Cisco Multicloud Defense FAQs

First Published: 2023-05-17

Last Modified: 2023-06-05

Cisco Multicloud Defense FAQs

What do Edge, Hub, Ingress, Egress mean?

Ingress: An application is running in a VPC. Traffic enters the VPC from outside (Internet) into the VPC. An Ingress Gateway is deployed to protect the application from external users.

Egress: Clients/EC2 Instances/Applications that require communication with the external world (Internet). To protect/control these clients from egressing traffic to the internet, it is necessary to restrict communication only with certain websites e.g. payment gateways, or approved source code repositories. An Egress Gateway is deployed to control the outgoing traffic.

Edge: The Gateways (Egress and Ingress) can be deployed in Edge or Hub mode. In the Edge mode, the Gateway is deployed in the same VPC as the application/s. If you have 5 VPCs running applications, 5 Gateways are deployed. This is best suited for a small number of VPCs.

Hub: Multicloud Defense creates a new VPC (called Service VPC) and deploys the Gateways inside this Service VPC. All VPCs that are running applications and the Service VPC containing Gateways are connected via AWS Transit Gateway. Multicloud Defense manages the orchestration of the Transit Gateway, VPC attachments and the routing automatically. The customer is required to edit the VPC route tables, making the Transit Gateway the default route destination. The Transit Gateway maybe new, or existing.

Multicloud Defense requires that you deploy different Gateways for Ingress and Egress use cases. A single Gateway cannot be used to protect Ingress and Egress traffic.

What is Forward Proxy and Reverse Proxy?

Forward Proxy rules and services are used by Egress Gateways. The Gateways act as proxy servers in both Ingress and Egress modes. In the Ingress case users access the proxy endpoint provided by the Multicloud Defense Gateway. In the Egress scenario, the proxy is transparent. The clients inside the VPC access external sites (internet) via routing through Multicloud Defense Gateways. Gateways respond to the clients. You will be asked to provide a root certificate that the Gateway uses to sign the external sites' certificates. The clients need to have this root certificate installed as a trusted source.

Reverse Proxy rules and services are used by Ingress Gateways. The service definition defines the port number the proxy listens on and the target application/host to forward the traffic toward.

What is URL filtering?

URL filtering is used in Egress Gateways only. A URL profile is a list of URLs and an action for each of the URLs. After creating a URL Profile, it is associated with a Policy Rule. When the traffic matches a rule that

has a URL Profile, URL filtering processing starts. The list is traversed in order and the action of the first item in the list that matches the traffic's URL is performed. The default policy is used if none of the URLs match to an allowed URL. There is an implicit (Default ALLOW) in the profile. The URLs can be provided as a string or a regular expression. You would not normally need an ALLOW rule, unless there is a regex match for a DENY in the list. For e.g if you want to allow https://website.com/news and DENY everything else from the same website, you can can define 2 items in the profile:

https://www.website.com/news ALLOW

https://www.website.com/.* DENY

When the URL is dropped, an event is logged in the URL Filtering events in the **Investigate** > **URL Filtering** menu of the Multicloud Defense Controller.

Can I use regular expressions in URL filtering?

Yes.

If the default action in URL filtering is to allow, why do we need ALLOW action? The default action in URL filtering if there is no match is to ALLOW. Specific action to ALLOW is useful if there is a very generic DENY further down in the list. To make a default action as DENY, add a rule

.* DENY 502

This causes all the URLs to be dropped. Now to open a specific URL to be allowed, add a rule above this to have an ALLOW. For example, if you want all the traffic to google.com to be allowed and rest all denied:

https://www.google.com ALLOW .* DENY 502

This can also be used to restrict the broader pages on a website but to allow a specific page to be ALLOWED

https://www.website.com/news ALLOW

https://www.website.com/.* DENY

What should a URL look like in the URL profile?

The URL in the URL list should be a complete string including http or https. You can use regular expressions like .* *(dot star) to define a generic scheme, like* .google.com.. This matches http or https and any prefix before google.com and any suffix after google.com.

What is L7 DOS?

L7 DOS is used in Ingress Gateways only. When the Gateway acts as an Ingress proxy targeting a backend application, you can enforce rate limits for the URLs. The limits can be set at a URL level for each of the HTTP actions (GET, POST etc). The rate limit is at a firewall instance level and not at the whole Gateway cluster level. So if the rate limit is set 1000 reqs/sec and the Gateway has 3 firewall instances, your application would receive 3000 reqs/sec.

How do I create a HUB mode Gateway and protect my VPCs?

HUB mode Gateway helps in a centralized security management of your cloud environment. If you have multiple Spoke VPCs running the applications, HUB mode is the preferred method of protecting all the VPCs. The security management is performed in a Service VPC. The Service VPC that hosts the Gateways is managed by the Multicloud Defense Controller. All VPCs must have non-overlapping CIDRs before they can be attached to the Transit Gateway. During the Gateway creation, choose either Ingress or Egress (the process is same) and select the Hub mode option. Select if you want to use a Transit Gateway that you already have, or choose to create a new Transit Gateway. Select a Service VPC if you have already created one, or choose to create a new one. While creating a new Service VPC, provide a CIDR that does not overlap with any of the Spoke VPCs that you plan to protect. Continue the Gateway creation process. You don't have to provide any other subnet or security group information. These are managed by Multicloud Defense. Provide the key pair and the firewall role that was created as part of account onboarding.

Once the Gateway is created, edit the Gateway to add the Spoke VPCs that you want to protect. In the Edit Gateway option, scroll down to the 'Protect VPCs' and select all the VPCs that you want to protect. Multicloud Defense creates Transit Gateway attachments to all the selected VPCs. It randomly picks a subnet from the VPCs to do the attachment. Once the VPCs are attached, change the VPC route table that's attached to the Application subnets and add/set the default route to the Transit Gateway. For Ingress Hub mode Gateway, you can be specific and set the route to the Service VPC CIDR instead of the default route. For Egress Gateway, the default route is the preferred option, but for SSH/management tasks you can set specific and more specific route to use an Internet Gateway.

How do I add my AWS account to the Valtix Controller?

The Multicloud Defense Controller needs access to your AWS account in order to create a Gateway, access to the Inventory and other tasks on the account. A CloudFormation Template (CFT) is provided by Multicloud Defense that creates the a cross account IAM role for use with the Multicloud Defense Controller. You will be provided with the Multicloud Defense account number as part of the onboarding process. The IAM role gives permissions to this account. Review the IAM Role documentation for the permissions assigned to this role.

What is a Gateway and firewall?

The terms Gateway and firewall are occasionally used interchangeably throughout the solution and the documentation. A Gateway is a cluster of firewall instances that are managed as a single entity. A Network Load Balancer (NLB) is created as part of the Gateway deployment that has all the firewall VM instances as the targets of this load balancer. A user never manages instances and Gateway independently. It's all managed by the Controller. The NLB ensures that the session traffic reaches the same firewall instance. A firewall instance is the security enforcer.

Does Valtix Gateway support HA, Auto-scaling?

Multicloud Defense Security platform is born in the cloud. HA and auto-scaling are built into the system from day 1. During the Gateway creation you are given the option to create instances in multiple zones, similar to running your applications in multiple zones (AZs). It is recommended to run Gateway instances in at least 2

zones. You also get to choose how many Gateway instances you wish to run. You can choose minimum and maximum instances. Review the next question for more details on auto-scaling.

What is auto-scaling, or how does Valtix scale with traffic?

During the Gateway creation you are given an option to choose the number of firewall instances to run. The minimum number is always 1. The maximum number can be up to 10. This is per Availability Zone (AZ). If you start with a minimum of 2 and you have 2 AZs, then 4 instances in total would be running in your account. The Controller keeps track of the usage of the instances and once the firewall gets busier, it automatically creates new instances until it reaches the maximum number. Once the traffic slows down, the instances are automatically deleted. You are creating resources on demand and pay only when they are used/required. If there is no usage for the instances, they are deleted and you are not charged.

How do I prepare my AWS environment to get started with Valtix?

Multicloud Defense security service works in Hub mode or Edge mode. Hub mode is used when you have multiple VPCs that you want to protect. AWS Transit Gateway is used to attach all the VPCs. For this mode, you need to provide a non-overlapping CIDR so Multicloud Defense can create a new Service VPC to deploy the Gateways. The Service VPC is completely managed by the Multicloud Defense Controller.

In Edge mode deployment, the Gateway is installed in the same VPC as your applications. For this deployment, Multicloud Defense needs 2 public subnets (management and datapath) and 2 security-groups (management and datapath). Both the security groups need rules to allow outbound traffic. datapath security-group can allow all the traffic or you can enable specific ports that you configure in services on Multicloud Defense Controller.

For both the modes of deployment, Multicloud Defense needs several IAM roles: Cross account IAM role for Controller to access your AWS account, IAM role assigned to Gateway instances to access KMS, Secrets Manager and S3 to write PCAP files.

Multicloud Defense provides a CloudFormation template that assists in creating the IAM roles and contains details about the permissions. This is detailed in the User Guide IAM Role documentation.

How do I prepare my Azure environment to get started with Valtix?

The Multicloud Defense solution works in Hub mode or Edge mode. Hub mode is used when you have multiple VNets that you want to protect. Azure UDRs are used for this purpose.

In Edge mode deployment, the Gateway is installed in the same VNet as your applications. For this deployment, Multicloud Defense needs 2 public subnets (management and datapath) and 2 network security-groups (management and datapath). Both the security groups need rules to allow outbound traffic. Datapath security-group can allow all the traffic or you can enable specific ports that you configure in services on Multicloud Defense Controller.

For both the modes of deployment, Multicloud Defense needs an Azure Active Directory Id (tenant id), Subscription id, an Application in Active Directory (AD) with a Client Key and Secret, a custom role assigned to the application that has permissions to create resources, access vault etc.

Please review the User Guide documentation for more information.

What is Sessionid in flow logs?

Multicloud Defense Gateways act as a proxy for both Ingress and Egress. In the Ingress scenario, an external user from the internet accesses the Gateway endpoint and the Gateway initiates a new session to the backend (target). These are 2 different traffic flows. Sessionid correlates these 2 flows and ties them together for display in the flow logs.

How do I provide the certificate for my proxied applications?

A TLS Decryption Profile needs to defined where there is an option to generate a self-signed certificate or import contents of an already generated certificate.

TLS Decryption Profile can be configured as listener decryption profile for reverse proxy for applications proxied for in the backend.

TLS Decryption Profile can be configured as rootCA decryption profile for forward proxy where the rootCA certificate and private key have been installed on the client which is egressing to the internet via the forward proxy.

How do I protect my private keys without giving to Valtix Controller?

In the definition of the TLS decryption profile there are multiple ways to import the private key.

- Import the contents in clear.
- AWS KMS encrypted private key.
- AWS Secrets Manager secret name.
- Credstash key name from the given credential store.
- Azure key name from the given key vault.

(b),(c),(d),(e) are the recommended choices if you do not wish to leave private keys with Multicloud Defense Controller.

What are all the different protocol options I see in the Reverse Proxy Service?

Table 1: What are all the different protocol options I see in the Reverse Proxy Service?

Proxy Type	Decryption Profile	Frontend Protocol	Backend Protocol
TCP-TCP	No	TCP	TCP
TLS-TLS	Yes	TCP	TCP
HTTP-HTTP	No	TCP	HTTP

Proxy Type	Decryption Profile	Frontend Protocol	Backend Protocol
HTTPS-HTTPS	Yes	TCP	HTTPS
HTTPS-HTTP	Yes	TCP	HTTP
WEBSOCKET-WEBSOCKET	No	TCP	WEBSOCKET
WEBSOCKETS-WEBSOCKETS	Yes	TCP	WEBSOCKET_S

How should I configure Reverse proxy for an SSH application?

Use proxy type TCP-TCP.

What's the difference between HTTPS and TLS in the Reverse Proxy Target?

In the TLS proxy, the TCP payload received from the client or server is preserved byte for byte during decryption followed by re-encryption. There are applications like RDP which depend on NTLM where this preservation of TCP payload bytes is mandatory.

In the HTTPS proxy, proxy terminates the HTTP connection and moves the HTTP payload from one leg of the proxy to the other leg with additional proxy headers attached to the HTTP PDU. HTTPS proxy lets you send responses at HTTP level for deep packet security related actions. It also lets specification of rate limiters at the URL level.

How do I apply the same policy rules to multiple Gateways?

Policy Rules are always defined in the context of a Policy Rule Set. A Policy Rule Set defines a set of rules. This Policy Rule Set can be associated with multiple Gateways. A Gateway can only have one Policy Rule Set.

My target Application IPs are different in each region/can change, how can I configure my backend target in the service?

Define a user-defined-tag which is associated with the instances where the application is running. Use this tag to define a backend address object. Associate this backend address object as the target of a service. Controller automatically maintains the membership of the set of IPs of the instances with that tag. Membership changes are also automatically handled by the Controller when the instances with that user- defined-tag come up and go down or the IP address on the instance with that user-defined-tag changes.

What is SNI in the Service Object?

SNI stands for Server Name Indication. There is a TLS client hello extension called server_name which contains the FQDN of the server. This can be then be used in the definition of the service object to route the

traffic to the appropriate backend using this. The set of SNIs defined in the service object can also be used to allow access from the clients to only those services.

Examples of SNIs in the definition of service object : service1.enterprise.com

This makes sense only for reverse proxy where the backend services and the associated FQDNs are well-defined.

My backend/target hosts multiple websites. I want them proxied on the same port by the Valtix Gateway. How can I achieve this?

Define a Service Object per website. Each Service Object uses the same listener port and the website SNI = website FQDN.

I have multiple web backends/targets that need to be proxied by the Gateway. How do I configure this?

Define a Service Object per web backend with the same listener port and

- SNI = web backend FQDN and
- target = backend FDQNs or ALB FQDN frontending the web backends

What's the relation between Decryption Profile and Certificate?

Decryption profile is one-one with a certificate. This decryption profile can be associated with Service Objects which in turn are used as part of the Policy Rules. This level of indirection helps in easier certificate management to renew expired certificates or to rotate certificates on a periodic basis, updating the decryption profile only, without having to update all the Policy Rules/Services dependent on this certificate.

I need different IPS protection rules for each of my backends. How do I do this?

Every Gateway can have only 1 IPS profile. Even though this is configured at the rule level, it is per Gateway. So it is not possible to have multiple IPS profiles using the same Gateway. You need to create multiple Gateways.

Where are the IPS rules, and how often do you update? Are the updates automatically pushed to the Gateway?

Cisco TALOS Rules are periodically polled at bi-weekly intervals and even shorter time periods based on critical rule update notifications. These updates are automatically made available in the Controller to customers who then have the option of choosing the right ruleset version to push to the Gateways.

IPS Profile has many configuration options. Can you elaborate?

IPS Profile lets the user choose the set of rules from the ruleset based on the SNORT policy, category or class-type.

In addition there is option to enable threat based PCAP file creation.

Rule suppression is provided for false positives based on trusted source CIDRs.

Rule level event filters can be enabled for chatty rules or a global profile level event filter across all rules.

I want to get PCAP (packet captures) files of every attack, is it possible?

Yes. Enable threat based PCAP checkbox in the Network Intrusion Profile or Web Protection Profile.

I have my own log analysis infrastructure. Can I forward the logs to it?

Yes. Syslog, Splunk and DataDog are supported. Review the User Guide documentation for full details.

I configured a Reverse Proxy to a backend application. What else do I need to do?

1. Change DNS record to point to the Multicloud Defense Gateway's FDQN.
2. Change existing application load balancer to private to avoid direct public access

What's the DNS profile and records and why would I use it?

Web based applications in AWS are typically referred to by an internal FQDN dynamically generated when creating a load balancer. In order for Multicloud Defense to be in the ingress path of that application for inspection, we would advise customers to update the DNS record of that application to refer to the Multicloud Defense Gateway.

For example, a DNS record for app.xyz.com points to the CNAME of the internal application load balancer. With Multicloud Defense Gateway to be in the ingress path of this application, we would update the DNS record to point to a CNAME of the Multicloud Defense Gateway endpoint. Multicloud Defense DNS profile allows one to specify the Route53 domain name associated with the application where you can configure this application's record and select the appropriate Multicloud Defense Ingress Gateway from the list of Gateways.

The documentation set for this product strives to use bias-free language. For purposes of this documentation set, bias-free is defined as language that does not imply discrimination based on age, disability, gender, racial identity, ethnic identity, sexual orientation, socioeconomic status, and intersectionality. Exceptions may be present in the documentation due to language that is hardcoded in the user interfaces of the product software, language used based on standards documentation, or language that is used by a referenced third-party product.