Lumen® SD-WAN

Deploying an SD-WAN branch in AWS (beta)
November 2020

LUMEN®

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.
- Cloud formation 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 3 subnets in the same availability zone, Internet gateway attached to 2 of these subnets, security groups, a key pair to be used on the appliance, and route tables as required for the 3 subnets. If the customer also requires Lumen MPLS connectivity, they will need to create a VPN gateway, attach it to the VPC, accept the connection and create the virtual interface and associated BGP session.

NOTE: Lumen will be able to provide the customer with a cloud formation 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.

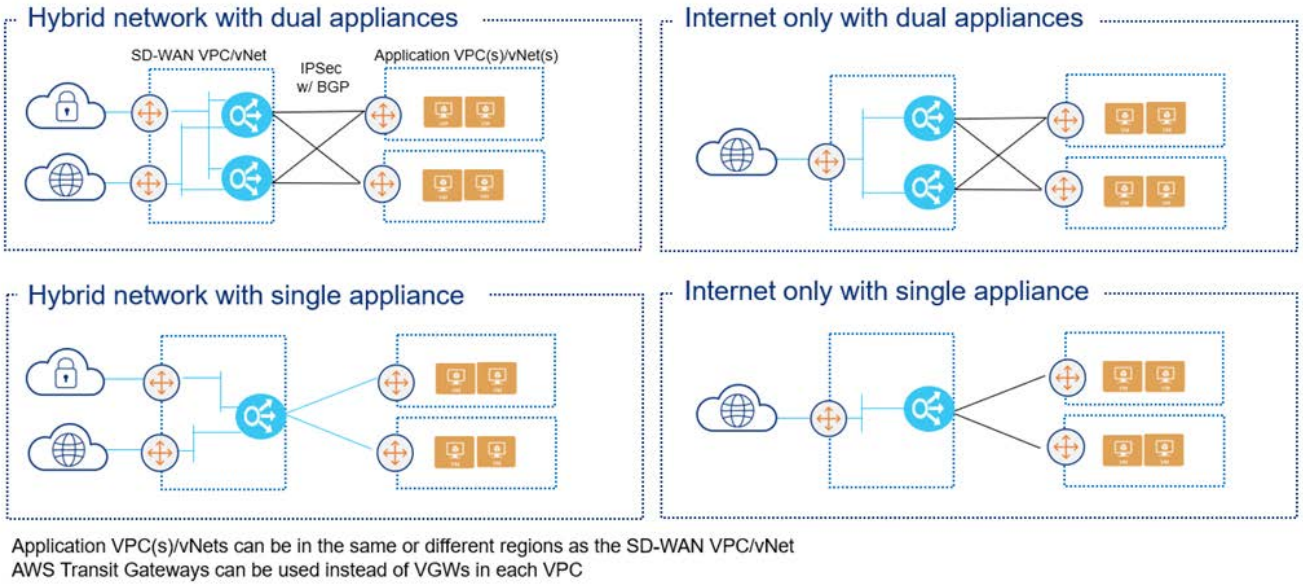

Figure 1: Cloud SD-WAN deployment topologies

Cloud formation templates

Lumen can provide the customer with cloud formation 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 SDWAN 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 SDWAN 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 SDWAN 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 SDWAN 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.

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. 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 cloud formation templates.
    - Key Pair Name
    - Availability Zone
    - Available image AMIID's 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 cloud formation templates.
    - Any IP information for interfaces not participating in DHCP and/or BGP neighbor information if required.

NOTE: Some of the information confirmed above will be possible to be created by cloud formation 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.

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.

AWS								
vCPE Size	Throughput BW	Cores	RAM	Disk	NICs	VM NIC Type	CPU Type	Cloud Provider Designation
Small-FlexVNF	200M	4	7.5G	80G	4	1 or 2	Intel Sandy or better	c4.xlarge
Medium-FlexVNF	600M	8	15G	80G	4	1 or 2	Intel Sandy or better	c4.2xlarge
Large-FlexVNF	1G	16	30G	80G	8	1 or 2	Intel Sandy or better	c4.4xlarge

* Throughput based on uni-directional streams. 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.

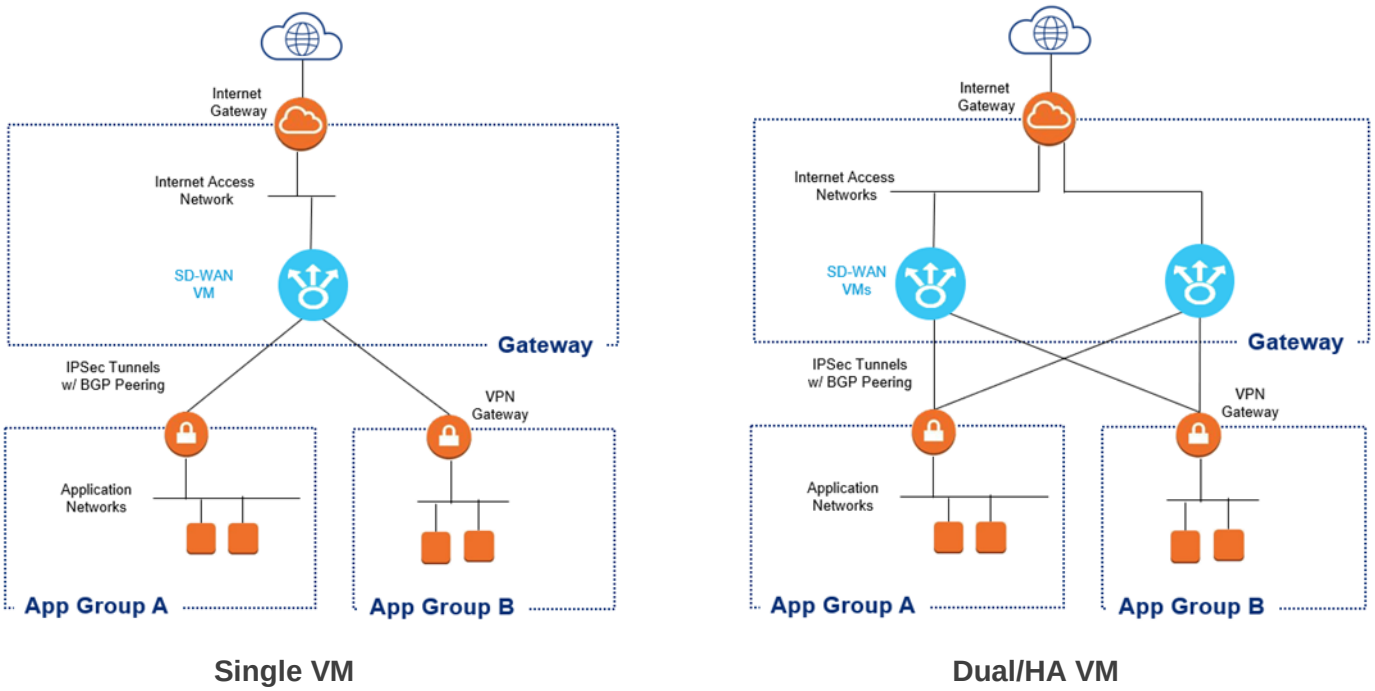

Figure 2

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 cloud formation 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 2 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 SDWAN 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 SDWAN VM instance.

a. Navigate to the VPC console, select 'Customer Gateways' in the left pane and click 'Create Customer Gateway.

b. Enter a unique name for the gateway, select Dynamic Routing, enter the BGP ASN that will be configured as the local-as on the SDWAN appliance and enter the public IP address of the Internet WAN port for the SDWAN appliance. This can be found in the EC2 Console under Instances by selecting the SDWAN VM and clicking on the eth1 interface and copying the Public IP address.

c. Repeat the above steps for the second SDWAN VM instance, if applicable.

4. Create and attach a Gateway – If the customer has a single host VPC and will be using the SDWAN 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 SDWAN 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 2 options depending on whether the customer is using a VPN Gateway or Transit Gateway. See the **Reference: VPN Connections (IPSec tunnels)** section of this document for instructions on creating the tunnels/connections.
6. After information from the VPN connections are provided back to Lumen SDWAN operations (from step 5), Lumen engineers will complete configuration steps on the SDWAN 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.

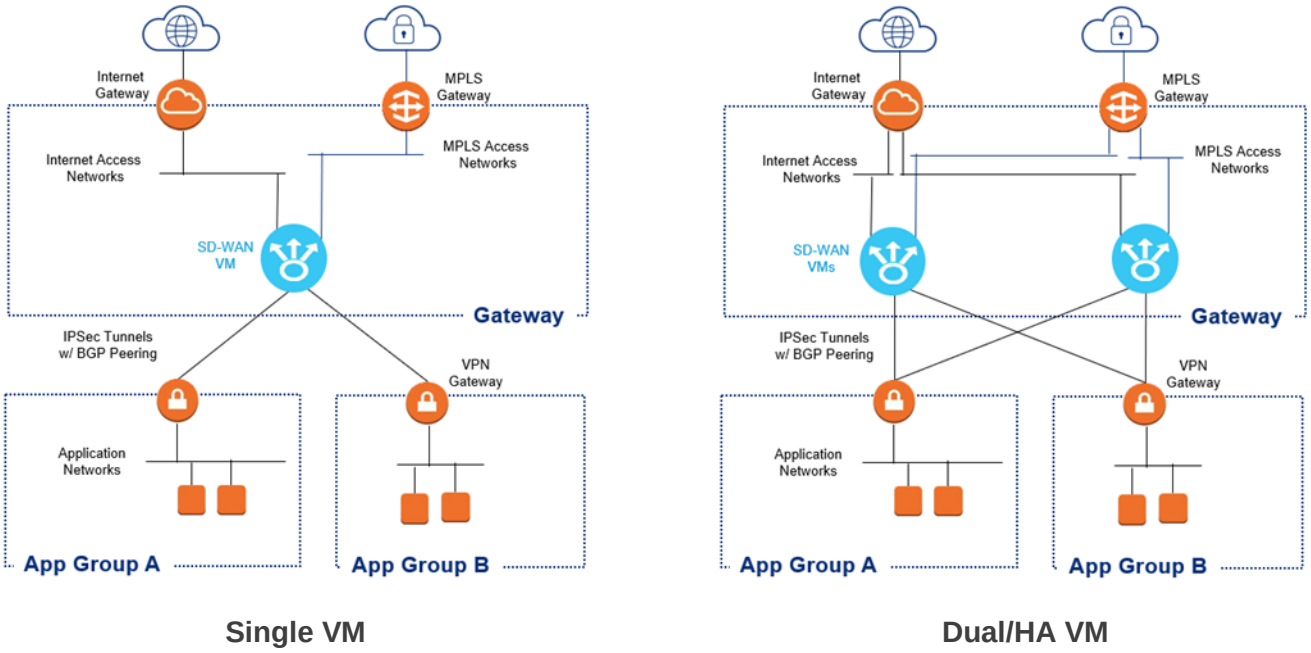

Figure 3

Summary of customer steps for deployment

Similar to 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 cloud formation 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 2 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 SDWAN 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 SDWAN VM instance.

a. Navigate to the VPC console, select 'Customer Gateways' in the left pane and click 'Create Customer Gateway.

b. Enter a unique name for the gateway, select Dynamic Routing, enter the BGP ASN that will be configured as the local-as on the SDWAN appliance and enter the public IP address of the Internet WAN port for the SDWAN appliance. This can be found in the EC2 Console under Instances by selecting the SDWAN VM and clicking on the eth1 interface and copying the Public IP address.

c. Repeat the above steps for the second SDWAN VM instance, if applicable.

4. Create and attach a Gateway – If the customer has a single host VPC and will be using the SDWAN 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 SDWAN 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 2 options depending on whether the customer is using a VPN Gateway or Transit Gateway. See the **Reference: VPN Connections (IPSec tunnels)** section of this document for instructions on creating the tunnels/connections.
6. After information from the VPN connections are provided back to Lumen SDWAN operations (from step 5), Lumen engineers will complete configuration steps on the SDWAN VMs.

Additional customer steps in AWS for MPLS connectivity

To enable MPLS connectivity, the customer must configure and attach a VPN gateway to their VPC, accept the connection, create and attach a virtual interface (VIF) to the VGW and create a BGP session between the VGW and Lumen MPLS Network. The following steps must be completed by the customer in their AWS Console for a hosted MPLS connection:

1. To support Native Controller access, the customer will need to add the /28 IP block 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.
2. Customer must have a VPC with an associated CIDR block(s), 4 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 connecting to the MPLS interface on the SDWAN appliance.

3. Customer must create and attach a VPN Gateway (VGW) to the VPC. 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. 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 1 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 a VPC by selecting only the newly created VGW, clicking the 'Actions' button and select 'Attach to VPC'. Select the VPC to attach the gateway to and click 'Attach'.
5. Modify the route table associated to the subnet for the MPLS 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'.
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'.
10. It may take up to 5 minutes for routes to propagate. You can confirm the routes by checking the MPLS route table or the VRF route table on the MPLS Provider router.

Common deployment elements

Reference: AWS gateway types

AWS supports two types of gateways that will support the topologies.

- VPN Gateway (VGW) - can only be attached to a single VPC and does NOT support ECMP. They do not allow traffic to pass directly from VPC to VPC. A deployment with multiple host VPCs would require a VPN gateway and 4 IPSec tunnels (2 to each SDWAN VM instance) for each host VPC. AWS bills a per hour charge for each connected customer gateway (SDWAN VM instance) and a per GB charge for traffic. There is no charge for the VGW or the attachment to the VPC. Steps to create:
    - From the VPC console, select 'Virtual Private Gateways' in the left pane and click 'Create Virtual Private Gateway'.
    - Enter a descriptive name for the gateway, select Custom ASN and enter the ASN for the VGW. This will become the peer-as on the SDWAN VM instances. If this is incorrectly set, the gateway will need to be deleted and re-added.
    - Once created, select the gateway from the 'Virtual Private Gateways' tab, click 'Actions' and select 'Attach to VPC' and select the appropriate host VPC from the dropdown
    - Route propagation must be configured for each subnet in the VPC that will use the SDWAN VM instances. From the VPC Console, select 'Route Tables' in the left pane, select the appropriate route table, and select the 'Route Propagation' tab in the lower pane. Click 'Edit route propagation', check the box under 'Propagate' and click save. Repeat for all subnets in the VPC that will use the SDWAN VM instances.
- Transit Gateway (TGW) - can be attached to multiple VPC 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 VM instance. AWS bills a per hour charge for each VPC connection and each customer gateway connection (SDWAN VM instances) as well as a per GB charge for traffic between VPC's or between VPC and VPN. Steps to create:
    - From the VPC console, select 'Transit Gateways' in the left pane and click on 'Create Transit Gateway'
    - Enter a name and description for the gateway. Enter the BGP ASN to be used on the gateway. This will become the peer-as on the SDWAN appliances. Leave VPN ECMP support checked. Default route table association and route table propagation will automatically associate and propagate routes into the default route table for each VPC the TGW is attached to. This may not be desirable, depending on the customer's AWS Infrastructure.

- Create a Transit Gateway Attachment by selecting 'Transit Gateway Attachment' in the left pane and clicking 'Create Transit Gateway Attachment'. Select the TGW from the dropdown, select VPC for the attachment type, enter an attachment name tag, select the host VPC and subnet within the VPC and click 'Create'. Repeat this step for each host VPC as required. Only one subnet can be selected, however, all subnets in a VPC can use this attachment.
- Create a TGW route table by selecting the 'Transit Gateway Route Tables' in the left pane and clicking 'Create Transit Gateway Route Table'. Enter a name and select the appropriate TGW ID and click 'Create'. Select the route table and the 'Associations' tab in the lower pane. Click 'Create Association' and select the Transit Gateway Attachment for the appropriate VPC. Repeat this step for each VPC as required. Select the 'Propagations' tab and click 'Create Propagation'. Select the VPC from the dropdown and click 'Create'. Repeat this step for each VPC as required.
- Routes from the TGW are not propagated directly into the routing table for each VPC subnet. To use the TGW routes from each subnet in each VPC, the routing table for each subnet in each VPC must have a route added with the next hop set to the TGW. This can be a default 0/0 route or individual routes based on the design of the customer's network. Navigate to 'Route Tables' in the left pane and select the appropriate route table. Select the 'Routes' tab in the lower pane, click 'Edit Routes' then 'Add Route'. In the new line, enter the route in the destination column, select 'Transit Gateway' in the target column and select the appropriate TGW. Click 'Save' to complete the addition. Repeat this step for each subnet in each VPC as required for the appropriate destinations.

NOTE: IPSec tunnels are created between the VGW/TGW and the SDWAN VM instance using the public Internet WAN IP of each VM instance. Additional per GB charges will apply for traffic egressing the Internet WAN of the VM instances toward the VGW/TGW and on to the host VPCs.

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 SDWAN appliance.
  - Select 'Site-to-Site VPN Connections' from the left pane and click 'Create VPN Connection'
  - Enter a descriptive name tag, select the appropriate VGW from the dropdown, select existing Customer Gateway and select the first SDWAN appliance from the dropdown. 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 IP's 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 SDWAN appliance.

- 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 SDWAN 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, similar to 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 SDWAN appliance.
  - Select 'Transit Gateway Attachment' from the left pane and click 'Create Transit Gateway Attachment'.
  - Select the appropriate TGW from the dropdown and select VPN for the attachment type. Select 'Existing' and select one of the Customer Gateways from the dropdown. 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 SDWAN appliance.
  - 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 SDWAN VPN connection and click 'Create Association'. Repeat this step for the second SDWAN VPN connection.
  - Select the 'Propagations' tab in the lower pane and click 'Create Propagation'. Select the first SDWAN VPN connection from the dropdown and click 'Create Propagation'. Repeat this step for the second SDWAN VPN connection.
  - 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 SDWAN appliances and must be shared with Lumen personnel.

Appendix

Subnets and security groups

These are listed for reference but should be created using the Cloud Formation 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
TCP	Any	Any	Any	4790
TCP	Any	Any	Any	4500
TCP	Any	Any	Any	500
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.

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: 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 in order 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 “Cloud Formation 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 a 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 cloud formation 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.