turnium

Best Practices in Hybrid Cloud and Networking

Carrier and Enterprise Solutions for
Multi-Cloud, Hybrid Cloud and Hybrid Networking

AWS | Azure | IBM Cloud | EQUINIX | Alibaba Group 阿里巴巴集团

Table of Contents

Table of Contents

Introduction

The cloud, while eliminating data center hardware and support requirements, increases the importance of networks. Carriers and enterprise require cost-effective, software-defined wide area networks to easily connect geographically dispersed branches, sites, and facilities to existing networks and applications stored in public and private cloud.

Solving multi-region network challenges often requires agreements and facilities with multiple national or local carriers. Providing and obtaining Day 2 support on these networks, once provisioned, adds complexity and challenge.

Turnium's wholesale, white-label SD-WAN platform offers carriers and enterprise a simple, cost-effective, multi-tenanted solution to delivering and managing private, secure over-the-top networks that support multi-path fixed-mobile and mobile-mobile converged networks at every site while using the carrier's or enterprise's existing IP address space. Turnium powered networks enable deploying:

- A Layer 3 that carriers and enterprise can fully host, brand, and control
- Managed network services to sites without fiber or where MPLS is not cost-effective
- An automated, managed network environment that eliminates manual routing configuration and testing, decreasing deployment time as the network expands or shrinks
- Built-in AES 128/256 encryption to secure data in-transit upstream and downstream
- Data obfuscation enabled by distributing packets across multiple circuits from multiple carriers to edge sites
- Economic models that support profitability, universal CPE deployments, IoT network, build your own SASE, equipment sparing cost management, and other scenarios in which managing SD-WAN cost is critical.

Traditional Telecom Networks vs. Turnium SD-WAN

Examining a typical use case can help clarify and demonstrate the benefits of Turnium. Let's imagine an organization moving applications into multiple clouds. To connect these clouds with their various branch offices, sites, or facilities, the organization first needs to establish a Wide Area Network between all their locations and the cloud facilities.

Solution Requirements

1. **Build** a network between multiple cloud providers and customer offices. Selecting the connectivity to be used will be based on the budget available for each office, a decision based on the job functions and staffing levels at each location.
2. **Source** multiple quotes from different carriers, determine failover options, obtain financial approval, sign multiple contracts, provision services, complete and test failover routing. service levels.
3. **Ensure** all network circuits support the required Quality of Service for voice, video and other applications.
4. **Maintain** the network while ensuring performance visibility and monitoring to ensure Service Levels are met and end-user experience is maintained.

Traditional Telecom Network Design Solution

A typical telecom network build includes a combination of different carrier circuits at each site for failover. Multiple carrier last-mile circuits are deployed at each site, usually in an active/standby mode with each carrier providing their own router as a demarcation point at the customer premise. These carrier networks and connected to the customer's core network. Internet-based VPN connections are used to reach sites where dedicated carrier circuits are not cost-effective or available.

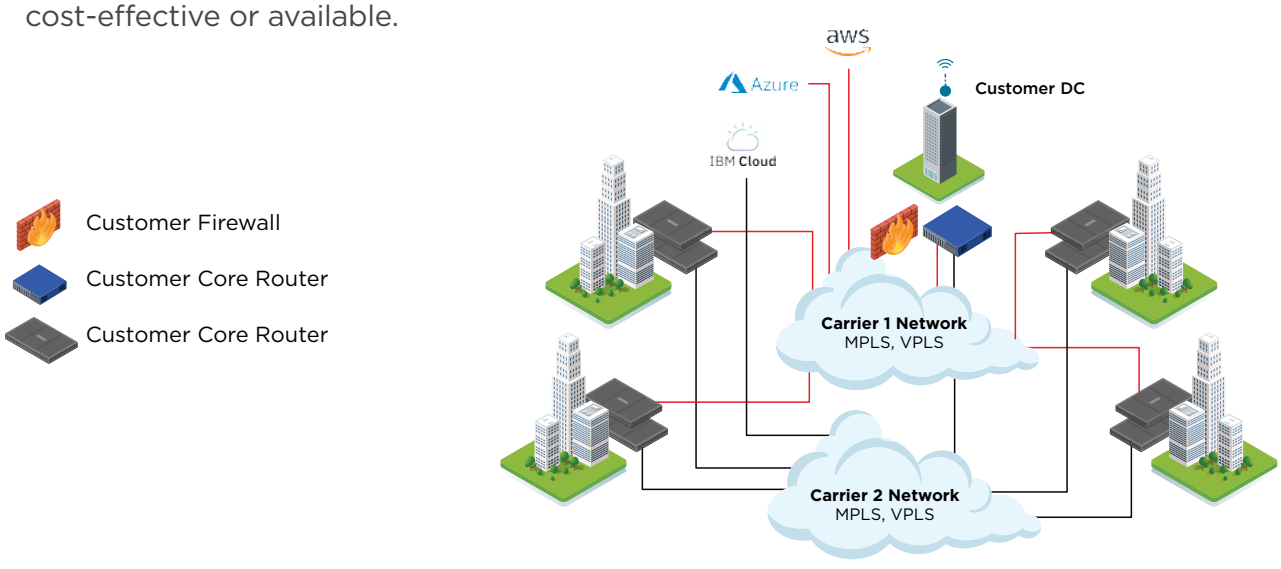

CPE routers managed by the carrier terminate MPLS/VPLS or Layer 2 Ethernet circuits at each customer site into a customer's routed environment and connect back to the Provider Edge (PE) router in the carrier's network. The PE allocates bandwidth and provides features such as QoS as the packets are transported across the carrier network using VRF or VLAN to other circuits connecting the carrier network to the targeted cloud providers and/or data centers.

Carrier demarc routers and the carrier networks to which they connect are "black boxes" to the organization or enterprise, meaning that there is no visibility to the router or the network behind it. Troubleshooting requires involving the carrier – in the worst case, it involves carrier staff coming to-site to perform testing.

In traditional carrier networks:

- The customer doesn't have visibility or control over the network – they rely on the carrier for information, support, and maintenance.
- Last-mile circuits charges are based on bandwidth reservation and features such as QoS.
- Often, carriers also charge based on the volume of data transmitted from each site.
- There are usually data ingress/egress charges applied by any Cloud provider.

SD-WAN Solution

SD-WAN provides a secure way to achieve the same goals cost-effectively while keeping data in-transit encrypted.

An SD-WAN network with centralized and secure Internet access can be deployed to connect multiple existing sites, including offices and data centers.

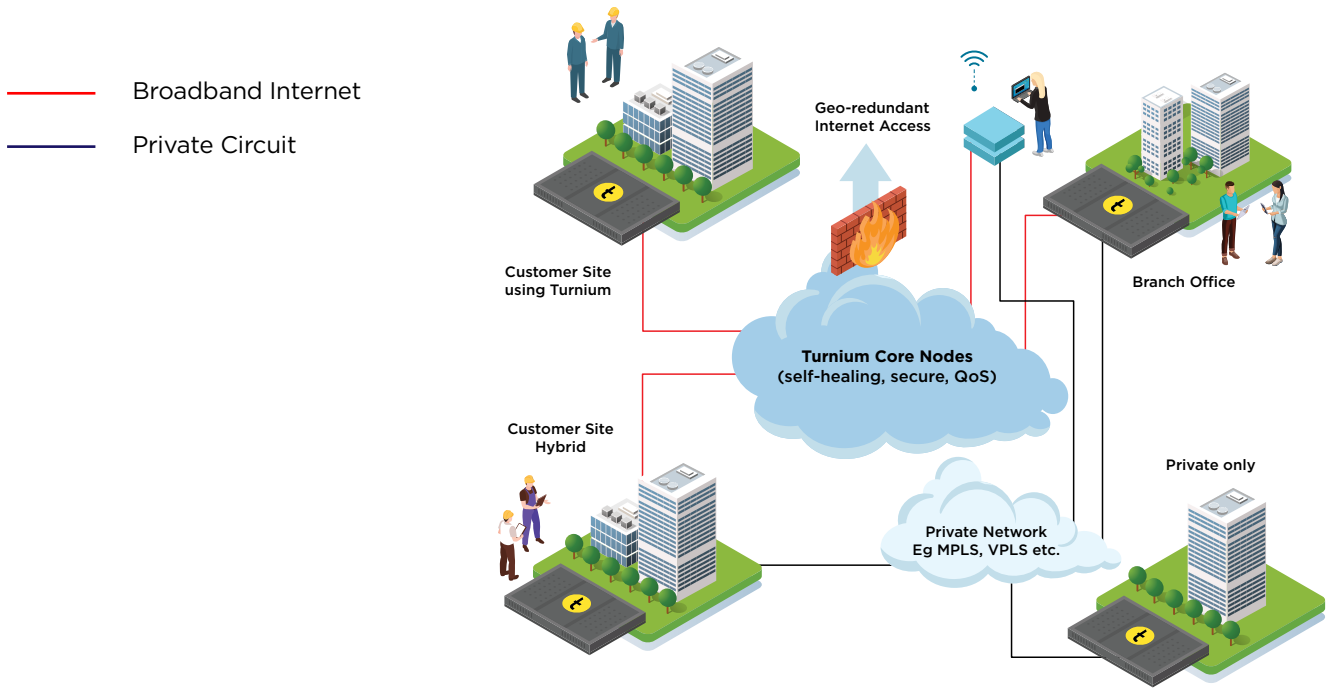

An SD-WAN network with centralized and secure Internet access can be deployed to connect multiple existing sites, including offices and data centers.

In addition, **public hyper-scale clouds** can easily be integrated into Turnium SD-WAN deployments, as shown below, with Core or Edge nodes deployed in public clouds, depending on the required functionality. Either option enables private, routed networking between the cloud, branch, and head office, ensuring **security, prioritized bandwidth, and application performance,** while using any combination of available last-mile circuits, or integrating into existing Layer 2 networks.

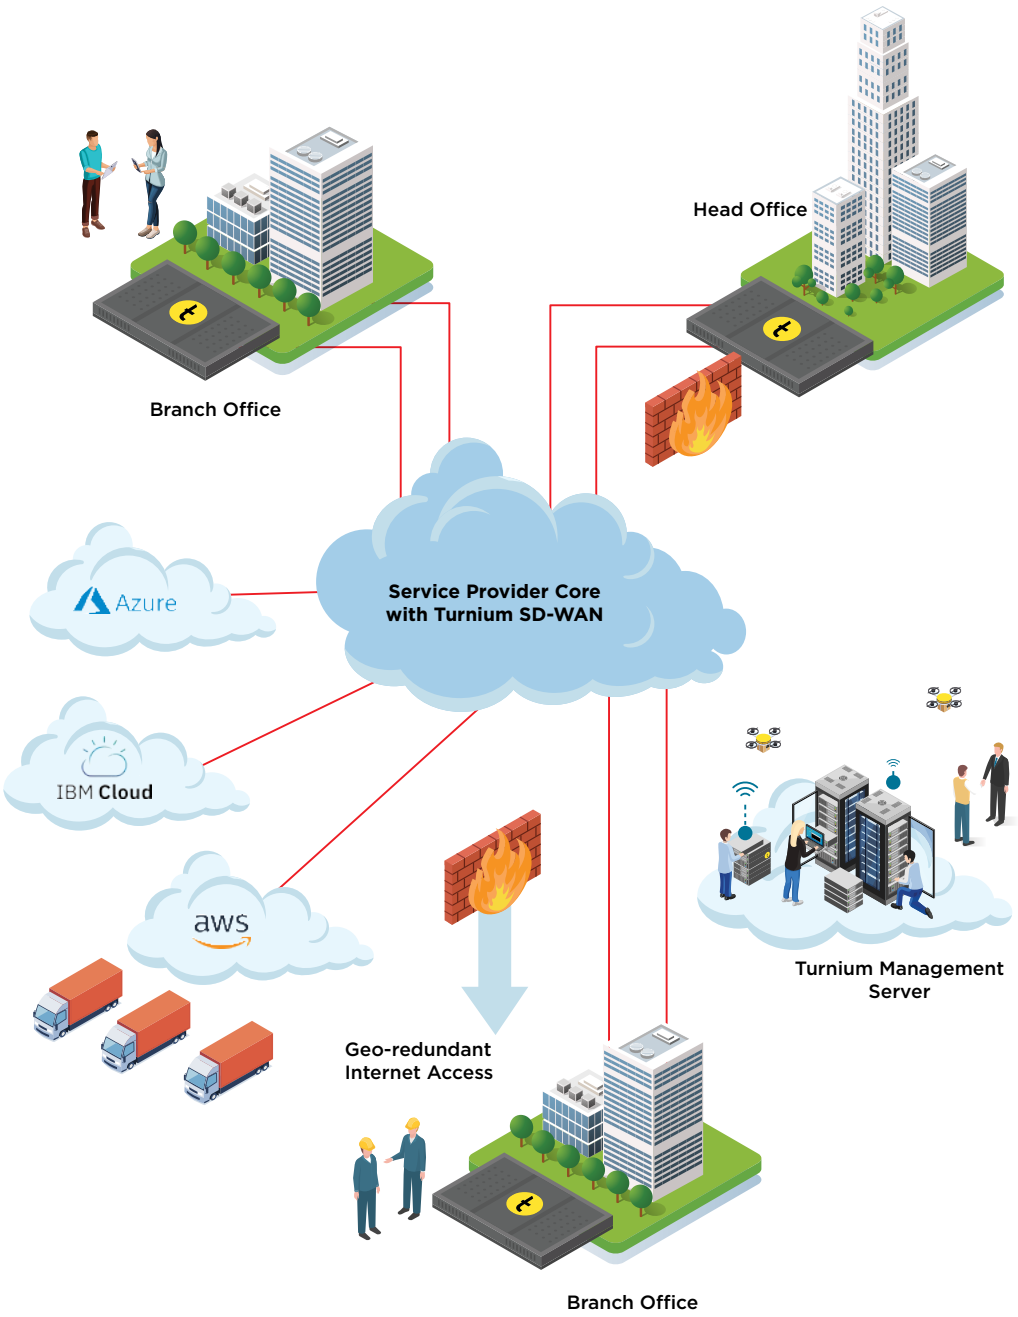

Turnium SD-WAN Architecture

Turnium SD-WAN enables **Service Providers** — including **Managed Service Providers, Telecoms, Voice and Contact Center Providers** — to deliver secure, dedicated, private routed networks for each of their customers from a multi-tenant core SD-WAN environment.

Using a simple point-and-click GUI, Turnium Partners can deploy multi-site networks on multi-tenanted core infrastructure that they, themselves, host and control. Turnium's licensing model:

- Provides Partners with unlimited, no-charge, core node licenses so Partner environments scale with no incremental cost
- Is based on a flat-fee, per active site, per month basis, giving Service Providers a predictable, simple-to-quote managed service to add to their sales.

Setting Up Turnium SD-WAN

Setting up Turnium SD-WAN involves deploying virtual or bare metal **Aggregators** within Service Provider infrastructure and virtual, bare metal, or containerized **edge nodes or CPE** at customer sites or in cloud. The same software image is deployed as an Aggregator or CPE, making control of images simple. When creating a node in the Turnium **Management Server or orchestrator, a node key** is generated that identifies the node as an Aggregator or edge CPE when the software is installed on the node.

Aggregators are **multi-tenanted**, capable of supporting multiple customers and multiple sites each within their own private spaces. Aggregator spaces can also be branded per-client, and delegated administration rights applied to enable customers, or downstream resellers, access to the environment for view-only, or limited MACD capabilities.

In terms of Aggregator density, Turnium recommends:

- Deploying 50 sites per Aggregator for customer sites to be configured with SD-WAN tunnels of between 100Mbps and 200Mbps per site.
- More sites per Aggregator can be deployed if site throughput is lower.
- Reducing the number of sites per Aggregator for higher speeds.

The CPE at each customer edge or site accepts LAN traffic and breaks it into streams of packets, which are then transmitted across the circuits that comprise the SD-WAN tunnel. Turnium SD-WAN can create an SD-WAN tunnel using multiple wired and wireless circuits at the same time, at each site. This enables Service Providers to deliver highly reliable multi-path fixed-wireless and wireless-wireless converged networks using combinations of:

- Broadband
- Dedicated Internet Access (DIA)
- Ethernet
- MPLS
- Fixed Wireless
- 4G/LTE and 5G

Satellite connections can also be used, but it is usually impossible to aggregate or bond Satellite with Landline connections due to the extreme differences in latency between the two types of transmission.

Aggregating multiple Satellite links has been more successful, and adding *Replify WAN Optimization* into that solution can also increase performance. *Replify WAN Optimization* is available at an extra cost per-instance required.

Hybrid Networks

Using Turnium's ability to create an SD-WAN tunnel from any available circuit, hybrid networks can be designed to support customers that need to integrate SD-WAN into their existing network architecture or are seeking to migrate to SD-WAN over time as their carrier contracts expire.

Link Aggregation

In Turnium SD-WAN, a site with one or more circuits connected to a CPE is often referred to as a **bond**. The term can also refer to the SD-WAN tunnel, which results from bonding the bandwidth of multiple circuits together. Bonding is also referred to as **link aggregation**.

Quality of Service

Using the GUI, a Service Provider creates **spaces** for each customer, adds sites and bonds per space and configures and tests the bonds. **Quality of service** profiles can also be created to define the class of service and priority of the application traffic and assign bandwidth appropriately.

Prioritization

Turnium SD-WAN's prioritization is elastic, meaning that bandwidth is only assigned to prioritized traffic when the associated packets are present in the data stream. When prioritized packets are not present, bandwidth is not reserved and can be used by other traffic classes, helping increase customer experience and ROI on bandwidth.

Other settings and configuration are done through the GUI, including encryption (AES128 or AES256), setting Bandwidth Adaptation (used to manage the circuits in the SD-WAN tunnel to minimize latency) and determining which of three algorithms will be used to allocate packets to the circuits in the SD-WAN tunnel.

Management Servers & Delegated Administration

Each Turnium Channel Partner deploys their own Management Server. The Management Server is used to control or orchestrate all end-customer or reseller SD-WAN. Permissions can be assigned to end-customer or reseller users, giving them rights to view, edit, or manage their networks. This enables Turnium Partners to set up and support downstream networks of resellers, each with their own environments, control, management, and branding in a self-serve manner. This capability meets the requirements of many customers and resellers for self-control, self-determination, and visibility.

In specific instances, Turnium Management Servers can also be provided to end-customers or downstream channel partners who require complete control for governance or compliance reasons.

Deploying Turnium Aggregator and CPE Software

It is recommended to run Aggregator and CPE on bare metal hardware for the best performance. However, Aggregators and CPE software can be run in virtual environments or in containers. CPE and Aggregators can also be deployed in Cloud IaaS environments. CPE software can be deployed in a Cloud to bring that Cloud on-net in a Private Wide Area Network.

In addition to bare-metal devices, the following virtualization and container technologies have been used during development and testing:

- LXC
- systemd-nspawn
- QEMU/KVM
- Red Hat ® OpenShift ®

During the release-candidate phase, Turnium SD-WAN is also tested in a production VMWare environment.

Container Environments

Turnium Edge nodes or CPE software can be deployed in container environments to maximize flexibility and management.

Deploying Edge Nodes or CPE In Container Environments

- Lightweight and more efficient than virtualization
- Managed or unmanaged
- OCF-compliant
- UBI based
- Certified OpenShift operator for cluster deployments
- Certified on OpenShift (runc/crio)
- Tested on K8's
- Tested on Podman (crio)
- Tested on nspawn
- IBM ROKS certified

Additional notes:

Container systems such as LXC and nspawn already use a reasonably efficient veth device by default. The primary advantage of these devices is that they do not have to emulate a physical device type, allowing the host and guest to pass packets relatively quickly via system memory.

Virtualization Best Practices

Turnium SD-WAN will operate in many types of virtualization for all host or node types:

- Management servers
- Aggregators or Core nodes
- CPE or Edge nodes

Virtualization makes it easy to provision and manage hosts, but performance is typically negatively impacted, even when the virtual machine is the only machine on a host.

The following best practices are intended for Private WAN routers (an older core node type), Aggregators and CPE. These nodes are a core part of your customer data network and very sensitive to resource availability and efficiency. Management Servers should be configured using practices generally accepted for web and database applications. *(For example, Management Server requirements focus on memory size and storage performance rather than CPU and network device performance.)*

General Recommendations for Virtualization

CPU

Due to the critical latency demands of networking, CPUs should be dedicated to the virtual machines. Sharing CPU cores negatively affects latency, which results in lower throughput and general instability of bandwidth.

Note that **disabling hyperthreading can yield performance improvements on Bonders that are CPU-limited.*

The processor is essential in virtualizing Turnium SD-WAN, and all resources should be reserved as the hypervisor is not fast enough to assign additional resources as-needed.

Storage is generally not as critical to virtualizing Aggregators or CPE as other resources, but care must be taken to avoid high disk read/write latency. If disk I/O operations take too long, service failures may occur.

If memory is low, the disk will be used to swap memory pages. If that occurs, the disk will be used more extensively, and the entire system performance will be negatively impacted.

Network

Network device virtualization methods incur overhead on network performance. A certain amount of CPU and memory is used to implement a virtual interface that copies network packets between the physical interface and the guest operating system.

Most virtualization systems have a relatively low overhead virtual device that should be used instead of full emulation. For example, VMWare offers a VMXNET3 device, while QEMU/KVM offers a VirtIO device.

**Note that these virtual devices are still not as efficient as using the card directly. Most modern server network devices have advanced offloading and acceleration features that are not always exposed via virtual devices. When the traffic load is very high, it may be desirable to consider passing dedicated network devices directly into the guest operating system.*

Storage	Network
When virtualizing the Turnium Management Server, disk I/O is very important, especially as the number of nodes increases. At over 250 nodes, a dedicated time series database server that meets or exceeds specific disk I/O requirements must be deployed. In virtual environments, memory must be reserved for the nodes. Generally, 2GB should be enough for most nodes, but this should be increased when using the TCP proxy or larger numbers of private WAN spaces.	

VMWare Tips

1. **VMWare Tools:** Install VMware tools. The open-source tools are acceptable; these can be installed from standard Debian repositories with:
   - apt-get install open-vm-tools -y
   - service bonding restart

2. **Segmentation:** If you are using Private WAN with encryption, you must disable TCP segmentation offload (TSO) on all the Aggregators running in VMWare. The VMWare VMXNET3 driver has an issue with TSO in combination with IPSEC that results in greatly reduced throughput.

3. **Latency Sensitivity:** Idle-wakeup latencies for guests may be reduced by setting the latency sensitivity option from Normal to High. This is found under VM Settings > Options tab > Latency Sensitivity.

Deploying Turnium SD-WAN CPE or Edge Nodes in AWS, IBM Cloud or Microsoft Azure

Turnium SD-WAN networks can integrate popular public cloud services including AWS, IBM Cloud and Microsoft Azure to bring corporate data or applications in these cloud services into a Turnium Private Wide Area Network (PWAN).

Deploying a Turnium Edge Node/CPE instance in these Clouds is extremely simple and can take as little as 15 minutes. Using cloud providers to launch services with Turnium gives you:

- The easiest way to bring cloud environments on-net and into your corporate WAN.
- The ability to ensure reliable, secure connectivity for applications wherever they are hosted, or
- wherever your end-users are located.

  Flexibility to create your own hybrid public/private cloud-based network.

  Scalability – by expanding or contracting SD-WAN resources in the cloud with bare metal, or virtual private server deployments.

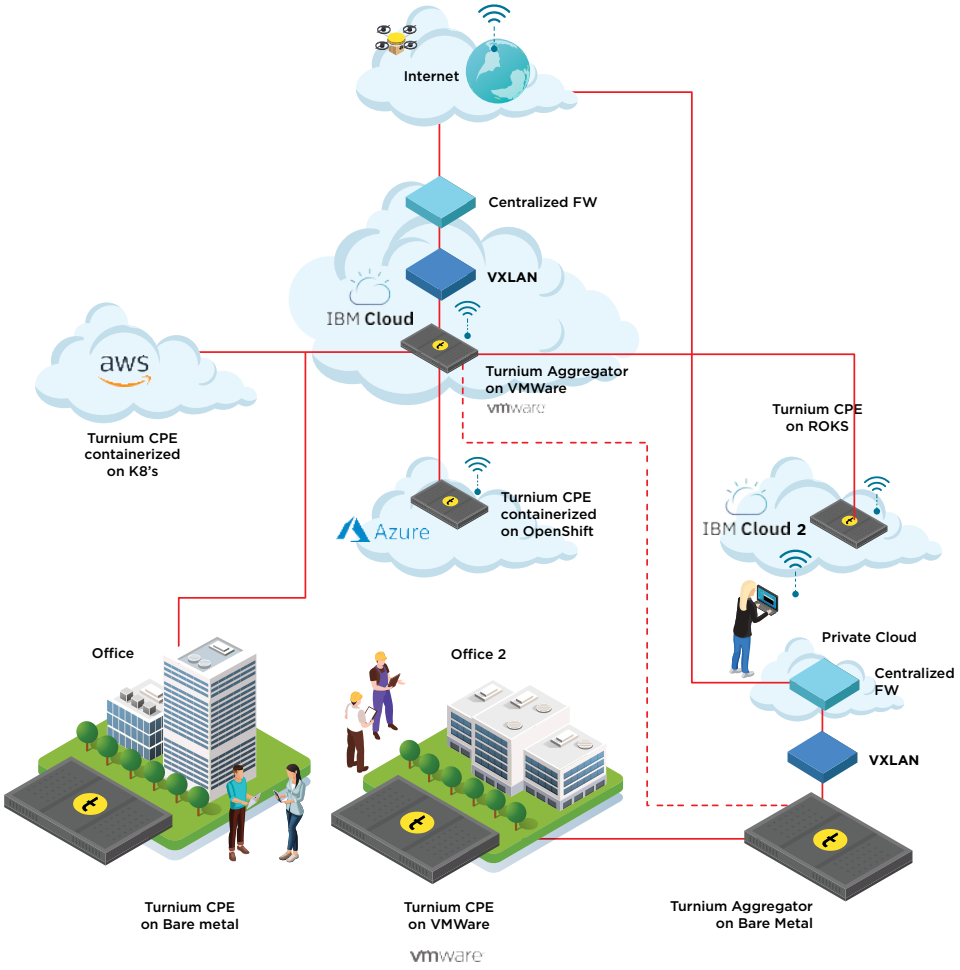

Deploying Turnium SD-WAN Aggregators Core Nodes in AWS, IBM Cloud or Microsoft Azure

Turnium SD-WAN networks can be deployed using popular public cloud services including AWS, IBM Cloud and Microsoft Azure to host and deliver Turnium Aggregators.

Networking solutions like SD-WAN generally appear easier to deploy in IBM Cloud (and some other options, such as OVH) than AWS and Azure, due to the way in which the latter allocate IP addressing and networking.

(In the next few sections, we'll examine scenarios that demonstrate methods found to deploy in these environments, but there are limitations.)

Turnium SD-WAN was built for network-centric deployments. SD-WAN is, at its core, a networking technology. Our solution was built for environments in which channel partners control their own IPs, core devices and networks.

The application-centric nature of AWS and Azure data centers raises challenges when deploying SD-WAN solutions in these cloud environments:

- It can be hard to get IP addresses in these environments
- Available IP addresses are NAT'd, which means:
  - SD-WAN core nodes don't get public IP addresses
  - Public IPs are routed to a private IP on the core node

It's been our experience that Channel Partners who start deployments in AWS or Azure realize the benefits of more network-centric cloud providers and switch early in the process.

Deploying Aggregators in Amazon Web Services (AWS)

Turnium SD-WAN can be deployed using AWS to host Aggregators.

Benefits	
• Allows for quick deployment and management of Internet-connected servers	• The most popular option for hosting websites

Caveats	
• When using IPv4 in AWS, there is no option for routing subnets to AWS hosts. It is only possible to allocate single IPs, and those IP addresses are randomly assigned, making it impossible to emulate subnet routing with IPv4. • Subnet routing can be emulated with IPv6 addresses, but each address in each range will need to be assigned to the aggregator interface, which is burdensome.	• Public IPv4 addresses in AWS are never assigned directly to instance interfaces. All public IPv4 addresses are assigned to a private address and translated using NAT. While aggregators are capable of handling tunnels in such a scenario, it makes bond routing complicated.

Despite these cautions, there are some scenarios that can work if set up in a particular fashion:

- CPE NAT IP to individual bonds
- Configure the CPE NAT IP
- Route CPE NAT IPs across a group of aggregators
- Deploy Turnium SD-WAN Private WAN in AWS

CPE NAT IP to Individual Bonds

What's Involved	Typical (Non-AWS) Deployment
This scenario involves providing individual per-bond access to the Internet, each with its own public IP address.	A subnet containing IPs available for use for CPE NAT IPs is routed to the aggregator network and the aggregator announces each IP via OSPF or iBGP. If dynamic routing is not available, static routes can be made to individual aggregators hosting the bonds, but this requires manual routing changes if a bond is moved.

Each public IP address is assigned to a private IP address that is configured on an interface. By default, when an AWS EC2 instance is created:

- An Elastic Network Interface (ENI) is created
- A dynamically chosen private IP address is assigned to and configured on the interface
- A public IP address is associated with the private IP address

An aggregator can use this address for its primary communication, but *each CPE NAT IP will need additional addresses*. The AWS Elastic IP (EIP) can be used to provide this, but it must be carefully configured to avoid losing the original assigned public IP address.

The procedure to do this is described in detail in Turnium SD-WAN Application Notes. Note that the aggregator's address does not need to be configured in the Aggregator Network Configuration in Turnium SD-WAN.

Configuring the CPE NAT IP

On a bond assigned to the aggregator, configure a private connected IP and add the CPE NAT IP using the private IP address added to the interface. The AWS gateway will NAT the public IP address to the private CPE NAT IP, and the bonder will NAT the private IP address to the connected IP.

Assuming the bonder is online, all of the routing will be in place.

Note that due to the lack of dynamic routing in AWS VPC, moving bonds between aggregators will not work unless the static associations are updated using the AWS console or API. As a result, aggregator failover will not work in this scenario.

Route CPE NAT IPs Across a Group of Aggregators

Since the public IPs are associated with interfaces on specific instances, each bond must be assigned to a specific aggregator or routing will break if moved. However, as of Turnium SD-WAN version 6.5, bonds can move to other aggregators without needing to update the association — at least not right away.

Typical (Non-AWS) Deployment	AWS Deployment
CPE NAT IP routes are distributed across the main physical segment of the Aggregators to a local router via a protocol such as BGP, OSPF, or Babel.	AWS does not allow such peering, so addresses must be routed statically via the AWS API. Aggregators can, however, peer with each other using BGP so bonds can move between aggregators without breaking routing.

In version 6.5, custom routing protocols can be implemented to allow an aggregator with an associated IP to route to the correct aggregator.

Note that the AWS VPC does not allow multicast traffic, so OSPF and Babel will not work unless tunneled. Additionally, CPE NAT IPs are always exposed in the main routing table, so the protocols should not have a space defined.

Deploying Turnium SD-WAN Private WAN in AWS

Private WAN can also be deployed in AWS, but the lack of layer-2 networking means private WAN traffic must transit across VXLANs to get to other aggregators, using **Managed or UnManaged Mesh** private WAN modes.

Since private WAN routing needs to be isolated, a unique VPC — created in the AWS Console — is required for each customer space.

Since private WAN routing needs to be isolated, a unique VPC — created in the AWS Console — is required for each customer space.

Managed mesh mode is easiest to implement, since it automatically detects which aggregators have bonds in a space and sets up VXLANs between them. However, these VXLANs will transit between the aggregators via their public IP addresses, so encapsulated traffic will have to hairpin at the AWS NAT gateway.

To configure Managed Mesh, select the Managed Mesh mode on the space's Private WAN tab.

Optional: Add Direct VXLAN Peering

It's also possible to add direct VXLAN peering between aggregators in AWS. This is optional for managed mesh, but **required** if unmanaged mesh is deployed. Direct VXLAN peering is useful if large amounts of traffic are being routed over the private WAN.

- Since the private IP addresses of the aggregators are known, a more direct path between aggregators can be provided by adding custom VXLANs within the AWS VPC.
- A unique VNI should be created for each space. Each VXLAN for the space should use that assigned VNI. Failure to do so will result in broken routing or unexpected cross routing.
- Be sure to choose a unique subnet within the private WAN to assign to the VXLAN interfaces. Each aggregator will need a unique address in the subnet.
- We recommend using Babel as the routing protocol. OSPF will work, but Babel is a more modern protocol, better suited to this kind of application and is simpler to implement.

Adding a Gateway to Private WAN in AWS

Typical (Non-AWS) Deployment	AWS Deployment
In a typical managed mesh or unmanaged private WAN deployment, gateways are assigned to VLAN interfaces connected to individual or multi-tenant routers or firewalls.	Since AWS does not support VLANs, a separate Elastic Network Interface (ENI) and an Elastic IP (EIP) is required for each gateway, and each gateway is assigned to a specific aggregator.

Note that AWS NAT gateways do not allow individual port forwarding. Hosts in a Private WAN deployed in AWS will not be able to act as servers to provide services to public Internet hosts. If this is required, separate NAT gateway appliances will need to be installed and configured.

Deploying Turnium SD-WAN Aggregators in IBM Cloud

IBM Cloud can support Turnium SD-WAN Aggregators using either IBM's Bare Metal Server or Virtual Server compute services. Turnium's Managed Mesh Private WAN mode is ideal for facilitating transit between bonds in a customer space, regardless of which region their respective Aggregators are located.

IBM Cloud offers two forms of secondary subnets that can be used for connected IPs, CPE NAT IPs, or routes:

Static	Portable
IP addresses are available to the resource identified as the routing endpoint.	IP addresses are available to all resources on a VLAN.
Routed to the specified resource (ie: an aggregator) and then will be further routed through the SD-WAN tunnel to the bond.	Can be used by any resource on a VLAN which, when combined with the proxy ARP setting on an aggregator, allows failover to work within a specified IBM Cloud VLAN.
Since the routing is statically defined, aggregator failover will not be effective.	For aggregator failover to be effective with proxy ARP, all primary/secondary aggregators must be assigned to the same VLAN in IBM Cloud.

Provisioning Turnium SD-WAN Aggregators in IBM Cloud

Provisioning Aggregators in IBM Cloud follows the standard deployment process for aggregators.

Managed Mesh Private WAN

Turnium's Managed Mesh Private WAN mode is easiest to implement, since it automatically detects which aggregators have bonds in a space and sets up VXLANs between them.

To configure this, just select the Managed mesh mode on the space's Private WAN tab. After doing so, all private WAN bonds in the space should have connectivity.

Optional: Add Direct VXLAN Peering

It is possible to provide a more direct path between aggregators deployed in IBM Cloud by adding custom VXLANs to the involved aggregators within the IBM Cloud VPC. This may only be necessary if you are passing large amounts of private WAN traffic between aggregators.

- A unique VNI should be created for each space. Each VXLAN for the space should use that assigned VNI. Failure to do so will result in broken routing or unexpected cross routing.
- Choose a unique subnet within the private WAN to assign to the VXLAN interfaces. Each aggregator will need a unique address in the subnet.
- Babel should be used as the routing protocol. OSPF will also work, but Babel is a more modern protocol that is better suited to this kind of application and simpler to implement.

Note that VLAN Spanning will need to be enabled in the IBM Cloud Dashboard, VLANs section, to enable aggregators on separate VLANs to communicate over the Private WAN.

Deploying Turnium SD-WAN Aggregators in Microsoft Azure

Microsoft Azure is a popular cloud computing platform similar to Amazon Web Services that allows for quick deployment and management of Internet-connected servers. It is most popular for hosting websites, but can be used to deploy Turnium SD-WAN aggregators as well, with some caveats:

- Currently, aggregators are only supported using Debian 8 Jessie. When creating the instance, ensure the Debian 10 image is used.
- There is no option for routing subnets to hosts -- only single IPs can be allocated and those IP addresses are randomly assigned, so it is not possible to even emulate subnet routing with IPv4. (It can be emulated with IPv6 addresses, but each address in each range will need to be assigned to the Aggregator interface, which is burdensome.)
- Public IPv4 addresses in Azure are never assigned directly to instance interfaces. All public IPv4 addresses are assigned to a private address and translated using NAT. While aggregators are capable of handling tunnels in such a scenario, it makes bond routing complicated

CPE NAT IP to Individual Bonds

What's Involved	Typical (Non-AWS) Deployment
This scenario involves providing individual per-bond access to the Internet, each with its own public IP address.	A subnet containing IPs available for use for CPE NAT IPs is routed to the aggregator network and the aggregator announces each IP via OSPF or iBGP. If dynamic routing is not available, static routes can be made to individual aggregators hosting the bonds, but this requires manual routing changes if a bond is moved.

Each public IP address is assigned to a private IP address that is configured on an interface. By default, when an AWS EC2 instance is created:

- A network interface is created
- A dynamically chosen private IP address is assigned to and configured on the interface
- A public IP address is associated with the private IP address

An aggregator can use this address for its primary communication, but each CPE NAT IP will need additional addresses, which can be provided through Azure public IP addresses using the Azure portal under the Networking menu for the virtual machine running the selected aggregator.

Note that it is not necessary to configure the private or public IP address in the Aggregator's network configuration in the Turnium SD-WAN Management Server.

Configuring the CPE NAT IP

On a bond assigned to the aggregator, configure a private connected IP and add the CPE NAT IP using the private IP address that was added to the interface. The Azure gateway will NAT the public IP address to the private CPE NAT IP, and the bonder will NAT the private IP address to the connected IP.

Assuming the bonder is online, all of the routing will be in place.

Note that due to the lack of dynamic routing for public IP addresses in Azure, moving bonds between aggregators will not work unless the static associations are updated using the Azure Portal or API. Aggregator failover will not work in this scenario.

Route CPE NAT IPs Across a Group of Aggregators

Since the public IPs are associated with interfaces on specific instances, each bond must be assigned to a specific Aggregator or routing will break if moved. However, as of Turnium SD-WAN version 6.5 or later, bonds can move to other Aggregators without needing to update the association — at least not right away.

Typical (Non-AWS) Deployment	Azure Deployment
CPE NAT IP routes are distributed across the main physical segment of the Aggregators to a local router via a protocol such as BGP, OSPF, or Babel.	Azure does not allow such peering, so addresses must be routed statically via the Azure API. Aggregators can, however, peer with each other using BGP so bonds can move between Aggregators without breaking routing.

In version 6.5, custom routing protocols can be implemented to allow an aggregator with an associated IP to route to the correct aggregator.

Note that the Azure virtual networks do not allow multicast traffic, so OSPF and Babel will not work unless tunneled. Additionally, CPE NAT IPs are always exposed in the main routing table, so the protocols should not have a space defined.

Deploying Turnium SD-WAN Private WAN in Azure

Private WAN can also be deployed in Azure, but the lack of layer-2 networking means private WAN traffic must transit across VXLANs to get to other Aggregators, using Managed or Unmanaged Mesh private WAN modes.

Note that Azure does not allow Aggregators to connect to multiple virtual networks, even with additional interfaces. To get around this, you must set up VXLAN peering from the Aggregator to a virtual machine hosted within the target resource group.

Due to the complications entailed in deploying this design, we recommend simply placing dedicated Aggregators for each customer into their own Azure resource groups, disabling Private WAN and setting up peering between Aggregators via BGP and static routing.

If you need to extend a space into a network to integrate with existing Aggregators -- or to use larger Aggregators for multiple spaces -- a new resource group will need to be added for the space in Azure and VXLAN endpoints and interfaces will need to be created.

Adding a VXLAN Endpoint for Turnium SD-WAN in Azure

To add a VXLAN endpoint for Turnium SD-WAN in Azure, the endpoint must be created in the target resource group. (This machine can be any Linux distribution or router with VXLAN support.)

The easiest option is to simply add an Aggregator to a private WAN in Managed Mesh mode, but without adding any bonds. The main interface will be used for the VXLAN traffic, while an additional interface will carry the internal private WAN traffic.

If the goal is to provide public internet access for the private WAN being deployed, using a bare Linux distribution or dedicated router distribution such as VyOS will enable VXLAN termination, Azure network routing and the ability to NAT the private WAN networks.

Since an Aggregator is being deployed to act as a VXLAN endpoint -- and therefore will not handle bonds -- this endpoint does not have the same CPU and memory requirements of a normal Aggregator, so any small instance should suffice.

It's important to remember that both the VXLAN endpoint and VXLAN peer need to have interfaces to each other created. The interface for the VXLAN peer can be created in the Turnium Management Server, if the VXLAN peer is being created using an Aggregator.

Routing Between Azure Virtual Networks and Turnium SD-WAN Private WAN Networks

Once VXLAN peer and endpoints are created with interfaces to each other, routes must be created on the private WAN and on the Azure portal, to ensure hosts can communicate.

The benefit of using Turnium Aggregators as the VXLAN endpoint in Azure is that it automatically injects the route for the subnet the interface's IP is in.

- If all hosts the Private WAN needs to access are within that subnet, no additional routes need to be configured in the Turnium Management Server
- If there are additional subnets in Azure that need to be accessed, a static protocol will need to be created to inject those routes
- Additionally, routes for the bonds in the private WAN will need to be added to the Azure portal

Optional: Adding Direct VXLAN Peering

It's also possible to add direct VXLAN peering between aggregators in Azure. This is optional for managed mesh, but **required** if unmanaged mesh is deployed or if large amounts of traffic are being routed over the private WAN.

- Since the private IP addresses of the aggregators are known, a more direct path between aggregators can be provided by adding custom VXLANs within the AWS VPC.
- A unique VNI should be created for each space. Each VXLAN for the space should use that assigned VNI. Failure to do so will result in broken routing or unexpected cross routing.
- Be sure to choose a unique subnet within the private WAN to assign to the VXLAN interfaces. Each aggregator will need a unique address in the subnet.
- We recommend using Babel as the routing protocol. OSPF will work, but Babel is a more modern protocol, better suited to this kind of application and is simpler to implement.

Deploying Hybrid L2-L3 Networks with Turnium SD-WAN

Use-Cases for Hybrid L2-L3 Networks

When deploying SD-WAN networks, you will often need to provide customers with a migration path to move from existing Layer 2 Ethernet or MPLS networks to a full SD-WAN deployment. In some cases, customers may also prefer to retain Ethernet/MPLS circuits for some sites and request a hybrid deployment that integrates both Layer 2 and Layer 3 SD-WAN technologies.

If a Hybrid deployment is required (either as a migration tactic or as a permanent solution), the advantage of deploying SD-WAN is that it provides a single, unified, virtual managed network that can incorporate Layer 2 circuits.

Deploying Hybrid L2-L3 Networks

When deploying Turnium SD-WAN for customers, you may need to use existing MPLS/L2 circuits as last-mile "legs" attached to CPE or Bonders at a client site.

For this work, the MPLS circuit must connect to a Turnium Aggregator's public IP address in the same manner as broadband, wireless or other circuits.

In this case, the following assumptions are made:

- The Aggregator has an interface to the MPLS network.
- The Aggregator's public IP is not behind NAT.
- Adding a route through dynamic routing to the MPLS network is possible.

Specifically, in this instance, the Aggregator has two interfaces, an eth0 with a public IP 203.0.113.1/24 and an eth1 on the MPLS network with an IP of 10.100.100.1/24. In the Aggregator record on the Turnium Management Server, the Aggregator will have the IP configured as 203.0.113.1. This (203.0.113.1) is the IP to which any circuits defined on bonds connecting to this Aggregator will connect.

To begin the integration with the MPLS/L2 network, a new ethernet interface for eth1, with the 10.100.100.1/24 IP defined as an IP of this interface is needed.

In this example, there is a basic OSPF setup for dynamic routing on the MPLS network that will be used on the Aggregator to get a route for the Aggregator's public IP made available through the MPLS network.

Step 1: Add a Filter

The first step is to add a filter to ensure that only a single IP will be exported into the MPLS network. The goal is not to add anything beyond what is strictly needed to get the edge devices able to connect using their MPLS circuits.

This filter will be set up with two rules, the first with a match for the network of 203.0.113.1/32 with an action of Accept, and the second rule simply having an action of Reject so that only the Aggregator's public IP is allowed.

It is essential that the filters and networks in the protocols are for single network addresses with a /32 prefix length and not entire networks. This limits exposure to the MPLS network.

Step 2: Make an OSPF Protocol

The next step is to make an OSPF protocol. Set the IPv4 export to the filter that was previously defined. Optionally, IPv4 import networks can be filtered if the IPs that will be associated with the last-mile legs are known (and can be used as a filter).

The rest of the OSPF protocol depends on the OSPF settings in the MPLS network but ensure that the interface for the MPLS network is chosen.

Step 3: Advertise Aggregator IP Address

Next, advertise the Aggregator IP address by adding a static protocol with the Aggregator IP as an interface route. In this example, the route will look like having a network of 10.87.100.51/32, a destination of Interface and an interface of eth0.

With these protocols and filters in place, the OSPF should connect, and any legs on MPLS circuits on this Aggregator should now be able to connect.

Routes for the MPLS network will be installed on the Aggregator and routable to all tunnels. However, the return path likely won't work. This can be mitigated by adding more filters if IPs used by legs are known and those routes are imported on the Aggregator.

The Turnium Mission: Connect Everything, Anywhere

Over 300 tier 1 and 2 channel partners across five continents rely on Turnium's SD-WAN software platform to connect their customers and IoT devices to cloud services, existing networks, branch offices and data centers globally.

Want to learn more? **Contact us** today with your questions or **book a demo** now.

About Turnium

Turnium Technology Group, Inc. delivers its disaggregated networking solution as a white-label platform that channel partners brand, host, manage, and price. Turnium SD-WAN transforms legacy networks, enabling new services such as managed hybrid network, hybrid cloud, uCPE, Edge Compute, fixed-wireless convergence, mobile-mobile convergence, seamless failover, and IoT networks. For more information, please visit **www.turnium.com** or follow us on Twitter @turnium.

Turnium Technology Group, Inc.
www.turnium.com
Contact Sales: +1.604.398.4314 option 1
or North America Toll-Free: +1-888-818-3361 option 1
or Email: **sales@ttgi.io**

LOCATIONS

Vancouver (Head Office)
Turnium Technology Group, Inc.
1127 West 15th Street
North Vancouver, British Columbia
V7P 1M7

Halifax
Turnium Technology Group, Inc.
331-1496 Lower Water Street