Manage Multicloud Defense Gateways

- Overview, on page 1
- Configure Multicloud Defense Gateway and VPC/VNets, on page 11
- Manage Your Gateway, on page 21

Overview

Multicloud Defense Gateway is a network-based security platform comprised of a network load balancer with a cluster of Multicloud Defense Gateway instances. It is an auto-scaling and self-healing cluster that scales out and in depending on the traffic load. Multicloud Defense Controller and gateway instances exchange constant and continuous information about the state, health and telemetry. The Multicloud Defense Controller makes the decision to scale out/in by measuring the telemetry data received from the gateway instances. The gateways can be configured to run in multiple availability zones for a highly available, resilient architecture. This ensures that a single availability zones failure from a cloud service provider does not compromise the security posture for running applications.

Once you have configured a gateway and any corresponding VPCs or VNets, you can use the **Gateway Details** page in the Multicloud Defense Controller to view and manage the state of them.

Multicloud Defense Gateways can be deployed in two ways; **Hub** mode and **Edge** mode.

Gateway Retry

The Multicloud Defense Gateway is a self-healing component Multicloud Defense. If at any point the deployment of your gateway fails or experiences issues, Multicloud Defense automatically attempts to redeploy the gateway with the **gateway retry**. This action happens infinitely until you manually disable or delete the gateway from the controller.

You can configure the retry action in terraform in two aspects: first, you can configure how many times Multicloud Defense retries to deploy the gateway. After the maximu number of attempts to redeploy are complete, Multicloud Defense stops retrying. Second, you can configure the time between retry attempts. As an example, you can configure three gateway retry attempts every hour. This means that every hour, Multicloud Defense retries to deploy the gateway three times and then stops. This action repeats until the gateway issues resolve or if you delete the gateway from the controller.

Tunnel Inspection in your Gateway

The Multicloud Defense Gateway automates GRE tunnel inspection by encapsulating the original packet by adding a new GRE header and an outer IP header. The encapsulated packet is then transmitted over the intermediate network and when the encapsulated packet reaches the destination endpoint of the GRE tunnel, the GRE header and the outer IP header are removed, revealing the original packet. The original packet is then forwarded to its final destination.

While GRE itself does not provide encryption, it can be combined with other protocols like IPsec (Internet Protocol Security) to secure the encapsulated traffic. IPsec can provide confidentiality, integrity, and authentication for the GRE tunnel. This is particularly useful for site-to-site VPN tunnel connections and can be used in conjunction with routing protocols to provide redundancy and failover capabilities.

Supported Gateway Use Cases

Egress

Egress/East-West Gateways

Deploying an Egress/East-West gateway to protect traffic leaving their public cloud networks. The egress gateway functions as a transparent forward proxy, performing full decryption and embedding advanced security features like intrusion prevention, antimalware, data loss prevention, and full-path URL filtering. Optionally, it can also operate in a forwarding mode, where it doesn't proxy or decrypt traffic but still applies security functionalities like malicious IP blocking and FQDN filtering.

The following diagram is an example of an AWS account with an egress gateway in a centralized mode:

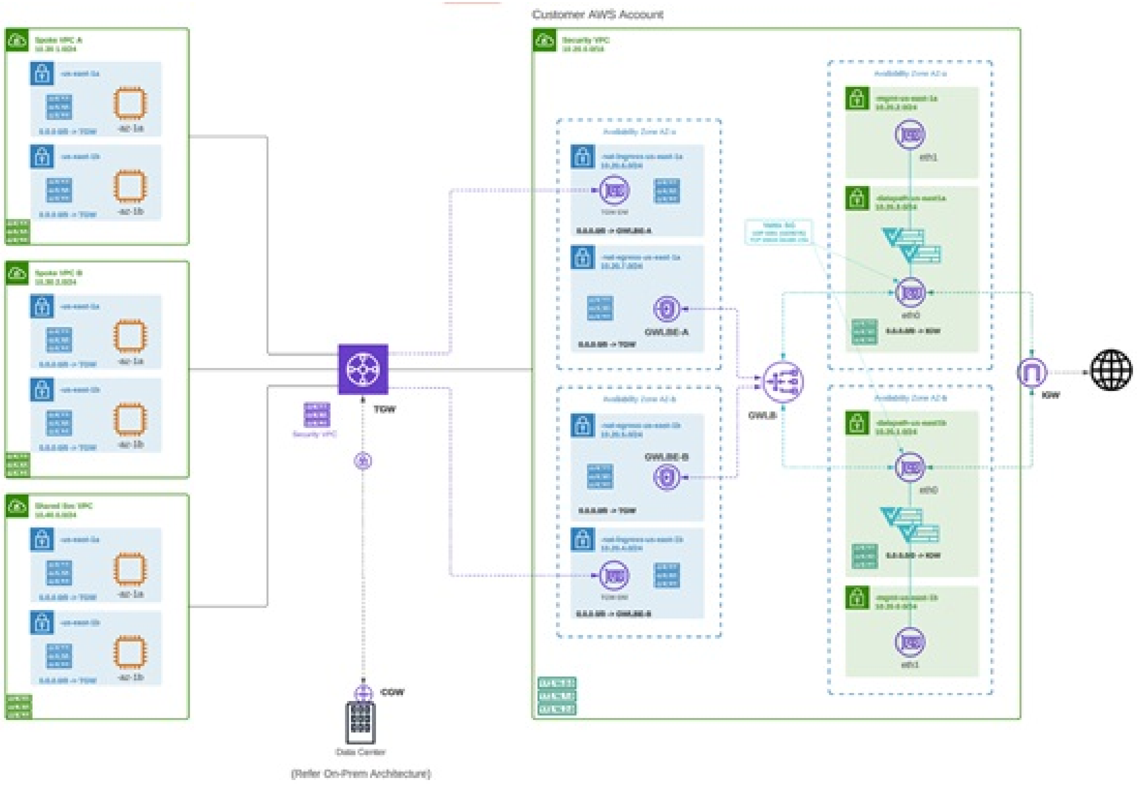

NAT Gateways in Egress

Note At this time, Multicloud Defense supports native gateways in an egress deployment for AWS and Azure ony.

Network Address Translation (NAT) gateways are gateways designed to originate from within your cloud service provider. Egress traffic appears from a single IP address, or at least one per availability zone. By building a gateway and hosting it from within your cloud environment you can potentially increase efficiency and reduce costs. Note that if the association between the VPC or VNet in Multicloud Defense and the gateway in your cloud ervice provider fails, Multicloud Defense system logs capture the instance for troubleshooting.

See the following supported configurations:

- Azure supports one subnet.
- You must have at least **one** public IP address configured in your NAT gateway.

The following diagram is an example of an Azure account with an egress gateway in a centralized mode:

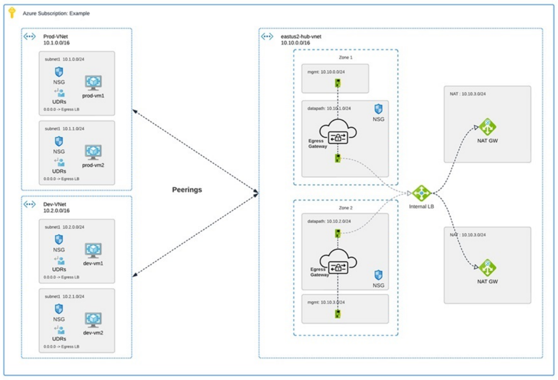

AWS CloudWAN

At this time, Multicloud Defense supports the inclusion of AWS' CloudWAN in egress gateways. CloudWAN is is an intent-driven managed wide area network (WAN) service that unifies your data center, branch, and AWS networks. While it is possible to create a global network by interconnecting multiple Transit Gateways across regions, CloudWAN offers built-in automation, segmentation, and configuration management features specifically designed for building and operating global networks based on your core network policy.

This option provides enhanced features such as automated VPC attachments, integrated performance monitoring, and centralized configuration, all managed within AWS Network Manager. This enables you to centrally manage and visualize your CloudWAN core network and Transit Gateway networks across AWS accounts, regions, and on-premises locations.

Key Benefits:

- **Simplified Network Management**: AWS CloudWAN provides a centralized dashboard through AWS Network Manager for managing network configurations, policies, and monitoring traffic. This greatly reduces the complexity of dealing with multiple, disparate networking solutions and offers a unified view of the network.

- **Scalability**: It enables customers to easily scale their network as their business grows. As organizations expand their cloud presence and global footprint, CloudWAN can accommodate the increased demand without requiring significant manual reconfiguration
- **Optimized Performance**: By leveraging AWS's global infrastructure, CloudWAN ensures high performance and low latency connectivity across various geographic locations, improving application performance and user experience.

CloudWAN Simplification:

- **Centralized Policy Management**: The core network policy, written in a declarative language, defines segments, AWS tegion routing, and how attachments should map to segments. With a single network policy, customers can manage their entire network's routing and security policies, reducing the need for manual configurations and minimizing errors.
- **Automated Operations**: CloudWAN automates many network management tasks, such as route propagation and policy enforcement, freeing up IT teams to focus on more strategic initiatives.
- **Seamless Integration**: It integrates with other AWS services and third-party solutions, enabling customers to build a cohesive and comprehensive network infrastructure with minimal friction.
- **Enhanced Visualization**: AWS Network Manager provides several dashboard visualizations, including world maps pinpointing network resources, monitoring with CloudWatch events, real-time event tracking, and topological diagrams of your network. This makes it easier to manage and monitor all aspects of your global network.

Security service insertion refers to the practice of integrating security services directly into the network path. Here are the benefits of implementing this with Multicloud Defense:

- **Enhanced Security Posture**: By inserting security services into the network, traffic can be inspected, monitored, and filtered in real-time, ensuring that threats are detected and mitigated before they can impact critical resources.
- **Consistent Security Policies**: Security service insertion ensures that consistent security policies are applied across the entire network, regardless of the underlying infrastructure or geographic location. This uniformity simplifies compliance and governance.
- **Improved Visibility and Control**: Integrating security services provides enhanced visibility into network traffic and potential threats. Administrators can leverage advanced analytics and monitoring tools to gain deeper insights and more effectively manage security risks.
- **Reduced Latency and Complexity**: By embedding security functions into the network path rather than routing traffic through separate security appliances, latency is minimized, and network complexity is reduced, leading to better performance and simpler network architecture.
- **Flexibility and Scalability**: Security service insertion with Multicloud Defense enables organizations to dynamically scale their security measures in response to changing network conditions and emerging threats, ensuring robust protection at all times.
- **Centralized Security**: Consolidates security resources, reducing management burden and saving on infrastructure costs.
- **Simplified Routing**: Easily steer network traffic to security appliances without complex routing configurations or third-party automation tools.

- **Multi-Region Security Inspection**: Simplifies multi-region deployments, allowing intra-region and inter-region traffic to pass through security infrastructure without complex configurations.

By leveraging AWS CloudWAN and Multicloud Defense for security service insertion, customers can achieve a high-performing, secure, and easily manageable network infrastructure that supports their business growth and operational resilience. Multicloud Defense allows users to create a security services VPC, attach it to an existing CloudWAN, create a Network Functions Group (NFG), and secure spoke segments by updating routing—all in an automated manner.

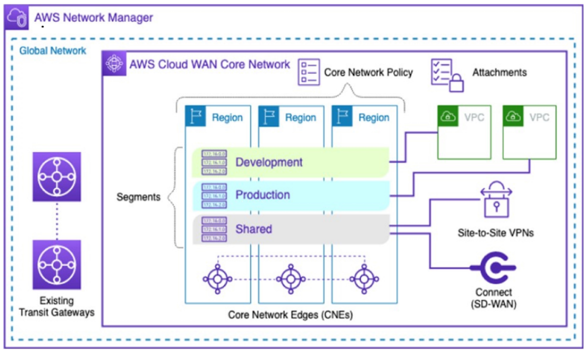

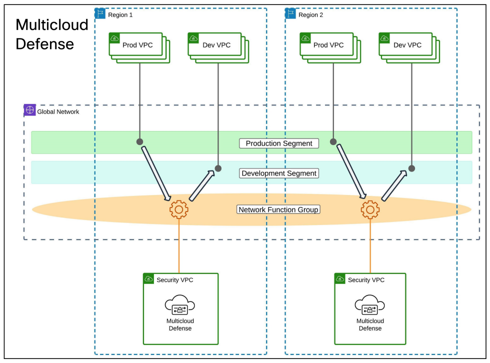

How to Create a Service VPC with AWS CloudWAN

To successfully create a service VPC with AWS CloudWAN, follow these steps:

- **Create Service VPCs**: Establish service VPCs in multiple CNEs with required gateways.
- **Create Network Function Groups (NFGs)**.
- **Attach Service VPCs as NFGs**: Use attachment policy rules to attach service VPCs.
- **Attach Workload VPCs**: Attach VPCs to respective segments using attachment policy rules.
- **Update Routing**: Modify policies and Workload VPCs to update the routing.
- **Update Core Network Policies**: Apply and execute the required changes in the Core Network policies.

Consider the following limitations before you create a service VPC with AWS CloudWAN:

- NAT gateways are mandatory for service VPCs.
- Dual-Hop and Edge Selection is currently **not** supported.
- Due to a limitation in AWS CloudWAN limitation that does not support SNAT-enabled traffic for forwarding, traffic drops for policy rulesets configured with SNAT. We **strongly** recommend you disable SNAT in your Multicloud Defense policy ruleset.
- To add an additional service VPCs in different regions (CNE) you one of two options:
  - Manual execution and application of policies are needed to update the routing for the new NFG attachment.
  - Manually update the routing tables of new service VPC datapath subnets with workload VPC routes through the Core Network.

Ingress

Deploying an Ingress gateway protects our public-facing applications. The Ingress gateway acts as a reverse proxy that carries out full decryption and applies advanced security functionalities such as intrusion prevention, antimalware, web application firewall (WAF), and full-path URL filtering.

The following diagram is an example of an AWS account with an ingress gateway in a centralized mode:

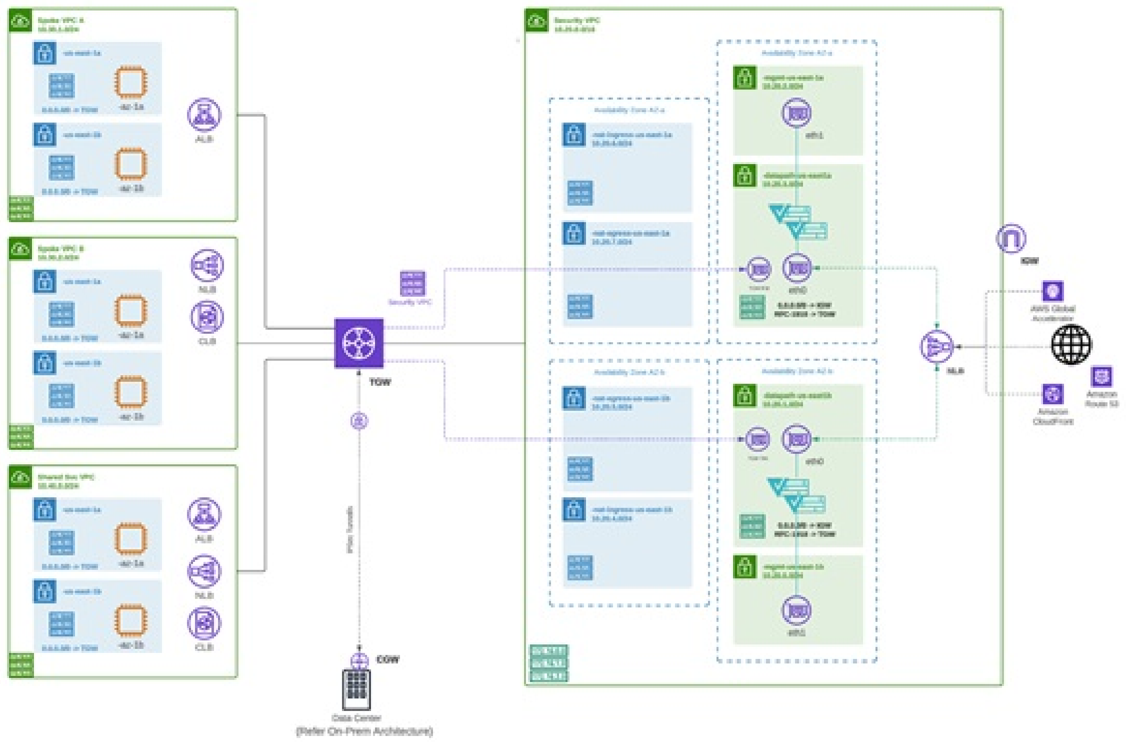

East-West

An Egress/East-West gateway deployment implements East-West L4 segmentation between subnets or VPCs/Vnets within their public cloud environments. The gateway functions in a forwarding mode with L4 firewall rules, allowing or denying traffic based on set parameters, with optional logging enabled.

The following diagram is an example of an AWS account with an east-west gateway in a centralized mode:

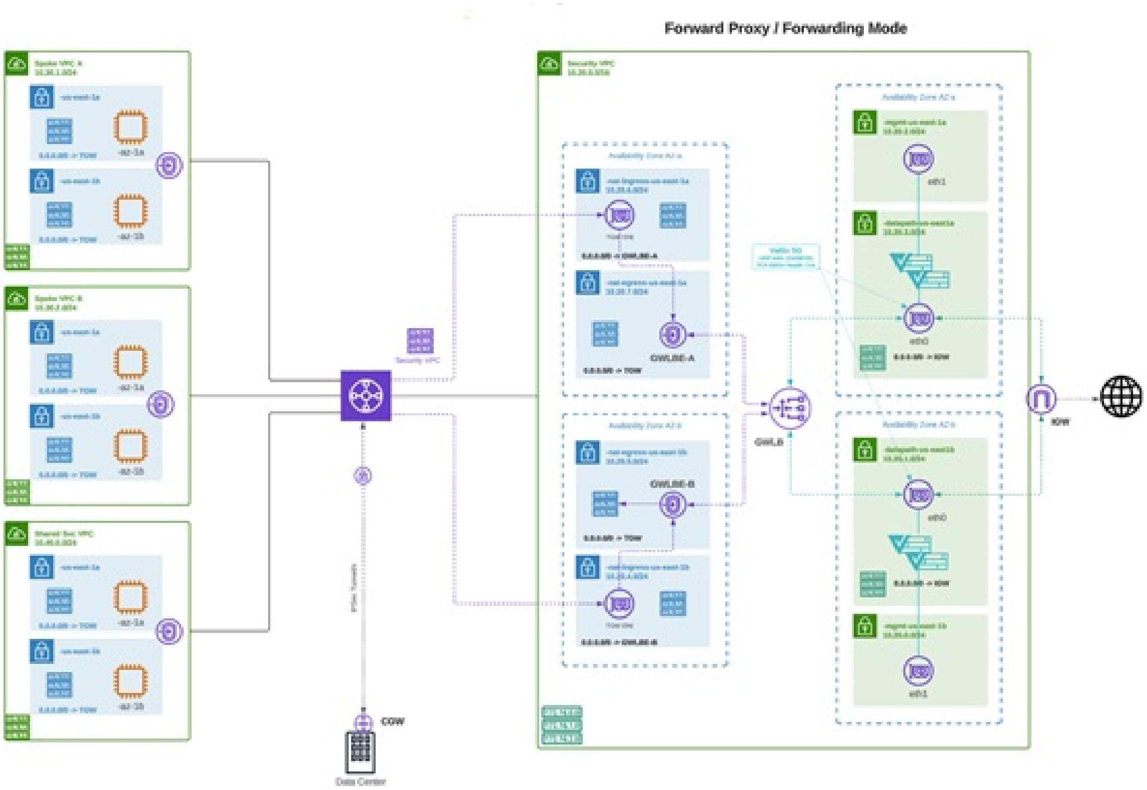

Distributed

You have applications running in multiple VPC/VNets. Deploy a Multicloud Defense Gateway in each of the VPCs/VNets.

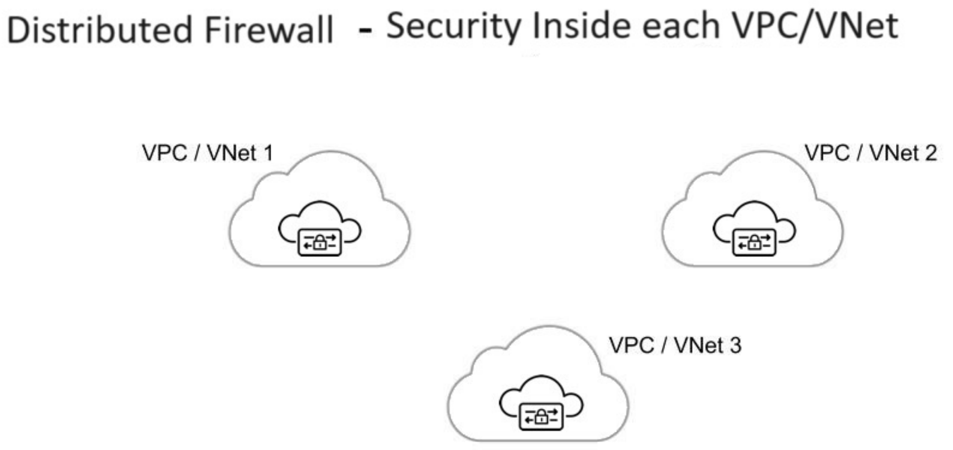

Centralized / Hub

You have applications running in multiple VPCs/VNet. You would like to secure all the applications through a centralized security services VPC/VNet. This model deploys the Multicloud Defense Gateway in a service VPC. You attach all the application VPCs (Spoke VPCs) and the Services VPC to the AWS Transit Gateway or VNet/VPC peering in Azure and GCP. Multicloud Defense provides an option to orchestrate the AWS Transit Gateway, Services VPC and the Spoke VPC Attachments. This is the recommended solution for ease of deployment, removing the complexity of multiple route tables and Transit Gateway attachments.

Figure 1: AWS - Using AWS Transit Gateway

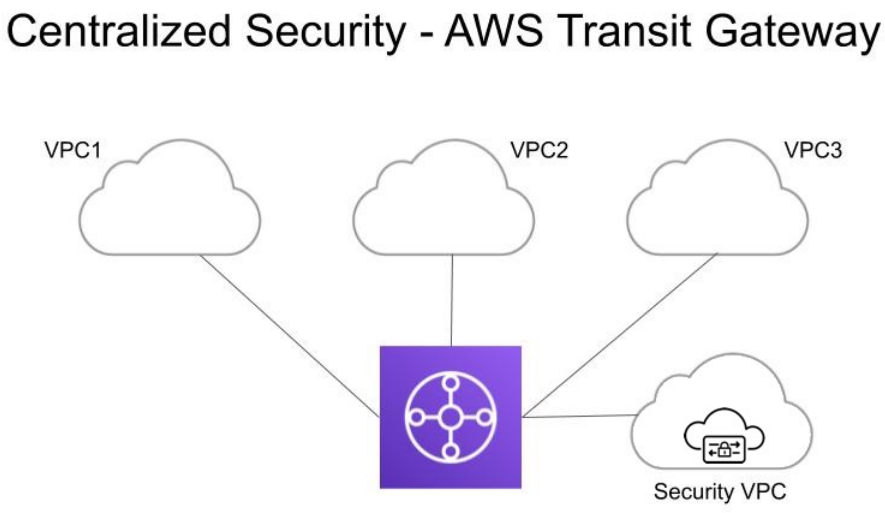

Figure 2: Azure - VNet Peering

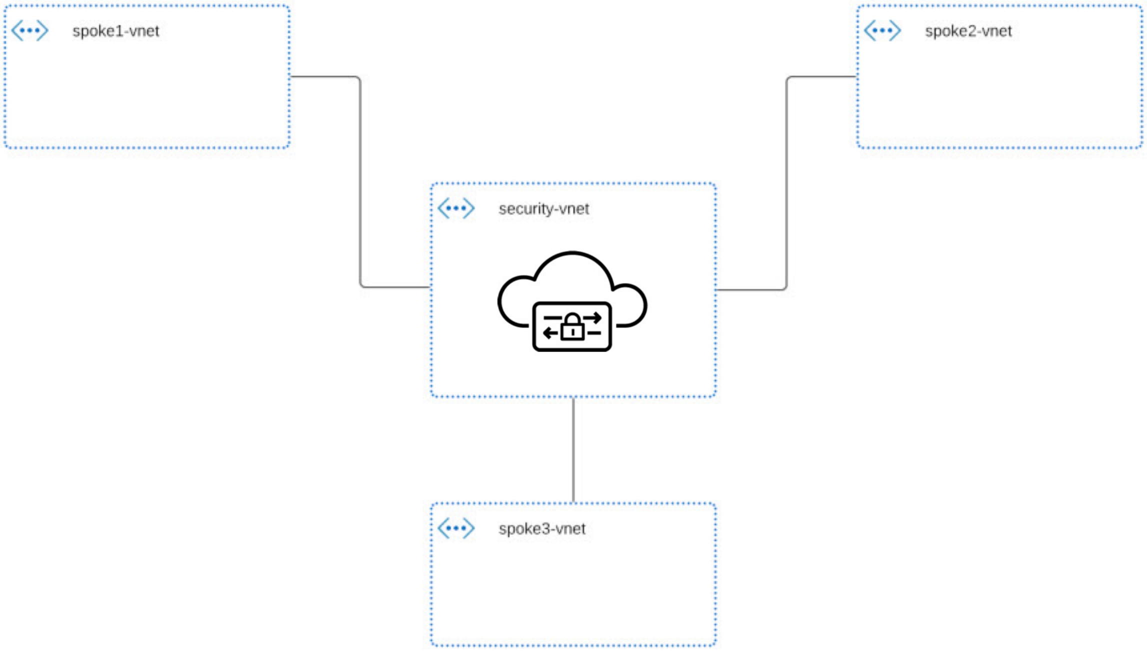

Figure 3: GCP - VPC Peering

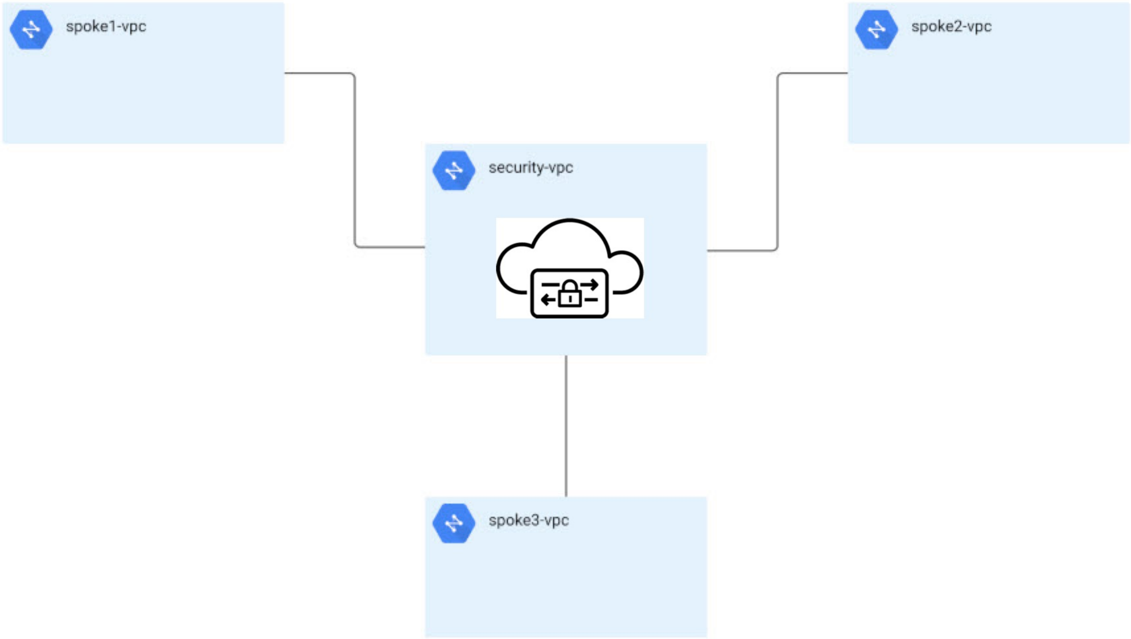

Advanced Gateway Configuration: Use Your Own Load Balancer

You can use a load balancer that is native to either AWS or Azure when creating a Multicloud Defense Gateway. Because AWS and Azure are different platforms, they do not use the same word for "load balancer" but the functionality mentioned below is identical in performance. Continue reading the appropriate information for the cloud service provider you currently have.

To configure your Multicloud Defense Gateway to use your own load balancer, see Add a Multicloud Defense Gateway, on page 18.

Note Note that both of these configurations support **ingress gateway** deployments only.

AWS Global Accelerator

Multicloud Defense can integrate with a set of one or more AWS global accelerators to use as an ingress point to load balance traffic across the Multicloud Defense Gateway instances. This is similar to the AWS network load balancer that is created and managed by Multicloud Defense when an ingress gateway is deployed, but offers an alternative ingress point for the ingress gateway to protect applications and workloads.

Accelerator manages the global accelerators' listener endpoint group to ensure the endpoint group has the active set of gateway onstances. Client IP addresses are preserved as they pass through the global accelerator to the Multicloud Defense ingress gateway.

In order to integrate Multicloud Defense with a global accelerator, you must first create the global accelerator within AWS, defined a desired listener and created an empty endpoint group (or an endpoint group that contains the existing Multicloud Defense ingress gateway instances). Once the AWS resources exist, then configure the Multicloud Defense ingress gateway to integrate with the global accelerator.

For any additional configuration information, see Amazon AWS documentation.

Azure Load Blanacer

If you have an Azure cloud service provider, you can now use your own load balancer from Azure as part of your Multicloud Defense Gateway. The Azure load balancer funnels and processes traffic from multiple proxy servers to a system-provided backendpool that contains at least one cluster of Multicloud Defense Gateway instances. This scenario is ideal if you want create a security policy for multiple proxy servers that handle non-HTTP traffic.

You must create a Multicloud Defense Gateway that defers to the Azure load balancer to be able ot use this capability. Beware the following prerequisites and limitations:

- You **must** have your Azure load balancer already configured.
- We **strongly** recommend creating and configuring a backend pool in Azure for your custom load balancer. The backend pool does not have to contain any resources at this time and can be modified later.
- If you opt to configure your Azure load balancer with a resource group, the Azure resource group and the Multicloud Defense Gateway's resource group must be configured for the same region.
- If you opt to configure your Azure load balancer with a resource group, the load balancer resource group and the Multicloud Defense Gateway resource group do **not** have to be the same.
- You can configure a health probe for your Azure load balancer but is not required.
- The Multicloud Defense Gateway's virtual network and the Azure load balancer's virtual network should be the same.
- The Multicloud Defense Gateway's datapth subnet and the Azure load balancer's subnet should be the same.

- You **must** attach your gateway to a VPC that has at least one availaibilty zone.

For any information on how to create, modify, or complete an Azure load balancer, see Microsoft Azure documentation.

Gateways Details

To view the **Gateway Details** page for already established gateways are available in **Manage** > **Gateways**. You can add and manage all gateways from this page. Managing a gateway allows you to edit, upgrade, enable, disable, export, or delete the instance. You must click the checkbox of the gateway you want to modify prior to making any changes.

Note You **must** be an Admin or SuperAdmin for these actions.

To filter and search the list of gateways, use the following criteria can be any of the following items:

- **Name** - The name of the gateway.
- **CSP Account** - The cloud service provider account that is associated with the gateway.
- **CSP Type** - The type of cloud service provider account.
- **Region** - The region of the cloud service provider that is associated with the gateway you are searching for.
- **State** - The current state of the gateway. Gateways can be active or inactive, or pending active or pending inactive.
- **Instance Type** - Each cloud service provider supports a number of instance types.
- **Mode** - Multicloud Defense Gateway instances can be depoyed in hub or edge mode.

Click **Switch to Advanced Search** to construct your own search. Use the drop-down option within the search bar to utilize some of the auto-generated search criteria if needed. For searches that have to repeated, you can **copy** or even **save** searches for future use.

Configure Multicloud Defense Gateway and VPC/VNets

Create a Service VPC or VNet

Use the following procedure to create a service VPC or service VNet, depending on the gateway you are creating this for.

Before you begin

Be aware the options listed in this procedure may be specific to your cloud service provider:

- If you opt to configure a VPC or VNet with a native gateway (NAT gateway), you must have a native gateway configured from your cloud service provider. See your cloud service provider documentation for more information.

- If you intend to deploy a service VNet with an Azure NAT gateway, confirm you have all of the permissions in your custom role within the Azure dashboard prior to creating and deployng. See Create a custom role to assign to the Application for the complete list of permissions.
- If you provide your own transit gateway, you are able to attach more than one VPC or VNet to it. It is even possible to replace an existing VPC or VNet with a new one without re-deploying the gateway.

If you intend to implement AWS CloudWAN as part of your service VPC, ensure the following is configured prior to this procedure:

- Global Network
- Core Network
- Core Network Edge (CNE) Regions
- Segments
- Workload VPCs
- (Optional) Network Function Groups (NFGs). Note that Multicloud Defense Controller allows you to create new ones as part of this procedure.

Procedure

Step 1 From the Multicloud Defense Controller, navigate to **Manage** > **Service VPCs/VNets**.

Step 2 Click **Create Service VPC/VNet**.

Step 3 Input parameter values:

- **Name** - Assign a name to the Service VPC/VNet.
- **CSP Account** - Select the CSP account to create the Service VPC/VNet.
- **Region** - Select the region the Service VPC will be deployed to.
- (Azure only) **CIDR Block** – The CIDR Block for Service VNet. This must not overlap with your Spoke(application) VNets.
- (AWS/GCP only) **Datapath CIDR Block** - The CIDR Block for the Multicloud Defense Gateway datapath Service VPC. This CIDR block must not overlap with address ranges in your Spoke (application) VPCs.
- (AWS/GCP only) **Management CIDR Block** - The CIDR Block for the Multicloud Defense Gateway management Service VPC. This CIDR block must not overlap with address ranges in your Spoke (application) VPCs.
- **Availability Zones** - If you are creating a VPC, you **must** configure **one** availabilitliy zone only. For a VNet, Multicloud Defense recommends to select at least two availability zones for resiliency.

  Note If you are attaching an AWS or Azure NAT gateway to this VPC, you must have at least one availability zone configured. Note that once you add availability zones to an AWS service VPC you cannot edit the zones to add or remove them if you deploy in an edge or centralized mode.

- (AWS CloudWAN only) **Network Type** - Select **CloudWAN**.
- (AWS CloudWAN only) **Network ID** - Expand the drop-down menu to select the core network that is associated with the global network in your AWS account.

- (AWS CloudWAN only) **Network Function Group** - Use the drop-down menu to select an existing network function group. This selection attaches the service VPC to the network function group in the core network. Alternatively, select **Create New** to create a new group for this VPC. If you create a new network function group, you will be prompted in this Service VPC window to enter a new name for the network function group.
- (Azure only) **Resource Group** - The resource group to deploy service VNet.
- (AWS only)**Transit Gateway** - The Transit Gateway connects virtual private cloud and on-premises networks through a central hub. Use the drop-down menu to select an existing gateway for this VPC. If there is no pre-existing gateway for you to select, choose **Create_new**. This option allows Multicloud Defense to create one as part of the VPC creation process.
- (AWS only) **Transit Gateway Name** - If you opted to create a new Transit Gateway, enter a name for the gateway in this field.
- (AWS only) **Auto accept shared attachments** - If you opted ot create a new Transit Gateway and intend to use this VPC for a multi-account hub gateway deployment, check this option.
- (AWS and Azure only) **Use NAT Gateway** - Enable this option if you want all egress traffic will go through NAT Gateway. If you are using a NAT gateway for an Azure account, confirm you have all of the permissions in your custom role within the Azure dashboard before finish creating this service VNet. See Create a custom role to assign to the Application for the complete list of permissions.

  Caution Do **not** enable this NAT Gateway option if you intend to deploy this Service VPC to deploy a Multicloud Defense VPN gateway in your AWS or Azure environment.

Step 4 Click **Save**.

What to do next

If you have just created a service VPC for an AWS or GCP account, you must first Manage the Service VPC/VNet, on page 15 and then Add a Multicloud Defense Gateway and associate the VPC or VNet with the gateway.

If you have created a service VNet for Azure, we strongly recommend you Add a Multicloud Defense Gateway.

Secure Spoke VPC or VNet

By securing the spoke VPCs, you create a more robust and resilient network that can respond to security threats. Securing spoke VPCs helps protect sensitive data that may be transmitted between the service VPC and the spoke VPCs; this can help reduce the overall attack surface as well as proper security measures in spoke VPCs support network segmentation, which is a key strategy in limiting the spread of potential security incidents.

We strongly recommend you secure you spoke VPCs for AWS and GCP accounts before you create or add a gateway.

Below is an example of how spoke VPCs interact with your network:

Figure 4: Azure Combined Hub - Multisubscriptions

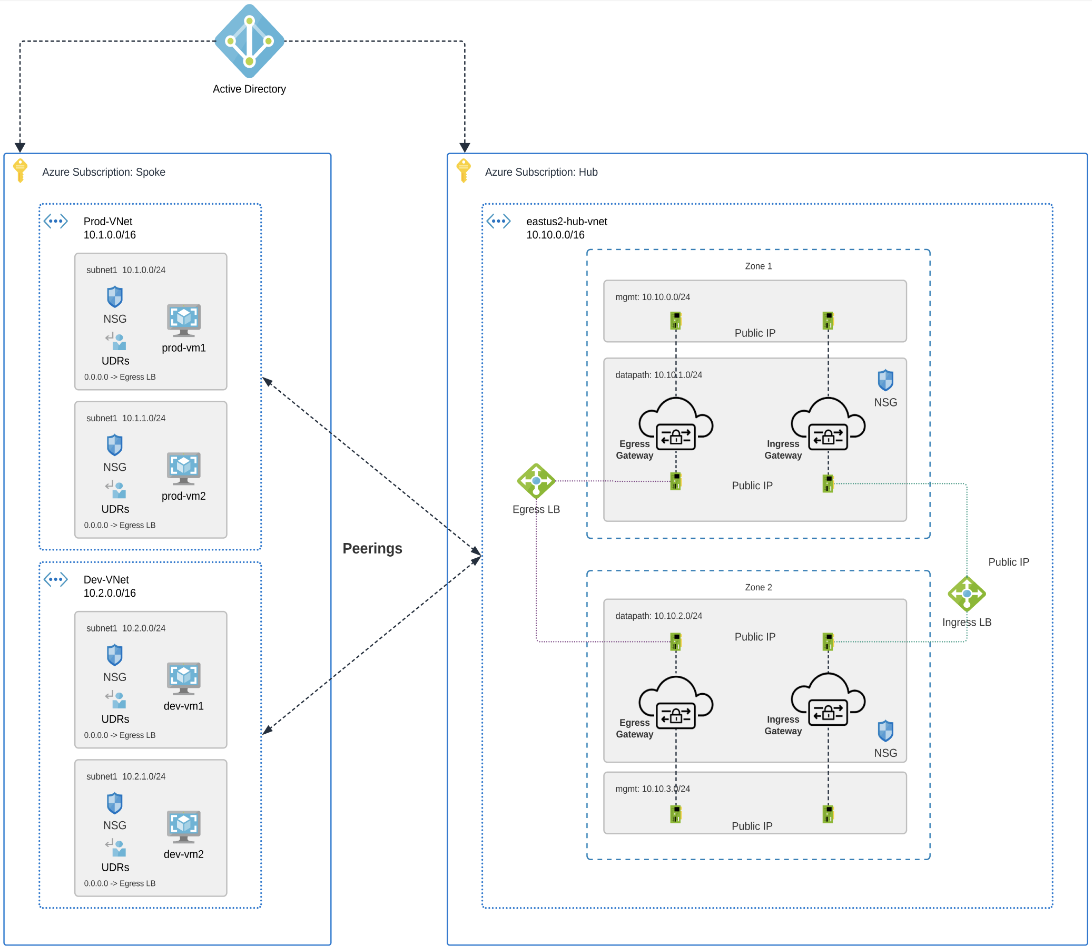

Prerequisites and Limitations

Be sure the following is completed prior to securing your spoke VPC or VNet:

AWS

- AWS does not support adding or removing availability zones from a service VPC after its creation for environments deployed in edge or centralized mode. If you need to modify the availabilty zones after creating a service VPC, you must create a new VPC with the correct zones included.
- AWS accounts with CloudWAN must have the following configured through the AWS Network Manager before you secure a spoke VPC or add a gateway:
  - For AWS accounts that are already onboarded, manually modify the permissions list in the AWS dashboard to include `networkmanager:*` to the MCDControllerRole IAM policy. See AWS' "Adding and removing IAM Identity permission" documentation for more information.
  - You must attach an egress/east-west gateway to the service VPC.
  - You must have at least one global network configured.
  - You must have at least one core network already created, does not have to contain segments already.

Azure

- VNet pairing is supported across accounts within the same CSP type. You can add spoke VPC/VNets within an account and across accounts. In Azure, for spoke VPCs peering across subscriptions, the CSP accounts should be onboarded using the same app registrations, and subscriptions should be within the same Active Directory.
- Azure environments require a route table attached **prior** to securing spoke VPC/VNet. See the "Associate a route table to a subnet" chapter in the Azure user guide for more.
- Azure does not support adding or removing availability zones from a service VPC after its creation for environments deployed in edge or centralized mode. If you need to modify the availabilty zones after creating a service VPC, you must create a new VPC with the correct zones included.

GCP

- GCP does not support adding or removing availability zones from a service VPC after its creation for environments deployed in edge or centralized mode; this also applies to Azure, GCP, and OCI for environments deployed in edge mode. If you need to modify the availabilty zones after creating a service VPC, you must create a new VPC with the correct zones included.

OCI

- OCI does not support adding or removing availability zones from a service VPC after its creation for environments deployed in edge or centralized mode. If you need to modify the availabilty zones after creating a service VPC, you must create a new VPC with the correct zones included.

Manage the Service VPC/VNet

Use the following procedure to manage a spoke VPC or spoke VNet:

Before you begin

When you protect an AWS service VPC that is configured to utilize the AWS CloudWAN, the table shown in this page has a separate row for each edge region. You can add/remove segments to secure the segment using the service VPC. Each segment can be edited with a list of VPCs that can be attached or dettached from the segment. Any traffic flowing through the segment will be protected by the network function group configured in the VPC. Anything forwarded from the segments seen in this table pass through the network function group configured in the VPC.

Procedure

Step 1 From the Multicloud Defense Controller dashboard, navigate to **Manage** > **Service VPCs/VNets**.

Step 2 Select Service VPC or Service VNet and click **Actions**.

Step 3 Click **Manage Spoke VPC/VNet**.

Step 4 To add a segment to a region that is attached to the VPC or VNet displayed in the table, click **Add**.

Step 5 Use the drop-down menu to select an available network segment. This action assigns an existing network segement to a service VPC or the network fucntion group inside your service VPC. Note that Multicloud Defense does not create network segments, you must create network segments as part of the core network in you AWS account.

Step 6 To **Remove** a network segment, select the segment and then click **Remove"**

Step 7 Click **+ Add VPC** to add a VPC and associate a user VPC to the network segment.

a) In the **Add VPC to Segment** window, select all spoke VPC or VNets in the left side of the window and click ">" to assign them to the segment. Altenatively, select any existing VPCs or VNets and click "<" to remove it from the segment.

b) Click **Save** to confirm the VPC changes.

Step 8 Click **Save** to confirm the network segment changes. Note that it may take up to 30 minutes for these changes to go into effect and for the affected VPC or VNet to become "Active".

Export a Spoke VPC or VNet

Use the following procedure to export the configuration of a spoke VPC or VNet:

Procedure

Step 1 From the Multicloud Defense Controller dashboard, navigate to **Manage** > **Service VPCs/VNets**.

Step 2 Select the Service VPC or Service VNet from the table and click **Actions**.

Step 3 Click **Export.**

Step 4 Multicloud Defensegenerates an export wizard.

Step 5 Either click **Download** to download the terraform locally or scroll down and click **Copy Code** to copy the JSON resource.

Step 6 Manually paste into the terraform script.

Step 7 Within the terraform prompt, execute the command provided in the lower half of the window.

Step 8 Follow the prompts within the terraform prompt to complete the task. Close the export window.

Delete a Spoke VPC or Vnet

Use the following procedure to delete a spoke VPC or VNet from your account configuration. Note that you may have to confirm the deletion through the dashboard of your cloud service provider.

Procedure

Step 1 From the Multicloud Defense Controller dashboard, navigate to **Manage** > **Service VPCs/VNets**.

Step 2 Select the Service VPC or Service VNet from the table and click **Actions**.

Step 3 Click **Delete.**

Step 4 Confirm the deletion of the service VPC or VNet and click **Yes**.

Before You Begin

The supported cloud service providers are separate entities that use their own vocabulary and gateway environment. Not every option available in the Multicloud Defense Controller is compatible with your cloud service provider. For example, AWS uses its own Transit Gateway and you can add VPCs to it while Azure utilizes a load- balancer to manage web traffic and applications and you can add VNets to it. Keep this in mind when proceeding.

Note For AWS environments, when securing spoke VPCs in centralized mode, Multicloud Defense attaches VPCs to the Transit Gateway that is associated to the service VPC. By default, Multicloud Defense will randomly select a subnet in each availability zone for Transit Gateway attachment. You can change this option when you add a VPC or you can modify a VPC that is already assigned to the gateway.

You can also orchestrate a transit gateway through the Multicloud Defense Gateway or attach an existing Transit Gateway.

Limitations

Be aware of the following limitations when creating a Multicloud Defense Gateway:

- If you deploy a Multicloud Defense Gateway that uses a site-to-site VPN tunnel containing an IPSec profile, you must deploy the gateway **with** a service VPC or service VNet and **without** a Network Address Translation (NAT) gateway on either side of the VPN connection.
- Autoscaling is not supported for gateways containing an IPSec profile.
- Policy rules within the gateway **must** be Forwarding only.
- If you intend to include an IPSec profile in a Multicloud Defense Gateway for an AWS or Azure account, the gateway instance **must** be configured with `core 8`. Multicloud Defense Gateway does not currently support gateways with core 2 or core 4 options.

Resources Created by Multicloud Defense

The following resources are created by Multicloud Defense when you create a gateway, VPC, or VNet. These are created as part of the process and do not require any additional actions from the user. Note that difference resources are created per each cloud service provider requirements.

GCP Resources

Multicloud Defense creates two service VPCs and four firewalls. See the following for the exact resource allocation:

Service VPC

- Management
- Datapath

Firewall Rules

- Management (ingress)
- MAnagement (egress)
- Datapath (egress)
- Datapath (egress)

Note The Service VPC CIDR **cannot** overlap with the Spoke VPC.

AWS Resources

Multicloud Defense creates three service VPCs to address the supported use cases (ingress, egress/ east-west). Created and affiliated with each of these VPCs is the following:

- Four subnets in each availability zone.
- One route table for each of the subnets.
- Two security-groups: management and datapath.
- One Transit Gateway.

  Note This Transit Gateway is created and attached to the gateway during the creation of the service VPC. This gateway can be reused with other service VPCs.

- A Transit Gateway route table.

  Note The route table is attache to the Service VPC as part of the creation process.

Note The AWS Gateway Load Balancer (GWLB) does not support add/remove of availability zones after initial deployment of a GWLB. You will need to redeploy the service VPC if you need to change availability zones. See AWS documentation for more information.

Azure Resources

Multicloud Defense created one Service VNet with the following resources:

- One VNet.
- Two network security groups.

The Service VNet CIDR value must not overlap with spoke VNet.

Add a Multicloud Defense Gateway

Use the following procedure to add a Multicloud Defense Gateway for your cloud service provider:

Before you begin

If you are planning on using an AWS global accelerator or Azure load balancer, be sure the load balancer is already configured prior to adding it to a Multicloud Defense Gateway. See Advanced Gateway Configuration: Use Your Own Load Balancer, on page 9 for more information.

Procedure

Step 1 Navigate to **Manage** > **Gateways**.

Step 2 Click **Add Gateway**.

Step 3 Select the cloud service provider you want to add the gateway to.

Step 4 Click **Next**.

Step 5 Enter the following information:

- **Instance Type** - Choose the type of cloud service provider. Note that there may be multiple variations of instances depending on which cloud service provider you are using.
- **Gateway Type** - Select either Ingress or Egress.

  Note Select **Egress** if you have an east-west network flow.

- **Minimum Instances** - Select the minimum number of instances that you plan to deploy.
- **Maximum Instances** - Select the maximum number instances that you plan to deploy. This is the maximum number that is used for auto-scaling in each availability zone.
- **HealthCheck Port** - Default is 65534. The port number used by Multicloud Defense load balancer to check the health of the instances. Datapath security groups assigned to the instance(s) must allow traffic on this port.
- (Optional) **Packet Capture Profile** - Packet Capture Profile for threat and flow PCAPs.
- (Optional) **Diagnostics Profile** - Diagnostics Profile used to store Technical Support information.
- (Optional) **Log Profile** - Log Forwarding Profile used to forward Events/Logs to a SIEM.
- (Optional) **NTP Profile** - Network Time Protocol (NTP) for time synchonization.
- (Optional) **BGP profile** - Border Gateway Protocol (BGP) used to support VPN Connections. If you intend on utizilign site-to-site VPN tunnels with a Multicloud Defense Gateway you **must** include this profile.

Step 6 Click **Next**.

Step 7 Provide the following parameters:

- **Security** - Select either Egress or Ingress.

  Note Select **Egress** if you have an east-west network flow.

- **Gateway Image** - Image to be deployed.
- **Policy Ruleset** - Select the policy ruleset to associate with this gateway.
- **Region** - Select the region this gateway will be deployed into.
- **Resource Groups** - Select the resource group to associate the gateway with.

- **SSHPublic Key** - Paste the SSH public key. This public key is used by the controller to access the CLI of the deployed gateway instances for debug and monitoring.
- **VNet ID** - Select the VNet to associate with the gateway.
- (Azure only) **User Assigned Identity ID** - Enter the cloud service provider identity to associate with this gateway. User-assigned managed identities can be used in place of credentials for resources. User-assigned managed identities can be used in place of credentials for resources for Azure services such as a private key stored in Azure Key Vault or to write PCAP files to an Azure Blob Storage.
- **Mgmt. Security Group** - Select the security group to associate with the management interface.
- **Datapath Security Group** - Select the security group to associate with the datapath interface.
- **Disk Encryption** - Select the appropriate option from the drop-down menu. For customer managed encryption key, the user will need to input the resource ID of the encryption key.

Step 8 Select the **Availability Zone**, the **Mgmt Subnet** and the **Datapath Subnet**. The available subnets will be based on the VPC or VNet selected above. For high availabilty purposes the gateway instances can be deployed in multiple availability zones. Click the plus button to add a new availability zone and select the parameters for the selected zones. Note that some cloud service provider regions do not support multiple availability zones. In such regions the gateway instances are deployed in only a single zone.

Note If your gateway is deployed in hub mode, availability zones cannot be edited after the initial deployment. Reconfirm your zones before deploying.

Step 9 (Azure only, optional) If you are deploying in distributed model with Multicloud Defense Gateway in the same VNet as application, ensure you complete the following:

- Add a route table in the Azure portal and associate the route table with all the subnets.
- Add a default route for 0.0.0.0/0 with **next-hop** as the IP address of the Gateway Network Load Balancer.

Step 10 Click **Next** to view the Advanced Settings.

Step 11 By default, the Multicloud Defense Gateway enables the use of the public IP of the router available. If you do not want this enabled, check the **Disable Public IP** box.

Step 12 (AWS and Azure only) **Attach Load Balancer**. Click **Add Load Balancer** to create a row for your custom load balancer. Alternatively, check any rows that are unecessary and click **Remove** to delete them from the gateway.

a) Expand the **Load Balancer** drop-down to select a load balancer from your AWS or Azure cloud service provider.

b) Expand the **Backend Pool** drop-down to select a backend pool to be associated with your gateway.

Step 13 Click **Save**.

What to do next

Multicloud Defense deploys the gateway.

You **must** attach at least one ruleset to the gateway before you secure a spoke VPC/VNet. See Rule Sets and Rule Set Groups for more information.

Manage Your Gateway

View your Multicloud Defense Gateways and statistic in **Manage** > **Gateways**. From this page you can search and filter your gateways, view the cloud service providers assocaited with each gateway, current instance count and type, and more.

For more information on the supported use cases for specific gateway environments, see Supported Gateway Use Cases, on page 2.

Edit a Multicloud Defense Gateway

You can edit a gateway in any state, whether it is enabled or disabled. Use the following procedure to edit an existing Multicloud Defense Gateway:

Procedure

Step 1 Navigate to **Manage** > **Gateways**.

Step 2 Select the Multicloud Defense Gateway you want to edit in the table so it is highlighted.

Step 3 Expand the **Actions** drop-down menu and click **Edit**.

Step 4 Modify the gateway configuration as needed.

Step 5 Click **Save** to confirm the changes. Alternatively, click **Cancel** to exit the changes.

Upgrade the Multicloud Defense Gateway

Multicloud Defense Gateways serve as an autoscaling self-healing Platform-as-a-Service (Paas), functioning as inline network-based security enforcement nodes. Unlike traditional firewalls, Multicloud Defense eliminates the need for customers to construct virtual firewalls, configure high-availability setups, or manage software installations.

Multicloud Defense Gateway instances operate on highly optimized software, incorporating a single pass datapath pipeline for efficient traffic processing and advanced security enforcement. Each gateway instance comprises three core processes: a "worker" process responsible for policy enforcement, a "distributor" process for traffic distribution and session management, and an "agent" process communicating with the controller. Gateway instances can seamlessly transition "in service" for a "datapath restart," enabling smooth upgrades without disrupting traffic flow.

New instances are spun up with new image. Once the instances are fully up, they are placed in the loadbalancer's (layer 4 sprayer of flows to gateway instances) target pool. The old instances are put in flow draining mode or flow timeout mode for the existing flows going through them. New flows will hit the new instances. Once the timeout (Azure) or the flows are drained (AWS), the old instances are reaped by the controller.

Use the following procedure to

Procedure

Step 1 Navigate to **Manage** > **Gateways**.

Step 2 Select the checkbox for the gateway you want to upgrade. You can make only one selection at this time.

Step 3 Select **Actions** > **Upgrade**.

Step 4 From the **Gateway Image** list, select the desired image.

Step 5 Click **Save**.

Step 6 Confirm the cloud service provider resource allocation necessary for the upgrade.

Step 7 Click **Yes** if the resource allocation is sufficient. Click **No** if the resource allocation is insufficient, increase the resource allocation in the cloud service provider, and return to continue the upgrade.

Note You can view the upgrade progress and new gateway instances being created from theinstances info for the gateway. Select the gateway and view the **Instances** in the Details pane.

Abort a Multicloud Defense Gateway

You can only abort a Multicloud Defense Gateway that is currently going through an in-progress gateway update.

Use the following procedure to abort an existing Multicloud Defense Gateway:

Procedure

Step 1 Navigate to **Manage** > **Gateways**.

Step 2 Select the Multicloud Defense Gateway you want to abort in the table so it is highlighted.

Step 3 Expand the **Actions** drop-down menu and click **Abort**.

Step 4 Confirm you want to abort the gateway and click **Yes**. To back out of the action, click **No**.

Enable a Multicloud Defense Gateway

You can only enable gateways that have been disabled. Use the following procedure to enable a

Procedure

Step 1 Navigate to **Manage** > **Gateways**.

Step 2 Select the Multicloud Defense Gateway you want to enable in the table so it is highlighted.

Step 3 Expand the **Actions** drop-down menu and click **Enable**.

Step 4 Multicloud Defense validates the gateway configuration. If the validation is successful, a table of current and required resources for an upgrade generate for review. If you approve of the gateway resource allocation, click **Yes** to confirm the action.

What to do next

Wait a few minutes for the Multicloud Defense Gateway to successfully enable.

If you've disabled a Multicloud Defense Gateway and deleted the site-to-site VPN tunnels affiliated with it, you **must** create a new site-to-site VPN tunnel connection, or recreate the previous VPN tunnel connection and then add it to the gateway. When a gateway is disabled, Multicloud Defense forgets the public IP address associated with the VPN tunnel. You must create a new tunnel connection to establish a new IP for the gateway instance.

Disable a Multicloud Defense Gateway

You can only disable a Multicloud Defense Gateway if it is currently enabled. You cannot disable gateways that are already disabled.

Use the following procedure to disable a Multicloud Defense Gateway:

Procedure

Step 1 Navigate to **Manage** > **Gateways**.

Step 2 Select the Multicloud Defense Gateway you want to disable in the table so it is highlighted.

Step 3 Expand the **Actions** drop-down menu and click **Disable**.

Step 4 Confirm you want to disable the gateway and click **Yes**. To cancel this action, click **No**.

What to do next

Wait a few minutes for the gateway to succesfully disable.

To completely disable the gateway, you **must** delete any site-to-site VPN tunnels affiliated with the gateway.

Export a Multicloud Defense Gateway

Use the following procedure to export the configuration of a Multicloud Defense Gateway:

Procedure

Step 1 Navigate to **Manage** > **Gateways**.

Step 2 Select the Multicloud Defense Gateway you want to export in the table so it is highlighted.

Step 3 Expand the **Actions** drop-down menu and click **Export**.

Step 4 Multicloud Defense generates an export wizard.

Step 5 Either click **Download** to download the terraform locally or scroll down and click **Copy Code** to copy the JSON resource.

Step 6 Manually paste into the terraform script.

Step 7 Within the terraform prompt, execute the command provided in the lower half of the window: `terraform import "ciscomcd_gateway"."object-name" <object name>.`

Step 8 Follow the prompts within the terraform prompt to complete the task. **Close** the export window in Multicloud Defense. There are no more steps in the dashboard.

Delete a Multicloud Defense Gateway

Use the following procedure to delete a Multicloud Defense Gateway. Note that this action is different from disabling the gateway.

Procedure

Step 1 Navigate to **Manage** > **Gateways**.

Step 2 Select the Multicloud Defense Gateway you want to delete in the table so it is highlighted.

Step 3 Expand the **Actions** drop-down menu and click **Delete**.

Step 4 Confirm the action and click **Yes**. To cancel the deletion action, click **Cancel**.

What to do next

We strongly recommend deleting any site-to-site VPN tunnel connections associated with this gateway after it is successfully deleted from the gateway table.