Lumen® SD-WAN

SD-WAN branch in AWS overview
August 2023

General disclaimer

Although Lumen has attempted to provide accurate information in this guide, Lumen does not warrant or guarantee the accuracy of the information provided herein. Lumen may change the programs or products mentioned at any time without prior notice. Mention of non-Lumen products or services is for information purposes only and constitutes neither an endorsement nor a recommendation of such products or services or of any company that develops or sells such products or services.

ALL INFORMATION PROVIDED IN THIS DOCUMENT IS PROVIDED “AS IS,” WITH ALL FAULTS, AND WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESSED OR IMPLIED OR STATUTORY. LUMEN AND ITS SUPPLIERS HEREBY DISCLAIM ALL WARRANTIES RELATED TO THIS GUIDE AND THE INFORMATION CONTAINED HEREIN, WHETHER EXPRESSED OR IMPLIED OR STATUTORY INCLUDING, WITHOUT LIMITATION, THOSE OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NON-INFRINGEMENT, OR ARISING FROM A COURSE OF DEALING, USAGE, OR TRADE PRACTICE.

LUMEN AND ITS SUPPLIERS SHALL NOT BE LIABLE FOR ANY DIRECT, INDIRECT, SPECIAL, CONSEQUENTIAL, OR INCIDENTAL DAMAGES INCLUDING, WITHOUT LIMITATION, LOST PROFITS OR REVENUES, COSTS OF REPLACEMENT GOODS OR SERVICES, LOSS OR DAMAGE TO DATA ARISING OUT OF THE USE OF THE GUIDE OR ANY LUMEN PRODUCT OR SERVICE, OR DAMAGES RESULTING FROM USE OF OR RELIANCE ON THE INFORMATION PROVIDED IN THIS GUIDE, EVEN IF LUMEN OR ITS SUPPLIERS HAVE BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

Any Internet Protocol (IP) addresses and other information used in this document are not intended to be actual addresses and phone numbers. Any examples, command display output, network topology diagrams, and other figures included in the document are shown for illustrative purposes only. Any use of actual IP addresses or phone numbers in illustrative content is unintentional and coincidental.

Many of the Lumen products and services identified in this guide are provided with, and subject to, written software licenses and limited warranties. Those licenses and warranties provide the purchasers of those products with certain rights. Nothing in this guide shall be deemed to expand, alter, or modify any warranty or license or any other agreement provided by Lumen with any Lumen product, or to create any new or additional warranties or licenses.

Overview

This document provides an overview of the steps a customer will need to perform in the customer owned AWS environment in the support of a Lumen SD-WAN branch VM deployment. The customer will also need to provide several pieces of information back to Lumen to facilitate the deployment.

Topics covered in this document:

- Customer cloud infrastructure in AWS.
- AWS root account ID and private images.
- API user in IAM.
- CloudFormation template to simplify customer AWS setup.
- Structure for single VM or a dual VM high availability design.

Customer is required to have their own AWS infrastructure account and will have to perform all AWS steps to support the VM deployment. Customer account should have, but not limited to, a VPC with associated CIDR block, at least three subnets in the same availability zone, internet gateway attached to two of these subnets, security groups, a key pair to be used on the appliance, and route tables as required for the three subnets. If the customer also requires Lumen MPLS connectivity, they will need to create a VPN gateway, attach it to the transit VPC, accept the connection and create the virtual interface and associated BGP session.

NOTE: Lumen will be able to provide the customer with a CloudFormation template described below that will cover the deployment of many of these requirements.

Design topologies

Our preferred deployment approach will be to establish a separate VPC to host the SD-WAN VMs in the customer AWS environment. Figure 1 below shows a brief overview of each deployment.

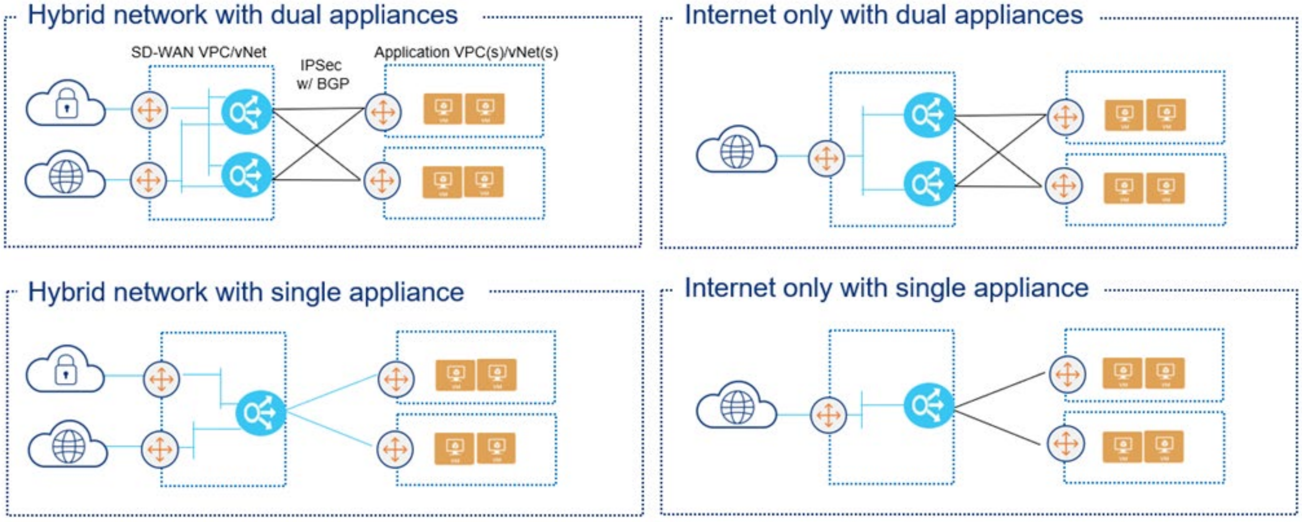
Hybrid network with dual appliances
SD-WAN VPC/vNet
IPSec w/ BGP
Application VPC(s)/vNet(s)
Internet only with dual appliances
Hybrid network with single appliance
Internet only with single appliance

Application VPC(s)/vNets can be in the same or different regions as the SD-WAN VPC/vNet
AWS Transit Gateways can be used instead of VGWs in each VPC

Figure 1: Cloud SD-WAN deployment topologies

CloudFormation templates

Lumen can provide the customer with CloudFormation templates for each of the design topologies in Figure 1. These templates will create the VPC, subnets, /24 CIDR block, route tables, routes, and security groups to support the deployment of the VM. Below is a summary of the 4 template options that the Lumen TDE can provide to the customer. These will be referred to later as a step in the process to follow.

1. AWS Single SDWAN-HA INET.json: Creates a VPC, three subnets (MGMT, INET and LAN) using a single /24 CIDR block, associated IGW, route tables, routes, and security groups to support deployment of a single SD-WAN appliance within a region supporting Internet connectivity only. For an HA configuration, this template must be run in a second region.
2. AWS Dual SDWAN-HA INET.json: Creates a VPC, six subnets (MGMT, INET and LAN primary and secondary) using a single /24 CIDR block in a redundant fashion using availability zones, associated IGW, route tables, routes and security groups to support deployment of dual SD-WAN appliances within a single region supporting Internet connectivity only.

3. AWS Single SDWAN-HA INET-MPLS.json: Creates a VPC, three subnets (MGMT, INET and LAN) using a single /24 CIDR block, associated IGW, route tables, routes, and security groups to support deployment of a single SD-WAN appliance within a region supporting internet and MPLS connectivity. Template requires a secondary /28 CIDR for the MPLS subnet and this CIDR will need to be within the 100.88.0.0/14 range to support native controller reachability. For an HA configuration, this template must be run in a second region which will require a second IP assignment out of the 100.88.0.0/14 range to support native controller reachability. This template does not create a VPN or transit gateway for MPLS support. The customer will need to create and attach one as well as configure any connectivity to the MPLS cloud.
4. AWS Dual SDWAN-HA INET-MPLS.json: Creates a VPC, six subnets (MGMT, INET and LAN primary and secondary) using a single /24 CIDR block in a redundant fashion using availability zones, associated IGW, route tables, routes, and security groups to support deployment of dual SD-WAN appliances within a single region supporting Internet and MPLS connectivity. Template requires a secondary /28 CIDR for the MPLS subnet and this CIDR will need to be within the 100.88.0.0/14 range to support native controller reachability. For an HA configuration, this template must be run in a second region which will require a second IP assignment out of the 100.88.0.0/14 range to support native controller reachability. This template does not create a VPN or transit gateway for MPLS support. The customer will need to create and attach one as well as configure any connectivity to the MPLS cloud.

Deployment requirements and overview

Customer prerequisites

- Customer must have an existing cloud infrastructure account.
- Customer will be required to provide the Lumen TDE with the AWS root account ID. This is necessary to load the Versa software image under Private Images.
- Customer must create an API user in IAM and add the user to a group to allow access to selected resources in the appropriate VPC. API user is necessary to create the VM from the Versa Director. Customer can disable API user after successful activation of the VM. Customer can review the AWS IAM creation process at the following link: https://docs.aws.amazon.com/IAM/latest/UserGuide/id_users_create.html
- The customer will be required to provide the following information from the creation of the API user:
    - Access key ID
    - Secret key
- Customer will be required to provide the additional information below for each SD-WAN VM image that will be deployed:
    - AWS region
    - VPC network name. *Created with the CloudFormation templates.
    - Key pair name, customer created.
    - Availability zone
    - Available image AMI IDs for SD-WAN Flex VNF(software image) after added in Private Images in AWS.
    - Confirm subnet name and security group name for management, WAN, and LAN interfaces. *Created with the CloudFormation templates.
    - Any IP information for interfaces not participating in DHCP and/or BGP neighbor information, if required. *Uncommon.

NOTE: Some of the information confirmed above will be possible to be created by CloudFormation templates that will be covered later in this document.

Overview of Lumen deployment steps

- Lumen TDE will ensure the correct version of the FlexVNF image is loaded in the Private Images in the customer AWS account.
- Lumen will create a CMS connector to the customer AWS environment to support deployment of the Versa VMs.

- Lumen will continue with completing the deployment templates to build and support the activation of the VM(s).

VM sizing options

The following table represents the VM sizes that are available and are standard AWS VM sizes. The size of the VM chosen should be based on the desired throughput and interfaces required. Customer can shut down the VM and re-size the instance and re-enable.

NOTE: Customer will start accruing additional charges directly from their AWS account based on the size and usage of the VM that is deployed based on their own AWS account guidelines. These charges are in addition to the Lumen SD-WAN service contracted through Lumen. AMD CPUs are not supported currently.

NOTE: AWS limits the number of specific EC2 instances on a per region basis. The customer must verify the deployment of two EC2 instances will not exceed their limits in the deployment region, or the deployment will fail. To check the limits on the AWS portal, navigate to the **EC2 Console** and select **Limits**.

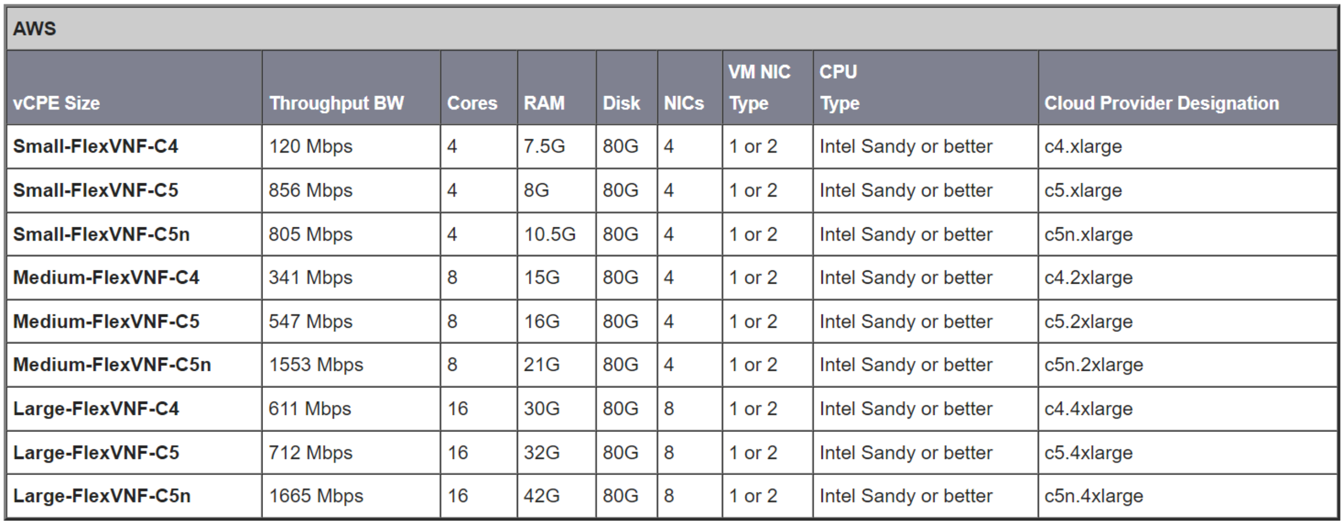

AWS

						VM NIC	CPU	
vCPE Size	**Throughput BW**	**Cores**	**RAM**	**Disk**	**NICs**	**Type**	**Type**	**Cloud Provider Designation**
Small-FlexVNF-C4	120 Mbps	4	7.5G	80G	4	1 or 2	Intel Sandy or better	c4.xlarge
Small-FlexVNF-C5	856 Mbps	4	8G	80G	4	1 or 2	Intel Sandy or better	c5.xlarge
Small-FlexVNF-C5n	805 Mbps	4	10.5G	80G	4	1 or 2	Intel Sandy or better	c5n.xlarge
Medium-FlexVNF-C4	341 Mbps	8	15G	80G	4	1 or 2	Intel Sandy or better	c4.2xlarge
Medium-FlexVNF-C5	547 Mbps	8	16G	80G	4	1 or 2	Intel Sandy or better	c5.2xlarge
Medium-FlexVNF-C5n	1553 Mbps	8	21G	80G	4	1 or 2	Intel Sandy or better	c5n.2xlarge
Large-FlexVNF-C4	611 Mbps	16	30G	80G	8	1 or 2	Intel Sandy or better	c4.4xlarge
Large-FlexVNF-C5	712 Mbps	16	32G	80G	8	1 or 2	Intel Sandy or better	c5.4xlarge
Large-FlexVNF-C5n	1665 Mbps	16	42G	80G	8	1 or 2	Intel Sandy or better	c5n.4xlarge

* Throughput based on unidirectional streams and are estimates. Throughput values will vary based on time of testing and load on the cloud instances. Single flow testing or single packet size testing was not performed. Tests were based on a mix of traffic.

Internet-only deployments

In these designs, the customer only requires internet WAN connectivity into their cloud environment. Figure 2 below shows an overview of the single VM or dual VM deployment topologies.

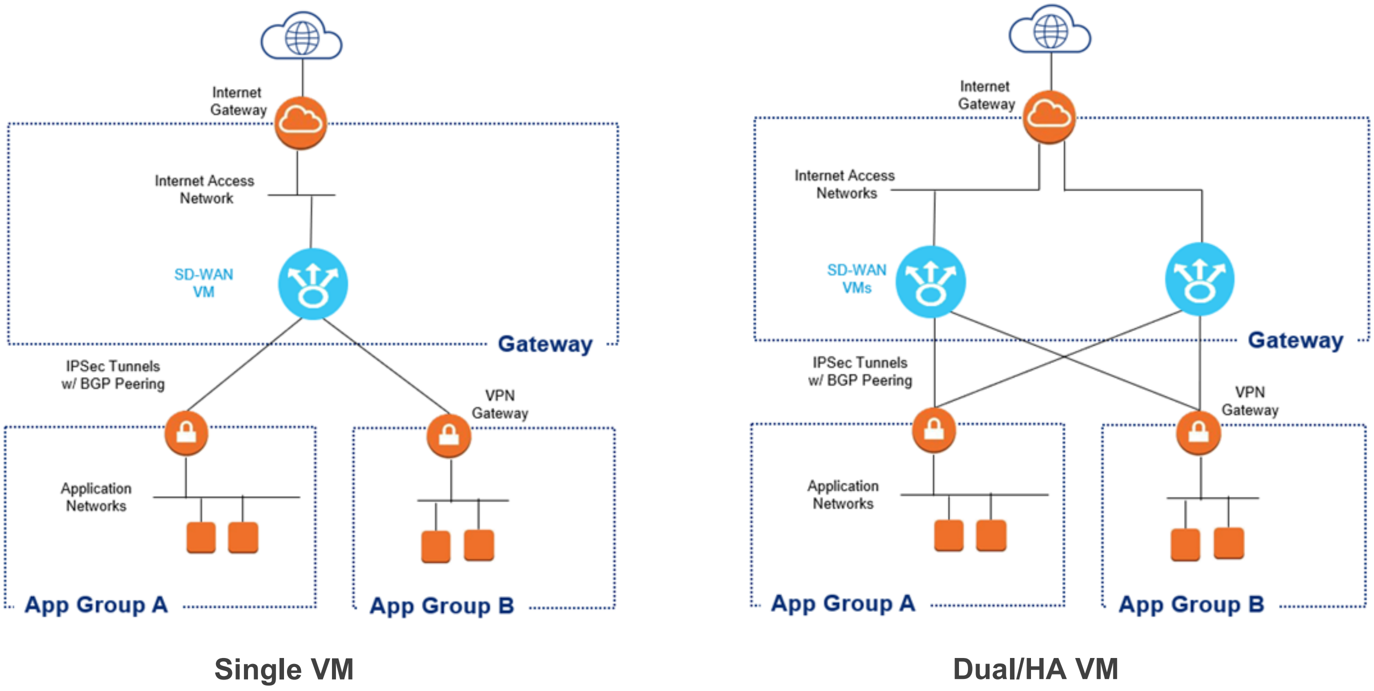

Figure 2: Internet Only Deployments

Summary of customer steps for deployment

For the Internet only deployment topologies, below is a summary of the steps required by the customer to complete this configuration/deployment.

1. Request the appropriate CloudFormation template from the Lumen TDE and run them in the customer AWS environment. Template #1 for a single VM in an AWS region (this can be run twice in if the design is for a single VM in two AWS regions, once per region). Template #2 for a dual VM design in the same AWS region.

   NOTE: Gateway VPC requirements are also covered in the appendix of this document for reference.

2. HOLD step – Customer will wait until Lumen has deployed the SD-WAN VM instances to proceed further. Steps are listed above in the overview of Lumen deployment steps.

3. Create customer gateways – The customer will need to create a customer gateway for each SD-WAN VM instance.

   NOTE: Customer can go ahead and create the transit gateway connect before the SD-WAN VM instance is deployed in step #2 above, and provide the GRE source and destination and BGP peer information to Lumen. (See **Reference: Transit Gateway Connect Deployment.)**

4. Create and attach a gateway – If the customer has a single host VPC and will be using the SD-WAN VM instance in active-backup mode (for HA deployments), a VPN gateway can be used. Lumen recommends the use of a transit gateway for flexibility and growth. **– OR –** If the customer has or plans to have multiple VPCs, VPC to VPC peering, or will be using the SD-WAN VMs in active-active mode, a transit gateway must be used. See the **Reference: AWS Gateway Types** section of this document for overview and steps on creating the gateways. **NOTE**: Customer should create either a VPN gateway (VGW) or transit gateway (TGW) at this step.

5. Create VPN connection or GRE connection – There will be two options depending on whether the customer is using a VPN gateway or transit gateway. See the **Reference: VPN Connections (IPSec tunnels)** – OR – **Reference: Transit Gateway Connect Deployment** section of this document for instructions on creating the tunnels/connections.

6. After information from the VPN connections or GRE connections are provided back to Lumen SD-WAN operations (from step 5), Lumen engineers will complete configuration steps on the SD-WAN VMs.

Hybrid deployments with MPLS

In these designs, the customer requires both Internet and MPLS connectivity into their cloud environment. Figure 3 below shows an overview of the single VM or dual VM deployment topologies.

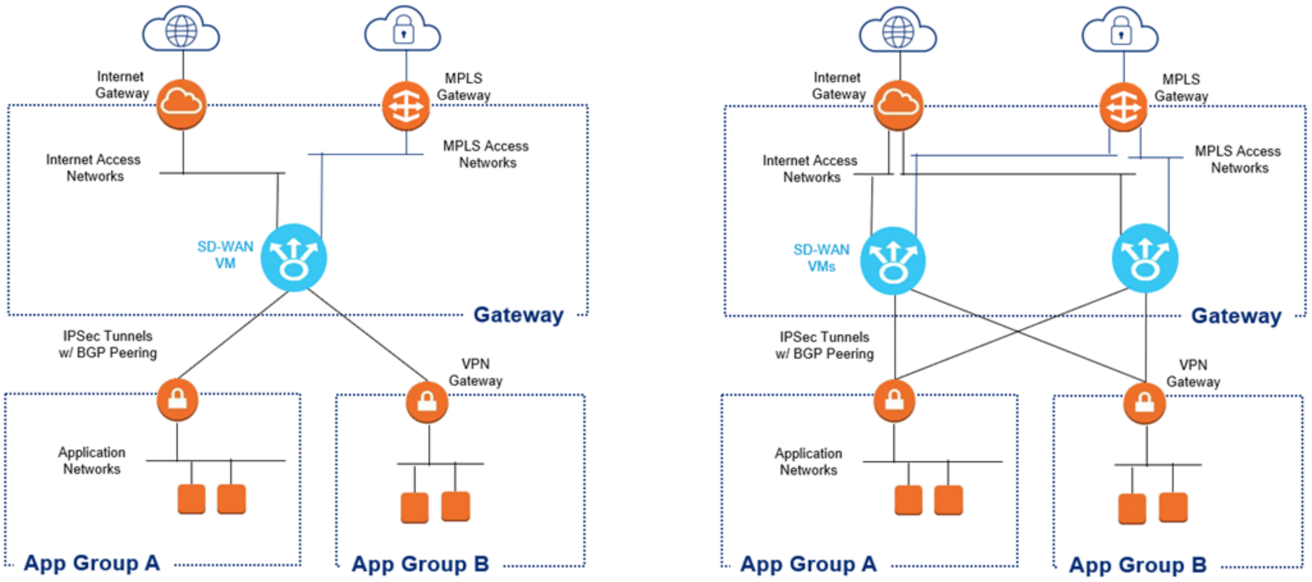

Single VM

Dual/HA VM

Figure 3: Hybrid Internet and MPLS deployments

Summary of customer steps for deployment

Like the internet-only deployment topologies, the customer will need to do many of the same steps and also additional steps related to setting up and attaching the MPLS gateway. Below is a summary of the steps required by the customer to complete this configuration/deployment.

1. Request the appropriate CloudFormation template from the Lumen TDE and run them in the customer AWS environment. Template #3 for a single VM in an AWS region (this can be run twice in if the design is for a single VM in two AWS regions, once per region). Template #4 for a dual VM design in the same AWS region.

   NOTE: Gateway VPC requirements are also covered in the appendix of this document for reference.

2. HOLD step – Customer will wait until Lumen has deployed the SD-WAN VM instances to proceed further. Steps are listed above in the overview of Lumen deployment steps.

Services not available everywhere. Business customers only. Lumen may change, cancel or substitute products and services, or vary them by service area at its sole discretion without notice.

3. Create customer gateways – The customer will need to create a customer gateway for each SD-WAN VM instance.

   NOTE: Customer can create the transit gateway connect before the SD-WAN VM instance is deployed in step #2 above, and provide the GRE source and destination and BGP peer information to Lumen. (See **Reference: Transit Gateway Connect Deployment.**)
4. Create and attach a gateway – If the customer has a single host VPC and will be using the SD-WAN VM instance in active-backup mode (for HA deployments), a VPN gateway can be used. However, Lumen recommends the use of a transit gateway for flexibility and growth. **– OR –** If the customer has or plans to have multiple VPCs, VPC to VPC peering, or will be using the SD-WAN VMs in active-active mode, a transit gateway must be used. See the **Reference: AWS Gateway Types** section of this document for overview and steps on creating the gateways. **NOTE**: Customer should create either a VPN gateway (VGW) or transit gateway (TGW) at this step.
5. Create VPN connections – There will be two options depending on whether the customer is using a VPN gateway or transit gateway. See the **Reference: VPN Connections (IPSec tunnels)** – OR – **Reference: Transit Gateway Connect Deployment** section of this document for instructions on creating the tunnels/connections.
6. After information from the VPN connections are provided back to Lumen SD-WAN operations (from step 5), Lumen engineers will complete configuration steps on the SD-WAN VMs.

Additional customer steps in AWS for MPLS connectivity

Connectivity between a customer's AWS Infrastructure and an MPLS service is supported. The following section details requirements and steps to connect to a Lumen service such as Lumen® Cloud Connect or VPNLynk. When a customer orders one of these services from Lumen, they will provide their AWS account ID and AWS region to provision the service to. When the service is provisioned on a Lumen PE, information is sent to AWS which includes the VLAN of the provisioned service and the AWS account to link it to. The customer must configure and attach a VPN gateway to the Transit VPC, accept the connection, create and attach a virtual interface (VIF) to the VGW, and create a BGP session between the VGW and Lumen PE. It is important to note, this VPN gateway is separate from the VPN or transit gateway providing access between the transit VPC and the customer's host VPCs and a transit gateway cannot be used for MPLS connectivity. The following steps must be completed by the customer in AWS Console for a hosted MPLS connection and assumes the CloudFormation template(s) supporting Internet and MPLS connectivity have been deployed in the customer's AWS instance.

1. To support Native Controller access, the customer will need to add the /28 IP block(s) from the 100.88.0.0/23 range as a secondary CIDR in the VPC. Lumen TDE will need to provide the customer with the /28 IP block(s) or two /28 IP blocks(s) in dual region deployment.
2. Customer must have a VPC with an associated CIDR block(s), four subnets, security groups for the management and Internet WAN subnets and route tables associated to each subnet. The above IP ranges from the 100.88.0.0/23 range will need to be used as the IP addresses for the MPLS subnet(s) connecting to the MPLS interface on the SD-WAN appliance. *Will be created by the CloudFormation template(s).

Services not available everywhere. Business customers only. Lumen may change, cancel or substitute products and services, or vary them by service area at its sole discretion without notice.

3. Customer must create and attach a VPN gateway (VGW) to the VPC, necessary for the MPLS connection. To create the gateway, select the appropriate region in the upper-right corner, navigate to the VPC Management Console and select Virtual Private Gateway in the left pane. Click the 'Create Virtual Private Gateway' button and provide a name for the VGW. An AS number must be associated to the VGW and MUST match the peer-as configured in the MPLS provider router (from Lumen provided cloud MPLS service). The default AWS AS number is 64512 and is typically not what is configured on the MPLS Provider. Select 'Custom ASN' and enter the peer-as configured on the MPLS Provider router. If the AS is not correctly assigned, the VGW must be deleted and a new one added. This can be a cumbersome process if the VIF and BGP session have been created as they will also need to be deleted and recreated. When ready, click the 'Create Virtual Private Gateway' which will create the VGW and return to the Virtual Private Gateway screen.

   NOTE: Only one VPN gateway is required per region and if the customer is deploying VMs in multiple regions, a VPN gateway would need to be created in each region.
4. Associate the VGW to the transit VPC by selecting only the newly created VGW, clicking the 'Actions' button and select 'Attach to VPC'. Select the transit VPC to attach the gateway to and click 'Attach'.
5. Modify the route table (associated to the subnet for the MPLS) in the transit VPC by selecting 'Route Tables' from the left pane, and selecting the appropriate routing table. It is recommended that a static default route be added under the 'Routes' tab with a next hop of the VGW. Route propagation must also be configured under the 'Route Propagation' tab. This allows routes learned by the VGW via the BGP session to the MPLS provider to be automatically installed in the route table for the MPLS subnet. It also allows routes in the VPC to be advertised to the MPLS provider via BGP.
6. The hosted connection must be accepted by navigating to the 'Direct Connect Console'. A list of connections based on the Cloud Connect or VPNLynk services ordered will appear. Click the link name to open the details of the connection, add tags if desired and click 'Accept' then click 'Confirm'. *May not apply for a dedicated connection.
7. A virtual interface and BGP session to the MPLS provider must be created. Click 'Create virtual interface' to open the 'Create virtual interface' console.
8. Select 'Private' for the type, and enter a VIF name, verify the correct connection is selected, select My AWS account as owner, select virtual private gateway, select the appropriate VGW in the dropdown, and enter the AS number of the MPLS provider router, which should be 3549.
9. Expand the 'Additional Settings' and confirm IPv4 is checked, enter the IP address and mask of the MPLS provider router and the IP address and mask of the gateway (IP of the BGP neighbor on the MPLS router), enter the BGP authentication key configured on the MPLS router, add any tags and click 'Create virtual interface'. Information from this step will need to be provided by Lumen.
10. It may take up to 5 minutes for routes to propagate. You can confirm the routes by checking the AWS MPLS route table or the VRF route table on the MPLS provider router.

Common deployment elements

Reference: AWS gateway types

An AWS VPC only allows communication between resources defined within the VPC. To facilitate communication from resources within a VPC to the Internet, MPLS or other VPCs, a gateway must be deployed and associated to the VPC. AWS supports four primary gateway types:

- Internet gateway (IGW) - provides access to the Internet from resources within a VPC and can only be attached to a single VPC. By default, all outgoing traffic is NAT'd using a one-to-many model. Elastic IPs (EIPs) can be associated to EC2 interfaces to allow a one-to-one NAT. Security groups must be applied to any interface with an EIP to provide security. A default route must be configured in the route table attached to a subnet with a next hop of the TGW to allow for Internet access. AWS bills a per-GB charge for egress traffic and EIP assigned.
- VPN gateway (VGW) - provides direct connectivity to an MPLS cloud or remote customer gateway using IPSec tunnels. A VPN gateway can only be attached to a single VPC and doesn’t support ECMP. This gateway type can also be used to establish connectivity to the FlexVNFs deployed in AWS. A deployment with multiple host VPCs would require a VPN gateway and four IPSec tunnels (two to each SD-WAN appliance) for each host VPC. AWS bills a per hour charge for each connected customer gateway (FlexVNF appliances) and a per-GB charge for traffic. There is no charge for the VGW or the attachment to the VPC. It is also important to note that routes from a VPN gateway are automatically propagated into the VPC it is attached to. AWS limits the number of propagated routes in a route table to 100 and this limit cannot be increased.
- Transit gateway (TGW) - can be attached to multiple VPCs and does support ECMP. Will allow traffic to pass directly from VPC to VPC. This is the recommended gateway to use. A deployment with multiple host VPCs would only require 1 TGW and associated IPSec tunnels to each SDWAN appliance. AWS bills a per hour charge for each VPC connection and each customer gateway connection (SDWAN appliances) as well as a per GB charge for traffic between VPCs or between VPC and VPN. A transit gateway can support up to 10,000 routes. These routes cannot be propagated directly into a VPC, so static routes in the host VPC to the transit gateway will be required.
- Transit gateway connect (TGW-Connect) - allows a standard transit gateway to be attached to the transit VPC and replaces the IPSec tunnels with GRE tunnels on the LAN side of the FlexVNF. An existing transit gateway can be configured to support TGW-Connect by adding a CIDR block to the TGW, creating the GRE tunnels and associated BGP peers.

Note: IPSec tunnels are created between the VGW/TGW and the SD-WAN appliances using the public internet WAN IP of each appliance. Additional per-GB charges will apply for traffic egressing the internet WAN of the appliances toward the VGW/TGW and on to the host VPCs.

The required type and number of gateways will depend on the deployment model used and the customer's existing AWS infrastructure. At a minimum, one IGW will need to be created and attached to the transit VPC in each region to allow for management access to the FlexVNFs and to provide connectivity to the internet transport domain. This gateway is created and attached, then the appropriate route table entries are created when the CloudFormation templates are used. No VPN or transit gateways are created by these templates. (See **Reference: Transit Gateway Connect Deployment**.)

Reference: VPN connections (IPSec tunnels)

- If the customer is using a VPN gateway (VGW), only VPN connections must be configured. The VGW is two separate, redundant VPN gateway devices, so the IPSec tunnels are created in pairs to each SD-WAN appliance.

  i. Select 'Site-to-Site VPN Connections' from the left pane and click 'Create VPN Connection'.

  ii. Enter a descriptive name tag, select the appropriate VGW from the list, select existing customer gateway and select the first SD-WAN appliance from the list. Ensure Dynamic Routing is checked. You may specify the IP addresses for use within the tunnels, but this must be within the 169.254.0.0/16 range. If you do not specify these IPs they will be randomly assigned by AWS within the same range. You can also specify a pre-shared key for use on each tunnel. If one is not specified, AWS will generate one unique for each tunnel. By default, multiple IKE/IPSec parameters are configured on the VGW. By selecting 'Edit Tunnel Options', you can disable the use of these parameters. Once both tunnels have been configured, click 'Create VPN Connection' to complete the process. Repeat these steps for the second SD-WAN appliance.

  iii. Once the VPN connections have been completed, select each one independently and click 'Download Configuration'. Chose 'Generic' for the vendor and save the file. Repeat this step for the second VPN connection. These files contain configuration information that is required to configure the IPSec tunnels on the SD-WAN appliances and will need to be shared with Lumen personnel.

- If the customer is using a transit gateway (TGW), the VPN connections will need to be attached to the TGW, associated, and propagated within the VGW Route table, like the process of associating and propagating a VPC. The TGW is two separate, redundant gateway devices, so the IPSec tunnels are configured in pairs to each SD-WAN appliance.

  i. Select 'Transit Gateway Attachment' from the left pane and click 'Create Transit Gateway Attachment'.

  ii. Select the appropriate TGW from the list and select VPN for the attachment type. Select 'Existing' and select one of the customer gateways from the list. Select 'Dynamic' for the routing options. When using the TGW, you can only specify the inside IPs and pre-shared keys for the tunnels. If you do not specify, these will be automatically generated by AWS. You cannot specify IKE/IPSec parameters. Click 'Create Attachment' to complete the VPN connection. Repeat this step for the second SD-WAN appliance.

  iii. Select 'Transit Gateway Route Table' in the left pane, select the appropriate VGW route table and select the 'Associations' tab in the lower pane. Click 'Create Association', select the first SD-WAN VPN connection, then click 'Create Association'. Repeat this step for the second SD-WAN VPN connection.

iv. Select the 'Propagations' tab in the lower pane and click 'Create Propagation'. Select the first SD-WAN VPN connection from the list and click 'Create Propagation'. Repeat this step for the second SD-WAN VPN connection.

v. Once the VPN connections have been configured, they will appear in the 'Site-to-Site VPN Connections' dashboard selected from the left pane. Select the first VPN connection and click 'Download Configuration', select generic and save the file. Repeat this step for the second VPN Connection. These files contain configuration information that is required to configure the VPN tunnels on the SD-WAN appliances and must be shared with Lumen personnel.

Reference: Transit gateway connect deployment

AWS has an option in the configuration of the transit gateway called transit gateway connect. This feature allows the transit gateway to be attached directly to the transit VPC (SD-WAN VPC) and employs the use of GRE tunnels with BGP from the LAN VR of the SD-WAN appliances deployed in the transit VPC. This eliminates the requirement for IPSec tunnels detailed in the **Reference: VPN connections (IPSec tunnels)** section above. This feature simplifies the SD-WAN appliance configurations removing the need to build templates with dummy values for IPSec tunnel and BGP configuration and replacing those values once the configuration of the transit gateway has been completed. The configuration of the Transit Gateway Connect and associated BGP peer connections can be completed by the customer prior to SD-WAN appliance deployment. The GRE tunnel source and destination and BGP peer information can be provided to Lumen personnel by the customer and used to configure the appliance during instantiation.

NOTE: Transit VPC, SD-WAN VPC, and gateway VPC are referring to the same VPC in the diagram below. **Figure 4.**

Use of transit gateway connect has the following caveats:

- Only supported on transit gateways. Standard virtual private gateways are not supported.
- Currently supported in N. Virginia[1], Oregon, N. California, and Ireland. Support in other regions is coming soon. (See AWS Transit Gateway Connect for the most recent information.)
- SD-WAN appliance pairs must be deployed in the same region transit VPC. Inter-region connectivity is not supported.
- A /24 CIDR block must be configured on the transit gateway.
- A single GRE tunnel from each SD-WAN appliance is created with two BGP neighbors.

The following diagram depicts an HA deployment using Transit Gateway Connect in AWS.

AWS SDWAN High Availability Topology using Transit Gateway Connect

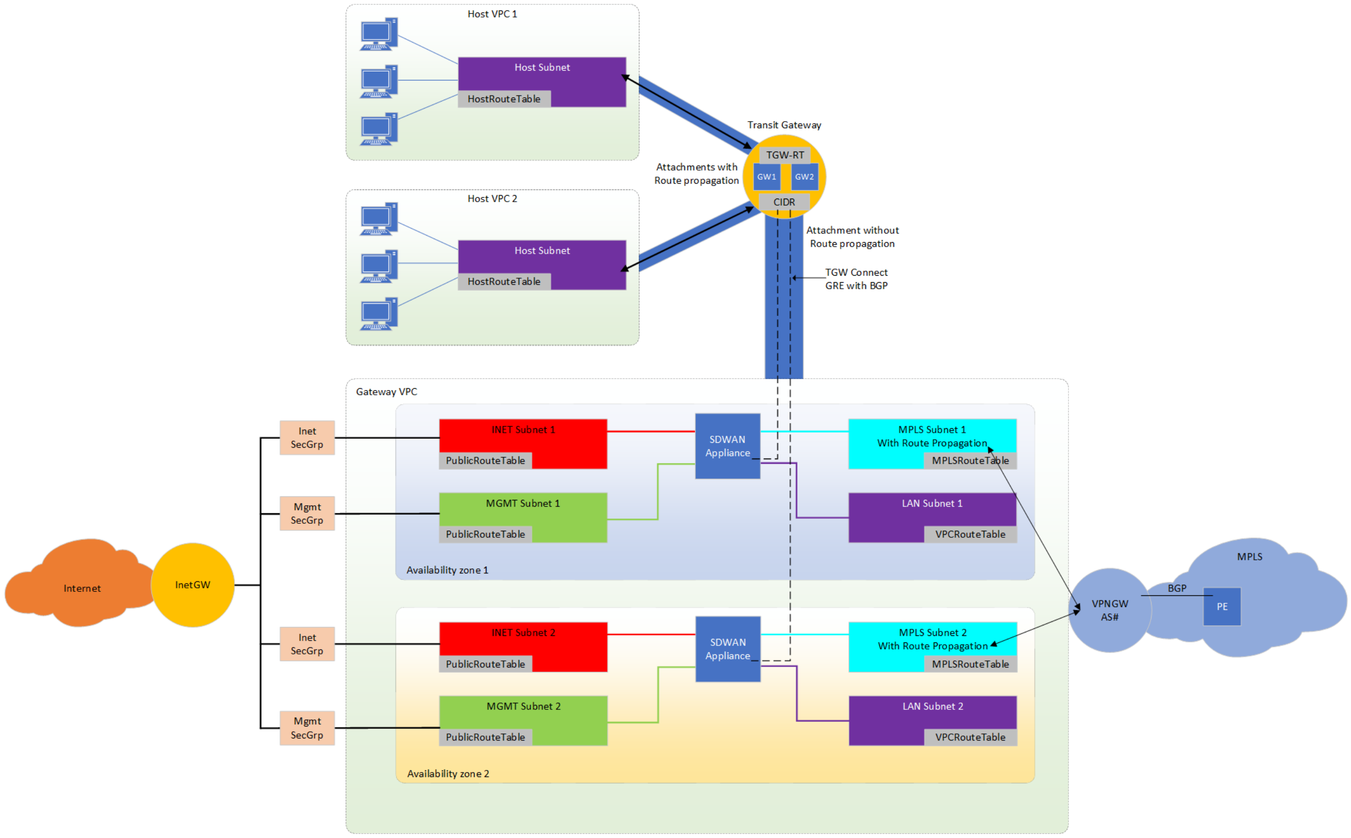

Figure 4: Transit Gateway Connect

Transit VPC(Gateway VPC) requirements

The transit VPC configuration and AWS Infrastructure requirements are the same as the IPSec deployment model with a few additions detailed below. The current CloudFormation templates can be used to deploy the basic infrastructure for the transit VPC. Only the following CloudFormation templates are supported:

- AWS dual SDWAN-HA Inet.json
  - Creates a VPC, six subnets (MGMT, INET and LAN primary and secondary) using a single /24 CIDR block in a redundant fashion using availability zones, associated IGW, route tables, routes and security groups to support deployment of dual SD-WAN appliances within a single region supporting internet connectivity only.
- AWS dual SDWAN-HA Inet-MPLS.json

- Creates a VPC, eight subnets (MGMT, INET, MPLS and LAN primary and secondary) using a single /24 CIDR block in a redundant fashion using availability zones, associated IGW, route tables, routes and security groups to support deployment of dual SD-WAN appliances within a single region supporting internet and MPLS connectivity. To support native controller reachability, two IP address assignments from the 100.88.0.0/14 range are required. This template does not create a VPN or transit gateway for MPLS support. The customer will need to create and attach one as well as configure any connectivity to the MPLS cloud.

Customer AWS configuration steps

Once the CloudFormation Template has been deployed in the customer's AWS instance, the following steps will be required to prepare for FlexVNF appliance deployment:

- 1. A transit gateway is required. If the customer does not have one created, they must create one and attach it to all host VPCs requiring connectivity to the SD-WAN overlay.

  Navigate to the **VPC console** and select **Transit Gateways**. Select **Create Transit Gateway** and provide a name tag, description, and Amazon side ASN. This ASN will be used as the peer AS in the FlexVNF appliance configuration.
- 2. The customer must associate a /24 CIDR block to the transit gateway.

  Navigate to the **VPC console** and select **Transit Gateways**. The CIDR block can be added as part of a new transit gateway creation or by selecting an existing transit gateway, clicking **Actions** and selecting **Modify**. This CIDR block must be added to the default route table of the transit VPC in step 8 below.
- 3. The transit gateway must be attached to the transit VPC.

  Navigate to the **VPC console** and select **Transit Gateway Attachments**. Click **Create Transit Gateway Attachment**, select the transit gateway ID, attachment type of **VPC** and select the VPC ID of the transit VPC. Add an attachment name tag (optional). Select the primary and secondary LAN subnets to attach to.
- 4. A transit gateway connect attachment to the transit VPC is also required.

  Navigate to the **VPC console** and select **Transit Gateway Attachments**. Click **Create Transit Gateway Attachment**, select the Transit Gateway ID, attachment type of **Connect** and select the Transit VPC attachment. Add an attachment name tag (optional).
- 5. Connect peers must be created for each FlexVNF appliance. This step will also create two BGP neighbors for each connect peer (FlexVNF appliance). This step must be repeated to create the peering for each FlexVNF appliance.

  Navigate to the **VPC Console** and select **Transit Gateway Attachments**. Select the **Connect Attachment** in the upper pane, select the **Connect Peers** in the lower pane and click **Create Connect Peer**. The following parameters are required:

    - Transit gateway GRE address can be auto assigned or populated manually. This address will be used as the SD-WAN appliances GRE tunnel destination.

  - Peer GRE address is the IP address of the LAN port on the SD-WAN appliance.
  - BGP inside CIDR blocks IPv4 is the IP address range that the BGP sessions between the SD-WAN appliance and the transit gateway and should be a /29 network assignment.
  - Peer ASN is the AS number configured on the SD-WAN appliance.

- 6. The information from the connect peers must be provided to Lumen personnel to be used in the configuration of the FlexVNF appliances.

  Navigate to the **VPC Console** and select **Transit Gateway Attachments**. Select the **Connect Attachment** in the upper pane and select the **Connect Peers** in the lower pane. Provide the configuration for transit gateway GRE address, Peer GRE address, transit gateway ASN, peer ASN, peer BGP address (IP to be used on the SD-WAN GRE tunnel interface/tvi-0/100) and both transit gateway BGP peer addresses (used as the BGP neighbor addresses in the FlexVNF appliance).
- 7. Route propagation from the transit VPC to the transit gateway should be disabled.

  Navigate to the **VPC Console** and select **Transit Gateway Route Table**. Select the appropriate transit gateway in the upper pane and select the **Propagations** tab in the lower pane. Select the VPC attachment ID to the transit VPC and click **Delete Propagation**.
- 8. A static route for the transit gateway CIDR must be added to the default route table in the transit VPC. This cannot be done until the transit gateway has been attached to the transit VPC.

  Navigate to the **VPC Console** and select **Route Tables**. Select the default transit VPC route table in the upper pane, select the **Routes** tab in the lower pane, click **Edit Routes** then **Add Routes**, enter the CIDR block for the transit gateway and select the **Transit Gateway** in the **Target** list.
- 9. Additional static routes may be required in the host VPC(s) representing the routes learned from the SD-WAN overlay. BGP routes learned from the FlexVNF appliances in the transit gateway are not injected into the host VPC(s). Transit gateways don't propagate their routes into the attached VPC(s), based on current AWS route propagation capabilities. Static routes will need to be created in each host VPC to represent the routes available in the SD-WAN overlay.

Appendix

Security groups

These are listed for reference but should be created using the CloudFormation templates.

AWS WAN security group configuration

Customer will be required to create and provide the name of a security group to be used for the WAN network. The security group must allow the following inbound connectivity at a minimum.

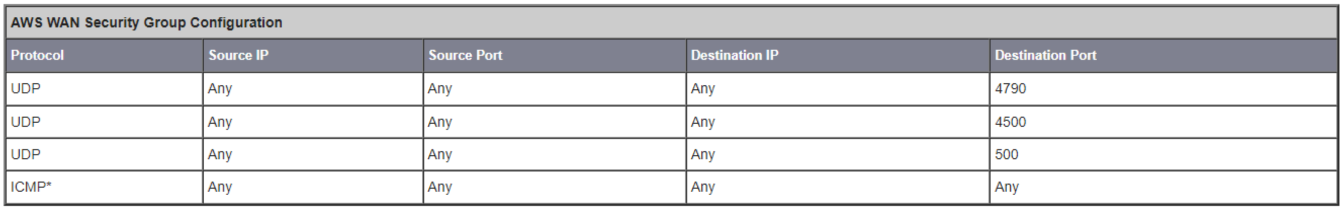

AWS WAN Security Group Configuration				
Protocol	Source IP	Source Port	Destination IP	Destination Port
UDP	Any	Any	Any	4790
UDP	Any	Any	Any	4500
UDP	Any	Any	Any	500
ICMP*	Any	Any	Any	Any

NOTE – ICMP can be locked down to specific hosts or ranges to prevent ICMP scanning or other similar attacks. Source or destination IPs can be limited based on customer requirements or needs.

AWS management security group configuration

Customer will be required to create and provide the name of a security group to be used for the Management network. The security group must allow the following inbound connectivity at a minimum.

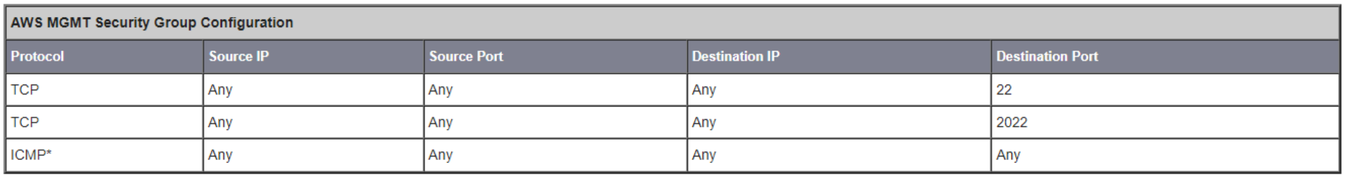

AWS MGMT Security Group Configuration				
Protocol	Source IP	Source Port	Destination IP	Destination Port
TCP	Any	Any	Any	22
TCP	Any	Any	Any	2022
ICMP*	Any	Any	Any	Any

Note: ICMP can be locked down to specific hosts or ranges to prevent ICMP scanning or other similar attacks.

Reference: Transit gateway VPC overview

The customer will need to create the 'Gateway VPC' within their AWS account. In an HA design, each appliance must be deployed in a different AWS availability zone to maintain high availability and to prevent both appliances from being affected by a single AWS maintenance. The customer must complete the following steps to create the gateway VPC:

Note: Most of these steps are covered in the “CloudFormation template” and are just a reference.

1. Create a new VPC and assign a CIDR block unique within the customer network.
2. Attach an IGW to the VPC allowing for connectivity to the internet.
3. Create a management subnet in availability zone 1 and another management subnet in availability zone 2.

4. Create an internet WAN subnet in availability zone 1 and another Internet WAN subnet in availability zone 2.
5. Create a LAN subnet in availability zone 1 and another LAN subnet in availability zone 2.
6. Optional if MPLS connected - Create an MPLS subnet in availability zone 1 and another MPLS subnet in availability zone 2. Create and attach a VPN gateway and enable route propagation as detailed in the ‘Customer steps in AWS for MPLS connectivity’ section of this document. (*This step is not completed by the CloudFormation template).
7. Create a route table with the IGW as the next hop for all traffic and attach to both management subnets and both Internet WAN subnets
8. Create a security group for management traffic.
9. Create a security group for internet WAN traffic.