AWS

Manage Service VPCs

AWS Service VPC

For the Centralized (hub) mode deployment using AWS Transit Gateway, the Multicloud Defense Gateway is deployed in a new VPC. This VPC is called a Service VPC. The Service VPC and the application (Spoke) VPCs are connected to the AWS Transit Gateway in a Hub-Spoke model as shown below:

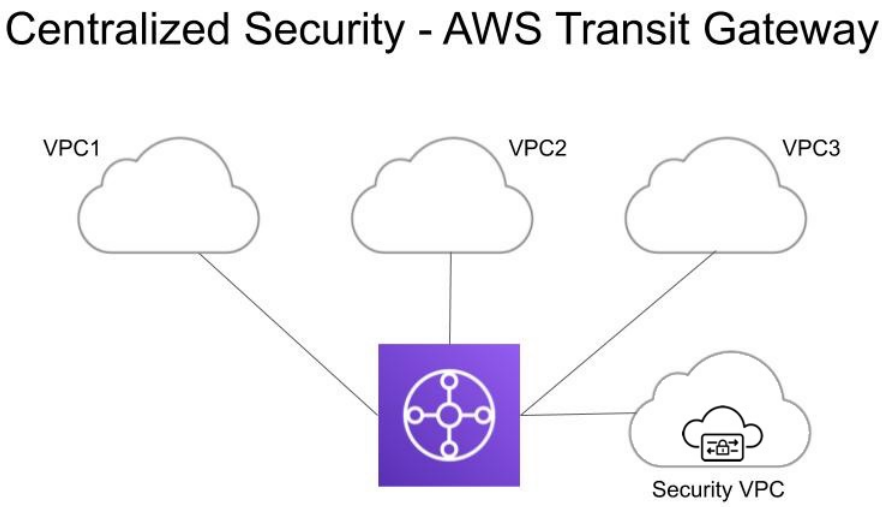

Multicloud Defense orchestrates creating the Service VPC, creating (or reusing) the AWS Transit Gateway (TGW), and attaching the Spoke VPCs and the Service VPC to the TGW. It updates the routing between the Service VPC and Spoke VPCs. Customers need to change the route tables associated with subnets in the Spoke VPCs to add a default route and set the destination to the TGW.

Note If the TGW is created using the AWS Terraform Provider, the attributes *default_route_table_association* and *default_route_table_propagation* MUST be set to *"disable"*.

If the TGW is created using the AWS Console, the attributes *Default association route table* and *Default propagation route table* MUST be set to *Disable*.

If the TGW is created using the Multicloud Defense Service VPC orchestration, the attributes are set appropriately.

If these values are not set properly, traffic will bypass the Service VPC and will not be protected by Multicloud Defense.

Create an AWS Service VPC

Step 1 Click **Manage** > **Service VPCs/VNets**.

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

Step 3 Provide a name for the service VPC (an example is *multicloud defense-service-vpc1*).

Step 4 Select the AWS account.

Step 5 Select the **Region** where the service VPC needs to be created (an example is*us-east-1*).

Step 6 Provide a CIDR block with mask minimum of /25 and maximum of /16. Make sure this does not overlap with any of the spoke VPC CIDRs that you plan to attach to the Transit Gateway (an example is *172.16.0.0/16*).

Step 7 Select the **Availability Zones**. It's recommended to select at least two (2) AZs for HA purposes (an example is *us-east-1a* and *us-east-1b*).

Step 8 Select a **Transit Gateway**. Alternatively, create a new one. You can reuse an existing transit gateway for all kinds of security types.

Step 9 Select the **Auto accept shared attachments**, if you are planning to use the transit gateway shared across multiple AWS accounts.

Step 10 Click **Save** to create the service VPC.

Note

- Multicloud Defense creates the following resources when a service VPC is created:
  - VPC
  - Four subnets in each AZ
  - One route table for each of the subnets
  - Two security-groups (management and datapath traffic)
- It's required to create a different Service VPC for each of the security types (Ingress, Egress and East-West).
- The transit gateway (created/selected during a service VPC creation) can be reused with other service VPCs.
- Review the transit gateway. If you opted to create a new TGW, it is included here.
- A transit gateway attachment to the service VPC is created.
- A transit gateway route table is created and associated with the attachment.
- AWS Gateway Load Balancer (GWLB) does not support add/remove of AZs after initial deployment of a GWLB. You will need to redeploy the service VPC if you need to change AZs.

Manage Spoke VPCs

Manage (Protect) Spoke VPCs in Hub Mode

When a Service VPC is created with a new Transit Gateway OR existing Transit Gateway, Multicloud Defense takes care of the orchestration of the Transit Gateway and Services VPC. It can also create Attachments for the Spoke VPCs and manage Transit Gateway route tables. This is a fully managed Transit Gateway solution that makes it very easy to use a Services VPC for Centralized security.

Note

- Wait for the Service VPC be created successfully and state is **ACTIVE** before proceeding with the following steps.
- Multicloud Defense Gateway can be deployed later in Service VPC that you just created.

To protect spoke VPCs, we need to associate spoke VPCs to the Service VPC. This allows Multicloud Defense to orchestrate the routing and create Attachments for spoke VPC's traffic to be inspected by Multicloud Defense.

When enabling Protected VPCs, Multicloud Defense Controller orchestrates the following:

- Creates Transit Gateway VPC Attachment for each of the Spoke VPCsAdds a Transit Gateway route table for each of the Attachments and associate with the Attachments.

- Adds a Transit Gateway route table for each of the Attachments and associate with the Attachments.
- Adds a default route in the TGW route table (associated with the Spoke VPC) to go to the Service VPC Attachment (and thus to the Service VPC).

Here is a sample routing setup after attaching two (2) Spoke VPCs:

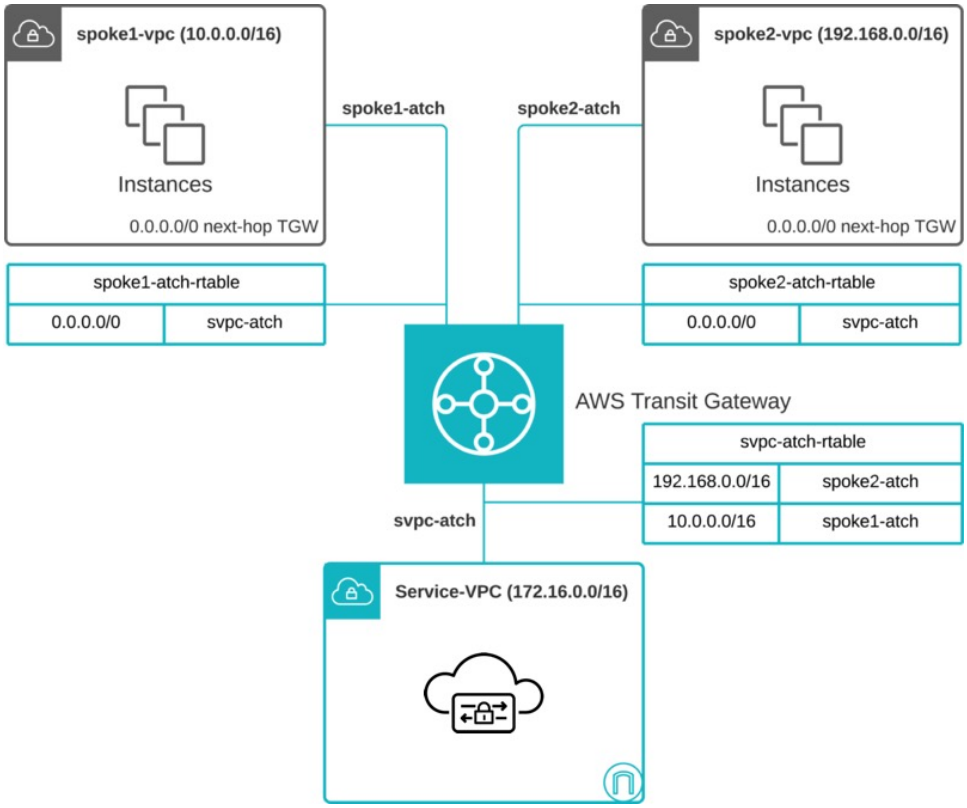

There are two ways to associate VPCs to the Service VPC.

- Add Spoke VPCs from Service VPC Menu, on page 4
- Add Spoke VPCs from Inventory Menu, on page 5

Add Spoke VPCs from Service VPC Menu

Step 1 Navigate to **Manage** > **Service VPCs**.

Step 2 Select a Service VPC and click on Manage Spoke VPCs.

Step 3 For the Spoke VPCs in the current account where the transit Gateway is created, add the VPCs under **Current Account VPCs to Protect**.

Step 4 Select the VPC from the dropdown, you cannot change the account and the region in this table. Click Add to add more VPCs.

Step 5 For the Spoke VPCs in the other accounts, add those under **External Account VPCs to Protect** table (The accounts must be added to the Multicloud Defense Controller prior to adding the VPCs. Please check the Add Cloud Account section on how to add a new Cloud account to the Multicloud Defense Controller).

a. Select the account, region and the VPCs in that region.

b. Multicloud Defense sets up automatic acceptance of the attachment invitations. So you don't need to do any manual steps to accept the attachments.

Step 6 Click on View/Edit link under the Route Tables column.

Step 7 Select the route table to update default route to Transit Gateway.

Step 8 (Optional) Select TGW Attachment Subnet to select which subnet to place the ENI.

Step 9 Click **Save**.

Add Spoke VPCs from Inventory Menu

Step 1 Navigate to **Manage** > **Cloud Accounts** > **Inventory**.

Step 2 Click on VPCs/VNets. This will list all the VPCs in your cloud accounts.

Step 3 Click on the **Secure** button to secureVPC.

Step 4 Select Service VPC.

Step 5 Select route table to update default route next hop to Transit Gateway.

Step 6 (Optional) Expand **Customize Transit Gateway Attachment Subnets** to customize Transit Gateway Subnet selection.

Step 7 Click **Save**.

Subnet Selection for Transit Gateway Attachment

When protecting spoke VPCs in centralized model (either through Service VPC Menu or Inventory Menu), Multicloud Defense attach VPCs to the Transit Gateway that is associated to the Service VPC. When attaching VPCs to the Transit Gateway, users can choose which subnet in each Avaliability Zone to place the ENIs. By default, Multicloud Defense will randomly select a subnet in each AZ for Transit Gateway attachment.

To customize the Transit Gateway subnet selection, please see Add Spoke VPCs from Service VPC Menu, on page 4 orAdd Spoke VPCs from Inventory Menu.

AWS Centralized Ingress Protection

The Multicloud Defense Gateway may be deployed in a central service VPC or distributed inside spoke VPCs to protect customer facing applications. The gateway acts as a **Reverse Proxy**. The users on the internet access the application via the Multicloud Defense Gateway. Configure the backend destination (the original application) as a proxy target on the Multicloud Defense Gateway. The proxy enables Multicloud Defense to decrypt TLS traffic and perform deep packet inspection. The proxied traffic to the backend/target can be sent as plain text HTTP, HTTPS, TCP or TLS.

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

Step 2 Click **Add Gateway**.

Step 3 Select the account you previously created.

Step 4 Click **Next** and ener the appropriate 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 Tpe** - AutoScaling.
- **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.

Step 5 Click **Next**.

Step 6 Provide the following parameters:

- **Security** - Ingress
- **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.
- **VPC** - Select the VPC in which the Multicloud Defense Gateway is deployed.
- **Key Pair** - Select the key pair to associate with this gateway.
- **IAM Role for Gateway** - Select the IAM role to associate with this Gateway.
- **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.
- **EBS Encryption** - EnableEBS encryption for the gateway instance. If enabled, the user will select either **AWS managed CMK** or **Customer managed encryption key**. For Customer managed encryption key, KMS key ARN needs to be provided.

Step 7 Select the **Availability Zone**, the **Mgmt Subnet** and the **Datapath Subnet**. The available subnets will be based on the 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.

Step 8 Click **Next**. The review page shows you the details of all the selected parameters. Review the available resources and see information about any AWS limits exceeded.

Step 9 Click **Finish**. The gateway deployment starts and takes approximately 5-7 minutes for the gateway to become **ACTIVE**.

Note On your AWS console, view the EC2 instances page and check the gateway instances created. The instances have a name tag that begins with **multicloud defense**.

The **Check Load Balancers** section and see that an internet facing Network Load Balancer is created. It does not yet have any listeners or target groups. The listeners and target groups (targeting the EC2 Multicloud Defense Gateway instances) are created when you add a service with the listener port and backend application.

Advanced Settings: 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, it will manage the global accelerators' listener endpoint group to ensure the endpoint group has the active set of gateway onstances. Client IP addresses will be preserved as they pass through the global acce

Accelerator to the Multicloud Defense ingress gateway.

In order to integrate Multicloud Defense with a global accelerator, the user must have first created 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 the Multicloud Defense ingress gateway can be configured to integrate with the global accelerator.

Parameter	Description
Global Accelerator	Select the Global Accelerator to attach to Gateway.
Listener Name	Friendlyname for the listener. This name will only exist in Multicloud Defense.
Listener	The listener in Global Accelerator.
Endpoint Group ARN	Multicloud Defense will automatically select the endpoint group ARN once listener is selected.

Note

- The AWS Network Load Balancer will still be deployed as part of Gateway deployment even if AWS Global Accelerator integration is enabled
- When configuring the Endpoint Group in the AWS Global Accelerator Listener, it is best to assign port TCP/65534 as the Health Check port. The Multicloud Defense Gateway is configured to respond to TCP/65534 to inform health status to the AWS Network Load Balancer and AWS Global Load Balancer. The same port can be used to inform health status to the AWS Global Accelerator.

AWS Centralized Egress / East-West Protection

The Multicloud Defense Gateway is deployed in a single VPC to protect the outgoing traffic of the applications running inside the VPC. The gateway acts as a forward proxy. For HTTP or TLS applications with SNI extension header, the Multicloud Defense Gateway can act as a transparent forward proxy. The applications access the internet without any change on their side. Multicloud Defense intercepts the traffic and considers that as proxied traffic. It creates a new session to the internet. For TLS traffic and the certificate to be trusted by the client applications, a trusted root/intermediate certificate must be configured on Multicloud Defense and the root certificate installed on all the client application instances.

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

Step 2 Click **Add Gateway**.

Step 3 Select the account you previously created.

Step 4 Click **Next** and enter the following parameters:

-
- **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 Tpe** - AutoScaling.
- **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.

Step 5 Click **Next**.

Step 6 Provide the following parameters:

- **Security** - Egress.
- **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.
- **VPC** - Select the VPC in which the Multicloud Defense Gateway is deployed.
- **Key Pair** - Select the key pair to associate with this gateway.
- **IAM RoleforGateway** - Select the IAM role to associate with this gateway.
- **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.
- **EBS Encryption** - Enable EBS encryption for the gateway instance. If enabled, the user will select either **AWS managed CMK** or **Customer managed encryption key**. For a customer-managed encryption key, **KMS key ARN** needs to be provided.

Step 7 Select the **Availability Zone**, the **Mgmt Subnet** and the **Datapath Subnet**. The available subnets will be based on the 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.

Step 8 Click **Next**. The review page shows you the details of all the selected parameters. Review the available resources and see information about any AWS limits exceeded.

Step 9 Click **Finish**. The gateway deployment starts and takes approximately 5-7 minutes for the gateway to become **ACTIVE**.

Note
- Check the AWS Console **Load Balancers** section and note that an internal network load balancer has been created. It does not yet have any listeners or target groups. The listeners and target groups (targeting the EC2 Multicloud Defense Gateway instances) are created when you add a service with the listener port and backend application.
- On your AWS console, check the **EC2 instances** page and check the gateway instances created. The instances have a name tag that begins with multicloud defense. Along with gateway instances, another helper/supporting instance is created. This is called a **NAT** instance. After the gateway is created and becomes **ACTIVE** change/add route in the route tables associated with the application subnets to have the default route's next-hop as the interface of the NAT instance. When the traffic exits the application subnets, it reaches the NAT instance. The destination IP in the packets is changed to the internal network load balancer's IP. This causes the traffic to reach the Gateway instance. The gateway inspects the SNI, or the HTTP host header, to find the destination address and sends the packet out. When the applications communicate over TLS, the gateway waits until the Client Hello reaches the gateway and then creates a new connection to the target (defined in the SNI field). The incoming certificate from the internet server is impersonated with the root/intermediate certificate installed on the Multicloud Defense Gateway and sent to the application.