# The 2015
# Guide to WAN Architecture & Design

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

Sponsors:






Produced by:

# Table of Contents

# Executive Summary

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

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

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

This e-book is part of a two part series that is focused on the WAN. This e-book describes a hypothetical company, referred to as NeedToChange, that has a traditional approach to WAN design. It then presents alternative scenarios directly from the e-book's sponsors that describe how NeedToChange should evolve its WAN. This e-book includes a summary of the key components of the alternative approaches to WAN architecture and design that were suggested by the sponsors and concludes with a call to action that outlines a project plan that network organizations can use to evolve their WAN.

The other e-book in the series is The 2015 State of the WAN Report. The goal of that e-book is to provide market research-based insight into the current state of the WAN. Towards that end, that e-book examines topics such as:

- What factors are driving change in the WAN?
- How are WAN budgets changing?
- How are IT organizations approaching WAN design?
- How receptive are organizations to new vendors of WAN functionality?
- What would drive or inhibit an organization from implementing a Software-Defined WAN?

# Introduction and Background

## Definition of WAN

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

## WAN Evolution

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

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

## WAN Services

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

**Table 1: Concerns with WAN Services**

| Concerns with MPLS | Concerns with the Internet |
|---|---|
| Cost | Security |
| Uptime | Uptime |
| Latency | Latency |
| Lead time to implement new circuits | Cost |
| Security | Packet loss |
| Lead time to increase capacity on existing circuits | Lead time to increase capacity on existing circuits |
| Packet loss | Lead time to implement new circuits |
| Jitter | Jitter |

## Traditional WAN Design

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

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

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

# Hypothetical Company: NeedToChange

Each of the sponsors of this e-book was given the description of a hypothetical company: NeedToChange. 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. The response guidelines that the vendors were given are contained in the appendix to this e-book.

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

Below is the description of NeedToChange's WAN that each sponsor received.

1. Data Centers

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

2. Traffic Prioritization

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

3. Business Critical Data Applications

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

4. Public Cloud Computing Services

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

5. Voice and Video

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

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

6. Internet Access

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

7. Remote Workers

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

8. Guest Workers

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

9. Branch Offices

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

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

10. Visibility

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

11. Regulations

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

12. Factors Driving Change

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

    - Improve application performance;
    - Reduce cost;
    - Increase uptime;
    - Reduce complexity;

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

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

## Sponsor Responses

Below is a description of how each of the 5 sponsors suggest that NeedToChange should evolve its WAN.

## Nuage Networks

nuagenetworks

# 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 data center servers. But now with Cloud IT, traffic patterns have changed. NTC has virtualized its Utah data center 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 data center. In disaster recovery situations, for example, it is relocated to a completely different data center. Unfortunately, outside the data center, 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, Data Center 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-data center/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 data center 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**

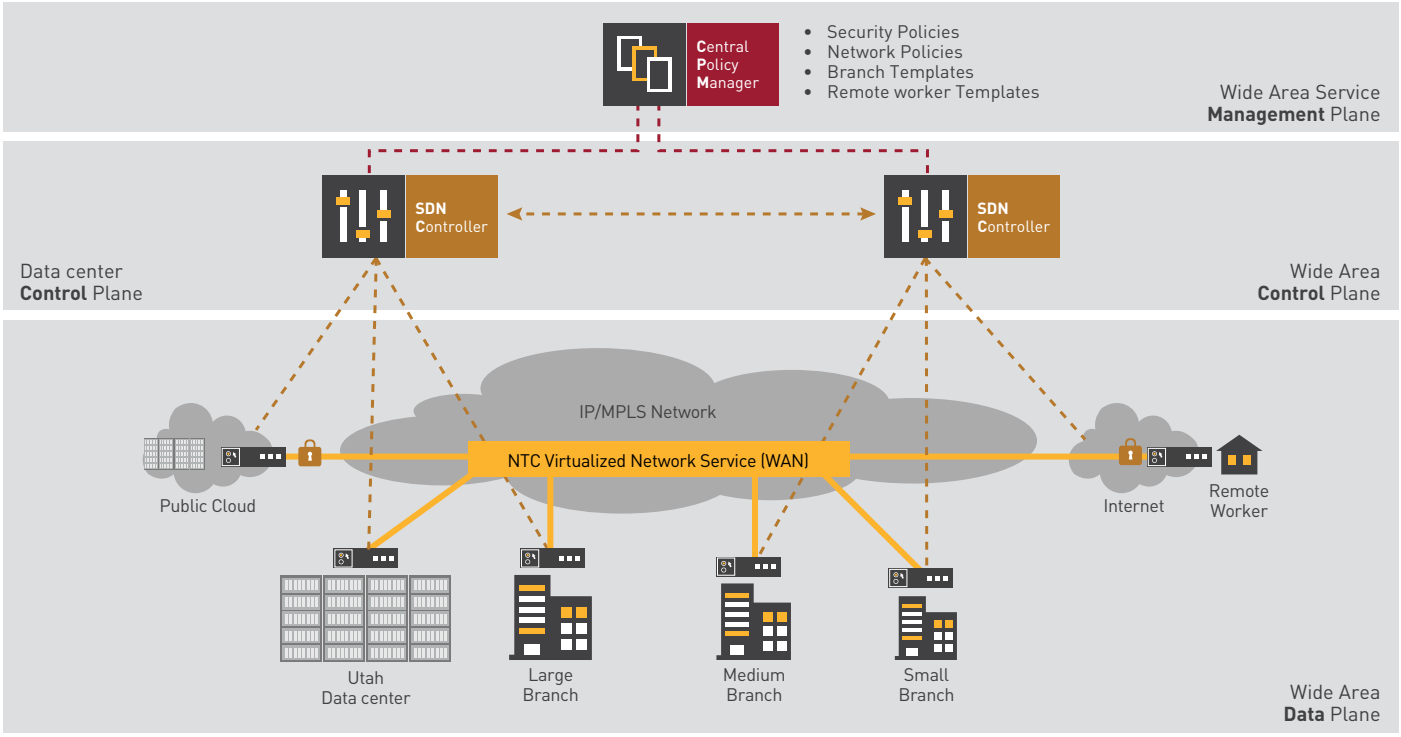


With Software Defined Networking (SDN) there are three key planes (or layers of network functionality) that will assist in this delivery (see figure 1):

- **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 SDN-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 data center and public cloud) of the NTC network.
- **WAN Data Plane:** Open compute (x86-based) branch equipment is deployed at the remote branch locations and data center 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 (SDN controllers).

## Any-to-Any Network Connections

NTC can implement a fully meshed network architecture to facilitate branch-to-data center 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 data center. This provides additional resiliency and enables NTC to improve network availability at the branch.

## Policy-Based Network Management

With the right SDN 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

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

**FIGURE 2. Using policies to intelligently offload traffic**

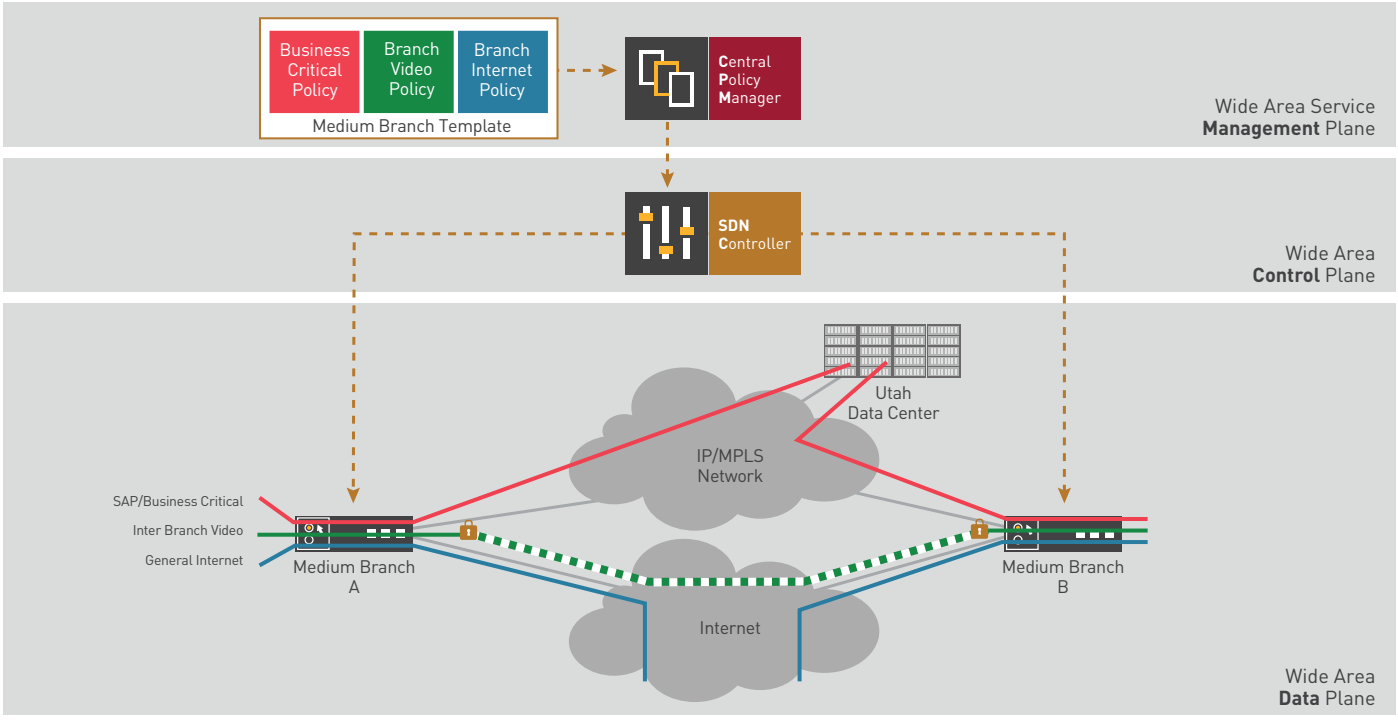


is relocated to the disaster recovery data center) or a new branch is added. These policies can be split into four key types:

- **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

SDN 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 SDN**

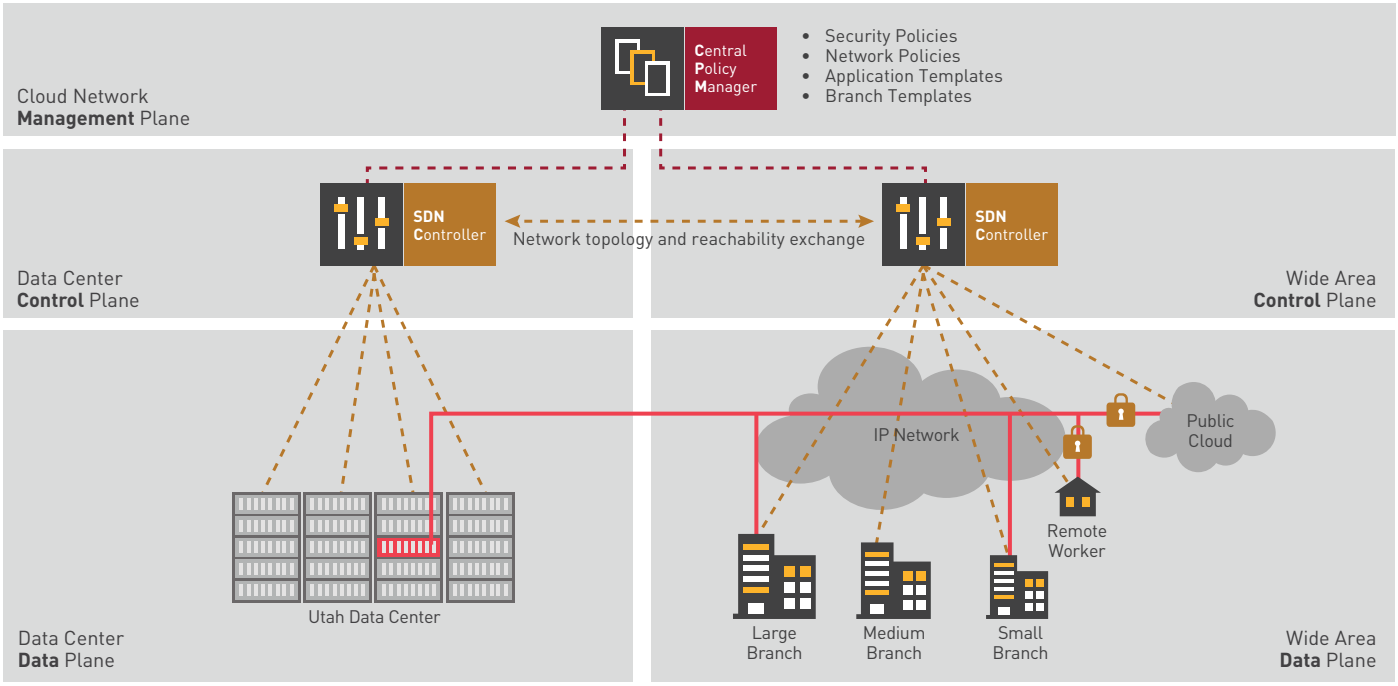


Comprehensive SDN-based WAN solutions use Network Functions Virtualization (NFV) to provide this enhanced functionality. Software features are "chained" into the traffic flows to and from the branches. By adopting this approach, NTC could enable a more robust and dynamic end-to-end policy that inserts 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 SDN 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 SDN, 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 data center and across the wide area 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. SDN provides the opportunity to achieve this. If SDN is controlling the network that underpins cloud applications and is managing the connectivity across the wide area 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 data center and beyond. To gain maximum benefit from the move to Cloud IT, NTC needs to centrally manage the data center 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

## Cisco

Cisco Systems Recommendations for NeedToChange:

# Modernizing the WAN for Mobility, Cloud, and IoT

## Introduction

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

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

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

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

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

## Steps to Modernizing the WAN

### Step 1: Migrate to a hybrid WAN overlay:

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

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

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

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

**Figure 1.** WAN Design for Private and Public Clouds

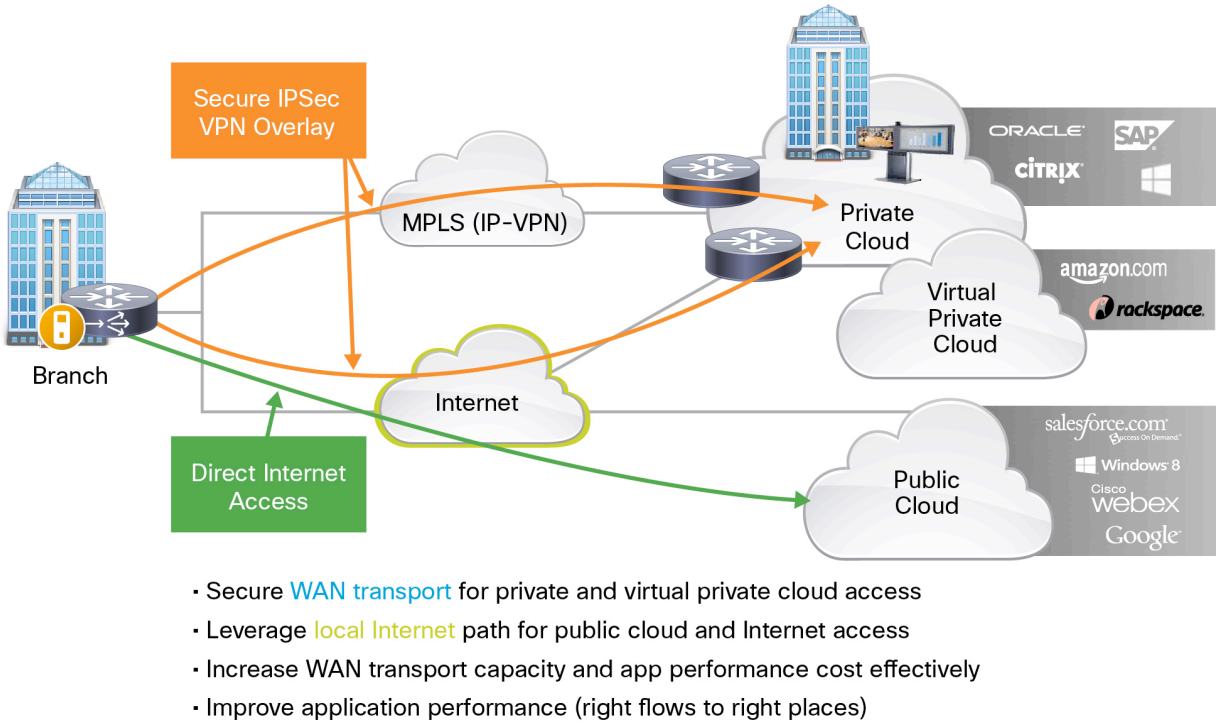


### Step 2: Protect and optimize application performance:

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

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

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

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

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

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

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

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

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

**Figure 2.** Latency and Bandwidth Impact on Page Load Time

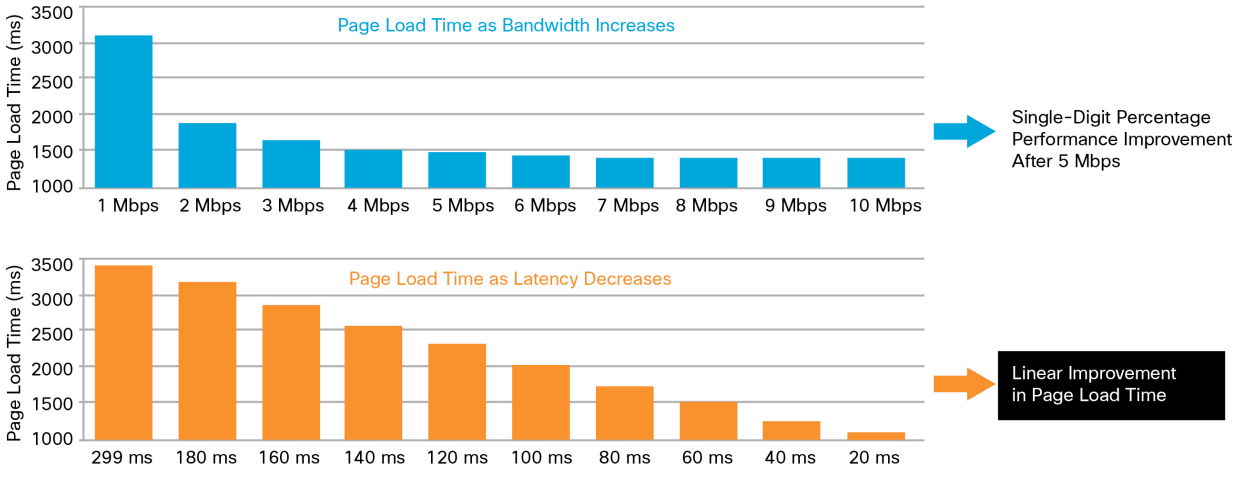


Based on Cisco and Akamai Lab Test Data

### Step 3: Enable a secure, scalable, and resilient Infrastructure:

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

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

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

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

**Step 4: Promote greater automation and orchestration:**
To promote greater agility, NeedToChange will require controlled-based architecture with open interfaces, and a software-defined networking (SDN) services plane that can abstract the device layer. This solution must automate and orchestrate WAN deployments in minutes with an intuitive browser-based GUI. A branch-office platform can be provisioned in just minutes without any knowledge of how to configure the devices (i.e., command-line interface or CLI). The application business priorities are translated by the controller into network policies using best practices and validated designs. The controller dramatically reduces the time required for configuration of advanced network services such as VPN, application visibility, path control, and QoS through simple, predefined work flows to deliver these services that align to business policies. The controller-based application offers an easily deployed solution that allows NeedToChange IT to get out of the complexity of managing low-level semantics such as VPN, QoS, and access list policies. Instead, NeedToChange IT can focus on the bigger picture: aligning network resources with the business priorities and delivering outstanding user experiences that result in better business outcomes.

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

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

## Summary

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

## Viptela

# A SOFTWARE-DEFINED WAN (SD-WAN) TRANSFORMATION

viptela

## Modified Enterprise Requirements

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

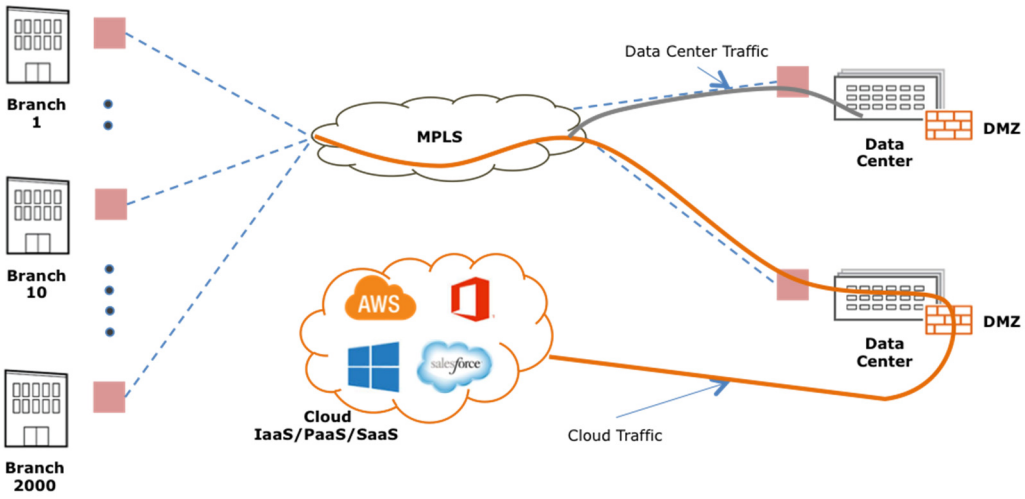


Figure 1: Present MPLS-based WAN

## Technology Requirements for Software-Defined WANs

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

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

## Network Transformation Steps

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

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

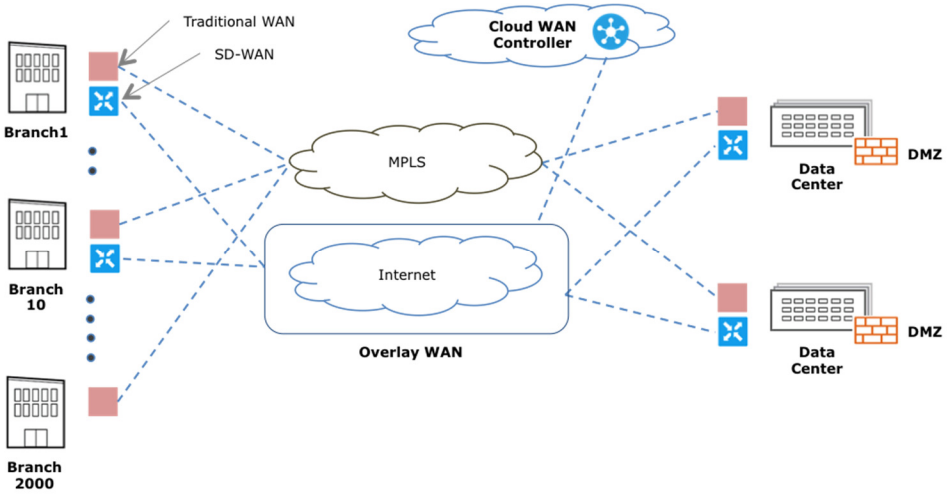


Figure 2: Insert SD-WAN at 10 sites over broadband or LTE

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

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

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

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

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

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

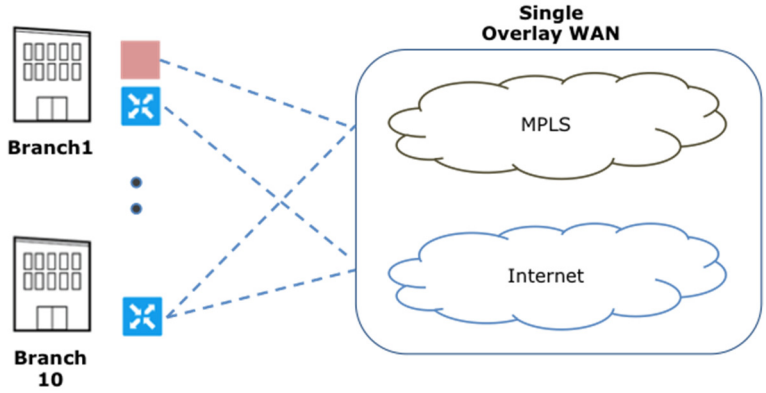


**Figure 2: Expand WAN Overlay to All Transport Links**

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

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

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

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

### Phase 4: Expand SD-WAN solution to all sites

*Goal: Achieve an enterprise-wide, single overlay WAN*

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

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

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

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

**Phase 6: Regional DMZs to optimize latencies of Cloud applications**

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

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

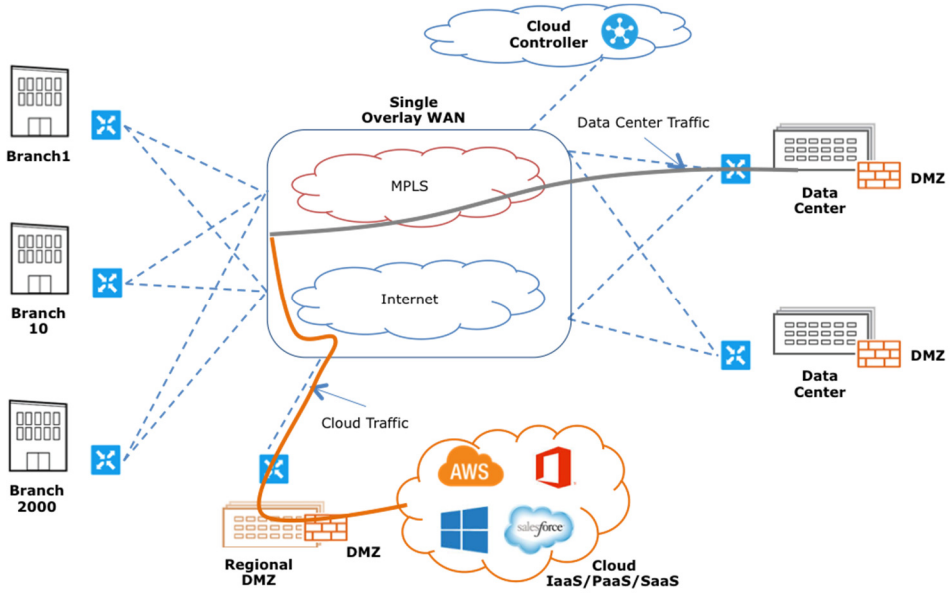


Figure 3: End state of a fully transformed SD-WAN network

## Summary

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

## Silver Peak

# NeedToChange Redesigns its WAN with Broadband using Silver Peak Unity

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

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

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

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

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

A Silver Peak proposed solution includes the following:

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

## Solution Overview

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

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

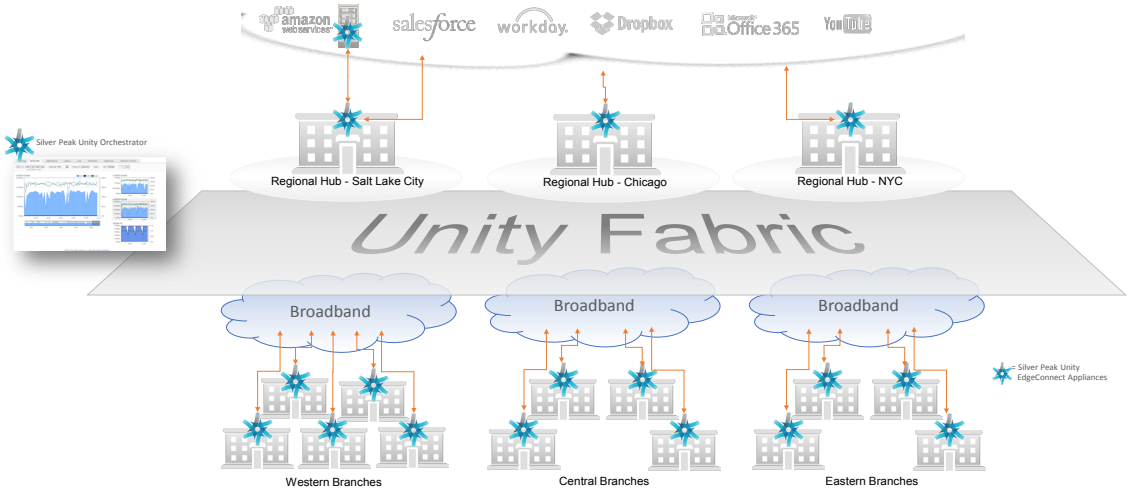


The Silver Peak Unity Fabric will be implemented using a phased approach. Initially, all sites will maintain their current bandwidth (MPLS and Internet) while the WAN fabric is deployed. For sites with multiple connections, Silver Peak's Dynamic Path Control (DPC) feature will be used to load-balance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

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

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

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

WAN Bandwidth Costs for all sites - Current (MPLS)

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

WAN Bandwidth Costs for all sites - with Silver Peak

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

Silver Peak Details

**Deployment and Management**

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

**Connectivity**

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

**Security**

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

**Visibility and Monitoring**

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

**WAN Optimization**

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

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

## Summary

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

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

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

## Silver Peak Benefit Highlights

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

## About Silver Peak

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

## Talari Networks

## NeedToChange's WAN Refresh

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 starting to experience slowdowns while accessing critical applications such as SAP and PDM, and complaints to the IT help desk were beginning to increase. There were complaints that calls were not connecting and the voice quality was garbled, particularly when the call volume was high. Worse, 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. And the outages that occurred were creating 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.

IT decided to address these problems and update their WAN to support the company's growth in the future, and, with the help of Talari's Software Defined THINKING WAN solution, were able to save money, dramatically improve their users' experience, improve overall productivity, and, best of all, stop their own team from lurching from one crisis to another.

## Moving to a Network with Talari

As NeedToChange began the WAN research process, two points were clear going in:

1. They would need more bandwidth at every branch office
2. They couldn't expand their MPLS commitment due to cost constraints

This meant that they would have to use Broadband Internet connections, as that was the only option for cost-effectively increasing available bandwidth and not increasing their MPLS commitment. The company quickly found a mixture of DSL and Cable providers and was able to get an extra connection the offices that didn't already have Internet links. Significantly, 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 VT500 for the small offices, and a mixture of Talari Appliance T510s and Talari Appliance T860Ls for the medium and large offices. In the Data Center they installed a high-availability pair of Talari Appliance T3010Ls.

They implemented in phases, starting with the Data Center and the offices that were reporting the most problems. The physical appliances were preconfigured with management IP addresses and the installation at the medium branch offices consisted of simply inserting the appliance between the LAN switch and the MPLS and DSL or Cable modem. In most cases this work was done by non-technical employees in the office and no one from IT had to travel to those offices. Once an appliance was available on the network, the Talari Network Control Node (NCN) pushed the appropriate configuration and policies to allow that appliance to participate in the SD-WAN.

With the introduction of the Talari SD-WAN solution, the company decided to stop backhauling Internet traffic. Now, all cloud and Internet traffic would be directed from each office to the Internet via a Talari Cloud Gateway CT800 located in the closest Amazon Web Services (AWS) region to that office. SaaS firewalls in those locations would replace the firewalls in the large offices, simplifying their network and eliminating the ongoing maintenance costs of the on-site legacy firewalls.

## The New Talari Software Defined THINKING WAN

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

To make path decisions, the Talari solution collects 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, is combined with the centrally defined policies regarding prioritization, bandwidth share and security to make intelligent decisions about individual applications. The WAN becomes a thinking network, able to accomplish the goals of the organization in the context of the actual network.

## Results for NeedToChange

The new THINKING 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 seen with the new WAN were:

### Cost savings coupled with greater capactiy

Upon installation of the Talari WAN, NeedToChange was able to eliminate two large ongoing expenses. One was the maintenance on the firewalls that had been in place at the large offices. And two was the elimination 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 more than able to make the Broadband Internet links enterprise quality.

Because of the success with eliminating the one MPLS connection, the company decided to consider eliminating MPLS at other locations, starting with their international sites as the MPLS costs at those consumed a large part of the telecom budget. They also decided to forgo MPLS at new locations in favor of two Broadband Internet links instead. These decisions, coupled with the fact that the Broadband Internet links provided more than enough bandwidth to accommodate future bandwidth demand, allowed them to significantly lower their cost for telecom services.

### Business continuity where outages go unnoticed

By monitoring every path through the WAN, including to and from the cloud, the Talari solution was able to detect link outages within a fraction of a second and 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.

As VoIP was a key and growing application and voice traffic is very latency and packet loss sensitive and required relatively low bandwidth, voice packets were duplicated across two diverse paths. This 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.

### Improved application performance guarantees a productive workforce

Directing direct traffic away from poor quality links and duplicating voice packets went a long way toward improving the performance of business applications. 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 was able to generally 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 virtual 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 a particular office and the prioritization policies would immediately come into play if that occurred, further preventing outages from impacting application availability and end user productivity

### Security protocols uninterrupted and reinforced

The new Talari SD-WAN was able to conform with the stringent security policies in place. The Talari solution does not store any packets, can work with source encrypted packets, and doesn't provide external access to packets, so the security protections that were already in place to achieve PCI compliance were able to be left intact. All data was also encrypted by the Talari solution prior to being sent across a WAN link. The Company chose 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.

With many options available for customizing security settings, the company's IT staff appreciated the fact that security was configured centrally and that changing security policies could be done from one location and then quickly pushed to the entire network with ease.

### 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 Cloud Appliance CT800 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 and eliminated any application performance problems caused by failed connections or poor quality Internet links.

By using the dynamic conduit feature of the Talari SD-WAN, 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.

### 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 were able to see the performance of individual links and the aggregated performance of their telecom providers. This information was invaluable to help them obtain support from the provider and to negotiate better rates!

The correlated information on applications available from the management interface also helped the IT staff ensure they were meeting application SLAs. By running reports on application performance across the WAN, they were able to show each business unit the quality scoring 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.

### Ongoing management is click-easy

One of the best results of the new Talari THINKING WAN was the ease of maintenance. Policies were centralized and changes could be made in one location and easily pushed through the network, even outside of maintenance windows. New applications automatically used default behaviors and then were customized as needed. And the increased visibility into network and application performance made it easy to identify areas for improvement.

## The End Result

With their new THINKING WAN installed, the IT department found themselves able to respond more quickly and positively to application requests from the business units. They even had time to think proactively about new ideas that could help the company grow. Confident that their WAN was up to the challenge, they were able to add more video communications options, expand their use of SaaS applications, and push much of the needed infrastructure growth to the cloud. And they had budget to invest in these ideas with the decrease in telecom spend.

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 such as SAP and PDM was always available and high quality. And the business could move quickly when opening new offices or acquiring new businesses, knowing that they had a thinking WAN to support them.

# Key WAN Architecture and Design Considerations

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

1. Location of key WAN functionality

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

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

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

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

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

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

2. The Use of Dynamic Multi-Pathing

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

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

3. The Use of Policy

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

4. Network Topologies

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

5. Support for Real-Time Applications

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

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

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

6. Optimization

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

**Table 2: Techniques to Improve Application Performance**

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

7. Security

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

8. Automation

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

9. Visibility

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

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

10. Customer Premise Equipment

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

# Call to Action

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

## Create an Effective Project Team

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

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

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

## Establish an Ongoing Dialogue with Senior Management

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

## Identify the WAN Challenges

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

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

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

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

| Table 3: Sample Weighting | | | |
|---|---|---|---|
| **Challenge** | **Team Member A** | **Team Member B** | **Average Weight** |
| Improving app performance | 25 | 55 | 40 |
| Increase availability | 25 | 25 | 25 |
| Reduce cost | 25 | 15 | 20 |
| Increase security | 25 | 5 | 15 |

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

## Agree on the Extent of the Analysis

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

1. Do a moderately detailed analysis of the solutions provided by their two incumbent vendors and by two other vendors to be chosen by the team.

2. Do a very detailed analysis of the solutions provided by all of the eight vendors that seem viable.

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

1. The loaded compensation (salary plus benefits) of each of the four project team members is $130,000 or roughly $2,500 per week.

2. Approach #1 will consume 10 weeks of work from each team member.

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

## Choose Vendors

As described above, the decisions that are made relative to the breadth and depth of the analysis of alternative solutions can have a dramatic impact on the amount of time and resources consumed by the process. That is just one of the reasons why the project team needs to choose potential vendors carefully. A reasonable strategy is to enter into a high level conversation with what the team determines to be a feasible set of vendors. If the content of

those conversations impresses the team, they can do a deeper analysis with a short list of vendors who they believe can best meet their needs. This approach balances off the desire to do a broad analysis of emerging solutions with the need to conserve IT resources.

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 5 key vendors.

## Rate Alternative Solutions

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

**Table 4: Evaluating Vendors**

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

As shown in **Table 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 app performance. That 9 was multiplied by the weight of that challenge (40) to arrive at a score of 360. That process was repeated for each challenge and the sum of the four scores (805) was determined. That process was also applied to Vendor B, whose total score of 730 is significantly lower than Vendor A's total score. If the scores were closer, it might be valuable to do a "what-if" analysis. For example, what-if reducing cost was weighted higher than 20? What-if Vendor B got an 8 for improving app performance?

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

[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".

## Manage existing contracts

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

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

## Build a business case

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

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

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

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

# Appendix

The sponsors were given the following guidelines to shape their response to how NeedToChange should evolve its WAN.

- Include in your response an architectural discussion of what functionality is needed and a discussion of where it should be located (i.e., in the cloud, on the customer premise, in a co-location facility), whether that functionality should/could be in hardware or software and any other important considerations. Also include in the description the key features that should be implemented.

- Highlight how the suggested architecture enables appropriate performance, as well as effective security and effective management, including visibility.

- Highlight how the suggested architecture responds to the factors that are driving NeedToChange to evolve its WAN.

- *(Optional)* Provide insight into what functionality is currently shipping and what isn't.

- *(Optional)* Provide insight into your vision for how branch office functionality can be provided.

- *(Optional)* Provide insight into the potential cost savings.

- Avoid making any references to specific products or services.

- Avoid marketing hyperbole (i.e., "This approach is the most scalable" or "The only way to solve this problem is…….")

- Avoid making negative comments either explicitly or implicitly about a competitor.

- Limit your response to 4 pages, not including a possible one page embellishment of the description of NeedToChange.

*To continue discussion of this e-book with your professional colleagues, check out the on-line version at Webtorials.*

## 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.