Cases

The vast majority of WAN use cases can be put into three broad categories:

- Connecting a distributed set of people and devices to centralized resources;
- Connecting multiple data centers;
- Providing peer-to-peer connectivity.

In many instances the WAN solution that is appropriate for one class of WAN use case is not appropriate for others. For example, a solution that is appropriate to connect multiple data centers is unlikely to be an appropriate solution for connecting mobile users to centralized resources.

Factors Impacting the WAN

The Survey Respondents indicated that the following factors were likely to have the most impact on their WAN over the next twelve months:

- Increase security;
- Reduce cost;
- Support real-time applications such as voice and/or video;
- Provide access to public cloud computing services;
- Prioritize business critical traffic.

Concerns with WAN Services

The following table identifies the concerns, listed in descending order of importance, that network organizations have with their use of MPLS and the Internet.

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

Some of the limitations that are associated with cellular services include variable signal coverage, link setup latency and constantly evolving specifications; i.e., 4G, LTE, XLTE, 5G.

Satisfaction with the Current WAN Architecture

Only a third of organizations are either very satisfied or completely satisfied with their current WAN architecture. This indicates that a large portion of the WAN marketplace would likely be receptive to alternative WAN architectures.

Location of WAN Functionality

In contrast to traditional WAN architectures, in the emerging WAN architectures there are a number of places to host functionality such as orchestration, control and security. **Figure 1** highlights the places where network organizations think such functionality should be located based on a survey question that allowed for multiple answers.

Figure 1: Location of WAN Functionality

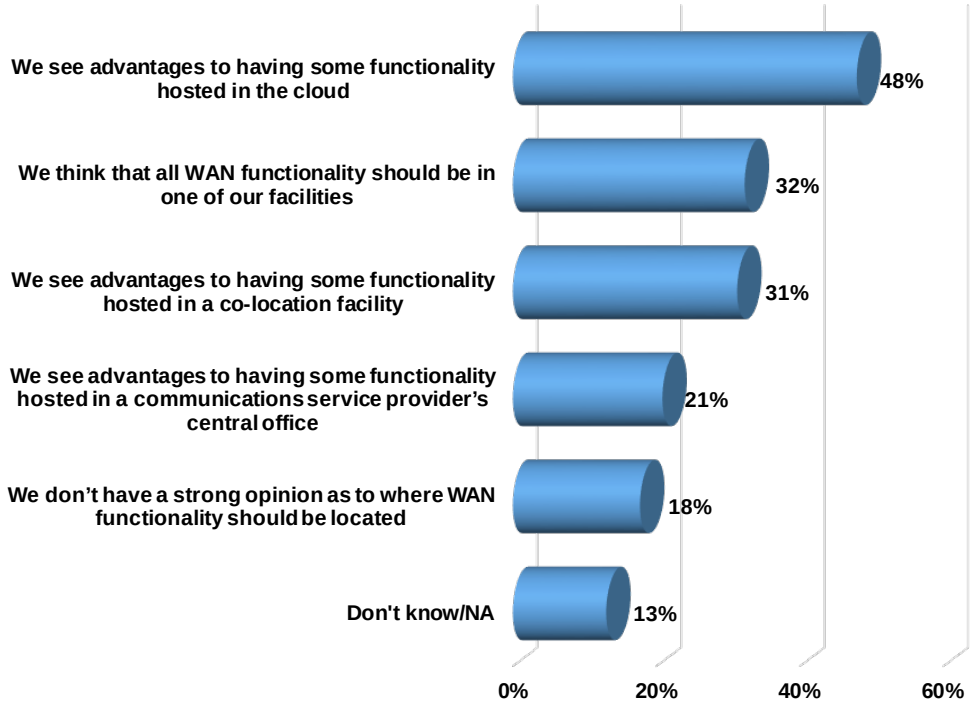

Choice of Implementation Options

When network organizations evaluate new WAN solutions they have a variety of implementation options to consider. The bullet list below indicates the options that network organizations prefer based on a survey question that allowed for multiple answers.

- Do-it-Yourself (DIY): 54%
- Managed Service: 42%
- Network-as-a-Service (NaaS): 27%

Choice of Vendors

Whenever there is a transition point in IT there is the potential that some vendors will gain market share and that some will lose market share. Based on a survey question that allowed for multiple answers, 27% of the respondents indicated that it was highly likely that their organization would stick with their incumbent vendor for a new WAN solution. 22% said that their organization would actively seek alternative vendors.

WAN Management

The visibility that network organization have into their WAN for troubleshooting problems related to network and/or application performance degradation is shown in **Figure 2.**

Figure 2: WAN Troubleshooting Visibility

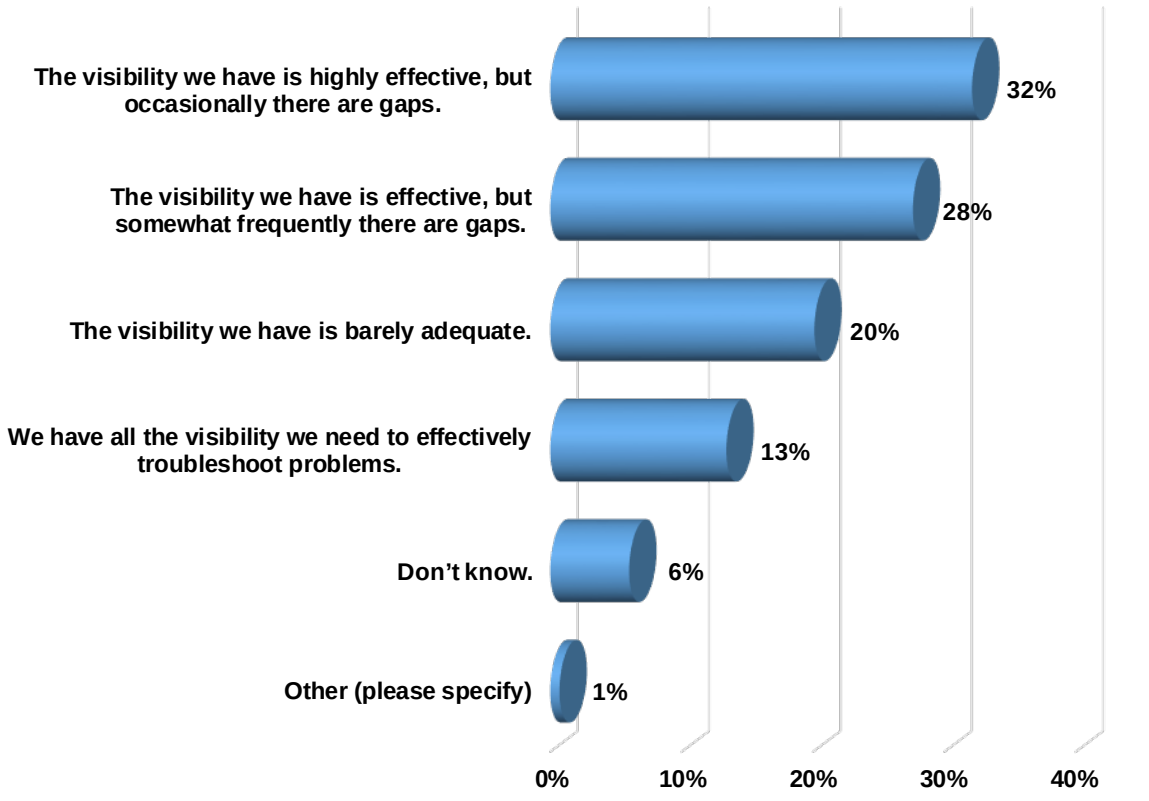

The deployment of new WAN solutions is an opportunity for network organizations to improve on their ability to troubleshoot the WAN and hence improve their ability to support the company's critical business processes.

Planning for a Successful Transition to a New WAN

Call to Action

Below is a brief outline of some of the key components of a project plan for evaluating new WAN solutions.

- **Identify the Focus and Extent of the Project as well as the WAN Challenges**

  As previously discussed, there are varying types of WAN use cases and as part of creating a project plan, the network organization needs to decide on which WAN use cases the project will address.

  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.

  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. The project 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.

- **Create an Effective Project Team and Choose Vendors**

  As part of evaluating alternative WAN designs, there are a number of components of each design that need to be analyzed. One viable option is to create a project team where each team member is a subject matter expert (SME) on one of the components.

  One way to choose vendors 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.

- **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 the organization has with their WAN service provider. This 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.

- **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, such as the reduction that results from cancelling an MPLS service and replacing it with a less expensive Internet circuit. Upon completion of a POC, network organizations should be able to accurately calculate these potential savings

  Soft savings, such as improving flexibility, while important, can be both harder to measure and more difficult to use as justification for upgrading the WAN.

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. 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.

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

  - Temporary networks;
  - In-vehicle networks;
  - The Internet of Things (IoT).

- **Location of Key WAN Functionality, the Use of Policy and Support for Real-Time Applications**

  In a traditional WAN, functionality such as optimization is typically provided onsite. However, as previously described currently there are a number of other options for where to house key functionality. In many instances network organizations will find that the best solution is for WAN functionality to be located in multiple types of sites.

  Functionality currently exists that enables dynamic load balancing over WAN links to be done based on a combination of policy and the characteristics of the WAN links. Since there are differences in terms of how this functionality is implemented, network organizations need to understand what those differences are and what the impact of those differences is.

  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**

  In many instances, optimization functionality can significantly improve application performance. Relevant optimization functionality includes:

  - Data Reduction:
    - Data Compression
    - Differencing (a.k.a., de-duplication)
    - Intelligent Caching
  - Mitigate packet loss:
    - Congestion Control
    - Forward Error Correction (FEC)
    - Packet Reordering

- **Security**

  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.

- **Customer Premise Equipment**

  The emerging set of WAN solutions offer alternatives for the customer premise equipment (CPE) that is deployed both at the branch office and at the data center. One alternative is whether the network organization wants to continue to use their existing routers or to replace them with a new device, either initially or over time. Another consideration is the ability of the CPE to support the dynamic insertion of multiple L4 – L7 services.

The State of the WAN

The role of the WAN and of a WAN Architecture

The primary objective of a WAN is to enable business operations in a frictionless, cost-effective manner. This includes supporting the existing business models as well as changes to those models, such as those brought about by the transformation to become a digital business. To accomplish that objective, the WAN must support the existing applications as well as new applications and the adoption of new application architectures, such as those based on cloud native applications.

Applications make varying demands on a WAN based on the application's:

- Location: On premise, cloud based or a combination
- Business criticality;
- Sensitivity to transmission impairments;
- Security risk;
- Time criticality;
- Compliance requirements;
- Bandwidth requirements;
- Type of user: fixed or mobile or a combination.

The role of a WAN architecture is to enable an organization to deploy a WAN that can adapt quickly to changing business and technical requirements and to respond appropriately to application demands. In order to be effective, a WAN architecture must:

- Ensure acceptable levels of application performance and availability;
- Provide monitoring and management functionality that enables the organization to plan for the deployment of new functionality and to perform rapid root cause analysis and remediation;
- Provide appropriate security;
- Be cost effective.

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. The Internet itself got commercialized in the 1990s with the advent of the World Wide Web.

In addition to the continued evolution of the Internet, the twenty-year period that began around 1984 saw the deployment of the following four distinct generations of wired WAN technologies and services:

- Mid to late-1980s: Integrated TDM-based WANs;
- Early 1990s: Frame Relay;
- Mid to late 1990s: ATM (Asynchronous Transfer Mode);
- Early 2000s: MPLS.

The early to mid-1980s also saw the beginning of the deployment of four generations of cellular services. **Figure 3** depicts the evolution of cellular services from the 1G services of the 1980s to the current generation of 4G LTE services. The next generation of cellular services, denoted 5G, should be in production in the 2018 to 2020 timeframe.

Figure 3: Evolution of Cellular Services

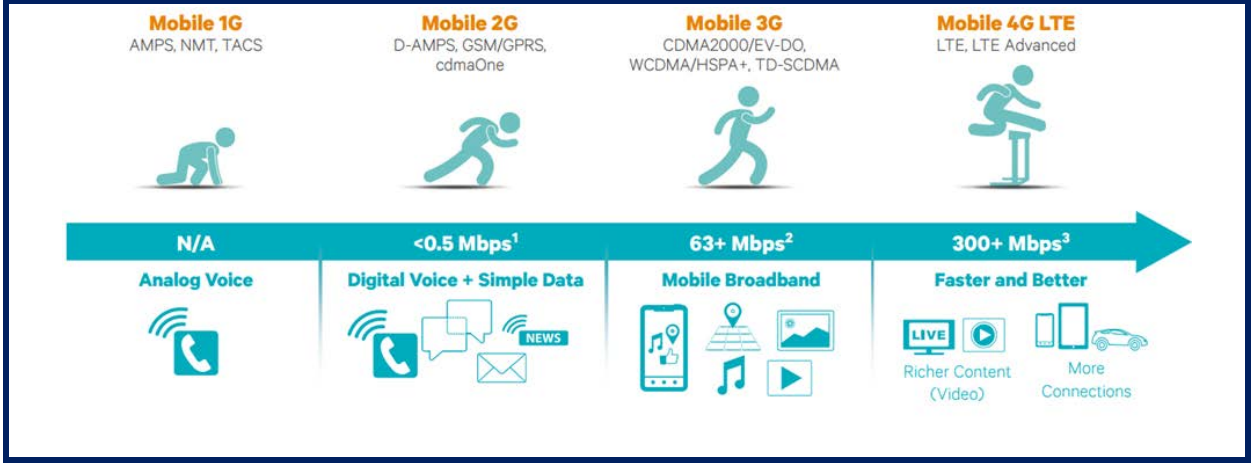

WAN services that were based on Ethernet technology, such as Carrier Ethernet, began to be deployed in the early 2000s primarily to support high speed connectivity in a metropolitan area. These services are also used in some instances for high speed Internet access and to interconnect data centers.

Why is this important?	Unlike virtually every other component of IT, there have been very few if any advances in wired WAN technologies and services for over a decade. Because the types of challenges that the WAN must respond to have evolved significantly during that time frame, there is a pent up demand for new WAN solutions.

WAN Use Cases

The vast majority of WAN use cases can be put into three broad categories:

- Connecting a distributed set of people and devices to centralized resources;
- Connecting multiple data centers;
- Providing peer-to-peer connectivity.

Connecting a distributed set of people and devices to centralized resources

Over the last twelve to eighteen months the vast majority of what has been written about the WAN has focused on providing connectivity between the users in a branch office and the resources they need to access, whether those resources are in a corporate data center or at a public cloud provider's facility. Some of the challenges of this use case are to minimize cost and to provide secure Internet access.

There are, however, other important use cases in this category. That includes supporting:

- Home users;
- Mobile employees;
- The IoT.

The challenges that are associated with the three use cases listed above are somewhat different than the challenges that are associated with providing branch office connectivity. This follows in part because in each of the use cases listed above it is more difficult, if not impossible, to implement distributed functionality to improve performance, management or security. In addition, similar to supporting mobile workers, in many instances supporting the IoT requires the use of cellular services which have notably different characteristics than do wireline WAN services.

Connecting multiple data centers

In the not too distant past, the primary use cases in this category were disaster recovery and business continuity. While those are still important use cases, another important use case, supporting the movement of workloads between data centers, has recently emerged.

This category of WAN use cases has a number of key characteristics that differ from the preceding category including the requirement for significantly more throughput and in many cases, for higher availability. This category of WAN use cases also introduces protocols that are not found in other categories and this category is often associated with WAN services, such as Carrier Ethernet, which have little relevance to the other categories.

Providing peer-to-peer connectivity

In contrast to the other categories of WAN uses cases, in a peer-to-peer WAN, tasks are partitioned between peers. Peers typically make a portion of their resources, such as processing power, disk storage or network bandwidth, directly available to other network participants, without the need for central coordination.

One key use case of a peer-to-peer WAN, file sharing, is often associated with illegal activities. However, there are legitimate instances of this use case such as Lion Share which enables academic institutions to share scholarly documents. A number of emerging applications also

use peer-to-peer WANs. This includes Spotify which uses a peer-to-peer network along with streaming servers to stream audio and video to their clients. It also includes Bitcoin and other alternative currencies such as Peercoin and Nxt.

Why is this important?	For the foreseeable future there will not be a WAN solution that is optimal for all organizations. The optimal WAN solution will depend on a number of factors, including the use case(s) it has to support. 

Factors Impacting the WAN

The Survey Respondents were presented with fifteen factors and asked to choose the three factors that would likely have the most impact on their WAN over the next twelve months. The factors that were the most important are shown in **Figure 4**.

Figure 4: Top Five Factors impacting WAN

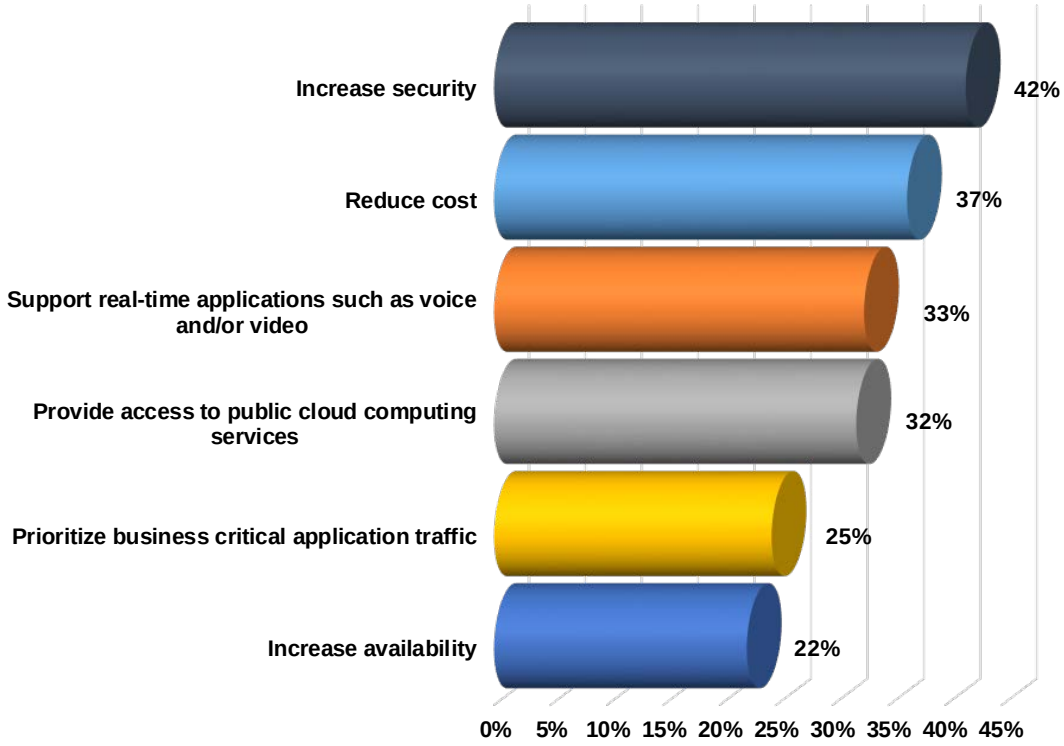

If there is a mild surprise in **Figure 4** it is that a third of The Survey Respondents indicated that providing access to public cloud services is one of the top factors impacting their WAN. This is a bit of a surprise only because unlike the other factors in **Figure 4**, until recently providing access to public cloud services was seldom mentioned as a factor driving change in the WAN.

It was not surprising that eighteen percent of The Survey Respondents indicated that supporting mobile users is one of the top factors impacting their WAN. However, an important and somewhat surprising result that is not shown in **Figure 4** is that sixteen percent of The Survey Respondents indicated that supporting the IoT was one of the top factors impacting their WAN. This is surprising only in that the vast majority of companies are just beginning to feel the impact of the IoT and this impact will likely increase significantly over the next few years.

Why is this important?

In order to justify the cost and the risk of implementing a new WAN solution, that solution must enable organizations to respond to at least some of the challenges shown in **Figure 4**.

Concerns with WAN Services

As discussed in *The 2015 Guide to WAN Architecture and Design*, network organizations currently make relatively little use of wired 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 2: 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

Wireline services are not the only WAN services that have limitations. Some of the limitations that are associated with cellular services include:

- Variable signal coverage;
- Link setup latency;
- Constantly evolving specs (3G, 4G, LTE, XLTE, 5G);
- Security;
- Supporting multiple carriers simultaneously.

Why is this important?

In order to provide value, any WAN solution that is comprised of multiple WAN services, whether they are wired or wireless services, must maximize the advantages of each service while simultaneously minimizing their disadvantages.

Satisfaction with the Current WAN Architecture

The Survey Respondents were asked to indicate how satisfied their organization was with their current WAN architecture. Their responses are shown in **Figure 5**.

Figure 5: Satisfaction with current WAN architecture

Why is this important?

As shown in **Figure 5**, only a third of organizations are either very satisfied or completely satisfied with their current WAN architecture. This indicates that a large portion of the WAN marketplace would likely be receptive to alternative WAN architectures.

Location of WAN Functionality

In contrast to traditional WAN architectures, in the emerging WAN architectures there are a number of places to host functionality such as orchestration, control and security. Those locations include:

- At the customer's branch offices;
- In a service provider's central office;
- At the customer's regional office or data centers;
- In a cloud site provided by a vendor;
- At a co-location facility;
- At a public cloud provider's facility.

The Survey Respondents were asked to indicate where their organization thinks that WAN functionality such as control, optimization and security should be located, and they were allowed to indicate multiple places. Their responses are shown in .

Figure 6: Location of WAN Functionality

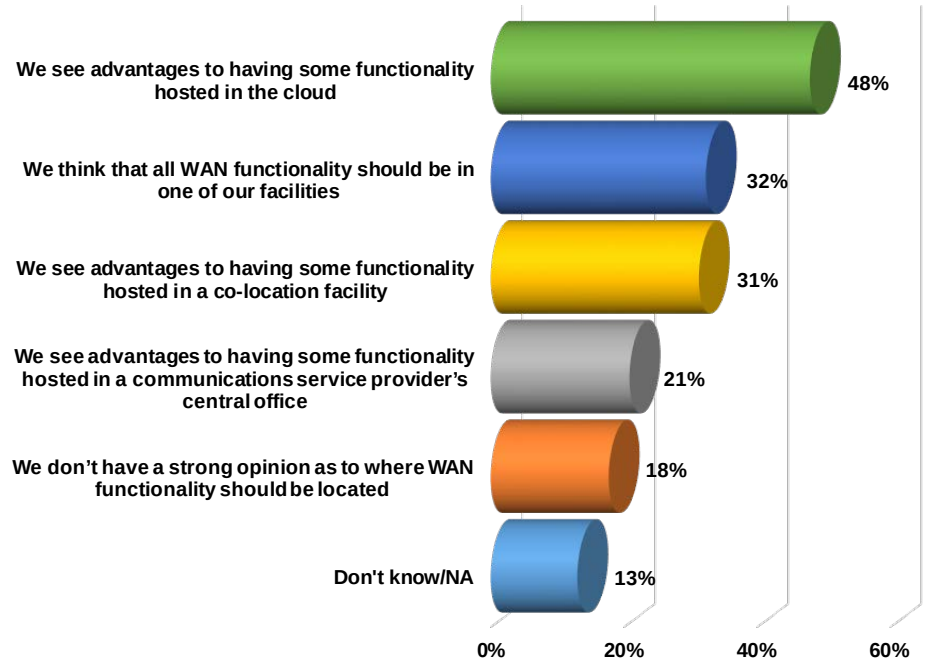

Why is this important?

indicates that a sizeable percentage of The Survey Respondents either didn't know where their organization believes that key WAN functionality should be hosted or they worked for an organization that didn't yet have a strong opinion. However, looking just at those organizations that have an opinion shows that many network organizations are receptive to a range of options relative to where WAN functionality is hosted. It also shows a strong interest in having some WAN functionality hosted in the cloud.

Choice of Implementation Options

When network organizations evaluate new WAN solutions they have a variety of implementation options to consider. This includes:

- Do-it-Yourself
  In the Do-it-Yourself (DIY) option, network organizations are responsible for all facets of the lifecycle of a WAN solution, including the planning, designing, implementing and ongoing management of the solution.
- Managed Service
  In this option a vendor such as a Communications Service Provider (CSP), systems integrator or value added reseller takes on the responsibility for all facets of the lifecycle of a WAN solution.

- Numerous CSPs have either already launched or have announced their intention to launch a Network-as-a-Service (NaaS) offering based on Software Defined Networking (SDN) and/or Network Functions Virtualization (NFV).

The Survey Respondents were asked to indicate which implementation option their organization was most likely to implement and they were allowed to indicate multiple choices.

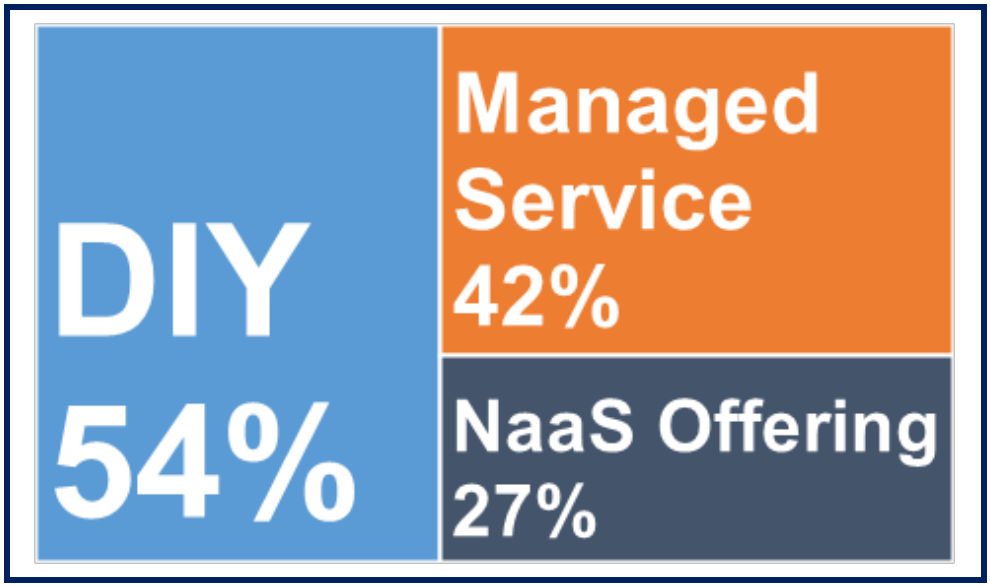

Why is this important?

One way to look at the survey results is that the DIY option is the preferred option by a relatively wide margin. However, another way to look at the survey results is to observe that the combination of a managed service and a NaaS solution are preferred over the DIY option by a relatively wide margin. In either case, the responses to this question provide further evidence that there isn't a WAN solution that is optimal for all organizations.

Choice of Vendors

Whenever there is a transition point in IT there is the potential that some vendors will gain market share and that some will lose market share. After more than a decade with little change in the available WAN products and services, the emergence or a broad range of new WAN related products and services marks the beginning of a major transition in the WAN market. The Survey Respondents were asked to indicate how their organization would likely approach the selection of a WAN vendor and they were allowed to indicate multiple choices. Their responses are shown in **Figure 7**.

Figure 7: Interest in Looking for New Vendors

Why is this important?	The fact that only 6% of The Survey Respondents indicated that they would definitely stick with their incumbent vendor(s) and that an additional 13% indicated that it was unlikely that they would stick with their incumbent vendor indicates that many network organizations are receptive to changing WAN vendors.

WAN Management

The Survey Respondents were asked to rate the visibility that their network organization has into their WAN for troubleshooting problems related to network and/or application performance degradation. Their responses are shown in **Figure 8.**

Figure 8: WAN Troubleshooting Visibility

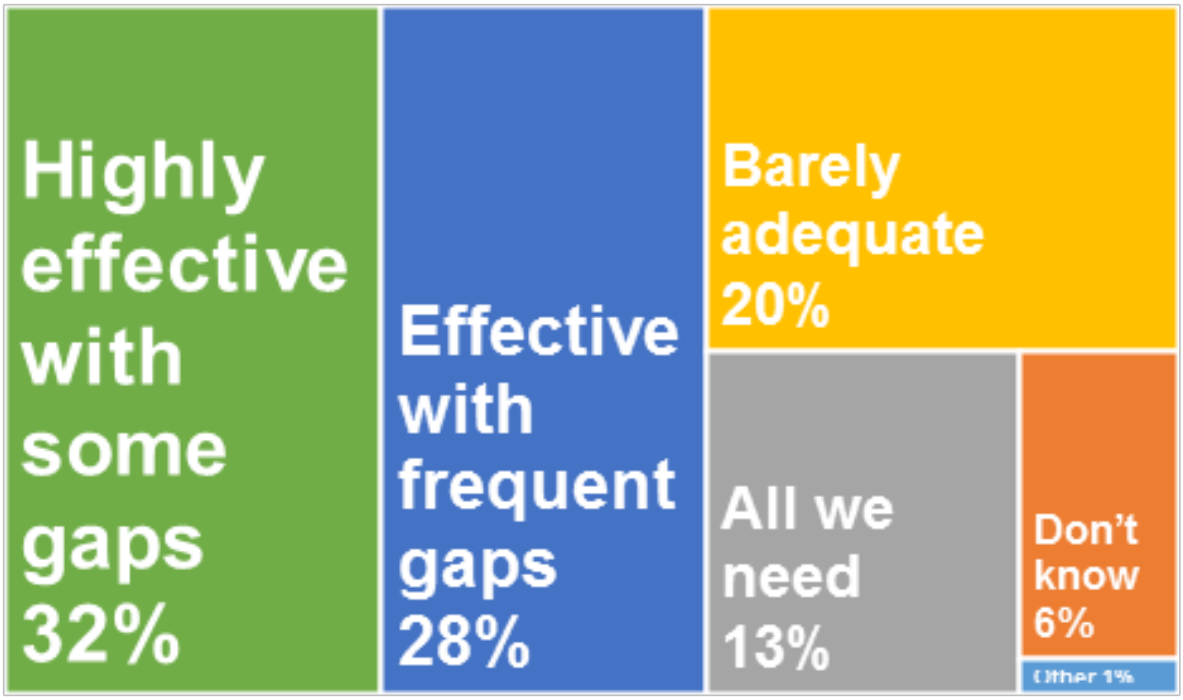

The survey results indicate that only a small percentage of network organizations have all of the visibility they need to effectively troubleshoot WAN performance problems.

As companies continually increase their reliance on the WAN in order to support critical business processes, the inability of the network organization to effectively trouble shoot the WAN will increasingly have a negative impact on those critical business processes. The deployment of new WAN solutions is an opportunity for network organizations to improve their ability to troubleshoot the WAN and hence improve their ability to support the company's critical business processes. The deployment of new WAN solutions also presents network organizations with a challenge. That challenge is that network organizations must have a tool that can effectively manage the new WAN solution throughout its lifecycle. Having such a tool significantly reduces the risk that is associated with adopting a new WAN solution.

As noted, if network organizations want to implement new WAN solutions they need an effective management tool before, during and after that implementation. To exemplify why that is the case, consider the situation in which a hypothetical network organization is interested in potentially adopting a Software Defined WAN (SD-WAN) solution. Prior to beginning its evaluation of SD-WAN solutions, the network organization needs to have an effective management tool that enables the organization to baseline the performance of its WAN and the performance of the business critical applications that transit the WAN. This is necessary so that

the organization has the performance data it needs so that it can evaluate the impact of implementing one or more SD-WAN solutions.

Before deciding to adopt an SD-WAN solution the network organization decides to run a proof of concept (POC) of one or more SD-WAN solutions. The primary goal of conducting a POC is to determine whether or not the solution will provide the promised benefits. The sites that are included in the POC must be chosen in such a way that if the solution is effective there then it will likely be successful in the remaining sites. An effective management tool can help the organization to choose the appropriate sites for the POC based on factors such as application and network usage. An effective management tool also provides insight that helps the network organization determine whether or not the solution provides the promised benefits. Because it provides this insight, the output of an effective management tool is a key input into the analysis that the network organization does to determine if it makes sense to adopt an SD-WAN solution.

While conducting a POC provides insight into the performance of an SD-WAN solution, the amount of insight increases as the network organization begins to implement the solution and more sites and more applications are supported by the solution. Using an effective management tool during the implementation phase of adopting an SD-WAN solution enables the network organization to fine tune its use of that solution. For example, the network organization may use the data generated by that tool to decide to change its policy about which WAN links an application can transit.

Unfortunately, the adoption of new WAN architectures, such as an SD-WAN, has the potential to further complicate the task of ongoing WAN management. As a result, adopting a new WAN architecture further increases the importance of having an effective management tool. One of the reasons why adopting an SD-WAN further complicates ongoing management is because SD-WANs introduce a new device into the WAN which must be managed. That device is referred to as a controller and its role is to support the central management of policy that enables network-wide policy definition and enforcement. One of the management challenges associated with the controller is that under heavy load the controller can add excessive delay. Another challenge is that the communications between the controller and the end devices must now be managed.

Another reason why the adoption of SD-WANs has the potential to further complicate the task of WAN management is that many SD-WAN solutions feature dynamic load balancing of traffic over multiple WAN links. Hence, network organizations that are trying to troubleshoot performance problems with an SD-WAN have a new management question they need to be able to answer. That question is: Which link or links did the traffic transit and how did that change over time?

Why is this important?	Having effective WAN management solutions significantly reduces the risk that is associated with adopting new WAN solutions and it enables network organizations to better support the company's critical business processes. 

Hypothetical Company: NeedsToChange

Each of the 7 sponsors was given the description of a hypothetical company: NeedsToChange. The goal was to present each sponsor with the description of a company that has a traditional WAN and ask them to provide their insight into how the company should evolve its WAN.

Even 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 NeedsToChange's WAN. In order to limit the size of the description of NeedsToChange's WAN and yet still bring out a wide array of important WAN options, each sponsor was allowed to embellish the description of NeedsToChange'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 NeedsToChange's WAN that each sponsor received.

1. **Data Centers**
   NeedsToChange has a class A data center in Salt Lake City, Utah. The site has two diversely routed T3 links into an MPLS network 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 NeedsToChange's business critical applications are SAP and Product Data Management (PDM). PDM is NeedsToChange's most bandwidth intensive application, however it is widely understood that NeedsToChange runs its business on SAP and so the performance of SAP is critical. In addition to the applications that NeedsToChange 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, NeedsToChange currently makes relatively modest use of public cloud computing services. However, the company has started to implement Office 365 and 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**
   NeedsToChange supports a modest but rapidly growing amount of real time IP traffic, including voice, traditional video and telepresence.

6. **Internet Access**
   NeedsToChange 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. NeedsToChange is also concerned that it is supporting non-business related Internet traffic that is negatively impacting business traffic.

7. **Mobile Workers**
   Roughly half of NeedsToChange's employees regularly work somewhere other than a company facility.

8. **Guest Workers**
   NeedsToChange's network organization is considering offering guest WiFi access from at least some of its facilities.

9. **Branch Offices**
   NeedsToChange 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. **Branch Office Availability**
    NeedsToChange wants to improve the availability of the WAN access at its branch offices and has established a goal of 99.99% availability.

11. **IoT**
    The company has begun a smart business initiative which the company believes is just the first in a number of initiatives that will quickly drive the need for them to support thousands, if not tens of thousands, of devices.

12. **Visibility**
    In the majority of instances in which the performance of one of NeedsToChange's business critical applications begins to degrade, the degradation is noticed first by the end users. In addition, the time it takes to identify and resolve performance problems has been increasing.

13. **Regulations**
    NeedsToChange is subject to PCI compliance. That is just one factor driving NeedsToChange to seek out ways to increase its security.

14, **Factors Driving Change**

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

- Improve application performance, notably for SAP;
- Reduce cost;
- Increase uptime;
- Reduce the time it takes to identify and remediate performance problems;
- Increase security;
- Reduce complexity;
- Provide access to public cloud computing services in general and Office 365 in particular;
- 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 NeedsToChange to seek alternative WAN designs is the fact that NeedsToChange will not be allowed to increase the size of its network organization.

Vendor Responses

Below is a description of how each of the 7 sponsors suggests that NeedsToChange should evolve its WAN.

Introducing the Next Evolution of Wide Area Freedom

Overview of the NeedtoChange Network Environment

The constructs for wide area networking at NeedToChange (NTC) have remained stagnant for over 20 years. Network connectivity (such as a managed MPLS-based VPN service) is purchased from a Service Provider via a multi-year contract. Then, the networking team rolls out routers to the branch and applies a site-specific configuration that creates the network topology based on a hub-and-spoke (HQ-to-branch) architecture.

The workflow for these network rollouts is rigorously managed with formal project management, specialist personnel and change control processes to ensure any deployment or augmentation to the WAN happens with minimal disruption to the business.

WAN bandwidth is expensive and thus in limited supply, so the skill in WAN management is squeezing the last drops of performance out of a finite resource. At NTC this has been achieved with advanced configurations within the branch routers or the addition of network appliances — both approaches that increase network complexity.

How Cloud-Based IT Consumption is Affecting the Branch

Today's IT environment is being hampered by the rigidity of the wide area network. Historically, traffic has been client-to-server, so a hub-and-spoke WAN design fitted NTC's needs well. Remote branches were clients to the Utah datacenter servers. But now with Cloud IT, traffic patterns have changed. NTC has virtualized its Utah datacenter and the critical Customer Relationship Management (CRM) and Product Data Management (PDM) applications reside on virtualized compute systems.

As the demand for these applications increases, the virtual compute environment flexes to accommodate the workload. This means that the application does not always reside in the same rack or row of the datacenter. In disaster recovery situations, for example, it is relocated to a completely different datacenter. Unfortunately, outside the datacenter, NTC's network architecture is static and cannot easily adapt to dynamic demand. To resolve this inflexibility with the current architecture NTC must either overbuild the network (inefficient and expensive) or reconfigure the network on the fly (manually intensive and high risk).

A similar shift in consumption is occurring on the client side of the network within the branch. Today any NTC employee connecting to the CRM is the client, but only for that application session. NTC has embarked on a new set of IP-based collaboration tools to improve workflow and communications across the organization, including instant messaging, desktop videoconferencing and IP voice. Now any employee in any branch can initiate a desktop video session to any employee in another branch. In this scenario, the employee's PC becomes the host or source of the traffic. This direct branch-to-branch communication is not handled efficiently in an HQ-to-branch (or hub-and-spoke) network architecture.

Unconstrained Networking, Datacenter to Branch

To address new business communication standards and Cloud-based IT, enterprises must re-examine what they need from the WAN. Some key areas of change that should be considered by NTC are:

- Change the network topology from hub-and-spoke to meshed network architecture to facilitate efficient branch-to-branch and branch-to-datacenter/cloud communications
- Manage premium bandwidth with secure Internet offload at the branch
- Reduce WAN operational overhead with centralized network policy enforcement
- Investigate alternative connectivity options on a per state, region or branch location basis
- Treat the datacenter and WAN at NTC as a single entity with common management, monitoring and reporting tools

In order to drive these benefits into its business NTC needs to deploy a virtualized network service WAN environment. This will deliver expansive wide area networking that matches the flexibility of cloud-based IT.

FIGURE 1. NTC virtualized network service architecture

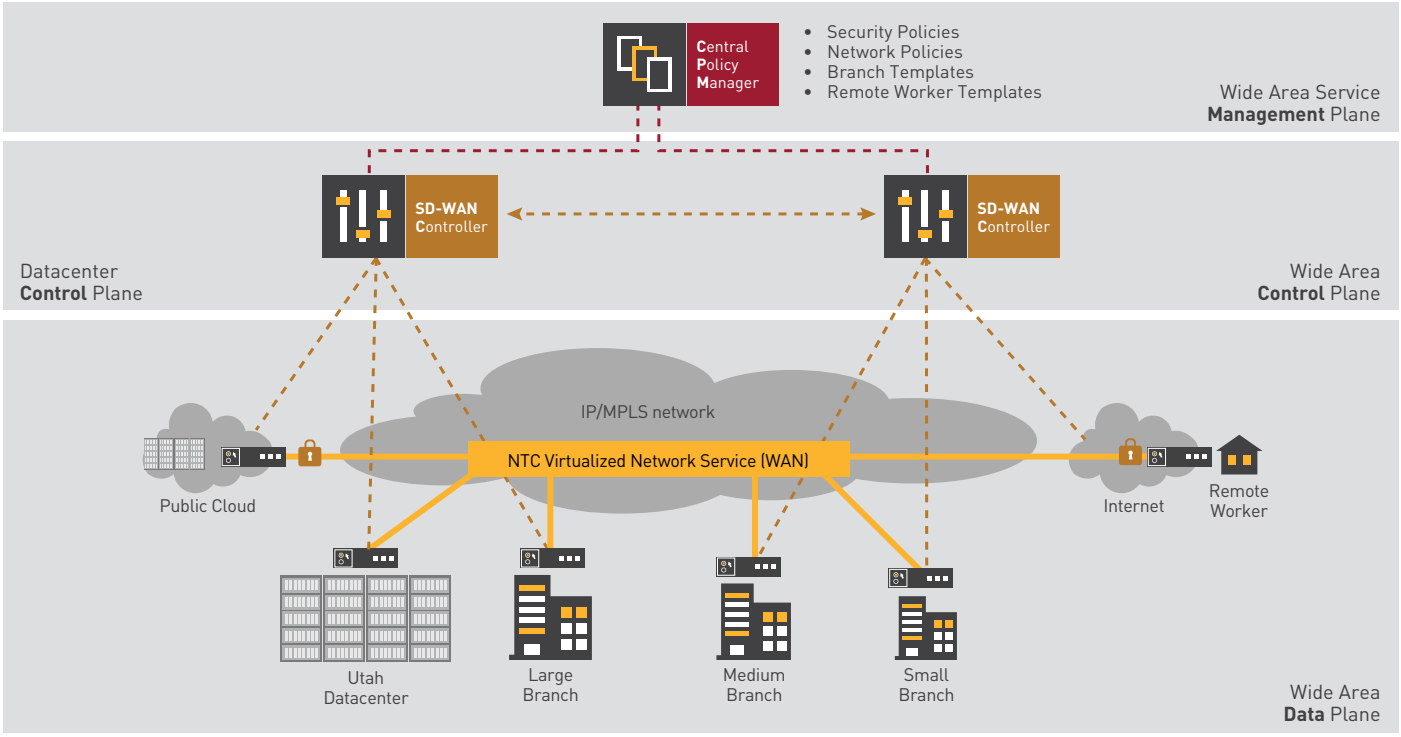

With software defined wide area networking (SD-WAN) there are three key planes (or layers of network functionality) that will assist in this delivery (see Figure 1):

- **Service Management Plane:** A policy system that centrally administers the network templates and policies. This layer should provide the visibility and control of the NTC network via an intuitive GUI. Templates can be created per branch type and automatically deployed when the branch equipment is deployed. All visibility and control aspects of the NTC WAN are managed via this WAN service management layer.
- **WAN Control Plane:** This layer contains the SD-WAN-based controllers that manage the control plane of the NTC WAN. Predominantly deployed in pairs, these controllers manage the network connections between the endpoints (branches, Utah datacenter and public cloud) of the NTC network.
- **WAN Data Plane:** Open compute (x86-based) branch equipment is deployed at the remote branch locations and datacenter connection points, and at the public cloud interconnect to provide enterprise-wide control of the network. These "branch devices" should support both a virtual deployment option (in a public cloud or on an existing branch server) and a dedicated hardware form factor. In either case (virtual or physical) management is provided by the service management planes with data forwarding control provided by the WAN control plane (SD-WAN controllers).

Any-to-Any Network Connections

NTC can implement a fully meshed network architecture to facilitate branch-to-datacenter and branch-to-branch communications. This provides the flexibility to transport inter-site traffic across the most efficient path. Rich IT communication tools can be deployed to enhance the collaboration between branches without the constraints of the rigid hub-and-spoke architecture of the past.

Intelligent Traffic Offload

Via the central policy system, the NTC network team implements the network policy that securely offloads any Internet traffic at the branch (see figure 2). There are three key benefits of this feature. First, the limited IP-VPN bandwidth is only used for business critical voice and data, which maximizes its availability for critical data. Second, via this policy a secured inter-branch tunnel can be created to force high-bandwidth usage across an encrypted Internet path. The third benefit of the intelligent traffic offload feature is the ability to use the Internet connection as a backup link in case the IP-VPN circuit fails. Using the same template-based policy push from the central management system, all branch traffic can be encrypted and sent over the Internet to the Utah datacenter. This provides additional resiliency and enables NTC to improve network availability at the branch.

Policy-Based Network Management

With the right SD-WAN solution, management and monitoring of the NTC WAN environment can be simplified via a policy-based manager. The policy manager can create policies for NTC traffic at many levels and these policies can be simply grouped together into templates. The templates can be deployed automatically when an application changes (for example, if CRM is relocated to the disaster recovery

> Traditional hub-and-spoke WAN designs inhibit the efficiency of today's rich collaboration tools

FIGURE 2. Using policies to intelligently offload traffic

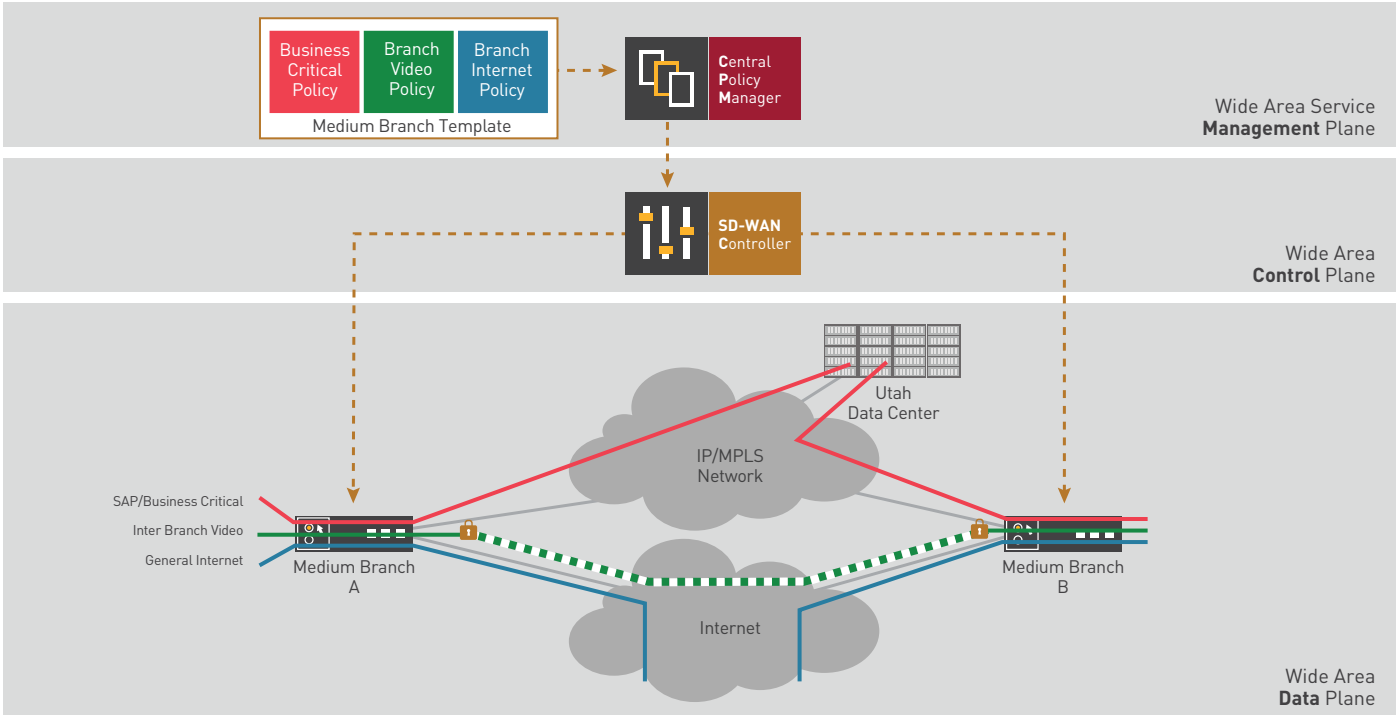

datacenter) or a new branch is added. These policies can be split into four key types:

- **Application policies:** These are the conditions each application needs to function across the network and can include specific security, quality of service and resiliency requirements. For instance, a policy for the CRM application may include QoS policies for interactive, batch and print traffic. This provides granular control of how individual flows are handled by the network. The CRM print traffic at the branch can be lower in priority to ensure that it doesn't affect the performance of the critical interactive traffic.
- **Branch policies:** These include the network functionality for specific or types of branches in the network. A branch may be a physical location or a virtual location, such as a public cloud interconnect where a new NTC application resides. NTC networking staff can deploy policies for the use of backup links, enforce encryption or automate equipment password changes across all branches.
- **Security policies:** User-based permission means network security can be managed by a specialty team. The security team can set the security policies on an application or branch level. For instance, the team can specify the mandatory time period for all branch device password changes or encryption keys exchanged. Once this policy is set it is called on by the operational team in the deployment of applications or branches. User-based permission functions ensure that the security policies are implemented, which guarantees compliance with NTC's security framework. And the single control point for policy enforcement reduces the complexity of regulatory/ industry auditing.
- **Network policies:** These are the network wide policies that control the flow of traffic across the NTC network. Examples include the overall quality of service policy that prioritizes CRM, PDM and voice traffic over general inter-office traffic.

Using these policies, templates for deployments can be created, such as the Intelligence Traffic Offload example provided earlier. Any number of policies can be grouped into a template. For example, a template could be designed for all medium-sized branches. It could include a policy on application forwarding (the three colored flows shown in Figure 2) plus a standard security policy for equipment so encryption keys and passwords are changed in accordance with any regulatory or business requirement. These templates could be called on whenever a new site is added to the network.

Relying on templates reduces the need for specialized personnel to visit the branch location. The branch equipment can be couriered to the branch manager with simple instructions to connect to the WAN links. Once connected the device will "call home" to the policy manager, authenticate and the template configuration will be sent over the WAN to the device.

Network Functions Virtualization

SD-WAN also provides the opportunity to reduce the reliance on external network devices at the branch. For many enterprises the only option to enhance network performance and security has been to deploy high CAPEX physical devices (such as firewalls and WAN accelerators) at the branch. These point solutions increase CAPEX up front and increase network complexity, which in turn drives up OPEX for maintaining the WAN environment.

FIGURE 3. Seamless interworking with SD-WAN

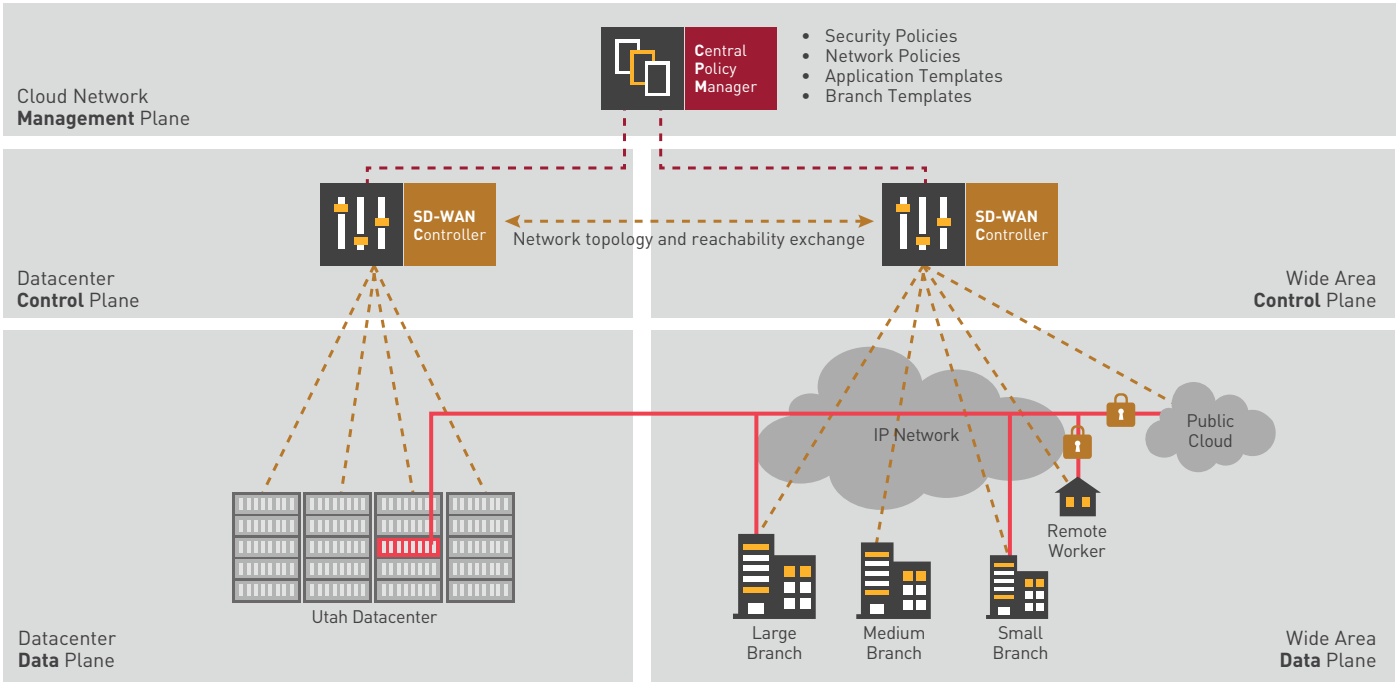

Comprehensive SDN-based WAN solutions use Network Functions Virtualization (NFV) to provide this enhanced functionality. Software features are "chained" into the traffic flows to and from the branches. By adopting this approach, NTC could enable a more robust and dynamic end-to-end policy that inserts the right network functions into the right locations to ensure data integrity at the branch, without the large CAPEX drain of physical devices.

Service Provider Independence

SDN provides the separation of the NTC WAN (overlay service) from the underlying IP transport (MPLS IP-VPN) network. With traditional WANs these are tightly integrated; with SD-WAN they can be completely separated. This separation delivers a new set of options for getting bandwidth across the WAN and into the branch. It means that NTC can procure the required IP connectivity services on a per-branch or per-region basis and use these links as an underlay network for the WAN. This gives NTC access to the world of competitive local carriers and alternative access technologies. If IP-VPN connections aren't available at a site then 4G/LTE mobile broadband, cable or DSL technologies can be deployed to provide the connection.

Summary

To gain maximum benefit from the move to SD-WAN, the operation and purpose of the network(s) in the enterprise need to be rethought. The network is there to connect the new cloud IT environment to business users regardless of their locations.

Implementing SDN in the datacenter and across the WAN is a great start. However, to drive a change across the whole business these two critical network islands need to operate in concert and that means removing any management boundaries that separate them.

The key to seamless interworking is the use of a single network policy framework that distributes business policies and network intelligence across both domains. SD-WAN provides the opportunity to achieve this. If SD-WAN is controlling the network that underpins cloud applications and is managing the connectivity across the WAN towards the applications' end users (employees and/or customers) then centralizing this intelligence onto an overarching policy and control framework makes sense.

With the right SDN-based WAN solution, NTC can achieve exactly this: unconstrained networking for the datacenter and beyond. To gain maximum benefit from the move to Cloud IT, NTC needs to centrally manage the datacenter and wide area networks with a single policy framework. This simplifies the overall network configuration. The enterprise can change a security policy once and have the network automatically roll that change out. Add a new application to the business and instantly deploy the updated network, branch and security policies. No more waiting for project rollouts, no more specialist personnel needed at the branch.

With this new network environment in place, NTC will get wide area networking on its terms.

Automated policy-based networking significantly reduces the complexity of regulatory and industry compliance

nuagenetworks
From Nokia

Cato Networks Re-Architects NeedToChange (NTC) WAN to Boost Capacity, Availability, Performance and Security

Current State of NTC's WAN

The Wide Area Network (WAN) was built to connect static and physical locations, not today's fluid and dynamic networks. Like many other companies, NTC depends on expensive and limited MPLS-based WAN for remote branch connectivity. NTC backhauls internet traffic as their small and medium remote sites don't have a security stack in place, resulting in the "trombone effect" (high latency and poor user experience) when accessing a business application hosted on SaaS and IaaS platforms. NTC has no control and visibility for employees working outside a company facility, and the plan to adopt SaaS applications and to connect thousands of IoT devices requires a new architecture to support this business transformation.

NTC's Network Challenges

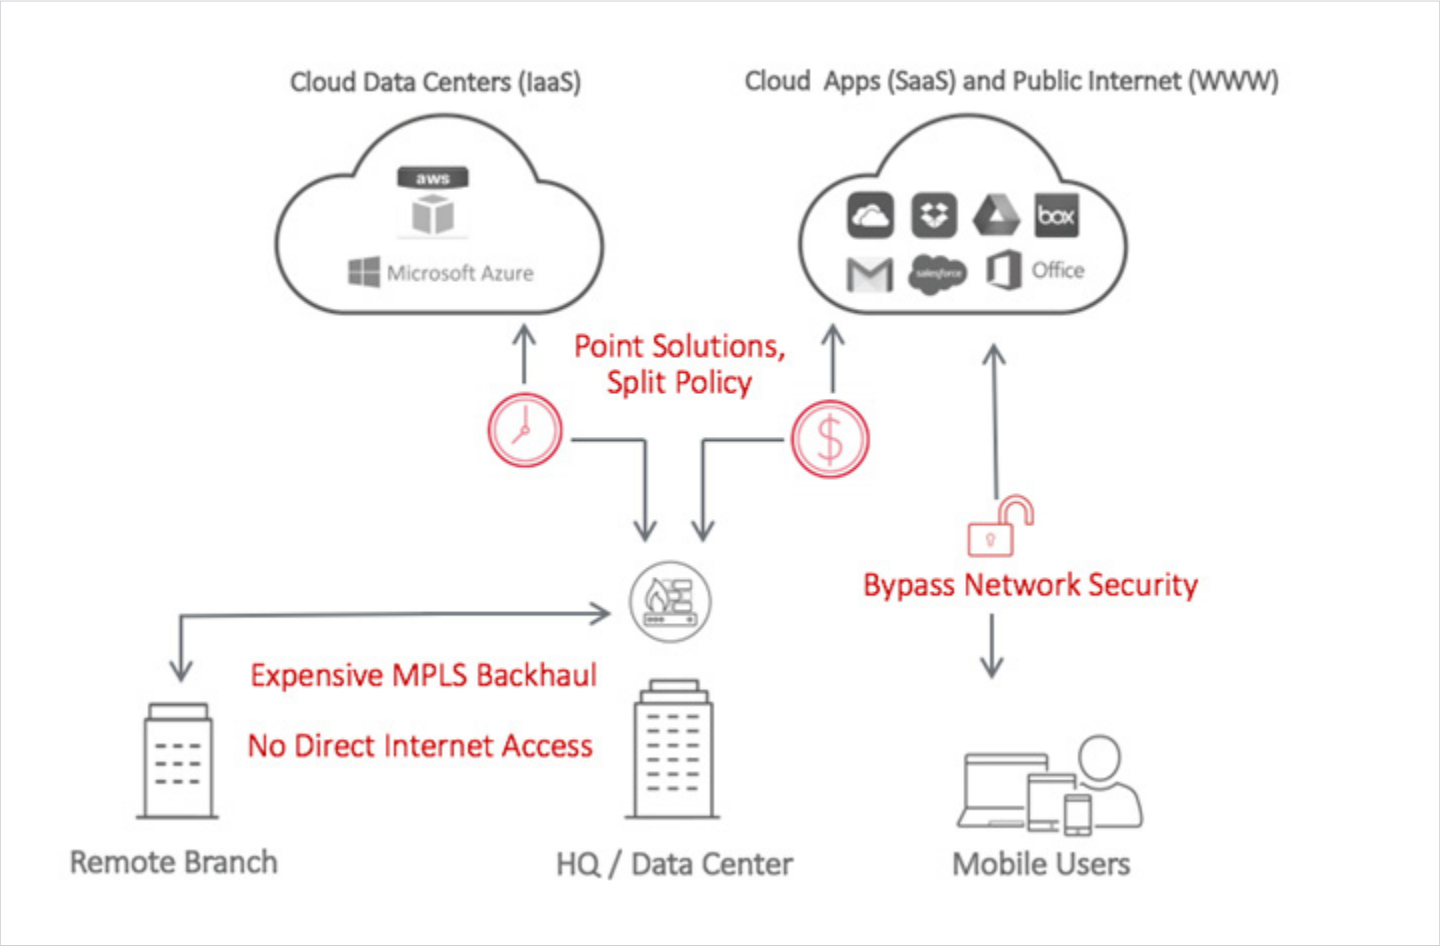

Cato Networks: Software-Defined and Cloud-Based Enterprise Network

Cato will enable NTC to efficiently and securely connect all branch locations, the mobile workforce, physical and cloud data centers, into a global software-defined and cloud-based secure enterprise network. All outbound traffic, both WAN and internet, is consolidated in the Cato Cloud, where a set of elastic and agile security services are applied to protect access to enterprise applications and data, regardless of their location. The Cato Cloud service is comprised of the following pillars:

Cato Cloud Network

A global, geographically distributed, low-latency and SLA-backed network of PoPs, interconnected by multiple tier-1 carriers. NTC will connect to Cato over optimized and secured tunnels. Physical locations use the Cato Socket; a small, zero touch, tunneling device that controls and splits traffic across WAN links based on business policy. Traffic transmitted via the internet is encrypted and optimized end to end. Cloud data centers, like Amazon VPC, use a virtual version of the Socket (Cato vSocket). Lastly, mobile users use the Cato Client to establish a secure tunnel for laptops, tablets and smartphones.

Cato Security Services

A fully integrated suite of enterprise-grade and agile security services directly built into the cloud network. The services include a NG firewall, URL filtering, anti-malware and more, have no capacity constraints and are continuously updated to introduce new capabilities and adapt to emerging threats. The integrated network and security stack enables NTC to enforce its corporate policy on all traffic, WAN and internet, from all locations and users.

Cato Cloud High Level Architecture

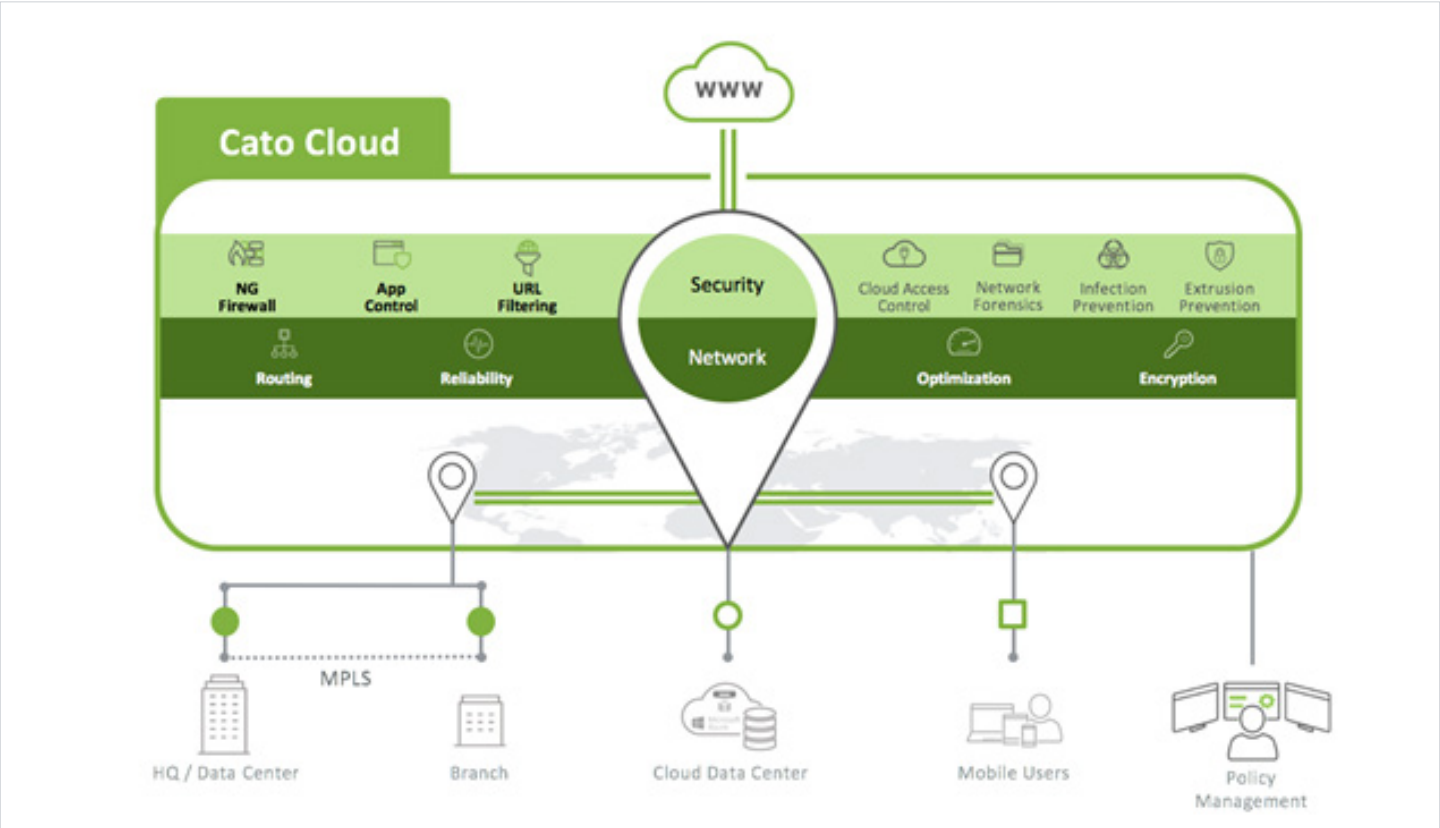

Recommendation:
Migrate to Cloud-Based SD-WAN with Built-in Security

To meet NTC's business needs and to future-proof the network, Cato recommends a cloud-based SD-WAN architecture that connects, secures, and simplifies NTC's global WAN following the 3 steps below.

Step 1: Expand WAN Capacity and Availability, and Add Policy-Based Routing to Meet Application Delivery Goals

Last Mile extension

NTC should deploy additional internet links in the locations currently served only by MPLS. Cato suggests NTC considers replacing MPLS with Cato, dual ISP links and optional 4G/LTE backup per below. Ultimately all sites will have either MPLS+Internet or 2 Internet links.

Policy-based routing

NTC will deploy a Cato Socket at each branch location and connect it to the available MPLS, internet and 4G links. Specifically, the internet links will connect the branch to the nearest available Cato PoP. Cato classifies and dynamically allocates traffic in real time to the appropriate link based on application policies and link quality (availability, utilization, latency, packet loss). NTC will specify these policies for SAP, PDM, Voice and Video to set prioritization and required service levels. With Cato, even the "internet leg" enjoys SLA-backed latency compared with the unmanaged public internet so it can offload more traffic off the MPLS link.

High availability, resiliency and quality

The Cato Socket can drive the WAN links in Active/Active mode to boost overall capacity and reach 99.99% availability. Forward Error Correction (FEC) is intelligently applied to reduce the impact of packet loss on latency and quality.

Latency control for WAN and cloud locations

Unlike appliance-based approaches, Cato's SLA-backed backbone guarantees latency and availability over the long haul WAN (for national and global locations). The Cato backbone is fully redundant across servers, PoPs and regions and is co-located with Microsoft Azure and Office 365 datacenters for optimized access.

Meeting application delivery goals

With all the enhancements above, NTC will improve access to SAP, PDM and Office 365 and is in a great position to eliminate MPLS even for latency sensitive applications like voice and video.

Step 2: Eliminate Internet-Bound Traffic Backhauling with Secure Direct Internet Access

With all branches connected to Cato Cloud, NTC employees can directly access the internet and cloud applications (i.e. office 365) behind Cato's enterprise grade and cloud-based security services. These services protect branch and mobile employees against threats, and can restrict access to critical applications as well as applications that violate corporate policies. All security capabilities are delivered without dedicated branch security appliances or regional co-location facilities.

Step 3: Extend the WAN and Security to Cloud Data Centers and the Mobile Workforce

NTC will use Cato vSockets to connects any IaaS platform (Such as AWS and Azure) to the Cato Cloud, making it an integrated part of the network. Instead of backhauling DR traffic over MPLS, NTC will use direct internet access to route traffic via the Cato Cloud between the data center and the DR location. NTC mobile users will deploy Cato Clients to connect Windows, Mac, iOS and Android devices to the the nearest Cato Cloud PoP. Users gain secure and latency-optimized access to NTC's physical and cloud datacenters as well as public cloud applications.

Transformation Done: NTC's New Secure and Software-Defined WAN

With full migration to the Cato Cloud, NTC will achieve the following:

- All NTC's data centers, branches and users are connected to a high capacity, redundant, optimized, affordable and secure WAN.
- Full protection of all traffic for both datacenter, cloud and internet resources that seamlessly scales to accommodate growth and adapt to emerging threats.
- Central management of all policies including full site-to-site mesh, network segmentation, access control, and security.
- Instant deployment of new sites with Cato Socket 10-minutes self provisioning.
- Full visibility into the network usage and security events for every location, application and user that simplify end-to-end troubleshooting of performance and security issues.

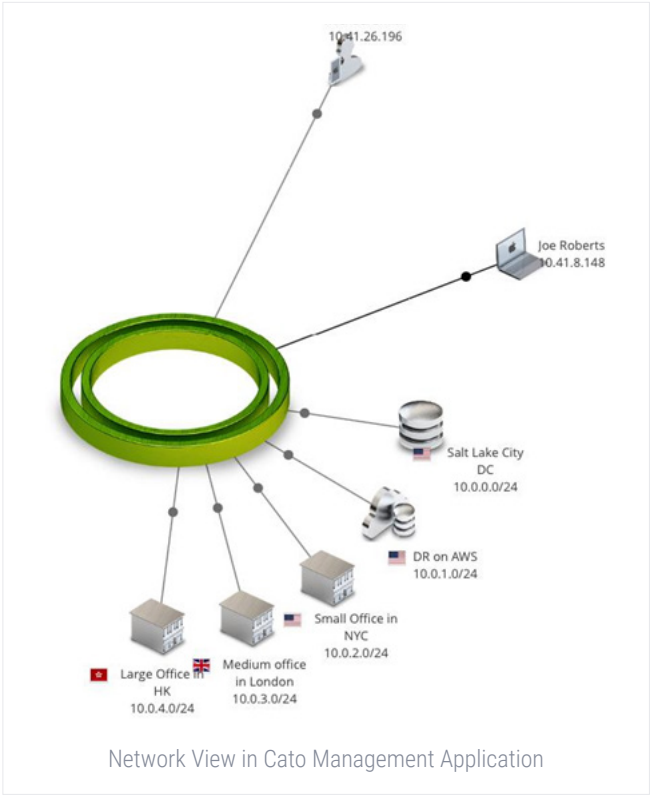

Network View in Cato Management Application

Summary

Cato provides NTC a flexible, software-defined WAN with built-in secure direct internet access, a SLA-backed global backbone, and seamless integration of cloud infrastructure and mobile users. By moving to Cato, NTC eliminates complexity, reduces costs, streamlines day-to-day operations and ensures scalability for the enterprise's future growth.

TALARI Networks

NeedToChange WAN Refresh
Delivering a Failsafe Software Defined WAN

Overview

The NeedToChange WAN had served them well throughout the 2000s, but with company growth and the introduction of new services and applications, it was starting to show its age. Users were experiencing slowdowns while accessing critical applications such as SAP, PDM and Office 365, and complaints to the IT help desk were beginning to rise. Also, with the heavy use of VoIP and traditional applications as well as the increased use of cloud applications, complaints emerged that calls were not connecting and the voice quality was garbled while access to key applications was not meeting user expectations. Beyond application performance issues, some offices had experienced network outages that left them unable to function for periods of time.

The IT staff was starting to be stretched thin by the effort to maintain the existing network. Every time a new application was introduced, they were forced to manually update the existing infrastructure. The outages that occurred created an atmosphere of "manage by crisis." The IT staff was starting to reject or slow roll the introduction of new applications – harming the ability of NeedToChange to maintain its leadership position in the market.

NeedToChange decided to address these problems and update their WAN to support the company's future growth with Talari's Software Defined WAN solution which delivered a failsafe WAN that saved money, dramatically improved their users' experience, increased overall productivity, and best of all, stopped their IT team from lurching from one crisis to another.

Moving to a Network with Talari

As NeedToChange began the WAN research process, four requirements were clear going in:

- They would need more bandwidth at every branch office
- They couldn't expand their MPLS commitment due to cost constraints
- Direct access to cloud applications from the branch was required
- They needed to increase branch security to support access to external resources

This meant they had to use broadband Internet connections since they were the only option to cost-effectively increase available bandwidth. Also, the Internet connections served as the primary method to access cloud applications. The company quickly found a mixture of DSL and Cable providers and established an extra connection for the offices that didn't already have an Internet link. While the broadband connections were far less expensive than MPLS, they offered significantly more bandwidth.

They purchased Talari for each of their physical offices, selecting the Talari Virtual Appliance VT800 for the small offices, and the Talari Appliance E100 for the medium and large offices. In the Data Center, they installed a high-availability pair of Talari Appliance T3010s. Talari Aware, the centralized management system that gives IT staff the ability to configure, monitor, and analyze a Talari SD-WAN, was also deployed in the data center location.

They implemented in phases, starting with the Data Center and the offices that were reporting the most problems. The physical appliances leveraged Talari's Easy Install capability which allowed plug and play device deployment at the branch locations. This work was done by non-technical employees and eliminated the need for IT staff to travel to those offices during the SD-WAN rollout.

With the introduction of the Talari SD-WAN solution, the company decided to stop backhauling Internet traffic. Since each Talari appliance supported a stateful zone-based firewall and network address translation capabilities, the Talari solution delivered an easy and cost-effective method to deploy the incremental capabilities required for secure local Internet access.

The New Talari Software Defined WAN

The Talari SD-WAN built secure, full mesh, on-demand virtual connections between the offices, the data center, and the cloud. These encrypted connections are tunnels that are abstracted from the underlying network links. Each application uses a virtual connection, with the Talari network controller utilizing policy-driven decisions to ensure the highest possible performance for each specific application.

To make path decisions, the Talari solution collected data with every packet to determine the loss, latency, jitter and congestion of every possible path through the network in each direction. This collected information, based on real network traffic and not probe data or round trip pings, was combined with the centrally defined policies regarding prioritization, bandwidth share and security to make decisions about individual applications. Thus, the WAN became an intelligent network, able to accomplish the goals of the organization in the context of the actual real-time state of each WAN link.

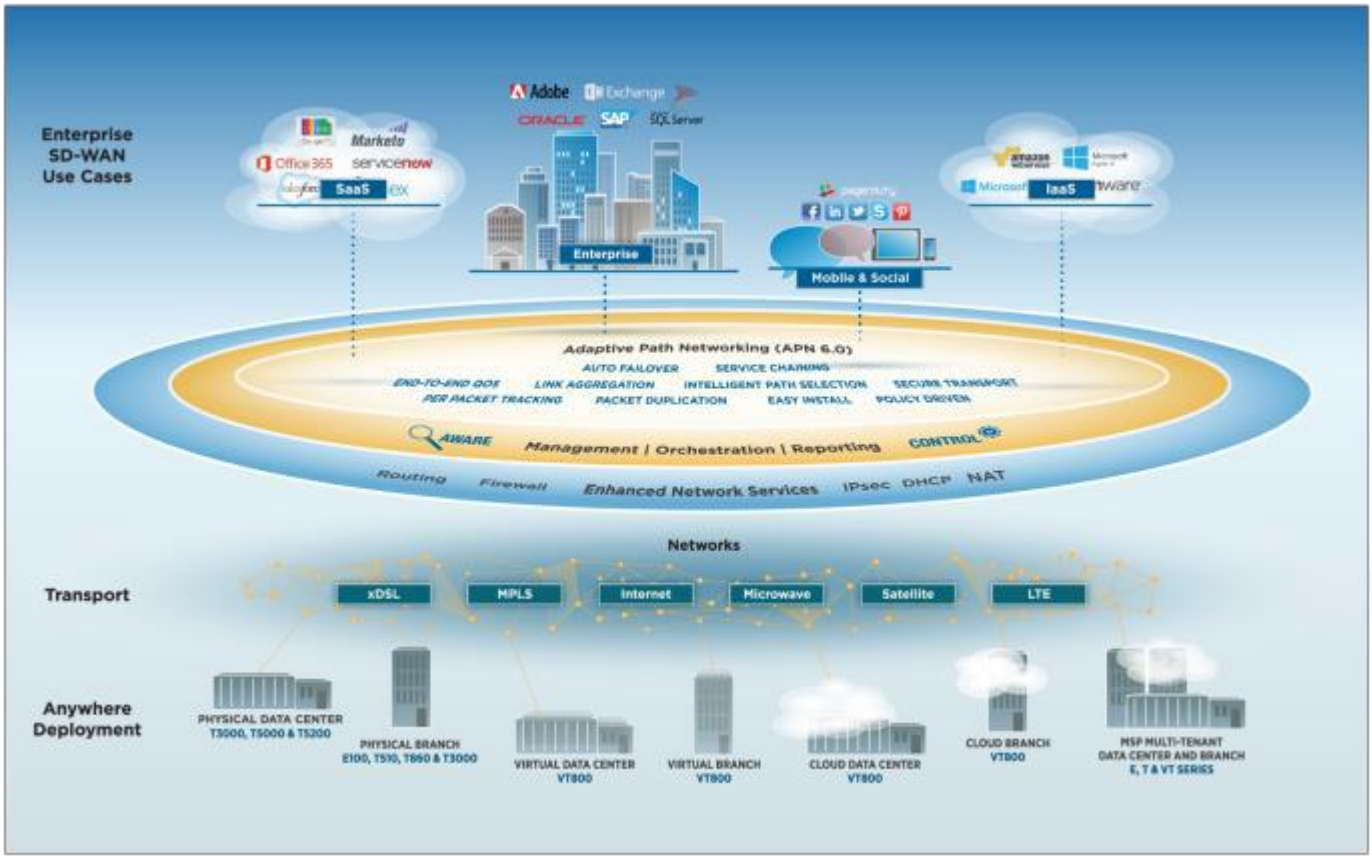

Figure #1, Talari Reference Architecture

Results for NeedToChange

The new WAN eliminated the problems that were plaguing the old network, improving the user experience while meeting the security goals of the organization and positioning them for the future. The key results realized with the new WAN included:

Cost Savings Coupled with Greater Capacity

Upon installation of the Talari SD-WAN solution, NeedToChange eliminated two large ongoing expenses. One was the maintenance on the firewalls in place at the large offices by replacing them with an integrated Talari Firewall. The second eliminated expense was due to the replacement of one of the MPLS circuits at the Data Center. While there had been much debate over whether they should scale back their MPLS commitment, the Talari solution proved it was more than able to convert broadband Internet links into a business class infrastructure.

Failsafe WAN Where Outages Go Unnoticed

By monitoring every path through the WAN, including to and from the cloud, the Talari solution detected link outages within a fraction of a second and automatically shift all WAN traffic to an alternate path. This prevented outages from disrupting offices; users now didn't even notice when an outage occurred. Because error detection was so fast, even small spikes in latency or loss were detected and avoided, improving overall application performance. This approach allowed NeedToChange to deliver a business class WAN with better than 99.99% availability over a hybrid network comprised of commodity Internet and MPLS links.

Since VoIP was a key and growing application and voice traffic is latency and packet loss sensitive, voice packets were duplicated across two diverse WAN paths which ensured basically no loss for voice packets and the lowest possible latency. This resulted in an immediate improvement in voice quality and an elimination of dropped calls. And because of the broadband Internet links that had been added, there was an abundance of bandwidth available to support all applications.

Improved Application Performance

By forwarding direct traffic away from poor quality links and duplicating voice packets, the performance of business applications significantly improved. But more was needed, particularly during times of congestion when applications had to compete for available bandwidth. A combination of application prioritization and bandwidth reservation was used to ensure that critical applications such as SAP and PDM were given a share of network resources and were never choked out by lower priority applications.

While the default categorization could prioritize traffic, and assign the appropriate SD-WAN services to each type of application, specific rules were developed for some applications via the centralized configuration system. This allowed performance tuning on SAP and PDM traffic, decreased the share of bandwidth assigned to guest Wi-Fi access and public websites, and directed Internet traffic through the firewalls. This assured that during times of congestion, non-critical traffic would not choke out critical traffic. While the addition of the Internet links had reduced the likelihood of congestion, a link failure could quickly reduce the amount of bandwidth available to an office. The prioritization policies would immediately come into play if that occurred, preventing outages from impacting critical application availability and end user productivity.

Security Protocols - Uninterrupted and Reinforced

The new Talari SD-WAN conformed with NeedToChange's established stringent security policies. Since the Talari solution did not store any packets, worked with source encrypted packets, and did not provide external access to packets, the security protections that were already in place to achieve PCI compliance were left intact. All Talari destined WAN traffic was secured with the company choosing to use 256-bit encryption, header encryption, rotating encryption keys and trailing authentication checksums to ensure that data sent across Internet or MPLS links could not be read or spoofed. Also, the Talari stateful zone-based firewall delivered packet filtering and network address translation functions to further secure the branch sites.

Cloud Access When and Where Desired

NeedToChange was beginning to invest heavily in the cloud, including hosting their disaster recovery data center in AWS and mandating cloud options for new applications. Talari's SD-WAN solution allowed them to seamlessly incorporate the cloud into their WAN. Using the Talari Virtual Appliance VT800 as a Cloud Gateway, all traffic to and from the Internet used Talari's secure conduits. This created a secure and reliable connection to the cloud or co-location sites and eliminated any application performance problems caused by failed links or poor quality Internet connections.

By using the dynamic conduit feature of the Talari SD-WAN, the configuration of cloud and Internet access directly from each office was easy. With Talari's dynamic conduits feature, a secure tunnel is built to the cloud on demand when it is needed with no requirement to preconfigure anything. As additional offices were added, they automatically had a secure and reliable connection to the cloud with just the check of a box.

Performance and Cost-effective Scale to Deliver IoT

As NeedToChange proceeds with their IoT deployment, they will need a robust and secure infrastructure to link the sites that contain the thousands of IoT devices. This infrastructure will need to deliver the required bandwidth that will allow the IoT devices to effectively exchange information with the central control infrastructure which can be on-premises or in the cloud. The core capabilities that the Talari SD-WAN solution delivered to support the reliable, secure and performance requirements of user applications can easily be extended to support the IoT deployment in any location.

Segmentation to Support Guest Access

Talari comes with supports for VLANs, Virtual routing and forwarding (VRFs) and intelligent path selection which allowed for the secure handling of guest WiFi services within branch locations. These capabilities ensured that guest user traffic was isolated and restricted to accessing a limited set of public resources, such as the Internet.

Visibility and Actionable Analytics

While users didn't notice intermittent quality errors and link failures, the Talari management interface collected and displayed this information to the IT staff. They could see the performance of individual links and the aggregated performance of their telecom providers. This information was invaluable in helping them obtain support from the provider and to negotiate better rates. It also assisted with the troubleshooting of WAN and link issues. The benefit of a Talari SD-WAN was that it automatically and rapidly remediated the WAN when impediments were found, allowing IT staff to troubleshoot issues in a lower stress environment since they could address the network issues without having to deal with key applications being unavailable.

The correlated information on applications available from the management interface also helped the IT staff ensure they were meeting application SLAs and identify the root cause of WAN issues. By running reports on application performance across the WAN, they could show each business unit the quality score for their specific applications. This led to a more collaborative environment between IT and the other business groups, and helped IT tune the WAN to support the company's application mix and priorities.

Fast and Simple Implementation and Administration

One of the best results of the new Talari WAN was the ease of maintenance. Policies were centralized and changes were made in one location and easily pushed through the network, even outside of maintenance windows. By using Talari templates, IT staff reduced the time required to perform deployment changes while helping them maintain a consistent configuration. Also, new applications automatically used default behaviors and then were customized as needed. In addition, the increased visibility into network and application performance made it easy to identify areas for improvement. Beyond the central administration of the network, enabling a branch site was easy using Talari Easy Install. With this capability, appliances were staged by IT staff at the central controller and only required a non-technical person at the remote site to unbox and plug in the device to bring it online.

The End Result

With their new Talari-based failsafe WAN installed, the IT staff found themselves able to respond more quickly and positively to application requests from the business units. This gave them time to think proactively about new ideas that could help the company grow. Confident that their WAN was up to the challenge, they could add more video communications options, expand their use of SaaS applications, and push much of the needed infrastructure growth to the cloud. Also, they had budget to invest in these ideas with the decrease in telecom costs.

Best of all, the company and its employees noticed the difference. Productivity was up at offices as outages stopped interrupting work and access to important applications was always available and high quality. Finally, the Talari SD-WAN solution allowed the business to move quickly when opening new offices or acquiring new businesses.

Talari Networks, Inc.,
1 Almaden Blvd, Suite 200
San Jose Ca, 95113

Phone: +1 408.689.0400

info@talari.com | www.talari.com

About Talari Networks

Talari Networks, the trusted SD-WAN technology and market leader, engineers the internet and branch for maximum business impact by designing failsafe WANs that deliver superior business-critical application reliability and resiliency, while unlocking the simplification and cost reduction benefits of branch consolidation.

Talari delivers a comprehensive solution, supporting a variety of network services in physical, virtual and cloud locations, which can be acquired through perpetual licensing, monthly subscription rates or as-a-service. Passionate and committed to their customers, Talari has incorporated eight years of innovation into five generations of product and is successfully deployed across thousands of sites in over 40 countries.

SD-WAN for People, Places and Things at NeedsToChange Corp.

Introduction

This document outlines Cradlepoint's approach to a new SD-WAN for NeedsToChange Corporation (referred to as NTC). The guiding principle was to meet the NTC's operational, performance, and security requirements while significantly reducing one-time capital costs and recurring operational expense.

Cradlepoint Overview

Cradlepoint is the global leader in software-defined and cloud-delivered network solutions for connecting people, places, and things over wired and wireless broadband. More than 15,000 enterprise and government organizations around the world—including 75 percent of the world's top retailers and 50 percent of the Fortune 100—rely on Cradlepoint to keep critical sites, workforces, vehicles, and devices always connected and protected.

Cradlepoint NetCloud™

Cradlepoint NetCloud™ is a software-defined and cloud-delivered platform that powers and extends a portfolio of LTE-enabled routers with unified management, overlay networking, and virtualized network services. The NetCloud platform consists of the following elements:

NetCloud Manager (formerly Enterprise Cloud Manager) is a single-pane-of-glass cloud management platform that goes beyond ease-of-use to provide the "ease-of-scale" needed to connect hundreds of thousands of people, places, and things distributed around the globe. NetCloud Manager capabilities include:

- Simplified configuration with mass templating
- Zero-touch deployment capability
- Schedulable software and configuration updates
- LTE SIM and carrier management
- End-to-end policy management
- Orchestration and automation
- WAN analytics and health monitoring

NetCloud Engine is a cloud-based Network-as-a-Service that provides a private virtual overlay fabric across the public Internet. Its SDN architecture consists of a distributed data plane that runs on standard virtual machines within public cloud datacenters throughout the world. Each of these data plane entities, called ServicePoints, can host one or more virtual overlay networks, which are called Virtual Cloud Networks (VCNs). The ControlPoint is a collection of micro-services that together comprise the SDN control plane and provide orchestration, oversight, and management of the service. The ControlPoint also gives a VCN its self-organizing, self-optimizing, and self-healing properties.

ServicePoints also enable the integration of virtual network services utilizing Cradlepoint's Network Service Virtualization (NSV) technology. NSV is a distributed, micro-services form of NFV that runs each service function as a discreet process within a VCN's packet path. NSV can also provide a "last-in-chain" egress to external VNFs or cloud services, like firewalls or secure web gateways. A ServicePoint also can act as a Secure Cloud Gateway (SCG) to provide a secure egress point from a VCN to the Internet for connecting to public cloud, SaaS, and the web.

NetCloud Services provides a library of virtual network services based on Cradlepoint and third-party technologies. These services run within the VCN overlay—using NSV—or at the Edge on **NetCloud OS**, the Linux-based open network operating system that powers Cradlepoint routers. The following is a summary of available NetCloud Services:

- Carrier-grade NAT
- PKI-as-a-Service
- Overlay DNS with Active Directory integration
- Distributed next-gen firewall
- Micro-segmentation at the user, app, or device level
- Threat management with IPS/IDS
- App control—1,500+ business and SaaS apps
- URL/content filtering
- Network Access Control (NAC)

NetCloud Client and **NetCloud Gateway** provide an on-ramp to VCN overlays for standalone and router-attached resources, including PC, mobile, and IoT devices. NetCloud Client provides LAN over WAN services for seamlessly connecting mobile users and devices anywhere to private and public cloud applications and resources. It supports Windows, Mac, Linux, Android, and iOS operating systems. NetCloud Gateway provides the same functionality within Cradlepoint routers for IP-attached users and devices.

Cradlepoint Routers

Cradlepoint offers an extensive portfolio of purpose-built, LTE-enabled routers for branch, in-vehicle, and IoT applications. For the NTC network, Cradlepoint has deployed the AER Series of converged, multi-WAN Edge routers and COR Series of IoT routers.

Cradlepoint SD-WAN Solution

BRANCH OFFICE WAN

For NTC's three branch offices, the Cradlepoint AER Series Advanced Edge Routers provide an "all-in-one" solution for WAN, LAN, guest WiFi, and IoT connectivity. The high-end AER3100 routers used in the NTC network cost less than $2,000 each and includes:

- Multi-WAN – MPLS, MetroE, Cable/DSL, WiFi, and LTE
- 12 Ethernet ports including four ports with PoE support
- 802.11ac WiFi – 3x3 MIMO, dual-band concurrent
- Up to two integrated 4G LTE modems, each with dual SIMs

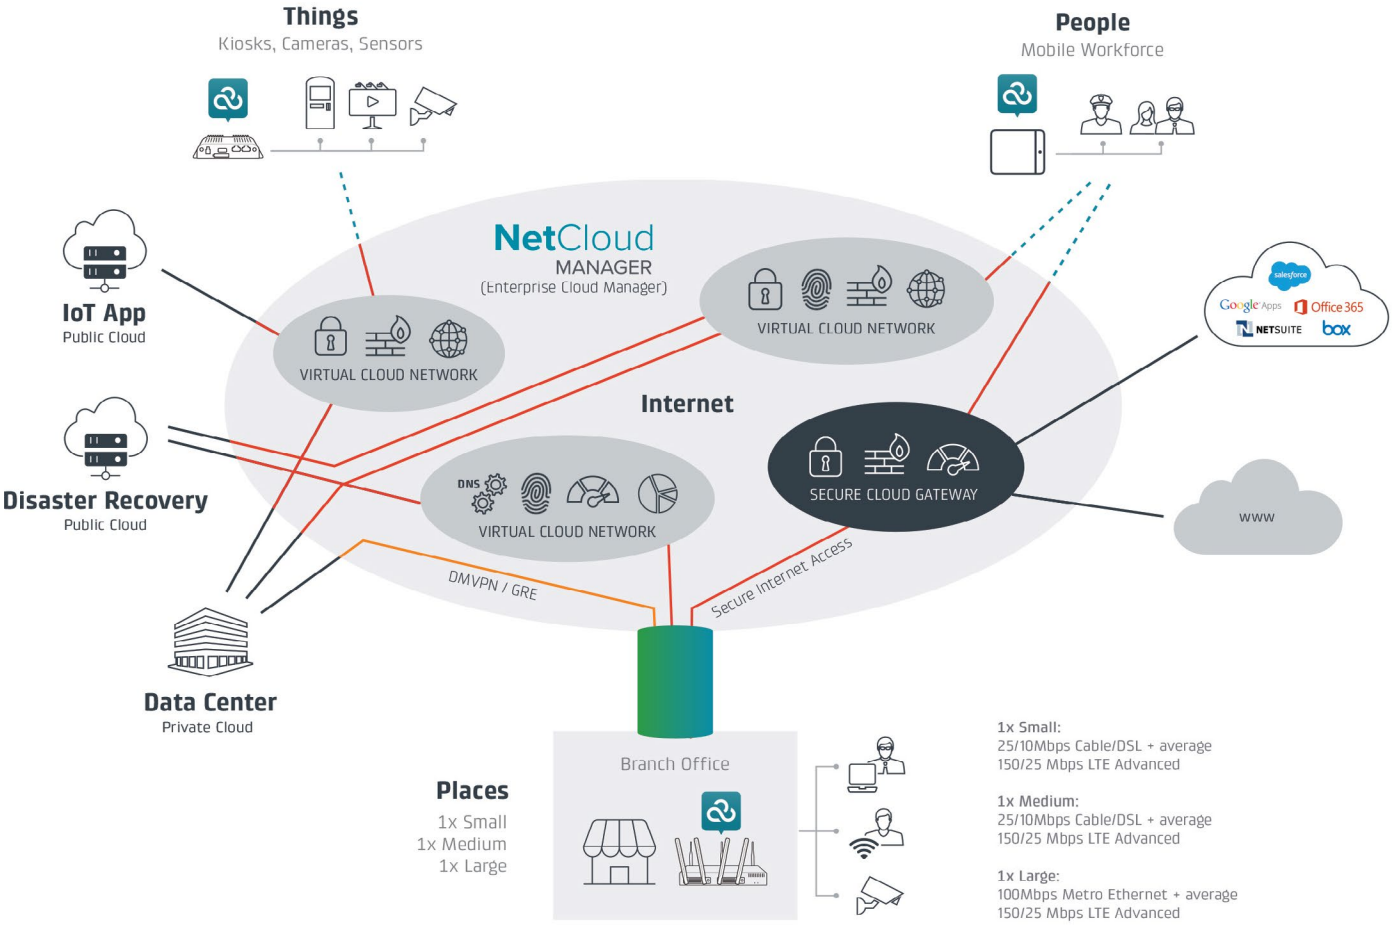

Replace MPLS with Broadband: Cradlepoint has replaced NTC's MPLS network with a hybrid WAN consisting of pooled wired and wireless broadband connections supporting both primary and failover requirements. The price for MPLS copper connections can range from $300 to $600 per Mbps/month, while business-class wired broadband access (cable or DSL) is typically less than $5 per Mbps/month—a savings of more than 90 percent. Given the price/performance advantage, a 25Mbps downstream and 10Mbps upstream wired broadband link is used for the small and medium NTC branches at $99/month a site. For the Large NTC branch, which has more than 100 users, a 100Mbps Metro Ethernet connection is used at a cost of $10 per Mbps/month.

LTE for Primary and Failover WAN: In addition to wired broadband, Cradlepoint has deployed Advanced 4G LTE as part of the hybrid WAN connection pool for both primary and failover uses. Cradlepoint's support of LTE Category 6 enables up to 300Mbps download and 50Mbps upload throughput speeds on networks such as Verizon's Advanced LTE. However, real-word speeds likely average 150/25Mbps.

Virtual Overlay Networks: Cradlepoint's SD-WAN capabilities include several forms of virtual overlay networks. For NTC's network, corporate intranet traffic flowing from the branch to the private cloud datacenter uses a DMVPN/GRE overlay. Internet-bound traffic headed to public cloud, SaaS applications, and the web utilizes a VCN overlay. While both overlay methods support point-to-point and meshed topologies, VCN has the added value of integrated DNS, LAN emulation, and AD integration.

Intelligent WAN Selection and Steering: Cradlepoint routers provide several policy-based traffic management and steering mechanisms that enable Quality of Service (QoS) and intelligent selection across both the wired and wireless broadband links that comprise the branch hybrid WAN connection pool. Using the mechanisms summarized below, specific QoS and traffic steering polices have been set to ensure the performance of NTC's business-critical SAP and PDM applications, control VoIP latency, and shape video traffic to avoid saturating links or over-consuming LTE data plans.

Intelligent WAN Selection & Steering: Cradlepoint provides several policy-based traffic management and steering mechanisms that enable quality of service (QoS) and intelligent selection across both the wired and wireless broadband links that compose a hybrid WAN connection pool.

- Policy-based QoS: App-enabled prioritization, bandwidth allocation, and traffic shaping for traffic traversing the router in each direction.
- WAN Diversity™: Ability to combine multiple wired and wireless WAN links into a hybrid WAN connection pool in primary and failover roles.
- WAN Affinity™: Traffic steering policies that control WAN link selection based on specific algorithms, including round-robin, load balancing, most available bandwidth, and LTE data usage.
- Intelligent LTE Failover: Complete policy control over the apps and traffic allowed to utilize the LTE link if one or more primary links fail.
- Data Plan Protection: Analytics-driven policies that automatically suspend or reduce LTE usage within the hybrid WAN connection pool as monthly data plan consumption reaches a set threshold.

High Availability: AER Series routers are configured with two LTE modems, each with dual SIM cards. This allows each router to be "dual-homed" on multiple LTE carriers in either redundant carrier (dual SIM) or concurrent carrier (dual modem) mode. With this approach, NTC can achieve 99.99% availability of its branch WAN. For the utmost in high availability and fault tolerance, NTC can deploy routers in tandem using VRRP to enable full hardware, WAN, and LTE carrier redundancy.

Guest WiFi: AER Series routers support advanced WiFi capabilities to enable secure guest access at each branch location. Guest network users can be micro-segmented from WiFi-to-WAN to isolate them from trusted branch networks. Moreover, the intelligent WAN selection and steering polices have been configured to ensure guest traffic does not interfere with business users and applications and does not utilize the LTE links. For added security and compliance, guest traffic also uses Secure Internet Access as described below.

Secure Internet Access: The new Cradlepoint SD-WAN achieves significant bandwidth savings for NTC by implementing Secure Internet Access (SIA) for branch employees and guest WiFi users. This direct access approach eliminates the backhauling of Internet traffic and avoids the cost and complexity of installing branch-based security appliances. Within each Edge router, the NetCloud Gateway provides an encrypted overlay to the nearest ServicePoint where traffic is securely routed through the NetCloud Engine SCG service to the Internet. NTC network admins can use NetCloud Manager to set the desired app, users, and device security policies, which in turn will automatically provision the appropriate virtual network services to be used for SIA traffic such as next-gen firewall (refer to NetCloud Services above for a listing of other available security functions).

Secure SaaS Access: SIA also provides branch employees with secure SaaS access from any device, including tablets and phones. NTC network admins can set user and device policies to allow or block the use of specific SaaS and web applications, such as Salesforce.com, Microsoft Office 365, or DropBox. For example, NTC may choose to allow access to all SaaS apps from any corporate-owned device but restrict access to only Salesforce.com for users of BYOD devices, like an Android tablet.

Public Cloud DR: The enclosed diagram illustrates how branch-level disaster recovery (DR) is provided using the NetCloud Engine service. At each branch, a separate disaster recovery VCN provides always-on connectivity to the public cloud DR site. In the event of a primary datacenter outage, or even loss of a single application or data store, the AER router can steer affected traffic over the DR-designated VCN.

MOBILE WORKFORCE WAN

Cradlepoint NetCloud extends the SD-WAN value proposition to NTC's mobile workforce, giving employees a secure, LAN-like connection to private and public cloud apps and files from anywhere and any device. As shown in the enclosed diagram, the NetCloud Client runs on each device and provides a persistent encrypted connection to a VCN overlay set up specifically for mobile access.

To address the security concerns around BYOD and public WiFi access, the mobile access VCN has been configured with NetCloud Services that provide NAC, micro-segmentation, next-gen firewall, and app control so that devices are isolated from one another and access is only granted to specific servers and applications at the datacenter and public cloud DR site. Mobile employees also are configured for SIA to provide secure and compliant access to SaaS applications and the web.

IoT WAN

Cradlepoint NetCloud and routers are optimized for IoT deployments in the field or within a branch. The ruggedized COR Series IoT router can support NTC's future field IoT deployments, such as kiosks, vehicles, digital signage, and surveillance cameras. It supports WiFi (for LAN or WAN), Ethernet and 4G LTE interfaces, DMVPN/GRE and VCN overlays, and the full suite of NetCloud Services. Within the branch, the AER3100 router with integral PoE can connect, protect, and power IoT devices such as cameras and sensors.

SINGLE-PANE-OF-GLASS MANAGEMENT

NetCloud Manager enables zero-touch branch and field deployments of AER Series and COR Series routers and utilizes a proprietary Stream management protocol that's 700 times more WAN-efficient than SNMP. Stream allows fine-grain management and control of routers, WAN interfaces, LTE carriers, and policies without the overhead of traditional management approaches, which can consume up to 30 percent of bandwidth.

Summary

With more than 15,000 customer deployments in some of the world's most demanding enterprise and IoT networks, and recognized leadership in 4G LTE solutions, Cradlepoint brings a unique pedigree to the SD-WAN market.

Cradlepoint NetCloud and router platforms provide a versatile SD-WAN solution that utilizes a single virtual overlay fabric to connect people, places and things, with advanced security and single-pane-of-glass management. For NTC, this all translates to a new software-defined and cloud-delivered WAN that makes its network more agile, secure, efficient, and extensible than ever before.

TO LEARN MORE, VISIT **CRADLEPOINT.COM**.

cradlepoint

NTC Tackles the Future with VeloCloud Cloud-Delivered SD-WAN: Fast. Agile. Secure.

Speed has become the currency of business, and security threats have multiplied and soared in sophistication. Internet-connected devices (IoT) are growing explosively and has been an integral part of NTC's network traffic growth. They have already investigated the cloud-based application industry trend and confirmed that NTC's IT strategy should leverage this direction to gain agility and cost savings.

NeedsToChange (NTC) is aggressively rolling out Office365 to boost productivity and enable mobile users—but it plays havoc with visibility into traditional traffic patterns. While NTC is still a reasonable-sized company—3000 employees across 50 sites—there is a distinct possibility that a near-term acquisition will double the number of sites and headcount, as well as add another data center.

IT staff is concerned that the current WAN architecture—static traffic management, distributed policy enforcement, fixed-bandwidth access links, MPLS contracts—cannot prepare NTC adequately to support digital business imperatives and general network upgrades at current funding and staffing levels. There is no time for a slow evolution of equipment and connectivity, and there is no funding or staffing for a replacement.

A Network for the Future

VeloCloud Cloud-Delivered SD-WAN, overlaying NTC's existing traditional WAN, immediately delivers many benefits and positions NTC favorably for growth, industry trends and the potential acquisition, while also leveraging existing WAN infrastructure investments. **Cloud-based applications are seamlessly integrated** and immediately rolled out with equal access to mobile and branch-site users. Traffic is routed via the **shortest path** to either the NTC data center or cloud-based applications. **Adding a broadband link per site** relieves bandwidth limitations, delivers the goal of direct Internet access, supports IoT, increases application performance, and improves branch uptime. **Security is simplified and strengthened** by inserting VNF firewalling and traffic inspection in each network site, and using VeloCloud built-in automated VPN technology. VeloCloud **cloud-delivered orchestration** provides network-wide dashboards, traffic and performance visibility, as well as centralized policy control.

A VeloCloud Cloud-Delivered SD-WAN for NTC

The illustration shows the VeloCloud overlay architecture. Transport-independence works across ***any*** combination of circuits that NTC deploys. Branch offices and data centers may be equipped with virtual or hardware-based VeloCloud platforms, replacing or augmenting legacy equipment.

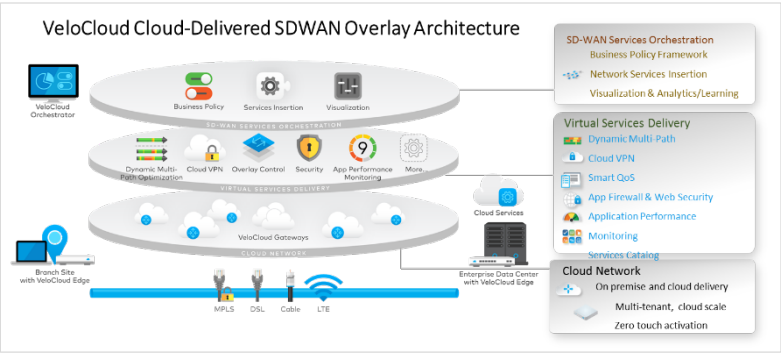

The network layer—consisting of VeloCloud Gateways, either on-premises or provider-based—enables connectivity to both enterprise data centers and IaaS/SaaS applications. A rich set of

virtual services, including those from ecosystems partners, are easily deployable from an application catalog. One essential service dynamically optimizes traffic over multiple links. At the top of the figure, the orchestration layer covers monitoring, configuration, policy coordination and unprecedented network visibility.

NTC Network Architecture: SD-WAN at Work

NTC's network architecture after implementing a VeloCloud Cloud-Delivered SD-WAN is shown below. Every site has broadband Internet and traditional MPLS connectivity. New sites do not require MPLS, and older sites may migrate to broadband when the MPLS contract expires, or both link types may co-exist indefinitely. NTC can make the most cost-effective decision per site. VeloCloud Dynamic Multi-path Optimization ensures a superior grade of service over any type of link.

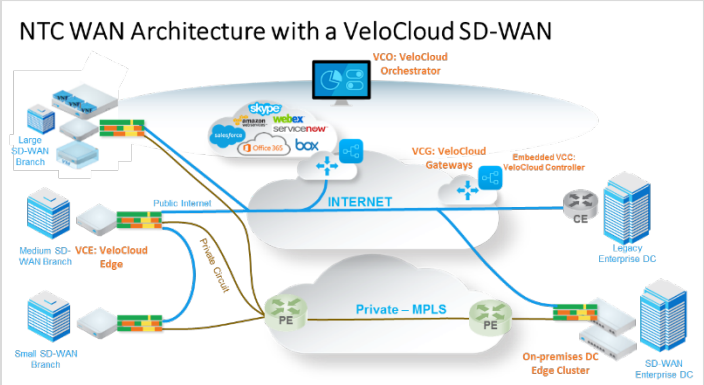

Most branch sites would be equipped with a hardware-based VeloCloud Edge platform (automatically deployed via the Zero-Touch capability), although larger sites that already support VM-hosting could be deployed with a VNF Edge. Data centers can be connected no-touch with legacy equipment, or with VeloCloud Hub Edges (virtual or hardware-based) when the time for site refresh is optimal.

Security services such as VNF firewalling are hosted in each site to secure the broadband connections. This eliminates the need to backhaul traffic to the data center—resulting in a better end-user experience and cost savings when the freed-up bandwidth is reused for PDM or other application traffic.

Essential Network Capabilities

Transport-Independent Branch Connectivity

VeloCloud's unique technology bundles traffic across multiple links and ensures enterprise-quality performance and security. A broadband link per site provides for NTC's desired Internet connectivity, relieving the backhauling of traffic through the data center, and also delivers critically-needed cost-effective bandwidth for IoT growth and convenient access for mobile users.

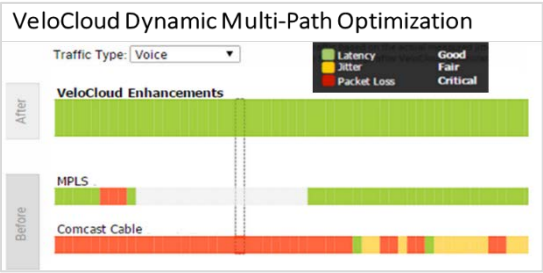

Links are auto-selected for traffic steering on a per-packet basis, based on (1) policy-driven priorities and business preferences, (2) auto-detection of application type, and (3) currently measured link performance. Mid-flow re-steering happens dynamically when changing link conditions are detected.

The illustration shows how VeloCloud technology achieves excellent performance across multiple links, each individually of lesser quality. The bundled multi-link connections (MPLS, broadband, LTE) provide increased branch uptime, headroom for IoT traffic, and elastic increases in cost-effective bandwidth.

Real-time Traffic—Superior QoS over Broadband

VeloCloud Dynamic Multi-path Optimization (DMPO) is a unique VeloCloud capability that assures application performance over any link types. A MOS score exceeding 3.6 (a good call) is maintained for 99% of calls with VeloCloud DMPO, while only 60% of calls achieve this without VeloCloud.

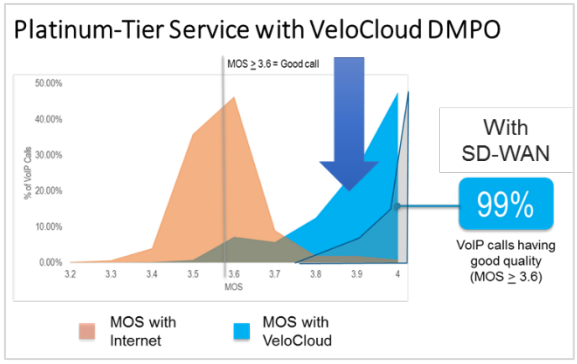

DMPO includes (1) *Continuous Monitoring and Analytics*: automatic capacity testing, link and path quality monitoring, tracking bandwidth, quality, packet loss, delay, and jitter; (2) *Dynamic Application Steering:* application-aware packet steering, aggregated bandwidth for single flows, maximizing throughput, sub-second reactions to network glitches ensuring no application impact; (3) *On-Demand Remediation*: error and jitter correction, automatic steering around brownouts/blackouts, link repair. DMPO maintains exceptional call quality and call success rates, and keeps video smooth and visible. IoT devices often reply on real-time-sensitive traffic that can be safely supported by deploying DMPO over broadband links.

Security

VPN deployment is significantly simplified with a VeloCloud SD-WAN. VPN any-site-to-any-site tunnels are automatically set up and secure all connections with strong PKI end-to-end encryption. Interoperable IPsec is supported directly to NTC's existing data center, and also connects to cloud-hosted data centers which may be a future NTC direction.

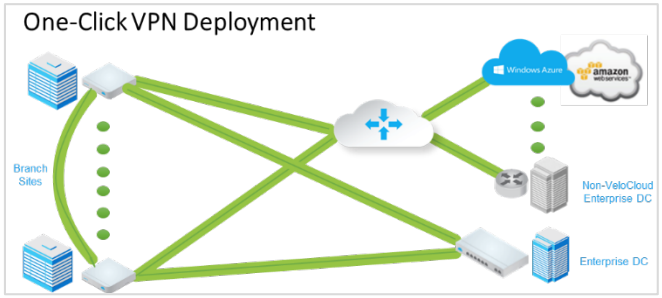

Scalability is achieved by eliminating the need for static hub-and-spoke VPN tunnels, and cost savings and simplicity derive from the automatic and dynamic set up of required tunnels.

Additional security services can be easily inserted with virtual instances of firewalls or other inspection engines. The cloud-delivered nature of a VeloCloud SD-WAN provides easy leveraging of cloud-based security providers offering sophisticated and cost-effective security with less demand on NTC IT staff.

Cloud Migration

Cloud applications can be accessed via the most direct path between the VeloCloud Edge and Gateway—a dynamic VPN tunnel ensures secure access with the Gateway dynamically bookending and aggregating connections on the cloud side. NTC can leverage VeloCloud or partner providers' multi-tenant Gateways already in place for IaaS and SaaS .

Orchestration

Cloud orchestration and controllers eliminate complexity and provide network-wide visibility and

control of policies and traffic patterns. Orchestration is hosted in the cloud or installed on-premises. As NTC commences rolling out SD-WAN capability, a hosted solution provides an easy entry point, while installing an on-premises solution may be preferable as deployment expands and the network hardens, especially if the expected acquisition takes place.

The VeloCloud orchestrator provides a consolidated dashboard for complete lifecycle management, including SLA measurements, remote diagnostics, link quality views, application analytics and bandwidth usage, network-wide business policy automation for traffic including voice and video, granular application and security policies, and multi-tenancy. The Zero-touch capability dramatically accelerates provisioning of new sites, while also significantly reducing IT staffing and costs.

Application Performance

Using the shortest path between the user and the accessed cloud site, in conjunction with VeloCloud Inbound Multi-source QoS and TCP Link Optimization capabilities, ensure optimal application performance. VeloCloud's Dynamic Application Steering and On-Demand Remediation capabilities allow sub-second packet steering around network problems to provide optimized performance for all traffic, and delivers a superior grade of service for real-time traffic to IaaS/SaaS sites as well as legacy data center applications.

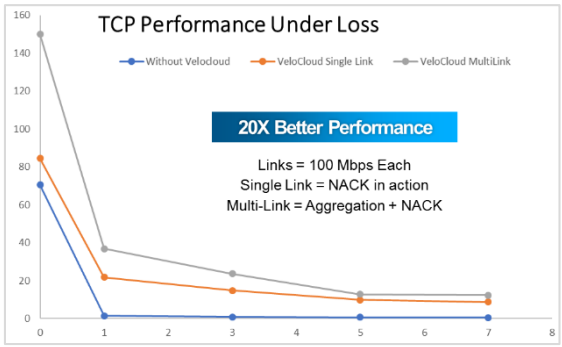

Conclusion

VeloCloud Cloud-Delivered SD-WAN offers an overlay network architecture allowing flexible, cost-effective incremental migration, and legacy interoperability. Cost reduction results from Zero-touch site deployment, multi-link bandwidth aggregation and optimization, technologies ensuring enterprise-grade service over broadband links, virtual and hardware options for Edge and Gateway services, and dramatic simplifications in cloud application access and VPN tunnel configurations. Security is always up-to-date with the options of virtual service insertion or using a cloud security provider. The transport independent connectivity (including MPLS, Broadband, LTE, and cable) allows for cost-effective link redundancy and increased branch uptime, agility in connectivity, ready access for mobile users, and economical incremental bandwidth for IoT.

VeloCloud SD-WAN customers have realized many benefits, including migrating successfully to 100% Internet links without network redesign, achieving enterprise-quality connectivity with capacity better than the prior MPLS service, reducing branch migration to thirty minutes, meeting PCI 3.0 compliance, significantly reducing OPEX costs with VNF and Zero-Touch deployment, maintaining exceptional voice quality with greatly improved call completion, one-click propagation of services and policies across the WAN, lowering bandwidth costs, and improving performance for real-time cloud applications

VeloCloud can help NTC address enterprise IT challenges with a fast, simple, secure Cloud-Delivered SD-WAN to achieve reduced complexity, easy migration to cloud applications and services, simplified traffic patterns, and distributed Internet access.

Planning for a Successful Transition to a New WAN

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 components[2]. For example, the team could include a SME from the organization's Network Operations Center (NOC). The role of that team member is to ensure that the NOC will be able to manage whatever technologies are eventually implemented.

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 3**. Also assume there are two viable alternative WAN designs, one from Vendor A and the other from Vendor B.

Table 3: Evaluating Vendors

Challenge	Weighting	Vendor A Scores	Vendor A Total	Vendor B Scores	Vendor B Total
Improving 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 3**, the team used a 10-point scale to evaluate how the two solutions responded to each of the WAN challenges[3]. The fourth column from the left demonstrates how the total score for vendor A was determined. The team gave Vendor A a 9 for improving 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 4** lists some of WAN characteristics that impact application delivery and identifies WAN optimization techniques that can mitigate the impact of those characteristics.

Table 4: Techniques to Improve Application Performance

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

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.

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.