AT&T Business

Next Generation Networking Solution Dimensions

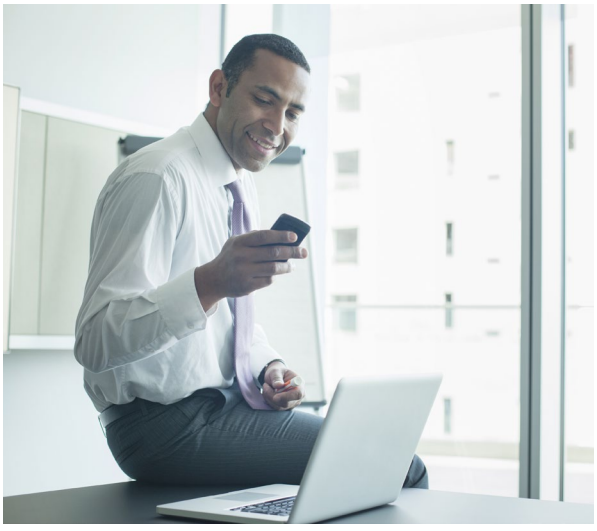

PART I
The four key NETWORK dimensions of Next Generation Networking (NGN)

Introduction

Welcome to this four-part series on advanced networking solutions. We'll take an in-depth look at the strengths and weaknesses of various networking components and processes and consider how they impact performance, security, and IT capabilities. Our goal is to provide your business with the guidance and insights needed to help you plan and make informed decisions about your network, operations, and internal business capabilities.

Gartner defines "Next Generation Network," or NGN, as "the evolution and migration of fixed and mobile network infrastructures from distinct, proprietary networks to converged networks based on IP."[1] NGNs represent a significant evolution towards greater reliability and security in corporate networks, as they incorporate the virtualized functions of Software-Defined Wide Area Networks (SD-WANs) and Secure Access Service Edge (SASE) to achieve more efficient, centrally-managed network operations. However, the benefits gained by NGN can increase complexity, and in this series we'll help guide your business towards adopting the right resources for your objectives.

1. Gartner Glossary, https://www.gartner.com/en/information-technology/glossary/ngn-next-generation-network

The series breaks the discussion of NGN into 11 key topics, or dimensions. These dimensions correspond to the considerations your business must weigh when evaluating a network transformation. The dimensions are divided into four categories, and each part of the series is dedicated to one of those four categories: Network (4 dimensions), Operations (3 dimensions), Customer (3 dimensions), and Planning and Management (1). Assessing and selecting a proper network infrastructure and the right components, and determining the best operational support model, are critical considerations in helping ensure that you obtain the greatest potential business benefits from NGN.

Key NGN dimensions

Planning and Management

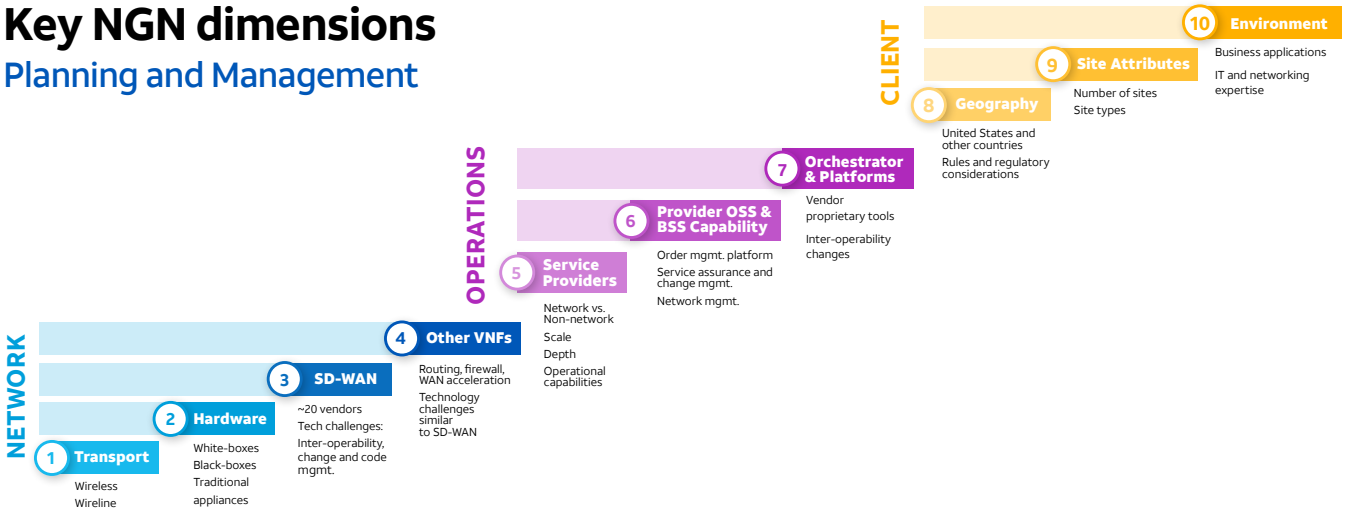

Before examining the four Network dimensions of NGN, let's first lay the groundwork by outlining some preliminary steps to approaching a network overhaul.

Initial considerations

While NGNs, particularly SD-WAN and SASE solutions, seem to represent the future of business networking, these networks come in multiple forms, each with its own strengths and weaknesses. As potential NGN customers, you need to establish your own hierarchy of priorities when designing your networks, as focusing on one aspect within an NGN solution will often require a trade-off against other aspects. Key aspects for you to consider are:

- Cost
- Performance
- Flexibility
- Stability
- Ease of Use
- Security

Your goal as an NGN business customer should be to find the best balance of these priorities that make the network work effectively for you. When considering the following dimensions, as a prospective implementer of NGNs, you should carefully consider the aspects that are most pertinent to your business and weigh them against potential trade-offs in other areas. These priorities should then be the lens through which all future considerations are filtered.

With these qualifications, AT&T Business offers our expertise to anyone seeking to implement Next Generation Networking for their business. We provide comprehensive assessments of NGNs with the understanding that you have other options beyond the AT&T solution set, and we recommend that these options should be similarly evaluated using the same key dimensions.

Network (Dimensions 1-4) overview

A Next Generation Network, as defined by AT&T Business within this paper, is comprised of four interrelated network dimensions acting in concert:

1. Physical transport
2. Customer-premises hardware
3. SD-WAN platforms
4. Virtual Network Functions (VNFs)

DIMENSION 1: Transport

For the dimension of transport, we are not focusing on the most recent advancements in wireless (5G) or wireline transport (100 GbE) — these are merely examples of transport types that Next Generation Networking (NGN) are able to leverage as they become available. Rather, NGNs now enable combinations of existing transport types in ways not previously available. They use virtualized circuits in a centralized control infrastructure to provide dynamic routing over multiple distinct transport paths. Logically, optimization of network performance will be contingent upon having multiple transport paths available; otherwise, dynamic traffic routing can only select a single path, effectively negating some of the intrinsic value of NGN capabilities.

When making decisions about transport, some considerations mirror traditional Multiprotocol Label Switching (MPLS) networking decisions required in the past:

- *What level of reserved bandwidth is required at each location?*
- *What security considerations must be addressed?*
- *What bandwidth growth projections should be considered?*
- *What type of access diversity (path, central office) is required?*

Because NGNs can now use new transport combinations, there are consequently new decision points for you to address:

- *Do we leverage wireless as primary or secondary transport (or both)?*
- *Should we maintain MPLS or VPN network connectivity?*
- *Do we use Dedicated Internet or Broadband (or both)?*
- *What type of Transport combinations will or will not be supported by SD-WAN platforms?*

Wireless transport

For many site types, the convenience, ease of ordering, ubiquitous nature, and generally reliable connectivity of wireless technologies are logical reasons to incorporate wireless transport into an overall NextGen solution. However, there are concerns and potential drawbacks to be considered due to performance unpredictability.

Reliability/stability

While a wireless coverage map may indicate good Long Term Evolution (LTE) coverage for a site, these coverage estimates are often overly optimistic about the actual signal strength being received by the network devices. Internal building structures/ materials and other equipment inside the site can adversely affect LTE signal strength, and the proper signal measurements should be taken from within the premises prior to provisioning LTE connectivity. Some poor internal signal strengths can be overcome by strategic placement of extensible antennas on accessible rooftops or at windows, or by using signal boosters, but these contingencies should be planned for and tested prior to provisioning the LTE.

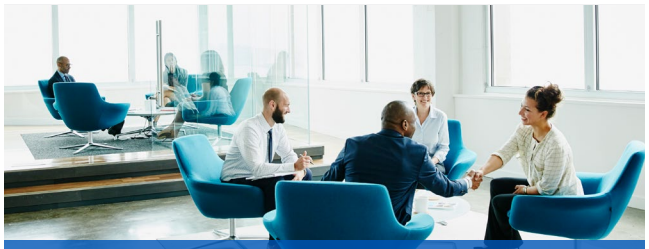

EXAMPLE #1
Real customer | Use case

Situation: Customer provisioned Long Term Evolution (LTE) as a backup wide area network (WAN) transport path without obtaining reliable signal strength data on interior locations prior to installation — testing externally to the building only.

Interior signal strength was unable to support and maintain the data transport path, rendering the requested backup WAN transport unavailable.

Learning: Ordered a subscriber identity module (SIM) from an alternative LTE provider that featured tested and confirmed sufficient signal strength into the interior location and performed secondary installation when these parts became available on-site.

Also, while wireless coverage maps may indicate a clear and stable signal for a given site, signals may be affected by external events, as they share radio access network (RAN) capacity from their assigned cell towers with other wireless components. Certain

large-scale events in close proximity, such as major business conferences or sporting events, may cause unexpected temporary shifts of the traffic distribution between RANs, while new/ongoing construction in the area may cause changes in the signal strength of a more permanent nature. Note that fixed wireless data links, such as LTE wide area network (WAN) circuits, are usually most susceptible to these capacity-driven interruptions. In addition, dynamic traffic control features of NGNs are dependent upon continuous monitoring and measurement of the transport circuits to determine their relative stability and reliability. While both wireline and wireless connections may consistently measure out to be stable at their given sites, useful stability calculations for these wireless links are limited to periods when traffic is carried across the link — intervals which are not continuous.

Remote monitoring and management

Wireless access is monitored from the equipment installed at your premises. Typical network monitoring parameters, such as latency and jitter, are neither proactively nor continuously analyzed — however, simple network monitoring parameters such as port Up/Down status are available. Additional monitoring and management capabilities are frequently made available, but these services tend to come at an incremental cost.

Additionally, wireless network capacity planning and deployment is typically evaluated based on the overall network health, rather than on individual device bandwidths or performance. In this regard, the wireless network is comparable to broadband access whereby the available bandwidth is competed for by all customers leveraging the same network capacity in the same geographic area. Performance may vary by time of day, and few performance guarantees should thus be expected.

Data security and integrity

Wireline access provides a more secure and consistent means of transmitting business data than does wireless access, particularly when using private networks for the most critical traffic in the NGN. Wireless is generally considered less secure than wireline, because data is broadcast over radio waves, which can be easily intercepted by anyone within range of the transmitting station. Modern advancements in encryption and use of Virtual Private Network (VPNs) within wireless transmissions can reduce some of the deleterious effects of potential hackers, but this reduction is never to a zero-risk state.

Most LTE wireless links are configured with very limited mechanisms for establishing Quality-of-Service (QoS) priorities for transmitted traffic, which can affect data integrity. Traffic to and from the FirstNet network — the nationwide, high-speed broadband communications platform dedicated to and purpose-built for America's first responders and the extended public safety community — is afforded better prioritization than non-FirstNet traffic, using the AT&T Dynamic Traffic Management (ADTM) protocols for QoS, Priority, and Preemption (QPP), but this unique network is a rare exception.

Also, consider that wireless communications may be subject to radio interference or jamming, whether accidental or deliberate, and this can be a cause for concern regarding data integrity on any wireless link. In addition to these expressed concerns for Wireless WAN links, Wireless LAN technologies such as Wi-Fi hotspots can further reduce overall data security while being subject to the same interference and jamming.

Cost/billing

Large data use over wireless connectivity, driven by sustained or high-volume traffic flows, can produce higher-than-expected charges. Most 4G LTE connections are billed on a usage-based model rather than upon a capacity-based billing model.

EXAMPLE #2
Real customer | Use case

Situation: Customer used an LTE connection to regularly update devices during off-peak hours.

Excessive wireless traffic caused an extremely large wireless bill the next month.

Learning: Clearly understanding network routing, data usage, and cost implications is critical when planning to leverage wireless connectivity.

Implementation considerations

It's important to have a secure process for controlling implementation of LTE-enabled devices and for preventing the unexpected deactivation of installed SIMs. If an activated SIM is installed into a device which is not successfully turned up onto the network within a short interval (usually about three business days) of activation, it may be automatically deactivated, requiring additional work to reactivate it.

Proper management of SIM activation and deactivation are critical in preventing fraudulent misuse of activated SIMs. There are two distinct cases for concern. First, a device shipped to a site may become lost or stolen, requiring its installed SIM to then be deactivated quickly to prevent fraudulent usage charges from impacting the account (note: just the SIM may itself become lost or stolen, with the same consequences). Second, an SD-WAN device that is configured for LTE-connectivity as WAN transport can be turned up anywhere that wireless signals can reach, even if the site is not the intended destination.

Assuming that the site is designed with a combination of wireline and wireless access, it's important to consider platforms that will allow the device to be shipped without the full wireless access port configuration, allowing that to be added by the Orchestrator. If possible, the SD-WAN platform should also be capable of verifying the path to the customer premises through the Edge devices from the Orchestrator. These precautions limit opportunities for installation into an unapproved site.

Wireless is evolving rapidly with the broad deployment of 5G, and the importance of wireless as local access transport will continue to grow. LTE Wireless is certainly a viable alternative transport for emergency backup or for handling infrequent traffic bursts; when coupled with wireline transport, it can also provide very desirable local access diversity. Also, dual LTE WAN access is becoming more practical due to its ability to accommodate fast deployment, and AT&T is working closely with its customers to resolve the various signal-strength and rate-plan-cost issues. Finally, wireless may be the best (or only) option in some remote locations where wireline connectivity may be cost-prohibitive or simply impractical.

Wireline transport

Even when used in conjunction with wireless transport, wireline connectivity is usually the cornerstone for an effective NGN. Ethernet-based wireline transport has become the clear standard for wireline network elements, supplanting the older Time Division Multiplexing technologies (such as T1 and T3 circuits). The following wireline transport types are frequently leveraged for NGN WANs:

Internet network connectivity (such as AT&T Dedicated Internet)

Switched ethernet networks (such as AT&T Switched Ethernet)

Broadband (such as AT&T U-verse)

MPLS-based networks (such as AT&T Virtual Private Network)

Dedicated/private vs. shared network connectivity (broadband)

Viewed solely from a cost perspective, broadband internet circuits demonstrate a clear advantage over private wide-area networks, but other factors must also be considered:

Security: The wide-open nature of internet traffic demands additional provisions be made to provide data security for your business traffic, usually in the form of a local firewall device or a cloud-based firewall service (or both). This adds a cost to the broadband-based transport that may not be immediately obvious.

Reliability: Broadband circuits traverse largely unregulated networks. Consequently, these links are inherently less stable and can be less reliable than alternatives. While dedicated and private networks provide robust Service Level Agreements (SLAs), broadband typically does not. If a local office loses connectivity on broadband, the customer should expect something like a "best effort" dispatch and restoration (often measured in days rather than hours), which is reflected in the lower provisioning cost of broadband.

Real time applications: Because multiple users essentially share and compete for bandwidth over broadband connections, enterprise-level performance and SLAs are generally not expected. This can significantly impact real-time applications such as voice and video conferencing.

VPN, dedicated internet, or switched ethernet

Virtual Private Network (VPN) is an Internet Protocol (IP) solution that is usually enabled via Multiprotocol Label Switching (MPLS). It represents a meshed communications network with the flexibility of IP access and the security and reliability of Frame Relay/Asynchronous Transfer Mode (ATM). Customer sites that are part of the same VPN communicate with each other regardless of the types of ports at those sites, and customers can choose and apply various classes of service to prioritize traffic over the network. AT&T Virtual Private Network is such an enterprise-grade VPN service which can be proactively monitored and managed as part of an AT&T-delivered NGN solution.

Dedicated Internet differs from broadband in that your connection does not compete for bandwidth with other subscribers. Bandwidth is guaranteed, and SLAs are available to support real-time applications. AT&T Dedicated Internet is an enterprise-grade service of this type which is proactively monitored and managed as part of an AT&T-delivered NGN solution.

Ethernet is the predominant computer networking transport technology today. It's easy to use and is mostly used as a "plug and play" to connect multiple networks as if they were on the same LAN, whether AT&T Dedicated Internet, AT&T Switched Ethernet, or AT&T Virtual Private Network. For the most part, Ethernet provides high performance with high bandwidth and is a cost-effective solution. However, AT&T standalone Switched Ethernet service is not proactively monitored and managed at the local access level, in contrast to the AT&T Virtual Private Network and AT&T Dedicated Internet services.

3rd-party networks

A 3rd-party network is a network that is not provided and maintained by the Managed Solution Provider (MSP) selected to implement and support the NGN. A Letter of Agency (LOA) is required for an MSP to interact directly with the third-party network provider on behalf of the business. This type of coordination typically involves an incremental cost that should be considered.

One caveat: As previously noted for NGNs, optimizing performance is contingent upon having multiple WAN paths available; otherwise, the dynamic traffic control inherent in many NGNs will have only one path over which to route the traffic. Nonetheless, some business customers have decided to invest in SD-WAN solutions for other reasons (such as lower capital costs alone) and have elected to install SD-WAN networks with single-WAN access; in doing so, they have not reaped the full benefits of an NGN.

In summary, new technologies provide more options and greater flexibility in building out your business networks, but these newer technologies potentially have both positive and negative impacts on traffic that must be evaluated against your operational priorities. NGN platforms allow integration of multiple WAN access technologies, allowing you to explore and find the proper balance of these access types to best serve your purposes.

DIMENSION 2: Hardware

The selection of site-level hardware, existing and new, is a critical component of any Next Generation Networking (NGN) solution. This choice can limit the available platform and tooling options, and, if not chosen carefully, can be an impairment to hardware and software interoperability during deployment and in the future.

Site-level hardware at a customer premises can be classified into three main categories

1. **White boxes**
   Generic devices designed to support a variety of virtual network functions (VNFs) from multiple vendors and combinations of vendors. White boxes fitted with additional network optimizations are often referred to as "gray boxes."
2. **Black boxes**
   Vendor-specific devices that support core functions specified by the device's vendor (typically monolithic SD-WAN images) or generalized devices that will support core functions with other VNFs, possibly from other vendors.
3. **Traditional (physical) network appliances**
   Examples include: local area network (LAN) switches, access points, IP handsets, modems, physical routers, firewalls, and WAN accelerators.

White boxes and black boxes, referred to hereafter as uCPEs (universal Customer Premises Equipment), incorporate hypervisors, virtual machines, and internal switches within their software and firmware. They are inherently more complex than the traditional appliances they replace; consequently, they generally need more frequent updates than their physical counterparts. This greater complexity drives the need for comprehensive testing of NextGen hardware, both for stability and for interoperability with other network components.

When considering an NGN, you will need to weigh the following hardware decisions:

1. What network functions will be needed within your network?
2. What uCPEs support those required networking functions?
3. What uCPEs support the applications used to effectively run your business?
4. If uCPEs from a single platform vendor can't support all the requirements, can the unsupported requirements be supported instead by traditional appliances in addition to those uCPEs?
5. Will the chosen hardware at each site support the throughput needed to run your business at peak usage? Will it support your required throughput at peak usage assuming normal business growth over the next 1-2 years?
6. Has the chosen hardware been approved for use/homologated within all the countries that will be represented by your network?
   - Are there hardware or software import/export requirements that may prove problematic for any of the countries that will be represented by your network?
   - Will there be a need for you to maintain any existing physical network devices from an existing Managed Service? If so, will the devices be compatible with the NextGen hardware?
   - Will there be a need for you to maintain any existing physical network devices from existing LANs? If so, will the devices be compatible with the Next Geneneration hardware?
   - Will out-of-band access be available for all hardware at your premises that will be managed by the Managed Services provider?
7. How challenging will it be to install the hardware and associated software?
   - In general, physical devices are quicker and easier to install, and are easier to troubleshoot as issues arise.
   - Virtual devices typically require additional time to upload software at Test and Turn Up and are often more complex to troubleshoot than physical appliances.

8. What are the one-time and recurring costs of the hardware options?
   - One uCPE with multiple software licenses versus two or more traditional devices?
   - White box hardware/licensing costs versus vendor-specific hardware/licensing costs?
   - What are the costs to maintain the hardware?
     a. Annual maintenance contracts on one uCPE device versus two or more traditional devices

To support some modes of connectivity (such as LTE wireless WAN transport) or some applications (such as VoIP), traditional hardware appliances must be incorporated into what could otherwise be a fully virtualized set of customer-premise devices at one or more sites. Traditional external/peripheral hardware may also be required at sites to support Out-Of-Band (OOB) connectivity for remote site management. These "mixed/hybrid" networks may or may not share common management platforms or tooling, often making it more difficult to successfully integrate mixed virtual and physical network devices into a single solution.

The network designer needs to consider the viability of pre-staging the hardware prior to it being shipped to the applicable site. Pre-staged devices can be pre-tested for stability, and any pre-staged configurations will help to reduce the resources (including time) needed onsite/online when hardware is ready for actual installation into the network. However, this option may not be workable where shipping the hardware with pre-loaded software violates country-specific laws (such as U.S. export laws) or vendor-specific licensing requirements.

Generally, the biggest issue with any network hardware is lifecycle management of the devices: handling updates, incorporating updated version offerings, and managing end-of-service (EoS) and end-of-life (EoL) states. In physical devices of sufficient maturity, these updates are infrequent and are announced well enough in advance to allow straightforward management of the device lifecycle. Within NGN virtualized platforms, which are still relatively immature and for which the software is far more complex, these updates are both more frequent and more disruptive.

The following challenges have been noted:

1. **Virtual solutions are more complex, requiring a greater level of expertise than native appliance-based solutions.**
   If the solution will allow it, using physical appliances rather than uCPE/VNF implementations for some or all the network elements can produce a simpler and more manageable solution. In addition, single-VNF virtual solutions may reduce complexity and provide more manageable solutions. Of course, this simplification can often result in an increased cost for the hardware.
2. **Platform and VNF vendors are not taking responsibility for ensuring full interoperability.**
   While it may not always be feasible for vendors to ensure all of the possible permutations of options and devices, many vendors have abdicated such responsibilities in favor of testing a limited number of their "preferred" solutions. This causes responsibility for testing and approval of hardware/VNFs to fall more heavily on you or upon your Managed Services Providers (MSPs). The vendors ultimately benefit from this testing, but it can seriously degrade your experience when MSP testing delays successful implementation. Both customers and MSPs need to drive vendors to more completely certify VNF and platform interoperability rather than taking on these issues themselves — often within active production environments. This vendor certification will also serve to reduce your dependency upon the MSPs to intelligently design elements within your Next Generation Networks.
3. **A uCPE with multiple VNFs requires changes to troubleshooting approaches due to a less clear delineation of component boundaries.**
   As an example, in a site using a traditional firewall behind a WAN-facing router, there are clear test points where technicians can divide and isolate components to perform local diagnostics with a laptop or a probe. In a uCPE solution, technicians may not have access to tools and/or insertion points to be able to troubleshoot problems as easily.
   - An on-site technician can no longer use the observation of external status lights for troubleshooting within a uCPE, making problem resolution more complicated.
   - Teams of specialists that have supported traditional solutions when the diverse functions were clearly separated into distinct appliances [e.g., managed security services (MSS) security

experts for Firewalls or Network Engineers for Routers] become an impediment to effective integration when the various functions now reside within a single device — having technicians highly skilled in multiple VNF types now becomes a more critical requirement.

4. **The tradeoff for having greater internal flexibility for virtual devices is a reduction in external hardware flexibility** for those same devices. Physical appliances such as routers have generally provided multiple options for external interfaces, but uCPE external interfaces tend to be simpler and much more restrictive in the current environment. As an example, some uCPEs will provide options for fiber adapters on designated WAN ports but are unable to support similar fiber connectivity on designated LAN ports, while physical router appliances will generally support both options easily. Additionally, with the convergence toward Ethernet connectivity as the industry standard, most SD-WAN platform solutions are not currently able to support legacy transport options such as time-division multiplexing (TDM), Frame-Relay, and ATM.

Summary

The obvious draw for uCPEs is the lower capital costs and the lower expenses to maintain the virtual devices vis-à-vis physical devices, but there is a tradeoff within the ability to effectively troubleshoot problems that may arise. This is compounded by the fact that most NGN virtual hardware is a mixture of proprietary and industry-standard elements, and there is often no certification of interoperability between vendors of uCPEs and the associated software.

EXAMPLE #3
Real customer | Use case

Situation: Unique combination consisting of four components and four unique manufacturers: 1) LAN switch, 2) uCPE Hardware, 3) SD-WAN, 4) Virtual Firewall

When network routes were dropped in the production network, was this caused by a routing table within the Firewall VNF, timeout parameters within the uCPE, flapping from an external switch, or by a software bug in the SD-WAN VNF?

Learning: Investigation required multiple weeks of coordinated troubleshooting and testing across expert engineering teams at AT&T and vendors. The root cause was a timeout parameter within the uCPE, which is now clearly identified and defined within AT&T best practices.

DIMENSION 3: SD-WAN

While SD-WAN virtualized devices may present clear benefits to you by reducing capital and maintenance costs, SD-WAN's greatest benefit often lies in the integrated platforms available from major SD-WAN vendors designed to stabilize your business networks, centralize the management of those networks, and make them more efficient using virtualized components and physical devices. Currently, five major competitors in the SD-WAN market have been selected by AT&T Business for standard SD-WAN offerings: VMware, Cisco, Silver Peak, Palo Alto Networks, and Fortinet, but other vendors may receive similar certification from AT&T in the future.

Virtualized integrated platforms open possibilities for many new and exciting networking solutions. Security management can now be integrated directly into the SD-WAN platform by the platform vendor, reducing capital and operational expenses while making security management simpler. Virtualization allows SD-WAN platforms to be more transport-agnostic, allowing multiple types of WAN transports to operate in parallel.

EXAMPLE #4
Real customer | Use case

Situation: An SD-WAN Orchestrator experienced significant loss of performance and an outage due to a high volume of sites being installed concurrently for a large customer within a single change window from the same Orchestration server.

After the problem was noted by AT&T Business engineers, the SD-WAN vendor required multiple updates to their Orchestrator operating code to enable a higher peak throughput sufficient to allow a reasonable number of concurrent updates.

Learning: The customer network rollout schedule had to be revised and extended as the change widow outages/delays prevented full deployment of required policies during allotted intervals. Subsequent customers are evaluated by our Service Delivery teams to ensure that throughput considerations are taken into account for large rollouts.

As new SD-WAN trends and applications emerge, many vendors are enhancing their platform offerings. It has become necessary for major service providers such as AT&T to evaluate the suitability of these platforms for you within their Managed Services offerings. Since most businesses' core competency is not focused on designing, engineering, implementing, or maintaining complex networks, customers often must rely on the judgement and experience of trusted Managed Services Providers (MSPs) to assess the suitability of the various SD-WAN platforms with respect to their customer's business requirements.

Based on the experience and knowledge AT&T possesses in managing and maintaining business data and voice networks, we recommend you address the following considerations when making decisions about SD-WAN platforms:

1. Are the requirements and priorities of your business network compatible with the strengths and capabilities of the platform vendor?
2. Are there preferred data routing/forwarding schemes to be carried across from your existing network that will make one platform vendor more desirable than the platforms of its competitors?
3. Are the desired transport characteristics within your network supported by the platform vendor (e.g., 10G/Subrated 10G WAN access)?
4. Will the vendor platform scale to your desired number of deployed devices, now and in the future, assuming reasonable business growth?

5. If you desire some degree of co-management of your sites with your MSP:
   a. Will the platform support shared management capabilities between the you and the MSP?
   b. Will the platform provide secure access to the relevant components for both the MSP and you at appropriate levels?
   c. Is the shared management you desire limited to self-service management by you of routine changes such as updates in policies or speed adjustments?
   d. Will the platform enable logging of user actions to support auditing?
6. Does the Next Generation Networking (NGN) platform vendor provide the appropriate traffic monitoring/management statistics or reports that your business requires?
7. How mature is the vendor platform within global and regional SD-WAN markets?
   a. Does the vendor offer consistent and repeatable processes behind a well-developed toolset?
   b. Has the SD-WAN platform been thoroughly tested by any of the MSPs you are considering?
   c. Does the platform vendor provide certification of interoperability with other vendors?
   d. Does the vendor present a well-organized set of update/release management procedures to the platform, as distinct from "bug fix" releases?
   e. Is the platform stable enough that update releases are infrequent and relatively non-disruptive?
8. How does client-edge software licensing/deployment work within the global arena? Is licensing renewal straightforward and non-disruptive? Can licensing and license renewal be managed by the selected MSP?
9. Does the vendor provide sufficient levels of data integrity and security features to meet the needs of your business? Are the security features integrated into the platform offering, or do they exist only as a standalone feature that can be added?

Different SD-WAN vendors try to incorporate many of the same elements and features within their platforms, with varying degrees of success, but each vendor tends to place more focus on certain strengths they can leverage than on other aspects.

For example:

- VMware emphasizes its Dynamic Multi-Path Optimization (DMPO) applications to maximize reliability and resiliency and focuses on the flexibility afforded by its open-source VMware base.
- Cisco emphasizes its integrated security and analytics features, in addition to its familiar routing and switching capabilities.
- Silver Peak emphasizes its Unity Boost WAN optimization for accelerated network performance and its simplified Business Intent Overlay templates, in conjunction with the security of built-in zone-based firewalls.

As much as practical, the vendors' emphases should align with the priorities you have in designing your business networks. AT&T Business is in an advantageous position to evaluate the fit between your requirements and the relative strengths and weaknesses of the different SD-WAN platforms in order to suggest better solutions.

Although the concept of software-driven networks (SDN) is founded on open-source applications, to remain competitive, many NGN vendors rely on the development of proprietary protocols and applications to differentiate themselves within the marketplace. This generates some significant challenges when creating solutions that incorporate multiple vendors:

1. **Customer-premises equipment (CPE) and VNF vendors may not be adequately confirming interoperability with other vendors.** NGN platform vendors are often not ensuring that their products are certified to integrate with other products. MSPs such as AT&T are forced into taking on a disproportionate share of this responsibility in supporting deployments, including the funding of certification testing to ensure that the vendor's products operate as claimed. This situation is then compounded by the impact of mergers and acquisitions (M&As) within the NGN marketplace. This contributes to high rates of new code deployment and of additional offerings, which can be risky when you are striving for network stability. Given our experience and expertise, we are in a position to note both successful and unsuccessful instances of interoperability, and we make best efforts to ensure that the selected products integrate successfully.

1. **Lack of product maturity and stability is driving high rates of change that are then reinforced by the incompleteness of vendor interoperability testing.** In the current SD-WAN environment, as already noted, product changes are not only frequent but are often significant enough in scope to require commensurate changes in operating procedures and in tooling — and yet the change notifications to MSPs and to end users are typically delivered with truly short lead times. The product immaturity and the short lead times for update announcements make interoperability testing by the MSP even more critical (and problematic).

2. **Maintenance releases lack rigor in distinguishing bug fixes from new feature introductions.** Traditionally, equipment vendors would announce major code releases and then follow these up with patch or maintenance releases; when these maintenance releases are based on the prior code release and only address a specific bug fix, the risk of introducing a new issue is relatively low. However, because SD-WAN feature development is currently fast-paced, vendors are interweaving the introduction of new features with bug fixes, creating a higher risk for the inadvertent introduction of new issues. Hybrid bug-fix + new-feature-introduction coding can force AT&T Business and other MSPs into an expensive and time-consuming code progression/regression test cycle just to implement a single bug fix.

3. **In the current environment, frequent new code releases create a mixture of different code levels within the network, causing problems in synchronizing Orchestrator and data code levels.** Current SD-WAN platforms specifically require code in the Data plane (i.e., edge devices) and in the Orchestration plane (orchestrators) at a minimum, and there are strict co-requirements on the variance between these two planes. Some SD-WAN platforms (such as Cisco) will also install separate control planes and/or management planes in between these devices, rather than rolling these additional functions into the Orchestration function, but these added functions are usually less subject to frequent updates.

Currently, diligent SD-WAN vendors require that the Orchestrator code be upgraded in advance of the code within the edge devices. These vendors usually engage in recognized best practices that serve to limit the release variability between these codes by releasing the codes in feature sets, representing a combination of edge and orchestrator features. These feature set upgrades are then driven through the Orchestrator, first upgrading the Orchestration plane and then the Data plane.

Summary

It must be emphasized that the high rate of code releases, the high level of M&A activity within the SD-WAN marketplace, and the scarcity of adequate interoperability testing by vendors in ensuring stable and compatible network components has created an environment in which MSPs must bear the brunt of the responsibility for keeping your networks operational and stable. Such an environment inevitably leads to various levels of code distributed across large customer networks over time, compounding the difficulties faced by these MSPs in properly synchronizing new releases via the Orchestrator. AT&T Business continues to strive for greater accountability by SD-WAN vendors in providing the necessary structure to establish and maintain the stability and interoperability that your business expects.

DIMENSION 4: Virtual Networking Functions (VNFs)

In modern Next Generation Networking (NGN) implementations, one of the differentiating features of these solutions is the concept of networking function virtualization (NFV). NFV allows various functions, previously provided by traditional physical hardware, to be managed as virtual functions driven by software on "virtual machines." Implementation of networked components as software reduces hardware and energy costs, carries the potential to improve reliability, and increases efficiency in resource management and in the overall throughput of your traffic.

The original view of NFV was that virtualized capabilities should be implemented in centralized locations (i.e., data centers) only. This approach indeed works for many cases, particularly for cloud-based applications and platforms (such as Content Delivery Networks) or for infrastructure-based functions (such as mobile network nodes). Such an approach does not, however, work for all cases. The current view of NFV now emphasizes great flexibility in determining the physical location of the virtualized functions; virtualized functions should be located where they can be the most effective and the least expensive. The service provider should be free to locate NFV components in all possible locations as needed — at the data center, in network nodes, or at the customer premises. This approach, known as Distributed NFV, became emphasized as NFV was being standardized, and it remains valid now. For purposes of this document, this dimension will focus primarily on components implemented at the customer premises — at your sites.

Often confused with NFV is the concept of the software defined network (SDN), but these two features are completely separate and are actually complementary. SDN involves the decoupling of the control plane and the data forwarding plane so that the control plane can be centrally located while the data forwarding planes can be distributed amongst multiple locations. SDN introduces the concept of northbound (data-plane to control-plane) and southbound (control-plane to data-plane) communication paths. NFV, on the other hand, involves the bundling of virtual network function components (VNFCs) into usable VNFs, which are then distributed across a network function virtualization infrastructure (NFVI) with management and orchestration functions. SD-WAN, then, represents a deliberate combination of the NFV and SDN concepts.

VNFCs are typically modular virtualized functions that can be used within various VNFs to perform repeatable functions, such as internal switching connectivity between a "main" VNF and other components (VNFs or external) or for the termination of tunneled traffic from outside the VNF containers.

Even with stable and repeatable VNFCs, a single VNF may not provide any advantage over its physical counterpart or even be fully functional; the functionality required of VNFs may necessitate the interconnection, or chaining, of multiple VNFs. The use of VNFs is designed to encourage Service Chaining, or the chaining of multiple virtual functions within a single device, including:

Routing/Forwarding (vRouters)

Firewall services

Intrusion Detection System (IDS) devices

WAN Acceleration/Optimization

Load balancers

SIP Border Controllers (SBCs)

EXAMPLE #5
Real customer | Use case

Situation: SD-WAN VNF software was constructed by the vendor with a one-year expiring performance license by default, rather than allowing a standard three-year licensing agreement.

Learning: The license expiry resulted in all network traffic for the affected site entering slowdown mode until the license could be renewed/extended.

As noted earlier for SD-WAN platforms, VNFs produced by different vendors also often reflect their inherent strengths. Each type of VNF may be available from many vendors, and selecting the optimal vendor requires the systematic evaluation of those strengths

(and weaknesses) against the requirements of your business — which often considers previous familiarity with the vendor's products. If compatible, various functions from different vendors may be supported within a single device on a common management platform, providing greater operational efficiency for the management of the site. The number and types of VNFs required may determine the selection of hardware and SD-WAN platforms, and vice versa. When implementing a new solution on an NGN or migrating an existing solution to SD-WAN, the following questions should be pondered when considering VNFs:

1. What network functions will be needed within your network? Which of these functions can be virtualized within the selected MSPs supported platforms?
2. Are the strengths of the selected VNF vendors aligned with your requirements for the desired functions?
3. Are sites that require multiple network functions capable of chaining the required functions together and making them compatible with the rest of the NFVI?
4. Will the solution require multiple VNF and SD-WAN platform vendors within a single site? Will the solution require VNF vendor components within hardware for a different vendor (e.g., Palo Alto vFirewall within a Cisco device)?
5. Have the desired VNFs been tested for compatibility with other enchained VNFs or with the SD-WAN platform?
6. If the desired VNFs cannot be made to work within the NFVI, are there alternate vendors that may provide a workable solution?

VNF incompatibility, often driven by the use of proprietary elements within the VNFs, may cause VNFs from different vendors not to work together properly nor support all desired features. This risk of incompatibility can be aggravated as the individual VNFs are subjected to ongoing vendor code releases that may further erode their interoperability. If severe enough, this disparity could bring down all or part of an active network.

There are challenges in developing VNFs which can run seamlessly on an NFVI, and which can be easily on- boarded while remaining scalable and flexible. The reasons for these challenges include:

1. Lack of standard procedures across the vendor community to develop and benchmark VNFs.
2. Lack of architectural guidelines for VNFs across vendors.
3. Lack of standard (open) protocols for VNFs across vendors.
4. Lack of standardized configuration policies for VNFs across vendors.
5. MSPs have their own workflow infrastructures; VNFs must be developed to integrate into these infrastructures.

Due to such challenges, manual efforts are frequently required to configure, update, and test VNFs in new scenarios. This is a huge roadblock to MSPs in realizing NFV success. As a result, the same issues that have been already defined for hardware and for NGN platforms are equally prevalent within VNFs:

1. CPE and VNF vendors are not adequately confirming interoperability with other vendors.
2. Lack of product maturity and stability is driving high rates of change that are then reinforced by the lack of interoperability testing.
3. Maintenance releases lack rigor in distinguishing bug fixes from new feature introductions.
4. In the current environment, frequent new code releases create a mixture of different code levels within the network, causing problems in synchronizing Orchestrator and Data code levels.
5. Virtual solutions are more complex, requiring a greater level of expertise than native appliance-based solutions. A uCPE with multiple VNFs requires changes to troubleshooting approaches due to a less clear delineation of component boundaries.

In Part II of our series on Next Generation Networking (NGN), we will discuss the three dimensions related to Operations.

To learn more about how AT&T Business can help your business build a tailored network solution, visit www.att.com/networkservices **or call 866-415-0949.**

PART II
The three key OPERATIONS dimensions of Next Generation Networking (NGN)

Introduction

In part I of our series on the 11 dimensions of Next Generation Networks (NGNs), we covered the dimensions pertaining to Network. In this installment, we'll discuss three dimensions related to Operations. AT&T Business is here to provide expert insights to help you choose the right solutions for your business goals. Read on to learn how NGN solutions can affect the way your organization operates key aspects of your infrastructure.

While the four Network dimensions focused on the physical and virtual components in NGN, the three Operations dimensions focus on the aspects of the network operations that are required to successfully deploy, manage, and monitor your network. As previously noted, one of the defining characteristics of NextGen networks is the centralization of the control plane and its separation from the data/forwarding planes across a network. With the increase in variability and complexity, you need to carefully consider the operational requirements for implementing and maintaining an NGN.

Once the appropriate level of design has been determined for your network through consideration of the previously cited four dimensions, you need to evaluate the operational elements that will determine your ability to deploy and maintain the solution. There are three operational dimensions for your business to take into consideration:

- Service providers: network and managed
- Provider operational and business support systems
- Orchestrators and platforms

DIMENSION 5: Services Providers

Selecting the service provider is usually one of the first major decisions you'll make when building a network. This may be the most important decision. The capabilities and expertise available from the provider will determine which Software Defined Wide Area Network (SD-WAN) or Next Generation Networking (NGN) solutions can be supported as well as defining the level of support available during and after the network implementation.

Managed Service Providers (MSPs) are responsible for implementing, monitoring, and maintaining a customer's network infrastructure. These providers can be categorized into two groups: Network Service Providers and Non-Network Service providers. Network Service Providers bring the ability to integrate and monitor not only the SD-WAN components but also the network itself. Some providers, such as AT&T Business, also provide services for additional network functions such as security, Wi-Fi, and managed local area network (LAN). The benefits of a more inclusive provider are a more integrated solution, singular ownership, and more rapid fault isolation for transport-related issues. While there are more than 150 established Internet Service Providers (ISPs) which may seek to support your business NGN, many of these will offer project management and onsite assistance but are not end-to-end management providers. Even fewer of these ISPs can provide full end-to-end network management within a global context; many are strictly regional.

An effective MSP should be able to:

1. Scale up highly trained resources when challenges arise (and challenges will arise).
2. Test and approve new solutions as vendors provide new NGN components or update code levels within existing components.
3. Engage manufacturers and vendors, up to executive levels, when critical support is required.
4. Display expertise on a wide variety of networking components and architectures to accommodate the full range of solution requirements that may be required.
5. Upgrade and adapt your business network to accommodate growth and technology changes.

A more restricted level of expertise may suffice when only a limited set of solutions is to be pursued. However, the adoption of such a limited solution set promises to be an ongoing liability within a constantly evolving NGN universe. An MSP that can adapt to changing conditions within the industry will be more valuable to your business when compared to an

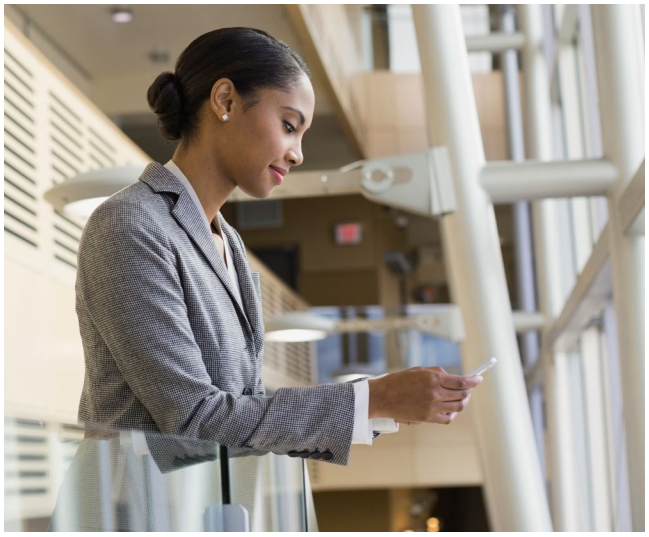

MSP that can provide only a fixed set of solutions, no matter how well-defined and inexpensive the fixed solution set may be.

In evaluating MSPs, there are two distinct areas to consider. The MSP needs to accommodate a solution that will match to your network requirements and to support that solution with monitoring, maintenance, and reporting once the solution has been successfully implemented. If you are looking to deploy an NGN, the following concerns should be addressed:

1. Can the MSP accommodate the network design priorities established for your business?
2. Can the MSP deploy an acceptable solution to all required locations globally?
3. Does the MSP possess sufficient expertise to anticipate issues or incompatibilities within the proposed design and provide workable alternatives?
4. Does the MSP have the resources to scale to the required number of sites on your network?
5. Can the MSP provide customization to accommodate non-standard requirements at one or more sites?
6. Does the MSP have the resources to directly engage vendors to resolve discovered issues?

7. As both business and network conditions change, does the MSP have the expertise to adapt your network to the new conditions?
8. Can the MSP test new combinations of network components or changes in the operating system coding if needed?
9. Can the MSP guarantee a minimum acceptable level of service for steady-state operations and for component failures once a site is implemented?
10. Can the MSP monitor your network to ensure compliance with established service level agreements (SLAs) or obligations?
11. Does the MSP have the necessary level of expertise to resolve issues when your network is out of compliance with service level agreements?
12. Can the MSP provide appropriate reports to give detailed snapshots of the status of your network?

Ideally, the MSP and your business need to exercise their due diligence in evaluating the suitability of any solution. Pre-sales and post-sales teams from the MSP need to ask questions to understand your expectations for your network. You need to carefully evaluate the claims of each potential MSP against your own expectations. Of the two sides of this relationship, your business unquestionably has the more difficult task, as you need to evaluate whether the MSP selected has the necessary levels of experience and knowledge to make the solutions work, rather than just a plan that looks and sounds good at face value.

Ultimately, the MSP is your first line of defense against anything that might go wrong, and the essential element that is needed in selecting an MSP is trust. OEMs (Original Equipment Manufacturers) are often not living up to expectations in ensuring inter-compatibility with other vendors, and the burden falls upon your MSP to test and resolve many issues. If an MSP is unwilling or unable to shoulder that responsibility, it quickly erodes your trust and places these burdens on you instead. This is true even in scenarios where you assume significant management responsibilities yourself — a frequent, and sometimes costly, occurrence within the current SD-WAN environment.

You need to evaluate MSPs against each other, and independent assessments by industry analysts provide a good start:

Do the industry experts consider the MSP an industry leader in NGN — and why?

Do the analysts consider the MSP as knowledgeable, both in the current modes of operation and in adapting to upcoming technological innovations?

Does the MSP have experience in SD-WAN over the course of the development of the modern SD-WAN environment or are they a latecomer to the industry?

Does the MSP form partnerships with equipment vendors that are themselves leaders in their fields?

Diligent research can help you build confidence and trust in the MSP you select and help ensure a positive experience.

EXAMPLE #1
Real customer | Use case

Situation: Service provider supporting the SD-WAN Orchestrator used a single set of user credentials for all employees accessing the device.

Significant security exposure that prevented identification of the actual user that created configuration issues within the Orchestrator and allowed all users the same high level of access regardless of experience or role.

Learning: Service provider migrated Orchestrator access to a server-based authentication, auditing, and logging (AAL) platform and assigned users individual roles with the privilege levels granted based upon the minimum access needed to perform their job. User IDs were uniquely assigned to individuals, with separate passwords for each user.

Independent market analysis provided by companies such as Gartner, IDC, and Frost & Sullivan provide insight into MSPs and their capabilities in deploying and supporting SD-WAN solutions. IDC in their Marketscape: Worldwide Managed SD-WAN 2020 Vendor Assessment identified AT&T Business as a leader in delivering SD-WAN capabilities and strategies. This report identified AT&T as capable of simplifying the operational experience, noting that AT&T is focused on enhancing the customer's experience, and identified the Global Services Organization as a key pillar for AT&T success in Managed SD-WAN. The same year, Frost & Sullivan's benchmark study, Frost Radar™: North American Managed SD-WAN Services Market 2020, stated that "AT&T leads the North American SD-WAN market and is a growth and innovation leader on the Frost Radar™." It noted the completeness of the AT&T offer, with the largest number of operational managed SD-WAN sites, and the Expert Engineering role that provides post-sales support, high touch, and deep network design verification.

Expert Engineering

AT&T Business NGN services, the Expert Engineering team, is comprised of highly experienced networking engineers specifically assembled to produce comprehensive E2E network designs and migration plans, to provide specific instructions for customer-requested changes, and to leverage our best practices and proven solution templates to ensure repeatable scalability for SD-WAN and large, complex FlexWare customers.

DIMENSION 6: Provider Operational Support Systems (OSS) and Business Support Systems (BSS) capability

Network providers provision and maintain customer business networks by leveraging a framework of tooling and automation that allows them to most effectively manage multiple network elements. This set of tooling is generally referred to as the Operational Support Systems (OSS). Additionally, providers require Business Support Systems (BSS) for effective order management, asset tracking, relationship management, fault and change management, and billing.

As network functions become increasingly software-defined, the line between BSS and OSS infrastructures continues to blur. Multiple BSS/OSS systems are typically required due to the innate complexity of Next Generation Networks (NGNs) and the fact that these networks are a mixture of disparate network products and vendors. Adding to this complexity is the fact that systems need to find a working balance between repeatable control of the managed components and the flexibility to adapt to new or changing customer requirements.

EXAMPLE #2
Real customer | Use case

Situation: Service provider was unable to scale resources to manage the volume of trouble tickets being generated to the provider's help desk from the platform vendor and properly inform the affected customer.

Extended customer outages were experienced.

Learning: Electronic bonding (E-bonding) was established between the vendor and the customer ticketing infrastructure to directly offload ticket notifications from the vendor.

As networking evolves, new challenges surface within the OSS and BSS infrastructures. Many providers that offer NGNs — SD-WAN in particular — have developed new systems and new automation to accommodate the processes of spinning up virtual devices and of managing the virtualized components. These new systems are not always fully compatible with the toolsets that may be used for managing the physical components of the affected network under previous offerings.

Additionally, the provisioning of a networking service, even at a simple branch site, is never just a single order. It is always a collection of interrelated orders to provision multiple transports, devices, and/or network functions. The maturity and capability of the service provider's systems and the expertise to effectively understand and coordinate the relationships amongst these products and services is critical.

Examples of potential OSS/BSS challenges include:

1. Changes to the network may require multiple distinct change orders, which is likely to increase the complexity of the coordination required to update existing sites.
2. Various network components may be managed and monitored by distinct systems and the provider must be prepared to account for issues such as trouble-ticket duplication, trouble-ticket correlation issues, and shared operational responsibilities for resolving component failures. Failure to accommodate these anomalies can increase the mean time to resolve issues caused by network component failures.
3. Cascaded connectivity, wherein multiple network components are serially connected to the same wide area network (WAN) link, can require workarounds to make each toolset recognize network components that may have been provisioned through a different toolset. This complexity may increase the time to provision new elements.

To be fair, by using an experienced service provider, most of these challenges will be completely invisible to the customer (and many of these challenges are the responsibility of the provider to resolve.) However, the customer should be aware that if these challenges are not managed appropriately, this can lead to customer frustration over these "unseen" issues.

DIMENSION 7: Orchestrators and Platforms

From an operational standpoint, some of the most challenging factors within current SD-WAN environments are the tooling systems and their associated automation. SD-WAN is often portrayed as an open-source environment leveraging white box hardware, but, in reality, most SD-WAN platforms use vendor-specific and often proprietary tools (orchestrators) to manage their devices. When engaging a Managed Services Provider to manage an SD-WAN solution, this situation is compounded by the additional toolsets needed by the provider to control the workflow and implementation of the various Next Generation Network (NGN) components.

Both vendor-based tooling systems (for the centralized setup and management of NGN components) and MSP tooling (for workflow management and component integration) must deal with a longstanding dilemma of system engineering: flexibility versus repeatability. In a rapidly changing environment, the flexibility to adapt to changes and incorporate new components into NGN solutions is often critical, and this requires equally flexible and adaptable tooling. On the other hand, if the solutions are going to be consistent and repeatable, the tooling needs rigid structures to ensure repeatability of processes within the solution. These two requirements are directly at odds with each other, and the synthesis of these two disparate elements into a workable toolset solution demands a careful balancing act.

Concerns with tooling cover three separate aspects:

- *Ability of vendor-provided tools to accommodate all features offered by that vendor*
- *Interoperability of vendor tools within a multi-vendor Next Generation Network*
- *Integration of one or more vendors into the provider's toolsets (such as the AT&T ECOMP platform)*

Ironically, the first aspect listed above is less of a problem for new vendors within the NGN environment than for the established competitors. The established competitors often have a fully developed environment that they are trying to adapt from their standard physical-appliance networks, while the relatively new companies are typically building SD-WAN/NGN solutions from scratch and have a less developed feature set upon which the tooling must act. There is therefore often a tradeoff between completeness of features and completeness of tooling. Of the three aspects listed above, this is easily the least critical.

The second aspect has been alluded to in prior dimensions, as interoperability of vendor components goes hand-in-hand with interoperability of the related tooling. Failure to achieve basic interoperability with components of other vendors often requires manually implemented workarounds (WAs) to drive the needed compatibility, and these WAs are typically not reflected in the vendor tooling. This is another result of the widespread failure of vendors to thoroughly test for compatibility with other vendors; the associated tooling is often unable to accommodate the resulting solution requirements without manual bypass of automated tools.

The third aspect is the most challenging, particularly when dealing with well-established providers. When the flexibility versus repeatability dilemma is brought back into the discussion, it should become apparent that the operational balances attained by the MSP and the balances attained by the vendors may be completely out of alignment with each other. Synchronizing the toolsets of the provider and the vendor to make them fully compatible may not always be possible, so even "simple" requests may have a custom element to them that requires additional investigation and research. In a constantly evolving environment, this lack of synchrony is even more pronounced, particularly when superimposed against inter-vendor incompatibilities.

As a result, unfortunately, some combinations of network elements and their associated tools cannot be directly supported through the business support systems and operational support systems of many providers because of proprietary components within the network (and other factors). As an example, many SD-WAN infrastructures for routing/forwarding may be incompatible with tooling requirements for VNFs (vFirewall, vWAN acceleration) from other vendors, and the MSP cannot possibly maintain adequate compatibility with both sets of tools simultaneously.

EXAMPLE #3
Real customer | Use case

Situation: The proposed firewall network function virtualization (NFV) software solution demanded accommodation of dual security zone features, but these capabilities were not available nor tested by the platform and VNF vendors.

A customized solution using SD-WAN hardware and firewall VNF policies required complex configurations that caused an extended outage of the customer production network as each site was installed.

Learning: The customer revised their design requirements to maintain the original hardware firewall solution.

Even more critical, as the centralized tooling for an SD-WAN platform rides upon its own specialized infrastructure elements, the presence of multiple orchestration infrastructures that must be synchronized with the provider tooling may make it impossible to mix and match SD-WAN platform vendors within a single network or within a single region on a network. Any solution that attempts to run multiple SD-WAN vendors on its NGN must certainly be considered a custom solution.

It is within this dimension that you need to start asking the tough questions that may not be otherwise answered by external research on your own. You may be able to identify vendors that have insufficient tooling and may be able to address some incompatibilities between vendors through independent research, but you will typically have no view into the MSP tooling systems without direct queries to each MSP:

1. Can the MSP tooling systems implement your requirements as a standard solution, or will it require custom processes/workflows?
2. Are there any known incompatibilities between hardware/SD-WAN vendors that the MSP must test and resolve for your design?
3. What are the impacts to your implementation timelines if a non-standard solution must be selected because of tooling limitations?
4. What are the impacts to your implementation timelines if the MSP must test and resolve vendor incompatibility issues?
5. Will customized solutions and incompatibility resolutions impact the ability to implement changes to the affected device in the future?

Because SD-WAN exists in a constantly evolving and changing environment, many associated tooling platforms are in a constant state of updates and fixes, requiring great diligence by the MSPs to avoid introducing new and unexpected bugs into the mix.

We at AT&T Business have years of experience with multiple best-in-breed SD-WAN vendors, giving AT&T Business Managed Solutions an advantage over less experienced MSPs in this crucial dimension. Some of the vendor incompatibilities and a significant number of customized processes have already been encountered and resolved by AT&T.

In Part III of our series on Next Generation Networking Solutions, we will cover three key customer considerations that can impact your NGN solution implementation.

To learn more about how AT&T Business can help your business build a tailored network solution, visit www.att.com/networkservices **or call 866-415-0949.**

PART III
The three key CUSTOMER dimensions of Next Generation Networking (NGN)

Introduction

In the previous installment of our 4-part series on Next Gen Networking (NGN) solutions, we covered three key operation factors affecting NGN. In part III, we'll outline three aspects of network deployment that require careful consideration by your network and IT operations teams. AT&T Business is committed to understanding the challenges facing your business, and we recognize the complexities that can arise during a transition to NGN solutions. In what follows, we'll provide a road map for navigating the customer components of this exciting terrain.

The location of sites within such a network is typically driven by the needs of your business rather than by the needs of the network itself. Unfortunately, some geographies present unique challenges, such as restrictive import and export laws for hardware and software; there may also be hidden pitfalls behind the arrangement of sites within the network, such as the need for hub/spoke global arrangements.

This installment will examine some of the subtle and not-so-subtle customer-level challenges that require customer awareness and frequently require action:

- Country location and government impacts
- Site attributes
- Customer environments and applications

DIMENSION 8: Country Location

For both network elements and operational considerations, the geographic location of each site and the distribution of those sites within the network globally are important factors in establishing a viable solution. It is important to consider each site through the lens of the local and national governments for the country in which it is located.

Importantly, some national governments will not permit specific types of transport, hardware components, or software components in-country. Additionally, some hardware and software components may fail to meet homologation (the granting of approval by an official authority) standards or national underwriting standards and may similarly be excluded within some geographies by the vendor. It is important that you work closely with your provider to identify all locations to be connected into your network and map those sites against these types of restrictions.

Import and export laws among countries may not only restrict availability of some components, but these laws may significantly affect the time intervals required to implement new sites. Laws in some countries impose restrictions that require in-country suppliers for hardware, software, and other resources. Similar requirements may force non-standard solutions to be developed and implemented from physical and operational perspectives.

Some governments may place limits on the ability to run specific applications or types of applications; the most prominent example of this is the severe regulation of customer data encryption within certain countries. Some countries may ban data encryption entirely, while others demand weak encryption standards or allow strong encryption only if

EXAMPLE #1
Real customer | Use case

Situation: Due to the remote location of edge sites, replacing hardware in the event of a failure would take an extended period of time. While software replacements can ideally be implemented in-band, the remote location and limited staffing created problems with virtual component replacement.

A hardware or software failure could cause an outage of more than a full workday.

Learning: Although not available as a standard offering, a warm standby solution was made available that would allow a fully pre-configured backup device to be swapped in for the in-production device if needed.

appropriate government offices are provided with the applicable encryption keys. It's important that both you and the MSP are aware of these restrictions so that appropriate solutions may be developed within the context of each country's requirements.

In addition to considering regulatory implications, MSPs must weigh other local factors that might influence performance. The distribution of sites within a country/region or between regions, and the ability of the MSP to support those site distribution patterns, must also be taken into consideration; this is doubly true in Next-Gen Networks. The locations of the centralized control elements for SD-WAN or for installed VNFs must be considered carefully — a global network requires careful planning to avoid the effects of high latency. The locations of critical application hubs within the network must be chosen carefully to similarly avoid high latency and round-trip-time (RTT) issues.

Customer networks with a limited geographical distribution, such as sites confined to a reasonably small core region, will likely not be affected by these considerations, as long as centralized control and infrastructure elements are available within the same region. Conversely, customers with a global network always need to consider the worst-case data flow scenarios and plan for the network to bring the performances for these worst-case scenarios within acceptable limits. This may involve planning for multiple centralized tooling locations and for multiple internet broadband (BB) termination locations. For an extreme example, if the customer has an internet BB site in Sao Paulo Brazil and a Multiprotocol Label Switching-Border Gateway Protocol (MPLS-BGP) site in Sydney, Australia, the latency issues for traffic between the sites can be greatly magnified if the only infrastructure sites for tunneling services available to that customer are all in Europe.

The prevalence of internet BB as a transport medium drives additional considerations. Additional infrastructure elements are typically needed to cross-connect internet BB tunnels to the network management domains or to other WAN transport services. The tunnel terminations within the network infrastructure must map out against customer sites and the centralized tooling locations. While customers may often be cognizant of these geographical considerations, providers with practical experience in dealing with global networks should have a more complete understanding of the potential issues.

Experience counts!

DIMENSION 9: Site Attributes

In addition to geographic and governmental considerations, the number of sites and site types play important roles in devising and deploying a network.

The number of sites connected to a Virtual Private Network (VPN) across an Next Generation Network (NGN) is significant for many reasons. There are limits to the number of routes that can be passed or stored by any given interface/router. With a sufficiently large number of any-to-any sites connected into a network, this can lead to overrunning virtual routing and forwarding (VRF) tables within network provider edge routers.

In an NGN environment, each site is connected to infrastructure components that aggregate "northbound" traffic destined for the centralized control tooling. These infrastructure components will also have a limited number of connections that can be supported at any given time.

The orchestration/tooling portals have limits for the number of components they can manage at a time.

These latter two examples hint that it is possible that existing and available infrastructure elements within the MSP framework may not be sufficient to support the network for a sufficiently large number of sites. As most SD-WAN platforms include these additional infrastructure elements in order to support their central control-plane functions, the number of sites supported within a Next-Gen Network may require the addition of more infrastructure elements than those that are already installed and available within the MSPs network. These additional resource requirements should be identified during the initial design phase of network provisioning so that they may be provisioned in a timely manner.

The distribution of the infrastructure sites within a global network again comes into play when assessing

the network for the availability of these infrastructure elements. Global networks should consider aggregate routes and summary routes for traffic traversing well-defined regions. If each region is properly assigned its own infrastructure elements, then interregional traffic aggregation can greatly reduce the number of advertised routes received by any network device. However, the distribution of sites inside regions within a global network may require additional regional infrastructure elements to support effective interregional traffic aggregation.

The other consideration for this dimension is the definition of site types across an extended Next-Gen Network. Typically, a site type is defined by at least three attributes: WAN access type(s), site resiliency, and site priority. Sites with all matching attributes are typically designed identically to each other. Variations within the WAN access types were covered in the first dimension (transport), but the other two primary attributes should be considered carefully as well.

Site resiliency applies to any site for which configurations have been employed to ensure that the failure of one or more key components within the site does not disable the site. Dual-access resiliencies ensure two or more paths through the network so that a failure of one access link does not isolate the site from the rest of the network. Dual customer-premises equipment (CPE) resiliencies exist so that if a CPE device (e.g., router/virtual router) fails, the site can route traffic through a second CPE device at the site. In addition, service interworking resiliencies are designed to allow failures affecting a site on one network to send the necessary traffic to an alternate network.

Resiliency should be distinguished from diversity in this context. Diversity indicates that failure of a single key component within the network between the customer premises and the network edge, external to the customer premises, will not isolate the applicable site.

The benefits of resiliency (and diversity) are fairly obvious — they increase the reliability of the site — but there is a tradeoff in capital costs. Multiple sets of CPE devices or multiple WAN circuits within a site necessarily increase the cost per location. Consequently, resiliency must always be weighed against the cost to your business if there is a network or component failure. The more critical the business traffic to or from a site, the greater the incentive for resiliency. For critical sites, such as headquarters sites or application server hubs, multiple instances of resiliency and diversity may be implemented.

Site priority is the remaining site characteristic that helps to establish the site type. In a typical business network, some sites are more important to the operation of the business than others. The most critical sites may be global or regional headquarters, or they may be server hubs that allow mission-critical applications to be run at all other sites. On the other end of the spectrum may be simple end-user (remote) sites that are primarily responding to requests from the critical sites as needed. In between these extremes, any number of levels of priority can be established and implemented.

Depending on your needs, site priorities may dictate either any-to-any connectivity between sites or else some form of hub-and-spoke connectivity, where “spoke” sites only communicate with their designated “hubs.” In the latter case, greater planning is needed to establish the proper hierarchy of sites to optimize data flow, particularly when multiple hubs are established within a single virtual private network (VPN). The often-unpredictable interactions within SD-WAN between headquarters/hub sites for networks with multiple hub sites must be specifically considered when planning out routing for individual sites and may be a limiting factor in establishing your platform requirements.

DIMENSION 10: Customer Environments and Applications

Next Generation Networks (NGNs) blur the line between traditional IT and networking. SD-WAN is application-aware; therefore, considerations such as traffic prioritization or port scanning applications that consume network capacity need to be understood by all involved parties. The business environment within your network may include an IT staff with some level of expertise in SD-WAN and NGNs. Your IT team requirements will likely overlap with the responsibilities of the service provider that is maintaining and managing their network. Openness and transparency between you and the provider help to ensure that your processes and expectations are aligned with the MSP's processes for site management, monitoring, and change management. This ultimately leads to smoother deployment and operation of the network.

Each customer's business processes, particularly those related to IT, are often driven by the specific nature of that business. For example, decisions made in highly regulated businesses, such as banks or government agencies, tend to have very strict requirements for monitoring and documentation. The business environment determines important aspects of network management that you may need to consider, such as:

1. What level of security is required for data traversing the NGN? Are all data flows through your network required to be encrypted, or just some of the data, in order to meet your business requirements?
2. Is partial or full co-management of some or all of the network hardware components something that you have the resources and procedures to handle?
3. Do you have the resources to monitor your own data traffic, in addition to or instead of the service provider? Are the monitoring techniques that would be so employed compatible with the NGN architecture?
4. Can precise demarcation and handoff of responsibilities between the network customer edge devices and your local area network (LAN) devices be negotiated to the satisfaction of both you and the provider?
5. Can your change management procedures be successfully integrated into the overall management of your network by the selected MSP?

Your business environment also determines the applications that must be supported across the network. In extreme cases, the requirements inherent within those applications may limit the choices of platforms or network components allowed in the network design. Standard Internet Protocol IP unicast data applications are generally straightforward and do not present any special problems in NGNs, but there are some applications that require special considerations.

Real-time (RT) applications

Most current NGNs are designed primarily for optimizing and sustaining the flow of Transmission Control Protocol (TCP)-based data traffic and may find the incorporation of User Datagram Protocol (UDP)-based real-time voice/video traffic difficult at best. Voice over Internet Protocol (VoIP) traffic also demands special restrictions on traffic routing that have been designed to help ensure consistent voice quality end-to-end, and these restrictions tend to be at odds with the WAN efficiencies enabled for data traffic within an NGN. Some virtual Session Border Controllers (vSBCs) are currently available from various vendors to efficiently deal with voice traffic within an SD-WAN context, but these tend to work best as standalone devices that process only real-time traffic and have not been widely certified to interact well with other NGN SD-WAN devices. Integrated voice border element components are not yet readily available nor tested broadly within the current SD-WAN environment.

Typical compromises within current NGN environments to accommodate voice/video real-time traffic involve providing separate WAN access ports for VoIP traffic and passing the VoIP traffic through the NGN as high-priority data traffic, with the complex voice routing and switching handled behind the LAN rather than within the WAN-facing SD-WAN components. In this scenario, Quality of Service (QoS) mechanisms are generally available to ensure adequate protection for the real-time traffic so as to allow the voice to maintain acceptable quality levels.

Multicast and anycast applications

Multicast (MC) and Anycast (AC) Applications — While most SD-WAN vendors are equipped to support basic MC/AC functionality, not all available MC/AC features have been fully integrated into the applicable vendor tooling or into the tooling of the major MSPs. You need to carefully consider your requirements for MC or AC connectivity and have the selected MSP verify that your required functions and features of the MC or AC can be supported through the selected vendors.

Non-IP protocol-based applications, such as Systems Network Architecture (SNA) or Data-Link Switching (DLS)

Most non-IP-based applications/protocols require that the non-IP application/protocol be encapsulated into IP for transport across the NGN, as most NGN components are not built for non-IP protocols such as SNA or DLSw. Nonetheless, there are still business customers using these legacy protocols and applications who may desire to keep these functions around as long as they can, and the MSP must evaluate the viability of migrating them onto an NGN.

Highly transactional IP applications, such as airline booking/registration applications

Highly Transactional Applications — Many of these applications behave poorly under attempts to optimize the traffic or under attempts to allow swapping between paths for WAN efficiency, and, as such, make poor candidates for standard SD-WAN connectivity solutions. Typically, such applications require very tight control of the QoS configurations to minimize the need to resend responses more than once.

EXAMPLE #2
Real customer | Use case

Situation: Because expected future growth was not factored into the original design, the customer outpaced the capabilities of their SD-WAN solution at a key site over time and eventually required expansion of the WAN port capacity to a 10 Gigabit Ethernet link that was not supported by the SD-WAN vendor.

Increased peak traffic caused customer-impacting performance issues at the affected site.

Learning: The site was moved back onto a traditional hardware solution that allowed for expansion to 10Gbps access ports.

Extremely bursty traffic or bulk transfer of very large files

These applications are not necessarily any more difficult for NGNs than for other networks, but extreme cases of these applications can overwhelm optimization mechanisms typical of SD-WAN and cause performance issues within the network. Therefore, these applications may require special handling (such as execution time) to avoid unacceptable delays to more time-sensitive applications.

Ultimately, for your business environment to be successfully served, mutual respect and trust between you and provider is required. Clarification of responsibilities and mutual transparency will enable communication that will deliver desired business outcomes.

In Part IV of our series on Next Generation Networking, we will discuss planning and management approaches to help your business successfully adopt NGN solutions.

To learn more about how AT&T Business can help your business build a tailored network solution, visit www.att.com/networkservices **or call 866-415-0949.**

PART IV
The PLANNING AND MANAGEMENT dimension of Next Generation Networking (NGN)

Introduction

In part III of our series, we discussed the customer dimensions of Next Gen Networking (NGN) solutions. In this final part of our series, we'll review the Planning and Management implications of adopting NGN. AT&T Business is here to assist you in choosing the best solutions to achieve your unique objectives. We draw on our expertise to deliver critical insights to guide you as you progress down the path towards implementing advanced networking resources. Read on for tips on how to maintain control over these powerful NGN solutions.

A transition to NGN solutions involves significant changes to existing infrastructure and procedures. Unless the Next Gen implementation is a new network, the infrastructure must somehow transition from an existing solution to the Next Gen solution. Such a transition requires effective planning to ensure that this migration occurs as smoothly and non-disruptively as possible. You and your provider need to consider how to maintain the continuity of communication and data flows between sites that have been migrated onto the NGN and those still awaiting migration. This is a critical consideration, and failure to address this dimension adequately can cause site failures and network traffic implications that lead to lost productivity and frustration.

Assuming a service provider is involved, we can consider two categories of network transitions: inter-provider and intra-provider. In the first category, inter-provider, the initial network is managed in whole or in part by one provider but will be transitioning to a new provider. In the second category, intra-provider, the existing network is a distinct managed service offering from the proposed NGN but will be transitioning without changing providers. For obvious reasons, the latter option is often simpler and less disruptive.

Inter-managed-service-provider (inter-MSP) transitions are typically limited by the requirements for maintaining the Network Management traffic paths within the pre-transition network, which may often be incompatible with the requirements for the post-transition network. Additionally, the management paths for a given device cannot be controlled from two distinct providers at the same time; as soon as a device transitions from one provider to another, the original management paths are no longer active. Overlapping billing and contracting costs between distinct providers during transition add to the logistical problems that must be overcome for this type of migration.
The simplest form of inter-MSP transition is known as a Walk-In Take-Over (WITO). In this scenario, if the existing devices are supportable by the new MSP, then the new provider supports the existing network while transferring network management capabilities to the new infrastructure. After the transfer, the transition to an NGN essentially becomes a simpler intra-provider migration. Where WITO is possible, it is typically a less disruptive inter-MSP transition; unfortunately, this is not always available as an option.

The most common form of inter-provider transition solutions involves providing one or more sites within the existing network at which a router from the "old" network and a router from the "new" network are collocated. These networks are then interlinked through a shared connection between the routers that provides a path for data to traverse. The greater the number of these dual-network sites that can be installed and maintained in parallel during the transition period, the less likely it is that the network will be overloaded (or will be crippled by a single point of failure) and will seriously affect network performance while the transition is in progress.

Network transition relies on the effective cross-connecting of the original data network and the new NGN. This cross-connection between the original network and the new Next Gen network to maintain the continuity of this data traffic can occur in several ways:

1. Directly within the network infrastructure (for an intra-MSP migration); this typically occurs in AT&T when migrating from the managed virtual private network (VPN) offering with Multiprotocol Label Switching-Border Gateway Protocol virtual private network (MPLS-BGP VPN) to Software Defined Wide Area Network (SD-WAN) or AT&T FlexWare with AT&T VPN (MPLS-BGP VPN).

EXAMPLE #1
Real customer | Use case

Situation: A customer security officer decided to turn on security scanning features at the company's corporate headquarters.

The scanning tool took down critical infrastructure components at the data center, impacting the full customer network.

Learning: The operations team quickly identified the cause of the infrastructure failure and was able to have the individual that initiated the issue shut down the scan.

2. Via an internetwork gateway device specifically deployed to allow route interexchange between the two networks (for an intra-MSP migration). This case would also include network-infrastructure-based firewalls designed to link internet connections to other circuit types (e.g., MPLS-BGP VPNs).
3. Route interexchange between customer premise equipment (CPE) WAN interfaces on devices at one or more locations (for an intra-MSP migration). This requires each such device to be directly connected to both the old and new network WAN circuits.
4. Route interexchange between two or more CPE devices at the same location via a direct cable or direct LAN connection (one device on each network). This solution may occur at more than one location within the combined networks.
5. Route interexchange between two or more CPE devices at the same location via a common (shared) LAN connection (one device on each network). This solution may occur at more than one location within the combined networks.

The following are considerations when planning provider transitions:

Technical

1. How will the NGN MSP ensure continued connectivity between migrated and unmigrated sites during the transition?
2. Can both you and the MSP ensure the ability to test success migration and certify the ability to roll back the changes at individual sites if necessary?

Project management

1. How will the MSP coordinate the various vendors, local customer contacts, and the customer's corporate management for turning up sites onto the NGN? You, as the customer, will typically have a plan for the sequencing of migrated sites, and this plan must be compatible with the processes used to maintain the connectivity between sites on both ends of the transition.
2. Will there be systematic tracking (and escalation as needed) of the WAN transport circuit delivery to allow the transitions to proceed in a timely and controlled manner? New access circuits are typically a limiting factor in processing orders for migrating sites. If not carefully tracked, the desired sequencing of migrations may be put in jeopardy.

Operational readiness

1. Does the provider have an appropriate staging/shipping infrastructure available to ensure that new hardware arrives on time and is ready for installation when you are ready? If the hardware arrives too early, it risks becoming misplaced/lost prior to the turn up of the affected site, or, in extreme cases, becoming obsolete or back-leveled prior to turn-up. If the hardware arrives too late, the access circuits may generate billing costs while being unusable by you, forcing either the MSP or you to absorb extra costs.
2. Can the MSP (and you) ensure the availability of personnel resources to perform the transitions at each site? These personnel could be resourced from the provider directly and/or your local technical contacts.

Each of these options and questions needs to be carefully considered to minimize the creation of single points of failure or to prevent overloading of interfaces that bridge the two networks, and multiple options might be considered. On the other hand, providing multiple network paths through CPE devices itself requires due diligence to avoid creating circular routing loops during the transition stages.

Intra-provider transitions are generally much simpler to coordinate. However, transition planning is still an important consideration — it's not safe to assume that because two offerings are maintained by the same provider, these two services can communicate across their respective infrastructures without deliberate configuration and planning. This improper assumption is often made by customers and sales teams alike, and it can cause significant problems for the transitioning customer. In the best-case scenario, the provider may be able to extend the existing customer network/VPN directly into the new NGN service, but other solutions are also possible.

Planning for circuit delivery

To complicate matters even further, there must also be similar planning efforts undertaken for new sites on new networks in order to properly assess the new challenges and to address old challenges that may have been exacerbated by NextGen Network features and technology. Most notably, there are challenges posed by the new ability to mix-and-match the provisioning of diverse circuit types into the same equipment. Timelines for the installation of network components are usually driven by the lead-time required to provision a new WAN transport circuit, but, for earlier network designs, only one type of transport circuit could typically be so provisioned into a terminating device.

With the new SD-WAN capabilities, service providers are being asked to provision multiple WAN transport circuits of different types, each with their own established lead-time ranges. This then requires intricate planning to ensure that all such WAN circuits are delivered within a short time interval of each other for precisely the same reasons that we plan for hardware delivery in operational readiness. Failure to plan for this aspect can lead to one of two undesirable outcomes: either the site Test and Turn Up is scheduled and attempted before all WAN circuits are ready, leading to impaired NGN capabilities for the affected site until the remaining transport circuits are ready, or else a WAN circuit may be delivered that cannot actually be used for an extended period of time, during which it also generates billing.

Conclusion

Next Generation Networks (NGNs) are a new and intriguing product that offers tangible benefits, and it is worthwhile to consider NGNs as the basis for a robust business data and voice network. The technologies involved in virtualization of the network and its components offer the potential for capabilities and levels of performance not previously available from traditional physical components. As with all new paradigms, it is still a new product and subject to a constantly evolving environment as it matures.

As such, we need a realistic outlook on what to expect. As a business customer, prospective adopters are bombarded with the pictures of the endless potential of this technology, but these customers also need to understand the innate complexity and the growing pains associated with those promises. Hopefully, this document will allow businesses to adjust expectations and set realistic goals and priorities for themselves and their business networks.

As an experienced service provider that has championed NGN technology for a number of years, AT&T Business would like to share these insights gained over that time to make business customers more aware of the true current state of the product. Our goal is to educate customers as to what questions to ask to make sure that regardless of which provider is selected, you can have your expectations met.

The eleven dimensions described in this series are not the full story; certainly, there are other considerations that could be highlighted. For example, network security impacts many of the discussed dimensions, and is referred to in a number of dimensions, but it could be rightly considered as a separate dimension. Certainly, in the future, the full realization of 5G technology will change the business networking landscape, as will the boom in Internet of Things (IoT) implementation across global networks — but including these aspects now would be more speculation than education. As products evolve and stabilize, there will be new questions to ask. AT&T Business will be there for you in consideration of these new aspects as well, walking you through the tough questions.

The issues of today will morph into new issues tomorrow. We cannot necessarily anticipate all future developments, but at AT&T Business, we can put forth the accumulated experience and knowledge of our employees as a pledge that we will be prepared to deal with whatever may come.

> "There are an estimated 32 million businesses in the U.S. That's potentially 32 million different network configurations. If you are one of those businesses, you want the one network build that's a spot-on fit for your specific needs. AT&T Business has the breadth and depth of networking solutions to make that happen."
>
> **Suzanne Galvanek**
> Vice President of Product Management
> Enterprise Networking Solutions at AT&T Business

AT&T Business is ready to assist you in harnessing Next Generation Networking solutions that help advance your specific business priorities. With so many network configurations available, the right choice depends upon identifying the unique challenges and opportunities facing your business. That's why we invest the time in understanding your goals before recommending an ideal solution.

To learn more about how AT&T Business can help your business build a tailored network solution, visit www.att.com/networkservices **or call 866-415-0949.**