Cisco Multicloud Defense User Guide

Cisco Multicloud Defense
Updated July 8, 2026

© 2023-2026 Cisco Systems, Inc. All rights reserved.

Full Cisco Trademarks with Software License

Full Cisco Trademarks with Software License

THE SPECIFICATIONS AND INFORMATION REGARDING THE PRODUCTS IN THIS MANUAL ARE SUBJECT TO CHANGE WITHOUT NOTICE. ALL STATEMENTS, INFORMATION, AND RECOMMENDATIONS IN THIS MANUAL ARE BELIEVED TO BE ACCURATE BUT ARE PRESENTED WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED. USERS MUST TAKE FULL RESPONSIBILITY FOR THEIR APPLICATION OF ANY PRODUCTS.

THE SOFTWARE LICENSE AND LIMITED WARRANTY FOR THE ACCOMPANYING PRODUCT ARE SET FORTH IN THE INFORMATION PACKET THAT SHIPPED WITH THE PRODUCT AND ARE INCORPORATED HEREIN BY THIS REFERENCE. IF YOU ARE UNABLE TO LOCATE THE SOFTWARE LICENSE OR LIMITED WARRANTY, CONTACT YOUR CISCO REPRESENTATIVE FOR A COPY.

The Cisco implementation of TCP header compression is an adaptation of a program developed by the University of California, Berkeley (UCB) as part of UCB's public domain version of the UNIX operating system. All rights reserved. Copyright © 1981, Regents of the University of California.

NOTWITHSTANDING ANY OTHER WARRANTY HEREIN, ALL DOCUMENT FILES AND SOFTWARE OF THESE SUPPLIERS ARE PROVIDED “AS IS" WITH ALL FAULTS. CISCO AND THE ABOVE-NAMED SUPPLIERS DISCLAIM ALL WARRANTIES, EXPRESSED OR IMPLIED, INCLUDING, WITHOUT LIMITATION, THOSE OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT OR ARISING FROM A COURSE OF DEALING, USAGE, OR TRADE PRACTICE.

IN NO EVENT SHALL CISCO OR ITS SUPPLIERS BE LIABLE FOR ANY INDIRECT, SPECIAL, CONSEQUENTIAL, OR INCIDENTAL DAMAGES, INCLUDING, WITHOUT LIMITATION, LOST PROFITS OR LOSS OR DAMAGE TO DATA ARISING OUT OF THE USE OR INABILITY TO USE THIS MANUAL, EVEN IF CISCO OR ITS SUPPLIERS HAVE BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

Any Internet Protocol (IP) addresses and phone numbers 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.

All printed copies and duplicate soft copies of this document are considered uncontrolled. See the current online version for the latest version.

Cisco has more than 200 offices worldwide. Addresses and phone numbers are listed on the Cisco website at www.cisco.com/go/offices.

Cisco and the Cisco logo are trademarks or registered trademarks of Cisco and/or its affiliates in the U.S. and other countries. To view a list of Cisco trademarks, go to this URL: https://www.cisco.com/c/en/us/about/legal/trademarks.html. Third-party trademarks mentioned are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (1721R)

Topics included

1 About Multicloud Defense....13

Multicloud Defense....14

Multicloud Defense Naming Conventions....14

Supported Regions....14

Recommended Versions of Multicloud Defense Components....14

Multicloud Defense in Cisco Security Cloud Control....15

Multicloud Defense Components....15

Multicloud Defense Controller Dashboard....17

My Profile Information....19

Multicloud Defense 90-Day Free Trial....20

Cloud Explorer....22

2 Setup with the Multicloud Defense Wizard....25

Connect Cloud Account....26

Connect an AWS account....26

Connect Azure account....27

Connect Google Cloud Platform account....28

Connect to an OCI Account....29

Enable Traffic Visibility....31

Enable traffic for an AWS account....32

Enable traffic for an Azure account....32

Enable traffic for a GCP project....33

Secure Your Account....34

Centralized Model: Add a VPC or VNet....34

Distributed Model....35

3 AWS....37

AWS Overview....38

Connect AWS Account to Multicloud Defense Controller from the Multicloud Defense Dashboard....38

CloudFormation Outputs....39

Roles Created by Multicloud Defense....40

4 Azure....45

Prepare Your Azure Account....46

Register Application in Microsoft Entra ID....46

Create a Custom Role to Assign to the Application....47

Connect an Azure Subscription to the Multicloud Defense Controller from the Multicloud Defense Dashboard....48

Accept Marketplace Terms....49
Connect an Azure Subscription to the Multicloud Defense Controller from the Multicloud Defense Dashboard....49
VNet Route Tables for your Azure Subscription....50
Roles Created by Multicloud Defense....51
Post-Onboarding Procedures....52
Azure VNet Setup....52

5 GCP....55
GCP Overview....56
Overview of Creating a GCP Controller Service Account....56
Create a GCP Firewall Service Account....57
Connect a GCP Project to the Multicloud Defense Controller from the Multicloud Defense Dashboard....57
Roles Created by Multicloud Defense....58

6 OCI....61
Prepare Your OCI Account....62
Connect the Oracle OCI Tenant to the Multicloud Defense Controller from the Multicloud Defense Dashboard....63

7 Remove a Cloud Service Provider From Multicloud Defense 65
Delete a GCP project from Multicloud Defense....66
Delete an AWS account from Multicloud Defense....67
Delete an Azure account from Multicloud Defense....67
Delete an OCI Account From Multicloud Defense....69

8 AI Defense....71
AI Defense Integration with Multicloud Defense....72

9 Asset and Inventory Discovery....73
Discovery Summary....74
Inventory....74
Applications....74
Discovered Assets....75
Enable Asset Discovery and Inventory....76
Security Insights....76
Types of Security Insights....77
Rules and Findings....80
Rules and Findings....80
Pre-Defined Rules....80
Custom Rules....81
Findings....81

10 Multicloud Defense Gateway and Service VPCs or VNets....83

Supported Gateway Use Cases....85
Firewall Threat Defense Virtual....85
Multicloud Defense Egress Gateways....90
Multicloud Defense Ingress Gateways....94
Multicloud Defense East-West Gateways....95
Multicloud Defense Distributed Gateways....96
Gateways Deployed in Centralized / Hub Mode....97
Advanced Gateway Configuration: Use Your Own Load Balancer....98
Configure Service VPCs and Service VNets....100
Create a Service VPC or VNet....101
Secure Spoke VPC or VNet....102
Configure Your Gateway....105
Before You Begin....105
Add a Multicloud Defense Gateway....108
Create a Firewall Threat Defense Virtual Gateway....111
Edit a Multicloud Defense Gateway....113
Upgrade the Multicloud Defense Gateway....114
Upgrade Your Firewall Threat Defense Virtual Device....114
Abort a Multicloud Defense Gateway....115
Enable a Multicloud Defense Gateway....115
Disable a Multicloud Defense Gateway....116
Export a Multicloud Defense Gateway....116
Delete a Multicloud Defense Gateway....117

11 Site-to-Site VPN Tunnel Connection....119
Prerequisites and Limitations for Site-to-Site VPN Tunnels....120
Enable VPN Within the Gateway....121
Create a Site-to-Site VPN Connection....122
Edit a Site-to-Site VPN Tunnel....123
Clone a Site-2-Site VPN Tunnel Connection....124
Delete a VPN Tunnel Connection....125

12 Configure a Virtual Network Connection for Azure Virtual WAN....127
Overview of Virtual WAN....128
Guidelines for Virtual WAN Connections to Azure vHub....128
Create a Service VPC with a Virtual WAN Attachment....128
Modify a Service VPC with a Virtual WAN Attachment....129

13 Rules and Rule Sets....131
Policy Management....132
Policy Rule Set Gateway and Management....132
Rules....132
Rule Sets and Rule Set Groups....133

Create Policy Rule Set....134
Create a Rule in a Rule Set....135
Create a Policy Rule Set Group....139

14 Objects....141
About the Multicloud Defense Connector....143
Import Objects From Security Cloud Control....143

15 Address Objects....145
Address Objects....146
Src/Dest....146
Reverse Proxy Target Address Object....150
System Objects....151
Create a Source/Destination Address Object....151
Create a Reverse Proxy Target Address Object....153
Edit Address Objects....154
Clone Address Objects....154
Delete Address Object....154
View Details....155

16 FQDN Objects....157
FQDN Match Object....158
Standalone vs. Group....158
Create Standalone FQDN Match Object....158
Create Group FQDN Match Object....159
Associate the Object....159

17 Service Objects....161
Reverse Proxy Service Object (Ingress)....162
Forward Proxy Service Object (Egress / East-West)....163
Forwarding Service Object (Egress / East-West)....164

18 Certificates and Keys....167
Certificates and Keys....168
Import Certificate....168
AWS - KMS....169
AWS KMS - Encrypt Private Key....169
AWS - Secrets Manager....169
Azure Key Vault....170
GCP - Secret Manager....170
Server Certificate Validation....170
Server Certificate Validation in the TLS Decryption Profile....171

Server Certificate Validation in the FQDN Service Object....172
Certificate for Custom Root CA....173

19 Certificate and Keys Tech Notes....175
Azure Key Vault....176
Create Azure Key Vault and User-assigned Managed Identities....176
Store Private Keys into Azure Key Vault Secrets....176
Configure Multicloud Defense to Retrieve Private Key....176
Generate a Self-Signed Root CA....177
Generate a Certificate Signed by your Self-Signed Root CA....177
Generate an Intermediate CA Signed by Your Root CA....178
App Certificate signed using the Intermediate CA....178
Install Root CA as Trusted CA on the Hosts....178

20 Types of Traffic....181
Enable DNS Logs....182
AWS: Enable DNS Logs....182
GCP: Enable DNS Logs....182
Azure: DNS Logs....183
Enable VPC Flow Logs....184
AWS: Enable VPC Flow Logs....184
GCP: Enable VPC Flow Logs....184
Azure: Enable VNet Flow Logs....186

21 Security Profiles....187
Decryption Profile....188
Create a Decryption Profile....189
Network Intrusion (IDS/IPS) Profile....190
Create an IPS/IDS Profile....190
Data Loss Prevention (DLP) Profile....192
Create a Data Loss Prevention Profile....192
Anti-Malware Profile....193
Create an Anti-Malware Profile....193
Web Application Firewall (WAF) Profile....194
Create WAF Profile....194
URL (Uniform Resource Locator) Filter Profile....197
Create the URL Filtering Profile....198
Fully Qualified Domain Name Filter Profile....199
Create a Standalone FQDN Filter Profile....201
Create a Group FQDN Filter Profile....201
Malicious IP Profile....202
Create a Malicious IP Profile....203
IP Reputation....203

22 Gateway Profiles.....205
Packet Capture Profile.....206
Create a Packet Capture Profile.....206
Log Forwarding Profile.....206
Create a Standalone Log Forwarding Profile.....207
Create a Log Forwarding Group.....207
Gateway Metrics Forwarding Profile.....208
Create a Standalone Metrics Forwarding Profile.....208
Create a Group Metrics Forwarding Profile.....209
(Preview Only) Network Packet Broker Profile.....209
Create Network Packet Broker Profile.....209
Network Time Protocol Profile.....210
Create a Profile.....211
IPSec Profile.....211
Create an IPSec Profile.....211
BGP Profile.....212
Create a BGP Profile.....213

23 Profile Actions.....215
View a Profile Details.....216
Edit a Standalone Metrics Forwarding Profile.....216
Edit a Group Profile.....216
Add a Gateway Association to a Profile.....217
Remove a Gateway Association.....217
Delete a Profile.....217

24 FQDN and URL Filtering Categories.....219
FQDN / URL Filtering Categories.....220
Malicious Categories.....220
Full List of Categories.....221
Associating a Filtering Profile with a Policy Ruleset Rule.....222
Cisco Talos Intelligence URL / IP Lookup Tool.....222

25 Flow Analytics.....223
Flow Analytics - Traffic Summary.....224
Flow Analytics - All Events.....226
Event Logs.....228
Firewall Events.....229
Network Threats.....230
Web Attacks.....232
URL Filtering.....233
FQDN Filtering.....235

HTTPS Logs....236
VPN Logs....237
EndUser Logs....238
HTTP2 Logs....239
Investigate summary page....240

26 Network Analytics....243
Stats....244
Total Bandwidth....244
CPU Usage....244
Memory Usage....245
Connection Rate....245
HTTP Request Rate....245

27 Audit and System Logs....247
Audit Logs....248
Search Filter....249
System Logs....250
Search Filter....251

28 Threat Research....253
Network Intrusion....254
Web Protection....254
Malicious Sources....255

29 Cloud Visibility Reports....257
Generate a Discovery Report....260
Generate a Threat And Cloud Analytics Report....260

30 Alerting Overview....263
Alert Services Overview....264

31 Alert Destinations / SIEMs....265
Datadog....266
Create an Alert Profile Service....266
Create an Alert Rule....266
Microsoft Sentinel....268
Create an Alert Profile Service....268
Create an Alert Rule....268
PagerDuty....269
Create an Alert Profile Service....270
Create an Alert Rule....270

ServiceNow....272
Create an Alert Rule....272
Create an Alert Profile Service....273
Slack....274
Create a Slack Alert Rule....274
Create an Alert Profile Service....275
Microsoft Teams....275
Create a Microsoft Teams Alert Profile Service....275
Create a Microsoft Teams Service Rule....276
Webex....277
Create an Alert Profile Service....277
Create an Alert Rule....278
Splunk....279
Create a Splunk Profile Service....279
Create a Splunk Rule....280

32 Log Forwarding Overview....283
Security Events and Traffic Logs....284
Create a Standalone Event or Traffic Log Profile....285
Edit a Standalone Event or Traffic Log Profile....285
Create a Group Event or Traffic Log Profile....286
Edit a Group Event or Traffic Log Profile....286
View an Event or Traffic Log Forwarding Profile....286
Delete an Event or Traffic Log Profile....287
Discovery Logs....287
Create a Standalone Discovery Log Profile....288
Edit a Standalone Discovery Log Profile....288
Create a Group Discovery Log Profile....288
Edit a Group Discovery Log Profile....289
View a Discovery Log Profile Details....289
Add a Discovery Log Profile with a Cloud Account....289
Remove a Discovery Log Profile from a Cloud Account....289
Delete a Discovery Log Profile....290
Gateway Metrics Forwarding Profile....290
Create a Standalone Metrics Forwarding Profile....290
Edit a Standalone Metrics Forwarding Profile....291
Create a Group Metrics Forwarding Profile....291
Edit a Group Profile....292
Delete a Profile....292
Add an Event, Traffic Log Forwarding Profile, or Metrics Forward Profile to a Gateway....293
Remove an Event, Traffic Log Forwarding Profile, or Metrics Forward Profile from a Gateway....293

33 Log Forwarding Destinations / SIEMs....295
AWS S3 Bucket....296

Datadog....297
GCP Logging....298
Microsoft Sentinel....301
Splunk....302
Sumo Logic....303
Syslogs....304
Webhook....306

34 Management....309
Management....310
API Keys....310
Account Level Settings....311
System....313
Metering....314
Alert Profiles....315
Services....315
Alerts....318

35 Manage Your Multicloud Defense Account....321
Account (Multicloud Defense Tenant)....322
Understand user roles in Security Cloud Control Firewall Management....322
Roles in Multicloud Defense....322

36 Cloud Accounts....325
Cloud Accounts....326
Add Account....326
Manage Inventory....326
Edit a Cloud Account....327
Update Log Profile for a Cloud Account....327
Export a Cloud Account....327
Delete a Cloud Account....328
Inventory....328

37 Compliance Certificates and Awards....329
Compliance Certificates....330

38 Troubleshoot Connecting Your Account....331
Manually Onboard an Account....332
Manually Onboard a GCP Project....332
Manually Onboard an Azure Subscription....340
Graceful Termination of Connections....343
Terraform Onboarding Scripts for Cloud Accounts....344

About Terraform.....344
Terraform Repository.....344
Exporting Configuration as Terraform Block.....345

1 About Multicloud Defense

Topics:

- Multicloud Defense
- Multicloud Defense Components
- Multicloud Defense Controller Dashboard
- Multicloud Defense 90-Day Free Trial
- Cloud Explorer

Learn about Multicloud Defense and the cloud security capabilities it provides for connected cloud environments.

Multicloud Defense

An overview of Cisco Multicloud Defense, its core components, and its support for AWS, Azure, Google Cloud Platform, and Oracle Cloud Infrastructure environments.

Multicloud Defense is a comprehensive security solution consisting of two primary components, the Multicloud Defense Controller and Multicloud Defense Gateway. These components collaborate to establish a secure multicloud environment.

Multicloud Defense currently supports Amazon Web Services (AWS), Azure, Google Cloud Platform (GCP), and Oracle OCI cloud accounts. The range of support for these platforms varies.

Multicloud Defense offers a sophisticated, streamlined security framework harmonizing controller orchestration, gateway communication, and optimized datapath processing for robust and efficient multicloud protection.

Multicloud Defense Naming Conventions

Find the Cisco Multicloud Defense naming limits for gateways and objects across supported cloud environments.

Multicloud Defense interacts with a variety of cloud service providers. To provide a universal experience across the platforms, Multicloud Defense limits the character count when you create gateways and objects.

Gateways and objects that exist outside of Multicloud Defense prepend `ciscomcd` as a prefix in their names, which may cause issues if the original gateway or object name is too long.

Consider the character limitations when naming your gateways or objects inside and outside of Multicloud Defense:

Table 1: Character Limitation for Naming Conventions

Multicloud Defense Feature	Character Limit
Gateway Instance	55
Object Name	63

Note

The values that are listed indicate the character limit for names **without** the prepended Multicloud Defense tag. You are not responsible for including the tag when you name the gateway or object.

Supported Regions

Find the cloud regions that Cisco Multicloud Defense supports for AWS, Azure, Google Cloud Platform, and Oracle Cloud Infrastructure.

We currently support all regions for any commercial cloud service provider region for AWS, Azure, GCP and OCI. Refer to your cloud service provider support documentation for detailed information.

If a region is not supported, or a new region needs support, contact Cisco Support to add support for the region.

Recommended Versions of Multicloud Defense Components

Learn why Cisco recommends keeping Multicloud Defense components updated and where to find current release information.

Ensure that all Multicloud Defense components are updated with the latest versions to benefit from enhancements, new features, and bug fixes. For more information about available updates and the issues addressed by each release, refer to the Cisco Multicloud Defense Release Notes.

Multicloud Defense in Cisco Security Cloud Control

Learn how to access and activate Cisco Multicloud Defense in Cisco Security Cloud Control, whether you are a new customer or an existing user.

Multicloud Defense is now hosted in Cisco Security Cloud Control. Security Cloud Control is a platform that allows you to manage your security products and achieve security outcomes from a single integrated interface. From Security Cloud Control platform, you can manage Multicloud Defense along with other Security products.

When you enroll in a Multicloud Defense, Security Cloud Control creates an account for your tenancy by default to better manage your enterprises across the board. The Security Cloud enterprise supports these cases:

1. You have purchased a license and already have a Multicloud Defense account, and
2. You have purchased a license but do **not** have a Multicloud Defense account.

New customers or users of Multicloud Defense can complete these steps in Security Cloud Control.

1. Create an organization. For more information, refer to the **Create an Organization** topic in the **Organizations and Regions** section of the Getting Started Guide for New Customers of Security Cloud Control.
2. Buy a subscription license. Once you purchase the license, you or the designated system administrator receives an email with a subscription **claim code**. Do not lose this email.
3. Claim the subscription. Enter the claim code in the **Claim Subscription** section of Security Cloud Control application and claim the product. You will receive a confirmation email that Multicloud Defense has been activated on Security Cloud Control. For more information, refer to the **Claim Your Product Subscriptions** section in the Getting Started Guide for New Customers of Security Cloud Control.

Multicloud Defense is provisioned and is available on the default landing page of Security Cloud Control. The navigation pane on Security Cloud Control contains the organization details and menu elements, and the central area contains the dashboard widgets.

> **Note**
>
> Customers who provision Multicloud Defense after August 6, 2025 and want to use **Object Sharing** will need to contact Cisco Technical Assistance Center (Cisco TAC) to enable this feature.

For existing users of Multicloud Defense migrating to Security Cloud Control, refer to the Getting Started Guide for Existing Customers of Security Cloud Control.

Multicloud Defense Components

An overview of the Cisco Multicloud Defense Controller and Gateway components that help you manage and protect applications and workloads across public cloud environments.

Multicloud Defense uses a common principle in public clouds and Software-Defined Networking (SDN) which decouples the control and data plane, translating to two solution components:

- Multicloud Defense Controller
- Multicloud Defense Gateway

Multicloud Defense Controller

The Multicloud Defense Controller is a highly reliable and scalable centralized controller that provides the management and control plane. This runs as software-as-a-service (SaaS) and is fully managed and maintained by Multicloud Defense. You can access a web portal to utilize the Multicloud Defense Controller. Alternatively, you may choose to use the Multicloud Defense provider for terraform to instantiate security into the DevOps/DevSecOps processes.

Multicloud Defense Gateway

The Multicloud Defense Gateway is an auto-scaling fleet of Multicloud Defense software deployed as Platform-as-a-Service (PaaS) into the customers public cloud account(s) by the Multicloud Defense Controller. This provides advanced, inline security protections that defend against external attacks, prevent egress data exfiltration and lateral movement of attacks. Multicloud Defense Gateways include functionality for Transport Layer Security (TLS) decryption, Intrusion Detection and Prevention (IDS/IPS), Web Application Firewall (WAF), antivirus filtering, Data Loss Prevention (DLP) and Fully Qualified Domain Name (FQDN) or Uniform Resource Locator (URL) filtering capabilities.

To facilitate the auto-scaling, the Multicloud Defense Gateway offers:

- **Autoscaling and Self-healing**: Operate as autoscaling, self-healing Platform-as-a-Service (PaaS), serving as inline network-based security enforcement nodes. For more information about auto-scaling and how it operates within the construct of the gateway, see below.
- **Simplified Management**: Eliminate the need for constructing virtual firewalls, configuring high-availability setups, or managing software installations.
- **Advanced Security Profiles**: Implement granular security profiles within a single-pass datapath pipeline, negating the need for traffic offloading to third-party engines.

Important

The Multicloud Defense Gateway does not currently support Internet Protocol (IP) fragmentation because of cloud service provider load-balancer limitations. We strongly recommend you adjust the Maximum Transmission Unit (MTU) size so it is consistent across the network to avoid the need for fragmentation.

Multicloud Defense SaaS Controller

Multicloud Defense offers a sophisticated, streamlined security framework, combining robust controller orchestration, gateway communication, and optimized datapath processing to provide efficient and comprehensive multicloud protection. This solution helps your organization secure your cloud workloads and applications from cyber threats while maintaining flexibility and scalability in your cloud environment.

- **SaaS-Based Management**: The Software-as-a-Service (SaaS) controller manages the gateway stack and includes an API Server for orchestration of CSP load balancers (LBs) and gateway instances.
- **Dynamic Scaling**: Facilitates dynamic horizontal scaling through instance additions and removals from the load balancer's "target pool," monitored for high availability.
- **Continuous Communication**: Engages in continuous communication with Cloud Service Provider (CSP) accounts to keep security policies up-to-date.

Communications

Multicloud Defense Gateways engage in continuous communication, approximately every 3 seconds, with the Multicloud Defense Controller, transmitting health status and policy updates. This enables proactive health reporting, gateway replacement, and scalability adjustments as needed.

Optimized Gateway Instances

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 updates without disrupting traffic flow.

Gateway Auto-scaling

Autoscaling in Multicloud Defense is triggered based on these usage thresholds:

- **CPU Usage**
  - Scale Out: Triggered when CPU usage exceeds 95%.
  - Scale In: Occurs when CPU usage falls below 40%.
- **Memory Usage**
  - Scale Out: Triggered when memory usage exceeds 85%.
  - Scale In: Occurs when memory usage falls below 40%.
- **Bandwidth Usage**
  - Scale Out: Triggered when bandwidth usage exceeds 75%.
  - Scale In: Occurs when bandwidth usage falls below 40%.
- **Connection Usage**
  - Scale Out: Triggered when connection usage exceeds 75%.
  - Scale In: Occurs when connection usage falls below 40%.
- **Load Sustain Period**: Autoscaling actions are triggered if the load conditions are sustained for 120 seconds.
- **Scale-In and Scale-Out Periods**: Both scale-in and scale-out actions have a 120-second assessment period to ensure consistent load conditions.

Advanced Security Profiles

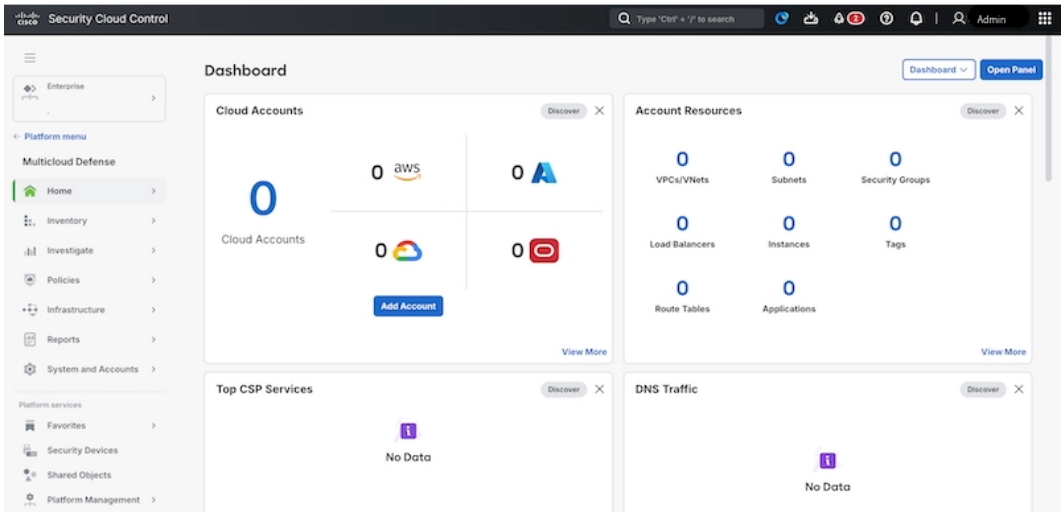
Multicloud Defense Controller implements granular security profiles within the single pass datapath pipeline, catering to evolving traffic needs. Customers have the flexibility to enable or disable **Advanced Security Profiles** as required. The pipeline's single pass architecture negates the need for traffic offloading to third-party engines. For instance, full TLS decryption is selectively triggered within the pipeline, ensuring efficient handling without unnecessary data transfers.

In essence, Multicloud Defense offers a sophisticated and streamlined security framework, harmonizing controller orchestration, gateway communication, and optimized datapath processing for a robust and efficient multicloud protection mechanism.

Multicloud Defense Controller Dashboard

Learn how to use the Cisco Multicloud Defense dashboard to monitor cloud accounts, resources, traffic, threats, security status, and system activity.

The dashboard of the Multicloud Defense Controller has a multitude of widgets to give you a quick snapshot of the current state of your accounts, account resources, and top-hitting policies or profiles.



You can drag-and-drop any of the widgets to customize and organize the dashboard to meet your needs. Click "x" at the top-right corner of a widget to remove it from your dashboard. You can click **View More** which is located within the widget, to go directly to the associated page to add accounts. At the top of each widget is a tag or label that indicates what function the widget serves, such as discovery, detection, deployment, or defending.

The dashboard generates these widgets by default.

Cloud Accounts

This provides a high-level view of all the cloud accounts you have connected to the Multicloud Defense Controller,and the number of accounts for each of the identified cloud service providers.

Click **Add Account** to open a connection wizard for onboarding a new cloud service provider.

Account Resources

This is a general list of allocated resources across all of your connected cloud accounts. It displays how many of these resources are currently used:

- VPC/VNets
- Subnets
- Security Groups
- Load Balancers
- Instances
- Tags
- Route Tables
- Applications.

Top CSP Services

The Multicloud Defense Controller connects to the cloud service providers and displays a top-down view that generalizes Domain Name System (DNS) traffic.

DNS Traffic

Similar to Top CSP Services, this DNS Traffic widget offers a limited view of current DNS traffic for the cloud service providers that are actively processing traffic. We recommend expanding the widget to the full discovery scope for more insight.

VPCs/VNets with Malicious Traffic

This widget displays any recent VPC or VNet that has encountered malicious traffic. For a comprehensive list of events and attacks, expand the widget and view the traffic.

Top Ports with Malicious Traffic

This small snapshot displays which ports within your cloud accounts have the most hits against malicious traffic.

Security Considerations

The Security Considerations widget summarizes the applications, VPCs or VNets, and associated gateways that are not protected by Multicloud Defense.

System Logs

The System Logs window provides a recent history of logs that catalog the accounts affected, the gateway associated, the severity of events or attacks and more. We strongly recommend utilizing this widget or the whole System Log page as a valuable resource.

Top Applications

This widget displays the topmost applications used across all cloud service providers.

Threats

This view provides a graph depicting the last seven days of traffic and the amount of the incoming traffic was categorized as threats.

Top Countries by Threat

This horizontal bar chart shows the top 10 countries with the most events and breaks the volume down by time increment.

Exfiltration Attempts

This view is a general display of egress data exfiltration that have occurred on the cloud service providers currently connected to Multicloud Defense.

My Profile Information

Learn how to view your Cisco Multicloud Defense profile information, including your account email, user role, tenant, and assigned cloud accounts.

The User Profile page details your user information. To access this page, select your username from the **Admin** drop-down list in the upper right corner of the Multicloud Defense dashboard to display the information:

- Your name.
- The email address associated with your Multicloud Defense account.
- The user role you currently have.
- The name of the tenant currently logged in.
- Assigned accounts.

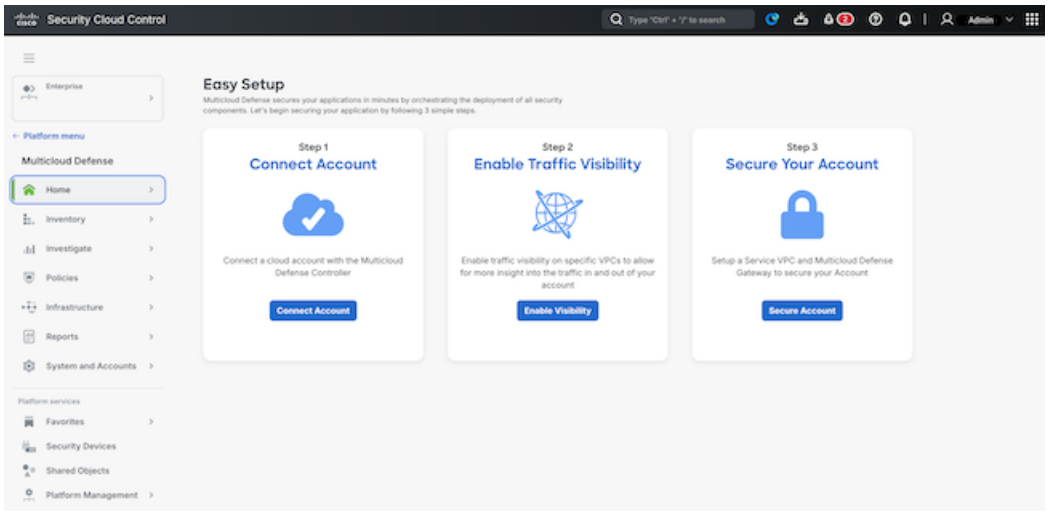
Use this page for general information or when seeking assistance from Cisco Support.

Multicloud Defense 90-Day Free Trial

Learn about the Multicloud Defense 90-Day Free Trial and how to evaluate its cloud security capabilities.

When you log in to your Security Cloud Control tenant, you will see a wizard that guides you through connecting your cloud accounts to Multicloud Defense so that you can manage them with a free 90-day trial of Multicloud Defense Controller. The 90-day trial experience offers the full functionality of a paid subscription to Multicloud Defense Controller.

Figure 1: Multicloud Defense Easy Setup



Click **Get Started** to begin your 90-day trial. This begins provisioning the Multicloud Defense Controller.

> **Note**
>
> Easy Setup supports the following Cloud Providers:
>
> - Account Onboarding: AWS, Azure, GCP, OCI
> - Enable Traffic Visibility: AWS (Flow Logs, DNS Query Logs), Azure (Flow Logs), GCP (VPC flow logs, DNS Query Logs)
> - Create Services VPC/VNet: AWS, Azure, GCP
> - Create Gateways: AWS, Azure, GCP
>
> Although Services VCN orchestration is not supported for OCI (requires the user to create the Services VCN using the Cloud Provider console), the gateway orchestration is supported in Multicloud Defense. Open Multicloud Defense and navigate to **Infrastructure** > **Gateways** > **Gateways**.

Connect Cloud Accounts

After Multicloud Defense Controller is provisioned, the **Connect a Cloud Account** page opens and you can connect any of the types of cloud accounts that are shown to the Multicloud Defense Controller.

The first step is to onboard a set of one or more cloud accounts. This allows the Multicloud Defense Controller to interact with each account by discovering inventory and traffic, orchestrating security deployment, and creating and managing policy.

To connect to your cloud accounts, follow these instructions:

- AWS on page 37
- Azure
- GCP on page 55
- OCI on page 61

Enable Visibility

After onboarding, to enable traffic visibility for your cloud account, click **Enable Visibility** in Multicloud Defense in Security Cloud Control.

Enabling traffic visibility provides awareness into the traffic flows within the Cloud Accounts by collecting VPC/VNet Flow Logs and DNS Query Logs. The Flow and DNS Query logs are used by Multicloud Defense to understand traffic flow, correlate with threat intelligence feeds, and provide insight into existing threats that can be protected using Multicloud Defense.

Enabling traffic visibility is a different process for every cloud account type, but typically you will need to identify account characteristics such as your cloud account's region, VPC/VNet that you want to monitor, network security groups, and a cloud storage account for logs.

View Traffic on your Cloud Account

After enabling traffic visibility, view the traffic moving through your cloud account and look for any malicious traffic that requires protection.

1. Log in to Security Cloud Control.
2. In the left pane, click **Multicloud Defense**.
3. In the upper-right corner, click **Multicloud Defense Controller** to open the controller in a new browser tab.
4. In the Multicloud Defense portal navigate to **Investigate** > **Traffic Analysis** > **Topology**.
5. Use the **Filters and Search** bar to find the cloud account you want to monitor.
6. In the **Global View**, add **Malicious Traffic** to your filtering.
7. Click through the malicious traffic bubble to view the information in the **Region View** about the affected country of origin, IP address, FQDN, Service, and Port.

Secure Your Cloud Account

Based on your known security needs, and after monitoring traffic, you can secure your account using either a centralized hub and spoke model or a distributed model.

Use the **Secure Your Account** wizard to configure a service VPC and Multicloud Defense Gateway to secure your account.

1. Log in to Security Cloud Control.
2. In the left pane, click **Multicloud Defense**.
3. In the upper-right corner, click **Multicloud Defense Controller** to open the controller in a new browser tab.
4. In the Multicloud Defense Controller dashboard, click **Setup** in the navigation panel.

5. Click **Secure Account** on the **Setup** page.
6. Click **Centralized** or **Distributed**.
7. Click **Next** and continue with the wizard setup.

Relaunching Easy Setup

After you complete the Easy Setup workflow on the Security Cloud Control dashboard, you cannot relaunch it to start your 90-day free trial. However, the Easy Setup wizard is available in Multicloud Defense Controller to help you connect and configure other cloud accounts at a later time.

1. In the left pane of the Security Cloud Control dashboard, click **Multicloud Defense**.
2. In the upper-right corner, click **Multicloud Defense Controller**.
3. On the Multicloud Defense dashboard, click **Setup** in the Multicloud Defense menu bar.

Cloud Explorer

Learn how to use Cloud Explorer to view and understand cloud resources in Multicloud Defense.

Cloud Explorer, also known as Explorer, is designed to provide users with a graphical visualization of their cloud assets and their relationships. It offers a comprehensive view of the network topology, including components such as service VPCs, VPC/VNets, subnets, and instances. This visualization helps users manage the structure and interconnections of resources within a single interface. Explorer is beneficial for users who want to visualize their cloud infrastructure, manage these assets and their interconnections, and protect them. To manage large amounts of cloud infrastructure and connections, you can create them manually or use Terraform.

The key features of Cloud Explorer are:

- Graphical representation: Displays a clear, visual topology of cloud resources, making it easier to comprehend the relationships between assets.
- Interactive asset management: Enables users to interact with the graphical interface to manage and secure their resources effectively.
- Integrated security posture: Enables users to build and enhance their security posture by protecting assets directly within the Cloud Explorer interface.

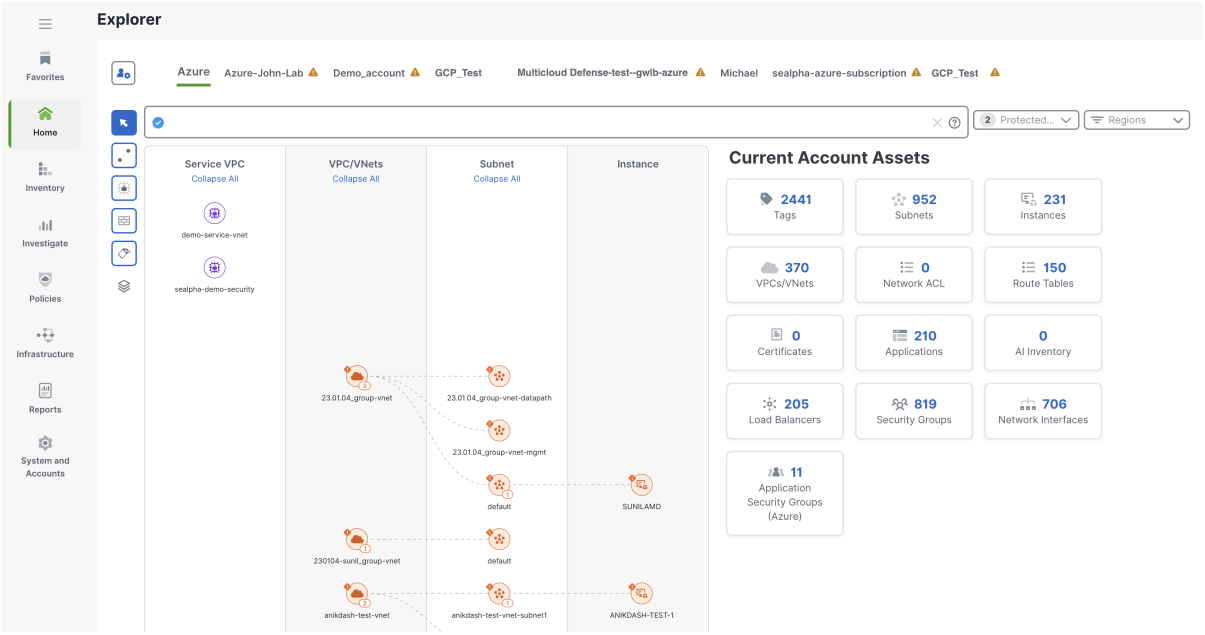
The benefits of Cloud Explorer are:

- Enhanced visibility: Provides a unified view of all cloud assets and their interconnections, simplifying resource management.
- Improved security: Allows users to secure their assets by creating service VPCs and gateways, and by protecting resources through easy drag-and-drop actions and drawing lines to connect assets.
- Ease of use: Simplifies complex configurations with an intuitive interface, enabling users to take immediate actions on their assets.
- Scenario-based utility: Particularly useful in scenarios where users need to understand resource relationships or enhance their security posture dynamically.

To access Cloud Explorer, navigate to **Home** > **Explorer**.

To create an account or select accounts to add to the viewer, click the icon ( ).

Figure 2: Cloud Explorer



Use these icons in the left pane to perform these actions:

- Move ( ): Move your assets between the swimlanes by drag-and-drop actions.
- Secure ( ): Secure your assets by forming connections between them and configuring them on the related screens. For example, click and drag on the service VPC icon to form a line and connect it to the VPC in the next swimlane. Configure and protect the VPC in the relevant screen that opens up.
- Add SVPC ( ): Click the icon and drop a service VPC into the **Service VPC** swimlane. This opens a drawer for **Create Service VPC** where you can provide details.
- Add Gateway ( ): Click the icon and drop a gateway into the swimlane. This opens a drawer for **Create Gateway** where you can provide details.
- Add Tags ( ): Add tags to selected objects. Click to select up to 3 tags to view in the display.

The central area or canvas contains swimlanes for Service VPCs, VPC/VNets, subnets, and instances. The assets that you include are displayed in the swimlanes. Click the asset and perform actions such as adding a gateway, protecting the asset, or creating a policy for the asset. You can also click an asset to expand or collapse the view of all the asset's connections.

2 Setup with the Multicloud Defense Wizard

Topics:

- Connect Cloud Account
- Enable Traffic Visibility
- Secure Your Account

Learn how to use the Multicloud Defense Wizard to set up cloud security for your environment.

The Multicloud Defense Controller provides a SaaS-delivered centralized control plane to deploy and manage Multicloud Defense and its security policy.

The **Setup** guides you through the process of setting up Multicloud Defense security using a series of simple steps:

- **Connect your Account**–This process onboards your cloud service provider account to Multicloud Defense, discovering regions and additional inventory and assets affiliated with your account.
- **Enable Traffic Visibility**–The easy setup method enables log collection to understand traffic flow.
- **Secure Your Account**–This procedure sets up a VNET or VPC, depending on the cloud account you have, and sets up Multicloud Defense Gateway to secure your experience.

For detailed instructions on getting started with Multicloud Defense, refer to the **Cisco Multicloud Defense Getting Started Guide**.

Connect Cloud Account

Learn how to connect AWS, Azure, Google Cloud Platform, or Oracle Cloud Infrastructure accounts to Cisco Multicloud Defense with the Setup wizard.

Onboard one or more cloud accounts into Multicloud Defense as the first step. This enables the Multicloud Defense Controller to interact with each account by discovering inventory, enabling traffic and logs, orchestrating security deployment, and managing policy.

Use these procedures to connect your cloud service provider account to the Multicloud Defense Controller.

Connect an AWS account

Connect an Amazon Web Services account to Cisco Multicloud Defense with the Setup wizard and an AWS CloudFormation template.

This task establishes connection between your AWS account and Cisco Multicloud Defense using Multicloud Defense easy setup wizard and CloudFormation template.

Before you begin

- An active Amazon Web Services (AWS) account is required.
- An Admin or Super Admin user role is required in your Security Cloud Control tenant.
- Multicloud Defense must be enabled for your Security Cloud Control tenant.

> **Note**
>
> Multicloud Defense Controller version 23.10 defaults to IMDSv2 in the AWS EC2 instance when using Multicloud Defense Gateway version 23.04 or newer. For more information about the difference between IMDSv1 and IMDSv2, refer to the AWS documentation.

Procedure

1. From the Multicloud Defense Controller dashboard, click **Setup** located to the left of the window.
2. Select **Connect Account**.
3. Select the AWS icon.
4. Enter the information in the modal:
   a) Click **Launch Stack** to download and deploy our CloudFormation template. This should open up another tab to deploy the template.
   b) Log in to AWS. Copy and paste the controller IAM role ARN from the CloudFormation stack output in the CloudFormation template.
   c) In the Multicloud Defense Controller easy setup modal, enter the information for the AWS account that you want to connect to. For more detailied instructions, refer to step 8 in Connect AWS Account to Multicloud Defense Controller from the Multicloud Defense Dashboard on page 38.
5. Click **Next** to onboard your account to the Multicloud Defense Controller.

Your AWS account is connected, triggering asset and inventory discovery.

What to do next

Before continuing, manually accept the AWS Marketplace Terms of Service in the AWS dashboard. Without the acceptance, the Multicloud Defense Controller cannot wholly communicate with the cloud service provider.

Once you connect the account, Multicloud Defense Controller automatically starts to discover assets and inventory associated with the cloud service provider account. Note that this is different from discovering traffic. Because Multicloud Defense Controller discovers account assets and inventory by default, the next step in this wizard is to Enable traffic visibility.

Connect Azure account

Connect an Azure subscription to Cisco Multicloud Defense with the Setup wizard and an Azure Cloud Shell onboarding script.

Use this procedure to connect to an Azure subscription through Multicloud Defense Controller's easy setup wizard:

Before you begin

- You must have an active Azure subscription.
- You must have an Admin or Super Admin role in your Security Cloud Control tenant.
- You must enable Multicloud Defense for your Security Cloud Control tenant.

Procedure

1. In the Security Cloud Control dashboard, click the **Multicloud Defense** tab located in the left navigation pane.
2. Click **Multicloud Defense Controller** located in the upper right window.
3. From the Multicloud Defense Controller dashboard, click **Setup** located to the left of the window.
4. Select **Connect Account**.
5. Select the Azure icon.
6. Enter the required information:
   a) Click the link to open an Azure Cloud Shell in bash mode.
   b) In the Azure account modal, click **Copy** to copy the onboarding script and execute it in the bash shell that was opened in step 1.
   c) In the Azure account modal, provide a name for this Azure account. You can choose to name this the same as your Azure subscription name. This name is visible on the Multicloud Defense Controller accounts page only.
   d) (Optional) Provide a description for the subscription. For detailed instructions, refer to step 8 in Connect an Azure Subscription to the Multicloud Defense Controller from the Multicloud Defense Dashboard on page 48.
7. Click **Next**.

Your Azure subscription connects with Cisco Multicloud Defense, providing integrated security and management capabilities for your network.

What to do next

Once you connect the account, Multicloud Defense Controller automatically starts to discover assets and inventory associated with the cloud service provider account. Note that this is different from discovering traffic. Because Multicloud Defense Controller discovers account assets and inventory by default, the next step in this wizard is to Enable traffic visibility.

Connect Google Cloud Platform account

Connect a Google Cloud Platform project to Cisco Multicloud Defense with the Setup wizard and Google Cloud Shell.

Use this procedure to onboard a singular GCP project as an account using the Multicloud Defense Controller's easy setup wizard:

Before you begin

- You must have an active Google Cloud Platform (GCP) project.
- Ensure you have the permissions required to create VPCs, subnets, and a service account within your GCP project. For more information, refer to the GCP documentation.
- You must have an Admin or Super Admin role in your Security Cloud Control tenant.
- You must enable Multicloud Defense for your Security Cloud Control tenant.

Procedure

1. In the Multicloud Defense Controller dashboard, click **Setup**.
2. Select **Connect Account**.
3. Select the GCP icon.
4. Click **Cloud Platform Cloud Shell** to launch the Cloud Shell. You can also log into your GCP account and launch the Cloud Shell from the project to connect to Multicloud Defense; note that the script automatically modifies the project name to the name of the project you launch the cloud shell from.
   a) Copy the command generated in the Multicloud Defense Controller easy setup modal and paste the command into the Cloud Shell. Execute it to initiate the onboarding process. This script automatically creates user accounts enabling Multicloud Defense Controller to communicate directly with your GCP project.
   b) If you have multiple GCP projects, you are prompted to select the project through a numbered list. Select the value for the project you want to connect and submit.
   c) When prompted with `Continue configuring this project? [y/n]` Type either "y" or "n". Then, press **enter**.

   If the GCP project you are connecting to Multicloud Defense has been previously onboarded, you may get an error about the GCP's cloud storage bucket already existing. If the cloud storage bucket does not meet your requirements, create a new storage bucket in your GCP account to handle the flow logs on this project after it is connected to Multicloud Defense.
5. Enter the required information for the GCP account that you want to connect to. For detailed instructions, refer to step 12 in Connect a GCP Project to the Multicloud Defense Controller from the Multicloud Defense Dashboard on page 57.
6. Click **Next**.

Successful connection of your GCP project to Multicloud.

What to do next

GCP does not automatically include the regions your project is configured for. After your project is connected to Multicloud Defense we strongly recommend navigating to **Inventory** > **Inventory** to manually modify and add any and all appropriate regions.

Once you connect the account, Multicloud Defense Controller automatically starts to discover assets and inventory associated with the cloud service provider account. Note that this is different from discovering traffic. Because Multicloud

Defense Controller discovers account assets and inventory by default, the next step in this wizard is to Enable traffic visibility.

Connect to an OCI Account

Learn how to prepare your Oracle Cloud Infrastructure account and required access before connecting it to Cisco Multicloud Defense.

Review the procedures to prepare your OCI account before you connecting it to Multicloud Defense.

Prepare your OCI account

Prepare an Oracle Cloud Infrastructure account for Cisco Multicloud Defense by creating the required group, policy, user, and API credentials.

This procedure automates the connection between Multicloud Defense and your OCI account and then create a policy with the correct permissions. Without the necessary permissions provided in this procedure, certain features are unavailable.

Execute this procedure to connect to an Oracle Cloud (OCI) account using Multicloud Defense's setup wizard:

Procedure

1. Log in to your OCI tenant.
2. Navigate to **Identity & Security** > **Groups**.
3. Click **Create Group**.
4. Enter the following:
   - **Name:** Multicloud Defense-controller-group
   - **Description:** Multicloud Defense Group
5. Click **Create**.
6. Create a Network Firewall Policy in OCI. For more information, refer to the OCI documentation. Include this information when creating the policy:
   - **Name:** Multicloud Defense-controller-policy.
   - **Description:** Multicloud Defense Policy.
   - **Compartment:** [Must be the "root" Compartment].

   a) Add the following permissions under the **Show Manual Editor** tab:

```
Allow group <group_name> to inspect instance-images in compartment
<compartment_name>
Allow group <group_name> to read app-catalog-listing in compartment
<compartment_name>
Allow group <group_name> to use volume-family in compartment
<compartment_name>
Allow group <group_name> to use virtual-network-family in compartment
<compartment_name>
Allow group <group_name> to manage volume-attachments in compartment
<compartment_name>
Allow group <group_name> to manage instances in compartment
<compartment_name>
Allow group <group_name> to {INSTANCE_IMAGE_READ} in compartment
<compartment_name>
Allow group <group_name> to manage load-balancers in compartment
```

```
<compartment_name>
Allow group <group_name> to read marketplace-listings in tenancy
Allow group <group_name> to read marketplace-community-listings in tenancy
Allow group <group_name> to inspect compartments in tenancy
Allow group <group_name> to manage app-catalog-listing in compartment
<compartment_name>
Allow group <group_name> to read virtual-network-family in tenancy
Allow group <group_name> to read instance-family in tenancy
Allow group <group_name> to read load-balancers in tenancy
```

- **group_name:** Multicloud Defense-controller-group.
- **compartment_name:**[Compartment where Multicloud Defense will be deployed].

> **Note**
>
> When replacing the **<compartment_name>** with the name of the compartment where the policy will apply, if the compartment is a sub-compartment, the name format is **compartment:sub-compartment** (for example, Prod:App1).
>
> If the **<compartment_name>** is specified as the root compartment (e.g., multicloud (root)), OCI will not accept the policy and will produce an error: *Invalid parameter*. The policy will need to be defined for a specific compartment and that compartment cannot be the root compartment.

b) Click **Create**.

7. Create a user in OCI. For more information, refer to the OCI documentation. Provide the configuration information when creating a user:
   - **Name:** *Multicloud Defense-controller-user*
   - **Description:** *Multicloud Defense User*
8. Create an API Key. For more information, refer to the OCI documentation.

   Download both the private key and the public key before adding the API Key.
9. Accept the **Terms and Conditions** for an OCI account. For more information, refer to the OCI documentation. Ensure to access the **Change image** section of the UI to add the following "community image" information specific to Multicloud Defense:
   a) **Check the box** for Multicloud Defense.
   b) **Check the box** for *I have reviewed and accept the Publishers terms of use, Oracle Terms of Use, and the Oracle General Privacy Policy*.
   c) Click **Exit** without deploying the image prior to connecting the account to Multicloud Defense.

   Repeat the steps for each compartment you plan to deploy a Multicloud Defense Gateway.

Your OCI account is now configured for Cisco Multicloud Defense integration, enabling enhanced functionalities and secure environment setup.

What to do next

Validate connections by testing access permissions and functionality within your configured environment.

Connect Oracle account

Connect an Oracle Cloud Infrastructure account to Cisco Multicloud Defense with the Setup wizard and Oracle Cloud Shell.

Use this procedure to connect to an OCI account using Multicloud Defense Controller's easy setup wizard:

Before you begin

- You must have an existing Oracle Cloud (OCI) account.
- Ensure your OCI account prerequisites are completed before onboarding. For more information, refer to Prepare your OCI account on page 29.
- You must have an Admin or Super Admin role in your Security Cloud Control tenant.
- You must enable Multicloud Defense for your Security Cloud Control tenant.

Procedure

1. In the Multicloud Defense Controller dashboard, click **Setup**.
2. Select **Connect Account**.
3. Select the OCI icon.
4. Click **Oracle Cloud Shell** to launch the native shell prompt.
5. Run the command from the Multicloud Defense Setup wizard in your cloud shell.

   This command automates the process of creating an IAM policy, OCI group, and an OCI user that facilitate the communication between your OCI account and the Multicloud Defense.
6. Enter the required information for the OCI account that you want to connect to. For detailed instructions, refer to step 6 in Connect the Oracle OCI Tenant to the Multicloud Defense Controller from the Multicloud Defense Dashboard on page 63.
7. Click **Next**.

The system initiates asset discovery and inventory associated with your cloud service provider account once connected.

What to do next

Once you connect the account, Multicloud Defense Controller automatically starts to discover assets and inventory associated with the cloud service provider account.

Enable Traffic Visibility

Learn how traffic visibility helps you understand cloud traffic and threats, and find the setup procedure for your AWS, Azure, or Google Cloud Platform account.

To understand traffic flows within your cloud account, enabling traffic visibility collects these types of logs:

- VPC or VNet flow logs
- DNS logs
- Route53 query logging

Flow and DNS query logs help Multicloud Defense to understand traffic flow, correlate data with threat intelligence feeds, and gain insights into existing threats protected by Multicloud Defense.

Enabling traffic visibility varies by cloud account type. Identify the characteristics of your cloud account, including the region, VPC/VNet for monitoring, network security groups, and a cloud storage account for logs.

> **Note**
>
> Multicloud Defense does not support traffic visibility for OCI at this time. Enable asset discovery as an alternative action for this procedure: which means Multicloud Defense identifies and collects metadata for assets from an external environment and the resulting data collected creates an inventory that can be used to assist migration. For more information, refer to Enable Asset Discovery and Inventory on page 76.

Enable traffic for an AWS account

Configure traffic visibility for an AWS account in Cisco Multicloud Defense to collect flow and DNS query logs from selected VPCs.

Use this procedure to enable traffic visibility for an AWS account using the Setup wizard:

Procedure

1. In the Multicloud Defense Controller portal click **Setup** in the left navigation bar.
2. In the setup wizard, click **Enable Traffic Visibility**.
3. Enter the required information into the modal:
   a) **CSP Account** - Use the drop-down menu to select the cloud service provider account to which Multicloud Defense Controller deploys the service VPC/VNet.
   b) **Region** - Use the drop-down menu to select the region where the cloud service provider you selected is located.
   c) **VPCs** - Scroll through the table of available VPCs that are applicable to the type of cloud service provider you selected and check the appropriate VPC. If the VPC does not appear, click the **Refresh** icon to update the list.
   d) **S3 Bucket** - Use the drop-down menu to select an existing S3 bucket from your account; stores DNS queries and VPC/VNet flow logs. This S3 bucket is already created in your account.
4. Click **Next**.

Traffic visibility is configured, enabling log collection for selected VPCs in your AWS account.

What to do next

Confirm that your account is secure and review the log data appropriately.

Enable traffic for an Azure account

Configure traffic visibility for an Azure account in Cisco Multicloud Defense to collect flow and DNS logs from selected virtual networks.

Use this procedure to enable traffic visibility for an Azure account from the Setup wizard:

Procedure

1. In the Multicloud Defense Controller portal click **Setup** in the left navigation bar.
2. In the setup wizard, click **Enable Traffic Visibility**.
3. Enter the required information into the modal:

a) **CSP Account** - Use the drop-down menu to select the cloud service provider account to which Multicloud Defense Controller deploys the service VPC/VNet.
b) **Region** - Use the drop-down menu to select the region where the cloud service provider you selected is located.
c) **Copy and run the script**. If you are re-onboarding an Azure account and are reusing a cloud storage bucket, the script does not automatically create a new storage bucket. Use the default, or preexisting storage bucket. Otherwise, create a new storage bucket in the Azure dashboard or manually edit the script command before executing it to include the storage bucket name for the flow logs.
d) **Virtual Network (VNet)** - Select at least one VNet for traffic to be visible on. Scroll through the table of available VNets applicable to the type of cloud service provider you selected, and check the appropriate one. Note that if you do not immediately see the VNet, click the **Refresh** icon to refresh the current list.

> **Note**
>
> You may see existing NSG Flow Logs, which will be supported by Multicloud Defense until it is deprecated by Microsoft Azure. You will no longer be able to create new NSG Flow Logs. Instead, you can create VNet Flow Logs.

e) **Storage Account** - Enter the full Resource ID in the selected region.

4. Click **Next**.

Traffic visibility for your Azure account is enabled, allowing DNS and flow logs collection.

What to do next

Ensure your account maintains security standards and review logs for new insights.

This task ensures traffic visibility for monitoring and enhances security by assimilating flow and DNS logs into Cisco Multicloud Defense for comprehensive analysis.

Enable traffic for a GCP project

Configure traffic visibility for a Google Cloud Platform project in Cisco Multicloud Defense to collect logs from selected VPCs.

Use this procedure to enable traffic visibility for a GCP account with the Setup wizard:

Procedure

---

1. From the Security Cloud Control Home page, click **Multicloud Defense** .
2. In the Multicloud Defense Controller portal click **Setup** in the left navigation bar.
3. In the setup wizard, click **Enable Traffic Visibility**.
4. Enter the required information into the modal:
   a) **CSP Account** - Use the drop-down menu to select the cloud service provider account to which Multicloud Defense Controller deploys the service VPC/VNet.
   b) **Cloud Storage** - Select an available cloud storage bucket that has already been assigned to the GCP project you selected.
   c) **Select VPC(s)** - Select at least one VPC for traffic to be visible on. Scroll through the table of available VPCs that apply to the type of cloud service provider you selected, and check the appropriate VPC. Note that if you do not immediately see the VPC, click the **Refresh** icon to refresh the current list.

d) **Copy and run the script**. Note that if you are re-onboarding a GCP project and reusing a cloud storage bucket, the script does not automatically create a new storage bucket. You can use the default or preexisting storage bucket. Otherwise, create a new storage bucket in the GCP dashboard or manually edit this script command before executing to include the storage bucket name for storing your GCP project's flow logs.

5. Click **Next**.

Your selected VPCs are configured for traffic visibility logs in your GCP project.

What to do next

Verify the security of your account settings.

Secure Your Account

Learn how to secure a cloud account with a Cisco Multicloud Defense gateway deployed in a centralized or distributed model.

Secure your account with a gateway deployed in either a centralized or a distributed model.

In a **Centralized** model, Multicloud Defense orchestrates the VPC or VNet and additional components, and deploys the gateway within this construct.

In a **Distributed** model, Multicloud Defense builds and deploys a gateway within the existing infrastructure that your network already has available.

Follow one of these procedures to secure your account.

Centralized Model: Add a VPC or VNet

Create a service VPC or VNet and deploy a Cisco Multicloud Defense gateway with the centralized deployment model.

Follow these steps to create and add a VPC or VNet for your gateway and secure your account:

Before you begin

You must have at least one cloud service provider connected to the Multicloud Defense Controller before you begin this wizard. This procedure may vary for different providers based on their required parameters.

Procedure

---

1. In the left pane of the Multicloud Defense Controller portal, choose **Home > Easy Setup**.
2. In the setup wizard, on the **Secure Your Account** section, click **Secure Account**.
3. Select **Centralized** so it is highlighted.
4. Click **Next**.
5. Add a service VPC/VNet:
   a) **Name** - Enter a name for the service VPC/VNet. Once created, this name is displayed in the **Infrastructure** > **Gateways** > **VPCs/VNets** page.
   b) (AWS only)**CSP Account** - Use the drop-down menu to select a cloud service provider account that is already connected to the Multicloud Defense Controller. Deploy the service VPC/VNet to the selected account.
   c) **Region** - Use the drop-down menu to select the region where the selected cloud service provider is located.
   d) **CIDR Block** - Enter the unique value for the Transit Gateway to which the service VPC/VNet is attaching.
   e) (GCP only) **Datapath CIDR Block** - Enter a valid CIDR block for datapath VPC, ensuring it does not overlap with spoke VPCs.

f) (GCP only) **Management CIDR Block** - Enter a valid CIDR block for the management VPC.
g) **Availability Zones** - Of the generated list, select at least one availability zone. Using two zones is **strongly** recommend for best results.
h) (Azure only) **Resource Group** - Use the drop-down menu to select a resource group to associate the gateway to. If no resource groups are listed, you can **Create Resource Group** from this screen.
i) (AWS only) **Transit Gateway** - Use the drop-down menu to select an available transit gateway for the VPC to assocaite with. If you do not have one available, click **create_new** to create a transit gateway from this window.
j) (AWS and Azure only) **Use NAT Gateway** - Check this option to direct all egress traffic through the NAT gateway. Multicloud Defense automatically creates a NAT gateway for each availability zone that is selected.
k) (Azure only) If you want to attach a Virtual WAN (VWAN), in **vWAN Attachment** set the toggle to **Enabled** .
l) (Azure only) From the **vHub** drop-down list, choose a hub.
m) (Azure only) In the **Associate Route Table** drop-down list, select a route table to associate.
n) (Azure only) In the **Propagate Route Tables** drop-down list, select route tables to propagate.

6. Click **Next**.

What to do next

Add a Gateway.

Distributed Model

Learn about the distributed gateway model and find the procedure for creating a Cisco Multicloud Defense gateway in your AWS, Azure, GCP, or OCI environment.

Use the procedures specific to your cloud service provider for a distributed gateway model.

Azure Distributed Model: Create a Gateway

Create a Cisco Multicloud Defense gateway in an existing Azure VNet with the distributed deployment model.

To create a gateway for your Azure account using the distributed model, follow these steps:

Procedure

---

1. In the Multicloud Defense Controller portal click **Setup** in the left navigation bar.
2. In the setup wizard, click **Secure Account**.
3. Select **Distributed** so it is highlighted.
4. Click **Next**.
5. Enter the gateway Information:
   a) **Account** - Use the drop-down menu to select an Azure account you want to deploy the gateway to.
   b) **Name** - Enter a name for the gateway. This name is displayed in the **Infrastructure** > **Gateways** > **Gateways** page.
   c) (Optional) **Description** - Enter a description for the gateway that might help identify it from other gateways.
   d) **Instance Type** - Use the drop-down menu to select the instance type that deploys the gateway.
   e) **Minimum Instances** - Select the minimum number of instances deployed in auto scaling group per availability zone.
   f) **Maximum Instance** - Select the maximum number of instances deployed in auto scaling group per availability zone.
   g) **HealthCheck Port** - Enter the healthcheck port number. Multicloud Defense Controller uses `65534` as the default value.
   h) **User Name** - Enter the user name used to access the gateway once created.

i) **Packet Capture Profile** - Use the drop-down menu to select where packets are stored in the cloud storage bucket. If options are not listed, click **Create Packet Capture Profile** to create one directly from this window.
j) **Log Profile** - Use the drop-down menu to select which cloud service provider is used to forward logging to.
k) **Metrics Profile** - Use the drop-down menu to select an entity to forward metrics to. If there are no option listed, click **Create Metrics Forward Profile** to create one from this window.
l) **NTP Profile** - Use the drop-down menu to select the NTP profile associated with the gateway. If there are no options listed, click **Create** to create one from this window.
m) **Security** - Select the type of traffic flow your gateway is expected to handle. Ingress security targets traffic that flows from the public internet to a private network; east-west and egress security targets outbound traffic from your private network and traffic moving between your data centers.
n) **Gateway Image** - Use the drop-down menu to select the gateway image to be deployed to the gateway.
o) **Policy Ruleset** - Use the drop-down menu to select a policy rulset to be deployed and start processing traffic. If there is not ruleset listed, click **Create new** to create a policy rulset from this window.
p) **Region** - Use the drop-down menu to select the region your gateway is deployed to.
q) **VPC/VNet ID** - Use the drop-down menu to select the VPC where the gateway is deployed to.
r) **Key Selection** - Select either an SSH Public key or an SSH Key Pair. Enter the value that is applied to the gateway in the next text field.
s) **Resource Group** - Use the drop-down menu to select an existing resource group that is applied to the gateway.
t) **User Assigned Identity ID** - Enter a valid value.
u) **Mgmt. Security Group** - Use the drop-down menu to select a security group used for the gateway management interface. Note that if you select a Multicloud Defense-created service VPC, a security group is created specifically for management.
v) **Datapath Security Group** - Use the drop-down menu to select a security group used for the gateway datapath interface. If selecting Multicloud Defense-created service VPC, a security group is created specifically for the datapath.
w) **Disk Encryption** - Enable disk encryption with either the Azure managed encryption or a customer-managed encryption key. Note that if you opt for a customer-managed encryption key, you need to create and deploy an IAM policy for successful deployment.
x) **Availability Zone** - Use the drop-down menu to select an availabililty zone.
y) **Mgmt. Subnet** - Use the drop-down menu to select a management subnet for the management interface.
z) **Datapath Subnet** - Use the drop-down menu to select a datapath subnet for the datapth interface.

   To add more instance types, click the "+" icon. Subseuqntly, you can remove additional instance types with the "-" icon.

6. Click **Next**.
7. Enter the details for advanced settings.
8. Click **Next**.
9. Review.

What to do next

3 AWS

Topics:

- AWS Overview
- Connect AWS Account to Multicloud Defense Controller from the Multicloud Defense Dashboard

Learn how to connect and manage AWS accounts in Multicloud Defense to protect AWS cloud environments.

AWS Overview

An overview of preparing an Amazon Web Services account for Cisco Multicloud Defense onboarding with the provided AWS CloudFormation template.

Multicloud Defense has created a CloudFormation template that you use when connecting an AWS account to the Multicloud Defense Controller.

To prepare a cloud account for integration with Multicloud Defense Controller, there are certain steps that need to be performed in the cloud account. Perform these prerequisite steps before connecting your AWS cloud account to the Multicloud Defense Controller. This provides an overview of the operation but is not intended to be performed manually. In the CloudFormation section, there are details of deployments and parameters information.

Overview of steps

1. Create a cross-account IAM role that the Multicloud Defense Controller uses to manage your cloud account.
2. Create an IAM role that is assigned to the Multicloud Defense Gateway EC2 instances that run in your account.
3. Create a CloudWatch event rule that transfers the management events to the Multicloud Defense Controller.
4. Create an IAM role that is used by the CloudWatch event rule to grant it the permissions to transfer the management events.
5. Optionally, create an S3 bucket in your account to store CloudTrail events, Route53 DNS query logs, and VPC Flow Logs.
6. Enable Route53 DNS Query Logging with the destination as the previously created S3 Bucket, and select the VPCs for which query logging must be enabled.
7. Enable CloudTrail to log all the management events to the S3 Bucket created above.
8. Enable VPC Flow Logs with destination as the S3 Bucket created above.
9. Accept the AWS Marketplace Terms of Service.

Connect AWS Account to Multicloud Defense Controller from the Multicloud Defense Dashboard

Connect AWS accounts to Cisco Multi-Cloud Defense using an automated CloudFormation template and ensure cloud resource security

Multicloud Defense has created a CloudFormation template that makes it easy to connect an AWS account to the Multicloud Defense Controller.

Before you begin

Read these requirements before you connect an AWS account to Multicloud Defense:

- You must have requested a Multicloud Defense Controller for your Security Cloud Control tenant before you begin.
- The name of the cloud storage bucket in your AWS account must be between 3 and 65 characters. Bucket names exceeding 65 characters will cause an error during the connection process.

Note

Multicloud Defense Controller version 23.10 defaults to IMDSv2 in the AWS EC2 instance when using Multicloud Defense Gateway version 23.04 or newer. For more information about the difference between IMDSv1 and IMDSv2, see AWS documentation.

Procedure

1. In the left pane, click Multicloud Defense.
2. Click Multicloud Defense Controller.
3. In the Cloud Accounts pane, click **Add Account**.
4. On the **General Information** page, choose AWS from the **Account Type** list box.
5. Click **Launch Stack** to download and deploy our CloudFormation template. This should open up another tab to deploy the template. Logging in to AWS is required.
6. Acknowledge that the AWS CloudFormation might create IAM resources with custom names.
7. Fill in these values:
   - **AWS Account Number**: Enter the AWS account number of the account you wish to secure. This number can be found in the output value CurrentAccount of the CloudFormation Template.
   - **Account Name**: Enter the name you want to give your account once it has been onboarded.
   - **Description:**(Optional) Enter an account description.
   - **External ID**: A random string for trust policy of the IAM role. This value will be used in the controller IAM role created. You can edit or regenerate the External ID.
   - **Controller IAM Role**: This is the IAM role created for theMulticloud Defense Controller during CloudFormation Template (CFT) deployment. Look for the output value MCDControllerRoleArn in CFT stack. It should be something similar to this: `arn:aws:iam::<Acc Number>:role/ciscomcdcontrollerrole`.
   - **Inventory Monitor Role**: This is the IAM role created for Multicloud Defense Inventory during CFT deployment. Look for the output value MCDInventoryRoleArn in CFT stack. Should be something similar to this: `arn:aws:iam::<Acc Number>:role/ciscomcdinventoryrole`.
8. Click **Save and Continue**.
   You return to the Multicloud Defense dashboard, where you will see that the you have a new AWS cloud account recorded.

What to do next

1. Add the correct permissions to your AWS account to ensure secure access and communication.
2. Enable traffic visibility.

CloudFormation Outputs

Learn which AWS CloudFormation stack outputs to save after onboarding an account for Cisco Multicloud Defense.

From the **Outputs** tab, copy and paste the following informationin to a text editor:

- CurrentAccount (This is your AWS Account ID where applications run and Multicloud Defense Gateways will be deployed)
  - MCDControllerRoleArn
  - MCDFirewallRoleName
  - MCDInventoryRoleArn
  - MCDS3BucketArn
  - MCDS3BucketName

Roles Created by Multicloud Defense

Describes roles created when onboarding a cloud account to multicloud.

When you onboard a cloud service account to Multicloud Defense Controller with the provided script, user roles are created within the parameters of the cloud service provider to ensure that communication between the services are protected. Depending on the cloud service provider, different roles and permissions are created.

The following roles are created when you onboard an account.

MCDControllerRole

Find details about the cross-account MCDControllerRole permissions required for Cisco Multicloud Defense to manage resources in your AWS account.

Cross-account IAM role that allows Multicloud Defense to access your cloud account and take necessary actions, for example, Create EC2 instances, create load balancers, and change Route53 entries. The service principal is the Multicloud Defense-controller-account with an external id applied. Here is the IAM policy applied to the role (e.g controller role name used in this example is ***Multicloud Defense-controller-role***):

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Action": [
        "acm:DescribeCertificate",
        "aacm:ListCertificates",
        "apigateway:Get",
        "bedrock:Get*",
        "bedrock:List*",
        "bedrock:Retrieve",
        "ec2:*",
        "elasticloadbalancing:*",
        "events:*",
        "globalaccelerator:*"
        "iam:ListPolicies",
        "iam:ListRoles",
        "iam:ListRoleTags",
        "logs:*",
        "route53resolver:*",
        "servicequotas:GetServiceQuota",
        "s3:ListAllMyBuckets",
        "s3:ListBucket",
        "wafv2:Get*",
        "wafv2:List*,
      ],
      "Effect": "Allow",
```

```
      "Resource": "*"
    },
    {
      "Action": [
        "iam:GetRole",
        "iam:ListRolePolicies",
        "iam:GetRolePolicy"
      ],
      "Effect": "Allow",
      "Resource": [
        "arn:aws:iam::<ciscomcd-account>:role/ciscomcd-controller-role"
      ]
    },
    {
      "Effect": "Allow",
      "Action": "s3:GetObject",
      "Resource": "arn:aws:s3:::<S3Bucket>/*"
    },
    {
      "Action": [
        "iam:GetRole",
        "iam:ListRolePolicies",
        "iam:GetRolePolicy",
        "iam:PassRole"
      ],
      "Effect": "Allow",
     "Resource": "arn:aws:iam::<customer- account>:role/ciscomcd_firewall_role"

    },
    {
      "Action": "iam:CreateServiceLinkedRole",
      "Effect": "Allow",
      "Resource": "arn:aws:iam::*:role/aws-service-role/*"
    }
  ]
}
```

Service Principal:

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Principal": {
        "AWS": [
          "arn:aws:iam::<ciscomcd-account>:root"
        ]
      },
      "Action": "sts:AssumeRole",
      "Condition": {
        "StringEquals": {
          "sts:ExternalId": "ciscomcd-external-id"
        }
      }
    }
  ]
}
```

MCDGatewayRole

Find details about the MCDGatewayRole permissions required for Cisco Multicloud Defense gateway instances to access supporting AWS services.

Role that is assigned to the Multicloud Defense Gateway (Firewall) EC2 instances. The role gives the gateway instance capabilities to access secretsmanager where the private keys for the application are stored, ability to decrypt keys using AWS KMS if the keys are stored in KMS, and save objects like PCAPs and technical support data onto a S3 bucket. The service principal of this role is ec2.amazonaws.com. Here is the IAM policy applied to the role:

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Action": [
        "s3:PutObject",
        "s3:ListBucket"
      ],
      "Effect": "Allow",
      "Resource": "arn:aws:s3:::*/*"
    },
    {
      "Action": [
        "kms:Decrypt"
      ],
      "Effect": "Allow",
      "Resource": "*"
    },
    {
      "Action": [
        "secretsmanager:GetSecretValue"
      ],
      "Effect": "Allow",
      "Resource": "*"
    }
  ]
}
```

Tip

You can download and edit the CloudFormation template to make the policy more restrictive e.g. restricting decrypt to use a specific key, or PutObject to a defined/specific S3 bucket.

MCDInventoryRole

Find details about the MCDCloudWatchEventRole required for Cisco Multicloud Defense to receive AWS CloudTrail inventory events from your cloud account.

This is the role used for dynamic inventory purposes and provides the capability for the CloudTrail events to be transferred to the Controller's AWS account. It does the following:

- Put events on the event bus in the AWS account where the Multicloud Defense Controller exists.
- Send events matching the rule to the Multicloud Defense Controller's webhook server directly from the customer's AWS account.

The Service Principal for this role is **events.amazonaws.com**. Here is the policy applied to the role:

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Action": "events:PutEvents",
      "Effect": "Allow",
      "Resource": [
        "arn:aws:events:*:<ciscomcd-account>:event-bus/default"
      ]
    }
  ]
}
```

InventoryMonitorRule

Find details about the MCDCloudWatchEventRule that identifies relevant AWS CloudTrail inventory events for Cisco Multicloud Defense.

Rule that is added to the MCDInventoryRole to put all CloudTrail inventory changes to EC2 and API gateways to be copied to the event bus on the AWS account where the Multicloud Defense Controller runs. The rule is required to match on specific event patterns that occur in the customer's AWS account. Once a match occurs, the rule states that the matched event should be sent to the webhook server (API based destination) of the controller. This rule is executed using the Multicloud DefenseMCDInventoryRole created in the previous section.

Custom Event Pattern:

```
{
  "detail-type": [
    "AWS API Call via CloudTrail",
    "EC2 Instance State-change Notification"
  ],
  "source": [
    "aws.ec2",
    "aws.elasticloadbalancing",
    "aws.apigateway"
  ]
}
```

Target:

```
Event Bus in another AWS Account (mcd-account) using the MCDInventoryRole
```

4 Azure

Topics:

- Prepare Your Azure Account
- Connect an Azure Subscription to the Multicloud Defense Controller from the Multicloud Defense Dashboard
- Post-Onboarding Procedures

Explains the range of services provided by Azure and key attributes that differentiate them.

Prepare Your Azure Account

Learn how to prepare an Azure account and subscription for Cisco Multicloud Defense onboarding, including application registration, access, and event notifications.

Prepare your Azure account and subscription(s) before you connect and onboard them to Multicloud Defense Controller with the following steps:

1. Acquire and register an Azure subscription. Ensure the subscription is associated to the Microsoft Entra ID. Review the list of **App Registrations** in your Azure portal to confirm whether the subscription is correctly linked to Multicloud Defense.
2. Create a custom role for your Azure subscription. This grants Multicloud Defense access to specific resources or actions that would otherwise be blocked.
3. Subscribe to the Azure Event Grid. This allows Multicloud Defense to receive real-time updates and can be configured to send events to subscribers (push) or subscribers can connect to Event Grid to read events (pull). See the "*Create event subscriptions*" chapter in the Azure User Guide for more information.
4. Connect an Azure Subscription to the Multicloud Defense Controller from the Multicloud Defense Dashboard on page 48. Azure subscriptions encapslulate "technical" resources such as virtual machines. Complete this step to use any Azure-based VMs with Multicloud Defense gateways or any depoyment action.
5. Accept Marketplace Terms. If this is the first time your Azure account is being onboarded to Multicloud Defense, you must accept Cisco marketplace terms. Without this agreement you cannot complete the onboarding action.
6. (Optional) User-assigned Managed Identity for Key Vault and Blob Storage access. Configured in the Azure environment, the key vault and blob storage access is intended to give you more flexibility to use the same identity across different resources, maintaining consistent permissions and identities across services.

If you find that you cannot use the automated script, refer to the alternative procedure to manually onboard your account here.

> **Note**
>
> If you have more than one subscription you want to configure with Multicloud Defense, use the procedure in Connect an Azure Subscription to the Multicloud Defense Controller from the Multicloud Defense Dashboard on page 48 for one subscription and then modify the policy in your Azure portal to add the other subscriptions. You must onboard these subscriptions individually, but you can associate them with Multicloud Defense in bulk.

Register Application in Microsoft Entra ID

Register the Cisco Multicloud Defense application in Microsoft Entra ID and collect the identifiers and credentials required to connect your Azure subscription.

Use the following procedure to register the Multicloud Defense application in your Entra ID.

Procedure

1. From your Azure portal, navigate to **Microsoft Entra ID**.
2. Select **App registrations**.
3. Click **New registration**.

4. Provide a name to reference the new app registration e.g. Multicloud Defense Controller In the *Supported account types* choose the second option *Accounts in any organizational directory*.
5. Choose the option appropriate to your organization. Note that the **Redirect URI** is not needed for the creation of the App registration.
6. Click **Register**.
7. Navigate to **App registrations** section and click on the newly created application.
8. Navigate to **Manage** section and click **Certificates & secrets**.
9. Click **+ New client secret**, and then enter the required information in the *Add a client secret* dialog box
   - **Description** - Add a description (for example, Multicloud Defense-controller-secret1)
   - **Expires** - Choose **24 months**. You can also make this selection at your convenience. You will need to create new secrets when the current one expires.
10. Click **Add**. The client secret is populated under the **Value** column.
11. Copy the **Client secret** into a notepad, as this is shown only once and is never displayed again.
12. In the left navigation bar click **Overview**.
13. Copy the **Application (client) ID** and **Directory (tenant) ID** into a notepad.

Create a Custom Role to Assign to the Application

Create and assign an Azure custom role that grants the Cisco Multicloud Defense application the required access to your subscription.

The CloudFormation template creates a **custom role** that will be assigned to the application created for the Multicloud Defense Controller. The custom role gives the application permissions to read inventory information and create resources, for example, VMs, load balancers, etc.

There are multiple ways to create a custom role but we recommend the following procedure:

Procedure

1. Navigate to **Subscription** and click **Access Control (IAM)**.
2. Click on **Roles** and on the top menu bar navigate to click **+Add** > **Add Custom Role**.
3. Give a name to the custom role, for example, `Multicloud Defense-controller-role`.
4. Keep clicking **Next** until you get to the JSON editing screen.
5. Click **Edit** on the screen and in the JSON text, under the **permissions > actions** section, copy and paste the provided code content in between the square brackets (adjust the indentation if needed). Note that you are also adding permissions in the Azure role to read role assignments and role definitions, as indicated in the code content.

```
"Microsoft.ApiManagement/service/*",
"Microsoft.Compute/disks/*",
"Microsoft.Compute/images/read",
"Microsoft.Compute/sshPublicKeys/read",
"Microsoft.Compute/virtualMachines/*",
"Microsoft.ManagedIdentity/userAssignedIdentities/read",
"Microsoft.ManagedIdentity/userAssignedIdentities/assign/action",
"Microsoft.Network/loadBalancers/*",
"Microsoft.Network/natGateways/*",
"Microsoft.Network/networkinterfaces/*",
"Microsoft.Network/networkSecurityGroups/*",
"Microsoft.Network/publicIPAddresses/*",
```

```
"Microsoft.Network/routeTables/*",
"Microsoft.Network/virtualNetworks/*",
"Microsoft.Network/virtualNetworks/subnets/*",
"Microsoft.Resources/subscriptions/resourcegroups/*",
"Microsoft.Storage/storageAccounts/blobServices/*",
"Microsoft.Storage/storageAccounts/listkeys/action",
"Microsoft.Network/networkWatchers/*",
"Microsoft.Network/applicationSecurityGroups/*",
"Microsoft.Compute/diskEncryptionSets/read",
"Microsoft.Insights/Metrics/Read"
"Microsoft.Network/locations/serviceTagDetails/read",
"Microsoft.Network/locations/serviceTags/read",
"Microsoft.CognitiveServices/*/read",
"Microsoft.CognitiveServices/accounts/listkeys/action",
"Microsoft.Network/virtualHubs/*",
"Microsoft.Network/virtualHubs/hubRouteTables/*",
"Microsoft.MarketplaceOrdering/offertypes/publishers/offers/plans/agreements/*"
"Microsoft.Authorization/roleAssignments/read",
"Microsoft.Authorization/roleDefinitions/read"
```

6. **Optional** - If you plan to use multiple subscriptions with Multicloud Defense, you must edit the JSON at `assignableScopes` to add another subscription line or change it to ***** (star) so the custom role can be used with all subscriptions.
7. Click **Save** at the top of the text box.
8. Click **Review + Create** and create the role.
9. Once the custom role is created, return to **Access Control (IAM).**
10. On the top menu bar, click **Add** > **Add role assignment**.
11. In the **Role** dropdown, select the custom role created above.
12. In the **Assign access to** radio button, leave it as the default (User, group, or service principal).
13. Click **Select members**. Select the required users, groups, or service principals and click **Save**.
14. Navigate to **Review + Create** tab and copy the subscription ID from the **Scope** field to the notepad.

Connect an Azure Subscription to the Multicloud Defense Controller from the Multicloud Defense Dashboard

Connect a prepared Azure subscription to Cisco Multicloud Defense from the dashboard with the Azure Cloud Shell onboarding script.

Once you prepared the Azure account and subscription as described in the previous sections, you can link it to the Multicloud Defense Controller.

Procedure

1. In the Multicloud Defense Controller dashboard, click **Add Account** in the Cloud Accounts pane.
2. On the General Information page, select **Azure** from the **Account Type** list box.
3. In step 1, follow the instructions to open an Azure Cloud Shell in bash mode.
4. In step 2, click the **Copy** button.
5. Run the onboarding script in the bash shell.

> **Note**
>
> - If there is another Azure subscription already connected to Multicloud Defense, this script may fail when creating an IAM role with the same existing name. There cannot be more than one IAM role. As a workaround, run the Bash script with the `-p` prefix.
> - To support spoke VNet protection across subscriptions, onboard subscriptions using Active Directory app registrations.

6. Provide a name for this Azure account. You can choose to name this the same as your Azure subscription name. This name is visible on the Multicloud Defense Controller accounts page only.
7. (Optional) Provide a description for the subscription.
   a) Enter the **Directory ID**, also referred as the Tenant ID.
   b) Enter the **Subscription ID** for the subscription being onboarded.
   c) Enter the **Application ID**, also referred to as the Client ID, created by the onboarding script.
   d) Enter the **Client Secret**, also referred to as the Secret ID.
8. Click **Save & Continue**.

The Azure subscription is onboarded and you are returned to the dashboard to see that the new device has been added.

What to do next

- Create a policy in the Azure portal.
- VNet Route Tables for your Azure Subscription on page 50
- Post-Onboarding Procedures on page 52.
- Enable traffic visiblilty.

Accept Marketplace Terms

Accept the Cisco Multicloud Defense virtual machine image terms for each Azure subscription before deploying gateways from Azure Marketplace.

Multicloud Defense Controller creates gateway instances using a Multicloud Defense virtual machine (VM) image from the Azure marketplace. The Terms and Conditions must be accepted for each subscription. Open the Azure cloud shell from the Azure portal website (on the top menubar towards the right side). Choose or switch to bash shell and execute the following command (replace the subscription-id with your subscription id copied in the previous section):

```
az vm image terms accept --subscription $sub_id --publisher valtix --offer
datapath --plan valtix_dp_image
```

Connect an Azure Subscription to the Multicloud Defense Controller from the Multicloud Defense Dashboard

Connect a prepared Azure subscription to Cisco Multicloud Defense from the dashboard with the Azure Cloud Shell onboarding script.

Once you prepared the Azure account and subscription as described in the previous sections, you can link it to the Multicloud Defense Controller.

Procedure

1. In the Multicloud Defense Controller dashboard, click **Add Account** in the Cloud Accounts pane.
2. On the General Information page, select **Azure** from the **Account Type** list box.
3. In step 1, follow the instructions to open an Azure Cloud Shell in bash mode.
4. In step 2, click the **Copy** button.
5. Run the onboarding script in the bash shell.

   Note

   - If there is another Azure subscription already connected to Multicloud Defense, this script may fail when creating an IAM role with the same existing name. There cannot be more than one IAM role. As a workaround, run the Bash script with the `-p` prefix.
   - To support spoke VNet protection across subscriptions, onboard subscriptions using Active Directory app registrations.

6. Provide a name for this Azure account. You can choose to name this the same as your Azure subscription name. This name is visible on the Multicloud Defense Controller accounts page only.
7. (Optional) Provide a description for the subscription.
   a) Enter the **Directory ID**, also referred as the Tenant ID.
   b) Enter the **Subscription ID** for the subscription being onboarded.
   c) Enter the **Application ID**, also referred to as the Client ID, created by the onboarding script.
   d) Enter the **Client Secret**, also referred to as the Secret ID.
8. Click **Save & Continue**.

The Azure subscription is onboarded and you are returned to the dashboard to see that the new device has been added.

What to do next

- Create a policy in the Azure portal.
- VNet Route Tables for your Azure Subscription on page 50
- Post-Onboarding Procedures on page 52.
- Enable traffic visiblilty.

VNet Route Tables for your Azure Subscription

Learn how to review and configure Azure VNet route tables for Cisco Multicloud Defense gateway deployments, including user-defined routing for egress traffic.

For egress deployments, you may need to create a user-defined routing (UDR) table to manually specify the direction for the spoke network. By default, both Azure and has the ability to automatically identify the routing values because of a private information exchange between the peers. This setup is ideal for ingress gateways; however, it is not suitable for egress gateways.

To override these values or the subnet routing table as a whole for your egress deployment, you must reassign the values in the Azure portal. See Azure documentation for more information.

What Kind of Routing Table Is My Gateway Using?

To determine if your routing table is based on a peer device's VNet or not, view the gateway assigned to your subscription in **Infrastructure** > **Gateways** > **Gateways** and click **View Details**. From this window navigate to the **Troubleshooting** > **Datapath Subnet** tab. If there is no routing table visible, then your subscription is utilizing the default routing table pulled from your peer device.

Roles Created by Multicloud Defense

Describes roles created when onboarding a cloud account to multicloud.

When you onboard a cloud service account to Multicloud Defense Controller with the provided script, user roles are created within the parameters of the cloud service provider to ensure that communication between the services are protected. Depending on the cloud service provider, different roles and permissions are created.

The following roles are created when you onboard an account.

Azure IAM Roles

Find the Azure custom role permissions required for the Cisco Multicloud Defense application to access your subscription and manage cloud resources.

You need to create a custom role that will be assigned to the application created for the Multicloud Defense Controller. The custom role gives the application permissions to read inventory information and create resources such as VMs, load balancers. The custom role can be created in multiple ways. One of the easiest ways is to navigate to your subscription and click Access Control (IAM).

When you add a custom role, you need to name the role and edit the JSON file. This file addresses all of the permissions required to allow communication and data transfer between the subscription and the Multicloud Defense Controller. This list contains all the required permissions:

```
"Microsoft.ApiManagement/service/*",
"Microsoft.Compute/disks/*",
"Microsoft.Compute/images/read",
"Microsoft.Compute/sshPublicKeys/read",
"Microsoft.Compute/virtualMachines/*",
"Microsoft.ManagedIdentity/userAssignedIdentities/read",
"Microsoft.ManagedIdentity/userAssignedIdentities/assign/action",
"Microsoft.Network/loadBalancers/*",
"Microsoft.Network/natGateways/*",
"Microsoft.Network/networkinterfaces/*",
"Microsoft.Network/networkSecurityGroups/*",
"Microsoft.Network/publicIPAddresses/*",
"Microsoft.Network/routeTables/*",
"Microsoft.Network/virtualNetworks/*",
"Microsoft.Network/virtualNetworks/subnets/*",
"Microsoft.Resources/subscriptions/resourcegroups/*",
"Microsoft.Storage/storageAccounts/blobServices/*",
"Microsoft.Storage/storageAccounts/listkeys/action",
"Microsoft.Network/networkWatchers/*",
"Microsoft.Network/applicationSecurityGroups/*",
"Microsoft.Compute/diskEncryptionSets/read",
"Microsoft.Insights/Metrics/Read"
"Microsoft.Network/locations/serviceTagDetails/read",
"Microsoft.Network/locations/serviceTags/read",
"Microsoft.CognitiveServices/*/read",
"Microsoft.CognitiveServices/accounts/listkeys/action",
"Microsoft.Network/virtualHubs/*",
```

```
"Microsoft.Network/virtualHubs/hubRouteTables/*",
"Microsoft.MarketplaceOrdering/offertypes/publishers/offers/plans/agreements/*"
"Microsoft.Authorization/roleAssignments/read",
"Microsoft.Authorization/roleDefinitions/read"
```

Post-Onboarding Procedures

Learn which tasks to complete after onboarding an Azure subscription to Cisco Multicloud Defense so you can prepare and secure your cloud environment.

Use the following procedures to wrap up and secure your Azure account with Multicloud Defense.

Azure VNet Setup

Learn how to prepare Azure VNet subnets and security groups for a Cisco Multicloud Defense gateway deployment.

This document describes the requirements and resources (subnets, security-groups) to be created in your VNet so that you can create Multicloud Defense Gateways in the VNet.

Subnets

Learn about the management and datapath subnet requirements for deploying Cisco Multicloud Defense gateways in an Azure VNet.

When configuring your gateway deployment, the Multicloud Defense Controller will prompt you for the **management** and **datapath** subnet information.

The **management** subnet is a public subnet that must be associated with the route table that has a default route to the Internet. The Multicloud Defense Gateway instance has an interface attached to this subnet that it uses to communicate with the Multicloud Defense Controller. This interface is used for policy pushes and other management and telemetry activities between the Multicloud Defense Controller and the Multicloud Defense Gateway instances. Customer application traffic **does not** flow through this interface and subnet. The interface is associated with the**management** security group, which is described in the Security Groups section below.

The **datapath** subnet is a public subnet that must be associated with the route table that has a default route to the Internet. The Multicloud Defense Controller creates a network load balancer (NLB) in this subnet. In addition, a Multicloud Defense Gateway instance has an interface attached to this subnet. The customer applications traffic **flows** through this interface. A security policy is applied to the trafficingressing through this interface. The interface is associated with the **datapath** security group, which is described in the Security Groups section.

Security Groups

Learn about the management and datapath security group requirements for Cisco Multicloud Defense gateway instances in an Azure VNet.

The management and datapath security groups are associated with the respective interfaces on the Multicloud Defense Gateway instance, as described in the subnets section above.

The **management** security group must allow outbound traffic that allows the gateway instance to communicate with the controller. Optionally, for inbound rules, enable port 22 (SSH) to allow SSH access to the gateway instance. SSH is not mandatory for the Multicloud Defense Gateway to function properly.

The **datapath** security group is attached to the datapath interface and allows traffic from the Internet to the Multicloud Defense Gateway. Currently, the Multicloud Defense Controller does not manage this security group. An outbound rule must exist, allowing the traffic to egress this interface. Inbound ports must be opened for each port that is configured in the Multicloud Defense Controller security policy and used by the Multicloud Defense Gateway.

For example, if an application is running on port 3000 and is proxied by the Multicloud Defense Gateway on port 443, port 443 must be opened on the datapath security group. This example also implies that port 3000 is open on the security group attached to your application.

Launch ARM Template

Deploy an Azure Resource Manager template to create a VNet and the network resources required for a Cisco Multicloud Defense gateway.

Use the provided template to create all of the resources described on this page.

This template creates a new VNet. This is very useful to get started on Multicloud Defense without touching your existing production environment.

The provided template will create the following resources:

- VNet.
- Management subnet.
- Datapath subnet.
- Management security group with outbound rules.
- Datapath security group with outbound rules and Inbound rules for port 443.

You can create additional subnets to run apps and create app-specific security groups, as needed.

Use the following steps to launch an ARM template:

Procedure

1. Log into your Azure account and Deploy a custom template.
2. Click **Build your own template in the editor**.
3. Copy the content from the ARM template and paste into the editor.
4. Click **Save**.
5. Select the *Subscription*, *Resource group* and the *Region*.
6. Click **Review+ create**.
7. Wait for a few minutes for all the resources to be created.

5 GCP

Topics:

- GCP Overview
- Connect a GCP Project to the Multicloud Defense Controller from the Multicloud Defense Dashboard

Explains key attributes and offerings of Google Cloud Platform as a cloud service.

GCP Overview

An overview of onboarding Google Cloud projects and folders to Cisco Multicloud Defense, including supported use cases and configuration options.

GCP Project and GCP Folders

Multicloud Defense currently supports both GCP projects and GCP folders although these components are supported separately. Note the following limitations and exceptions for both of these options.

A GCP project has the potential to contain GCP resources like virtual machines, storage buckets, databases, and more. It can be used to create, enable, and use all Google Cloud services.

- Projects can be onboarded with Terraform, manual onboarding, and scripted onboarding.
- Projects are ideal for environments requiring orchestration, including discovery and investigation.
- You can interact with each project individually through the Multicloud Defense dashboard.

As of Version 23.10, you can connect a GCP folder with Terraform. A GCP folder contains projects, other folders, or a combination of both. Organization resources can use folders to group projects under the organization resource node in a hierarchy.

- Consider folders without the `roles/compute.admin` permission enabled as empty, and do not use them.
- Projects associated with onboarded folders are used for asset and traffic discovery only.
- Projects associated with onboarded folders do not accommodate orchestrating service VPC or gateway creation.
- Permissions made to folders from the GCP console must be made at the folder level. Therefore, perform, Multicloud Defense actions at the folder level.
- When creating a new gateway, do not use 172.16.0.0/16 subnet.

If you want to onboard a GCP folder, see refer to the Terraform repository.

Overview Procedure

The following is an overview of how to connect your GCP project. A shell **script** is provided by Multicloud Defense and facilitates an easy connective process as part of a wizard. The script automates the following steps so you don't have to:

1. Create two service accounts.
2. Enable the following APIs (Compute Engine, Secret Manager).
3. Create the two following VPCs (management, datapath).
4. Create firewall rules to allow traffic to the Multicloud Defense Gateway (app traffic) in the datapath VPC.
5. Create firewall rules to allow management traffic from Multicloud Defense Gateway to the Multicloud Defense Controller in the management VPC.

If you find that the script does not work, or if you need to manually change your settings, these actions can be executed using the GCP cloud console web UI, or using the gcloud CLI. See the alternative method of connecting your project here.

Overview of Creating a GCP Controller Service Account

Create a Google Cloud controller service account and key so Cisco Multicloud Defense can access and manage resources in your project.

The Multicloud Defense Controller uses the controller service account to access and manage resources in your GCP project. You must create the account and generate a key. The key is added to the controller as part of Account onboarding to the controller.

Procedure

1. In your GCP dashboard, open IAM in your GCP project.
2. Click **Service Accounts**.
3. Create **Service Account**.
4. Provide a name and ID, such as `multicloud-firewall`, and click **Create**.
5. Add **Compute Admin** and **Service Account User** roles.
6. Click **Continue**.
7. Click **Done**.
8. Click on the newly created account, scroll down to Keys and in the dropdown for Add Key and select **Create New Key**.
9. Choose JSON (default option) and click **Create**.
10. A file is downloaded to your computer. Save this file to your local drive.

Create a GCP Firewall Service Account

Create a Google Cloud firewall service account with the access required by Cisco Multicloud Defense gateway instances in your project.

The firewall service account is used by the Multicloud Defense Gateway instances running inside your GCP project. The gateways may need to access the private keys stored in the SecretManager for TLS decryption and access storage to store PCAP files etc. (if configured by the user). Also, the gateways many need log writer permissions to send logs from Multicloud Defense Gateway to the GCP logging instance (if configured by the user).

Follow these steps to create a controller service account:

Procedure

1. In your GCP dashboard, open IAM in your GCP project.
2. Click **Service Accounts**.
3. Create **Service Account**.
4. Provide a name and ID, such as `multicloud-firewall`, and click **Create**.
5. Add **Secret Manager Secret Accessor** and **Logs Writer** roles.
6. Click **Continue**.
7. Click **Done**.

Connect a GCP Project to the Multicloud Defense Controller from the Multicloud Defense Dashboard

Connect a prepared Google Cloud project to Cisco Multicloud Defense from the dashboard with the Google Cloud Shell onboarding script.

Once you prepared the GCP project as described in the previous sections, you can link it to the Multicloud Defense Controller.

Before you begin

You must already have a Google Cloud Platform (GCP) project created and have permissions to create VPCs, subnets, and a service account.

Procedure

1. In the left pane of Security Cloud Control, click Multicloud Defense.
2. Click the Multicloud Defense Controller button.
3. In the **Cloud Accounts** pane, click **Add Account**.
4. On the **General Information** page, select **GCP** from the Account Type list box.
5. Login to the Multicloud Defense Dashboard.
6. Click **Manage** and then **Accounts**.
7. Click **Add Account**.
8. Click the link to open the Google Cloud Platform Cloud Shell.
9. Click **Copy**.
10. Run the bash script in the Google Cloud Platform Cloud Shell.
11. Enter the required information:
    a) Enter the **GCP Account Name**. You can choose to name this the same as your GCP project name. This name is visible on the Multicloud Defense Controller only..
    b) (Optional) Enter a **Description**.
    c) Enter the **Project ID** for the GCP project.
    d) Enter the **Client Email** for the service account created for Multicloud Defense Controller.
    e) Enter the **Private key** of the service account.
    f)
12. Click **Save & Continue**.

What to do next

Enable traffic visibility.

Roles Created by Multicloud Defense

Describes roles created when onboarding a cloud account to multicloud.

When you onboard a cloud service account to Multicloud Defense Controller with the provided script, user roles are created within the parameters of the cloud service provider to ensure that communication between the services are protected. Depending on the cloud service provider, different roles and permissions are created.

The following roles are created when you onboard an account.

GCP IAM Roles

Find the Google Cloud IAM roles and permissions required for Cisco Multicloud Defense Controller and gateway service accounts.

This document explains the details of the service accountes created by the CloudFormation template used in the previous section.

The CloudFormation template creates the following accounts:

- **ciscomcd-controller service account** - This account is used by the Multicloud Defense Controller to access your GCP project to create resources (Multicloud Defense Gateway), load balancers for gateways, and read information about the VPCs, subnets, security group tags, and more. See Overview of Creating a GCP Controller Service Account on page 56 for more information.
- **ciscomcd-firewall service account** - This account is assigned to the Multicloud Defense Gateway (compute VM instances). The account provides access to the secret manager (private keys for TLS decryption) and storage. Also, the gateways many need permissions to send logs from Multicloud Defense Gatewayto the GCP logging instance (if configured by the user). See Create a GCP Firewall Service Account on page 57 for more information.

6 OCI

Topics:

- Prepare Your OCI Account
- Connect the Oracle OCI Tenant to the Multicloud Defense Controller from the Multicloud Defense Dashboard

Learn how to connect and manage OCI accounts in Multicloud Defense to protect Oracle Cloud Infrastructure environments.

Prepare Your OCI Account

Learn how to prepare an Oracle Cloud Infrastructure tenant for Cisco Multicloud Defense onboarding, including regional access, image terms, and required account resources.

Before you onboard an Oracle Cloud Infrastructure (OCI) tenant to Multicloud Defense, your OCI account needs to be properly set up. This topic covers the general steps required to prepare the tenant.

Note

Multicloud Defense supports both Ingress and Egress/East-West protection for OCI. Traffic visibility is not supported.

In order to onboard the OCI tenant, you need to subscribe to the US West (San Jose) region. If this region is not subscribed, then the onboarding of the OCI tenant will result in an error.

In order to deploy a Multicloud Defense Gateway into OCI, the terms and conditions for the Multicloud Defense compute image **must** be accepted in each OCI compartment. Otherwise the deployment will show an unauthorized error.

This procedure shows how to prepare your OCI environment to successfully connect with Multicloud Defense. For OCI-specific documentation on how to accomplish these requirements, see OCI documentation.

Overview of Automated Steps

Multicloud Defense provides a script that automates the preparation of your OCI account. The automation does include the required group, policy, permissions, and users included in the manual procedure listed here.

1. Open your Oracle Cloud Shell or any linux-based shell prompt.
2. Enter and execute the following command:

```
bash <(curl -sSL
https://raw.githubusercontent.com/valtix-security/cli-oci-setup/main/oci_onboarding.sh)
```

3. Once successfully finished, continue to Connect the Oracle OCI Tenant to the Multicloud Defense Controller from the Multicloud Defense Dashboard on page 63.

Overview of Manual Steps

Perform the following procedure to manually prepare your OCI account:

1. Create a Group.
2. Create a Policy. Note that the policy must have the `root` Compartment selected and the following permissions are included:

```
Allow group <group_name> to inspect instance-images in compartment
<compartment_name>
Allow group <group_name> to read app-catalog-listing in compartment
<compartment_name>
Allow group <group_name> to use volume-family in compartment <compartment_name>
Allow group <group_name> to use virtual-network-family in compartment
<compartment_name>
Allow group <group_name> to manage volume-attachments in compartment
<compartment_name>
```

```
Allow group <group_name> to manage instances in compartment <compartment_name>
Allow group <group_name> to {INSTANCE_IMAGE_READ} in compartment
<compartment_name>
Allow group <group_name> to manage load-balancers in compartment
<compartment_name>
Allow group <group_name> to read marketplace-listings in tenancy
Allow group <group_name> to read marketplace-community-listings in tenancy
Allow group <group_name> to inspect compartments in tenancy
Allow group <group_name> to manage app-catalog-listing in compartment
<compartment_name>
Allow group <group_name> to read virtual-network-family in tenancy
Allow group <group_name> to read instance-family in tenancy
Allow group <group_name> to read load-balancers in tenancy
Allow group <controller-group> to manage cloudevents-rules in tenancy
Allow group <controller-group> to manage ons-family in tenancy
```

3. Create a User.
4. Add the User to the Group.
5. Create an API Key for the User.
6. Record the *user* and *tenancy* OCIDs.
7. Accept the Terms and Conditions.

What to do next:

Connect the OCI account to your Multicloud Defense using Connect the Oracle OCI Tenant to the Multicloud Defense Controller from the Multicloud Defense Dashboard on page 63.

Connect the Oracle OCI Tenant to the Multicloud Defense Controller from the Multicloud Defense Dashboard

Connect a prepared Oracle Cloud Infrastructure tenant to Cisco Multicloud Defense from the dashboard with Oracle Cloud Shell and your tenant credentials.

Before you begin

Review the requirements in Prepare Your OCI Account on page 62.

Procedure

1. In the Cloud Accounts pane, click **Add Account**.
2. In the General Information page, select **OCI** in the Account Type list box.
3. Click **Oracle Cloud Shell** to launch the native shell prompt.
4. Copy the command provided in the Multicloud Defense Setup wizard and paste it into your cloud shell. Execute the command.

   This command automates the process of creating an IAM policy, OCI group, and an OCI user that facilitate the communication between your OCI account and the Multicloud Defense.
5. Fill in the following fields:
   - **OCI Account Name**- Used to identify this OCI Tenant within the Multicloud Defense Controller.
   - **Tenancy OCID** - Tenancy Oracle Cloud Identifier obtained from the OCI User.

- **User OCID** - User OCID obtained from the OCI User.
- **Private Key** - API private key that was assigned to the OCI User.

---

7 Remove a Cloud Service Provider From Multicloud Defense

Topics:

- Delete a GCP project from Multicloud Defense
- Delete an AWS account from Multicloud Defense
- Delete an Azure account from Multicloud Defense
- Delete an OCI Account From Multicloud Defense

Remove a cloud service provider from Multicloud Defense when you no longer want to manage its cloud environment.

Use the following procedures to terminate communications and permissions between Multicloud Defense and your cloud service provider. This action includes removing any gateways or Vnets that have been created within the Multicloud Defense Controller as well as any roles or permissions you have set up within your cloud service provider. You must perform **all** of the steps for a complete cleanup of every Multicloud Defense instance.

Note that some of these procedures do not occur in the Multicloud Defense Controller and that you may need access to the cloud service provider's dashboard to execute these procedures.

Delete a GCP project from Multicloud Defense

Remove a Google Cloud project from Cisco Multicloud Defense and delete the associated gateways, Service VPCs, roles, and service accounts.

Use the following procedure to delete a GCP account from the Multicloud Defense Controller and remove all instances of Multicloud Defense from your GCP project. You must delete any subnets, Service VPCs/VNets, or gateways created in the Multicloud Defense Controller prior to deleting Multicloud Defense from your account.

> **Note**
>
> This procedure requires you to remove orchestration preparation from both the Multicloud Defense UI and the GCP dashboard.

Procedure

1. Delete any current gateways or Service VPCs/VNets from Multicloud Defense:
   a) In the Multicloud Defense Controller, navigate to **Infrastructure** > **Gateways** > **Gateways**.
   b) Select the gateway associated with the account so its checkbox is checked.
   c) Expand the **Actions** drop-down menu and select **Delete.**
   d) Confirm the deletion.
   e) In the Multicloud Defense Controller, navigate to **Infrastructure** > **Gateways** > **VPCs/VNets**.
   f) Select the VPCs associated with the account so the checkbox is checked.
   g) Expand the **Actions** drop-down menu and select **Delete.**
   h) Confirm the deletion.

   > **Note**
   >
   > You do not have to delete any affiliated subnets after you delete the VPC and gateway.

2. Delete the GCP project from Multicloud Defense Controller.
   a) In the Multicloud Defense Controller, navigate to **Systems and Accounts** > **Account** > **CSP Accounts**.
   b) Select the GCP account so the checkbox is checked.
   c) Expand the **Actions** drop-down menu and select **Delete.**
   d) Confirm the deletion.
3. Delete the Multicloud Defense Controller service account from GCP.
   a) Log into the GCP dashboard.
   b) Open IAM in your GCP project.
   c) In the navigation pane to the left, click **Service Accounts**.
   d) Select the project associated with the Multicloud Defense.
   e) Under the **View by Principals** tab, search for the `ciscomcd-controller`.
   f) Ensure the row's checkbox is checked, then click **Delete**.
4. Delete the Multicloud Defense firewall service account from GCP.
   a) Log into the GCP dashboard.

b) Open IAM in your GCP project.
c) In the navigation pane to the left, click **Service Accounts**.
d) Select the project associated with the Multicloud Defense.
e) Under the **View by Principals** tab, search for the `ciscomcd-gateway`.
f) Ensure the row's checkbox is checked, then click **Delete**.

The Google Cloud project is entirely disassociated from Cisco Multicloud Defense, including its accounts and roles.

What to do next
Validate deletion impacts across connected services; alternate configuration as needed.

Delete an AWS account from Multicloud Defense

Remove an Amazon Web Services account from Cisco Multicloud Defense and delete its associated gateways, service VPCs, and cloud resources.

Use the following procedure to completely remove an AWS account from your Multicloud Defense.

After you delete the AWS account, it may take up to 24 hours for the cloud service provider to clean up all objects within the S3 bucket that is associated with your account.

Procedure

---

1. Log into Security Cloud Control and launch the Multicloud Defense Controller.
2. Navigate to **Infrastructure** > **Gateways** > **Gateways**.
3. Locate the gateway associated with your account and select the checkbox, then click the **Actions** drop-down menu.
4. Select **Disable**. This action automatically removes all the associated gateway instances.
5. Make sure the gateway's checkbox is still selected and click the **Actions** drop-down menu again.
6. Select **Delete**. This action removes the load balancers associated with the AWS account.
7. Navigate to **Systems and Accounts** > **Account** > **CSP Accounts**.
8. Locate the AWS account in the list and select it so the checkbox is check.
9. Click the **Actions** drop-down menu and select **Delete**.
10. Confirm you want to delete the account.

The AWS account along with all associated resources will be successfully removed from Cisco Multicloud Defense.

What to do next
Monitor deletion progress and verify successful object removal after a potential 24-hour cleanup period.

Delete an Azure account from Multicloud Defense

Remove an Azure account from Cisco Multicloud Defense and delete its associated gateways, service VNets, certificates, and cloud resources.

This task involves manually deleting all instances of the Azure account from Multicloud Defense:

Before you begin

You must delete any subnets and VNets created in the Multicloud Defense Controller prior to deleting Multicloud Defense from your Azure account.

> **Note**
>
> This procedure requires you to remove orchestration preparation from both the Multicloud Defense UI and the Azure portal.

Procedure

1. Log into Security Cloud Control and launch the Multicloud Defense Controller.
2. If you did not create a user-assigned Managed Identity for the key vault, continue to step 4. If you **did** create a key for the Azure account, do the following:
   a) Navigate to **Policies** > **Security Policies** > **Certificates**.
   b) Select the certificate associated with the account and then open the **Actions** drop-down menu.
   c) Select **Delete** and confirm the deletion of the certificate for the key vault.
3. In the Multicloud Defense Controller, delete any gateways or VNets associated with the account.
   a) Navigate to **Infrastructure** > **Gateways** > **Gateways** to delete any gateways previously created.
   b) Select the gateway associated with the account so its checkbox is checked.
   c) Expand the **Actions** drop-down menu and select **Delete.**
   d) Confirm the deletion.
   e) In the Multicloud Defense Controller, navigate to **Infrastructure** > **Gateways** > **VPCs/VNets** to delete any VNets previously created.
   f) Select the VNet associated with the account so the checkbox is checked.
   g) Expand the **Actions** drop-down menu and select **Delete.**
   h) Confirm the deletion.
   i) In the Multicloud Defense Controller, navigate to **Systems and Accounts** > **Account** > **CSP Accounts**.
   j) Select the Azure account so the checkbox is checked.
   k) Expand the **Actions** drop-down menu and select **Delete**.
   l) Confirm the deletion.
4. Delete the Multicloud Defense Controller role in Azure.
   a) Log into the Azure portal.
   b) Navigate to **App Registrations**.
   c) Select the **Owned Appications** tab.
   d) Select the **ciscomcd-controller-app** application.
   e) Once selected, click **Delete** at the top of the window.
   f) Confirm the deletion.
   g) Navigate to, or search for, **Subscriptions** and click **Access Control (IAM)**.
   h) Select the **Roles** tab at the top of the window.
   i) Search for **ciscomcd-controller-role-rw** and select it so the checkbox is checked.
   j) Click **Remove** at the top of the window.

The Azure account resources are successfully removed from Cisco Multicloud Defense, preventing further associations and interactions.

What to do next
Validate that all associated resources are offline and securely handled.

Delete an OCI Account From Multicloud Defense

Remove an Oracle Cloud Infrastructure account from Cisco Multicloud Defense and delete its API keys, users, groups, policies, and associated resources.

Use the following procedure to remove an OCI cloud environment from Multicloud Defense:

Procedure

1. Log into the OCI console.
2. Delete the API key. For more information, refer to the Deleting API Signing Keys from a Roving Edge Infrastructure Device.
3. Delete Multicloud Defense Users. For more information, refer to the Deleting a User.

   Note

   When you remove the user from the OCI account, this does not delete the audit data of the user from when it was valid.

4. Delete the Multicloud Defense Group. For more information, refer to the **Deleting Groups** chapter in the Oracle Cloud Infrastructure Documentation.
5. Delete any and all Multicloud Defense access policies. For more information, refer to the **Deleting an Access Policy** chapter in the Oracle Cloud Infrastructure Documentation.
6. Delete the OCI account from Multicloud Defense Controller. .
   a) In the Multicloud Defense Controller, navigate to **Systems and Accounts** > **Account** > **CSP Accounts**.
   b) Select the OCI account so the checkbox is checked.
   c) Expand the **Actions** drop-down menu and select **Delete.**
   d) Confirm the deletion.

8 AI Defense

Topics:

- AI Defense Integration with Multicloud Defense

Learn about AI Defense and how it helps you protect AI applications and services in your cloud environment.

When you enable the integration with AI Defense you can secure your AI assets including the associated activity, types of connections, and number of identities accessing unsanctioned models. By utilizing this function with your Multicloud Defense tenant, you can discover models in your environment and apply AI Defense runtime protection. Once discovered, models can also be tested for safety and security using AI Defense validation. For more information on AI Defense and what it can do to improve your safety and security with AI, see AI Defense documentation.

AI Defense Integration with Multicloud Defense

Learn how to integrate Cisco AI Defense with Cisco Multicloud Defense to discover AI assets and apply runtime protection to supported cloud traffic.

Support and Limitations

Depending on how integrated you want your Multicloud Defense tenant and AI Defense to be, here are the requirements and limitations:

- You **must** have a Security Cloud Control account prior to accessing either AI Defense or Multicloud Defense.
- Only **egress** Multicloud Defense Gateways are currently compatible with AI Defense.
- If you want the full AI Defense experience with AI Runtime monitoring of LLM prompts and responses, you **must** "Secure Your Account" and add a service VPC or VNet to your gateway.
- Profiles and rulesets created in Multicloud Defense directed to support your AI Defense integration **must** be modified in the Multicloud Defense Controller; you cannot delete or modify a Multicloud Defense policy or ruleset in the AI Defense dashboard.
- You must have an AI Defense license. See Administration for more information on AI Defense licenses.
- AI Discovery of assets is done for AWS and Azure.

Overview

The following list is an overview of the procedure to enable both aspects of these products to allow a secure integration:

1. Log into your Multicloud Defense tenant.
2. Generate an API Key with the Multicloud Defense dashboard.
3. Connect your Multicloud Defense tenant to AI Defense.
4. Onboard a cloud service provider to Multicloud Defense. Be sure to add the correct permissions to your AWS account to allow secure access and communication.
5. Enable Traffic Visibility.
6. Secure Your Account.
7. Attach your profile to the policy ruleset of an egress gateway.

9 Asset and Inventory Discovery

Topics:

- Discovery Summary
- Inventory
- Security Insights
- Rules and Findings

Learn how Multicloud Defense discovers cloud assets and inventory to help you understand your cloud environment.

Discovery is an important component of Multicloud Defense's approach of **" Discover, Deploy, and Defend"** .

Discovery provides real-time visibility into the current resources deployed in any onboarded cloud accounts. In addition, it provides an interface into VPC flow logs and DNS logs to give a complete picture of your cloud deployment. The Multicloud Defense Controller, through the permissions granted to the IAM role (AWS and OCI), AD app registration (Azure) or the service account (GCP), periodically crawls your cloud resources, and also keeps tab on the changes, to maintain an "evergreen" inventory model of the resources.

Using the **Discovery** tab, you have the ability to see the attributes of your resources and how they are interconnected. Multicloud Defense collates this information into a succinct view of the security posture of all your resources with respect to the configuration and with context to the traffic flows.

Discovery Summary

Learn how to use the Cisco Multicloud Defense Discovery Summary to review cloud traffic, inventory, applications, and security findings.

The Discovery Summary page is a collection of widgets that summarize the available traffic and inventory. You can use the **Filter** at the top of the page to change the history of the widgets.

Traffic Summary Widgets

Currently, Multicloud Defense presents a condensed block of the traffic in two widgets: one for DNS traffic and one for VPC and VNet flow logs. These windows into traffic differentiate between malicious traffic and DNS or VPC/VNet traffic, respectively. Click inside either of these widgets to zoom into a specific time frame.

You can enable or disable logs on either of these widgets from this summary page by simply clicking the **Logs** toggle. For more information on either of these types of logs and the traffic that is compiled, see Types of Traffic on page 181

Discovery Summary

The Discovery summary is a series of windows of inventory recovered by Multicloud Defense as part of the discovery process when connecting your cloud service provider. These statistics are condensed here for a quick preview. For more details, refer to Inventory on page 74

Inventory

An overview of Cisco Multicloud Defense inventory and how continuous cloud asset discovery supports visibility and security planning across connected accounts.

Through permissions granted to the IAM role (AWS and OCI), AD app registration (Azure) or the service account (GCP), Multicloud Defense continuously maintains an "evergreen" inventory model of the cloud resources as well as real-time discovery that exists in your cloud service provider accounts, subscriptions and projects that are relevant to apply advanced network security. Once discovered, the resources are available in workflows that enable administrators to quickly deploy security rules to mitigate risks of exposed applications. Any activity is reported immediately through the Multicloud Defense Controller.

When inventory is enabled, Multicloud Defense Controller will perform a full inventory discovery periodically. The default is 60 minutes, but is tunable). Real-time inventory discovery is enabled on regions where the CloudFormation template was deployed.

Part of the discovery process highlights the logs of each cloud service provide. Note the following types of logs per service provider:

- **AWS** - VPC flow logs, Route53 flow logs, and DNS logs.
- **Azure** - VNet flow logs.
- **GCP** - VPC flow logs.

Note that Multicloud Defense does note provide the same level of support for all cloud service providers.

Applications

Learn how to find discovered cloud applications, define application tags, and start protection workflows in Cisco Multicloud Defense inventory.

Application shows all load balancers and API gateways for the cloud accounts. Under the Applications section of **Inventory**, there are three filter buttons: **Known Tags**, **Tags**, and **Applications**. Within **Applications**, users can invoke a workflow to create and apply protection for the specific application.

Application Tags

Create a list of **Application Tags** used to identify applications. During the inventory discovery stage, all discovered load balancers that have the specified tags are treated as applications.

As an example, you can assign the **Application Tags** to all load balancers that act as applications. The value of this tag is shown as the **Application Tags** in the discovered inventory. See the table below as a visual example:

Load Balancer	Tag	Value
Load Balancer 1	ApplicationName	Billing
Load Balancer 2	ApplicationName	UserManagement

The discovered inventory will show the **Billing** and **UserManagement** applications in the discovered application assets.

To create a list of **application tags**, click **Create**.

Parameter	Description
Name	Pre-populated.
Description	User-specified description.
Value	The tag value that will be used assign to the load balancers.

For more information on application tags, see Application Tags on page 311.

Known Tags

Known Tags show applications identified by application Load Balancers in your cloud account that the administrator has identified by a known tag. These known tags are listed in **Systems and Accounts** > **Account** > **Account**.

Tags

Tags shows all applications identified by application load balancers with fields showing the tag keys and tag values and whether these applications are secured by Multicloud Defense Gateways.

Discovered Assets

Learn how to view and filter cloud assets discovered by Cisco Multicloud Defense across connected AWS, Azure, GCP, and OCI accounts.

When you enable inventory discovery in regions for your cloud account, the Multicloud Defense Controller continuously discovers cloud assets. To view the discovered assets, navigate to **Inventory** > **Inventory**. The default views show the discovered assets for all cloud accounts. To filter to a specific cloud account, use the **Select Account** to specify a particular cloud account and view discovered assets.

The discovered asset categories and what they refer to are as follows:

- Security Groups - AWS Security Groups (SGs), Azure Network Security Groups (NSGs), Azure Virtual Networks (VNets).
- Network ACL - AWS Network Access Control Lists (NACLs).
- Subnets.
- Route Tables.

- Network Interfaces.
- VPCs/VNets - AWS VPCs, Azure VNets and GCP VPCs.
- Applications - Applications are identified by AWS Application Load Balancers (ALBs).
- Load Balancers.
- Instances - AWS Instances, Azure Virtual Machines and GCP Compute Instances.
- Tags - AWS Tags, Azure Tags and GCP Labels.
- Certificates - AWS Certificates Manager (ACM) certificates.

Enable Asset Discovery and Inventory

Enable Cisco Multicloud Defense asset discovery for selected cloud regions and set how often the full inventory refreshes.

To enable discovery of assets in your cloud account:

Procedure

1. Navigate to **Inventory** > **Inventory** .
2. Select the cloud account and click **Manage Inventory**.
3. Select the **Regions** where you have cloud assets that you would wish Multicloud Defense to discover. The refresh interval is the time in minutes after which the inventory is refreshed (recommended default of 60 min). Multicloud Defense also performs continuous discovery using the cloud service provider's APIs and events instead of a regular poll. The refresh time interval specified here is for a full re-crawl; this reconciles all assets for any missed events during real time discovery.

   Note that different refresh intervals can be defined for different regions by adding a new row and selecting the desired regions. A region can belong to a single refresh interval only.
4. Click **Save & Continue**.

   Note

   The Multicloud Defense Controller will request the asset inventory for the newly added region immediately after saving.

What to do next

To review the discovered assets, navigate to **Inventory** > **Inventory** .

Security Insights

Connect a cloud account, enable inventory, and view Cisco Multicloud Defense security insight findings for discovered AWS, Azure, and GCP assets.

Insights are a rules-based evaluations of assets discovered in AWS, Azure and GCP that are presented asfindings. Insights can be used without deploying Multicloud Defense Gateways since they operate on the periodic and real-time inventory monitoring accommodated by the Multicloud Defense Controller.

Procedure

1. In the Multicloud Defense Controller interface, click **Add Account**. As an alternative, we strongly recommend using the Easy Setup wizard to connect to an account. Go through the steps to connect the account.
2. Once the account is connected and onboarded, Enable Inventory.
3. Navigate to **Inventory** > **Summary**. This page displays a summary view of all discovered assets and the **Insight Findings**.

Types of Security Insights

An overview of security insights that help you understand threats, compliance, data protection, performance, and cloud resource configuration.

Security insights refer to the analysis and understanding of data related to the security of a system, such as a cloud-based gateway. These insights are gathered from monitoring and analyzing network traffic, user behavior, threat intelligence, and other security-related data. By leveraging the following apsects of security insights, a cloud-based gateway can enhance its overall security, ensuring data protection, reliability, and trustworthiness for users and clients.

- **Threat Detection** helps identify potential threats and vulnerabilities in real-time. By analyzing patterns and anomalies, they can alert administrators to suspicious activities or breaches.
- **Compliance** within your insights assists witht compliance maintenance for industry standards and regulations by providing detailed reports and audits of security measures and incidents.
- **Incident Responses** enable quicker response times to security incidents by providing detailed information about the nature and scope of the threat, helping to minimize damage.
- **Resource Optimization** provides a better understanding the types of threats most likely to affect their systems, organizations can allocate resources more effectively, focusing on the most significant risks.

Read through the following types of security insights to understand what the dashboard can do.

Security Groups

Find discovered security groups, review their associated resources and network exposure, and create insight rules from search criteria.

Customers often struggle with the proliferation of **Security Groups**. Security groups are often shared amongst resources that could present risk. Changes made to a security group intended for a specific resource could impact a larger group of resources.

Security groups provides a list of all security group, their details and the set of resources utilizing the security group. The **Is Inbound Public** and **Is Outbound Public** fields indicate security groups configured with `0.0.0.0/0`.

In the search window, define the search criteria based on fields and their values with the option to create a rule based on the search criteria.

Rules

Rules provide a view of security groups based on their configured Inbound and outbound rules.

Ports

Ports provide a view of security groups based on their configured inbound and outbound ports.

Application Security Groups

Learn how Azure application security groups support discovered resource membership and security management in Cisco Multicloud Defense.

Application Security Groups are an Azure construct similar to the AWS security group. Azure application security groups have a member of the security group that contains that system and its interfaces. It has both membership and security controls. As a result, Multicloud Defense uses this membership construct to build dynamic policies. Create and use an application security group within an Azure environment, Multicloud Defense recognizes the change and adapts the policy to incorporate it.

For more information about Azure's application security groups and how they operate, see the Microsoft Azure documentation.

Network ACL

Find discovered network access control lists and review their inbound rules, outbound rules, and public network exposure.

Network access control list (ACL) provides a list of all network ACLs and their details. The **Is Inbound Public** and **Is Outbound Public** fields indicate network ACLs configured with 0.0.0.0/0.

Rules

Rules provide a view of network ACLs based on their configured inbound and outbound rules.

Subnets

Find discovered cloud subnets and review whether each subnet allows automatic assignment of public IP addresses.

Subnets provides a list of all subnets and their details. The **Is Public** field indicate subnets that are publicly accessible based on whether auto-assign public IP is enabled.

Route Tables

Find discovered cloud route tables and review whether their routes provide public internet access.

Route Tables provides a list of all route tables and their details. The **Is Inbound Public** and **Is Outbound Public** fields indicate route tables that are configured to provide default access the internet.

Network Interfaces

Find discovered network interfaces and review their security groups, route tables, and public network exposure in Cisco Multicloud Defense.

Network Interfaces provides a list of all network interfaces and their details. The **Is Inbound Public** and **Is Outbound Public** fields indicate network interfaces that are configured with a security group that is open (0.0.0.0/0), or route tables that allows default access to the internet.

VPCs\VNets

Find discovered virtual private clouds and virtual networks and review their details in Cisco Multicloud Defense.

VPCs/VNets provides a list of all VPCs/VNets and their details.

Applications

Find discovered application load balancers, review their protection status, and start a Cisco Multicloud Defense workflow to secure an application.

Applications provides a list of all deployed application load balancers and their details. The **Secured** field identifies whether a Multicloud Defense Gateway and security policy is applied to secure the application and offers an ability to invoke a workflow to protect the application.

Load Balancers

Find discovered cloud load balancers and review their internet exposure, protection status, and cloud web application firewall status.

Load Balancers generally improve application performance by increasing response time and reducing network latency. They perform several critical tasks such as distributing the load evenly between servers to improve application performance and redirecting client requests to a geographically closer server to reduce latency.

Load Balancers and Supported Cloud Service Providers

At this time you can configure load balancers for an AWS gateway.

When you configure a load balancer in Multicloud Defense, the **Public** field shows whether resource is an internet-facing load balancer. The **CSP WAF Enabled** shows whether a CSP WAF has been enabled for the application load balancer.

Instances

Find discovered cloud instances and review their assigned interfaces, security groups, and public network exposure in Cisco Multicloud Defense.

Instances provides a list of all instances along with summary information on the number of security groups and interfaces that are assigned and configured for the resource. The **Is Inbound Public** and **Is Outbound Public** fields indicate instances that have network interfaces that are configured with a security group that is open (0.0.0.0/0), or route tables that allows default access to the internet.

Tags

Learn how to review tags on discovered cloud resources and use them to organize assets and support security management.

Tags in a cloud-based environment are metadata labels that can be attached to resources such as virtual machines, storage accounts, and network interfaces. When appropriately utilized, tags can help categorize and organize resources based on various attributes like environment, department, or application. Likewise, they can also be used to enforce security policies or assist in defining and managing access controls. By tagging resources, administrators can set permissions that align with the roles and responsibilities associated with those tags. Within the Multicloud Defense Controller, tags provide a list of all VPCs/VNets, subnets, security groups, instances and load balancers that are configured with tags.

Certificates

Find discovered AWS Certificate Manager certificates and review their issuer, domain name, and expiration details in Cisco Multicloud Defense.

Certificates provides a list of all certificates available in AWS certificates manager along with summary information on issuer, domain name and expiry date.

Topology

Learn how to use the Cisco Multicloud Defense topology map to explore cloud assets, traffic, accounts, and regions with configurable filters.

This tab shows a high-level map view by region of cloud assets in cloud accounts. You can finetune the visuals with the **Filter** bar at the top of the screen. From here you can determine what cloud service provider accounts you want to pull data drom, which region of the world, specific VNet or VPCs, instances, and a period of time in history.

The **Global View** of the world map allows you to scroll in for a closer look at specific regions that are dictated by the Filter bar mentioned above. Immediately to the left of the map you can dictate which types of traffic and inventory you want to view. Check and uncheck the boxes appropriately for what you want to see .

Insights

Learn how Cisco Multicloud Defense rules evaluate discovered AWS, Azure, and GCP assets and present matching resources as security findings.

Insights are a rules-based evaluations of assets discovered in AWS, Azure and GCP that are presented as findings.

Rules

Rules are a set of evaluations to identify findings in discovered assets. Multicloud Defense provides a set of default rules. New rules can be created by selecting an inventory category (e.g., security groups, applications, load balancers, tags, etc.), defining a search criteria, selecting **Add Rule** and specifying additional required information. Navigate to **Insights** > **Rules** to view the new rule. From there you can operate against existing and newly discovered assets.

Findings

Findings is a list of discovered assets that match the defined set of rules.

Rules and Findings

Learn how Cisco Multicloud Defense rules evaluate discovered cloud resources and present findings that help you identify security concerns.

Rules can be configured to place checks and guardrails on your cloud resources.

Rules and Findings

Learn how Cisco Multicloud Defense rules evaluate discovered cloud resources and present findings that help you identify security concerns.

Rules can be configured to place checks and guardrails on your cloud resources.

Pre-Defined Rules

Find the default Cisco Multicloud Defense rules that identify potential security concerns in discovered cloud applications, networks, instances, and certificates.

Multicloud Defense Controller has some basic pre-defined rules:

- Application load balancers with no cloud service provider WAF enabled.
- Security groups with few instances (< 5) that have ingress open. Lots of low utilization security groups can create gaps that are hard to see and may make it easy to exploit.
- Instances with two or more network interfaces.
- Security groups with open outbound (0.0.0.0/0) access.
- Public subnets - all AWS subnets with **Auto-Assign Public IP** enabled.
- Security groups with with too many egress ports (25 or more) open to the internet.
- Security ports with too many ingress ports (5 or more) open to the internet.
- Security groups with 65,535 ports open for ingress with public access enabled.

- Certificates expiring in 30 days - AWS Certificate Manager only.

The cloud resources that match the rules, will be flagged as findings with a matching severity.

For information on custom rules, see Custom Rules on page 81.

Custom Rules

Create custom Cisco Multicloud Defense rules to evaluate discovered cloud resources and optionally generate alerts for matching security findings.

The user can configure additional rules for a resource.

1. Navigate to **Inventory** and select a resource e.g. load balancers.
2. Create a rule criteria in the text area and select **Add Rule**.
3. Enter content for the following entries and the number of finding meeting the rule criteria.
   - Name
   - Description
   - Severity
   - Default Action
   - Type
   - Account
4. Click **Save**.

The default action of the rule can be either **info** or **alert**. If a rule is configured with a default action of alert, then any new findings for the rule results in an alert notification from the Multicloud Defense Controller. The following configurations are required if you want a default action of alert.

- Configure **Alert Profile** to indicate if the user wants ServiceNow, PagerDuty, or Webhook notifications.
- Configure **Alert Rule of type Discovery** and sub-type **Insights Rule** with the level of severity specified.

Findings

Learn how to review cloud resources that match predefined or custom Cisco Multicloud Defense rules from the dashboard or inventory summary.

Based on the pre-defined and custom rules, you can view the findings for the resources. For easy access, the **Findings Summary** is located in the dashboard, and also in the Summary view in the Inventory tab.

10 Multicloud Defense Gateway and Service VPCs or VNets

Topics:

- Supported Gateway Use Cases
- Configure Service VPCs and Service VNets
- Configure Your Gateway

Learn about Multicloud Defense gateways and service VPCs or VNets for securing cloud network traffic.

Multicloud Defense Gateway is a network-based security platform comprising 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 or 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 or 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 zone's failure from a cloud service provider does not compromise the security posture for running applications.

> **Note**
>
> If the Multicloud Defense Gateway loses connectivity to the Multicloud Defense Controller for any reason, the Multicloud Defense Gateway will continue to function with existing policies.

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.

Gateway Details

To view the **Gateway Details** page for already established gateways are available in **Infrastructure** > **Gateways** > **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.

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.

Gateway Actions

The following gateway actions occur in the background and do not need any additional tasks or items from you:

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 maximum 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

Find the supported Cisco Multicloud Defense gateway deployment models and traffic protection use cases for cloud environments.

Firewall Threat Defense Virtual

Explains what Secure Firewall Threat Defense Virtual devices are, their main features, deployment considerations, licensing models, and scaling characteristics.

Cisco Secure Firewall Threat Defense virtual brings next-generation firewall functionality to virtualized environments, enabling consistent security policies to follow workloads across your physical, virtual, and cloud environments.

You can provision a tenant in Security Cloud Control to onboard and manage your Firewall Threat Defense virtual device and navigate between Multicloud Defense and Cloud-delivered Firewall Management Center.

Guidelines for deploying Firewall Threat Defense Virtual Gateway

By deploying Firewall Threat Defense virtual gateway using Multicloud Defense, you gain advanced network security and management features that are not available in standard Internet Service Provider gateways. You can also protect your network when you deploy an Firewall Threat Defense virtual gateway using Multicloud Defense.

Because this is a multiproduct task, you must navigate between both Multicloud Defense and Cloud-delivered Firewall Management Center to complete the steps. Multicloud Defense deploys and registers the Firewall Threat Defense virtual device including interfaces, gateway configuration, NAT rules, platform settings, whereas you edit your access policy, rules, and objects in your Cloud-delivered Firewall Management Center account.

Use these guidelines when you create a Firewall Threat Defense virtual that is managed by Cloud-delivered Firewall Management Center:

- When you create a service VPC and apply a gateway to your Firewall Threat Defense virtual environment, Multicloud Defense automatically creates a subnet and a corresponding security group for the secondary interface.
- Based on your requirements, deploy an ideal number of Firewall Threat Defense virtual gateways in the Cloud-delivered Firewall Management Center tier within a tenant, for optimal efficiency. During upgrades, additional capacity is temporarily required to bring up new instances before the earlier instances can be decommissioned.
- For Azure deployments, the default SSH username for Firewall Threat Defense virtual is centos, and access is authenticated using the SSH key that is provisioned during the gateway deployment.

  Alternatively, you can use Azure Serial Console to log in to Firewall Threat Defense virtual using the admin username and the password provided to you during gateway deployment.
- In AWS, for East-West traffic inspection, in the access control policy, use VNI for both the source security zone and destination security zone. Because this is U-turn traffic, the source and destination security zones are the same.
- In AWS, for egress traffic, in the access control policy configure the VNI zone as the source security zone and the Outside zone as the destination security zone.
- In Azure, for East-West traffic inspection, in the access control policy configure the Inside zone as the source security zone and destination security zone. Because this is U-turn traffic, the source and the destination security zone are the same.
- In Azure, for egress traffic, in the access control policy configure the Inside zone as the source security zone and the Outside zone as the destination security zone.

Configurations by Multicloud Defense

This section describes the cdFMC objects and policies that Cisco Multicloud Defense automatically creates in your cdFMC instance during FTDv gateway deployment.

Multicloud Defense automatically creates all the following objects in cdFMC when you create and enable an FTDv gateway. You do not need to pre-create these – Multicloud Defense handles them during deployment.

Policies (per gateway)

Policy	Name	Description
Access Policy	`<gatewayName>-policy`	Default action: Block. You add your own access rules to this policy. Multicloud Defense does not delete this policy when the gateway is removed–your rules are preserved.
NAT Policy	`<gatewayName>-NATPolicy`	For Egress gateways, Multicloud Defense adds an Auto NAT masquerade rule. For Ingress gateways, you add your own Manual NAT rules.
Platform Settings Policy	`<gatewayName>-PlatformSettings`	Configures the health probe HTTP listener. Do not modify this policy.

Device Group (per gateway)

- Multicloud Defense creates a **Device Group** named after your gateway.
- All FTDv instances for that gateway are automatically added to this group.
- Policies are assigned at the Device Group level, so they apply uniformly to all instances.

Security Zones (shared across all gateways)

These zones are shared tenant-wide in cdFMC and are created once:

Zone	When Created
`ciscomcd-inside`	Always
`ciscomcd-outside`	Always
`ciscomcd-VNI`	AWS gateways with Egress or East-West security type only

Network Objects (shared)

Object	Value	Cloud
`ciscomcd-Any`	`0.0.0.0/0`	All
`ciscomcd-AzureUtilityIP`	`168.63.129.16`	Azure only
`ciscomcd-GCPHealthProbeNetworkGroup`	Google health probe CIDRs (`35.191.0.0/16`, `130.211.0.0/22`, `209.85.152.0/22`, `209.85.204.0/22`)	GCP only

Network Objects (per gateway)

Object	What it represents
`<gatewayName>-InsideGatewayIP[-<az>]`	Inside subnet gateway IP (per Availability Zone)
`<gatewayName>-OutsideGatewayIP[-<az>]`	Outside subnet gateway IP (per Availability Zone)
`<gatewayName>-InsideUTurn`	Network group used for u-turn routing (Egress/East-West)
`<gatewayName>-VPCNetworkGroup`	AWS only—services VPC CIDR blocks (used for health probe routing)

Configurations (Per FTDv Instance)

When each FTDv instance registers with cdFMC and is available online, Multicloud Defense automatically:

1. **Assigns physical interfaces** to the correct security zones (`ciscomcd-inside`, `ciscomcd-outside`).
2. **Configures static routes** for health probe return traffic and default routing.
3. **Deploys all assigned policies** (access, NAT, platform settings) to the instance.

For **AWS Egress / East-West gateways**, Multicloud Defense also configures GENEVE (GWLB) encapsulation—VTEP and VNI interface—automatically.

For **GCP gateways**, Multicloud Defense creates a loopback interface and VRF per instance to terminate load balancer health probes.

User Configurations

Ingress Gateways—Manual NAT Rules

Multicloud Defense does **not** create Manual NAT rules for ingress gateways. You must create these in cdFMC:

- Add **Manual NAT rules** to `<gatewayName>-NATPolicy` specifying the destination ports you want to forward.
- Multicloud Defense reads these rules after each deployment and automatically provisions the corresponding NLB listener ports on AWS.
- A maximum of **40 listener ports** can be configured. Exceeding this limit will trigger an alert and block the NLB update.

Access Policy Rules

Multicloud Defense creates the Access Policy with a default **Block** action. Add your own **Allow** rules for the traffic you want to permit. Multicloud Defense will not modify or delete rules you add.

Optional: Provide Your Own Access Policy

If you prefer to use an existing Access Policy, you can provide its cdFMC UID when creating the gateway (`FmcAccessPolicyUid`). Multicloud Defense will verify the policy exists and use it instead of creating a new one.

Security Type Behavior

Security Type	Auto NAT Rule	Outside Interface	AWS VNI/VTEP	NLB Port Sync
Ingress	No	Yes	No	Yes—reads your Manual NAT rules
Egress	Yes (masquerade)	Yes	Yes (AWS only)	No
East-West	No	No	Yes (AWS only)	No
Ingress + Egress	Yes (masquerade)	Yes	Yes (AWS only)	Yes

Cloud Behavior

Feature	AWS	Azure	GCP
VNI/VTEP (GENEVE for GWLB)	Yes (Egress/E-W)	No	No
Health probe source	Services VPC CIDRs	`168.63.129.16`	Google health probe ranges
Loopback interface per instance	No	No	Yes
VRF per instance	No	No	Yes

When you delete a gateway:

- The **NAT Policy**, **Platform Settings Policy**, and **Device Group** are deleted by MCD.
- The **Access Policy** is intentionally **not deleted**–your access rules are preserved.
- Shared objects (`ciscomcd-inside`, `ciscomcd-outside`, `ciscomcd-Any`, and so on) are not deleted–they are reused by other gateways.

Naming Convention

All objects use the gateway name as a prefix: `<gatewayName>-<suffix>`. Names that exceed cdFMC's character limit are automatically shortened with a hash suffix to remain unique.

Prerequisites for deploying Firewall Threat Defense Virtual Gateway

- You must provision a tenant on Security Cloud Control. For more information, refer to Create a Security Cloud Control Tenant section.
- You must have an active Cloud-delivered Firewall Management Center account.
- You must have an active Multicloud Defense account.

Limitations for deploying Firewall Threat Defense Virtual Gateway

Review these limitations when you create a Firewall Threat Defense virtual gateway that is managed by Cloud-delivered Firewall Management Center:

- You cannot deploy Firewall Threat Defense virtual gateways in cluster mode.
- VPN (Site-to-site and Remote Access) are not supported.
- Only the centralized model is supported.
- You must create a **new** service VPC. VPCs created before this feature was introduced do not support this functionality. New VPCs support both Multicloud Defense gateways and Firewall Threat Defense virtual gateways.
- To use **Cisco Smart Licensing**, you must purchase a license through your Cisco seller or partner.
- Perform software upgrades through the Multicloud Defense dashboard.
- Firewall Threat Defense virtual configuration changes made to the Firewall Threat Defense virtual gateway must be done through the Multicloud Defense dashboard. You must only edit security policies for the devices deployed, from Multicloud Defense. Do not perform any changes directly in Cloud-Delivered Firewall Management Center. It must be used only to manage access control policies and NAT policies.

- Changes to the access control policy (such as adding and removing security rules) must be done through Cloud-delivered Firewall Management Center.
- At this time, Jumbo frames are not supported.

Licensing requirements for deploying Firewall Threat Defense Virtual Gateway

The Multicloud Defense Gateway supports Multicloud Defense licensing, Smart Licensing, and Pay-As-You-Go (PAYG) Licensing.

Multicloud Defense Licensing is subscription-based access to software delivered via the cloud, with benefits such as ease of access, scalability, and regular updates. Multicloud Defense licensing is typically subscription-based with a recurring fee that often includes maintenance, updates, and support as part of the package. Most Multicloud Defense licensing models offer scalability, enabling your organization to easily adjust the number of users or features based on current needs. Because this licensing model is based in the cloud, you can access the latest features and updates immediately. For more information, refer to the Cisco Multicloud Defense Ordering Guide.

Smart Licensing is more suited for utilizing existing licenses for cost savings in specific deployment scenarios. Opting for a smart license allows your organization to use pre-purchased licenses with their cloud or hybrid deployments, leveraging licenses that are already owned by the organization to be more cost effective. Smart licensing is often used for specific deployment types, such as cloud or hybrid scenarios, allowing organizations to migrate existing workloads without acquiring new licenses.

Pay-As-You-Go (PAYG) Licensing is a usage-based model which enables you to use the license on demand. This licensing model is available only on AWS.

When creating your gateway for your Firewall Threat Defense virtual you have the opportunity to select one of these three options, depending on your resources.

> **Note**
>
> Once you deploy a gateway with an Firewall Threat Defense virtual device, you cannot change the licensing model. You can change the performance tier of the licensing model you selected. If you change the licensing model, you will need to deploy a new gateway.

The different types of licenses are:

- **Base licensing** - Enables basic firewall and networking functionality, such as stateful firewalling, routing and NAT features. This is a standard foundational license that is mandatory for all deployments.
- **Carrier Licensing** - Enables the inspection of protocols such as GPRS, M3UA, Stream Control Transmission and so on. For more information, refer to Carrier License.
- **Threat Licensing (Threat Protection)** - Provides access to Intrusion Prevention System (IPS) features, enabling threat detection and prevention. It also includes signature-based detection for known vulnerabilities and threats.
- **Malware Licensing (Malware Defense)** - Enables advanced malware protection through Cisco Advanced Malware Protection (AMP) and includes file trajectory, sandboxing, and retrospective malware detection capabilities.
- **URL Filtering Licensing** - Allows URL filtering to control and monitor web traffic and provides access to Cisco's global threat intelligence for web categorization.

To accompany these license types, there is tiered licensing intended to align performance and throughput that can potentially assist in cost optimization and scalability:

- **FTDv5**: Up to 100 Mbps. This is only supported in Azure and GCP.
- **FTDv10**: Up to 1 gigabit per second (Gbps). This is only supported in Azure and GCP.

- **FTDv20**: Up to 3 Gbps
- **FTDv30**: Up to 5 Gbps
- **FTDv50**: Up to 10 Gbps
- **FTDv100**: Up to 16 Gbps

Create and deploy a Firewall Threat Defense Virtual Gateway

Follow this set of procedures to successfully create and deploy the Firewall Threat Defense virtual gateway:

1. Enable Cloud-delivered Firewall Management Center on Your Security Cloud Control Tenant.
2. Onboard a CSP
3. Create a Service VPC.
4. Create an FTDv Gateway.
5. Configure your access policy in Cloud-delivered Firewall Management Center.

Auto scaling in Firewall Threat Defense Virtual

Auto scaling helps ensure that applications have the necessary resources to maintain performance while optimizing cost efficiency.

Note that auto scaling in Firewall Threat Defense virtual gateways is not the same as auto scaling in Multicloud Defense Gateway. For more information, refer to Gateway Auto-scaling in Multicloud Defense Gateway Components for native auto scaling abilities. The information provided here addresses how auto-scaling impacts Firewall Threat Defense virtual gateways.

- **Monitored Metrics** - To ensure auto scaling is performed appropriately and timely for your specific environment, the metrics relating to system memory, Snort CPU, and data plane are monitored.
- **Scale Out** - If any of the metrics register thresholds that are above the accepted threshold, the environment scales up or out to accommodate the load or resources associated with the instances to handle traffic spikes. This triggered event is for the availability zone and not the designated region.

  > **Note**
  >
  > Scaling out your environment increases deployment time.

- **Scale In** - If any of the metrics register below the allowed thresholds, the environment scales down or in to to ensure efficiency and reduce cost. This triggered event is for the availability zone and not the designated region.

Multicloud Defense Egress Gateways

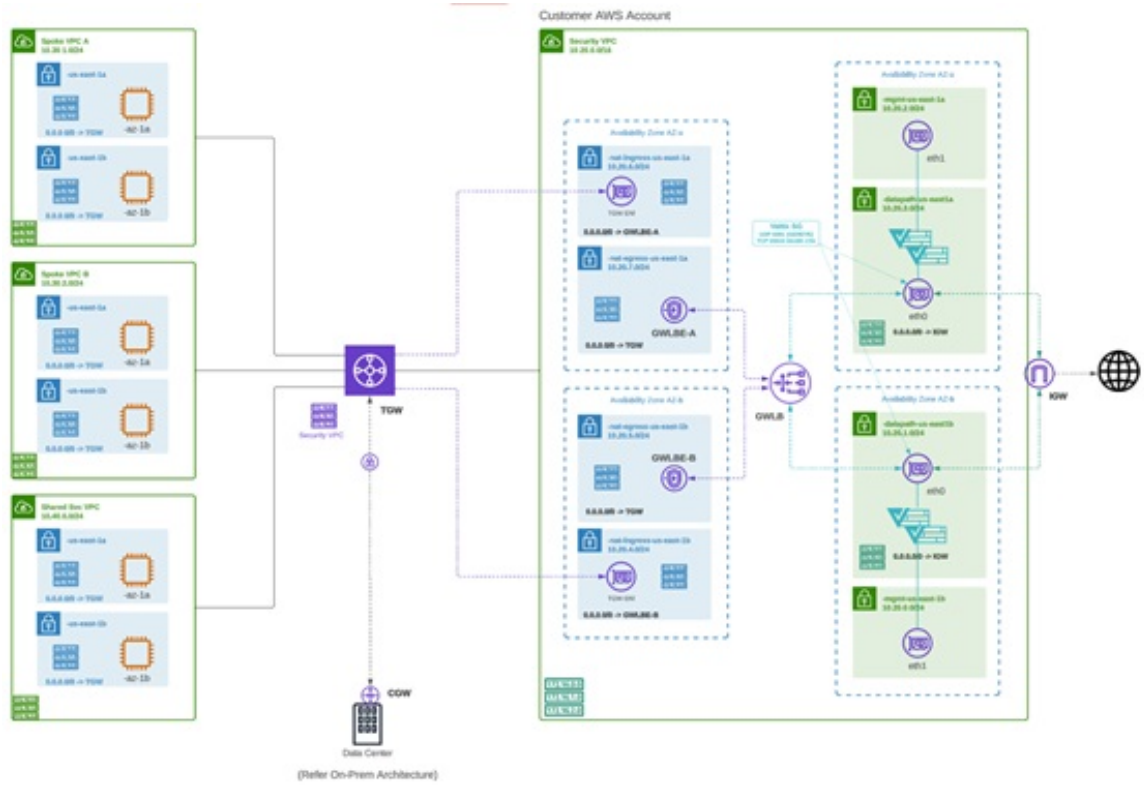
Learn how Cisco Multicloud Defense egress gateways protect outbound cloud traffic with proxy or forwarding deployment options.

Egress/East-West

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 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.

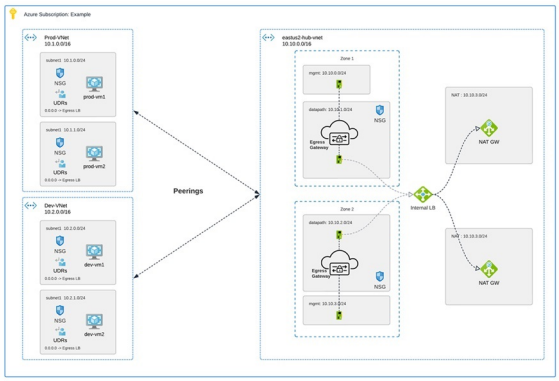
Network Address Translation (NAT) gateways are gateways designed to originate from within your cloud service provider. Egress traffic appears from a single IP address, or at least one per availability zone. By building a gateway and hosting it from within your cloud environment you can potentially increase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

These are the configurations that are supported:

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

> **Note**
>
> Multicloud Defense Egress Gateway and Cisco FTDv, when orchestrated by Multicloud Defense, can perform NAT using AWS Elastic IP (Public IP) addresses associated with their instances. These Public IPs may change during events such as redeployment, upgrade, or autoscaling. To ensure all outbound traffic from your AWS environment translates to specific, static IP addresses, an AWS NAT Gateway (NGW) is required.

The 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 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:

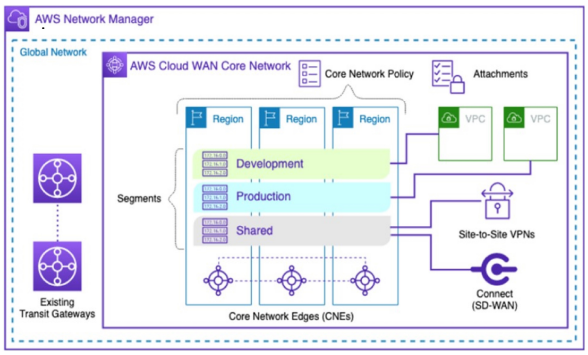
- **Centralized Policy Management**: The core network policy, written in a declarative language, defines segments, AWS region routing, and how attachments should map to segments. With a single network policy, customers can manage their entire network's routing and security policies, reducing the need 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.

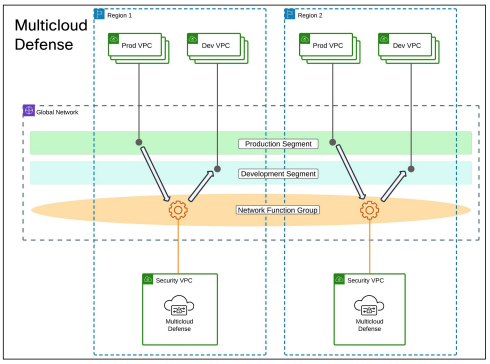
- **Seamless Integration**: It integrates with other AWS services and third-party solutions, enabling customers to build a cohesive and comprehensive network infrastructure with minimal friction.
- **Enhanced Visualization**: AWS Network Manager provides several dashboard visualizations, including world 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:

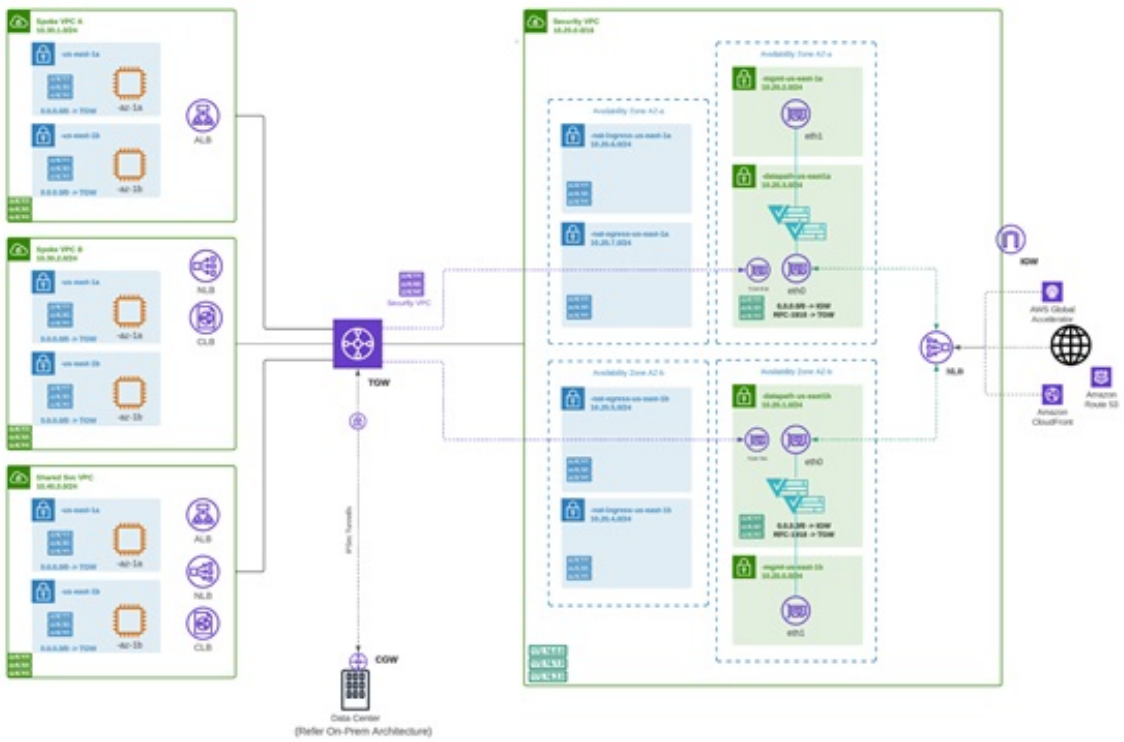
- NAT gateways are mandatory for service VPCs.
- Dual-Hop and Edge Selection is currently **not** supported.
- Due to a limitation in AWS CloudWAN limitation that does not support SNAT-enabled traffic for forwarding, traffic drops for policy rulesets configured with SNAT. We **strongly** recommend you disable SNAT in your Multicloud Defense policy ruleset.
- To add an additional service VPCs in different regions (CNE) use one of the 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.

Multicloud Defense Ingress Gateways

Learn how Cisco Multicloud Defense ingress gateways protect public-facing cloud applications and inspect incoming traffic.

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 diagram is an example of an AWS account with an ingress gateway in a centralized mode.

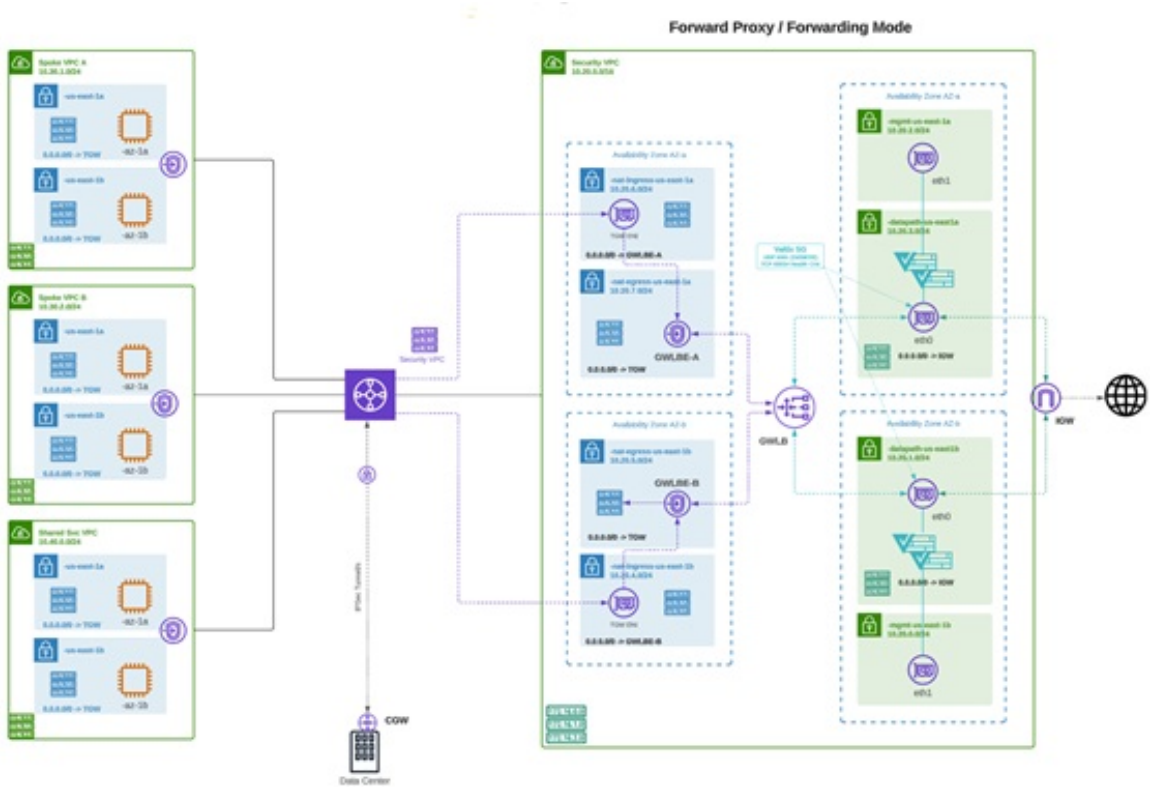


Multicloud Defense East-West Gateways

Learn how Cisco Multicloud Defense east-west gateways help protect traffic moving between cloud subnets, VPCs, and VNets.

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.

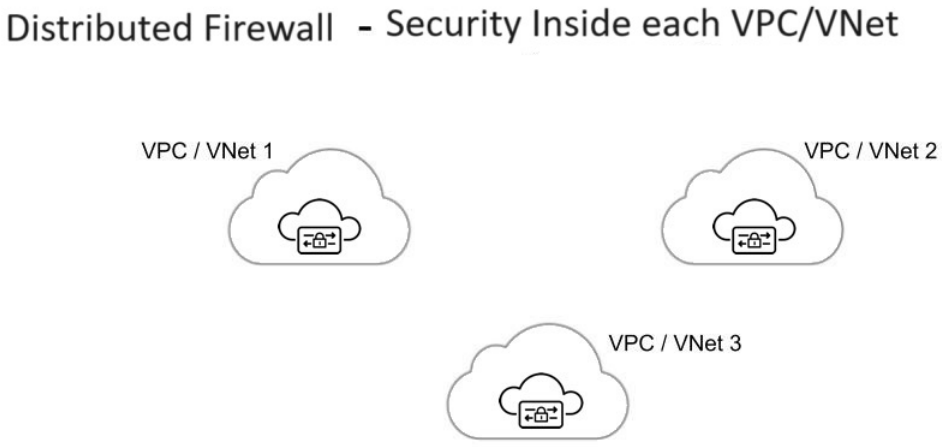
The diagram shows an example of an AWS account with an east-west gateway in a centralized mode:



Multicloud Defense Distributed Gateways

Learn how distributed Cisco Multicloud Defense gateways protect applications by deploying a gateway in each VPC or VNet.

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



Gateways Deployed in Centralized / Hub Mode

Learn how a centralized Cisco Multicloud Defense gateway can protect applications across multiple spoke VPCs or VNets through a shared security network.

You have applications running in multiple VPCs/VNet. You would like to secure all the applications through a centralized security service VPCs or VNets. This model deploys the Multicloud Defense Gateway in a service VPC. You attach all the application VPCs (Spoke VPCs) and the Service VPC to the AWS Transit Gateway or VNet and VPC peering in Azure and GCP. Multicloud Defense provides an option to orchestrate the AWS Transit Gateway, service 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 3: AWS - Using AWS Transit Gateway

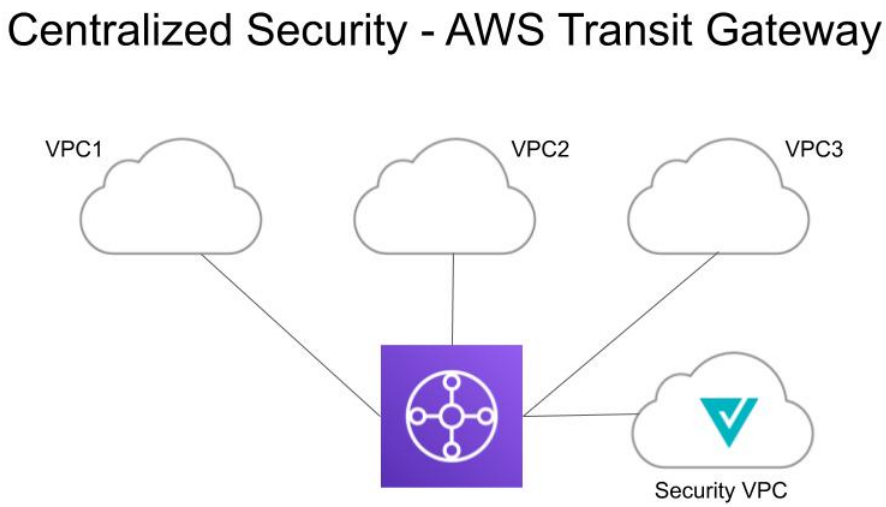


Figure 4: Azure - VNet Peering

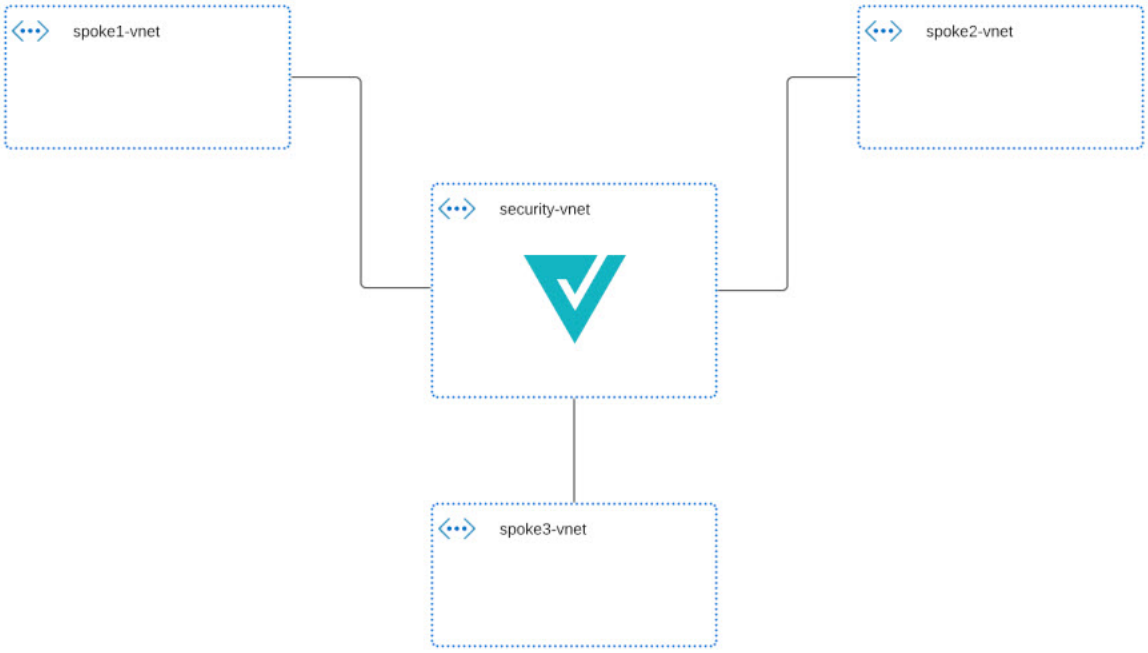
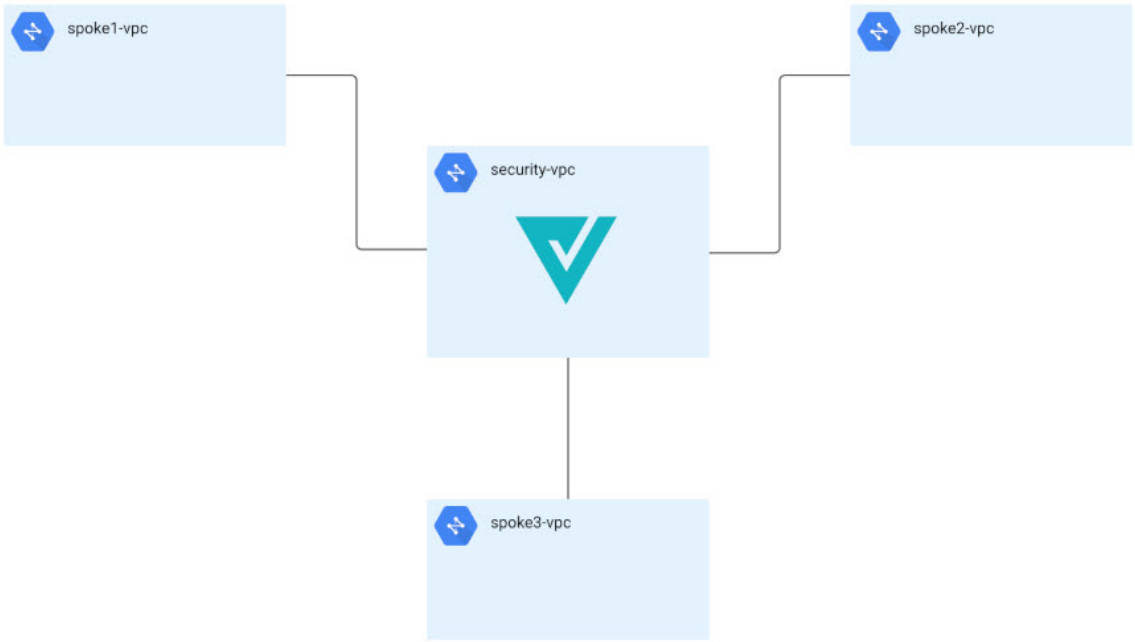


Figure 5: GCP - VPC Peering



Advanced Gateway Configuration: Use Your Own Load Balancer

Learn how to use an existing AWS or Azure load balancer with a Cisco Multicloud Defense ingress gateway.

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 Gateway.

> **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 instances. 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 the Amazon AWS documentation.

Azure Load Balancer

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 backend pool that contains at least one cluster of Multicloud Defense Gateway instances. This scenario is ideal if you want to 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 to use this capability. Consider the following prerequisites and limitations:

- You **must** have your Azure load balancer already configured.
- For the load balancer to work, you **must** enable the floating IP for your load balancer rules.
- 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 choose to configure your Azure load balancer with a resource group, note that 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.

Azure UDP Fragmentation

While it is recommended to avoid fragmentation in general, there may be scenarios where fragmentation does occur and the Azure load balancer drops packets.

To enable support for fragmentation, contact Cisco Support.

Configure Service VPCs and Service VNets

Learn how service VPCs and service VNets support secure, organized Cisco Multicloud Defense gateway deployments across cloud environments.

While not strictly mandatory, we do **strongly** recommend creating and attaching a VPC or VNet to your Multicloud Defense Gateway as part of the deployment process. VPCs and VNets provide the necessary framework to secure, organize, and efficiently manage your network resources while integrating with a firewall gateway.

When you attach a VPC or VNet, you open your network to the following bonuses:

- **Isolation and Security** - A VPC/VNet allows you to create a logically isolated network within a cloud provider, ensuring that your resources are segregated from other users. You can also define security rules that control inbound and outbound traffic to and from your resources, using the firewall gateway to enforce these rules, thus controlling access.
- **Customizable Network Architecture** - The ability to create subnets within a VPC or VNet and organize and segment your resources, while also managing IP address ranges as well as customizing routing tables to direct traffic efficiently within your network.
- **Scalability and Flexibility with Resource Management** - Easily add or remove resources, scale your network, and adjust configurations to meet changing demands.

Without it, your environment faces increased security risks and reduced control.

Before You Begin

Before you create a VPC or VNet for your gateway, we recommend looking over the following prerequisites. Some of these are specific to the cloud service provider you use.

Prerequisites

- If you opt to configure a service 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 on page 47 for the complete list of permissions.
- If you provide your own transit gateway, you are able to attach more than one service VPC or VNet to it. It is even possible to replace an existing service VPC or VNet with a new one without re-deploying the gateway.
- If you create a service VPC for an FTDv gateway, only AWS and Azure accounts are supported.

Shared VPCs in GCP

If you intend to create, or have already created, a shared VPC in your GCP environment, you must do an additional step to enable inventory in the Multicloud Defense Controller. Without these permissions, asset discovery fails and the Inventory page is not reliable. Access the IAM page of your GCP host project (shared VPC) and grant access to the Multicloud Defense Controller service account from the service project . This access is required to allow the service project to interact with shared network resources. You need to grant the following IAM roles for **every** GCP project that is affiliated with the shared VPC:

- `Compute Viewer`
- `Compute Network User`

Important

If you create an environment in GCP where there is a shared VPC setup on one GCP project and the instances are attached to a different GPC project, the Multicloud Defense Gateway ignores all received logs from GCP. This is a default action because Multicloud Defense expects DNS logs to come from a singular project where the instances **and** network are both located.

Create a Service VPC or VNet

Create a service VPC or VNet for a Cisco Multicloud Defense gateway in a supported cloud environment.

Be sure to review the prerequisites and use cases that affect specific cloud service providers in the Before You Begin section for VPCs and VNets. Use the following procedure to create a service VPC or VNet.

Procedure

1. From the Multicloud Defense Controller, navigate to **Infrastructure** > **Gateways** > **VPCs/VNets**.
2. Click **Create Service VPC/VNet**.
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.
   - (AWS/Azure) **CIDR Block** - The CIDR Block for service VNet. This must not overlap with your spoke(application) VNets.
   - (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.
   - (GCP only) **Datapath 2 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.
   - (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 attaching this service VPC to an AWS or Azure NAT gateway, 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. For a a service VNet, Multicloud Defense recommends to select at least two availability zones for resiliency.
   - (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 to go through the 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 on page 47 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.

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 104 and then Add a Gateway and associate the VPC or VNet with the gateway.

If you are creating a service VPC for an FTDv gateway, continue with Create a Firewall Threat Defense Virtual Gateway on page 111

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

Secure Spoke VPC or VNet

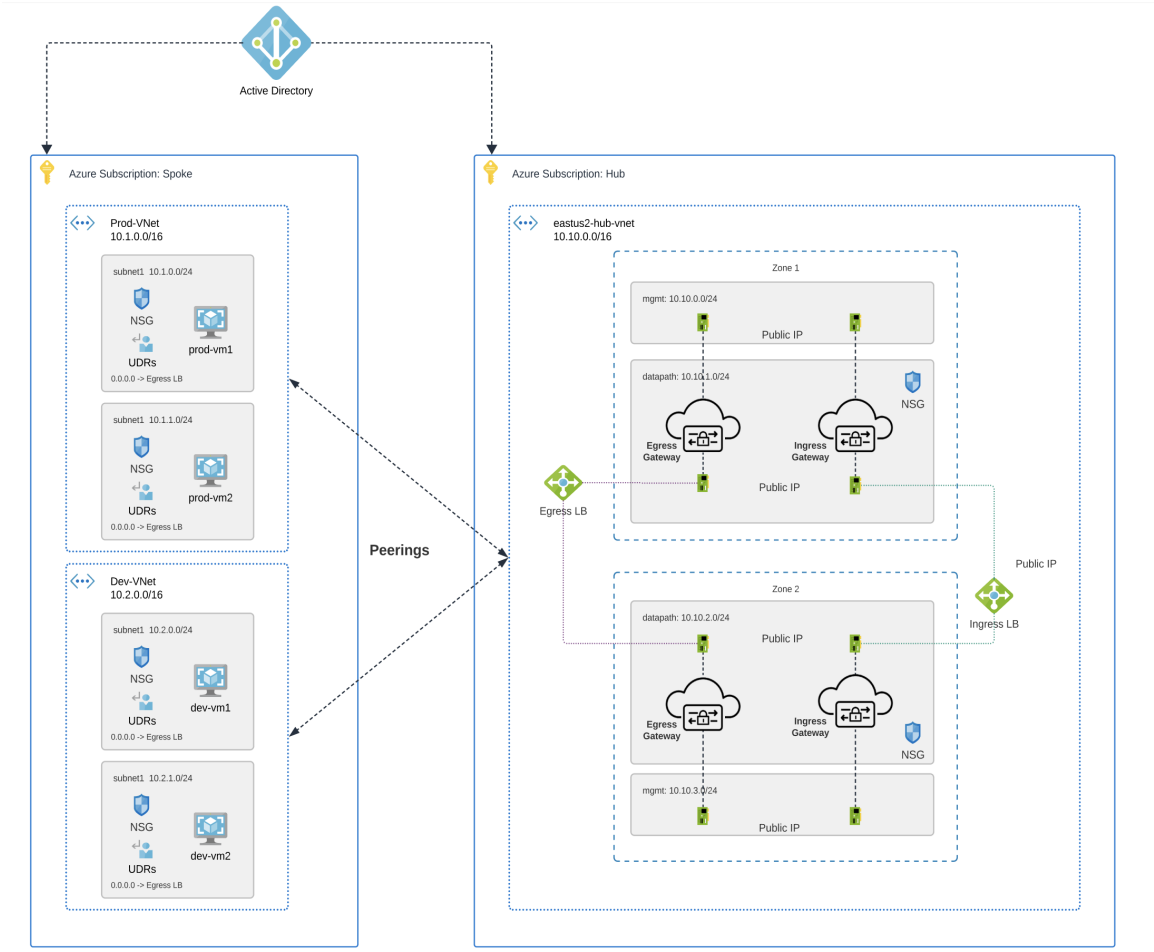
Learn how securing spoke VPCs or VNets helps protect cloud traffic, support network segmentation, and reduce exposure before gateway deployment.

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 your 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 6: Azure Combined Hub - Multisubscriptions



Prerequisites and Limitations

Complete these requirements before you secure 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

Manage spoke VPCs, VNets, and AWS Cloud WAN segments attached to a Cisco Multicloud Defense service network.

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

1. From the Multicloud Defense Controller dashboard, navigate to **Infrastructure** > **Gateways** > **VPCs/VNets**.
2. Select Service VPC or Service VNet and click **Actions**.
3. Click **Manage Spoke VPC/VNet**.
4. To add a segment to a region that is attached to the VPC or VNet displayed in the table, click **Add**.
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.
6. To **Remove** a network segment, select the segment and then click **Remove**
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.
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

Export a spoke VPC or VNet configuration from Cisco Multicloud Defense for use with Terraform.

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

Procedure

1. From the Multicloud Defense Controller dashboard, navigate to **Infrastructure** > **Gateways** > **VPCs/VNets**.
2. Select the Service VPC or Service VNet from the table and click **Actions**.
3. Click **Export.**
4. Multicloud Defensegenerates an export wizard.
5. Either click **Download** to download the terraform locally or scroll down and click **Copy Code** to copy the JSON resource.
6. Manually paste into the terraform script.
7. Within the terraform prompt, execute the command provided in the lower half of the window.
8. Follow the prompts within the terraform prompt to complete the task. Close the export window.

Delete a Spoke VPC or Vnet

Delete a spoke VPC or VNet from a Cisco Multicloud Defense service network configuration.

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

1. From the Multicloud Defense Controller dashboard, navigate to **Infrastructure** > **Gateways** > **VPCs/VNets**.
2. Select the Service VPC or Service VNet from the table and click **Actions**.
3. Click **Delete.**
4. Confirm the deletion of the service VPC or VNet and click **Yes**.

Configure Your Gateway

Learn how to find, filter, and review Cisco Multicloud Defense gateways, cloud providers, instance counts, and configuration details.

View your Multicloud Defense Gateways and statistic in **Infrastructure** > **Gateways** > **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 85.

Before You Begin

Review the cloud provider prerequisites, limitations, and network requirements before configuring Cisco Multicloud Defense gateways, service VPCs, or service VNets.

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

Multicloud Defense Gateway Prerequisites and Limitations

Prerequisites

The supported cloud service providers (AWS, Azure, GCP, OCI) 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.

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.

FTDv Gateway Prerequisites and Limitations

Multicloud Defense orchestrates only the process of creating, deploying and maintaining the gateway. Any policy or rule creation occurs in Cloud-delivered Firewall Management Center. Consider the following prerequisites, limitations, and recommendations to support the integration of both managerial products before you create a gatway.

Prerequisites

You must have the following completed and configured before you create a Multicloud Defense Gateway for your FTDv:

- You must create a **new** service VPC. VPCs created before this feature do not support this functionality; note that when you create a new VPC it can still be used for both Multicloud Defense gateways or FTDv gateways.
- You must have a cloud service provider onboarded to your Multicloud Defense tenant.
- If you are using an AWS account as your designated cloud service provider for the FTDv gateway, you **must** manually accept the AWS Marketplace Terms of Service. Without it, the Multicloud Defense Controller cannot send the required API requests.
- You must have at least one license purchased through your Cisco seller or partner.
- You must have a subscription to Cloud-delivered Firewall Management Center.

For limitations and requirements for this environment, please refer to Firewall Threat Defense Virtual on page 85.

Resources Created by Multicloud Defense

Find the cloud resources Cisco Multicloud Defense creates when you deploy gateways, service VPCs, or service VNets in supported environments.

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 (ingress)
- 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.

FTDv Resources

Multicloud Defense creates the following resources for Service VPCs that are used with Secure Firewall Firewall Threat Defense virtual (FTDv) devices:

- Management subnet.
- Datapath 1 subnet.
- Datapath 2 subnet.
- Three security groups corresponding to subnets.

Add a Multicloud Defense Gateway

Add and configure a Cisco Multicloud Defense gateway for a supported cloud provider and deployment model.

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

Before you begin

Review the Before You Begin on page 105 for information or requirements pertaining to your specific environment before you create a gateway.

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 98 for more information.

Procedure

1. Navigate to **Infrastructure** > **Gateways** > **Gateways**.

2. Click **Add Gateway**.
3. Select the cloud service provider you want to add the gateway to.
4. Enter your gateway information.
    - **Instance Type** - Choose the type of cloud service provider. You may encounter different instance types, depending on your cloud service provider.
        - (Azure only) If you are creating a VPN, and you use Red Hat Enterprise Linux 9 (RHEL 9) as a base operating system, we recommended that you use a Dsv3-Series V3 instance type instead of a Dsv5-Series V5 instance type, as the V5 instance type may not be compatible.
    - **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 synchronization.
    - (Optional) **BGP profile** - Border Gateway Protocol (BGP) used to support VPN Connections. If you intend on utilizing site-to-site VPN tunnels with a Multicloud Defense Gateway you **must** include this profile.
5. Click **Next**.
6. 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. The controller uses this key to access the CLI of deployed gateway instances for debugging 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. For Azure services, these identities can access a private key stored in Azure Key Vault or write PCAP files to Azure Blob Storage.
- (AWS only) **Gateway IAM Role** - Expand the drop-down menu and then you can either create a new IAM role or select an existing IAM role that allows the gateway to perform READ and WRITE operations on your AWS account. Multicloud Defense creates this role for you when you save and deploy the 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.
- **Disk Encryption** - Select the appropriate option from the drop-down menu. If you use a customer-managed encryption key, input the encryption key's resource ID.

7. Select the **Availability Zone**, the **Mgmt Subnet** and the **Datapath Subnet**. The available subnets will be based on the VPC or VNet you selected. For high availability purposes, you can deploy the gateway instances in multiple availability zones. Click the plus button to add a new availability zone and select the parameters for the selected zones. Some cloud service provider regions support only single availability zone and not 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.

8. (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.
9. Click **Next** to view the Advanced Settings.
10. 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.
11. (AWS and Azure only) **Attach Load Balancer**. Click **Add Load Balancer** to create a row for your custom load balancer. You can also remove unnecessary rows by selecting them and clicking **Remove**.
    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.
12. 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 on page 133 for more information.

Create a Firewall Threat Defense Virtual Gateway

Configure and deploy a Virtual Gateway for traffic regulation across AWS, Azure, or GCP platforms.

Use the following procedure to create a Firewall Threat Defense virtual gateway:

Create a Firewall Threat Defense virtual gateway for AWS, Azure, or GCP for Ingress, Egress, and East-West traffic.

Before you begin

Review the Before You Begin on page 105 for information or environment limitations before you create a gateway.

Procedure

1. Navigate to **Infrastructure** > **Gateways** > **Gateways**.
2. Click **Add Gateway**.
3. Select the cloud service provider you want to add the gateway to.
4. Enter the following **Gateway Information**. Note that if you do not have the feature enabled, the options for "Gateway Type" may differ and you should refer to Add a Gateway.
   - **Account** - Expand the drop-down menu and select a cloud service provider account that is already onboarded to Multicloud Defense Controller. At this time, only AWS and Azure accounts are supported.
   - **Gateway Type** - Expand the drop-down menu and select **FTDv Gateway**.
   - **Name** - Enter a name for the gateway as it will be displayed in the Multicloud Defense Controller.
   - (Optional) **Description** - Enter a description for the gateway. We recommend using unique identifiers to differentiate this gateway from others that may have a similar name or purpose.
   - **Instance Type** - Choose the type of cloud service provider. This selection should match the cloud service provider that is selected in the before-mentioned "Account" field. Note that there may be multiple variations of instances depending on which cloud service provider you are using.
   - **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.
5. Click **Next**.
6. Provide the following parameters where applicable. Note that several fields are auto-filled in based on the configuration of the service VPC you created for this gateway. See Create a Service VPC or VNet on page 101 for more information.
   - **Security** - Select Ingress, Egress, or East-West based on your requirement.
     - AWS - You can select Ingress, Egress, or East-West.
     - Azure - You can select Ingress, Egress, or East-West.
     - GCP - You can select Ingress, Egress, or East-West.
   - **FTDv Version** - Select the version of software to run on the FTDv device when Multicloud Defense creates and deploys the device.
   - **Policy Ruleset** - Select the access control policy ruleset to associate with this gateway. If you do not already have an access policy ready you can either use the default policy or create a new policy from this menu.

- **Admin Password** - Enter a password for the admin of the FTDv device. Follow the on-screen prompts for a strong password.
- **License Model** - Click the toggle to highlight your preferred licensing model:
  - **Multicloud Defense Licensing** - The Multicloud Defense licensing model allows you to use all the Firewall Threat Defense virtual features without procuring individual feature licenses in Cisco Smart Licensing account.
  - **Smart Licensing** - Select this option if you intend to purchase a license or have an unused license already purchased through your Cisco Smart Account.
  - **Pay-As-You-Go (PAYG) Licensing** - Select this option if you intend to use the license on demand.
- **Performance Tier** - Expand the drop-down menu and select the appropriate performance tier for your device. Note that FTDv50 is auto-selected. See the tiers listed in Licensing.
- **License Types** - Expand the drop-down menu and select the appropriate license type that you have purchased or will purchase in the future. Note that the **Base** license is auto-selected. See the different licensing types listed in Licensing.
- **Region** - Select the region this gateway will be deployed into.
- **VPC/VNet ID** - Select the ID of the service VPC or VNet to associate with the gateway. Identifying a service VPC or VNet in this step confirms the management and datapath security groups as well as the availability zones. To modify these values, create a new service VPC and add it to this gateway.
- (AWS only) **Key Pair** - Expand the drop-down menu and select the key pair that is associated with the cloud account you selected in the previous screen.
- (Azure and GCP only) **Key Selection** - Select the type of key, its size, and its activation and expiration dates. Choose either **SSH Public Key** or **SSH Key Pair**. Based on your selection, enter the appropriate information in the text field when prompted.
- **Resource Groups** - Select the resource group to associate the gateway with.
- (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 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 1** - Select the security group to associate with the datapath 1 interface.
- **DataPath Security Group 2** - Select the security group to associate with the datapath 2 interface.
- (GCP only) **Datapath VPC** - Select a VPC.
- (GCP only) **Datapath Network Tag 1** - Choose a security group from the drop-down list.
- (GCP only) **Datapath Network Tag 2** - Choose a security group from the drop-down list.
- (GCP only) **Management VPC** - Select a Management VPC from the drop-down list.
- (GCP only) **Management Network Tag** - Select a Management Network tag from the drop-down list.
- (AWS only) **EBS Encryption** - Expand the drop-down menu and select the appropriate EBS encryption for your specific AWS account.
- (Azure) **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.

7. The **Instance Details** are auto-populated based on the VPC you select. Review the VPC configuration before you deploy the gateway. Click **Next**.
8. (Optional) In the **Advanced Settings** window, check the **Disable Public IP** checkbox; if you opt to disable this step you reduce the exposure of your network to external threats. Private IP addresses can help protect against unauthorized access and potential attacks. They also offer better control of internal traffic.
9. Click **Next**.
10. Review the configuration of the gateway. If you are satisfied with the gateway and want to deploy it, click **Finish**. If you want to modify the settings, click **Back**.

What to do next

Clicking **Finish** at the end of this procedure deploys the gateway; creating and deploying the gateway may take up to 30 minutes. While Multicloud Defense deploys the gateway, it also registers the gateway instance in Cloud-delivered Firewall Management Center for your convenience, creates and applies subnets and security groups for the appropriate interfaces, and applies the access policy you selected in this procedure.

Once the gateway deploys and is displayed in the Cloud-delivered Firewall Management Center we recommend making any necessary updates or inclusions to the policy. Any network objects associated with the policy are shared and displayed in Multicloud Defense's Object page as **network** objects for visibility. Policy orchestration and management is done through Cloud-delivered Firewall Management Center.

We strongly recommend the following actions once the gateway is deployed:

- Secure and Protect your VPC. Securing a spoke VPC for your Firewall Threat Defense virtual centralizes security while maintaining scalability and isolation.
- Access your Cloud-delivered Firewall Management Center account and modify the access control policy. For more information, refer to Introduction to Access Control.
- Create a NAT policy in Cloud-Delivered Firewall Management Center. For more information, refer to the Cloud-Delivered Firewall Management Center Configuration Guide.

Edit a Multicloud Defense Gateway

Edit the configuration of an enabled or disabled Cisco Multicloud Defense gateway.

You can edit a gateway in any state, whether it is enabled or disabled.

> **Note**
>
> If you edit a gateway that is attached to an FTDv device, note that you cannot change the licensing model, only the performance tier of the smart license.

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

Procedure

1. Navigate to **Infrastructure** > **Gateways** > **Gateways**.
2. Select the Multicloud Defense Gateway you want to edit in the table so it is highlighted.
3. Expand the **Actions** drop-down menu and click **Edit**.
4. Modify the gateway configuration as needed.

If you are modifying the license peformance tier for a gateway that is asssociated with an FTDv device, note that existing instances are automatically updated in Cloud-delivered Firewall Management Center; Multicloud Defense creates new instances to support the new license configuration and sync to Cloud-delivered Firewall Management Center for you.

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

Upgrade the Multicloud Defense Gateway

Upgrade Cisco Multicloud Defense gateway instances to a selected software version while maintaining protected cloud traffic.

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

1. Navigate to **Infrastructure** > **Gateways** > **Gateways**.
2. Select the checkbox for the gateway you want to upgrade. You can make only one selection at this time.
3. Select **Actions** > **Upgrade**.
4. From the **Gateway Image** list, select the desired image.
5. Click **Save**.
6. Confirm the cloud service provider resource allocation necessary for the upgrade.
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.

Upgrade Your Firewall Threat Defense Virtual Device

Configure and upgrade your virtual device affiliated with a cloud service provider and a gateway.

This procedure explains how to upgrade the Firewall Threat Defense virtual devices that are affiliated with a cloud service povider and a Multicloud Defense Gateway. We strongly recommend upgrading your FTDv device through the Multicloud Defense Controller.

Procedure

1. Navigate to **Infrastructure** > **Gateway** > **Gateways** and select the gateway associated with your Firewall Threat Defense virtual device.
2. Click **Actions** and select **Upgrade**.
3. Expand the **FTDv Version** drop-down menu and select the version you want to upgrade to. If there are no versions listed, click the refresh icon located to the right of the drop-down menu or confirm with the Cisco Secure Firewall Threat Defense release notes to see if a new version is publicly available.
4. Click **Save**. Multicloud Defense initiates the upgrade and redeploys that gateway.

Abort a Multicloud Defense Gateway

Abort an in-progress Cisco Multicloud Defense gateway update from the gateway management page.

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

Complete these steps to abort an existing Multicloud Defense Gateway:

Procedure

1. Navigate to **Infrastructure** > **Gateways** > **Gateways**.
2. Select the Multicloud Defense Gateway you want to abort in the table so it is highlighted.
3. Expand the **Actions** drop-down menu and click **Abort**.
4. Confirm you want to abort the gateway and click **Yes**. To back out of the action, click **No**.

The current gateway update process is successfully aborted, mitigating issues that may arise during the update procedure.

What to do next
Review system logs for confirmation and assess gateway for any necessary troubleshooting.

Enable a Multicloud Defense Gateway

Enable a disabled Cisco Multicloud Defense gateway after reviewing its validated cloud resource requirements.

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

Procedure

1. Navigate to **Infrastructure** > **Gateways** > **Gateways**.
2. Select the Multicloud Defense Gateway you want to enable in the table so it is highlighted.
3. Expand the **Actions** drop-down menu and click **Enable**.
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

Disable an active Cisco Multicloud Defense gateway and remove associated site-to-site VPN tunnels when required.

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

Complete these steps to disable a Multicloud Defense Gateway:

Before you begin
Ensure gateway is active before starting the disabling process.

Procedure

1. Navigate to **Infrastructure** > **Gateways** > **Gateways**.
2. Select the Multicloud Defense Gateway you want to disable in the table so it is highlighted.
3. Expand the **Actions** drop-down menu and click **Disable**.
4. Confirm you want to disable the gateway and click **Yes**. To cancel this action, click **No**.

The gateway begins the disabling process which may take several minutes.

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

Export a Cisco Multicloud Defense gateway configuration for use with Terraform.

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

Procedure

1. Navigate to **Infrastructure** > **Gateways** > **Gateways**.
2. Select the Multicloud Defense Gateway you want to export in the table so it is highlighted.
3. Expand the **Actions** drop-down menu and click **Export**.
4. Multicloud Defense generates an export wizard.
5. Either click **Download** to download the terraform locally or scroll down and click **Copy Code** to copy the JSON resource.
6. Manually paste into the terraform script.
7. Within the terraform prompt, execute the command provided in the lower half of the window: `terraform import "ciscomcd_gateway"."object-name" <object name>`.

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

Permanently delete a Cisco Multicloud Defense gateway and clean up any associated site-to-site VPN connections.

Distinct from disabling, this task involves permanently removing the gateway from the management portal. Ensure no VPN connections require this gateway.

Complete these steps to delete a Multicloud Defense Gateway.

Procedure

1. Navigate to **Infrastructure** > **Gateways** > **Gateways**.
2. Select the Multicloud Defense Gateway you want to delete in the table so it is highlighted.
3. Expand the **Actions** drop-down menu and click **Delete**.
4. Confirm the action and click **Yes**. To cancel the deletion action, click **Cancel**.

Multicloud Defense Gateway is removed from the management portal permanently.

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.

11 Site-to-Site VPN Tunnel Connection

Topics:

- Prerequisites and Limitations for Site-to-Site VPN Tunnels
- Enable VPN Within the Gateway
- Create a Site-to-Site VPN Connection
- Edit a Site-to-Site VPN Tunnel
- Clone a Site-2-Site VPN Tunnel Connection
- Delete a VPN Tunnel Connection

Learn about site-to-site VPN tunnel connections between Cisco Defense Orchestrator and Multicloud Defense.

A site-to-site VPN tunnel connects networks in different geographic locations. You can create site-to-site IPsec connections between two different Multicloud Defense Gateways or between a Multicloud Defense Gateway and a cloud service provider that complies with all relevant standards. After the VPN connection is established, the hosts behind the local gateway can connect to the hosts behind the remote gateway through the secure VPN tunnel.

Typically, the dynamic peer must be the one that initiates the connection as the other peer would not know the IP address of the dynamic peer. When the remote peer attempts to establish the connection, the other peer validates the connection using the preshared key, IKE settings, and IPsec configurations.

Because the VPN connection is established only after the remote peer initiates the connection, any outbound traffic that matches access control rules that allow traffic in the VPN tunnel will be dropped until that connection is established. This ensures that data does not leave your network without the appropriate encryption and VPN protection.

At this time, Multicloud Defense supports site-to-site VPN tunnel connections with the following platforms or products:

- AWS
- Azure
- GCP
- ASA device
- FTD device
- Extranet or a third party firewall

Prerequisites and Limitations for Site-to-Site VPN Tunnels

Review supported endpoints, prerequisites, and limitations before creating site-to-site VPN tunnel connections with Cisco Multicloud Defense gateways.

Supported VPN Tunnel Connection Endpoints

You can create a VPN tunnel connection with any of the following setups:

- Multicloud Defense Gateway to a Multicloud Defense Gateway.
- Multicloud Defense Gateway to a cloud service provider (AWS, Azure, GCP).
- Multicloud Defense Gateway to a third party firewall.
- Multicloud Defense Gateway to an ASA device hosted in Security Cloud Control.

Multicloud Defense Gateway Prerequisites and Limitations

You must have the following prerequisites completed prior to creating a VPN tunnel regardless of the type of device or platform involved:

- Multicloud Defense Gateway and Multicloud Defense Terraform Provider **must** be running version 24.04 or later.
- At least one cloud service provider or third party device must be already connected to Multicloud Defense.
- Your cloud service provider or third party device must be configured to allow and create VPN tunnel connections. See the service or platform documentation for more information.
- You must have at least **one** IPSec profile. This profile must be attached to the VPN tunnel connection.
- You **must** add the service VPC or VNet to the gateway prior to deploying.
- The VPC and VNet must be deployed **without** Network Address Translation gateway on both sides.
- You **must** create at least one BGP profile. This profile must be attached to the gateway instance associated with the VPN tunnel connection. VPN tunnels can be more effective when paired with a BGP profile as the profile offers additional control over how traffic flows in your networks. See BGP Profile on page 212 for more information.

> **Note**
>
> When you create a BGP profile, the BGP profile must be enabled for traffic and same value to be used in the tunnel as in the BGP Profile

Be aware of the following limitations when creating a VPN tunnel connection:

- The Multicloud Defense Gateway you select **must** be an egress/east-west gateway.
- AWS and Azure gateways must be **8 core** instance type. 2-core and 4-core are not supported at this time.
- Enable AWS inventory for the region the gateway deploys to. Without this enabled, not all traffic flow is passed.
- Site-to-site VPN connections only support up to 10 VPN peers.
- VPCs and VNETs for either AWS or Azure environment must be created with a **single** availability zone. Multiple availability zones are not supported at this time.
- Site-to-site VPN tunnels **do not** support forward-proxy firewall rules at this time.

- Your bandwidth must be at least 800 Mbps.

> **Note**
>
> If you disable or enable a gateway, you **must** delete the site-to-site connection associated with the gateway and recreate the VPN connection.

Limitations for VPN Tunnel Between Multicloud Defense and an ASA Device

Be aware of the following limitations when creating a VPN tunnel connection between the Multicloud Defense Gateway and an ASA device:

- When choosing the endpoints for the VPN tunnel, ensure at least one endpoint is an ASA device and the one endpoint is an Multicloud Defense Gateway (step 4-6).
- If you create a site-to-site VPN tunnel for a third-party or an on-premises device, the table of VPN connections only displays the status of the IPSec profile on Multicloud Defense's endpoint of the connection.
- Autoscaling is not currently supported.

For more information on VPN Tunnels to an ASA device that is hosted in Security Cloud Control, see ASA Site-to-Site VPN Configuration.

> **Note**
>
> If you are using a third-party device or an On-Premises Firewall Management Center, only the Multicloud Defense's side of the IPSEC status is displayed at this time.

Limitations for VPN Tunnel Between Multicloud Defense and an FTD Device

Be aware of the following limitations when creating a VPN tunnel connection between the Multicloud Defense Gateway and an FTD device:

- Both IPsec IKEv1 & IKEv2 protocols are supported.
- Automatic or manual pre-shared keys for authentication.
- IPv4 and IPv6. All combinations of inside and outside are supported.
- IPsec IKEv2 site-to-site VPN topologies provide configuration settings to comply with Security Certifications.
- Static and dynamic interfaces.
- Support for the dynamic IP address for the extranet device as an endpoint.

For more information on VPN Tunnels to an FTD device that is hosted in Security Cloud Control, see Configure Site-to-Site VPN for an FDM-Managed Device.

Enable VPN Within the Gateway

Enable VPN support on a Cisco Multicloud Defense gateway and assign a BGP profile before creating a site-to-site tunnel connection.

Use this procedure to enable the VPN for a gateway in the Multicloud Defense Controller dashboard:

Before you begin

Before you can establish a VPN connection between two devices using Multicloud Defense, you must enable the gateway to utilize a BGP profile. See Prerequisites and Limitations for Site-to-Site VPN Tunnels on page 120 for the complete list of prerequisites for creating and deploying a gateway.

> **Note**
>
> If you opt to use a BGP profile, the BGP profile is run over the IPSEC tunnel with the remote peer.

Procedure

1. Navigate to **Infrastructure** > **Gateways** > **Gateways**.
2. Click **Add Gateway** to create a new gateway or select an existing gateway. Choose **Edit** it in the **Actions** drop-down menu. To configure a gateway, see Configure Your Gateway.

   For Azure VPN, we recommend that you use a Dsv3-Series V3 instance type instead of a Dsv5-Series V5 instance type, as the V5 instance type may not be compatible.
3. When you create or edit the gateway, scroll to the bottom of the window and select a **BGP profile** from the drop-down menu when prompted.
4. Locate the **VPN Connection** options under **Advanced Settings**. Check the **Enable VPN** option to opt into VPN tunnel connection.
5. Expand the **BGP Profile** drop-down menu and select a profile that has already been created.

What to do next

Create a site-to-site tunnel connection.

Create a Site-to-Site VPN Connection

Create a site-to-site VPN tunnel between a Cisco Multicloud Defense gateway and a supported cloud platform, Cisco device, or third-party firewall.

This procedure allows you to create a site-to-site VPN tunnel connection between your gateway and an ASA device, Azure, AWS, and GCP cloud service providers or even a third party firewall of your choice.

Note

When entering the **virtual interface IP address**, we strongly recommend using an IP from the 169.254.xx/16 range, excluding the threat defense reserved range 169.254.1.x/24.

For the net mask, we recommend using /30; this allows you to only use two IP addresses for the endpoints of the virtual tunnel interface connection. For example, 169.254.100.1/30.

Use the following procedure to create an site-to-site VPN tunnel using the Multicloud Defense Controller dashboard:

Before you begin

You must have at least one IPSec profile already created prior to creating a VPN connection tunnel.

We **strongly** recommend creating a BGP profile and adding it to your Multicloud Defense Gateway before you create a VPN tunnel connection. See BGP Profile on page 212 for more information.

Procedure

1. Navigate to **Infrastructure** > **Network** > **VPN Connections**.
2. Click **Create VPN Connection**.
3. Enter a **Name** for the connection.
4. Expand the **Device 1** drop-down menu to select a Multicloud Defense Gateway or manually enter a public IP address of a remote endpoint.
5. Enter the **Device 1 Virtual Interface IP** address. Read the **Note** at the beginning of this procedure for guidance on how to optimize this field.
6. Expand the **Device 2** drop-down menu to select your Multicloud Defense Gateway or manually enter a public IP address of a remote endpoint. Do not use the same device or gateway for both device 1 and device 2.
7. Enter the **Device 2 Virtual Interface IP** address. Read the **Note** at the beginning of this procedure for guidance on how to optimize this field.
8. Enter the **Authentication Value** for the tunnel. At this time, PreShared Key is the preferred authentication method.
9. Expand the **IPSec Profile** drop-down menu to select a profile that has already been created.
10. Click **Save**.

What to do next

View the connection status to review the statistics for incoming and outgoing bytes at both ends of the connection.

If you want to associate a BGP profile with your VPN tunnel connection, create a gateway or edit an existing one and add the desired BGP profile. Note that the IPSec profile of the VPN connection remains the primary profile used and the BGP profile is executed on top of the IPSEC tunnel with the remote peer.

Edit a Site-to-Site VPN Tunnel

Edit the endpoints, virtual interface addresses, authentication value, or IPSec profile for an existing Cisco Multicloud Defense site-to-site VPN tunnel.

Use the following procedure to edit an existing site-to-site VPN connection using the Multicloud Defense Controller dashboard:

Procedure

1. Navigate to **Infrastructure** > **Network** > **VPN Connections**.
2. Select a VPN connection so it is highlighted.
3. In the **Actions** drop-down menu, select **Edit**.
4. Modify any of the following information:
   - Name.
   - Device 1.
   - Device 1 Virtual Interface IP.
   - Device 2.
   - Device 1 Virtual Interface IP.
   - Authentication Value.
   - IPSec profile selection.
5. Click **Save**. You can **Cancel** at any time.

Clone a Site-2-Site VPN Tunnel Connection

Clone an existing site-to-site VPN tunnel connection in Cisco Multicloud Defense and update its endpoints, IP addresses, profile, or authentication value.

Use the following procedure to clone a VPN Tunnel connection using the Multicloud Defense Controller dashboard:

Procedure

1. Navigate to **Infrastructure** > **Network** > **VPN Connections**.
2. Select a VPN connection so it is highlighted.
3. In the **Actions** drop-down menu, select **Clone**.
4. Enter a **Name** for the connection. It must be different from the connection that is being cloned.
5. Modify any of the following information that is cloned:
   - Device 1.
   - Device 1 Virtual Interface IP.
   - Device 2.
   - Device 1 Virtual Interface IP.
   - IPSec Profile selection.
6. The Authentication type is cloned, but the key value for it is not. Enter the **Authentication Value** for the tunnel.
7. Click **Save.**

Delete a VPN Tunnel Connection

Delete a site-to-site VPN tunnel connection from Cisco Multicloud Defense and remove related BGP profiles when they are no longer needed.

Use the following procedure to delete a VPN Tunnel connection using the Multicloud Defense Controller dashboard:

Procedure

1. Navigate to **Infrastructure** > **Network** > **VPN Connections**.
2. Select a VPN connection so it is highlighted.
3. In the **Actions** drop-down menu, select **Delete**.
4. Confirm the deletion action and click **Delete**.

What to do next

We strongly recommend deleting any BGP profiles that you created for the VPN tunnel you just deleted.

12 Configure a Virtual Network Connection for Azure Virtual WAN

Topics:

- Overview of Virtual WAN
- Guidelines for Virtual WAN Connections to Azure vHub
- Create a Service VPC with a Virtual WAN Attachment
- Modify a Service VPC with a Virtual WAN Attachment

Configure a virtual network connection between Multicloud Defense and Azure Virtual WAN.

Overview of Virtual WAN

Learn how Cisco Multicloud Defense integrates with Azure Virtual WAN to protect egress traffic between service VNets, spoke VNets, and a virtual hub.

If you use Azure cloud services, you can create a Virtual WAN (VWAN) to orchestrate and simplify network connectivity between your on-premises networks, branch offices, and remote users. You can integrate Multicloud Defense with Azure VWAN by orchestrating virtual network connections and route propagation between Service VNets and a virtual hub (vHub).

Typically, within a vHub, Multicloud Defense is not supported as a Network Virtual Appliance (NVA). Instead, you can use VWAN route orchestration to address this. To use Multicloud Defense to protect your application on Azure, you can orchestrate the creation of a virtual network connection from a Service VNet to a vHub inside a VWAN. You can propagate routes between the vHub and Multicloud Defense. Multicloud Defense supports VWAN for egress mode only. Note that Azure VWAN is supported only for Multicloud Defense gateways. Azure does not support east-west traffic for VWAN due to challenges in routing traffic between secured hubs.

Guidelines for Virtual WAN Connections to Azure vHub

Review the prerequisites and limitations for connecting a Cisco Multicloud Defense service VNet to an Azure Virtual WAN vHub.

Prerequisites

- An Azure subscription with VWAN and vHub must be configured.
- Service VNet and spoke VNets must be set up in Azure.
- Multicloud Defense gateway must be deployed in the service VNet.
- Permissions must be available to create and manage virtual network connections and route tables in Azure.
- Permissions must be available to enable and disable vHub connections.

Limitations

- Multicloud Defense is not supported as an NVA within a vHub.
- Classless Inter Domain Routing (CIDR) selection is only available during the edit phase, not during VNet creation.
- Route propagation is dependent on the configuration of the egress or ingress gateway.

Create a Service VPC with a Virtual WAN Attachment

Create a Cisco Multicloud Defense service VNet and attach it to an Azure Virtual WAN vHub with the required route tables.

You can create a service VPC with VWAN attachment using the **Easy Setup** wizard when you secure your account. For details, see Centralized Model: Add a VPC or VNet.

Use the following procedure to create a service VPC and attach a VWAN.

Procedure

1. From the Multicloud Defense Controller, navigate to **Infrastructure** > **Gateways** > **VPCs/VNets**.
2. Click **Create Service VPC/VNet** to create a service VPC.

3. Enter a **Name**.
4. From the **Region** drop-down list, choose a region.
5. From the **CSP Account** drop-down list, choose an account.
6. Enter details for **CIDR Block**.
7. From the **Availability Zones** drop-down list, choose a zone.
8. From the **Resource Group** drop-down list, choose a resource group.
9. Check the **Use NAT Gateway** check box to direct traffic through the NAT gateway.
10. In the **vWAN Attachment** section, set the toggle to **Enabled**.
11. From the **vHub** drop-down list, choose a hub.
12. From the **Associate Route Table** drop-down list, select a route table to associate.
13. From the **Propagate Route Tables** drop-down list, select route tables to propagate.
14. Click **Save**.

    A service VPC is created with a vHub connection to an Azure VWAN. You can also view the configuration changes made, on your Azure account.

> **Note**
>
> If you delete a service VPC in Multicloud Defense, the vHub connection between the VWAN and the Azure service VPC is also deleted.

Modify a Service VPC with a Virtual WAN Attachment

Modify an existing Cisco Multicloud Defense service VNet to connect it to an Azure Virtual WAN vHub and manage route propagation.

Procedure

1. From the Multicloud Defense Controller, navigate to **Infrastructure** > **Gateways** > **VPCs/VNets**.
2. Choose a Service VPC that you want to edit, from the list.
3. Click **Edit**.
4. In the **vWAN Attachment** section, set the toggle to **Enabled**.
5. From the **vHub** drop-down list, choose a hub.
6. From the **Associated Route Table** drop-down list, choose a route table to associate.
7. From the **Propagate Route Tables** drop-down list, choose route tables to propagate.
8. To propagate all spoke CIDRs to the vHub, set the toggle to **Always**.

> **Note**
>
> To add multiple spoke VPCs to the route tables, use the list builder to move the spoke VPCs from the **Available** section to the **Selected** section. Moving the VPC to the **Selected** section ensures that the VPC is added.

9. Click **Save**.

   The service VPC connects to a VWAN containing a vHub and spoke VPCs in Azure. Note that any changes made to route tables are also updated in Azure.

13 Rules and Rule Sets

Topics:

- Policy Management
- Rules
- Rule Sets and Rule Set Groups

Learn how rules and rule sets organize security policy configuration in Multicloud Defense.

Policy Management

Learn how Cisco Multicloud Defense keeps security policies current across gateways and supports consistent protection in multicloud environments.

Policies are created in the Multicloud Defense dashboard or through orchestration using the Multicloud Defense Terraform provider. The policies are stored and retained as part of the Multicloud Defense Controller database. The gateway retrieves the policy or any policy changes through a periodic heartbeat where the gateway provides the controller health and telemetry information, while also requesting if there are any policy changes that need to be applied. The gateway to controller communication is fully encrypted and established through a mutual TLS session. The heartbeats occur every five seconds to ensure that policies on the gateway are synchronized with the policies created or modified by the user.

Policy Rule Set Gateway and Management

Learn how to change the policy rule set assigned to an active Cisco Multicloud Defense gateway without interrupting protected traffic.

Policy Rule Management

A policy rule set assigned to a gateway can be changed dynamically to a different policy rule set. If there is a requirement to swap in a different policy rule set to an active gateway, this operation can be initiated in a non-impactful way. The assignment of the new policy rule set operates similarly to a gateway update/upgrade process. New gateway instances are instantiated with the new policy rule set. New traffic sessions are redirected to the new gateway instances once they are active and healthy. Old traffic sessions are flushed from the old gateway instances. The old gateway instances are deleted. The operation completes in a matter of minutes. This change is initiated as part of the gateway configuration settings. Navigate to **Infrastructure** > **Gateways** > **Gateways**. The change can be initiated using the Multicloud Defense portal or the Multicloud Defense Terraform Provider.

Policy Rule Set Gateway Status

The status of the connection between the policy rule and the gateway it is associated with can be one of the two options:

- **Updated** - The policy is active on the gateway and is synchronized with the controller.
- **Updating** - The gateway is actively processing a policy change. The policy change is known to the gateway, but is not yet active. The gateway is still process traffic using the current policy.

Rules

Consider these guidelines when creating rules for gateways in different environments.

In general, rules specify the rights of a user, group, role, or organization to access objects of a specified type and state within a domain. Multicloud Defense supports a variety of cloud service providers and each of these environments have their own requirements or methods for their rules. The Multicloud Defense Controller might handle rules created in your cloud account differently than those created within it. Some rules are applied to gateways and instances by default so the environments have a basic level of protection as you continue to add and modify the rules and policies for optimal performance and coverage.

Rule **types** are important when considering the type of gateway environment you are catering to. Not all rules or rule types are completely compatible with every gateway environment. Gateway types supported in Multicloud Defense Controller are ingress, egress, and east-west.

For information about rules and rule sets, or how to create or modify rules and rule sets for policies and groups, read the rest of this chapter.

Rule Sets and Rule Set Groups

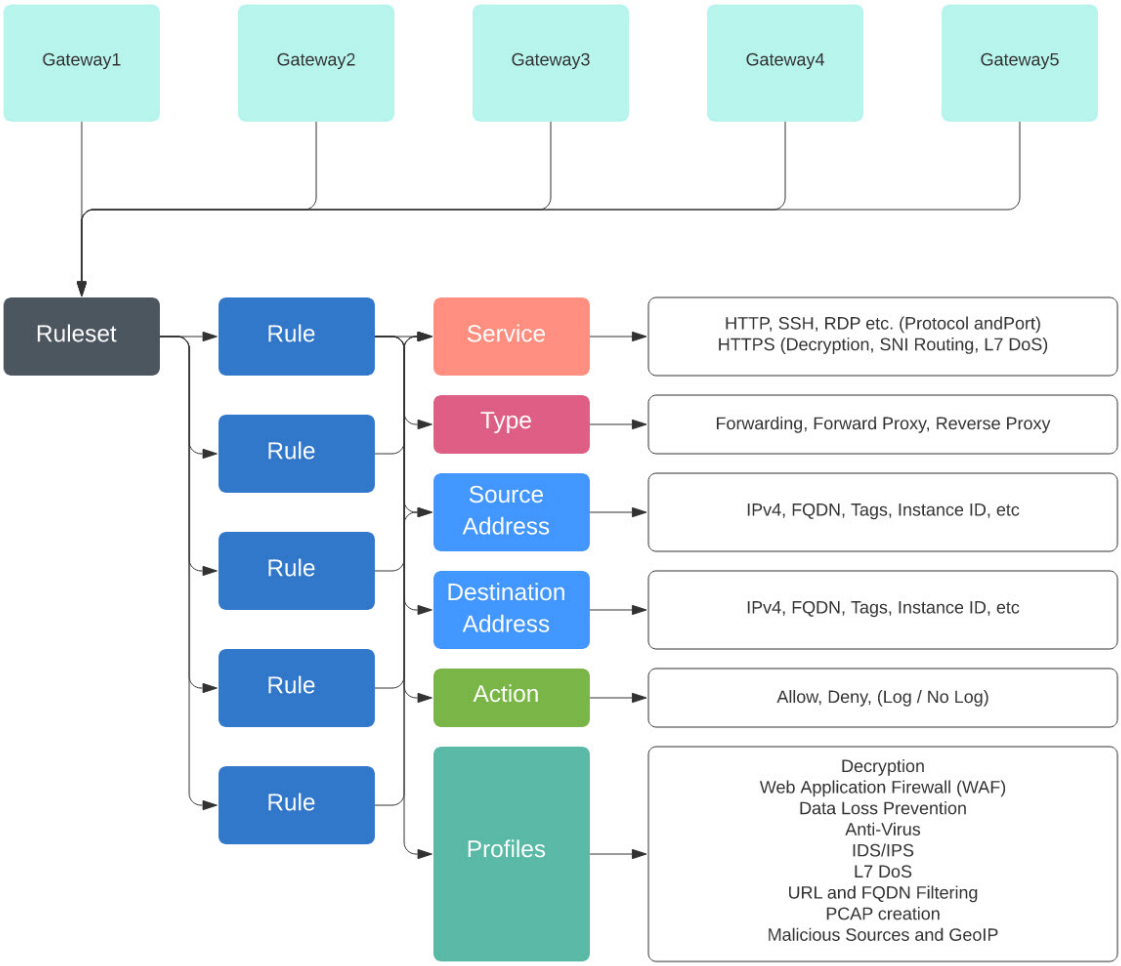
An overview of Cisco Multicloud Defense rule sets and rule set groups used to organize and apply security policies across gateways.

Rule Sets

Rule sets consist of a set of rules that define a segmentation and advanced security policy that are applied to a set of one or more gateways to accommodate application and workload protection. The rules are organized as a priority list where traffic is processed by a matched rule, a general action is taken to allow or deny, and further inspection is accommodated through advanced security.

Rule sets must be associated with at least one Multicloud Defense Gateway. The following limitations apply to all rule sets:

- Rule sets are cloud agnostic and can be applied to one or more gateways operating across multiple cloud environment.
- A gateway can only be associated with a single rule set, although more than one rule set can be applied using a rule set group.
- Rules within a rule set can use discovered cloud asset information to form a dynamic policy, or a policy that adapts in real time to changes.
- A rule set can include rules that only apply to specific cloud accounts and/or cloud regions, although the rule set is applied to gateways that cross cloud environments. Here is an example:
  - A dynamic tag-based address object used in a rule within a rule set that is applied to two gateways across two clouds can resolve to a set of IP addresses that are associated with a gateway in one cloud, while resolving to a different set of IP addresses that are associated with a gateway in another cloud.
- Rule sets can be created from the **Policies** > **Security Policies** > **Rule Sets** page or from within the gateway creation workflow. The following diagram is of a a single rule set applied to multiple gateways:



Another supported use case is of multiple rule sets associated with multiple gateways.

Policy Rule Set Groups

A policy rule set group is a collection of standalone rule sets. Users can combine multiple standalone rule sets into a policy rule set group and associate the group to one or more Multicloud Defense Gateways. Policy rule set groups allow organizations to separate policies in an organized fashion and combine them to an overarching policy.

Note

- A policy rule set group can only consist of rule set members.
- Ensure all rule sets associated with a policy rule set group do not have conflicting rules.
- A policy rule set group can have a maximum of 100 rule set members.

Create Policy Rule Set

Create a Cisco Multicloud Defense policy rule set and add standalone security rules for gateway protection.

To create a policy rule set:

Procedure

1. Navigate to **Policies** > **Security Policies** > **Rule Sets**.
2. Click **Create**.
3. Add a name and description for the policy rule set.
4. Click **Save.**

What to do next

Once the policy rule set is created, add standalone rules to the rule set.

Create a Rule in a Rule Set

Describes how rule sets are structured and used for managing rules effectively.

-

Add or Edit a Forwarding Rule in a Rule Set

Add or modify a forwarding rule within a defined rule set.

Use the following procedure to add existing rules to a policy rule set or to edit rules that are already included in a policy rule set:

Before you begin

You can create a new rule within the Multicloud Defense Gateway. Note the following limitations before you add or edit rules to your rule set:

- A single policy rule set can have a maximum of 2047 rules.
- A policy rule set group can have a maximum combined set of 2047 rules.

Procedure

1. Navigate to **Policies** > **Security Policies** > **Rule Sets**.
2. Click the policy rule set name to view the policy rule set.
3. Click **Add Rule** to create a new rule or add an existing rule. This generates a prompt.
4. Enter the following properties:
   - **Name** - a unique name used to reference the rule.
   - (optional) **Description** - A brief description of the rule.
   - **Type** - Select **Forwarding**.
5. Enter the following Object information:
   - **Service** - The service object used to determine the protocols and ports for which the rule will apply.
   - **Source** - The address object used to determine the resources for which the rule will apply.
   - **Destination** - The address object used to determine the destination resources for which the rule will apply. For a **ReverseProxy** rule type, the destination is always the Multicloud Defense Gateway. For **ForwardProxy** rule types, the destination is always any.

- **FQDN** - Use the drop-down menu to select a set of FQDNs used for SNI match. Note this applies only to **Forwarding** rule types.

6. Enter the Details:
   - **Action** - The action defines whether the traffic should be allowed or denied, and whether the traffic should be logged or not logged in events. Traffic is always logged in traffic summary, no matter whether the action is set to **Log** or **No Log**. The advanced security profiles evaluate traffic allowed by the rule. Note that each advanced security profile has its own action that will either use or override this action.
   - **Reset On Deny** - If enabled, the Multicloud Defense Gateway will send a TCP Reset packet for the sessions that match this policy and are dropped by the gateway. Note this only applies to **Forwarding** rule types.
7. Enter the following Profiles information:
   - (Optional) **Network Intrusion** - The Network Intrusion (IPS) profile to be used for advanced security.
   - (Optional) **Anti-malware** - The Anti-malware profile to be used for advanced security. If you do not already have an Anti-malware profile created, click **+ Create Anti Malware** here.
   - (Optional) **Data Loss Prevention** - The Data Loss Prevention (DLP) profile to be used for advanced security. Note that this applies only to **ForwardProxy** rule types.
   - (Optional) **FQDN Filtering** - The FQDN Filtering (FQDN) profile to be used for advanced security.
   - (Optional) **Malicious IPs** - The Malicious IPs (MIP) profile to be used for advanced security.
   - (Optional) **PCAP** - Check this box to enable. Whether packet capture is enabled or disabled for the rule. Whenever traffic matches a rule with PCAP enabled, a packet capture of the session traffic will occur and the PCAP will be stored in the location specified by the PCAP profile. The PCAP profile is configured on the Multicloud Defense Gateway.
8. After specifying the configuration for the rule, click **Save**.
9. Continue adding more rules. Once all desired rules have been added, click **Save Changes**. You will be presented with a before and after view of all changes made to the rule set. If satisfied with your changes, click **Save**. If you need to make further changes, click **Cancel** to return to editing your rule set.

Add or Edit a Reverse Proxy Rule in a Rule Set

Configure and modify reverse proxy rules to manage network traffic securely.

Use the following procedure to add existing rules to a policy rule set or to edit rules that are already included in a policy rule set:

Before you begin

You can create a new rule within the Multicloud Defense Gateway. Note the following limitations before you add or edit rules to your rule set:

- A single policy rule set can have a maximum of 2047 rules.
- A policy rule set group can have a maximum combined set of 2047 rules.

Procedure

---

1. Navigate to **Policies** > **Security Policies** > **Rule Sets**.
2. Click the policy rule set name to view the policy rule set.

3. Click **Add Rule** to create a new rule or add an existing rule. This generates a prompt.
4. Enter the following properties:
   - **Name** - a unique name used to reference the rule.
   - (optional) **Description** - A brief description of the rule.
   - **Type** - Select **ReverseProxy**.
5. Enter the following Object information:
   - **Service** - The service object used to determine the protocols and ports for which the rule will apply.
   - **Source** - The address object used to determine the resources for which the rule will apply.
   - **Destination** - The address object used to determine the destination resources for which the rule will apply. For a **ReverseProxy** rule type, the destination is always the Multicloud Defense Gateway.
   - **Target** - The address object used to specify the destination for which the Multicloud Defense Gateway will establish a gateway to server connection.
6. Select the preferred rule **Action**. This defines whether the traffic should be allowed or denied, and whether the traffic should be logged or not logged in events. Traffic is always logged in traffic summary, no matter whether the action is set to **Log** or **No Log**. For traffic that is allowed by the rule, the advanced security profiles are evaluated. Note that each advanced security profile has its own action that will either use or override this action.
7. Enter the following Profiles information:
   - (Optional) **Network Intrusion** - The Network Intrusion (IPS) profile to be used for advanced security.
   - (Optional) **Anti-malware** - The Anti-malware profile to be used for advanced security. If you do not already have an Anti-malwre profile created, click **+ Create Anti Malware** here.
   - (Optional) **Web Protection** - The Web Protection (WAF) profile to be used for advanced security. Note that this applies only to **ReverseProxy** rule types.
   - (Optional) **URL Filtering** - The URL Filtering (URL) profile to be used for advanced security. Note that this applies only to **ForwardProxy** and **ReverseProxy** rule types.
   - (Optional) **Malicious IPs** - The Malicious IPs (MIP) profile to be used for advanced security.
   - (Optional) **PCAP** - Check this box to enable. Whether packet capture is enabled or disabled for the rule. Whenever traffic matches a rule with PCAP enabled, a packet capture of the session traffic will occur and the PCAP will be stored in the location specified by the PCAP profile. The PCAP profile is configured on the Multicloud Defense Gateway.
8. After specifying the configuration for the rule, click **Save**.
9. Continue adding more rules. Once all desired rules have been added, click **Save Changes**. You will be presented with a before and after view of all changes made to the rule set. If satisfied with your changes, click **Save**. If you need to make further changes, click **Cancel** to return to editing your rule set.

Add or Edit a Forward Proxy Rule in a Rule Set

Add or edit a forward proxy rule in a Cisco Multicloud Defense policy rule set to define traffic handling and security inspection.

Use the following procedure to add existing rules to a policy rule set or to edit rules that are already included in a policy rule set:

Before you begin

You can create a new rule within the Multicloud Defense Gateway. Note the following limitations before you add or edit rules to your rule set:

- A single policy rule set can have a maximum of 2047 rules.
- A policy rule set group can have a maximum combined set of 2047 rules.

Procedure

1. Navigate to **Policies** > **Security Policies** > **Rule Sets**.
2. Click the policy rule set name to view the policy rule set.
3. Click **Add Rule** to create a new rule or add an existing rule. This generates a prompt.
4. Enter the following properties:
   - **Name** - a unique name used to reference the rule.
   - (optional) **Description** - A brief description of the rule.
   - **Type** - Select **ForwardProxy**.
5. Enter the following Object information:
   - **Service** - The service object used to determine the protocols and ports for which the rule will apply.
   - **Source** - The address object used to determine the resources for which the rule will apply.
   - **Destination** - The address object used to determine the destination resources for which the rule will apply. For **ForwardProxy** rule types, the destination is always any.
   - **FQDN** - Use the drop-down menu to select a set of FQDNs used for SNI match. Note this applies only to **Forwarding** rule types.
6. Enter the preffered rule **Action**. This defines whether the traffic should be allowed or denied, and whether the traffic should be logged or not logged in events. Traffic is always logged in traffic summary, no matter whether the action is set to **Log** or **No Log**. For traffic that is allowed by the rule, the advanced security profiles are evaluated. Note that each advanced security profile has its own action that will either use or override this action.:
7. Enter the following Profiles information:
   - (Optional) **Network Intrusion** - The Network Intrusion (IPS) profile to be used for advanced security.
   - (Optional) **Anti-malware** - The Anti-malware profile to be used for advanced security. If you do not already have an Anti-malwre profile created, click **+ Create Anti Malware** here.
   - (Optional) **Data Loss Prevention** - The Data Loss Prevention (DLP) profile to be used for advanced security. Note that this applies only to **ForwardProxy** rule types.
   - (Optional) **URL Filtering** - The URL Filtering (URL) profile to be used for advanced security. Note that this applies only to **ForwardProxy** and **ReverseProxy** rule types.
   - (Optional) **FQDN Filtering** - The FQDN Filtering (FQDN) profile to be used for advanced security.
   - (Optional) **Malicious IPs** - The Malicious IPs (MIP) profile to be used for advanced security.
   - (Optional) **PCAP** - Check this box to enable. Whether packet capture is enabled or disabled for the rule. Whenever traffic matches a rule with PCAP enabled, a packet capture of the session traffic will occur and the PCAP will be stored in the location specified by the PCAP profile. The PCAP profile is configured on the Multicloud Defense Gateway.

8. After specifying the configuration for the rule, click **Save**.
9. Continue adding more rules. Once all desired rules have been added, click **Save Changes**. You will be presented with a before and after view of all changes made to the rule set. If satisfied with your changes, click **Save**. If you need to make further changes, click **Cancel** to return to editing your rule set.

Disable, Edit, Clone, or Delete Rules in a Rule Set

Disable, edit, clone, or delete rules within a Cisco Multicloud Defense policy rule set.

Use the following procedure to edit or clone an existing rule that is configured for a rule set. You can also disable a rule if you do not need it active for your currenty policies or rule set. You can delete a rule if you do not need it now or for any future deployment.

Note that you can only edit or clone one rule at a time. You can disable or delete multiple rules simultaneously.

Procedure

---

1. Navigate to **Policies** > **Security Policies** > **Rule Sets**.
2. Locate the rule set that contains the rule you want to disable, edit, clone, or delete and click the rule set name.
3. Check the checkbox of the standalone rule.
4. Expand the **Actions** button.
5. Select your actionable item:
   - **Disable** - This option keeps the rule in the rule set but disables the rule and the configured rule action from affecting traffic.
   - **Edit** - This option launches the Properties window and allows you to edit the configuration of the rule. Click **Save** to keep the changes you made.
   - **Clone** - This option creates a duplicate of the rule and opens the Properties window for you to name the cloned rule, or make any additional changes to the rule's configuration. Click **Save** to confirm the configuration. Saving a cloned rule automatically adds it to the rule set you are viewing.
   - **Delete** - This option permanently removes the rule from the rule set. Note that this also removes the rule from the gateway.
6. Click **Save Changes** to confirm the changes you made to the rule and, indirectly, do the rule set. If you do not want to save the changes, click **Cancel**. Confirm that losing any changes made to the gateway is OK.

Create a Policy Rule Set Group

Create a Cisco Multicloud Defense policy rule set group that combines multiple rule sets for gateway assignment.

To create a policy rule set group:

Procedure

---

1. Navigate to **Policies** > **Security Policies** > **Rule Sets**.
2. Click **Create**.
3. Add a name and description for the policy rule set group.
4. Select **Type** as the group.

5. Expand the drop-down menu to add rule sets in the **Rule Set List** section. If you want to add more rule sets, click **Add Rule Sets** to add another row.

14 Objects

Topics:

- About the Multicloud Defense Connector
- Import Objects From Security Cloud Control

Learn about shared objects that you can use across Multicloud Defense security policies.

In an environment where you may have cloud-based managers such as AWS or GCP interacting with on-premises datacenters, it is crucial to be able to share objects within policies to protect your environment. Shared objects make it easy to maintain policies because you can modify an object in one place and that change affects all other policies that use that object. Without shared objects, you would need to modify all the policies individually that require the same change.

Multicloud Defense shows you a combined or "flattened" view of the elements of the object in the details pane. Notice that in the details pane, the network elements are flattened into a simple list and not directly associated with a named object.

Note that sharing objects is only supported when you deploy an access control policy that allows traffic from your cloud-based datacenter. Ensure that your policy includes, or excludes, instances or attributes from your third-party datacenter.

Multicloud Defense has the capability to communicate with either a datacenter or a cloud platform, ensuring your policies for security can be managed anywhere.

Static Objects

Static objects specifies unchanging IP addresses, subnets, or specific firewall rules to provide predictable and stable configurations which can be important for compliance and security purposes. In a cloud environment, this allows you to create and share objects that maintain the same IP address or FQDN within a hybrid environment.

If you choose to delete a shared object, Multicloud Defense deletes it only from its system. The object continues to exist within Security Cloud Control.

Dynamic Objects

In contrast, dynamic objects do not have to specify an IP address at all. Dynamic objects are adaptable configurations that automatically adjust to varying conditions or environments. They allow firewalls to respond to real-time events without requiring manual intervention.

You can also **tag** resources and use them as objects to create a more fine-tuned ruleset within your policy. This level of fleibiltity within a cloud environment allows the system to adjust for yoou based on real-time data and can result on reduced maintenance.

Sharing Objects with Security Cloud Control

When you share objects with Security Cloud Control they are automatically translated into **network objects**. This does not affect the original state of the object in Multicloud Defense.

To configure object sharing between Multicloud Defense and Security Cloud Control you **must** create a connector in Security Cloud Control and attach the connector to an applicable policy to enable this feature and then import objects

to see them in the Multicloud Defense Controller. See About the Multicloud Defense Connector on page 143 for more information.

If you happen to share dynamic objects there is the option to preserve the original values of the object by creating an override value. An object override allows you to override the value of a shared network object on specific devices. See Object Overrides for more information.

Note

Objects cannot be shared with .

About the Multicloud Defense Connector

Learn how the Multicloud Defense Connector lets you share address objects with cloud-delivered Firewall Management Center for use in access control policies.

You can optionally send address objects from Cisco Multicloud Defense to the configured Cloud-Delivered Firewall Management Center using a Dynamic Attributes Connector. A connector is responsible for gathering dynamic data (such as IP addresses) and streaming them to the Cloud-Delivered Firewall Management Center so they can be used in access control policies.

For more information about Multicloud Defense objects, see the Address Objects chapter and address object API documentation.

For more information about the Multicloud Defense Connector, see the *Managing Firewall Threat Defense with Cloud-delivered Firewall Management Center in Security Cloud Control*.

Import Objects From Security Cloud Control

Import Cisco Defense Orchestrator address objects into Cisco Multicloud Defense from the Controller dashboard, and confirm that the shared object count updates.

> **Note**
>
> You do not have to enable dynamic sharing in the Security Cloud Control dashboard to import objects to Multicloud Defense.

Use the following procedure to manually import Security Cloud Control objects into Multicloud Defense using the Multicloud Defense Controller dashboard:

Procedure

1. Log into Security Cloud Control and in the left pane, click Multicloud Defense.
2. Click **Multicloud Defense Controller** located in the upper right to cross-launch into the Controller dashboard.
3. Navigate to **Policies** > **Security Policies** > **Addresses**.
4. Click **Import Objects**.
5. From the pop-up window of Security Cloud Control objects, scroll or use the search bar to locate an individual object.

   Note: Objects with names that contain "**.**" are not supported by Multicloud Defense at this time. Attempting to share or import objects with periods in their name results in an error message.
6. Select the object so it is highlighted and click **Import**. You can click **Cancel** at any point to back out of the action.

What to do next

Allow a few minutes for Multicloud Defense to communicate with Security Cloud Control and synchronize the object you imported. From the Security Cloud Control dashboard you will be able to see an updated shared object count in the "Multicloud Defense Shared Object" widget.

15 Address Objects

Topics:

- Address Objects
- Create a Source/Destination Address Object
- Create a Reverse Proxy Target Address Object
- Edit Address Objects
- Clone Address Objects
- Delete Address Object
- View Details

Learn how to use address objects in Multicloud Defense when defining network destinations and sources for security policies.

Address Objects

An overview of Cisco Multicloud Defense address objects that represent IP addresses, networks, domains, and discovered cloud resources in security policies.

An **Address object** represents a set of one or more IPs, CIDRs or FQDNs for use as a **source** or **destination** in a **security policy rule set rule**, or as a **target backend address** in a **reverse proxy service object**, depending on how it is defined. The address object can be configured statically using traditional constructs or dynamically using cloud constructs.

An address object represents a set of one or more IPs, CIDRs or FQDNs within a **Source**, **Destination**, or **Reverse Proxy Target** field within a security policy rule or rule set. You can also define it as a target backend address within a reverse proxy service object. This section focuses on source and destination objects.

As of Version 24.04 and later, you can configure an address object to **exclude** specific IP addresses or an IP address range.

Currently, IPv6 is not supported when you configure Address Groups for AWS, Azure, and GCP environments on Multicloud Defense, as the controller inhibits the discovery of IPv6 addresses in address objects. Do not use IPv6 addresses in Address Groups.

Src/Dest

Learn how source and destination address objects identify traffic by IP address, network, domain, or cloud resource in security policy rules.

These objects are used to define match criteria that maps explicitly to IP addresses or CIDRs. The objects are referenced inside a policy rule and are evaluated against traffic entering a gateway instance when a policy rule is processed.

Source and destination address objects are useful when IP Addresses and CIDRs are explicitly needed to match application traffic entering a gateway instance. These objects are referenced inside the source and destination fields of a policy rule definition. The type of address object used to populate each of these fields depends on the traffic flow, application type, and use-case.

Source or Destination Address Objects

A source or destination address object specifies a source or destination for a rule inside a security policy rule set. It is used by the rule to match traffic based on its source or destination IP address. The different types of address objects are defined as follows:

IP/CIDR/FQDN (Static) Address Objects

An IP/CIDR/FQDN address object is configured as a set of IP addresses, CIDR blocks or FQDNs. Examples of IP/CIDR address objects include:

- Destination IPs for DNS servers.
- Destination IPs for SMTP Relay Servers.
- Destination IPs for NTP servers.
- Source IPs or subnets for application workloads.

FQDN address objects define an explicit set of FQDNs for allowing or blocking IPs based on DNS resolution. When you define an FQDN inside an FQDN address object and reference it in a policy rule, the gateway instance performs a DNS resolution to retrieve the corresponding IP address(es) to match incoming traffic. By default, caching is not enabled. In this case, the DNS resolution is done every 60 seconds, and the gateway instance uses the retrieved resolution for 60 seconds. If the FQDNs specified inside the FQDN address object are resolving to a large set of IP addresses (i.e. more

than 400 each), then caching can be enabled. In this case, the DNS resolution interval can be specified, along with the cache size and cache TTL.

FQDN address objects are useful to match on application traffic that is either UDP based (ex. NTP) or TCP traffic for which host information does not exist in the request packet (ex. SMTP). In either case, it is recommended to use an FQDN address object to match on this kind of application traffic instead of manually defining a list of IP addresses for all appropriate NTP servers or SMTP servers, for example, your internal workloads are required to connect to.

Dynamic Cloud Constructs

Learn how dynamic address objects represent discovered cloud resources and user-defined tags across supported cloud environments.

Cloud-Native address objects are dynamic cloud resources discovered by the Multicloud Defense Controller through either periodic inventory collection (API-based) or real-time event tracking (GCP Pub/Sub integration). These resources can be individual resources such as VPCs/VNETs, Instance IDs, security groups, Subnet IDs or a set of resources referenced through user-defined tags. The Multicloud Defense Controller uses a combination of real-time event tracking and targeted API calls to dynamically populate the IP addresses associated with the cloud resource. Therefore, any subsequent changes made to a cloud-native resource with be automatically reflected inside the address object referencing this resource.

> **Note**
>
> Using cloud-native constructs to define source or destination address objects allows you to create a truly dynamic cloud policy across both single and multi-cloud environments. As cloud resources are added, deleted, or changed within a cloud environment, the address objects are dynamically updated to reflect these changes, making sure your security posture is automatically updated across all applications and functions in your environment.

User-Defined Tags in VNet and VPC Environments

Tags map the IP addresses or CIDR for a cloud resource defined with a set of tags to an address object. In GCP, labels are key-value pairs that are often used to categorize resources dedicated to different environments (i.e., development, staging, production, etc.). Inside a source or destination address object, user-defined tags can be used to reference resources including instances, VPCs/VNETs, subnets, and security groups. Most commonly, organizations use tags to categorize instances.

Tag based policy rules are a very powerful component of dynamic cloud policies. Granular policy rules can be defined for groups of instances with specific tags. With these policy rules in place, anytime a new instance is deployed with the appropriate tags, it automatically inherits the desired security policy defined for the category of instances it belongs to. This is because the Multicloud Defense Controller does not only discover a new instance has been deployed, but also the tags that have been assigned to that instance. It will then dynamically update the source or destination address object referencing this instance-based tags with the new instance's IP address. If an instance is deployed with the incorrect tags or no tags, it will not be allowed to communicate to any other resources because the appropriate policy rule is not matched against.

In VNets and VPCs, tags map the CIDR associated with the VPC to an address object CIDR. Provides a contextual way of creating a rule that matches any instance deployed within a VPC or VNET. It can use the name of a discovered VPC or VNET to define match criteria instead of having to manually figure out what CIDR is associated with a particular VPC or VNET. Any changes to the VPC or VNET will be dynamically updated in the policy rule with no intervention. If a VPC or VNET is removed and a new VPC/VNET is created in its place, the rule will no longer apply even if reusing the CIDR.

Instance ID

Instance IDs map the IP addresses associated with an instance to a list of IP addresses inside an address object. This provides a contextual way of creating a policy rule for a specific instance without manually figuring out how the instance

is configured. The policy rule reflects any changes to the instance or its removal. Note that the policy rule cannot apply to any other instance, even if the instance is deleted and replaced with a new instance with the same configuration.

Security Group

Security Groups map the IP addresses of network interfaces associated with a security group to a list of IP addresses inside an address object. Any interface related changes, such as fields that are added or removed to the security group, are dynamically reflected in the list of IP addresses inside the address object. This provides an organization with the ability to align existing security groups with the advanced security capabilities of the gateway data path pipeline.

Subnet IDs

Subnet IDs map the CIDR associated with a subnet to an address object CIDR. This provides a contextual way of creating a policy rule for all resources associated with a specific subnet ID without manually figuring out how the subnet is configured. A VPC or VNET is typically divided into multiple subnets and resources deployed within these subnets may serve different purposes. For example, instances in one subnet may require a specific set of advanced security profiles or may have a different traffic flow requirement. To simplify the process of creating different security rules for each subnet, Multicloud Defense gives you the capability to define a policy rule using the subnet's name as match criteria. Therefore, each subnet can have a unique policy rule, with unique security profiles. Any changes to the subnet and any instance deployed within the subnet is dynamically reflected in the policy rule.

Geo IP

Learn how Geo IP address objects help you allow or block cloud traffic associated with selected countries.

These objects are used to allow or block traffic that is coming from or going to IP addresses based on their geographic location (country). Multicloud Defense integrates with the MaxMind GeoIP2 Database for maintaining a list of updated GeoIPs.

To review a full list of country names and codes, or IP address to GeoIP country codes, go to the GeoNames website.

Group

Learn how address object groups combine source or destination objects to simplify reusable traffic criteria in security policies.

A group address object is configured as a set of sour or destination address objects. A group provides flexibility by defining individual address objects and then grouping them together, simplifying the number of rules necessary to match traffic based on the members of the group. The group inherits the set of IPs, CIDRs, or FQDNs from the members of the group, whether the members are static, dynamic, or a combination of the two.

Source or Destination Address Object Parameters

Find the required and optional parameters for static and dynamic Cisco Multicloud Defense source or destination address objects.

Type	Mode: Dynamic or Static	Parameter	Required or Optional	Notes
IP/CIDR/FQDN	Static	Value	Required	The total number of FQDNs per Address Object is limited to 200 where each FQDN can resolve to at most 400 IPs. The Multicloud Defense Gateway will perform DNS resolution every 60 seconds, regardless of the DNS record TTL.
VPC/VNet ID	Dynamic	CSP Account	Required	
		Region	Required	
		Resource Group	Optional	Azure Only
		VPC/VNet ID	Required	
Security Group	Dynamic	CSP Account	Required	
		Region	Required	
		VPC/VNet ID	Required	
		Resource Group	Optional	Azure Only
		Security Group ID	Required	
Application Security Group	Dynamic	CSP Account	Required	Azure Only
		Region	Required	
		Resource Group	Required	
		Application Security Group	Required	
Instance ID	Dynamic	CSP Account	Required	
		Region	Required	
		VPC/VNet ID	Required	
		Resource Group	Optional	Optional
		Instance ID	Required	

Type	Mode: Dynamic or Static	Parameter	Required or Optional	Notes
Subnet ID	Dynamic	CSP Account	Required	
		Region	Required	
		VPC/VNet ID	Required	
		Resource Group	Optional	Azure Only
		Subnet ID	Required	
User Defined Tag	Dynamic	CSP Account	Optional	
		Region	Optional	
		VPC/VNet ID	Optional	
		Resource Group	Optional	Azure Only
		Resource/Tag/Value	Required	List of Resources and Tag Key-Value Pairs.Resources can be Instance, VPC/VNet, Subnet, Load Balancer, Security Group, Security Group (Azure).
Geo IP		Value	Required	
Group		Address	Required	

Reverse Proxy Target Address Object

Learn how reverse proxy target address objects identify application backends for Cisco Multicloud Defense reverse proxy services.

A reverse proxy target address object is specified as a backend target address in a reverse proxy service object. It is used by the service object to establish a backend connection to an application. The application can be the address of one or more application load balancers or instances in the form of IPs or FQDNs. The different types of reverse proxy target address objects are defined as follows:

Static IP/FQDN Address Object

An IP/FQDN address object is configured as a set of IP addresses or FQDNs. When more than one IP or FQDN is configured, the gateway handles the addresses without priority amongst the configured fields when setting up a backend connection. When an FQDN is configured, the gateway resolves the FQDN with DNS to determine the IP address to use when setting up a backend connection.

Dynamic Applications Address Object

An applications address object is configured as an individual application load balancer cloud resource determined by its applications tag. The configuration dynamically populates a set of IPs or FQDNs represented by the cloud resources, obtained from the cloud account using the Multicloud Defense real-time inventory discovery. Any changes to the cloud

resources will be automatically reflected in the address object. When the configuration results in more than one IP or FQDN, the gateway handles the fields with no priority amongst the set when setting up a backend connection. When the configuration result is an FQDN, the gateway will resolve the FQDN with the DNS to determine the IP address to use when setting up a backend connection.

Reverse Proxy Target Address Object Parameters

Find the required and optional parameters for static and dynamic Cisco Multicloud Defense reverse proxy target address objects.

Type	Mode: Dynamic or Static	Parameter	Required or Optional	Notes
IP/FQDN	Static	Value	Required	
Applications	Dynamic	CSP Account	Required	
		Region	Required	
		VPC/VNet ID	Required	
		Resource Group	Optional	Azure Only
		Tag/Value	Required	Single Tag Key-Value pair

System Objects

Find the predefined Cisco Multicloud Defense address objects available for common public, private, and unrestricted IPv4 traffic criteria.

Multicloud Defense provides a list of pre-defined address objects to simplify policy creation. All system objects cannot be deleted or modified. Users can choose to clone system objects if modification is needed.

Name	Description
Any	This represents the entire IPv4 address space.
any-private-rfc- 1918	This represents all IPv4 private address as defined in RFC-1918.
Internet	This represents the entire IPv4 public address space, minus the private IPv4 addresses (RFC1918).

Create a Source/Destination Address Object

Create a Cisco Multicloud Defense source or destination address object for use as traffic criteria in security policy rules.

For information on what this object is, see Source or Destination Address Object Parameters on page 148. Use the following procedure to create a src/dst address object in Multicloud Defense:

Procedure

1. Go to **Policies** > **Security Policies** > **Addresses**.

2. Click **Create**.
3. Select **Src/Dest**.
4. Enter a unique **Name** to identify the address object.
5. (Optional) Enter a description for the object. This may provide context to help differentiate the object from other objects.
6. Select the **Object Type**. For information on object types and what they are, see Address Objects on page 146. Select one of the following types:
   - IP/CIDR/FQDN
   - VPC/VNet ID
   - Security Group
   - Application ID (Azure only)
   - Instance ID
   - Subnet ID
   - User-Defined Tag
   - Geo IP
   - Service End Point (Cloud Service IP)
   - Group

   > **Note**
   >
   > If you select **Group**, you can include a specific IP address or a range of IP addresses to either include or exclude.

7. Depending on which type you selected in step 6, enter the following paramters:
   - **Value** - Enter a valid IP, CIDR, or FQDN IP address.
   - **CSP Account** - Use the drop-down menu to select a cloud service provider account that has already connected to the controller.
   - **Region** - Select the region your cloud service provider is located in.
   - **VPC** - Use the drop-down menu to select the VPC or VNet. Note that options available may change depending on the cloud service provider account your choose.
   - **Subnet** - Use the drop-down menu to select the subnet that applies to your VPC or VNet.
   - (Azure only) **Resource Group** - Use the drop-down menu to select the resource group that is compatible with your selections.
     - **Resource Level** - Use the drop-down menu to select a value.
     - **Resource Tag** - Use the drop-down menu to select a keyword as the resource tag.
     - **Value** - Enter a valid value for the resource group. Note that this is different from the Value entry expected for IP/CIDR/FQDN objects.
   - **Geo IP** - Use the drop-down menu to select a specific IP that is associated with the gelocation of your choice.

- **X-Forwarded-For Match Enabled** - Check this box to allow the gateway to match against XFF HTTP header fields.
- **Address** - Select an existing object. This selection determines the group of addresses that
- **Include Addresses** - This option is only applicable if you select "Group" as the type in step 6. Enter a specific IP address or a range of IP addresses to include. You can also use `any` to include all valid addresses.
- **Exclude Addresses** - This option is only applicable if you select "Group" as the type in step 6. Enter a specific IP address or a range of IP addresses to exclude. You can also use `any` to include all valid addresses. Note that there is no validation from the Multicloud Defense Controller for address exclusion.

8. (Optional) Include a **Matching Expression**. This represent the set of conditions which must be matched for the object to execute.
9. Click **Save** when complete.

Create a Reverse Proxy Target Address Object

Create a Cisco Multicloud Defense reverse proxy target address object to identify an application backend for a reverse proxy service.

For more information on what this object is, see Reverse Proxy Target Address Object Parameters on page 151. Use the following procedure to create a reverse proxy target address object in Multicloud Defense:

Procedure

1. Navigate to **Policies** > **Security Policies** > **Addresses**.
2. Click **Create**.
3. Select **Reverse Proxy Target**.
4. Enter a unique **Name** to identify the address object.
5. (Optional) Enter a description for the object. This may provide context to help differentiate the object from other objects.
6. Select the **Object Type**. For information on object types and what they are, see Address Objects on page 146. Select one of the following types:
   - IP/CIDR/FQDN
   - Applications
7. Depending on which type you selected in step 6, enter the following paramters:
   - **Value** - Enter a valid IP, CIDR, or FQDN IP address.
   - **CSP Account** - Use the drop-down menu to select a cloud service provider account that has already connected to the controller.
   - **Region** - Select the region your cloud service provider is located in.
   - **VPC** - Use the drop-down menu to select the VPC or VNet. Note that options available may change depending on the cloud service provider account your choose.
   - **Subnet** - Use the drop-down menu to select the subnet that applies to your VPC or VNet.
   - (Azure only) **Resource Group** - Use the drop-down menu to select the resource group that is compatible with your selections.

8. Use the drop-down menus to select both an existing **Applications Tag** and its **Value** for this object.
9. Click **Save** when complete.

Edit Address Objects

Edit the supported configuration parameters of an existing Cisco Multicloud Defense address object.

If you need to modify a parameter that cannot be modified, you will need to Clone the address object and then change the parameters as desired.

Use the following steps to edit an address object. Note that not all parameters can be edited.

Procedure

1. Navigate to **Policies** > **Security Policies** > **Addresses**.
2. Check the box next to the address object you would like to **Edit**.
3. Click **Edit**.
4. Modify the parameters as desired.
5. Click **Save** when complete.

Clone Address Objects

Clone a Cisco Multicloud Defense address object to reuse or modify its configuration and update the required policy associations.

If the desire is to use the clone in place of the original, you will need to replace all associations of the original with the clone. The associations will be in a set of one or more security policy rule set rules or reverse proxy service objects. The associations can be seen by viewing the Address Object Details.

Use the following steps to clone an existing address object:

Procedure

1. Navigate to **Policies** > **Security Policies** > **Addresses**.
2. Check the box next to the address object you would like to **Clone**.
3. Click **Clone**.
4. Specify and modify the parameters as desired.
5. Click **Save** when complete.

Delete Address Object

Delete an unused Cisco Multicloud Defense address object after removing its security policy and reverse proxy service associations.

If an address object is actively used in a policy rule set or a reverse proxy service object, it will have one or more associations, and you will be unable to delete the address object. In order to delete an address object, you must first

remove all associations, then the address object can be deleted. The associations can be seen by viewing the Address Object Details.

Procedure

1. Navigate to **Policies** > **Security Policies** > **Addresses**.
2. Check the box next to the address object you would like to **Delete**.
3. Click **Delete**.
4. Click **Save** to confirm the delete.

View Details

Learn how to review an address object's resolved values, configuration, and associations with security policies or reverse proxy services.

You can view the address object **Details** by clicking the **Name** of an object from the **Policies** > **Security Policies** > **Addresses** page. The **Details** will display the IPs, CDIRs and FQDNs populated based on its type and configuration. It will also display the associations with policy rule sets and any object services.

16 FQDN Objects

Topics:

- FQDN Match Object

Learn how to use FQDN objects in Multicloud Defense to define domain-based destinations for security policies.

FQDN Match Object

Learn how Cisco Multicloud Defense FQDN match objects identify domain-based HTTP and encrypted TLS traffic for security policy rules.

A Fully Qualified Domain Name (FQDN) Match Object evaluates the Server Name Indication (SNI) associated with TLS-encrypted traffic or the Host header for unencrypted HTTP traffic. It uses the results of the evaluation for rule matching. If the traffic matches all match objects (Address, FQDN, Service) associated with a rule, then the rule is used to process the traffic. To evaluate the FQDN, traffic must be TLS encrypted and contain an SNI in an unencrypted TLS Hello header or be unencrypted HTTP and contain a Host header. The FQDN can be evaluated for traffic that is processed by either a **Forwarding** or **Forward Proxy** rule. The set of FQDNs in the profile is specified as strings representing the full domain or as strings represented by a Perl Compatible Regular Expression (PCRE).

Note

The FQDN match object is organized as a table containing user-specified rows (FQDNs).

The rows do not contain log-related actions to perform. This is because FQDN match object is a first-level matching criteria. When you have a clear list of FDQNs that you want to allow, you can use FQDN match objects. After a rule match, if you have categories that you want to allow based on criteria, use FQDN filtering. For more information, see Fully Qualified Domain Name Filter Profile on page 199.

The limits for each FQDN match object are as follows:

- Maximum user-specified rows: 254 (Standalone or Group of Standalones)
- Maximum FQDNs per row: 60
- Maximum FQDN character length: 255

When specifying a multilevel domain (for example, `www.example.com`), it's important to escape the `.` character (for example,`www\.example\.com`), otherwise it treats it as a wildcard for any single character.

Standalone vs. Group

Learn when to use standalone or group FQDN match objects to organize reusable domain criteria for Cisco Multicloud Defense policies.

An FQDN Match Object can be specified as type Standalone or Group.

A FQDN Match Standalone Object contains FQDNs. The object will be applied directly to a set of one or more Policy Ruleset Rules or associated with a FQDN Match Group Object.

A FQDN Match Group Object contains an ordered list of Standalone FQDN Objects that can be defined for different purposes and combined together into a Group Object. The Group Object can be applied directly to a set of one or more Policy Ruleset Rules. Each team can create and manage specific Standalone Profiles. These Standalone Profiles can be combined together into a Group Profile to create hierarchies or different combinations based on use case. An example combination could be a global FQDN list that would apply to everything, a CSP-specific list that would apply to each different CSP, and an application-specific list that would apply to each different application.

Create Standalone FQDN Match Object

Create a standalone Cisco Multicloud Defense FQDN match object with domain patterns and optional decryption exceptions.

Procedure

1. Navigate to **Policies** > **Security Policies** > **FQDN**.
2. Click **Create**.
3. Provide a Profile Name and Description.
4. Specify the Type as Standalone.
5. Click **Add** to create a new row.
6. Specify individual FQDNs (e.g., www.twitter.com,.*.google.com).
   a) Each FQDN is specified as a PCRE (Perl Compatible Regular Expression).
   b) Consider escaping the . character, or else it will be treated as a single character wildcard.
7. (Optional) Specify Decryption Exception for any FQDNs where decryption is not desired or possible. Possible reasons for considering Decryption Exception include:
   - The desire to not inspect encrypted traffic (financial services, defense, healthcare, etc.).
   - SSO authentication traffic where decryption is not possible.
   - NTLM traffic that cannot be proxied.
8. Click **Save**.

Create Group FQDN Match Object

Create a group FQDN match object from reusable standalone profiles for use in Cisco Multicloud Defense security policies.

Procedure

1. Navigate to **Policies** > **Security Policies** > **FQDN**.
2. Click **Create**.
3. Provide a Profile Name and Description.
4. Specify the Type as Group.
5. Select an initial Standalone Profile (at least one Standalone Profile is required).
6. Specify additional Standalone Profiles.
7. Click **Add FQDN Profile** to create a new row.
8. Select a Standalone Profile.
9. Click **Save** when completed.

Associate the Object

Learn where to associate Cisco Multicloud Defense FQDN match objects with security policy rules.

To access and associate objects to policy rules, see Rules.

17 Service Objects

Topics:

- Reverse Proxy Service Object (Ingress)
- Forward Proxy Service Object (Egress / East-West)
- Forwarding Service Object (Egress / East-West)

Explains the use and function of service objects in network policies.

Service objects are used to define layer 4 match criteria (destination ports, protocol) within a policy rule. They are also used to define the connection type for an incoming traffic flow. For an Egress or East-West Gateway, the connection type is either set to Forwarding or Forward Proxy.

Reverse Proxy Service Object (Ingress)

Create a reverse proxy service object for Cisco Multicloud Defense to protect incoming application traffic and route it to backend targets.

Ingress service objects are used in the Ingress/Reverse proxy rules. The object defines a listener port that the Multicloud Defense gateway listens for the traffic it receives and forwards to the target/backend address. Listener port can be configured with a decryption profile that has a TLS certificate configured. When the traffic hits the listener port, Multicloud Defense Gateway returns the TLS certificate configured. consider the following configurable options:

- An SNI can be configured on this port. This enables a single listener port (e.g 443) to be proxied to multiple backend targets based on the SNI.
- Target defines the backend address object and port to forward the traffic. The proxied traffic can be forwarded as HTTP, HTTPS, TCP, or TLS.

Use the following procedure to create and add a reverse proxy service object:

Procedure

1. Navigate to **Policies** > **Security Policies** > **Services**.
2. Click **Create**.
3. Click **Reverse Proxy**.
4. Provide a **Name** and **Description**.
5. Configure proxy parameters as defined below:

Option	Description
Decryption Profile	Assign a decryption profile, which also includes the server certificate, to be used for the proxy service.
Dst Port	Assign a destination port. For most web-based services, the destination port will be 443. This is the port Multicloud Defense Gateway listens on for the incoming traffic.
Protocol	TCP is the default.
SNI	Enter the list of SNIs.
Target Backend Port	Enter the Target/Backend application port number.
Protocol	Select the backend protocol.
Address	Select a backend IP address. The IP address in most cases will be the frontend IP of an internal load balancer.

Note

If the proxy service is required to run on multiple ports, you can add more entries. However all the ports serve the same certificate and are proxied to the same backend destination address object.

Forward Proxy Service Object (Egress / East-West)

Create a forward proxy service object for Cisco Multicloud Defense to inspect and protect outbound HTTP or TLS traffic.

Forward Proxy services are specifically used for HTTP based traffic. The object defines a listener port that the Multicloud Defense Gateway listens for the traffic it receives and forwards to the address/host that's available in the TLS SNI extension header or HTTP Host Header.

Note

We recommend using this for egress/East-West traffic.

Use the following procedure to create and add a forward proxy service.

Procedure

1. Navigate to **Policies** > **Security Policies** > **Services**.
2. Click **Create**.
3. Click **Forward Proxy**.
4. Provide a name and description.
5. Optionally, select the Application IDs to match.
6. Configure proxy parameters as defined below.

Option	Description
Decryption Profile	Assign a decryption profile, which includes the certificate to be used. Multicloud Defense impersonates the external certificate by signing it with the certificate provided in this profile. The root certificate is assumed to be installed on all the client application instances.
Dst Port	Assign a destination port. For most web-based services, the destination port will be 443.
Protocol	HTTP or HTTPS.

Note

- Multicloud Defense listens on the **Dst Port** and waits for the HTTP Host header or TLS SNI Header packet. Once Multicloud Defense receives this packet, it connects to the host using the protocol. If the protocol is HTTPS, the received certificate data from the external host is signed by the certificate in the decryption profile and sent to the client. The root certificate **must** be installed on the client app instances to avoid a certificate error.
- For a given destination port, there can be only one decryption profile (root CA certificate) association in a policy rule set across all service objects.
- During a forward proxy session, Multicloud Defense Gateway performs a DNS lookup on the destination with DNS request timeout of 30 seconds and cache age-out of TTL seconds.

Forwarding Service Object (Egress / East-West)

Create a forwarding service object for Cisco Multicloud Defense egress or east-west traffic that does not require proxying or decryption.

Forwarding service objects are used in the forwarding rules. The traffic that matches this type of rule/service is not proxied, and is forwarded as-is. This means there is no deep packet inspection and no Application ID on *encrypted* traffic.

Note

We **strongly** recommend using this for East-West traffic.

Use the following procedure to create and add a forwarding service object:

Procedure

1. Navigate to **Policies** > **Security Policies** > **Services**.
2. Click **Create**.
3. Click **Forwarding**.
4. Provide a name and description.
5. Multicloud Defense supports source NAT on a per service level. For traffic that requires source IP preservation (e.g. East-West traffic), disable SNAT.

   For Egress traffic, SNAT **must** always be enabled.
6. Configure port parameters as defined below.

Option	description
Dst Port	Assign a destination port or a range of destination ports as `start-end`.

Option	description
Protocol	TCP, UDP, ICMP

Note

In a forwarding policy, deep packet inspection operations **only** occur on non-encrypted traffic.

18 Certificates and Keys

Topics:

- Certificates and Keys
- Server Certificate Validation

Certificates and Keys

Learn how certificates and private keys support TLS inspection for ingress and egress proxy traffic, and review options for importing or securely storing private keys.

TLS certificates and keys are used by the Multicloud Defense Gateway in proxy scenarios. For ingress (reverse proxy) users access the application via Multicloud Defense Gateway and it presents the certificate configured for the service. For egress (forward proxy) cases, the external host's certificate is impersonated and signed by the certificate defined. The gateway instance must act as a certificate authority itself. This means that it must have its own self-signed certificate with the associated private key to sign impersonation certificates generated locally for any external websites that the client is initiating a connection to. The impersonation server certificate superficially resembles the actual external server's certificate and is generated on the go, when the client issues a TLS Hello request to connect to a unique website.

The certificate body is imported to the Multicloud Defense Controller. The private key can be provided in the following ways:

- Import the private key contents.
- Store in AWS Secrets Manager and provide the secret name.
- Store in AWS KMS and provide the cipher text contents.
- Store in GCP Secret Manager and provide the secret name.
- Store in Azure Key Vault and secret and provide the keyvault and secret name.

For testing purposes you can also generate a self-signed certificate on the Multicloud Defense Controller. This is similar to importing the private key contents from your local file system.

Note

Certificates are **NOT** editable once created. If you need to replace the existing certificate, you will need to create a new certificate, edit the decryption profile to reference the new certificate, and then delete the old certificate.

When importing the certificate and private key, the Multicloud Defense Controller/UI can detect if there is a mismatch. However, when using any other import method where the private key is stored within the cloud service provider, the Multicloud Defense Controller / UI will not be able to detect if there is a mismatch. This is by design to ensure the private key remains private and within your cloud service provider. When the private key is needed by the Multicloud Defense Gateway, it is accessed and used, and if there is a mismatch, an error is generated.

Import Certificate

Import a TLS certificate, private key, and optional certificate chain into Cisco Multicloud Defense for use with proxy traffic inspection.

Procedure

1. Navigate to **Policies** > **Security Policies** > **Certificates**.
2. Click **Create**.
3. When prompted with the **Method**, choose **Import your Certificate and Private Key**.
4. Copy the contents of the certificate file in the **Certificate Body**. This can include the certificate and the chain.
5. Copy the contents of the private key in **Certificate Private Key**.
6. (Optional) Import the chain into the **Certificate Chain** if your certificate and the chain are in different files.

7. Click **Save**.

AWS - KMS

Import a TLS certificate into Cisco Multicloud Defense and use AWS Key Management Service encrypted ciphertext to provide the associated private key.

Procedure

1. Navigate to **Policies** > **Security Policies** > **Certificates**.
2. Click **Create**.
3. In the Method choose *Import AWS - KMS*.
4. Select the cloud account and the region.
5. Copy the contents of the Certificate file in the *Certificate Body*. This can include the certificate and the chain.
6. Copy the AWS KMS encrypted cipher text in the *Private Key Cipher Text*. To generate AWS KMS encrypted cipher text, see AWS KMS - Encrypt Private Key on page 169.
7. Click **Save**.

AWS KMS - Encrypt Private Key

Explains AWS KMS encryption options to secure private keys.

Use the following command to encrypt the private key stored in the local file *privatekey.pem*.

```
aws kms encrypt --key-id 12345678-1234-1234-1234-12345678900 \
    --plaintext fileb://privatekey.pem \
    --output text --query CiphertextBlob > encrypted-cipher-text.txt
```

Copy the content of *encrypted-cipher-text.txt* and provide this as the content for the **Private Key** on the Multicloud Defense Controller while using the AWS - KMS option.

For more information, see AWS KMS Encyption help.

AWS - Secrets Manager

Import a TLS certificate into Cisco Multicloud Defense and retrieve its associated private key from AWS Secrets Manager for proxy traffic inspection.

Procedure

1. Navigate to **Policies** > **Security Policies** > **Certificates**.
2. Click **Create**.
3. In the Method choose *Import AWS - Secret*.
4. Select the cloud account and the region.
5. Copy the contents of the Certificate file in the *Certificate Body*. This can include the certificate and the chain.
6. Provide the Secret Name where the private key is stored. The private key contents must be stored as *Other type of Secrets* > *Plain Text* in the AWS Secrets Manager.
7. Click **Save**.

Azure Key Vault

Import a TLS certificate into Cisco Multicloud Defense and retrieve its associated private key from an Azure Key Vault secret.

Procedure

1. Navigate to **Policies** > **Security Policies** > **Certificates**.
2. Click **Create**.
3. In the Method, choose *Import Azure - Key Vault Secret*.
4. Select the cloud account and the region.
5. Copy the contents of the Certificate file in the *Certificate Body*. This can include the certificate and the chain.
6. Provide the Key Vault Name and the Secret Name where the private key is stored.
7. Click **Save**.

GCP - Secret Manager

Import a TLS certificate into Cisco Multicloud Defense and retrieve its associated private key from a Google Cloud Secret Manager secret.

Procedure

1. Navigate to **Policies** > **Security Policies** > **Certificates**.
2. Click **Create**
3. In the Method, choose *Import GCP - Secret*
4. Select the cloud account
5. Provide the Secret Name (full path) and the Secret Version
6. Copy the contents of the Certificate file in the *Certificate Body*. This can include the certificate and the chain.
7. Click **Save**.

Server Certificate Validation

Learn when to validate server certificates for secure forward proxy traffic and where to configure validation actions for TLS-decrypted connections.

When the gateway acts as a forward proxy, server certificate validation is automatically included in traffic processing. A designated server certificate validation **action** is not required in order to process traffic but it can improve the general security. By default, server certificate validation is not enabled and traffic going to servers that may have an invalid server certificate passes. Enable a server certificate validation action to prioritize rules for traffic that should not be allowed, or for specific traffic that should be trusted even regardless of its server certificate validations state.

> **Note**
>
> This validation process is **only** applicable for forward proxy environments and when **decryption** is enabled.

We recommend that you enable server certificate validation actions primarily in the TLS decrpytion profile for general rule actions. You can modify FQDN service objects to enable validation actions if you need to override the TLS decryption selection. You can include and enable a server certificate validation in two methods:

- TLS Decryption
- FQDN Service Object

Server Certificate Validation in the TLS Decryption Profile

Configure server certificate validation in a TLS decryption profile to define how Cisco Multicloud Defense handles connections with invalid certificates across associated rule sets.

When you select an action for server certificate validation within a TLS decryption profile, this action is used in all the rule sets that use this decryption profile. By default the validation action is configured to allow all traffic regardless of whether the server certificate is valid or not, and Multicloud Defense does not generate an alert within the HTTPs logs.

> **Note**
>
> If you enable the validation check to **Log**, locate the logs in **Policies** > **Profiles** > **Decryption**.

Use the following procedure to enable the server certificate validation in the TLS decrpytion profile:

Procedure

1. From the Multicloud Defense Controller, navigate to **Policies** > **Profiles** > **Decryption**.
2. Select the TLS decryption profile you want add the server certificate validation to. If you do not have a profile ready, create one here. See Decryption Profile on page 188 for more information.
3. **Edit** the decryption profile.
4. Under the **Profile Properties** section, expand the **Invalid Server Certificate Action** drop-down.
5. Select one of the following options:
   - **Deny Log** - This option automatically drops connections that do not provide a validated server certificate and logs the incident.
   - **Deny No Log** - This option automatically drops connections that do not provide a validated server certificate and **does not** log the incident.
   - **Allow Log** - This option allows connections that do not provide a validated server certificate to pass and logs the incident.
   - **Allow No Log** - This option allows connections that do not provide a validated server certificate to pass and **does not** log the incident. This is the default action selection.

6. Click **Save**.

What to do next

Ensure the TLS decryption profile is correctly associated with a forward proxy service object. See Forward Proxy Service Object (Egress / East-West) on page 163 for more information.

Once the TLS decrpytion profile is included in a service object, confirm that the rule order within the policy is ordered in a way that supports how you want traffic processed.

Server Certificate Validation in the FQDN Service Object

Configure server certificate validation in an FQDN service object to allow or block traffic for a specific server and override the TLS decryption profile when needed.

Invalid server certificate validation within the FQDN service object is optional. If specified it will override the behavior designated in the TLS decryption profile. If you do not specify a selection here, no additional action or override action is taken. You can use the invalid server certificate validation within the FQDN service object to block or allow traffic for a specific server that may otherwise be blocked or allowed by the TLS decryption profile.

Note that when you enable the validation check to **Log**, these logs are located in **Investigate** > **Flow Analytics** > **HTTPS Logs**.

Use the following procedure to include a server certificate validation action in a FQDN service object:

Procedure

---

1. From the Multicloud Defense Controller, navigate to **Policies** > **Security Policies** > **FQDN**.
2. Select the FQDN service object you want to modify.
3. **Edit** the selected FQDN service object.
4. In the list of FQDN service objects included in the ruleset, expand the **Invalid Server Certificate Action** drop-down menu and select one of the following options:
   - **Deny Log** - Automatically drop connections that do not provide a validated server certificate and logs the incident.
   - **Deny No Log** - Automatically drop connections that do not provide a validated server certificate and **does not** log the incident.
   - **Allow Log** - Allow connections that do not provide a validated server certificate to pass and logs the incident.
   - **Allow No Log** - Allow connections that do not provide a validated server certificate to pass and **does not** log the incident.
5. Click **Save**.

What to do next

Ensure the FQDN service object is correctly associated with a rule or rule set. See Rule Sets and Rule Set Groups on page 133 for more information.

Once the FQDN service object is successfully assocaited with a rule or rule set in your policy, confirm that the rule order within the policy is ordered in a way that supports how you want traffic processed.

Certificate for Custom Root CA

Learn how custom root CA certificates establish trust for TLS decryption and support secure traffic inspection in Cisco Multicloud Defense.

The Custom Root Certificate Authorities (CAs) store is a repository or storage location used in computing environments, such as operating systems or network devices, to store custom root CA. This allows organizations to add their own root certificates or third-party root certificates that are not included by default in the standard CA store provided by operating systems or browsers.

When you have a **decryption profile** enabled in your tenant, the Multicloud Defense Controller automatically initiates the validation of certificates for a custom root CA. The decryption profile relyies on it to enable secure traffic inspection, necessitating careful management and trust establishment across the network.

19 Certificate and Keys Tech Notes

Topics:

- Azure Key Vault
- Generate a Self-Signed Root CA
- Generate a Certificate Signed by your Self-Signed Root CA
- Generate an Intermediate CA Signed by Your Root CA
- App Certificate signed using the Intermediate CA
- Install Root CA as Trusted CA on the Hosts

Find technical guidance for certificates and keys used with Multicloud Defense cloud security deployments.

Azure Key Vault

Learn how Cisco Multicloud Defense retrieves securely stored certificate private keys from Azure Key Vault secrets for use with TLS decryption.

Multicloud Defense integrates with many secret management systems offered by public cloud providers to store the private keys corresponding to the certificates to be used in TLS decryption. In this section, you can view the steps to retrieve a stored private key in Azure Key Vault for a certificate. Multicloud Defense supports retrieval of private keys through Azure Key Vault secrets.

Create Azure Key Vault and User-assigned Managed Identities

Create an Azure Key Vault and user-assigned managed identity, then grant the identity permission to retrieve certificate private keys from the vault.

Multicloud Defense uses a user-assigned managed identity to access the Azure Key Vault secret. The following steps show how to create the Key Vault and how to grant the managed identity the necessary permissions to access the secret:

Procedure

1. Create a new User-assigned Managed Identity.
2. Create a new **Key Vault**, for example, `mcd-key-vault`.
3. In the newly created Key Vault, go to **Access Policies**.
4. Select **Add Access Policy**.
5. Assign Secret permissions: `Get`.
6. Select **Principal**.
7. Select the **Managed Identity**.
8. Staying on the Key Vault configuration, navigate to **Access Control (IAM)** > **Select Add** > **Role Assignment**.
9. Select **Reader Role**.
10. Search for the Managed Identity in the **Select** field and select it.
11. Click **Save**.

Store Private Keys into Azure Key Vault Secrets

Store a certificate private key in an Azure Key Vault secret by importing the key file through Azure Cloud Shell.

Import your private key for your secret into the secrets folder of the Azure Key Vault that you have created. Since the private key is a multi-line string, this can be imported via the Azure Cloud Shell.

Start the Azure Cloud Shell and run the following command:

```
az keyvault secret set --file <privatekey file> --encoding ascii --vault-name
 <key vault name> --name <secret name>
```

The key file is created by simply creating a new file with a text editor (vi) in Azure Cloud Shell, pasting the contents of the private key and saving it.

Configure Multicloud Defense to Retrieve Private Key

Configure Cisco Multicloud Defense to retrieve a certificate private key from Azure Key Vault and use the certificate in a TLS decryption profile.

Before you begin, ensure that the Multicloud Defense Gateway deployed in Azure is configured with the user- managed identity as specified in Azure Key Vault and User-assigned Managed Identities.

Paste the certificate body into the field associated with the private key in the Azure Key Vault secret

Once you have created the Multicloud Defense certificate, it can be used with a Decryption Profile when defining services in Multicloud Defense Controller.

Generate a Self-Signed Root CA

Generate a self-signed root certificate authority for forward and reverse proxy deployments, and prepare the certificate for installation on client systems.

Generate a self-signed root certificate authority (CA).

```
openssl genrsa -out myca.key 2048
# password protect key: openssl genrsa -out myca.key -des3 2048
openssl req -x509 -new -key myca.key -sha384 -days 1825 -out myca.crt \
  -subj "/C=US/ST=CA/L=Santa
Clara/O=MyOrg/OU=SecurityOU/CN=rootca.myorg.com/emailAddress=rootca@myorg.com"
```

This root CA must be installed as a trusted root CA on the users (client) machines

Note

Generating a self-signed certificate using **macOS** will not produce a certificate suitable for forward and reverse proxy scenarios. The certificate must have the *Is CA* option set to *True* and the certificate generated using macOS does not. It is recommended that the self-signed certificate be generated from within the Multicloud Defense UI (Certificates > Create > Generate) or using **Linux**.

Generate a Certificate Signed by your Self-Signed Root CA

Generate an application certificate signed by your self-signed root certificate authority for use with secure application traffic in proxy deployments.

Generate a certificate signed by the above root certificate authority (CA). You can use this certificate in applications.

```
openssl genrsa -out app1.key 2048
# password protect key: openssl genrsa -out -des3 app1.key 2048
openssl req -new -key app1.key -out app1.csr \
  -subj "/C=US/ST=CA/L=Santa
Clara/O=MyOrg/OU=AppOU/CN=app1.myorg.com/emailAddress=app1@myorg.com"
openssl x509 -req -in app1.csr -CA myca.crt -CAkey myca.key -out app1.crt
-sha384\
  -days 365 -CAcreateserial -extensions SAN \
  -extfile <(printf "
[SAN]\nbasicConstraints=CA:false\nsubjectAltName=DNS:app1.myorg.com,DNS:app1-
 1.myorg.com,IP:192.168.10.21,IP:192.168.10.22")
```

Generate an Intermediate CA Signed by Your Root CA

Generate an intermediate certificate authority signed by your root CA so you can issue application certificates without using the root CA directly.

If you don't want to use the root certificate authority (CA) to sign app certs, then create an intermediate CA signed by the root CA, then sign the app certs using the intermediate CA. Append the intermediate cert to the app cert. At this point the app crt has 2 certs (as a chain).

```
openssl genrsa -out interca.key 2048
# password protect key: openssl genrsa -out -des3 interca.key 2048
openssl req -new -key interca.key -out interca.csr \
  -subj "/C=US/ST=CA/L=Santa
Clara/O=MyOrg/OU=InterSecurityOU/CN=interca1.myorg.com/emailAddress=interca1@myorg.com"
openssl x509 -req -in interca.csr -CA myca.crt -CAkey myca.key -out interca.crt
 - sha384 \
  -days 365 -CAcreateserial -extensions SAN \
  -extfile <(printf "[SAN]\nbasicConstraints=CA:true")
```

App Certificate signed using the Intermediate CA

Generate an application certificate signed by an intermediate certificate authority, create the certificate chain, and prepare it for use with applications and trusted client hosts.

```
openssl genrsa -out app1.key 2048
# password protect key: openssl genrsa -out -des3 app1.key 2048
openssl req -new -key app1.key -out app1.csr \
  -subj "/C=US/ST=CA/L=Santa
Clara/O=MyOrg/OU=AppOU/CN=app1.myorg.com/emailAddress=app1@myorg.com"
openssl x509 -req -in app1.csr -CA interca.crt -CAkey interca.key -out app1.crt
 - sha384 \
  -days 365 -CAcreateserial -extensions SAN \
  -extfile <(printf "
[SAN]\nbasicConstraints=CA:false\nsubjectAltName=DNS:app1.myorg.com,DNS:app1-
 1.myorg.com,IP:192.168.10.21,IP:192.168.10.22")
```

Append the files app1.crt and interca.crt to create a combined certificate, and use it in your application. The root CA must be installed as a trusted root CA on your client machines.

Install Root CA as Trusted CA on the Hosts

Install a root CA certificate as a trusted certificate on Ubuntu, CentOS, or Windows hosts so clients trust certificates issued for proxy traffic.

OS	Command
Ubuntu	Copy crt file to `/usr/local/share/ca-certificates`. Run the command `sudo update-ca- certificates`.
CentOS	Copy crt file to `/etc/pki/ca-trust/source/anchors`. Run the command `sudo update-ca- trust extract`.

OS	Command
Windows	Double click the file and add the cert to Trusted Root, or run the command `certutil -addstore "Root" <crt-file>`.

20 Types of Traffic

Topics:

- Enable DNS Logs
- Enable VPC Flow Logs

Explains traffic log interactions and enabling traffic visibility on cloud accounts.

When enabled, traffic logs are generated whenever traffic hits a rule. These log interactions record information about incoming and outgoing traffic, including the source and destination IP addresses, port numbers, and protocols used. Logs can be incredibly useful to audit the network; monitor activity, investigate potential security breaches, or simply keep an eye on what is happening with your firewall. Traffic visibility can be enabled at any time but we strongly recommend enabling traffic immediately after onboarding a cloud service provider account and assigning a gateway policy.

Enabling traffic visibility is a different process for every cloud account type, but typically you will need to identify account characteristics such as your cloud account's region, VPC/VNet that you want to monitor, network security groups, and a cloud storage account for logs.

If you did not onboard an account with the Easy Setup wizard or if you did not enable traffic visibility from the Easy Setup wizard, we strongly recommend enabling the following logs:

- VPC/VNet Flow Logs
- DNS Logs
- Route53 Query Logging.

> **Note**
>
> You can download logs for flows and events. In the Time Range section, select a time range and click the download icon. A maximum of 10,000 records are downloaded in a single instance. You will need to repeat the step to download larger sets of records.

Enable DNS Logs

Learn how DNS query logs improve visibility into domain activity, suspicious requests, resolution problems, and network performance across supported cloud environments.

Domain Name System (DNS) logs are a series of recorded queries and responses handled by a network. These logs capture details such as the domain names queried, the IP addresses returned, timestamps of the queries, and the devices making the requests. These logs can help ientify suspicious activities, such as attempts to access malicious domains or unusual spikes in DNS traffic; they can also help assist in diagnosing network issues by providing insight into DNS resolution problems, helping to identify whether domain resolution failures are due to network issues. Utilizing these logs as a part of your moitoring routing gives you the advatage when it comes to performance and efficiency of DNS queries, helping to optimize network performance and reduce latency.

AWS: Enable DNS Logs

Enable Amazon Route 53 query logging for selected VPCs so Cisco Multicloud Defense can provide DNS traffic insights across your AWS environment.

If you provided a S3 bucket during the stack creation from the CloudFormation template in the previous section, a S3 bucket is created by the template that acts as the destination for the route53 Query Logs. The VPCs that are monitored for the DNS query logs must be added manually.

Procedure

1. In AWS Console go to the Route53Query Logging .
2. Select the **Query Logger** created by the template. Locate the logger with the prefix name provided in the template.
3. Select all the VPCs for which you want to get the traffic insights and click **Add**.
   a. Under the " VPCs that queries are logged for" section, click **Log queries for VPCs** or **Add VPC**.
   b. Select all the VPCs and click **Choose**.

GCP: Enable DNS Logs

Enable GCP DNS query logs, route them to Cloud Storage, and grant Cisco Multicloud Defense access for traffic visibility and analysis.

To enable GCP DNS query logs, follow the steps below.

Procedure

1. Navigate to VPC network in GCP console.
2. Open Google cloud shell and execute this command:
   ```
   gcloud dns policies create POLICY_NAME --networks=NETWORK --enable-logging
   ```
3. Navigate to **Cloud Storage** section and create a storage bucket. You can leave everything as default when creating storage bucket.

> **Note**
>
> *Both DNS and VPC logs can share the same cloud storage bucket.*

4. Navigate to **Logs Route** section.
5. Click on **Create Sink**.
6. Provide a sink name.
7. Select "Cloud Storage bucket" for sink service.
8. Select the cloud storage bucket that was created above.
9. In "Choose logs to include in sink" section, put in this string: `resource.type="dns_query"`.

   The following steps are the same as mentioned in the VPC flow log for GCP. If you are sharing cloud storage bucket, you only need to perform below steps once.

10. Click **Create Sink**.
11. Navigate to **IAM** > **Roles**.
12. Create a custom role with this permission: **storage.buckets.list**.
13. Create another custom role with following permission:

    storage.buckets.get storage.objects.get storage.objects.list.

14. Add both custom role to the service account created for Multicloud Defense Controller. When adding the second custom role, put this condition:

```
(resource.type == "storage.googleapis.com/Bucket" || resource.type ==
 "storage.googleapis.com/Object") &&
resource.name.startsWith('projects/_/buckets/<cloud storage name>')
```

15. Navigate to **Pub/Subs**.
16. Click on **Create Topic**.
17. Provide a Topic name and click **create**.
18. Click on **Subscriptions**. You will find that there is a subscription created for the topic that was just created.
19. Edit the subscription.
20. Change Delivery type as **Push**.
21. Once **Push** is selected, enter in the endpoint URL: `https://prod1-webhook.vtxsecurityservices.com:8093/webhook/<tenant name>/gcp/cloudstorage`. Tenant name is assigned by Multicloud Defense. To view tenant name, navigate to Multicloud Defense Controller and click on your username.
22. Click **Update**.
23. Create a cloud storage notification by opening a Google cloud shell and execute this command: `gsutil notification create -t <TOPIC_NAME\> -f json gs://<BUCKET_NAME>`.

Azure: DNS Logs

Learn about DNS logging limitations for Azure cloud accounts when configuring traffic visibility in Cisco Multicloud Defense and planning alternative monitoring coverage.

Azure currently does not expose DNS log queries. Multicloud Defense Controller cannot enable logs for this cloud service provider.

Enable VPC Flow Logs

Learn how VPC flow logs provide visibility into cloud network traffic for security monitoring, threat investigation, compliance, troubleshooting, and forensic analysis.

A VPC (Virtual Private Cloud) log refers to the logging of network traffic and events within a Virtual Private Cloud environment. These logs provide detailed records of the activities and communications between resources in a VPC and the external internet or other VPCs. In the context of a cloud-based firewall,VPC logs help identify unauthorized access attempts, suspicious activities, or potential security threats by providing visibility into network traffic patterns and anomalies. They can also demonstrate compliance with regulatory requirements by maintaining detailed records of network activity and access, or diagnosing connectivity issues, configuration errors, or bottlenecks within the VPC by providing insights into the traffic flow and network behavior. VPC logs serve as a crucial tool for forensic analysis, helping to trace the source and impact of the incident.

AWS: Enable VPC Flow Logs

Enable AWS VPC Flow Logs for selected VPCs and send complete traffic records to the designated Amazon S3 bucket for security analysis.

If you provided a S3 bucket during the stack creation from the CloudFormation template in the previous section, a S3 bucket is created by the template that acts as the destination for the VPC flow logs. Flow logs must be enabled for each of the VPCs.

To enable AWS VPC flow logs, follow the below steps:

Procedure

1. In the AWS Console, go to the VPCs section.
2. Select the VPC and select the **Flow Logs** tab for that VPC.
3. Select **All** as the filter.
4. Select **Send to an Amazon S3 bucket** as the destination.
5. Provide the S3 bucket ARN copied from the tutputs of the CloudFormation template stack.
6. Select Custom **Custom Format** as the log record format.
7. Select all the fields from the log format dropdown.
8. Click **Create Flow Log**.

GCP: Enable VPC Flow Logs

Enable GCP VPC Flow Logs for selected subnets, route them to Cloud Storage, and grant Cisco Multicloud Defense access for traffic visibility.

To enable GCP VPC flow logs, follow the below steps.

Procedure

1. In the GCP console, navigate to **VPC network**
2. to enable the VPC flow log, select the **subnet**.
3. Ensure that flow logs is turned **On**. If it is off, click the **Edit**option and turn flow logs on.

4. Turn on flow log on all subnets where you want to enable flow log.
5. Navigate to **Cloud Storage** section and create a storage bucket. You can leave everything as default when creating storage bucket.

   > **Note**
   >
   > Both DNS and VPC logs can share the same cloud storage bucket.

6. Navigate to the **Logs Route** section.
7. Click **Create Sink**.
8. Enter a name for the sink.
9. Select **Cloud Storage bucket** for sink service.
10. Select the cloud storage bucket that was created above.
11. In the **Choose logs to include in sink** section, enter this string: `logName:(projects/<project-id>/logs/compute.googleapis.com%2Fvpc_flows)`

    If you are sharing a cloud storage bucket, you only need to perform the remaining steps of this procedure once.

12. Click **Create Sink**.
13. Navigate to **IAM** > **Roles**.
14. Create one custom role with this permission: `storage.buckets.list`.
15. Create one custom role with following permission: `storage.buckets.get storage.objects.get storage.objects.list`.
16. Add both custom roles to the service account created for Multicloud Defense Controller. When adding the second custom role, enter the following condition:

```
(resource.type == "storage.googleapis.com/Bucket" || resource.type ==
"storage.googleapis.com/Object") &&
resource.name.startsWith('projects/_/buckets/<cloud
storage name>')
```

17. Navigate to **Pub/Subs**.
18. Click **Create Topic**.
19. Provide a **Topic** name and click **Create**.
20. Click **Subscriptions**. A subscription is created for the topic created in step 18.
21. **Edit** the subscription.
22. Change the **Delivery** type to **Push**.
23. Enter this as the endpoint URL:`https://prod1-webhook.vtxsecurityservices.com:8093/webhook/<tenant name>/gcp/cloudstorage`.

    Multicloud Defense automatically assigns the tenant name. To see tenant name, navigate to Multicloud Defense Controller and click on your username.

24. Click **Update**.
25. Open a Google cloud shell and execute the following command: `gsutil notification create -t <TOPIC_NAME> -f json gs://<BUCKET_NAME>`.

Azure: Enable VNet Flow Logs

Enable Azure VNet flow logs, store them in an Azure storage account, and configure event delivery for Cisco Multicloud Defense traffic visibility.

To enable Azure VPC/VNet flow logs, follow the below steps.

Procedure

1. Go to the **Resource Groups** section in Azure portal.
2. Click the **Create** button.
3. Choose the subscription and provide a name for this new resource group.
4. Select a **Region**. (example: (US) East US).
5. Click the **Review + create** button.
6. Go to the **storage accounts** section and click the **Create** button.
7. Select the **Subscription** and **Resource** group that was just created.
8. Select the same **region** as the resource group.
9. Provide a name for the storage account.

   Note that **Redundancy cannot** be locally-redundant storage(LRS)

10. Click the **Review + create** button. This creates a storage account where VNet flow logs are stored.
11. Go to the **Subscription** section and find the subscription that was recently created.
12. Navigate to **Resource Providers**.
13. Ensure that the `microsoft.insights` and `Microsoft.EventGrid` providers are registered. If they are not registered, click **Register**.
14. Go to the **Network Watcher** section.
15. Click **Add** and add the regions that you want VNet flow logs to be enabled for.
16. Go to **Network Watcher** > **Virtual Network flow logs**.
17. Create flow logs for the VNet where you want to enable VNet flow log. Provide the storage account created above. Set the **Retention days** as 30.
18. Navigate to the storage account created and click **Events**.
19. Click **Event Subscription**.
20. Provide a name for this event subscription.
21. Select the resource group that was created above.
22. Provide a **System Topic Name**.
23. For **Filter to Event Types**, the default value is **Blob Created** and **Blob Deleted**.
24. For **Endpoint Type**, select **Web Hook**.
25. Click the **Select an endpoint** link.

    The Subscriber Endpoint is `https://prod1-webhook.vtxsecurityservices.com:8093/webhook/<tenant_name>/azure`. Tenant name is assigned by Multicloud Defense. You can find the tenant name by clicking on the username in Multicloud Defense Controller.

21 Security Profiles

Topics:

- Decryption Profile
- Network Intrusion (IDS/IPS) Profile
- Data Loss Prevention (DLP) Profile
- Anti-Malware Profile
- Web Application Firewall (WAF) Profile
- URL (Uniform Resource Locator) Filter Profile
- Fully Qualified Domain Name Filter Profile
- Malicious IP Profile

Learn how security profiles define protection settings for cloud traffic in Multicloud Defense.

A security profile typically refers to a set of rules and configurations applied to network traffic to enforce security policies. These profiles include the following protective measures:

- Firewall Rules
- Intrusion Prevention Systems (IPS)
- Antivirus/Antimalware
- Web Filtering
- Data Loss Prevention (DLP)
- Application Control

These particular profiles are generally added to a policy rule, policy rule set, or a policy rule set group and ordered by priority.

Decryption Profile

Learn how a decryption profile enables TLS traffic inspection for forward and reverse proxy services and lets you select certificates, TLS versions, and cipher suites.

A decryption profile is used by the Multicloud Defense Gateway in a reverse proxy **or** forward proxy scenario. When a connection is proxied, the front-end session is terminated on the gateway and a new back-end session is established to the server. The intention of this termination is to decrypt and inspect the traffic to protect against malicious activity. In order to decrypt encrypted traffic, a decryption profile is necessary.

TLS Versions in your Decryption Profile

The Multicloud Defense Gateway supports all TLS versions: TLS 1.0, TLS 1.1, TLS 1.2, and TLS 1.3. Users can specify a minimum TLS version to use and Multicloud Defense Gateway will negotiate a TLS version that is equal to or higher than the specified minimum TLS version. The gateway always uses the highest TLS version possible combined with the strongest cipher suite during the TLS negotiation. In the case where the Multicloud Defense Gateway cannot negotiate a version that meets the minimum TLS version specified, the gateway drops the session and logging a `TLS_ERROR` event.

> **Note**
>
> Only a single minimum TLS version can be applied to a gateway. A consistent minimum TLS version must be used across all decryption profiles referenced by all service objects that are used within a policy ruleset or policy ruleset group. If different minimum TLS versions are specified, the minimum TLS version that will be applied cannot be predetermined.

Cipher Suites

The Multicloud Defense Gateway supports a set of default and user-selectable cipher suites. The default set are PFS cipher suites that are always selected. The user-selectable set are Diffie-Hellman and PKCS (RSA) cipher suites that can be selected by the user. The combined set of cipher suites (default and user-selected) are used by the gateway for establishing a secure front-end encrypted session. The client will send an ordered list of preferred cipher suites. The gateway will respond with a cipher suite chosen from the ordered set submitted by the client and the set available by the gateway. If the client allows the server to define the order, then the cipher suite chosen is from the ordered set available by the gateway and the set submitted by the client.

With version 24.10 and later, the Multicloud Defense Controller assists the creation of your decryption profile by auto-selecting the strongest cipher suites once you've chosen the minimum TLS version. Note that both the Multicloud Defense Controller and the Multicloud Defense Gateway must be running at least version 24.10. Older gateway versions do not support this automated help feature and you cannot edit the ciphers suites of an existing decrpytion profile. We strongly recommend updating your gateway to version 24.10 or later to take advantage of this functionality.

The following is an ordered list of cipher suites supported by the gateway and available in a decryption profile:

Category	Cipher Suite	Key Exchange	Cipher	Hash	Default
PFS	ECDHE-RSA-AES256-GCM-SHA384	ECDHE-RSA	AES256-GCM	SHA384	
PFS	ECDHE-RSA-AES256-CBC-SHA384	ECDHE-RSA	AES256-CBC	SHA384	
Diffie-Hellman	DH-RSA-AES256-GCM-SHA384	DH-RSA	AES256-GCM	SHA384	
PFS	DHE-RSA-AES256-GCM-SHA384	DHE-RSA	AES256-GCM	SHA384	
PFS	DHE-RSA-AES256-CBC-SHA256	DHE-RSA	AES256-CBC	SHA384	

Category	Cipher Suite	Key Exchange	Cipher	Hash	Default
PFS	DHE-RSA-AES256-CBC-SHA	DHE-RSA	AES256-CBC	SHA	
Diffie-Hellman	DH-RSA-AES256-SHA256	DH-RSA	AES256-CBC	SHA256	
Diffie-Hellman	DH-RSA-AES256-SHA	DH-RSA	AES256-CBC	SHA160	
PKCS (RSA)	AES256-GCM-SHA384	PKCS-RSA	AES256-GCM	SHA384	
PKCS (RSA)	AES256-SHA256	PKCS-RSA	AES256-CBC	SHA256	
PKCS (RSA)	AES256-SHA	PKCS-RSA	AES256-CBC	SHA160	
PFS	ECDHE-RSA-AES128-GCM-SHA256	ECDHE-RSA	AES128-GCM	SHA256	
PFS	ECDHE-RSA-AES128-CBC-SHA256	ECDHE-RSA	AES128-CBC	SHA256	
Diffie-Hellman	DH-RSA-AES128-GCM-SHA256	DH-RSA	AES128-GCM	SHA256	
PFS	DHE-RSA-AES128-GCM-SHA256	DHE-RSA	AES128-GCM	SHA256	
PFS	DHE-RSA-AES128-CBC-SHA256	DHE-RSA	AES128-CBC	SHA256	
Diffie-Hellman	DH-RSA-AES128-SHA256	DH-RSA	AES128-CBC	SHA256	
Diffie-Hellman	DH-RSA-AES128-SHA	DH-RSA	AES128-CBC	SHA160	
PKCS (RSA)	AES128-GCM-SHA256	PKCS-RSA	AES128-GCM	SHA256	
PKCS (RSA)	AES128-SHA256	PKCS-RSA	AES128-CBC	SHA256	
PKCS (RSA)	AES128-SHA	PKCS-RSA	AES128-CBC	SHA160	
PFS	ECDHE-RSA-DES-CBC3-SHA	ECDHE-RSA	DES-CBC3	SHA	
PFS	ECDHE-RSA-RC4-SHA	ECDHE-RSA	RC4	SHA	
PKCS (RSA)	RC4-SHA	PKCS-RSA	RC4	SHA160	
PKCS (RSA)	RC4-MD5	PKCS-RSA	RC4	SHA160	

Create a Decryption Profile

Create a decryption profile by selecting a certificate, minimum TLS version, and optional cipher suites, then attach the profile to a security policy ruleset.

Use the following procedure to create a decryption profile.

Procedure

1. Navigate to **Policies** > **Profiles** > **Decryption**.
2. Click **Create**.
3. Specify a **Profile Name** and a **Description**.
4. For **Certificate Method** choose **Select Existing**.
5. For **Certificate** choose the desired certificate.

6. For **Min TLS Version** choose the lowest TLS version that is accepted by the decryption profile. The default is TLS 1.0.
7. If using non-default (non-PFS) cipher suites, select the set of desired cipher suites from the Diffie- Hellman or PKCS (RSA) menus. Note that the Multicloud Defense Controller filters the available cipher suites so only the strongest ciphers are listed, dependent on the min TLS version you selected in step 6.
8. Click **Save**.

What to do next

Attach the profile to a policy rule set. See Rule Sets and Rule Set Groups on page 133 for more information.

Network Intrusion (IDS/IPS) Profile

Learn how a Network Intrusion profile uses IDS and IPS rules to detect, report, and block suspicious or malicious network traffic.

Network intrusion profiles are a collection of Intrusion Detection and Protection (IDS/IPS) rules that can be used to evaluate transactions to ensure the traffic is not malicious.

An Intrusion Detection System (IDS) is defined as a solution that monitors network events and analyzes them to detect security incidents and imminent threats, specifically suspicious or abnormal activity such as malicious transactions, and sends immediate alerts when it is observed. IDS searches for and against hosts and networks.

An Intrustion Protection System (IPS) actively analyzes network traffic and compared it against known attack patterns and signatures. When the system detects suspicious traffic, it blocks it from entering the network. IPS rules cover both network-bsed IPs and host-based IPs.

Multicloud Defense combines both of these systems within a singular profile to create an easy-to-configure network intrusion profile made to detect malicious probes or new network patterns from a compromised system that both detects, rejects, and reports supsicious traffic. Preemptive blocking and reporting can mitigate any downtime on your network and further improve blocking activity in the future. A network intrusion profile in Multicloud Defense comprises the following rule sets:

Table 2: Multicloud Defense supports the following IDS/IPS Rule Sets

Rule Sets	Description
Talos Rules	The Talos rules are a premium set of rules from Cisco based on intelligence gathered from real-world investigations, penetration tests and research that provide an advanced level of protection for applications and frameworks.

Note that the IDS/IPS profile does not include web applications that might be malicious. See Web Application Firewall (WAF) Profile on page 194 for more information.

Create an IPS/IDS Profile

Create an IDS/IPS Profile, select Cisco Talos or custom rulesets, define update and detection actions, and attach the profile to a security policy ruleset.

Use the following procedure to create and add an IPS/IDS profile to a ruleset:

Procedure

1. Navigate to **Policies** > **Profiles** > **IPS/IDS**.

2. Click **Create**.
3. Click into the **General Settings** tab.
4. Enter a unique **Profile Name**.
5. (Optional) Enter a **Description**. This may help differntiate between other profiles with a similar name.
6. Toggle the Threat PCAP option file if the IDS/IPS Profile detects malicious activity. Note that if you toggle this option **on**, you must have a PCAP profile attached to the gateway.
7. In the **Rule Set** section of the general settings, note that at least one ruleset from a rules library (Talos, Custom) is required to be specified in the IDS/IPS profile. If Talos rules and custom rulesets are used, at least one of the two must be enabled. If the desire is to disable the entire IDS/IPS Profile, remove the IDS/IPS Profile from any policy ruleset so the IDS/IPS profile will not be evaluated. Use the drop-down menu to select one of the followingsettings that are applied to all rulsets within this profile:
   - **Disabled** - Specify whether to disable the use of Talos rules.
   - **Manual** - Specify the Talos rule's version.
   - **Automatic** - Specify the number of days from publish date to delay automatic update to the latest Talos rule's version.

   Use the other drop-down menu to select when the rules within this profile are updated. You can opt to update the rule set **Immediately** after Talos sends out an update, or any number of days after the update.
8. Click **Talos Rules: Policy** and choose from the table which policy profile to use as a base. You can only select one profile.

   Unless your window view is maximized, scroll to the right of the window and assign an **action** for the selected profile:
   - **Rule Default** - Allow or Deny the requests based on the action specified in each triggered Rule and log an Event.
   - **Allow Log** - Allow the requests and log an event.
   - **Allow No Log** - Allow the requests and do not log an event.
   - **Deny Log** - Deny the requests and log an event.
   - **Deny No Log** - Deny the requests and do not log an event.
9. Click the **Talos Rules: Category** tab and choose at least one category from the table to the profile.
10. Click the **Talos rules: Class** tab and choose at least one class from the table to the profile.
11. At the top of the screen click into the **Advanced Settings** tab.
12. Under **Rule Supression** click **Add** and enter a valid **Source IP/CIDR List** of IP addresses and a corresponding **Rule ID List**. To remove a row of lists simply click the minus icon to the right of the row.
13. Under **Event Filtering: Profile Event Filtering**, enter the following information:
    - **Type** - You can opt for either Rate or Sample. Generated events are rate- or sample-limited based on the specified **Number of Events** triggered over a **Time** evaluation interval (in seconds).
    - **Number of Events** - Manually enter a value of allowed number of events.
    - (Available for the Rate type) **Time (Seconds)** - enter a numerical value in seconds.
14. Under **Event Filtering: Rule Event Filtering**, click **Add**. Enter the following information:
    - **Rule ID List** - Specify a comma-separated list of rule IDs.
    - **Number of Events** - Manually enter a value of allowed number of events.
    - (Available for the Rate type) **Time (Sec)** - enter a numerical value in seconds.

- **Type** - Select either Rate or Sample. Generated events are rate- or sample-limited based on the specified **Number of Events** triggered over a Timeevaluation interval (in seconds).

15. Under the **Rule Setting List** section of the advanced settings, click **Add** and enter the following:
    - **Source IP/CIDR List** - provide a comma-separated list of IPs or CIDRs
    - **Rule ID List** - provide a comma-separated list of rule IDs. Note that for high number rules, only the rule ID is necessary. For low number rules, the GID and ID need to be specified for the rule ID as GID:ID. An example is 119:3.
    - **Action** - Select an action for when the source IP/CIDR list or rule ID list is triggered on. Note that if a rule is suppressed, no action is taken and no logs are sent or captured.
        - **Allow Log** - Allow the requests and log an event.
        - **Allow No Log** - Allow the requests and do not log an event.
        - **Deny Log** - Deny the requests and log an event.
        - **Deny No Log** - Deny the requests and do not log an event.

What to do next

Attach the profile to a policy rule set. See Rule Sets and Rule Set Groups on page 133 for more information.

Data Loss Prevention (DLP) Profile

Learn how a Data Loss Prevention profile detects sensitive data in egress traffic and supports compliance protection with predefined or custom patterns.

The DLP (Data Loss Prevention) profile provides Multicloud Defense customers with the ability to specify policy rules to detect and take action upon finding exfiltration patterns in the data when the Multicloud Defense solution is deployed in the forward proxy (egress) mode.

Multicloud Defense allows customers to specify common pre-packaged data patterns such as Social Security Numbers (SSN), AWS secrets, credit card numbers etc., in addition to custom PCRE based regular expression patterns. This makes it easy to enforce protections for PCI, PII, and PHI data to meet compliance requirements. This feature is integrated with the existing Multicloud Defense feature set requiring no separate DLP services.

Create a Data Loss Prevention Profile

Create a Data Loss Prevention profile with predefined or custom patterns, match thresholds, and actions to protect sensitive data in egress traffic.

Procedure

---

1. Navigate to **Policies** > **Profiles** > **Data Loss Prevention**.
2. Click **Create Intrusion Profile**.
3. Select **Data Loss Prevention**.
4. Provide a unique **Name** and enter a description for the profile.
5. Enter the **DLP Filter List** in the table.
6. Click **Add** to insert more rows as needed.

7. Provide a **Description** for the filter.
8. Choose a predefined static pattern (e.g CVE Number) from the dropdown list or provide a custom Regular expression.
9. Provide a **count** to define the number of times the pattern must be seen in the traffic.
10. Select an **Action** to take if the pattern matches the count number of times.

> **Note**
>
> There are cases where the pre-defined pattern for AWS Access Key and AWS Secret Key doesn't match in DLP inspection due to pattern being more restrictive. Use the following relaxed custom pattern in DLP profile to detect AWS Access Key and AWS Secret Key. Be aware that this could generate false positives log events.
>
> `AWS Access Key: (?<![A-Z0-9])[A-Z0-9]{20}(?![A-Z0-9])`
>
> `AWS Secret Key: (?<![AZa-z0-9/+=])[A-Za-z0-9/+=]{40}(?![A-Za-z0-9/+=])`

What to do next

Attach the profile to a policy rule set. See Rule Sets and Rule Set Groups on page 133 for more information.

Anti-Malware Profile

Learn how an Anti-Malware Profile scans incoming data for viruses, ransomware, and other malicious threats before they can infect protected systems.

An anti-malware profile prevents malware attacks by scanning all incoming data to prevent malware from being installed and infecting a computer. Anti-malware programs can also detect advanced forms of malware and offer protection against ransomware attacks. Currently, the Talos ClamAV virus detection engine is a large portion of the profile. ClamAV® is an antivirus engine for detecting trojans, viruses, malware and other malicious threats.

The current version of ClamAV in use is 1.4.

If you opt to create an anti-malware profile, we **strongly** recommend immediately adding it to a policy by configuring it to a rule.

Create an Anti-Malware Profile

Create an Anti-Malware Profile, choose how Cisco Talos ruleset updates are applied, and select the action for detected virus signatures.

Procedure

1. Navigate to **Policies** > **Profiles** > **Anti Malware**.
2. Select **Anti-malware**.
3. Provide a unique **Name** and enter a description.
4. Select one of the following modes for Talos ruleset:
   - **Manual Mode** - select the Talos Ruleset Version from dropdown. The selected ruleset version is used by the Multicloud Defense datapath engine on all Gateways which use this profile and is not automatically updated to newer ruleset versions.

- Automatic Mode - select how many days to delay the deployment by, after the ruleset version is published by Multicloud Defense. New rulesets are published daily by Multicloud Defense and the gateways using this profile are automatically updated to the latest ruleset version which is **N** days or older, where **N** is the "delay by days" argument selected from the dropdown. For example, if you select to delay the deployment by 5 days on Jan 10, 2024, the Multicloud Defense Controller will select a ruleset version which was published on Jan 5th or before. Note that Multicloud Defense may not publish on certain days if internal testing with that ruleset version fails.

5. Select the desired **Action** to take when a match for a virus signature is found.

What to do next

Attach the profile to a policy rule set. See Rule Sets and Rule Set Groups on page 133 for more information.

Web Application Firewall (WAF) Profile

Learn how a Web Application Firewall profile uses core, premium, or custom rules to detect and block common attacks against web applications.

Web protection profiles are a collection of Web Application Firewall (WAF) rules that can detect and block known web application attacks. You can configure WAF profiles to use signatures and constraints to examine web traffic. You can also enforce an HTTP method policy, which controls the HTTP method that matches the specified pattern. It typically protects web applications from attacks such as cross-site forgery, cross-site-scripting (XSS), file inclusion, and SQL injection, among others.

Table 3: Supported WAF rule sets

Rulesets	Description
Core Rules	The core rules are a standard set of rules from ModSecurity CRS (Core Rule Set) that provide a base level of protection for any web application.
Trustwave Rules	The Trustwave rules are a premium set of rules from ModSecurity based on intelligence gathered from real-world investigations, penetration tests and research that provide an advanced level of protection for specific web applications and frameworks.
Custom Rules	The custom rules are a particular set of rules written by customers that provide a specialized level of protection for custom web applications.

Note that the WAF profile does not include malicious IPs. See Malicious IP Profile on page 202 and Network Intrusion (IDS/IPS) Profile on page 190 for more information.

Create WAF Profile

Create a WAF profile, select rulesets and response actions, configure optional threat records, and attach the profile to a security policy ruleset.

Use the following procedure to create a WAF profile.

Note

If core Rulesets are specified, the core rules cannot be disabled. In order to disable the core rules, remove all core rulesets from the WAF profile so they will not be evaluated.

Procedure

1. Navigate to **Policies** > **Profiles** > **WAF**.
2. Click **Create**.
3. Specify the following general settings:
   a) Enter a unique **Profile Name**.
   b) (Optional) Enter a **Description**. This may help differentiate between profiles with a similar name.
   c) Specify the action:
      - **Rule Default** - Allow or deny the requests based on the action specified in each triggered rule and log an event.
      - **Allow Log** - Allow the requests and log an event.
      - **Deny Log** - Deny the requests and log an event.
   d) Specify whether to generate a Threat HAR file if the WAF profile detects malicious activity. The gateway should have a Pcap profile attached, for this to work.
   e) Specify whether to generate a HTTP Request HAR file if the WAF profile detects malicious activity.
   f) In the **RULE SETS** section, in the vertical tab located to the left, click **Core Rules**. You must specify at least one ruleset from a rules library (Core, Trustwave, Custom):
      - Specify the following:
        - **Manual** - Specify the core rules version to use.
        - **Automatic** - Specify the numbers of days from publish date to delay automatic update to the latest core rules version.
      - Identify the rules you want to add to the profile and click **Add to Profile**. The selections appear in the table located to the right.
   g) In the vertical tab located to the left, click **Trustwave Rules**.
      - Specify the following:
        - **Disabled** - Specify whether to disable the use of Trustwave rules.
        - **Manual** - Specify the Trustwave rules version to use.
        - **Automatic** - Specify the number of days from publish date to delay automatic update to the latest Trustwave rules version.
      - Identify the rules you want to add to the profile and click **Add to Profile**. The selections appear in the **Profile Selections** table located to the right.
   h) In the vertical tab located to the left, click **Custom Rules**.
      - Specify one of the following options:

- **Disabled** - Specify whether to disable the use of custom rules.
- **Manual** - Specify the custom rules version to use.
- **Automatic** - Specify the number of days from publish date to delay automatic update to the latest custom rules version.

- Identify the rules you want to add to the profile and click **Add to Profile**. The selections appear in the **Profile Selections**table located to the right.

4. Scroll to the top of the window and click the **Advanced Settings** tab:
   a) Under "Rule Suppression", click **Add** to add one or more rows for rules. Rules can be suppressed for a specific IP or a list of CIDRs:
      - For **Source IP/CIDR List**, provide a comma-separated list of IPs or CIDRs.
      - For **Rule ID List**, provide a comma-separated list of rule IDs.
   b) Under "Event Filtering" provide the following information:
      - **Type** - **Rate** or **Sample**
      - **Number of Events**
      - **Time (Seconds)**
   c) Under "Rule Event Filtering" click **Add** to add one or more rows for rules. For every new row you create, enter a valid **Rule ID List**, **Number of Events**, **Time (Sec)**, and choose either Type or Sample as the **Type**.
   d) Under "Core Rule Set", select a value for both the **Request Anomaly** and **Response Anomaly**. Note that using a value less than 3 for the "Request Anomaly" results in a huge volume of alerts.
   e) Select the **Paranoia Level**. Your options range from 1-4.
5. Click **Save**.

What to do next

Attach the profile to a policy rule set. See Rule Sets and Rule Set Groups on page 133 for more information.

Event Filtering

Configure WAF event filtering to rate-limit or sample security events at the profile or rule level without changing web threat detection or protection.

To reduce the number of security events that are generated when the WAF Profile is triggered, the Event Filtering under **Advanced Settings** can be configured to rate limit or sample the events. The configuration does not alter the detection or protection behavior.

When specifying Type as **Rate**, the generated events are rate limited based on the specified *Number of Events* triggered over a *Time* evaluation interval (in seconds). For example, if *Number of Events* is specified as 50 and *Time* is specified as 5 seconds, only 10 events per second will be generated.

When specifying Type as **Sample**, the generated events are sampled based on the specified *Number of Events*. For example, if *Number of Events* is specified as 10, only 1 event will be generated for every 10 events triggered.

Profile Event Filtering

Profile Event Filtering applies to all rules that are configured in the WAF Profile:

- Specify the Type as **Rate** or **Sample**:

- **Rate**- Specify the *Number of Events* and the *Time* evaluation interval (in seconds).
- **Sample**- Specify the *Number of Events*.

Rule Event Filtering

To reduce the number of security events that are generated when the WAF profile is triggered, event filtering can be configured to rate limit or sample the events. The configuration does not alter the detection or protection behavior.

Rule event filtering applies to specific rules that are configured in the WAF profile.

Procedure

1. Click **Add** under Rule Event Filtering.
2. For **Rule ID List**, specify a comma-separated list of **Rule IDs**.
3. Specify Type as **Rate** or **Sample**.
   - **Rate**- Specify the **Number of Events** and the **Time** evaluation interval (in seconds).
   - **Sample**- Specify the **Number of Events**.

What to do next

Associate WAF Profile with a Policy Rule

URL (Uniform Resource Locator) Filter Profile

Learn how a URL filtering profile lets you allow or deny HTTP requests by URL, category, and method when traffic uses a Forward Proxy rule.

A URL filtering profile evaluates the URL of an HTTP request and applies an action to either allow or deny the traffic. In order to evaluate the URL, the traffic must be processed by a **Forward Proxy** rule. The set of URLs in the profile can be specified as strings representing the full path or as strings representing a Perl Compatible Regular Expression (PCRE). If only domain filtering is required, it is best to use an FQDN filtering profile. An FQDN filtering profile can also be used in conjunction with URL filtering, where the domain is evaluated using the FQDN filtering profile and the URL is evaluated using the URL filtering profile.

The URL filtering profile can use a set of pre-defined categories. To view more information on categories, please see FQDN / URL Filtering Categories on page 220.

Note

The URL filtering is organized as a table containing user-specified rows (URLs and Categories) along with two default rows (**Uncategorized** and **ANY**). Categories and URLs can be combined within each row if desired.

The limits for each URL filtering profile are as follows:

- Maximum user-specified rows: 254 (Standalone or a group of standalones)
- Maximum Categories and URLs per row: 60
- Maximum URL character length: 2048

When specifying a multi-level domain (e.g., `www.example.com`), it's important to escape the `.` character (e.g., `www\.example\.com`) otherwise it will be treated as a wildcard for any single character

Uncategorized

- The penultimate row in a URL filtering profile, which is represented as **Uncategorized**.
- Specifies the policy action to take for URLs that do not match the user-specified URLS or do not have a category.
- If a standalone profile is used in a group profile and the group profile is applied to a policy ruleset, the **Uncategorized** row will be taken from the group profile. The **Uncategorized** row of a standalone profile is only applicable if the standalone profile is directly applied to a policy ruleset.

Default (ANY)

- The final row in a URL filtering profile, which is represented as **ANY**.
- Specifies the policy action to take for URLs that do not match the user-specified URLs or categories, or are not uncategorized.
- If a standalone profile is used in a group profile and the group profile is applied to a policy ruleset, the **ANY** row will be taken from the group profile. The **ANY** row of a standalone profile is only applicable if the standalone profile is directly applied to a policy ruleset.

Create the URL Filtering Profile

Create a URL filtering profile to allow or deny HTTP requests by specified URLs, categories, and methods, with optional event logging.

Use the following procedure to create a standalone URL filtering profile:

Procedure

1. Navigate to **Policies** > **Profiles** > **URL Filtering**.
2. Click **Create**.
3. Provide a unique **Name**.
4. (Optional) Enter a **Description**. This may help differenatiate between other profiles with similar names.
5. Click **Add** to create a new row.
6. Specify individual URLs (e.g., https://www.google.com):
    - Each URL is specified as a PCRE (Perl Compatible Regular Expression).

- Each URL must be specified as a full path.
- Consider escaping the decimal "." character else it will be treated as a single character wildcard.

7. Specify **Categories** (e.g., Gambling, Sports, Social Networking).
8. Specify the HTTP methods to which the policy is applied.
9. Select on of the following as a subset of methods:
   - Delete
   - Get
   - Head
   - Options
   - Patch
   - Post
   - Put
10. Specify **All** for all methods.
11. Specify the policy **Action** for the user-specified URLs/Categories, Uncategorized and ANY rows:
    - **Allow Log** - Allow the requests and log an event.
    - **Allow No Log** - Allow the requests and do not log an event.
    - **Deny Log** - Deny the requests and log an event.
    - **Deny No Log** - Deny the requests and do not log an event.
12. Specify the **Return Status Code**.
13. Specify an integer value **greater than or equal to 100 and less than 600**. The value represents the HTTP status that will be returned to the client making the request. A common return code is **503**.
14. Click **Save**.

What to do next

Attach the profile to a policy rule set. See Rule Sets and Rule Set Groups on page 133 for more information.

Fully Qualified Domain Name Filter Profile

Learn how an FQDN filter profile lets you allow or deny encrypted traffic for specified domains and categories in Forwarding or Forward Proxy rules.

A Fully Qualified Domain Name (FQDN) filter profile evaluates the FQDN associated with traffic and applies an action to either allow or deny the traffic. In order to evaluate the FQDN, traffic must be TLS encrypted and contain an FQDN in the SNI field of a TLS Hello header. The FQDN can be evaluated for traffic that is processed by either a **Forwarding** or **Forward Proxy** rule. The set of FQDNs in the profile can be specified as strings representing the full domain or as strings represented by a Perl Compatible Regular Expression (PCRE). If only domain allowlisting is required, it is best to use an FQDN filtering profile. An FQDN filtering profile can also be used in conjunction with a URL filtering profile, where the domain is evaluated using the FQDN filtering profile and the URL is evaluated using the URL filtering profile.

Use FQDN filtering to filter categories that you want to allow or deny based on criteria, after a rule match. You can set filters at a granular level. The FQDN filter rows contain log-related actions such as deny or allow that you can use.

The FQDN filtering profile can also use a set of pre-defined categories. To view more information on categories, see FQDN / URL Filtering Categories on page 220.

Note

The FQDN filtering profile is organized as a table containing user-specified rows (FQDNs and categories) along with two default rows (Uncategorized and ANY). Categories and FQDNs can be combined within each row if desired.

The limits for each FQDN filter profile are as follows:

- Maximum user-specified rows: 254 (standalone or group of standalones)
- Maximum categories and FQDNs per row: 60
- Maximum FQDN character length: 255

When specifying a multi-level domain (e.g., 'www.example.com'), it's important to escape the `.` character (e.g.,`www\.example\.com`) otherwise it will be treated as a wildcard for any single character.

Standalone vs. Group

A FQDN filter profile can be specified as standalone or group.

A standalone FQDN filter profile contains FQDNs and categories. The profile will be applied directly to a set of one or more policy rulesets or associated with a FQDN group profile.

A FQDN filter group profile contains an ordered list of standalone profiles that can be defined for different purposes and combined together into a group profile. The group profile can be applied directly to a set of one or more policy rulesets. Each team can create and manage specific standalone profiles. These standalone profiles can be combined together into a group profile to create hierarchies or different combinations based on the use case. An example combination could be a global FQDN list that would apply to everything, a CSP-specific list that would apply to each different CSP, and an application-specific list that would apply to each different application.

Uncategorized

- The second-to-last row in an FQDN filter profile which is represented as **Uncategorized**.
- Specifies the policy action to take for FQDNs that do not match the user-specified FQDNs or do not have a category.
- If a standalone profile is used in a group profile and the group profile is applied to a policy ruleset, the **Uncategorized** row will be taken from the group profile. The **Uncategorized** row of a standalone profile is only applicable if the standalone profile is directly applied to a policy ruleset.

Default (ANY)

- The final row in an FQDN filter profile, which is represented as **ANY**.
- Specifies the policy action to take for FQDNs that do not match the user-specified FQDNs or categories, or are not **Uncategorized**.
- If a standalone profile is used in a group profile and the group profile is applied to a policy ruleset, the **ANY** row will be taken from the group profile. The **ANY** row of a standalone profile is only applicable if the standalone profile is directly applied to a policy ruleset.

Create a Standalone FQDN Filter Profile

Create a standalone FQDN filter profile to allow or deny domain traffic by specified FQDNs and categories, with optional logging and decryption exceptions.

Use the following procedure to create a standalone FQDN filter profile:

Procedure

1. Navigate to **Policies** > **Profiles** > **FQDN Filtering**.
2. Click **Create**.
3. Provide a unique **Name**.
4. (Optional) Enter a **Description**. This may help differentiate between profiles with a similar name.
5. Specify the Type as **Standalone**.
6. Click **Add** to create a new row.
7. Specify individual FQDNs (for example, google.com).
   a) Each FQDN is specified as a PCRE (Perl Compatible Regular Expression).
   b) Consider escaping the "**.**" character else it will be treated as a single character wildcard.
8. Specify a **Category** (for example, Gambling, Sports, Social Networking).
9. Specify the policy **Action** for the user-specified FQDNs/Categories, Uncategorized and ANY rows.
   - **Allow Log** - Allow the requests and log an event.
   - **Allow No Log** - Allow the requests and do not log an event.
   - **Deny Log** - Deny the requests and log an event.
   - **Deny No Log** - Deny the requests and do not log an event.
10. (Optional) Specify **Decryption Exception** for any FQDNs where decryption is not desired or possible. Possible reasons for considering decryption exception include:
    - Desire to not inspect encrypted traffic (for example, financial services, defense, health care, etc.).
    - SSO authentication traffic where decryption is not possible.
    - NTLM traffic that cannot be proxied.
11. Click **Save** when completed.

What to do next

Attach the profile to a policy rule set. See Rule Sets and Rule Set Groups on page 133 for more information.

Create a Group FQDN Filter Profile

Create a group FQDN filter profile that combines standalone profiles and defines actions or decryption exceptions for uncategorized and other domain traffic.

Use the following procedure to create a group FQDN filter profile with at least two standalone profiles:

Procedure

1. Navigate to **Policies** > **Profiles** > **FQDN Filtering**.

2. Click **Create**.
3. Provide a unique **Name**.
4. (Optional) Enter a **Description**. This may help differentiate between profiles that may have a similar name.
5. Specify the Type as **Group**.
6. Select an initial standalone profile (at least one standalone profile is required).
7. Click **Add FQDN Profile** to create a new row for additional profiles.
8. Select a standalone profile.
9. Specify the policy **Action** for uncategorized FQDNs.
10. Specify the policy **Action** for **ANY** FQDNs (default).
11. (Optional) Specify the **Decryption Exception** for uncategorized or ANY if decryption is not desired or possible. Possible reasons for considering decryption exception include:
    - Desire to not inspect encrypted traffic (financial services, defense, health care, etc.).
    - SSO authentication traffic where decryption is not possible.
    - NTLM traffic that cannot be proxied.
12. Click **Save**.

What to do next

Attach the profile to a policy rule set. See Rule Sets and Rule Set Groups on page 133 for more information.

Malicious IP Profile

Learn how a Malicious IP Profile helps protect cloud traffic from known malicious addresses identified by Cisco Talos threat intelligence.

Additional security protections can be enabled to prevent communication from and to known malicious IPs. These malicious IPs are defined and powered by Cisco Talos, and integrated into Multicloud Defense as a security profile ruleset. The ruleset is updated frequently as updates are made available by Talos. The updates can be either dynamically or manually applied to a policy ruleset using the automatic update configuration or manual update configuration. For more information, see Create a Malicious IP Profile on page 203.

> **Note**
>
> Malicious IP are identified by Talos based on various learned behavior:
>
> - Malicious attackers identified from web honeypots
> - Botnet C&C (command and control) hosts
> - TOR exit nodes
> - Other learned behavior.

Create a Malicious IP Profile

Create a Malicious IP Profile, enable IP Reputation, and choose how Cisco Talos ruleset updates are applied to associated security policies.

Use the following procedure to create a malicious IP profile:

Procedure

1. Navigate to **Policies** > **Profiles** > **Malicious IPs**.
2. Click **Create**.
3. Provide a unique **Profile Name**.
4. (Optional) Enter a **Description**. This can help differentiate between other profiles with similar names.
5. Check the box to enable **IP Reputation**.
6. Choose a ruleset version from the drop-down menu:
   - **Manual** - The selected ruleset version is used by the Multicloud Defense datapath engine on all gateways which use this profile. The profile will not be automatically updated to newer ruleset versions.
   - **Automatic** - Select the number of days to delay the update, after the ruleset version is published by Multicloud Defense. New rulesets are published frequently by Multicloud Defense. The gateways using this profile are automatically updated to the latest ruleset version which is **N** days or older, where **N** is the "delay by days" argument selected from the dropdown. For example, if you select to delay the deployment by 5 days on Jan 10, 2021, the Multicloud Defense Controller will select a ruleset version which was published on Jan 5th or before. Note that Multicloud Defense may not publish on certain days if internal testing with that ruleset version fails.
7. Click **Save**.

What to do next

Attach the profile to a policy rule set. See Rule Sets and Rule Set Groups on page 133 for more information.

IP Reputation

Learn how the IP Reputation setting enables malicious IP protection when a Malicious IP Profile is attached to a security policy ruleset.

The IP reputation checkbox is used as a means to **enable** or **disable** the profile. When checked and the profile is attached to a policy ruleset, malicious IP protection will be enforced. When unchecked and the profile is attached to policy

rules, malicious IP protection will not be enforced. Our recommendation is to always check the IP reputation checkbox. If you want to disable the malicious IP profile, then remove its association from the policy rules rather than uncheck the checkbox.

22 Gateway Profiles

Topics:

- Packet Capture Profile
- Log Forwarding Profile
- Gateway Metrics Forwarding Profile
- (Preview Only) Network Packet Broker Profile
- Network Time Protocol Profile
- IPSec Profile
- BGP Profile

Learn about gateway profiles and how they define security settings for Multicloud Defense gateways.

A gateway profile is typically associated with the configuration of a network gateway by way of a device that connects different networks and routes traffic between them. Gateway profiles are used to manage the behavior and functionality of network gateways, ensuring efficient and secure communication between different parts of the network. These profiles generally apply to the following protective methods:

- Routing Policies
- Network Address Translation (NAT)
- Virtual Private Network (VPN) Settings
- Quality of Service (QoS)
- Authentication and Access Control

These profiles are generally applied to either a Multicloud Defense Gateway or a VPN tunnel that is associated with a gateway.

Packet Capture Profile

Learn how a Packet Capture Profile records cloud network packets for security investigations, incident reconstruction, and troubleshooting connectivity, latency, or packet loss.

Packet Capture (PCAP) captures data packets that are transmitted across the network, allowing for detailed analysis of the network traffic. PCAP can be used to monitor network traffic for signs of malicious activity by analyzing the captured packets, security systems can detect and respond to potential threats in real-time and allows you to reconstruct the sequence of events leading up to the incident and identify the source and nature of the attack. This information can be helpful in diagnosing a timeline or to troubleshoot events such as connectivity problems, latency, and packet loss.

Create a Packet Capture Profile

Create a Packet Capture Profile and select the cloud account and storage destination for packet captures generated by associated gateways.

Use the following procedure to create a pack capture profile:

Procedure

1. Navigate to **Infrastructure** > **Profiles** > **Packet Capture**.
2. Click **Create**.
3. Specify a unique **Name**.
4. (Optional) Enter a **Description**. This may help differentiate between other profiles with a similar name.
5. Specify a **CSP Account**.
6. The type of cloud service provider may determine the parameters for the storage bucket. Be aware of the following requirments per cloud service provider:
   - **AWS** - S3 Bucket.
   - **Azure** - Storage Account Name, Blog Container , and Storage Access Key.
   - **GCP** - Storage Bucket.
7. Click **Save**.

What to do next

Attach the profile to a policy rule set. See Rule Sets and Rule Set Groups on page 133 for more information.

Log Forwarding Profile

Learn how a Log Forwarding Profile sends gateway, VPC, and VNet logs to one or more third-party destinations for monitoring and analysis.

A log forwarding profile allows you to send a collection of gateway, VPC, and VNet logs to a third party. The communication between Multicloud Defense and the third party of your choice contains the log type that needs to be forwarded and the destination server profiles the logs will be sent to. You can have a single profile or a profile group that sends logs to multiple endpoints simultaneously.

Note that this profile does not include metrics. See Gateway Metrics Forwarding Profile on page 208 for more information about forwarding log metrics.

Create a Standalone Log Forwarding Profile

Create a standalone Log Forwarding Profile and configure a secure third-party monitoring destination for gateway, VPC, or VNet security logs.

Use the following procedure to create a standalone profile to forward logs with:

Procedure

1. Navigate to **Infrastructure** > **Profiles** > **Log Forwarding**.
2. Click **Create**.
3. Enter a unique **Profile Name**.
4. (Optional) Enter a **Description**. This may help differentiate from other profiles with a similar name.
5. Expand the **Type** drop-down menu and select **Standalone**.
6. Expand the **Destination** drop-down menu and select the third-party application to send logs to.
7. Based on the type of destination you select in step 7, enter the appropriate information when prompted to secure the final endpoint where the logs are forwarded to. Note that not all options are available based on the type of destination.
8. Click **Save**.

What to do next

Attach the profile to a policy rule set. See Rule Sets and Rule Set Groups on page 133 for more information.

Create a Log Forwarding Group

Create a Log Forwarding Group that combines standalone profiles to send cloud and gateway security logs to multiple third-party monitoring destinations.

Use the following procedure to create a profile group to forward logs with:

Before you begin

- You must have at least one third party application to forward the metric to prior to creating this profile.
- You must have at least two standalone metrics forwarding profiles already created. See Create a Standalone Log Forwarding Profile on page 207 for more information.

Procedure

1. Navigate to **Infrastructure** > **Profiles** > **Log Forwarding**.
2. Click **Create**.
3. Enter a unique **Profile Name**.
4. (Optional) Enter a **Description**. This may help differentiate from other profiles with a similar name.
5. Expand the **Type** drop-down menu and select **Group**.
6. Under **Group Details**, click **Add** for every new row you need to add to the profile.
7. Expand the drop-down menus for each row to select a profile to add to the group. If you want to remove a profile at any point prior to saving, select the profile's checkbox so it is highlighted and select **Remove**.
8. Click **Save**.

What to do next

Attach the profile to a policy rule set. See Rule Sets and Rule Set Groups on page 133 for more information.

Gateway Metrics Forwarding Profile

Explains the purpose and capabilities of the Gateway Metrics Forwarding Profile.

This profile is intended to forward gateway metrics generated by the Multicloud Defense Gateway for data monitoring and analysis.While the metrics are generated by the gateway, it is the Multicloud Defense Controller that forwards the metrics to the third party analysis application. This forwarding profile allows you to monitor, analyze, and organize your gateway metrics without logging into Multicloud Defense. Use this information to gauge the performance and behavior of your gateway environment; you can also utilize this information for environmental troubleshooting.

> **Note**
>
> As of Multicloud Defense Controller version 23.09, only Datadog is supported as a third-party analytics application.

For the majority of analytics applications available, for example, Datadog, you must already be an authorized user to access the tool's APIs and rendered data.

Create a Standalone Metrics Forwarding Profile

Configure and establish a profile for metrics forwarding to a third-party application.

Create a standalone profile and forward metrics to be processed by a third party.

Before you begin

You must have at least one third-party application to forward the metrics to, prior to creating this profile.

Procedure

1. Navigate to **Infrastructure** > **Profiles** > **Metrics Forwarding**.
2. Click **Create**.
3. Enter a unique profile **Name**.
4. (Optional) Enter a **Description**. This may help differentiate from other profiles with a similar name.
5. Expand the **Type** drop-down menu and select **Standalone**.
6. Expand the **Destination** drop-down menu and select the third-party application to process and analyze the metrics.
7. Enter the **Endpoint** to be used as the endpoint location for the metrics.
8. (Optional) Send a **Test** message to the destination. Enter the content of the test message in the text field when prompted and clikc **Validate**. If the destination does not receive the test message, confirm the destination type and configuration in step 7.
9. Click **Save**.

   If you select Datadog as your analytics application, the **Endpoint** is filled in by default with an HTTPS webhook. This entry, if defaulted, can be modified prior to saving the profile.

What to do next

Attach the profile to a policy rule set. See Rule Sets and Rule Set Groups on page 133 for more information.

Create a Group Metrics Forwarding Profile

Configure a profile that aggregates multiple standalone metrics profiles to be assigned to a single gateway.

In this process, you create a profile and then assign it to a specific gateway. A group profile combines up to five standalone metrics forwarding profile that can then be assigned to a single gateway. Use the following procedure to create a grouped metrics forward profile:

Before you begin

- You must have at least one third party application to forward the metric to prior to creating this profile.
- You must have at least two standalone metrics forwarding profiles already created. See Create a Standalone Metrics Forwarding Profile on page 208 for more information.

Procedure

1. In the Multicloud Defense Controller interface navigate to **Infrastructure** > **Profiles** > **Metrics Forwarding**.
2. Click **Create**.
3. Enter a unique **Profile Name**
4. (Optional) Enter a **Description**. This may help differentiate between profiles with a similar name.
5. Expand the **Type** drop-down menu and select **Group**.
6. Under **Group Details**, click **Add** for every new row you need to add to the profile.
7. Expand the drop-down menus for each row to select a profile to add to the group. If you want to remove a profile at any point prior to saving, select the profile's checkbox so it is highlighted and select **Remove**.
8. Click **Save**.

What to do next

Attach the profile to a policy rule set. See Rule Sets and Rule Set Groups on page 133 for more information.

(Preview Only) Network Packet Broker Profile

Learn how a Network Packet Broker Profile directs selected network traffic to monitoring, security, and performance tools for inspection and analysis.

A Network Packet Broker (NPB) profile is a configuration set within a Network Packet Broker device that defines how network traffic should be managed, and directed; traffic is filtered based on various criteria such as IP addresses, protocols, ports, and application types. These profiles are specialized devices used to optimize the flow of network traffic to various monitoring, security, and performance management tools.

Create Network Packet Broker Profile

Create a Network Packet Broker Profile to direct selected network traffic to a preferred monitoring destination in the required capture format.

Use the following procedure to create a Network Packet Broker (NPB) profile:

Procedure

1. Navigate to **Policies** > **Profiles** > **Network Packet Broker**.
2. Click **Create**.
3. Provide a unique **Name**.
4. (Optional) Enter a **Description**. This may help differenatiate between other profiles with similar names.
5. Under **Destinations** expand the drop-down menu and select the preferred destination for your NPB profile should direct matching traffic to.
6. Define th capture format to ensure that network data is effectively captured, stored, and analyzed, leading to better network management and troubleshooting capabilities. Expand the **Capture Format** drop-down menu and select one of the following formats:
   - **None** - Select this option if you have limited resources or operate within a highly dynamic environment that utilizes multiple formats. You can also select this option if you intend to change this field later on.
   - **VXLAN** - Virtual Extensible LAN, is not a capture format itself but rather a network virtualization technology. It extends Layer 2 networks over a Layer 3 infrastructure, enabling the creation of a virtualized network overlay and contains Ethernet frames within UDP packets, allowing them to be transmitted over IP networks
   - **Netflow** - Select this option if your environment is configured to export flow records to a designated collector or analysis tool, rather than capturing raw packets in formats like PCAP.
   - **ERSPAN** - Select this protocol to mirror traffic from a source to a destination over a network and extend the capabilities of traditional SPAN (Switched Port Analyzer) by containing mirrored traffic into GRE packets that are then sent over Layer 3 networks. This is ideal for environments where network traffic needs to be monitored from a central location.
7. Expand the **Slicing** drop-down menu and select how each network packet is sectioned and captured, rather than the entire packet. If you opt to configure this, be sure to also configure the following options:

   Offset value - This field allows you to configure the number of bytes that are skipped from the start of each packet before beginning to capture data. This is set to **4** by default.

   Strip Encrypted Payload - Check this option to remove an encrypted portion of a packet's data payload during capture or analysis.Enable this to concentrate on metadata and network behavior while respecting privacy and optimizing resource use.
8. Click **Save**.

What to do next

Attach the profile to a policy rule set. See Rule Sets and Rule Set Groups on page 133 for more information.

Network Time Protocol Profile

Learn how a Network Time Protocol Profile synchronizes gateway clocks to support accurate event correlation, monitoring, troubleshooting, and secure network operations.

Network Time Protocol synchronizes computer clocks to each other and to international standards via telephone modem, radio and satellite. As a profile, especially within distributed systems, synchronized time is essential for coordinating actions and ensuring that distributed processes work together seamlessly. Consistent time across devices is ideal in network management tasks, such as monitoring and troubleshooting. It ensures that logs from different devices can be correlated accurately and ensures the smooth and secure operation of the network.

Create a Profile

Create a Network Time Protocol profile with the NTP servers that gateways use to maintain synchronized time across your cloud network.

Use the following procedure to create an NTP profile:

Procedure

1. Navigate to **Infrastructure** > **Network** > **NTP**.
2. Click **Create**.
3. Specify a unique **Name**.
4. (Optional) Enter a **Description**. This may help differentiate between other profiles with a similar name.
5. Specify the **List** of NTP servers.
6. Click **Save**.

What to do next

Attach the profile to a policy rule set. See Rule Sets and Rule Set Groups on page 133 for more information.

IPSec Profile

Learn how an IPSec Profile defines secure connection settings for site-to-site VPN tunnels and protects data exchanged between network peers.

The use of Internet Protocol Security (IPSec) profiles for a virtual tunnel interface can simplify the configuration process when you need to provide protection for remote access. An IPSec profile contains the required security protocols and algorithms required to ensure a secure, logical communication path between two site-to-site VPN peers. t is a required component when creating a tunnel as the VPN depends on IPsec tunnels for network-to-network, host-to-network and host-to-host communications. The IPSec profile allows you to configure both IKE and IPSEC parameters in one place for additional security and encryption protection.

If you choose to include an IPSec profile within your site-to-site tunnel configuration, the profile provides robust network security by encrypting and authenticating data as it travels between points on the network as well as the flexibility of being compatible with site-to-site, client-to-site, and client-to-client tunnels.

Create an IPSec Profile

Create an IPSec Profile with the IKE and IPSec authentication, encryption, key lifetime, and integrity settings required for a secure VPN tunnel.

Use the following procedure to create an IPSec profile from the Multicloud Defense Controller dashboard:

Procedure

1. Navigate to **Infrastructure** > **Network** > **IPSec**.
2. Click **Create**.
3. Enter a unique **Profile Name**.
4. (Optional) Enter a **Description**. This may help differentiate from other profiles with a similar name.
5. Enter the appropriate IKE information when prompted:

a) **DH Group** - Diffie-Hellman (DH) groups determine the strength of the key used in the key exchange process. Expand the drop-down menu to select the appropriate groups for the profile.
b) **Authentication** - Expand the drop-down menu to select the types of authentication you want for this tunnel.
c) **Encryption** - Intercepted stacks require encrypting and decrypting. Expand the drop-down menu to select your method of encryption.
d) **Hash** - SHA1 is a one-way hashing algorithm that produces a 160-bit digest. Use the drop-down menu to select the appropriate option.
e) **Key Lifetime** - Enter a time value in seconds for how long the key lasts. Avaialble values are between 60 sec and 86400 sec.
f) **IKE Version** -The Internet Key Exchange (IKE) is a protocol that is used to set up a security association in the IPSec protocol suite that provides robust authentication and encryption of IP packets. Use the drop-down menu to select either IKE version 1 or version 2. There are significant differences between the versions so be sure to select the one most appropriate for your environment.

6. Enter the appropriate IPsec information when prompted:
   a) **Authentication** - Expand the drop-down menu to select an authentication method: None, SHA256, SHA, or Null.
   b) **Encryption** - Expand the drop-down and select a type of key: AES GCM 256, AES GCM 192, or AES GCM. This generates a unique key exchange between the connected devices, so that each device can decrypt the other device's messages.
   c) **Mode** - Expand the drop-down menu to select the IPSec policy authentication protocol. You can select more than one.

What to do next

Attach the profile to a policy rule set. See Rule Sets and Rule Set Groups on page 133 for more information.

BGP Profile

Learn how a BGP profile supports dynamic route exchange between Multicloud Defense gateways and neighboring routers for site-to-site VPN connections.

Border Gateway Protocol (BGP) is an Internet Engineering Task Force (IETF) standard, and the most scalable of all routing protocols. BGP is the routing protocol for the global Internet and service provider private networks. BGP enables the VPN gateways and your BGP neighbors to exchange routes that inform the gateways on both sides of the connectoin of the availability of the gateways or routers involved.

You **must** create and add a BGP profile to your gateway if you are establishing a site-to-site VPN tunnel connection to another platform or device. Deploying with a BGP profile depoys a gateway that uses dynamic routing with BGP between your networks and cloud service providers.

BGP Neighbors and Path Selection

BGP profiles utilize a property called "neighbors"; a neighbor refers to another BGP router with which a BGP session is established. The purpose of configuring neighbors in a BGP profile is to facilitate the exchange of routing information between autonomous systems (ASes) or within a singular AS.

Important

We **strongly** recommend adding at least one nighbor to your BGP profile.

Within the neighbor section of the BGP profile, you have the chance to opt to **Route Map In** or **Route Map Out**. The route maps provide us a mechanism to only advertise (outbound) or accept (inbound), based on whats identified in that route map.

Allowing a route map **in** enables the following actions:

- **Incoming Route Filtering:** Control which routes are accepted from a BGP neighbor. Filter out unwanted routes to optimize the routing table to ensure only the relevant routes are considered.
- **Attribute Modification**: Adjust attributes of incoming routes, such as the local preference or metric, to influence the path selection process within your network. This helps prioritize certain routes over others based on your deployed network policies.
- **Security and Policy Compliance**: Prevent routes that do not comply with your network policies from being accepted to enhance security and ensure policy compliance.

Adversely, allowing a route map **out** enables the following actions:

- **Outgoing Route Filtering**: Control which routes are advertised to a BGP neighbor. This helps manage the visibility of your network to external peers and has the potential to prevent the advertisement of specific internal routes.
- **Attribute Setting**: Modify route attributes before they are sent to a neighbor.
- **Traffic Engineering**: Influence inbound traffic paths by adjusting route attributes like AS path length to guide traffic through preferred routes.

BGP may receive multiple advertisements for the same route from different sources. BGP selects only one path as the best path. When this path is selected, BGP puts the selected path in the IP routing table and propagates the path to its neighbors. BGP uses the following criteria, in the order presented, to select a path for a destination:

- If the path specifies a next hop that is inaccessible, drop the update.
- Prefer the path with the largest weight.
- If the weights are the same, prefer the path with the largest local preference.
- If the local preferences are the same, prefer the path that was originated by BGP running on this router.
- If no route was originated, prefer the route that has the shortest AS_path.
- If all paths have the same `AS_path` length, prefer the path with the lowest origin type (where IGP is lower than EGP, and EGP is lower than incomplete).
- If the origin codes are the same, prefer the path with the lowest MED attribute.
- If the paths have the same MED, prefer the external path over the internal path.
- If the paths are still the same, prefer the path through the closest IGP neighbor.
- If both paths are external, prefer the path that was received first (the oldest one).
- Prefer the path with the lowest IP address, as specified by the BGP router ID.
- If the originator or router ID is the same for multiple paths, prefer the path with the minimum cluster list length.
- Prefer the path that comes from the lowest neighbor address.

Create a BGP Profile

Create a BGP profile with local autonomous system and neighbor settings, then enable it for dynamic routing across site-to-site VPN connections.

Use the following procedure to create a BGP profile from the Multicloud Defense Controller dashboard:

Before you begin

> **Note**
>
> When you create a BGP profile, the profile must be enabled for traffic and same value to be used in the tunnel as in the BGP profile.

Procedure

1. Navigate to **Infrastructure** > **Network** > **BGP**.
2. Click **Create**.
3. In the **General Settings** tab of the creation window enter a unique **Profile Name**.
4. (Optional) Enter a **Description**. This may help differentiate from other profiles with a similar name.
5. Enter the **LocalAS** value. This value represents the local autonomous system (AS) in which the BGP4 device resides.
6. Click the **Neighbors** tab at the top of the window to switch views. For more information on neighbors and what this configuration can do for you environment, see BGP Neighbors and Path Selection on page 212.
7. Click **Add Neighbors**.
8. Expand the **Neighbor 1** space.
9. Manually enter a singular address or a range of IP addresses and BGP peer groups in the **IP Address** text box. If you are adding multiple addresses, separate each address with a space.
10. Enter the **Autonomous System** value, the LocalAS for where the neighbor resides.
11. If you opt to include **Route Map In**, note that enabling this option applies the route map on all matching traffic in the inbound direction on the interface. If you opt to include **Route Map Out**, enabling this option applies the routemap to all matching traffic in the outbound direction on the interface. Check the appropriate option for your environment and then enter the following information:
    a) **Local Preference** - By default, this value is "100". Optionally, enter a 32-bit unsigned integer value between 0 to 4,294,967,295. Note that with a higher local preference value indicating a more preferred route within an autonomous system.

    Local preference is only exchanged between BGP routers within the same autonomous system (iBGP) and not advertised externally (eBGP).

    b) **AS Path Prepend** - Maually enter a value for this; if you enter more than one separate each value with a **space**. This value influences the path selection process by artificially lengthening the AS path attribute of a route. While it is unconventional to include this for inbound traffic, prepending additional AS numbers to incoming routes can make these routes appear less preferable to your internal BGP speakers when selecting paths for routing traffic.
    c) Click **Add** to include an IP address or a network and enter an IP address, a range of IP addresses sepearated by commas, or a network comprised of both the IP and netmask. These are routes or networks that you want to allow in or out within the bgp session(s). At any point, click **Remove** to remove an IP address from the neighbor.
12. Click **Save**.

What to do next

Add your BGP profile to a Multicloud Defense Gateway. You can either create a new gateway or edit an existing gateway to include the new profile.

23 Profile Actions

Topics:

- View a Profile Details
- Edit a Standalone Metrics Forwarding Profile
- Edit a Group Profile
- Add a Gateway Association to a Profile
- Remove a Gateway Association
- Delete a Profile

View a Profile Details

View the configuration details of a Packet Capture Profile to confirm its cloud account, storage destination, and related profile settings.

Use the following procedure to view the details of a Packet Capture profile.

Procedure

1. Navigate to **Infrastructure** > **Profiles** > **Packet Capture** and select the appropriate profile filter, if required.
2. Select the profile you want to view the details of.
3. View the profile's details.

Edit a Standalone Metrics Forwarding Profile

Edit a standalone profile that has already been created.

Use the following procedure to edit a standalone profile that has already been created.

Procedure

1. Navigate to **Infrastructure** > **Profiles** > **Metrics Forwarding** and select the appropriate profile **Type**.
2. Check the box next to the profile you want to edit.
3. Click **Edit**.
4. Modify the parameters as desired.
5. Click **Save**.

Edit a Group Profile

Change existing grouped profiles by modifying, adding, or removing profile settings.

Use the following procedure to edit a set of grouped profiles that has already been created:

Procedure

1. Navigate to **Infrastructure** > **Profiles** > **Metrics Forwarding** and select the appropriate profile **Type**.
2. Check the box next to the profile you want to edit.
3. Click **Edit**.
4. Modify, add or remove group profiles.
5. Click **Save**.

Add a Gateway Association to a Profile

Associate a gateway with a Packet Capture Profile so the gateway can store captured network traffic in the profile's configured cloud destination.

Use the following procedure to add a gateway association to the desired packet capture profile:

Procedure

1. Navigate to **Infrastructure** > **Gateways** > **Gateways**.
2. Check the box next to the gateway you want to associate the profile to.
3. Click **Edit**.
4. Expand the profile's drop-down menu and select the desired **Profile** from the menu.
5. Click **Save**.

Remove a Gateway Association

Remove an existing gateway association from a packet capture profile without deleting the gateway or profile.

Use the following procedure to remove an existing gateway that is associated with a packet capture profile. Note that this process only removes the gateway association from the profile. This does not delete the gateway or the profile from Multicloud Defense.

Procedure

1. Navigate to **Infrastructure** > **Gateways** > **Gateways**.
2. Check the box next to the gateway you want to disassociate from a packet capture profile.
3. Click **Edit**.
4. Scroll towards the bottom of the page and click the '**X**' within the appropriate profile drop-down menu to remove the association.
5. Click **Save**.

Delete a Profile

Delete a packet capture profile along with existing gateway associations.

Use the following procedure to delete a packet capture profile. This process includes removing any and all existing gateway associations as well as deleting the profile.

Procedure

1. Navigate to **Infrastructure** > **Profiles** > **Packet Capture** and select the appropriate profile **Type**.
2. View the profile details and examine the associated gateways.
3. Remove all gateway associations. See Remove a Gateway Association for more information.

4. Navigate to **Infrastructure** > **Profiles** > **Packet Capture** and select the same profile type that you selected in step 1.
5. Check the box next to the profile you want to delete.
6. Click **Delete**.
7. Click **Yes** or **No** to either confirm or cancel the delete action.

24 FQDN and URL Filtering Categories

Topics:

- FQDN / URL Filtering Categories
- Malicious Categories
- Full List of Categories
- Associating a Filtering Profile with a Policy Ruleset Rule
- Cisco Talos Intelligence URL / IP Lookup Tool

Learn about FQDN and URL filtering categories for defining web traffic controls in Multicloud Defense security policies.

FQDN / URL Filtering Categories

Learn how Cisco Talos Intelligence categories help Cisco Multicloud Defense evaluate outbound FQDN and URL traffic by site type and web reputation.

Multicloud Defense also uses threat intelligence from Cisco Talos Intelligence to categorize web sites based on their risk score. This includes fully qualified domain names (FQDNs), sometimes referred to as domain names, and URLs. This provides sites across 84 categories when traffic from your public cloud environment makes outbound connections (egress) to these sites:

- FQDNs (domains) - 1+ billion categorized FQDNs (domains)
- URLs - 45+ billion categorized URLs

To improve efficiency in recognizing and processing traffic, the gateway will pre-load a cache of the top 1 million FQDNs/URLs and their categories. The gateway will also utilize a runtime cache of 10k FQDNs/URLs and their Categories that are not part of the top 1 million. If traffic contains any of the cached FQDNs/URLs, then the categories will be known immediately. If the FQDN/URL is not found in the cache, the gateway will query the Multicloud Defense Controller to resolve the category via Talos. This operation is expected to complete in no more than 200ms. If it completes within the expected time, then the traffic will be processed based on the learned category and the profile will operate on the traffic based on the policy defined for the category. If the operation does not complete within the expected time, then the traffic will be processed as Uncategorized and the profile will operate on the traffic based on the policy defined for Uncategorized. Once the resolution returns, the learned category will be added to the cache for subsequent resolutions, even if the resolution occurs after the expected time and the traffic has already been processed. If the run-time cache is exhausted, the gateway will purge the oldest accessed FQDNs/URLs and their categories in batches of 10 entries to ensure space is available for more recently accessed FQDNs/URLs and their categories.

Note

FQDN filtering with categories happens for:

1. SNI in TLS client hello
2. DNS queries for FQDN lookups
3. HTTP hostname header (for cleartext HTTP traffic)

Malicious Categories

Find details about FQDN and URL categories that Cisco Multicloud Defense identifies as malicious, including malware, phishing, botnet, spam, and anonymizer sites.

Multicloud Defense considers the following categories to be particularly malicious:

Table 4: Malicious Categories Multicloud Defense considers the following categories to be particularly malicious

Category Name	Category Description
Malware Sites	Sites hosting malicious content including executables, drive-by infection sites, malicious scripts, viruses, trojans, and code.

Category Name	Category Description
Phishing and Other Frauds	Phishing, pharming, and other sites that pose as a reputable site, usually to harvest personalinformation from a user. These sites are typically quite short-lived, so they don't last long in terms of uptime.
Proxy Avoidance and Anonymizers	Proxyservers and other methods to gain access to URLs in any way that bypasses URL filtering or monitoring. Web-based translation sites that circumvent filtering.
Keyloggers and Monitoring	Softwareagents that track a user's keystrokes or monitor their web surfing habits. Often used for collecting sensitive data such as usernames and passwords.
SPAM URLs	Sites known to distribute unsolicited email (spam) messages.
Spyware and Adware	Spywareor Adware sites that provide or promote information gathering or tracking that is unknown to, or without the explicit consent of, the end user or the organization, also unsolicited advertising popups and programs that may be installed on a user's computer.
Bot Nets	These are URLs, often IP addresses, which are determined to be part of a Bot network, from which network attacks are launched. Attacks may include SPAM messages, DOS, SQL injections, proxy jacking, and other unsolicited contacts.

Multicloud Defense offers traffic analysis when viewing traffic via **Investigate** > **Traffic Analysis** > **DNS** and **Investigate** > **Flow Analytics** > **Traffic Summary**, where a pre-defined *Malicious Categories* filter can be selected to show instances and VPCs communicating with these Malicious Category FQDNs and URLs.

The complete list of categories is shown below.

Full List of Categories

Find the complete list of Cisco Talos Intelligence categories available for classifying FQDNs and URLs in Cisco Multicloud Defense filtering profiles.

Category Name	Category Name	Category Name	Category Name
Abortion	Games	Motor Vehicles	Sex Education
Abused Drugs	Government	Music	Shareware and Freeware
Adult and Pornography	Gross	News and Media	Shopping
Alcohol and Tobacco	Hacking	Nudity	Social Networking
Auctions	Hate and Racism	Online Greeting Cards	Society
Bot Nets	Health and Medicine	Open HTTP Proxies	SPAM URLs
Business and Economy	Home and Garden	Parked Domains	Sports
Cheating	Hunting and Fishing	Pay to Surf	Spyware and Adware
Computer and Internet Info	Illegal	Peer to Peer	Streaming Media

Category Name	Category Name	Category Name	Category Name
Computer and Internet Security	Image and Video Search	Personal sites and Blogs	Swimsuits and Intimate Apparel
Confirmed SPAM Sources	Individual Stock Advice and Tools	Personal Storage	Training and Tools
Content Delivery Networks	Internet Communications	Philosophy and Political Advocacy	Translation
Cult and Occult	Internet Portals	Phishing and Other Frauds	Travel
Dating	Job Search	Private IP Addresses	Uncategorized
Dead Sites	Keyloggers and Monitoring	Proxy Avoidance and Anonymizers	Unconfirmed SPAM Sources
Dynamically Generated Content	Kids	Questionable	Violence
Educational Institutions	Legal	Real Estate	Weapons
Entertainment and Arts	Local Information	Recreation and Hobbies	Web Advertisements
Fashion and Beauty	Malware Sites	Reference and Research	Web Hosting
Financial Services	Marijuana	Religion	Web-based Email
Gambling	Military Search Engines	Services	

Associating a Filtering Profile with a Policy Ruleset Rule

Find guidance for creating or editing FQDN and URL Filtering Profiles before associating them with a security policy ruleset rule.

- Refer to FQDN Filtering to create/edit FQDN Filtering Profiles
- Refer to URL Filtering to create/edit URL Filtering Profiles

Cisco Talos Intelligence URL / IP Lookup Tool

Look up an FQDN, URL, or IP address with Cisco Talos Intelligence to find its assigned web category and reputation information.

Cisco Talos Intelligence (Talos) offers an online URL / IP lookup tool that can be used to understand what category a particular FQDN / URL is classified as, along with its web reputation.

25 Flow Analytics

Topics:

- Flow Analytics - Traffic Summary
- Flow Analytics - All Events
- Firewall Events
- Network Threats
- Web Attacks
- URL Filtering
- FQDN Filtering
- HTTPS Logs
- VPN Logs
- EndUser Logs
- HTTP2 Logs
- Investigate summary page

Learn how to use Flow Analytics in Multicloud Defense to understand network traffic across cloud environments.

Flow Analytics - Traffic Summary

Find details about active and closed proxy sessions in the Flow Analytics Traffic Summary, including endpoints, traffic volume, applications, policies, and connection status.

This view provides detailed visibility, filtering and analysis for events recorded by Multicloud Defense from either a forward or reverse gateway proxy. Traffic Summary events contribute to one of three event types: `Firewall Events`, `Network Events` and `Web Attacks`.

The summary page shows both closed and ongoing sessions; select "All events" in the traffic summary page and this selection will list all events including closed and the periodic 5 minute events for long running sessions.

Traffic Summary

Tables and Fields available in Session Summary are as follows:

Event Details	Description
Date and Time	ISO 8601 format: YYYY-MM-DD T HH:MM:SS:S Example: `2020-11-22T10:58:46.820`
CSP Account	Multicloud Defense CSP Account
Gateway	Multicloud Defense Gateway
Region	Region of the Multicloud Defense Gateway
Level	INFO
Session ID	..

Client-side Connection	Description
Src IP	Source IP Address
Src Port	Source Port
Dest IP	Destination IP Address
Dest Port	Destination Port
Protocol	UDP, TCP

Client-side Stats	Traffic between client and Multicloud Defense Gateway
Received Bytes	Number of bytes received from client
Transmitted Bytes	Number of bytes sent to client
Received Packets	Number of packets received from client
Transmitted Packets	Number of packets sent to client

Policy Match Info	Description
Dest Address Group	Destination Address Group configured in the matched policy rule

Policy Match Info	Description
Src Address Group	Source Address Group configured in the matched policy rule
Request SNI	Server Name Indication in the request
Service Type	Service Type. Example: `PROXY`
Src Country	Country that the request originated from on the client-side
Dest Country	Country that the request was destined to on the server-side. Example: `United States`

Server-side Connection	Description
Src IP	Source IP Address
Src Port	Source Port
Dest IP	Destination IP Address
Dest Port	Destination Port
Protocol	UDP, TCP

Server-side Stats	Traffic between Multicloud Defense Gateways and server
Received Bytes	Number of bytes received from server
Transmitted Bytes	Number of bytes sent to server
Received Packets	Number of packets received from server
Transmitted Packets	Number of packets sent to server

Application Info	Description
Client App Name	Application name associated with client side of the session. Example: `Advanced Packaging Tool`
Payload App Name	HTTP application name associated with webserver host. Example: `Facebook`
Service App Name	Application name associated with server side of the session. Example: `HTTP`

Action	Description
Action	ALLOW, DENY

Cloud Service	Description
Cloud Service	Name of the destination cloud service accessed with the request. Example `AMAZON, EC2`

Src Instance Info	Description
Instance ID	Client instance ID
Instance Name	Client instance name (and provides ability to see tags)
VPC ID	Client VPC ID

HTTP Request	Description
Host	Host portion of URL
Method	GET, PUT, POST, HEAD, DELETE, PATCH, OPTIONS
URI	URI Identifier RFC 3986

Rule	Description
ID	ID number/description of Multicloud Defense Rule. Example `59 (egress-prod-apt-80).`

FQDN	Description
FQDN	Fully Qualified Domain Name
Category Name	Category classification of the FQDN. Example: `Social Media`
Reputation	Reputation score of the FQDN

Flow Analytics - All Events

Find details about all network and security events in Flow Analytics, including session, application, policy, traffic, and threat information for investigation.

Flow Analytics – All Events provides overall visibility into network and security events from the entire Multicloud Defense solution.

Tables and Fields available in All Events are as follows:

Event Details	Description
Date and Time	ISO 8601 format: YYYY-MM-DD T HH:MM:SS:S Example: `2020-11-22T10:58:46.820.`
Type	APPID, AV, DLP, DPI, FLOW_LOG, FQDNFILTER, L4_FW, L7DOS, MALICIOUS_SRC, SNI, TLS_ERROR, TLS_LOG, URLFILTER.
CSP Account	Multicloud Defense CSP Account.
Gateway	Multicloud Defense Gateway.
Region	Region of the Multicloud Defense Gateway.

Event Details	Description
Level	DEBUG, INFO, NOTICE, WARNING, ERROR, CRITICAL, ALERT, EMERGENCY.
Session ID	..

Service	Description
Src IP	Source IP Address.
Src Port	Source Port.
Dest IP	Destination IP Address.
Dest Port	Destination Port.
Protocol	UDP, TCP.

Application Info	Description
Client App Name	Application name associated with client side of the session. Example: `Advanced Packaging Tool`.
Payload App Name	HTTP application name associated with webserver host. Example: `Facebook`.
Service App Name	Application name associated with server side of the session. Example: `HTTP`.

Action	Description
Action	ALLOW, DENY.
State	ESTABLISHED, CLOSE, CLOSED, CLOSE_WAIT, TIME_WAIT, FIN_WAIT, LAST_ACK.

HTTP Request	Description
Host	Host portion of URL.
Method	GET, PUT, POST, HEAD, DELETE, PATCH, OPTIONS.
URI	URI Identifier RFC 3986.

Rule	Description
ID	ID number/description of Multicloud Defense Rule. Example `59 (egress-prod-apt-80)`.

FQDN	Description
FQDN	Fully Qualified Domain Name.

FQDN	Description
Category Name	Category classification of the FQDN. Example: `Social Media`.
Reputation	Reputation score of the FQDN.

Event Logs

Describes the components and function of event logs that monitor and detail network traffic.

Event logs contain details of all traffic that flows through the Multicloud Defense Gateway.

After inspection, Multicloud Defense generates sessions and events that are based on what is in the packet and what is defined in the policy. The analysis, related details of events, and actions that are taken are all captured in the form of logs, available under **Investigate** > **Flow Analytics** > **All Events**. The system retains these logs for 30 days.

Event types that the logs capture:

Table 5: Event Types and Descriptions

Event Type	Event Name	Description
FQDN FILTER	Fully Qualified Domain Name (FQDN) Filtering	The related logs generate with details of the FQDN, source, destination IP and so on. The FQDN filtering event only generates in case the policy has an FQDN filtering profile.
SNI	Server Name Indication (SNI)	SNI allows multiple host names to be served over HTTPS. This generates when Multicloud Defense observes the SNI in the TLS handshake.
APPID	App ID (APPID)	APPID logs generate when the event matches known applications in the database.
L4_FW	L4 Firewall	An L4 Firewall event generates when the event matches the policy in the ruleset.
URL FILTER	URL Filtering	URL filtering is used to filter out network traffic based on the URL. This event log generates when it matches the URL filtering profile.
IPS	Intrusion Prevention System (IPS)	An IPS event generates when the network traffic matches the IPS ruleset.
DLP	Data Loss Protection (DLP)	A DLP event generates when the network traffic matches the DLP profile that is configured. The logs record these incidents, along with details of transmission such as endpoint, domain, username, rules, source, destination, action taken, and so on.
WAF	Web Application Firewall	A WAF event generates when the network traffic matches the WAF profile that is configured.
AV	Antivirus (AV)	An AV event generates when the event matches an AV ruleset in the network traffic.

Event Type	Event Name	Description
DPI	Deep Packet Inspection (DPI)	A DPI event generates when the network traffic matches a rule that has an advanced security configured.
MALICIOUS_SRC	Malicious Source	A Malicious Source generates when the network traffic matches a malicious IP.
TLS_ERROR	TLS Error	A TLS error generates when there is an error during the TLS handshake.
TLS_LOG	TLS Log	A TLS log generates when the network traffic uses TLS. This captures the TLS handshake information such as cipher suites and TLS version.

Firewall Events

Find details about firewall events in Flow Analytics, including connection endpoints, applications, actions, session states, HTTP requests, and matching security policy rules.

This view provides detailed visibility, filtering and analysis for events recorded by the Multicloud Defense Firewall configuration and summarized in `Firewall Events` category.

Tables and Fields available in Firewall Events are as follows:

Event Details	Description
Date and Time	ISO 8601 format: YYYY-MM-DD T HH:MM:SS:S Example: `2020-11-22T10:58:46.820`
Type	APPID, L4_FW, MALICIOUS_SRC, SNI
CSP Account	Multicloud Defense CSP Account
Gateway	Multicloud Defense Gateway
Region	Region of the Multicloud Defense Gateway
Level	DEBUG, INFO, NOTICE, WARNING, ERROR, CRITICAL, ALERT, EMERGENCY
Session ID	..

Service	Description
Src IP	Source IP Address
Src Port	Source Port
Dest IP	Destination IP Address
Dest Port	Destination Port
Protocol	UDP, TCP

Application Info	Description
Client App Name	Application name associated with client side of the session. Example: `Advanced Packaging Tool`
Payload App Name	HTTP application name associated with webserver host. Example: `Facebook`
Service App Name	Application name associated with server side of the session. Example: `HTTP`

Action	Description
Action	ALLOW, DENY
State	ESTABLISHED, CLOSE, CLOSED, CLOSE_WAIT, TIME_WAIT, FIN_WAIT, LAST_ACK

HTTP Request	Description
Host	Host portion of URL
Method	GET, PUT, POST, HEAD, DELETE, PATCH, OPTIONS
URI	URI Identifier RFC 3986

Rule	Description
ID	ID number/description of Multicloud Defense Rule. Example 59 (egress-prod-apt-80)

FQDN	Description
FQDN	Fully Qualified Domain Name
Category Name	Category classification of the FQDN. Example: `Social Media`
Reputation	Reputation score of the FQDN

Network Threats

Find details about antivirus, data loss prevention, and intrusion events in Flow Analytics to investigate detected threats and affected network sessions.

This view provides detailed visibility, filtering and analysis for threats recorded by the Multicloud Defense threat analysis engine and summarized in `Network Threats`.

Network Threats

Tables and Fields available in Network Threats are as follows:

Event Details	Description
Date and Time	ISO 8601 format: YYYY-MM-DD T HH:MM:SS:S Example: `2020-11-22T10:58:46.820`
Type	AV, DLP, DPI
CSP Account	Multicloud Defense CSP Account
Gateway	Multicloud Defense Gateway
Region	Region of the Multicloud Defense Gateway
Level	DEBUG, INFO, NOTICE, WARNING, ERROR, CRITICAL, ALERT, EMERGENCY
Session ID	..

Service	Description
Src IP	Source IP Address
Src Port	Source Port
Dest IP	Destination IP Address
Dest Port	Destination Port
Protocol	UDP, TCP

Application Info	Description
Client App Name	Application name associated with client side of the session. Example: `Advanced Packaging Tool`
Payload App Name	HTTP application name associated with webserver host. Example: `Facebook`
Service App Name	Application name associated with server side of the session Example: `HTTP`

Action	Description
Action	ALLOW, DENY
State	ESTABLISHED, CLOSE, CLOSED, CLOSE_WAIT, TIME_WAIT, FIN_WAIT, LAST_ACK

HTTP Request	Description
Host	Host portion of URL
Method	GET, PUT, POST, HEAD, DELETE, PATCH, OPTIONS
URI	URI Identifier RFC 3986

FQDN	Description
FQDN	Fully Qualified Domain Name
Category Name	Category classification of the FQDN. Example: `Social Media`
Reputation	Reputation score of the FQDN

Rule	Description
ID	ID number/description of Multicloud Defense Rule. Example 59 (egress-prod-apt-80)

Web Attacks

Find details about WAF and Layer 7 denial-of-service events in Flow Analytics to investigate security attacks against protected web applications.

This view provides detailed visibility, filtering and analysis for threats recorded by the Multicloud Defense web protection engine. The `Web Attacks` event types include WAF and L7 DOS.

Tables and Fields available in Web Attacks are as follows:

Event Details	Description
Date and Time	ISO 8601 format: YYYY-MM-DD T HH:MM:SS:S Example: `2020-11-22T10:58:46.820`
Type	L7DOS, WAF
CSP Account	Multicloud Defense CSP Account
Gateway	Multicloud Defense Gateway
Region	Region of the Multicloud Defense Gateway
Level	DEBUG, INFO, NOTICE, WARNING, ERROR, CRITICAL, ALERT, EMERGENCY
Session ID	..

Service	Description
Src IP	Source IP Address
Src Port	Source Port
Dest IP	Destination IP Address
Dest Port	Destination Port
Protocol	UDP, TCP

Application Info	Description
Client App Name	Application name associated with client side of the session. Example: `Advanced Packaging Tool`
Payload App Name	HTTP application name associated with webserver host. Example: `Facebook`
Service App Name	Application name associated with server side of the session Example: `HTTP`

Action	Description
Action	ALLOW, DENY
State	ESTABLISHED, CLOSE, CLOSED, CLOSE_WAIT, TIME_WAIT, FIN_WAIT, LAST_ACK

HTTP Request	Description
Host	Host portion of URL
Method	GET, PUT, POST, HEAD, DELETE, PATCH, OPTIONS
URI	URI Identifier RFC 3986

FQDN	Description
FQDN	Fully Qualified Domain Name
Category Name	Category classification of the FQDN. Example: `Social Media`
Reputation	Reputation score of the FQDN

Rule	Description
ID	ID number/description of Multicloud Defense Rule. Example 59 (egress-prod-apt-80)

URL Filtering

Find details about URL Filtering events in Flow Analytics, including requested URLs, categories, HTTP methods, actions, applications, and matching security policy rules.

This view provides detailed visibility, filtering and analysis for events recorded by the Multicloud Defense URL Filtering configuration. URL Filtering events contribute to one of three event types: `Firewall Events`, `Network Events` and `Web Attacks`.

Event Details	Description
Date and Time	ISO 8601 format: YYYY-MM-DD T HH:MM:SS:S Example: `2020-11-22T10:58:46.820`

Event Details	Description
Type	URLFILTER
CSP Account	Multicloud Defense CSP Account
Gateway	Multicloud Defense Gateway
Region	Region of the Multicloud Defense Gateway
Level	DEBUG, INFO, NOTICE, WARNING, ERROR, CRITICAL, ALERT, EMERGENCY
Session ID	..

Service	Description
Src IP	Source IP Address
Src Port	Source Port
Dest IP	Destination IP Address
Dest Port	Destination Port
Protocol	UDP, TCP

Application Info	Description
Client App Name	Application name associated with client side of the session. Example: `Advanced Packaging Tool`.
Payload App Name	HTTP application name associated with webserver host. Example: `Facebook`
Service App Name	Application name associated with server side of the session Example: `HTTP`

Action	Description
Action	ALLOW, DENY
State	ESTABLISHED, CLOSE, CLOSED, CLOSE_WAIT, TIME_WAIT, FIN_WAIT, LAST_ACK

HTTP Request	Description
Host	Host portion of URL
Method	GET, PUT, POST, HEAD, DELETE, PATCH, OPTIONS
URI	URI Identifier RFC 3986

Rule	Description
ID	ID number/description of Multicloud Defense Rule. Example 59 (egress-prod-apt-80)

FQDN	Description
FQDN	Fully Qualified Domain Name
Category Name	Category classification of the FQDN. Example: `Social Media`
Reputation	Reputation score of the FQDN

FQDN Filtering

Find details about FQDN Filtering events in Flow Analytics, including domains, categories, reputation scores, actions, session data, and matching policy rules.

This view provides detailed visibility, filtering and analytical options for events recorded from the FQDN Filtering configuration. FQDN Filtering events contribute to one of three event types: `Firewall Events`, `Network Events` and `Web Attacks`.

Event Details	Description
Date and Time	ISO 8601 format: YYYY-MM-DD T HH:MM:SS:S Example: `2020-11-22T10:58:46.820`.
Type	FQDNFILTER.
CSP Account	Multicloud Defense CSP Account.
Gateway	Multicloud Defense Gateway.
Region	Region of the Multicloud Defense Gateway.
Level	DEBUG, INFO, NOTICE, WARNING, ERROR, CRITICAL, ALERT, EMERGENCY.
Session ID	..

Service	Description
Src IP	Source IP Address.
Src Port	Source Port.
Dest IP	Destination IP Address.
Dest Port	Destination Port.
Protocol	UDP, TCP.

Action	Description
Action	ALLOW, DENY.

Action	Description
State	ESTABLISHED, CLOSE, CLOSED, CLOSE_WAIT, TIME_WAIT, FIN_WAIT, LAST_ACK.

HTTP Request	Description
Host	Host portion of URL.
Method	GET, PUT, POST, HEAD, DELETE, PATCH, OPTIONS.
URI	URI Identifier RFC 3986.

FQDN	Description
FQDN	Fully Qualified Domain Name.
Category Name	Category classification of the FQDN. Example: `Social Media`.
Reputation	Reputation score of the FQDN.

Rule	Description
ID	ID number/description of Multicloud Defense Rule. Example 59 (egress-prod-apt-80).

HTTPS Logs

Find details about HTTPS events in Flow Analytics, including TLS errors, connection endpoints, applications, certificate information, actions, and related policy rules.

This view provides detailed visibility, filtering and analytical options for events recorded from HTTPS Logs. HTTPS logs may contribute to one of three event types: `Firewall Events`, `Network Events` and `Web Attacks`.

Event Details	Description
Date and Time	ISO 8601 format: YYYY-MM-DD T HH:MM:SS:S Example: `2020-11-22T10:58:46.820`
Type	TLS_ERROR, TLS_LOG.
CSP Account	Multicloud Defense CSP Account.
Gateway	Multicloud Defense Gateway.
Region	Region of the Multicloud Defense Gateway.
Level	DEBUG, INFO, NOTICE, WARNING, ERROR, CRITICAL, ALERT, EMERGENCY.
Session ID	..

Service	Description
Src IP	Source IP Address.
Src Port	Source Port.
Dest IP	Destination IP Address.
Dest Port	Destination Port.
Protocol	UDP, TCP.

Application Info	Description
Client App Name	Application name associated with client side of the session. Example: `Advanced Packaging Tool`.
Payload App Name	HTTP application name associated with webserver host. Example: `Facebook`.
Service App Name	Application name associated with server side of the session Example: `HTTP`.

Action	Description
Action	ALLOW, DENY.
State	ESTABLISHED, CLOSE, CLOSED, CLOSE_WAIT, TIME_WAIT, FIN_WAIT, LAST_ACK.

HTTP Request	Description
Host	Host portion of URL.
Method	GET, PUT, POST, HEAD, DELETE, PATCH, OPTIONS.
URI	URI Identifier RFC 3986.

FQDN	Description
FQDN	Fully Qualified Domain Name.
Category Name	Category classification of the FQDN. Example: `Social Media`.
Reputation	Reputation score of the FQDN.

VPN Logs

Find details about VPN connections in Flow Analytics, including activity, usage, errors, security events, gateway information, and downloadable event reports.

Virtual Private Network (VPN) logs are records of activities and events that occur within a VPN and can provide detailed information about the usage, performance, and security of the connection. VPN logs include connection, usage, activity,

error, and security logs. Note that the display shown on this page is directly dependent on the selected event details. Click the **Edit** icon to modify the display and select from the following options:

Event Details	Description
Date and Time	ISO 8601 format: YYYY-MM-DD T HH:MM:SS:S Example: 2020-11-22T10:58:46.820.
CSP Account	Name of your cloud service account.
Region	Region of the Multicloud Defense Gateway.
Gateway	The Multicloud Defense Gateway involved in the event.
Text	A preview of the text included in the event message. Click an individual message to expand.
Gateway Security Type	Designation of the Multicloud Defense Gateway.
Instance Name	Identifier for a VPN session or connection instance.

EndUser Logs

Find details about end-user application activity in Flow Analytics to review user sessions, application performance, connection endpoints, and potentially suspicious behavior.

An EndUser log is a record of interactions and activities performed by end users within an application or service. These logs act as a snapshot of valuable information about user behavior, system usage, and application performance. Implementing this within your network can create insight as to how users interact with their applications, can track application performance from the user's perspective, and has the potential to help detect suspicious activities with proper investigation.

Event Details	Description
Date and Time	ISO 8601 format: YYYY-MM-DD T HH:MM:SS:S Example: 2020-11-22T10:58:46.820.
CSP Account	Name of your cloud service account.
Region	Region of the Multicloud Defense Gateway.
Gateway	The Multicloud Defense Gateway involved in the event.
Session ID	The unique identifier assigned to a user's session when they interact with an application or system.
Text	A preview of the text included in the event message. Click an individual message to expand.
Level	The severity or importance of a logged event. This can help categorize and prioritize entries for easy analyzing.
Src IP	Identifier for a VPN session or connection instance.
Dst IP	The destination IP address of a network connection of communication.

Event Details	Description
Dst Port	The numerical destination port of a network connection or communication.
Payload App Name	The application or component of a system where the user action took place.
Action	The specific operation or event performed by the user within the system.
Policy Name	The name of the policy matching against the user.
Instance Name	The unique identifier or label assigned to that particular instance within an environment.
First Name	The identified first name of the user performing this action.
Last Name	The identified last name of the user performing this action.
Group	The identified group the user performing this action is associated with.
Department	The identified department of the user performing this action.
Method	The action that a client uses to communicate with a server: **GET**, **POST**, **PUT**, **DELETE**, **HEAD**, and so on.
URI	The identifying string of the resource being requested from the server.
FQDN	The FQDN of where the logged event originates from.
Category Name	The the name of the category the logged event is associated as.

HTTP2 Logs

Find details about HTTP/2 activity in Flow Analytics to analyze application behavior, connection performance, resource use, and data transmission issues.

The HTTP2, or HTTP/2, log refers to logging information specifically related to HTTP/2. It is designed to improve the performance and efficiency of data transfer between clients and servers. When logging HTTP/2 activity, the logs capture details about how HTTP/2 protocol features are utilized. This has the potential to provide insights into the performance and behavior of web applications using this protocol This can help developers and administrators optimize resource usage, troubleshoot issues, and ensure efficient data transmission.

Event Details	Description
Date and Time	ISO 8601 format: YYYY-MM-DD T HH:MM:SS:S Example: 2020-11-22T10:58:46.820.
CSP Account	Name of your cloud service account.
Region	Region of the Multicloud Defense Gateway.

Event Details	Description
Gateway	The Multicloud Defense Gateway involved in the event.
Session ID	The unique identifier assigned to a session.
Text	A preview of the text included in the event message. Click an individual message to expand.
Src IP	Identifier for a VPN session or connection instance.
Dst IP	The destination IP address of a network connection of communication.
Dst Port	The numerical destination port of a network connection or communication.
Payload App Name	The application or component of a system where the HTTP2 communication took place.
Action	The specific operation or type of event that was logged regarding HTTP2 communication.
Policy Name	The name of the policy matching against the HTTP2 log.
Instance Name	The unique identifier or label assigned to that particular HTTP2 activity.
First Name	The identified first name of the user associated with the HTTP2 activity.
Last Name	The identified last name of the user associated with the HTTP2 activity.
Group	The identified user group associated with this HTTP2 activity.
Department	Wwhich department within an organization is responsible for or associated with the logged activity.
FQDN	The FQDN of the server or host involved in the communication.
Category Name	The the name of the category the logged communication is associated as.

Investigate summary page

Explains the concept of traffic summary and logs in diagnosing policy effectiveness and threats.

The Investigate tab of the Multicloud Defense Controller offers a collection of traffic, events, and logs that can assist in diagnosing policy effectiveness and threats.

Flow Analytics

Flow Analytics provides overall visibility into the traffic seen, processed and protected by the Multicloud Defense Gateway. The traffic is organized into two main categories: traffic summary logs and security events. Traffic summary

logs provide information related to each traffic session that is being processed by the gateways. Security events provide information related to how the gateway datapath protects each traffic session.

Network Analytics

Network Stats provides information on the performance of the gateway. The generated graph has the potential to display how gateways and instances associated with the gateways autoscale to combat with the capacity threshold. This can be a useful tool in troubleshooting gateway behavior, trends or spikes, and gateway management.

System Status

System Logs detail which user logged into the Multicloud Defense Controller by time and time range, as well as actions performed.

26 Network Analytics

Topics:

- Stats

Learn how to use Network Analytics in Multicloud Defense to investigate cloud network traffic and activity.

Stats

View and filter Network Analytics statistics for selected cloud accounts and gateways to assess bandwidth, connections, HTTP requests, CPU usage, and memory usage.

This view provides detailed visibility into the bandwidth and connections of selected Multicloud Defense gateway/s, both instantaneously and over selected timeframes.

Procedure

1. Navigate to **Investigate** > **Network Analysis** > **Stats**.
2. Initially, statistics are displayed for `All CSP Accounts` and `All Gateways` with timeframe default set to the `Last 1 hour`.
3. Graphically, the X and Y axes are auto-scaled based on timeframe selection/bandwidth, and auto- updated while viewing. Statistics are refreshed every 5 seconds while viewing this page.
4. Use the drop-down options in the filter bar to finesse the display and view the stats of a specific **Account**, **CSP Type**, or **Instance Type**.

   Note that if you select **Instance Type**, you see two additional stats: CPU usage and memory usage.
5. Select a **Timeframe** from the pulldown as shown below. Options are: `Last 15 mins`, `Last 1 hour`, `Last 1 day`, `Last 7 Days`, `Last 30 days`.

Total Bandwidth

Learn how Total Bandwidth combines inbound and outbound gateway traffic measurements to help you assess network capacity and data transfer rates.

The total network bandwidth is a measurement indicating the maximum capacity of a wired or wireless communications link to transmit data over a network connection in a given amount of time. This value is a compilation of **total speed** (addition of inbound and outbound bandwidth of selected gateways), **inbound** bandwidth (bandwidth ingressing a gateway), and **outbound** bandwidth (bandwidth egressing a gateway).

CPU Usage

Learn how CPU Usage statistics help you monitor gateway processing demand, identify unusual activity, and assess performance trends for selected instance types.

> **Note**
>
> This statistic is only available if you include an **Instance Type** in your selection from the filter bar located at the top of the page.

This view provides information about gateway instances that may have higher than normal memory use. You can use this information to monitor and optimize the performance of the gateway activities based on CPU capacity. You can also use these stats to help assess trends in traffic and the effort expressed by the CPU over the behaviors.

Memory Usage

Learn how Memory Usage statistics help you identify gateways with unusual resource consumption and assess performance capacity for selected instance types.

> **Note**
>
> This statistic is only available if you include an **Instance Type** in your selection from the filter bar located at the top of the page.

This view provides information about gateway instances that may have higher than normal memory use. You can use this information to monitor and optimize the performance of the gateway activities based on memory usage capacity.

Connection Rate

Learn how Connection Rate metrics show active gateway connections and connections per second to help you assess network demand and gateway activity.

Connection rates refer to the percentage of successfully connected calls out of the total attempted calls. Specifically, it equates to the **connection** (total number of current active connections) and **connections per second** (bandwidth of both inbound and outbound connections to a gateway).

HTTP Request Rate

Learn how HTTP Request Rate measures requests per second to help you understand demand on web services protected by Cisco Multicloud Defense gateways.

An **HTTP request rate** typically measures how much demand is being placed on your system, expressed as a high-level system-specific metric. For a web service, this measurement is usually HTTP requests per second.

27 Audit and System Logs

Topics:

- Audit Logs
- System Logs

Provides a unified view and advanced features to manage logs efficiently.

Audit Logs and System Logs are available in a combined view under **Investigate** > **System Status** > **Audit and System Logs**. The unified page offers advanced search capabilities, clear identification of audit and system logs, and easier correlation of events and logs, all in one place.

Audit Logs

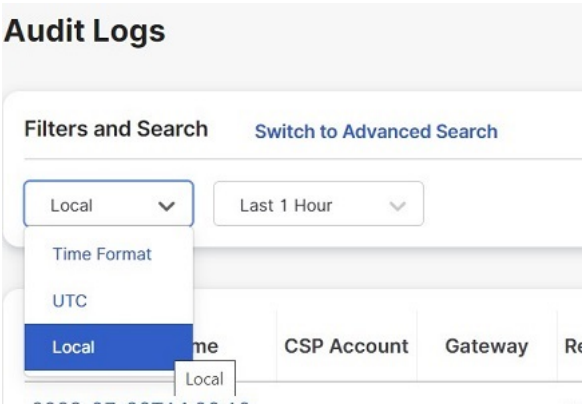
Learn how Audit Logs help you review user activity and configuration changes, select a time format and timeframe, and support operational investigations.

Audit logs contain details of actions performed by Users. This includes, but is not limited to, actions such as login/logout activity, creating, deleting, updating, enabling, and disabling of Profiles, Rules, Gateways, or any user activity that relates to the configuration and operation of the Multicloud Defense solution. Multicloud Defense supports log forwarding to AWS S3.

Time Format

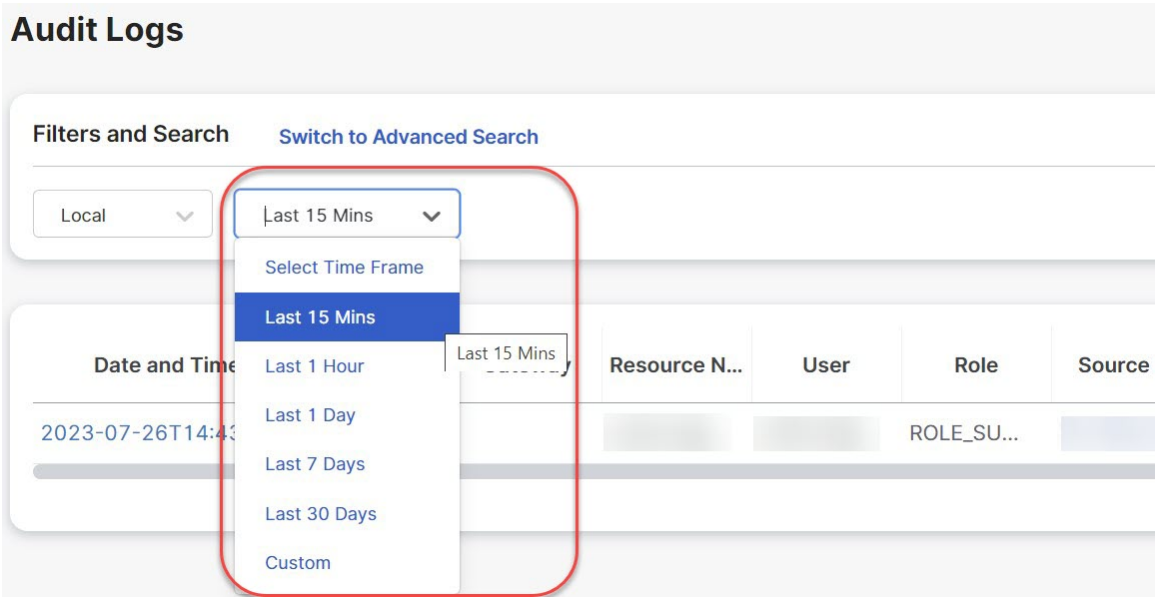
Logs can be displayed in UTC (Coordinated Universal Time) or Local time format. Local means the time zone of the user as configured e.g. USA/Pacific. Date and Time of logs will be displayed in ISO 8601 format (complete date plus hours, minutes, seconds and a decimal fraction of a second - YYYY-MM-DD T HH:MM:SS:S). Example: `2020-11-22T10:58:46.820`

To select or switch between different time formats, click the drop-down and select a format.



Timeframe

Logs can be displayed in increment options from 15 minutes to 30 days, or Custom timeframes. To select or switch between timeframes, click the drop-down menu and select a timeframe as shown:



For Custom timeframes, select **Custom**, the **Start** and the **End** date or time by clicking the calendar objects followed by **Save**.

Search Filter

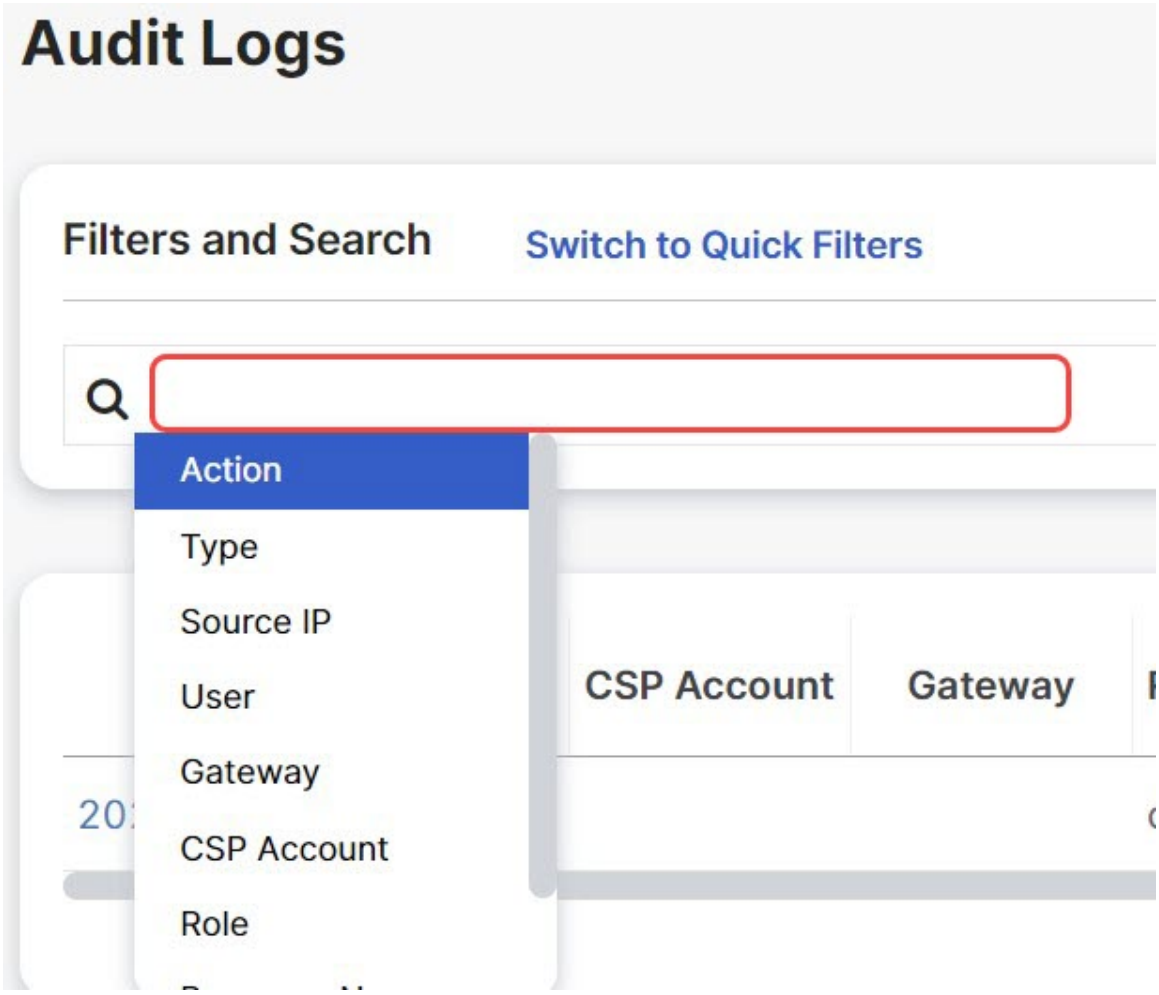
Search and filter Audit Logs by action, source IP address, user, gateway, cloud account, or role to investigate specific user activity.

Logs can be filtered using the Search function and audit log fields. The audit log fields are `Action Type`, `Source IP`, `User`, `Gateway`, `CSP Account`, `Role`.

To filter audit logs on one, or multiple, fields:

Procedure

1. Left mouse-click in the Search field to access the pull down menu.



2. Select a field.
3. Type a desired search string.
4. Add additional fields to the search criteria as required.

   Example: Filter for Actions = **" DELETE"** and performed by user with string containing **" steve"** would appear in the filter criteria and results.

System Logs

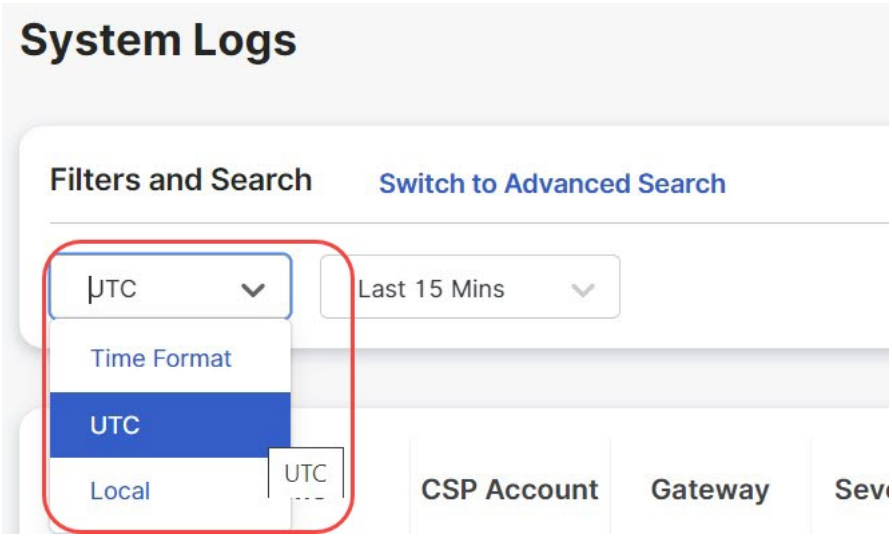
Learn how System Logs help you review system messages, gateway events, resource changes, and operational issues by timeframe and severity level.

System logs contain details of actions that the Multicloud Defense solution performs. This includes system messages, gateway events, instance creation or deletion, and other configuration and operation modifications of the Multicloud Defense solution and more. The system stores these logs for a duration of 1 year. Multicloud Defense supports log forwarding to AWS S3.

Time Format

Logs display in UTC (Coordinated Universal Time) or local time format. Local means the time zone of the user as configured. For example, in the USA/Pacific time zone. Date and Time of logs display in ISO 8601 format (Complete date plus hours, minutes, seconds, and a decimal fraction of a second - YYYY-MM-DD T HH:MM:SS:S). Example: `2020-11-22T10:58:46.820`

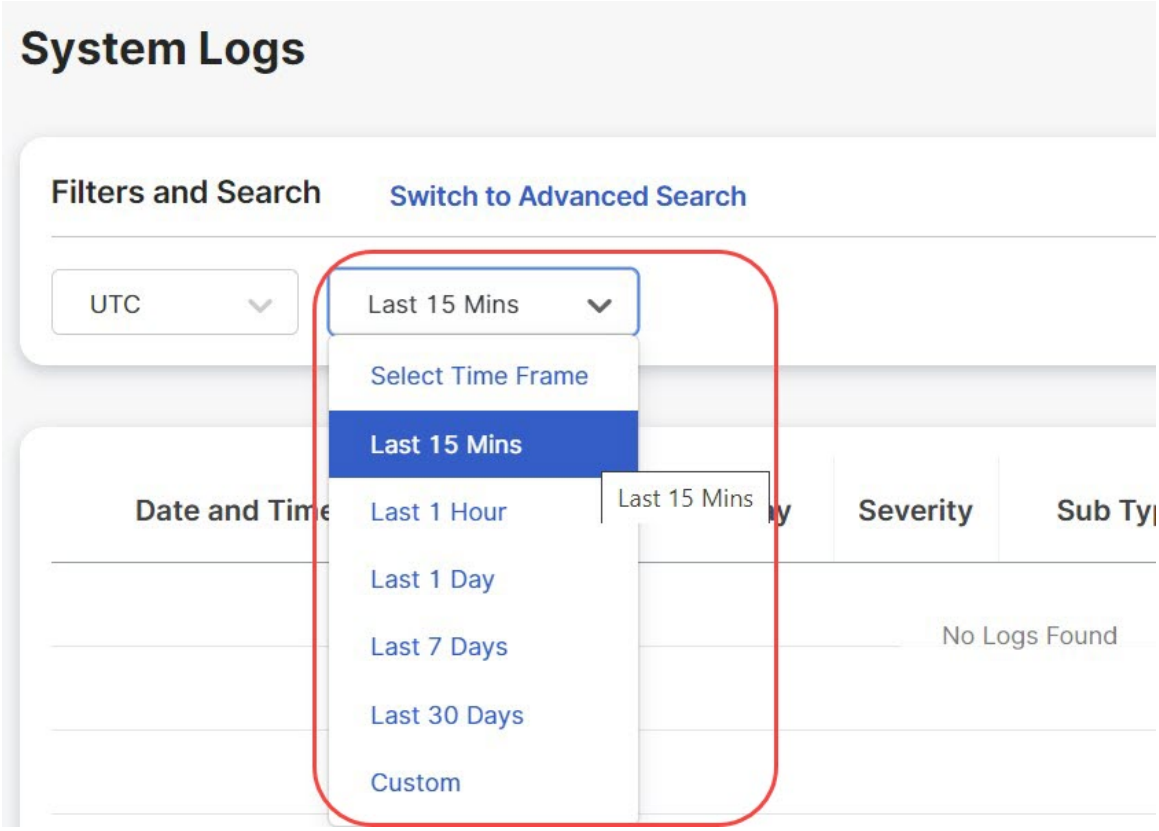
To select or switch between different time formats, click the radio button as shown:



Timeframe

You can display logs in increment options from 15 minutes to 30 days, or Custom timeframes.

To select or switch between timeframes, click the drop-down-menu and select a timeframe as shown:



For Custom timeframes, select **Custom**, the **Start**, and the **End** date or time by clicking the calender objects followed by **Save**.

Severity Levels

The severity levels of system logs are:

- `Info` - Informational details such as sign in, sign out, password changes, configuration changes and so on. These contain events that do not qualify as other severity levels.
- `Warning` - Notifications that inform you of a possible system action or change, for example, password updates.
- `Medium` - Issues that are medium in severity such as package upgrades and so on.
- `High` - Serious issues such as network disconnections with external devices and so on.
- `Critical` - Major issues that are critical in nature such as hardware failures and so on.

Search Filter

Search and filter System Logs by gateway, cloud account, or message content to find specific system events and operational changes.

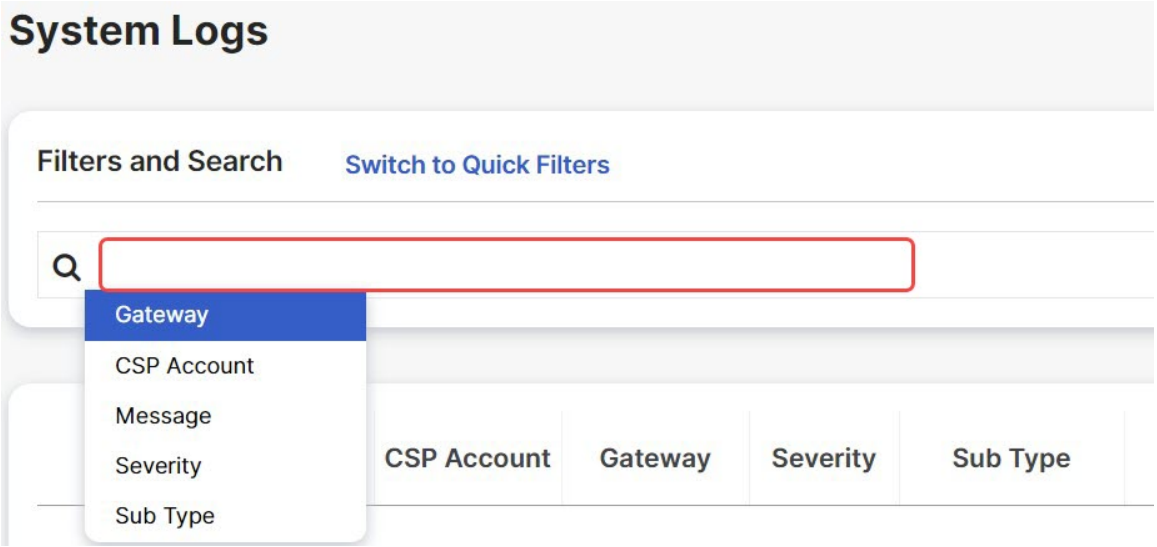
Logs can be filtered using the Search function and System log fields.

The System log fields are `Gateway`, `CSP Account`, and `Message`.

To filter System logs on one or multiple fields:

Procedure

1. Left mouse-click in the Search field to access the pull down menu.



2. Select a field, for example `Gateway`.
3. Type a desired search string, for example `ingress`.
4. Add additional fields to the search criteria, as required.

   Example: Filter for a Gateway = **"ingress"** and Messages containing **"created"** would appear in the filter criteria and results.

28 Threat Research

Topics:

- Network Intrusion
- Web Protection
- Malicious Sources

Explains how threat researches identify and organize rules for detecting malicious activity.

Threat Research is generated from a set of rules that are applied to the inspection engine to detect threats and malicious activity. This page allows you to view these rules. Once a day, Multicloud Defense searches for new or modified rules for network intrusion and includes or removes rules and known malicious sources from the internal library. This action is automated. Included in this function is the act of downloading and validating the new list of IP addresses as sources and implementing them in new rulesets. These rulesets are then deployed.

The rules have a variety of ways in which they are organized such as policy, class, application, ruleset library date, and other parameters. If you are interested in understanding more about a rule that has tripped (e.g., detected a threat or malicious activity), use the **Threat Research** page to view more details about the rule. The following parts of each of the page are available for your use:

Search Bar

The search bar at the top of the window allows you to search each page under threat research for any singular dientifying facet: a known IP address, action, rule name, gateway name, attack type, or profile name. If you find a specific field value by scrolling, you can **Add to Search** to facilitate an easier search experience.

Note that the searches are isolated to each page, and you cannot cross-search the different types of threat research. See the section below for more details.

View Details

Each of the facets under threat research offer the ability to **View Details** of a singular incident or attack. The values provided in these details differ between the types of threat research, but can be valuable if you want to finetune your policies, security profiles, rules or rulesets.

Add to Search

For any of the types of research available here, you can click on any one value within a row and automatically have the option to **Add to Search**. This automatically applies the selected value to the search bar at the top of the window and filters the viewing window to the content in the search bar. You can do this multiple times and the values you select compound into a complex search request.

Network Intrusion

Find details about detected network intrusions, including affected gateways, triggered profiles, attack classifications, impacted services, severity, rule content, CVSS scores, and CVEs.

Network intrusion refers to any unauthorized activity on your network. Note that this table does not include the built-in rules to the IDS/IPS engine or any affiliated information from these rules; these rules are designated for detection only; the remainder of the IDS/IPS rules are configured to protect and perform actions based on the varying levels of intrusion or attack.

The Network Intrusion page displays the following:

- **Gateway Names** - the names of the affected gateways that processed the malicious source.
- **Profile Names** - the names of the security profiles triggered by the malicious source.
- **IPS Policy** - the policy within Multicloud Defense triggered by the event or attack.
- **IPS Class** - the type of attack as determined by the database of attack signatures traffic is compared against.
- **IPS Category** - the IPS signature category triggered by the event or attack.
- **Rule ID** - the rule ID as documented internally within Multicloud Defense that was triggered by the event or attack.
- **Services Impacted** - the type of web service affected by the event or attack.
- **Impact** - the severity level of impact, known or assumed, by the event or attack.
- **Message** - the contents of the event that has been identified as an attack.
- **Rule Content** - content of the rule triggered by the event or attack.
- **CVSS Score** - Common Vulnerability Scoring System (CVSS) is a framework that assigns a numerical score to the severity of an information security vulnerability. CVSS scores range from 0 to 10, with 10 being the most severe.
- **CVEs** - Common Vulnerabilities and Exposures (CVE) is a glossary that classifies vulnerabilities. If there is a CVE associated with the type of attack or event, the internal library automatically generates its value here.
- **References** - If publicly available, this link directs you to the original announcement and categorization of the CVE.

Web Protection

Find details about web threats detected by Web Application Firewall profiles, including affected gateways, rule categories, inspection methods, attack types, platforms, and development languages.

The Web Application Firewall (WAF) research is displayed as 'Web Protection.' This lets you secure your devices against web threats and helps you regulate unwanted content. The Web Protection page displays the following:

- **Gateway Names** - the names of the affected gateways that processed the malicious source.
- **Profile Names** - the names of the security profiles triggered by the malicious source.
- **CRS Category** - the Core Rule Set (CRS) category identified per set of generic attack detection rules.
- **Inspection Type** - the type of inspection Multicloud Defense peformed on the traffic that encapsulated the attack or event.
- **Attack Type** - the type of unauthorized attack traversed over a network.
- **Platform** - the platform type identified from the attack or event.
- **Lanuage** - the noted web development language detected from the event.

Malicious Sources

Find details about malicious sources detected by Cisco Multicloud Defense, including affected gateways, triggered security profiles, response actions, impact, and source IP addresses.

Malicious sources are any type of code or packet that causes harm to a network. The Malicious Sources page dispalys the following:

- **Gateway Names** - the names of the affected gateways that processed the malicious source.
- **Profile Names** - the names of the security profiles triggered by the malicious source.
- **Malicious Source Action** - the action taken when the malicious source was identified.
- **Impact** - the impact of the malicious material deteremined by how it is ranked within the library.
- **Malicious Source IP** - the IP address where the malicious source originates from.

29 Cloud Visibility Reports

Topics:

- Generate a Discovery Report
- Generate a Threat And Cloud Analytics Report

Explains the various types of reports available for cloud visibility and what insights these can offer into the network's health.

Reports provide valuable statistical information that you can use as insights to the network and its general health, and decide accordingly. Multicloud Defense enables you to generate the following reports:

Discovery

The Discovery Report is generated by taking out-of-band traffic information from DNS queries and VPC flow logs, and correlating the data with threat intelligence and cloud inventory information. These logs are available only if you configure the VPC of your cloud service provider to send logs to an S3 bucket. These logs are then transferred directly to the Multicloud Defense Controller.

The report contains:

- **Top of Discovery Report** - Network and cloud asset analytics, presented in volume and distinct counts of field values. You can derive insights on what is happening in your cloud environment by quantifying network behavior.
- **Network Traffic – Bytes** - This graph displays the volume of bytes by traffic direction. You can view where the volume of bytes is going - Ingress, Egress, or East-West.
- **Network Traffic – Packets** - This graph displays the volume of packets by traffic direction. You can view where the volume of packets is going - Ingress, Egress, or East-West.
- **Network Traffic – Events** - This graph displays the volume of events by traffic direction. You can view where the volume of events is going - Ingress, Egress, or East-West.
- **Ingress Account Summary** - This summary displays the distinct count of cloud assets with ingress network traffic, by CSP. You can view the flow of assets communicating into a CSP environment.
- **Egress Account Summary** - This summary displays the distinct count of cloud assets with egress network traffic, by Cloud Service Provider (CSP). You can view review the flow of assets communicating out of a CSP environment.
- **Ingress Network Events by Country** - This geographic heatmap shows the volume of ingress traffic by country. You can view countries that communicate with the cloud environments.
- **Egress Network Events by Country** - This geographic heatmap shows the volume of egress traffic by country. You can view countries that communicate with the cloud environments.
- **Top 10 Source Countries** - This graph displays the top 10 source countries by volume of events, with other network analytics. This is a summary of the top source countries that the cloud environment is communicating with.

- **Top 10 Destination Countries** - This graph displays the top 10 destination countries by volume of events, with other network analytics. This is a summary of the top destination countries that the cloud environment is communicating with.
- **Top 10 Ingress Source IP Addresses** - This graph displays the top 10 source IP addresses by volume, with other network analytics. You can view the entities that create the most inbound events.
- **Top 10 Egress Destination IP Addresses** - This graph displays the top 10 destination IP addresses by volume, with other network analytics. You can view entities that the cloud environment mostly communicates with.
- **Top 10 FQDN Category Names by Volume** - This graph displays the category names by volume for FQDNs. You can view the top category types based on the FQDNs being requested by the cloud environment.
- **Top 10 FQDNs by Volume** - This graph displays the top 10 FQDNs by volume. You can view the top FQDNs that are requested by the cloud environment.
- **Top 10 Malicious FQDNs by Volume** - This graph displays the top 10 malicious FQDNs by volume. If a malicious or suspicious category name is found, the top FQDNs in that category name are displayed here.
- **FQDN Category Name Mapped to MITRE ATT&CK** - This graph displays the top 10 malicious category names mapped to MITRE ATT&CK. This view provides more context on how the FQDN category name relates to an attack chain by using the Enterprise MITRE ATT&CK framework.

Threat Indicators Snapshot

The Cloud Visibility Report - Threat Indicators Snapshot is a compilation of data on the gateway instance. You can use this report to determine the gateway's endurance under duress by examining traffic patterns, when and how thresholds are met, trends of attacks, and specific instances. The report includes the following points:

- **IDS/IPS Detection** - This data shows how many attacks are detected, the type of attack, the time of the detected attacks, and the top ten IDS/IPS signatures over the time range selected.
- **WAF Detection** - This data shows how many attacks are detected by WAF rules, the time of the detected attacks, and the top ten WAF signatures over the time range selected.
- **Relocation of Threats by Volume** - This choropleth map shows the volume of attacks for both WAF and IDS/IPS events by country in volume.
- **Top Ten Attacking Countries by Volume and Time** - This horizontal bar chart depicts the volume of the top 10 countries that during the entirety of the timespan produced the most events, then displayed breaking that volume across the time increments for which the events occurred during the timespan.
- **Policy and Prevention** - This data chart shows the action that is taken by the gateway security type in whichever CSP environment it is deployed in. This includes the type of action, how many events generated from the action, the gateway security type and more.

You **must** have Web Application Firewall (WAF), intrusion detection and prevention (IDS/IPS) rules that are enabled in your policy for the Multicloud Defense Gateway to collect and poll data.

For Additional Information:

Generate a Discovery Report

Generate a Discovery Report from DNS queries and VPC flow logs for a selected reporting period, then save the report for offline review.

A discovery report is generated by taking DNS queries and VPC flow logs that have been sent to an S3 bucket prior to getting processed by the Multicloud Defense Controller.

Procedure

1. In the Multicloud Defense Controller page, navigate to **Reports** > **Discovery**.
2. Under Threat & Cloud Analytics Report, select the **Frequency** from the drop-down list for the data that is pulled: daily, weekly, monthly, quarterly, or yearly.
   - **Daily** - From 12 a.m. for 24 hours. This is in UTC time.
   - **Weekly** - From Monday to Sunday.
   - **Monthly** - Generally from the beginning to the end of the month.
   - **Quarterly** - From the beginning to end of a quarter. Quarters are generally defined as from January 1 - March 31, April 1 - June 30, July 1 - September 30, and October 1 - December 31.
   - **Yearly** - From January 1 to December 31 of the year selected.
3. Select a date. Use the drop-down **Calendar** to select the time range or specific days that you want to collect data on. Days that are grayed out have no data to compile. If you have no data available to generate a report, confirm your policies contain WAF and IDS/IPS rules.
4. Click **Generate Report**. The Discovery Report is generated in a new tab.
5. To save the report locally, click **Print Report**. Navigate to a location on your local server and save the report.

Generate a Threat And Cloud Analytics Report

Generate a Threat and Cloud Analytics Report from inspected traffic to review threat indicators for a selected reporting period and save the results.

The Threat and Cloud Analytics Report is a **Threat Indicator Snapshot** that is generated by using the traffic collected and inspected by Multicloud Defense Gateway. This provides a more comprehensive report as Multicloud Defense is now in the datapath and compliments the discovery report.

Note that reports cannot be generated for the day, since a qualitative summarization of events cannot be made until end of day, end of month, end of quarter, or end of year.

Note

You **must** have Web Application Firewall (WAF), intrusion detection and protection (IDS/IPS) rules enabled in your policy in order for the Multicloud Defense Gateway to collect and poll data. For more information, see the following links respecitvely:

- Web Application Firewall
- Network Intrusion (IDS/IPS) Profile on page 190

Use the following procedure to generate a Threat And Cloud Analytics with the threat indicators snapshot:

Procedure

1. In the Multicloud Defense Controller page, navigate to **Reports** > **Threat Indicators Snapshot**.
2. Under Threat & Cloud Analytics Report, select the **Frequency** from the drop-down list for the data that is pulled: daily, weekly, monthly, quarterly, or yearly.
   - **Daily** - From 12 AM for 24 hours. This is in UTC time.
   - **Weekly** - From Monday to Sunday.
   - **Monthly** - Generally from the beginning to the end of the month.
   - **Quarterly** - From the beginning to end of a quart. Quarters are generally defined as from January 1 - March 31, April 1 - June 30, July 1 - September 30, and October 1 - December 31.
   - **Yearly** - From January 1 to December 31 of the year selected.
3. Select a date. Use the drop-down **Calendar** to select the time range or specific days that you want to collect data on. Days that are grayed out have no data to compile. If you have no data available to generate a report, confirm your policies contain WAF and IDS/IPS rules.
4. Click **Generate Report**.
5. The report is generated. To save the report locally, click **Print Report**. Navigate to a location on your local server and save the report.

30 Alerting Overview

Topics:

- Alert Services Overview

Explains the purpose and key attributes of system alerts.

Alert Services Overview

Learn how Alert Service integrations send critical Cisco Multicloud Defense system alerts to Microsoft Sentinel, PagerDuty, ServiceNow, or Slack for timely response.

To integrate with widely deployed alerting services, Multicloud Defense integrates with Microsoft Sentinel, PagerDuty, ServiceNow and Slack to forward critical system level alerts. This enables cloud operations teams to be alerted, and respond to, user-defined system events and severity levels detected by the Multicloud Defense Cloud Controller. This is accomplished within Multicloud Defense Controller using an Alert Service Profile, together with an Alert Rule, for a given integration.

To configure integrations with supported alerting services, navigate to: **System and Accounts** > **Service Alerts** > **Services**

Integration with these services require either an API URL, API key, or both. Generally, the API Keys and URLs need to be generated by your Organization's Administrator of these services.

> **Note**
>
> For ServiceNow integrations, a Webhook must be configured to enable ServiceNow to receive and display alerts from the Multicloud Defense Controller.

31 Alert Destinations / SIEMs

Topics:

- Datadog
- Microsoft Sentinel
- PagerDuty
- ServiceNow
- Slack
- Microsoft Teams
- Webex
- Splunk

Explains how service rules with alert destinations are used to notify web services or applications.

A service rule that includes an alert destination typically defines how and where notifications or alerts should be sent when certain conditions are met. Service rules monitor specific conditions or events, such as system thresholds being exceeded, errors occurring, or security incidents being detected and when the predefined conditions are met, the service rule triggers an alert. Alerts are sent to web services or applications via webhooks for further processing or automation. This ensures that critical events do not go unnoticed.

Datadog

Learn how the Datadog integration sends Cisco Multicloud Defense alerts to Datadog through an Alert Service Profile and associated Alert Rule.

Once configured, Multicloud Defense alerts will sent to Datadog using the defined Alert Service Profile and Alert Rule.

Create an Alert Profile Service

Create a Datadog Alert Service Profile with an API key so Cisco Multicloud Defense can authenticate and send alerts to Datadog.

Before you begin

In order to send alerts to Datadog, the following information is required:

- Datadog account
- API Key

> **Tip**
>
> - To Sign up for a Datadog account, refer to Datadog Account (https://www.datadoghq.com/).
> - To create a Datadog API Key, refer to Datadog API Key (https://app.datadoghq.com/account/login?next=%2Faccount%2Fsettings#api).

Procedure

1. Navigate to **System and Accounts** > **Service Alerts** > **Services**.
2. Click **Create**.
3. **Name** - Enter unique name for the alert integration. Example `Multicloud Defense-Datadog-profile`.
4. **Description** (optional) - Enter a description for the alert integration.
5. **Type** - Using the pulldown, choose **Datadog**.
6. **API Key** - Specify the Datadog API Key used to authenticate the communication.
7. Click **Save**.

What to do next

Create an alert rule with this new profile.

Create an Alert Rule

Create a Datadog Alert Rule that uses an existing alert profile and selects the event types, subtypes, and severity levels sent to Datadog.

Before you begin

In order to send alerts to Datadog, the following information is required:

- Datadog account
- API Key

Tip

- To Sign up for a Datadog account, refer to Datadog Account (https://www.datadoghq.com/).
- To create a Datadog API Key, refer to Datadog API Key (https://app.datadoghq.com/account/login?next=%2Faccount%2Fsettings#api).

Procedure

1. Navigate to **System and Accounts** > **Service Alerts** > **Alert Rules**.
2. Click **Create**.
3. **Profile Name** - Enter unique name for the integration. Example `Multicloud Defense-Datadog-alert-rule`.
4. **Description** (optional) - Enter a description for the alert rule.
5. **Alert Profile** - Using the pulldown, choose a PagerDuty Alert Profile. As example, select profile created above `Multicloud Defense-Datadog-profile`.
6. (Optional)**Description** - Enter a description for the aler trule.
7. **Alert Profile** - Expand the drop-down menu and select a Microsoft Teams alert profile.
8. **Type** - Expand the drop-down menu and select one of the following types:
   - System Logs
   - Audit Logs
   - Discovery

   If you select **Audit Logs**, there are no other configurable items. Click **Save** to finalize the rule.
9. If you select either System Logs or Discovery as your Type, then expand the **Sub Type** drp-down menu and select one of the following options:
   - Gateway
   - Account
   - Controller
10. Expand the **Severity**drop-down menu and select one of the following labels. Note that the options below are dependent on the **Type** you selected in step 7.
    - Info
    - Warning
    - Medium
    - High
    - Critical

11. **Enabled** - This option is checked by default to enable and implement this alert immediately after saving. Unceck this box if you do not to immediately apply the rule to your environment.
12. Click **Save**.

Microsoft Sentinel

Learn how the Microsoft Sentinel integration sends Cisco Multicloud Defense alerts to Sentinel through an Alert Service Profile and associated Alert Rule.

Once configured, Multicloud Defense alerts will sent to Microsoft Sentinel using the defined Alert Service Profile and Alert Rule.

Create an Alert Profile Service

Create a Microsoft Sentinel Alert Service Profile with Azure Log Analytics workspace credentials and table details so Cisco Multicloud Defense can send alerts.

Before you begin

In order to send alerts to Microsoft Sentinel, the following information is required:

- Create an Azure Log Analytics Workspace.
- Define an Azure Log Table.

Procedure

1. Navigate to **System and Accounts** > **Service Alerts** > **Services**.
2. Click **Create**.
3. **Name** - Enter unique name for the alert integration. Example `mcd-mssentinel-profile`.
4. **Description** (optional) - Enter a description for the alert integration.
5. **Type** - Using the pulldown, choose **Microsoft Sentinel**.
6. **API Key** - Specify the Shared Key created in Azure for the Azure Log Analytics Workspace.
7. **Azure Log Table Name** - Specify the name of the Azure Log defined when creating the Azure Log Analytics Workspace.
8. **Azure Log Analytics Workspace ID** - Specify the ID of the Azure Log Analytics Workspace.
9. Click **Save**.

What to do next

Create an alert rule with this new profile.

Create an Alert Rule

Create a Microsoft Sentinel Alert Rule that selects the system, audit, or discovery events and severity levels sent through an existing Sentinel profile.

Before you begin

In order to send alerts to Microsoft Sentinel, the following information is required:

- Create an Azure Log Analytics Workspace.
- Define an Azure Log Table.

Procedure

1. Navigate to **System and Accounts** > **Service Alerts** > **Alert Rules**.
2. Click **Create**.
3. **Profile Name** - Enter unique name for the integration. Example `mcd-mssentinel-alert-rule`.
4. (Optional)**Description** - Enter a description for the aler trule.
5. **Alert Profile** - Expand the drop-down menu and select a Microsoft Teams alert profile.
6. **Type** - Expand the drop-down menu and select one of the following types:
   - System Logs
   - Audit Logs
   - Discovery

   If you select **Audit Logs**, there are no other configurable items. Click **Save** to finalize the rule.
7. If you select either System Logs or Discovery as your Type, then expand the **Sub Type** drp-down menu and select one of the following options:
   - Gateway
   - Account
   - Controller
8. Expand the **Severity**drop-down menu and select one of the following labels. Note that the options below are dependent on the **Type** you selected in step 7.
   - Info
   - Warning
   - Medium
   - High
   - Critical
9. **Enabled** - This option is checked by default to enable and implement this alert immediately after saving. Unceck this box if you do not to immediately apply the rule to your environment.
10. Click **Save**.

PagerDuty

Learn how the PagerDuty integration sends Cisco Multicloud Defense alerts to a PagerDuty API endpoint through an Alert Service Profile and associated Alert Rule.

Once configured, Multicloud Defense alerts will sent to a PagerDuty API gateway using the defined Alert Service Profile and Alert Rule.

Create an Alert Profile Service

Create a PagerDuty Alert Service Profile with an API key so Cisco Multicloud Defense can authenticate and send alerts to PagerDuty.

Before you begin

In order to complete the steps in this guide, you will need:

- A PagerDuty account with an API Key configured.

> **Tip**
>
> - Sign up for a PagerDuty account (https://www.servicenow.com/my-account/sign-up.html).
> - Setup the API Key (https://developer.pagerduty.com/api-reference).

Procedure

1. Navigate to **System and Accounts** > **Service Alerts** > **Services**.
2. Click **Create**.
3. **Name** - Enter unique name for the alert integration. Example `mcd-pagerduty-profile`.
4. **Description** (optional) - Enter a description for the alert integration.
5. **Type** - Using the pulldown, choose **PagerDuty**.
6. **API Key** - Copy the PagerDuty API key generated above, or other PagerDuty API Key as desired.
7. Click **Save**.

What to do next

Create an alert rule with this new profile.

Create an Alert Rule

Create a PagerDuty Alert Rule that selects the system, audit, or discovery events and severity levels sent through an existing PagerDuty profile.

Before you begin

In order to complete the steps in this guide, you will need:

A PagerDuty account with an API Key configured.

Tip

- Sign up for a PagerDuty account (https://www.servicenow.com/my-account/sign-up.html).
- Setup the API Key (https://developer.pagerduty.com/api-reference).

Procedure

1. Navigate to **System and Accounts** > **Service Alerts** > **Alert Rules**.
2. Click **Create**.
3. **Profile Name** - Enter unique name for the integration. Example `mcd-pagerduty-alert-rule`.
4. **Description** (optional) - Enter a description for the alert rule.
5. (Optional)**Description** - Enter a description for the aler trule.
6. **Alert Profile** - Expand the drop-down menu and select a Microsoft Teams alert profile.
7. **Type** - Expand the drop-down menu and select one of the following types:
   - System Logs
   - Audit Logs
   - Discovery

   If you select **Audit Logs**, there are no other configurable items. Click **Save** to finalize the rule.
8. If you select either System Logs or Discovery as your Type, then expand the **Sub Type** drp-down menu and select one of the following options:
   - Gateway
   - Account
   - Controller
9. Expand the **Severity**drop-down menu and select one of the following labels. Note that the options below are dependent on the **Type** you selected in step 7.
   - Info
   - Warning
   - Medium
   - High
   - Critical
10. **Enabled** - This option is checked by default to enable and implement this alert immediately after saving. Unceck this box if you do not to immediately apply the rule to your environment.
11. Click **Save**.

ServiceNow

Learn how the ServiceNow integration sends Cisco Multicloud Defense alerts to ServiceNow through an Alert Service Profile, webhook, and associated Alert Rule.

Once configured, Multicloud Defense alerts will sent to a ServiceNow API gateway using the defined Alert Service Profile and Alert Rule.

Create an Alert Rule

Create a ServiceNow Alert Rule that selects the system, audit, or discovery events and severity levels sent through an existing ServiceNow profile.

Before you begin

In order to complete the steps in this guide, you will need:

- A ServiceNow account with an Incoming Webhook URL.
- An API Key configured.

> **Tip**
>
> - Sign up for a ServiceNow account (https://www.servicenow.com/my-account/sign-up.html)
> - Setup Webhook and API Key (https://docs.servicenow.com/search?q=setup%20webhook)

Procedure

1. Navigate to **System and Accounts** > **Service Alerts** > **Alert Rules**.
2. Click **Create**.
3. **Profile Name** - Enter unique name for the integration. Example `mcd-servicenow-alert-rule`.
4. (Optional)**Description** - Enter a description for the aler trule.
5. **Alert Profile** - Expand the drop-down menu and select a Microsoft Teams alert profile.
6. **Type** - Expand the drop-down menu and select one of the following types:
   - System Logs
   - Audit Logs
   - Discovery

   If you select **Audit Logs**, there are no other configurable items. Click **Save** to finalize the rule.
7. If you select either System Logs or Discovery as your Type, then expand the **Sub Type** drp-down menu and select one of the following options:
   - Gateway
   - Account
   - Controller

8. Expand the **Severity**drop-down menu and select one of the following labels. Note that the options below are dependent on the **Type** you selected in step 7.
   - Info
   - Warning
   - Medium
   - High
   - Critical
9. **Enabled** - This option is checked by default to enable and implement this alert immediately after saving. Unceck this box if you do not to immediately apply the rule to your environment.
10. Click **Save**.

Create an Alert Profile Service

Create a ServiceNow Alert Service Profile with an API key and incoming webhook URL so Cisco Multicloud Defense can send alerts to ServiceNow.

Before you begin

In order to complete the steps in this guide, you will need:

- A ServiceNow account with an Incoming Webhook URL.
- API Key configured.

> **Tip**
>
> - Sign up for a ServiceNow account (https://www.servicenow.com/my-account/sign-up.html)
> - Setup Webhook and API Key (https://docs.servicenow.com/search?q=setup%20webhook)

Procedure

1. Navigate to **System and Accounts** > **Service Alerts** > **Services**.
2. Click **Create**.
3. **Name** - Enter unique name for the alert integration. Example `mcd-servicenow-profile`.
4. **Description** (optional) - Enter a description for the alert integration.
5. **Type** - Using the pulldown, choose **ServiceNow**.
6. **API Key** - Specify the ServiceNow API key generated above, or other ServiceNow API Key as desired.
7. **API URL** - Specify the ServiceNow Webhook URL generated above, or other ServiceNow Webhook URL as desired.
8. Click **Save**.

What to do next

Create an alert rule with this new profile.

Slack

Learn how the Slack integration sends Cisco Multicloud Defense alerts to an incoming webhook through an Alert Service Profile and associated Alert Rule.

Once configured, Multicloud Defense alerts will sent to a Slack Incoming Webhook URL using the defined Alert Service Profile and Rule.

Create a Slack Alert Rule

Create a Slack Alert Rule that selects the system, audit, or discovery events and severity levels sent through an existing Slack profile.

Before you begin

You must create a Slack alert profile before you create an alert rule.

Procedure

1. Navigate to **System and Accounts** > **Service Alerts** > **Alert Rules**.
2. Click **Create**.
3. **Profile Name** - Enter unique name for the integration. Example `mcd-slack-alert-rule`.
4. (Optional)**Description** - Enter a description for the aler trule.
5. **Alert Profile** - Expand the drop-down menu and select a Microsoft Teams alert profile.
6. **Type** - Expand the drop-down menu and select one of the following types:
   - System Logs
   - Audit Logs
   - Discovery

   If you select **Audit Logs**, there are no other configurable items. Click **Save** to finalize the rule.
7. If you select either System Logs or Discovery as your Type, then expand the **Sub Type** drp-down menu and select one of the following options:
   - Gateway
   - Account
   - Controller
8. Expand the **Severity**drop-down menu and select one of the following labels. Note that the options below are dependent on the **Type** you selected in step 7.
   - Info
   - Warning
   - Medium
   - High
   - Critical
9. **Enabled** - This option is checked by default to enable and implement this alert immediately after saving. Unceck this box if you do not to immediately apply the rule to your environment.

10. Click **Save**.

Create an Alert Profile Service

Create a Slack Alert Service Profile with an incoming webhook URL so Cisco Multicloud Defense can send system alerts to Slack.

Before you begin

In order to complete the steps in this guide, you will need:

- A Slack account with an incoming webhook URL configured.

> **Tip**
>
> 1. Create a Slack account (https://slack.com/get-started#/create).
> 2. Create an incoming Webhook (https://slack.com/help/articles/115005265063-Incoming-webhooks-for-Slack#set-up-incoming-webhooks).

Procedure

1. Navigate to **System and Accounts** > **Service Alerts** > **Services**.
2. Click **Create**.
3. **Name** - Enter unique name for the alert integration. Example `mcd-slack-profile`.
4. **Description** (optional) - Enter a description for the alert integration.
5. **Type** - Using the pulldown, choose **Slack**.
6. **API URL** - Specify the Slack Webhook URL generated above, or other Slack Webhook URL as desired.

What to do next

Create an alert rule with this new profile.

Microsoft Teams

Learn how the Microsoft Teams integration uses an alert profile and service rule to deliver Cisco Multicloud Defense alerts to a Teams incoming webhook.

Create an alert profile and then use that profile in a service alert rule to generate alerts specifically for Microsoft Teams using a uique incoming webhook from the service.

Create a Microsoft Teams Alert Profile Service

Create a Microsoft Teams Alert Service Profile with an incoming webhook URL so Cisco Multicloud Defense can send alerts to Teams.

Before you begin

You must do the following items before you finalize the alert profile for your Microsoft account:

- You **must** create an Incoming Webhook in the Microsoft Teams UI. See Microsoft documentation for more information.
- You **must** save the unique API URL generated from creating an incoming webhook for the procedure below.

Procedure

1. Navigate to **System and Accounts** > **Service Alerts** > **Services**.
2. Click **Create**.
3. **Name** - Enter unique name for the alert integration.
4. (Optional) **Description** - Enter a description for the alert integration.
5. **Type** - Using the pulldown, choose **Microsoft Teams**.
6. **API URL** - Enter the API URL that is generated from creating an incoming webhook. Copy the URL from your Microsoft Teams UI and paste it in this text field.
7. Click **Save**.

What to do next

Create an alert rule with this new profile.

Create a Microsoft Teams Service Rule

Create a Microsoft Teams Service Rule that selects the system, audit, or discovery events and severity levels sent through an existing Teams profile.

Before you begin

You must create a Microsoft Teams service profile before you create a service rule.

Procedure

1. Navigate to **System and Accounts** > **Service Alerts** > **Alert Rules**.
2. Click **Create**.
3. **Profile Name** - Enter unique name for the integration. An example would be similar to `mcd-microsoft-alert-rule`.
4. (Optional)**Description** - Enter a description for the aler trule.
5. **Alert Profile** - Expand the drop-down menu and select a Microsoft Teams alert profile.
6. **Type** - Expand the drop-down menu and select one of the following types:
   - System Logs
   - Audit Logs
   - Discovery

   If you select **Audit Logs**, there are no other configurable items. Click **Save** to finalize the rule.
7. If you select either System Logs or Discovery as your Type, then expand the **Sub Type** drp-down menu and select one of the following options:
   - Gateway

- Account
- Controller

8. Expand the **Severity**drop-down menu and select one of the following labels. Note that the options below are dependent on the **Type** you selected in step 7.
   - Info
   - Warning
   - Medium
   - High
   - Critical
9. **Enabled** - This option is checked by default to enable and implement this alert immediately after saving. Unceck this box if you do not to immediately apply the rule to your environment.
10. Click **Save**.

Webex

Learn how the Webex integration sends Cisco Multicloud Defense alerts to a Webex space through an Alert Service Profile and associated Alert Rule.

Once configured, Multicloud Defense alerts will sent to a Webex API gateway using the defined Alert Service Profile and Alert Rule.

Create an Alert Profile Service

Create a Webex Alert Service Profile with an incoming webhook URL so Cisco Multicloud Defense can send alerts to a Webex space.

Use the following procedure to create an alert profile for the Webex service:

Before you begin

In order to complete the steps in this guide, you will need:

- A Webex account with an Incoming Webhook URL.
- API Key configured.

Note

1. Create or access a Webex account.
2. Create a Webex Incoming Webhook.
3. Accept the Incoming Webhook permissions.
4. Provide a Name and select a Webex Space.
5. Copy the Webex Webhook URL to use in the configuration of the Alert Service Profile.

Procedure

1. Navigate to **System and Accounts** > **Service Alerts** > **Services**.
2. Click **Create**.
3. **Name** - Enter unique name for the alert integration.
4. (Optional) **Description** - Enter a description for the alert integration.
5. **Type** - Using the pulldown, choose **Webex**.
6. **API URL** - Specify the Webex Webhook URL generated as part of the prerequisites, or other Webex Webhook URL as desired.
7. Click **Save**.

What to do next

Create an alert rule with this new profile.

Create an Alert Rule

Create a Webex Alert Rule that selects the system, audit, or discovery events and severity levels sent through an existing Webex profile.

Procedure

1. Navigate to **System and Accounts** > **Service Alerts** > **Alert Rules**.
2. Click **Create**.
3. **Profile Name** - Enter unique name for the integration. An example is `mcd-servicenow-alert-rule`.
4. (Optional)**Description** - Enter a description for the aler trule.
5. **Alert Profile** - Expand the drop-down menu and select a Microsoft Teams alert profile.
6. **Type** - Expand the drop-down menu and select one of the following types:
   - System Logs
   - Audit Logs
   - Discovery

   If you select **Audit Logs**, there are no other configurable items. Click **Save** to finalize the rule.

7. If you select either System Logs or Discovery as your Type, then expand the **Sub Type** drp-down menu and select one of the following options:
   - Gateway
   - Account
   - Controller
8. Expand the **Severity**drop-down menu and select one of the following labels. Note that the options below are dependent on the **Type** you selected in step 7.
   - Info
   - Warning
   - Medium
   - High
   - Critical
9. **Enabled** - This option is checked by default to enable and implement this alert immediately after saving. Unceck this box if you do not to immediately apply the rule to your environment.
10. Click **Save**.

Splunk

Learn how the Splunk integration sends Cisco Multicloud Defense alerts to Splunk through a configured Alert Service Profile and associated Alert Rule.

Once configured, Multicloud Defense alerts will sent to an API gateway using the defined Alert Service Profile and Alert Rule.

Create a Splunk Profile Service

Create a Splunk Alert Service Profile with HTTP Event Collector connection and authentication details so Cisco Multicloud Defense can send alerts to Splunk.

Use the following procedure to create an alert profile for the Splunk service:

Before you begin

You must have the following configured and ready:

- Create an API Key in Multicloud Defense and store both the key and secret. See Create an API Key in Multicloud Defense on page 310 for more information.
- Set up the HTTP Event Collector (HEC) in Splunk Web. See Configure HTTP Event Collector on Splunk Cloud for more information.
- Your Splunk HEC must have the following configured:
  - HEC must be **enabled**.
  - You must have at least one active HEC token available.
  - You must use an active token to authenticate into HEC.

- You must format the data that goes to HEC in a certain way. See Format events for HTTP Event Collector.

Procedure

1. Navigate to **System and Accounts** > **Service Alerts** > **Services**.
2. Click **Create**.
3. **Name** - Enter unique name for the alert integration.
4. **Description** (optional) - Enter a description for the alert integration.
5. **Type** - Using the pulldown, choose **Splunk**.
6. **API Key** - Copy the Splunk API key generated above, or other PagerDuty API Key as desired.
7. Check the **Skip Verify Certificate** box if your server doesnt have certificates with SAN field matching with domain. If you server does have ceritficats with SAN fields matching the domain, leave this unchecked.
8. **Index(default - main)** is Splunk's default index where all the processed data is stored. This is provided when you configure the Splunk HEC.
9. Enter the **API URL** for the Splunk HTTP Event Collector. We recommend this URL `https://<host>:<port>/services/collector .`
10. Click **Save**.

What to do next

Create an alert rule with this new profile.

Create a Splunk Rule

Create a Splunk Alert Rule that selects the system, audit, or discovery events and severity levels sent through an existing Splunk profile.

Use the following procedure to create a rule containing the splunk alert service:

Procedure

1. Navigate to **System and Accounts** > **Service Alerts** > **Alert Rules**.
2. Click **Create**.
3. **Profile Name** - Enter unique name for the integration. Example `mcd-mssentinel-alert-rule`.
4. (Optional)**Description** - Enter a description for the aler trule.
5. **Alert Profile** - Expand the drop-down menu and select a Microsoft Teams alert profile.
6. **Type** - Expand the drop-down menu and select one of the following types:
   - System Logs
   - Audit Logs
   - Discovery

   If you select **Audit Logs**, there are no other configurable items. Click **Save** to finalize the rule.
7. If you select either System Logs or Discovery as your Type, then expand the **Sub Type** drp-down menu and select one of the following options:
   - Gateway

- Account
- Controller

8. Expand the **Severity**drop-down menu and select one of the following labels. Note that the options below are dependent on the **Type** you selected in step 7.
   - Info
   - Warning
   - Medium
   - High
   - Critical
9. **Enabled** - This option is checked by default to enable and implement this alert immediately after saving. Unceck this box if you do not to immediately apply the rule to your environment.
10. Click **Save**.

---

32 Log Forwarding Overview

Topics:

- Security Events and Traffic Logs
- Discovery Logs
- Gateway Metrics Forwarding Profile
- Add an Event, Traffic Log Forwarding Profile, or Metrics Forward Profile to a Gateway
- Remove an Event, Traffic Log Forwarding Profile, or Metrics Forward Profile from a Gateway

Explains how logs are forwarded to various systems for analysis, monitoring, and storage.

Security Events and Traffic Logs

Learn about security events and traffic logs in Multicloud Defense and how they support investigation and monitoring. Find the available event categories, forwarding destinations, and profile requirements for sending log data to a SIEM.

Security Information Event Management (SIEM) systems are solutions that specialize in combining security information and security event information together into a single management platform. The security and event information will originate from third party security solutions that are configured to forward this information to the SIEM.

Multicloud Defense supports viewing security event information directly within the UI. These events are available under the **Investigate** > **Flow Analytics** section. The events are categorized and viewable as follows:

Category	Type	Description
Flow Logs	FLOW_LOG	Information related to the different stages of a traffic flow
Firewall Events	APPID	Traffic matched based on Application ID (OpenAppID)
	GEOIP	Traffic sourced from or destined to a Geo IP (MaxMind)
	L4_FW	Traffic matched based on layer4 information (Source/Dest IP/Port and Protocol)
	MALICIOUS_IP	Traffic sourced from or destined to a malicious IP (Talos)
	SNI	Traffic matched based on SNI information
Network Threats	AV	Traffic where a virus has been detected (ClamAV)
	DPI	Traffic where an IDS/IPS threat has been detected (Talos)
	DLP	Traffic where sensitive data is being exfiltration
Web Protection	WAF	Traffic where a web application threat has been detected (ModSecurity)
	L7DOS	Traffic that is contributing to a layer7 DOS attack
URL Filtering	URLFILTER	Traffic that matches a URL category or URL (Talos)
FQDN Filtering	FQDNFILTER	Traffic that matches a FQDN category or FQDN (Talos)
HTTPS Logs	HTTP_REQUEST	Information related to web-based traffic (HTTP)
	TLS_ERROR	Information related to TLS errors
	TLS_LOG	Information related to TLS behavior
Traffic Summary Logs	SESSION_SUMMARY	Summary information on each processed traffic session

> **Note**
>
> Flow Logs are deprecated in 2.10 and later gateway releases. The information contained within each flow Log is made available as part of the session information available in **Traffic Summary** > **Logs**.

Each of the event categories can be sent to a SIEM using a log forwarding profile. The SIEMs currently supported by Multicloud Defense are:

- AWS S3 Bucket
- Datadog
- GCP Logging
- Microsoft Sentinel
- Splunk
- Sumo Logic
- Syslog
- Webhooks

Use these steps to operate a log forwarding profile:

Create a Standalone Event or Traffic Log Profile

Create a standalone event or traffic log profile for a supported SIEM destination, and then associate the profile with the gateways whose logs you want to forward.

Procedure

1. Navigate to **Infrastructure** > **Profiles** > **Log Forwarding**.
2. Click **Create**.
3. Specify a Profile Name and Description.
4. Specify *Type* as Standalone.
5. Fill in the appropriate parameters (refer to the SIEM-specific documentation).
6. Click **Save**.
7. Add the desired Gateway Associations (refer to Add a Gateway Association).

Edit a Standalone Event or Traffic Log Profile

Edit a standalone event or traffic log profile in Multicloud Defense to update its forwarding destination settings for a supported SIEM integration.

Procedure

1. Navigate to **Infrastructure** > **Profiles** > **Log Forwarding**.
2. Check the box next to the Profile you want to *Edit*.
3. Click **Edit**.

4. Modify the parameters as desired (refer to the SIEM-specific documentation).
5. Click **Save**.

Create a Group Event or Traffic Log Profile

Create a group event or traffic log profile to manage multiple standalone forwarding profiles together, and then associate the group profile with the required gateways.

Procedure

1. Navigate to **Infrastructure** > **Profiles** > **Log Forwarding**.
2. Click **Create**.
3. Specify a Profile Name and Description.
4. Specify *Type* as Group.
5. Add as many rows as needed to accommodate for the number of standalone profiles you want to group.
6. Click **Save**.
7. Add the desired **gateway associations** (refer to Add a Gateway Association).

Edit a Group Event or Traffic Log Profile

Edit a group event or traffic log profile to add, change, or remove its associated standalone profiles and keep the forwarding group current.

Procedure

1. Navigate to **Infrastructure** > **Profiles** > **Log Forwarding**.
2. Check the checkbox next to the profile that you want to edit.
3. Click **Edit**.
4. You can modify, add, or remove standalone profiles.
5. Click **Save**.

View an Event or Traffic Log Forwarding Profile

View an event or traffic log forwarding profile in Multicloud Defense to review its current forwarding configuration and destination information.

Procedure

1. Navigate to **Infrastructure** > **Profiles** > **Log Forwarding**.
2. Select the Profile link you want to view the *Details*.
3. View the *Details* information.

Delete an Event or Traffic Log Profile

Delete an event or traffic log profile from Multicloud Defense after removing its gateway associations, so obsolete forwarding configurations no longer appear on the dashboard.

Use the following procedure to delete the profile from your dashboard:

Before you begin

You **must** remove the association between the event or profile and the gateway before you delete the profile from your dashboard. See Remove a Gateway Association for more information.

Procedure

1. Navigate to **Infrastructure** > **Profiles** > **Log Forwarding**.
2. Check the box next to the Profile you want to *Delete*.
3. Click **Delete**.
4. Confirm the *Delete* operation by clicking **Yes** or **No**.

Discovery Logs

Learn about discovery logs in Multicloud Defense and how they help you review cloud security activity. Find the supported log categories, forwarding destinations, and profile requirements for sending discovery data to a SIEM.

Discovery logs may be forwarded to Security Information Event Management (SIEM) systems to aggregate into a single management platform.

Multicloud Defense supports viewing security event information directly within the UI. These events are available under the **Investigate** > **Traffic** section. The events are categorized and viewable as follows:

Category	Type	Description
DNS Logs	DNS_LOG	Correlation of Threat Intelligence with DNS Log information gathered from cloud provider
VPC Logs	VPC_LOG	Correlation of Threat Intelligence with VPC/VNet Flow Log information gathered from cloud provider

Each of the categories can be sent to a SIEM using a Log Forwarding Profile and attaching the Profile to the onboarded Cloud Account. The Log Forwarding destinations currently supported by Multicloud Defense are:

- AWS S3 Bucket
- Datadog
- GCP Logging
- Microsoft Sentinel
- Splunk
- Sumo Logic

- Syslog

To forward Discovery Logs, use the steps below:

Create a Standalone Discovery Log Profile

Create a standalone discovery log profile for a supported SIEM destination, and then associate the profile with the cloud accounts whose discovery logs you want to forward.

Procedure

1. Navigate to **Infrastructure** > **Profiles** > **Log Forwarding**.
2. Click **Create**.
3. Specify a Profile Name and Description.
4. Specify **Type**as Standalone.
5. Fill in the appropriate parameters (refer to the SIEM-specific documentation).
6. Click **Save**.
7. Associate the Log Profile to the desired Cloud Accounts (refer to Add a Cloud Account Association).

Edit a Standalone Discovery Log Profile

Edit a standalone discovery log profile in Multicloud Defense to update its forwarding destination settings for a supported SIEM integration.

Procedure

1. Navigate to **Infrastructure** > **Profiles** > **Log Forwarding**.
2. Check the box next to the profile you want to *Edit*.
3. Click **Edit**.
4. Modify the parameters as desired (refer to the SIEM-specific documentation).
5. Click **Save**.

Create a Group Discovery Log Profile

Create a group discovery log profile to organize standalone forwarding profiles together, and then associate the group profile with the required resources.

Procedure

1. Navigate to **Infrastructure** > **Profiles** > **Log Forwarding**.
2. Click **Create**.
3. Specify a Profile Name and Description.
4. Specify *Type* as Group.
5. Add a row for to associate a Standalone Profile.
6. Click **Save**.
7. Add the desired Gateway Associations (refer to Add a Gateway Association).

Edit a Group Discovery Log Profile

Edit a group discovery log profile to add, change, or remove its associated standalone profiles and keep the forwarding group current.

Procedure

1. Navigate to **Infrastructure** > **Profiles** > **Log Forwarding**.
2. Check the box next to the Profile you want to *Edit*.
3. Click **Edit**.
4. Modify, Add or Remove Standalone Profiles.
5. Click **Save**.

View a Discovery Log Profile Details

View the details of a discovery log profile in Multicloud Defense to review its current configuration and forwarding destination information.

Procedure

1. Navigate to **Infrastructure** > **Profiles** > **Log Forwarding**.
2. Select the Profile link you want to view the *Details*.
3. View the *Details* information.

Add a Discovery Log Profile with a Cloud Account

Associate a discovery log profile with a cloud account in Multicloud Defense to forward the account's discovery logs to the selected destination.

Procedure

1. Navigate to **Systems and Accounts** > **Account** > **CSP Accounts**.
2. Check the checkbox next to the cloud account that you want to associate the *Profile* with.
3. Click **Actions** > **Update Log Profile**.
4. Select the **Log Profile** object for cloud logs forwarding profile.
5. Click **Save & Continue**.

Remove a Discovery Log Profile from a Cloud Account

Remove a discovery log profile from a cloud account when you no longer want Multicloud Defense to forward that account's discovery logs through the profile.

Procedure

1. Navigate to **Systems and Accounts** > **Account** > **CSP Accounts**.
2. Check the checkbox next the cloud account that you want to disassociate the *Profile* with.

3. Click **Actions** > **Update Log Profile**.
4. For the *Cloud Logs Forwarding Profile* parameter, click the 'X' next to the *Profile* to remove it.
5. Click **Save & Continue**.

Delete a Discovery Log Profile

Delete a discovery log profile from Multicloud Defense after removing its cloud account associations, so obsolete forwarding configurations no longer appear on the dashboard.

Use the following procedure to delete the profile from your dashboard:

Before you begin

You **must** remove the association between the profile and the gateway before you delete the profile from your dashboard. See Remove a Cloud Account Association for more information.

Procedure

1. Navigate to **Infrastructure** > **Profiles** > **Log Forwarding**.
2. Check the box next to the Profile you want to *Delete*.
3. Click **Delete**.
4. Confirm the *Delete* operation by clicking **Yes** or **No**.

Gateway Metrics Forwarding Profile

Explains the purpose and capabilities of the Gateway Metrics Forwarding Profile.

This profile is intended to forward gateway metrics generated by the Multicloud Defense Gateway for data monitoring and analysis.While the metrics are generated by the gateway, it is the Multicloud Defense Controller that forwards the metrics to the third party analysis application. This forwarding profile allows you to monitor, analyze, and organize your gateway metrics without logging into Multicloud Defense. Use this information to gauge the performance and behavior of your gateway environment; you can also utilize this information for environmental troubleshooting.

> **Note**
>
> As of Multicloud Defense Controller version 23.09, only Datadog is supported as a third-party analytics application.

For the majority of analytics applications available, for example, Datadog, you must already be an authorized user to access the tool's APIs and rendered data.

Create a Standalone Metrics Forwarding Profile

Configure and establish a profile for metrics forwarding to a third-party application.

Create a standalone profile and forward metrics to be processed by a third party.

Before you begin

You must have at least one third-party application to forward the metrics to, prior to creating this profile.

Procedure

1. Navigate to **Infrastructure** > **Profiles** > **Metrics Forwarding**.
2. Click **Create**.
3. Enter a unique profile **Name**.
4. (Optional) Enter a **Description**. This may help differentiate from other profiles with a similar name.
5. Expand the **Type** drop-down menu and select **Standalone**.
6. Expand the **Destination** drop-down menu and select the third-party application to process and analyze the metrics.
7. Enter the **Endpoint** to be used as the endpoint location for the metrics.
8. (Optional) Send a **Test** message to the destination. Enter the content of the test message in the text field when prompted and clikc **Validate**. If the destination does not receive the test message, confirm the destination type and configuration in step 7.
9. Click **Save**.

   If you select Datadog as your analytics application, the **Endpoint** is filled in by default with an HTTPS webhook. This entry, if defaulted, can be modified prior to saving the profile.

What to do next

Attach the profile to a policy rule set. See Rule Sets and Rule Set Groups on page 133 for more information.

Edit a Standalone Metrics Forwarding Profile

Edit a standalone profile that has already been created.

Use the following procedure to edit a standalone profile that has already been created.

Procedure

1. Navigate to **Infrastructure** > **Profiles** > **Metrics Forwarding** and select the appropriate profile **Type**.
2. Check the box next to the profile you want to edit.
3. Click **Edit**.
4. Modify the parameters as desired.
5. Click **Save**.

Create a Group Metrics Forwarding Profile

Configure a profile that aggregates multiple standalone metrics profiles to be assigned to a single gateway.

In this process, you create a profile and then assign it to a specific gateway. A group profile combines up to five standalone metrics forwarding profile that can then be assigned to a single gateway. Use the following procedure to create a grouped metrics forward profile:

Before you begin

- You must have at least one third party application to forward the metric to prior to creating this profile.

- You must have at least two standalone metrics forwarding profiles already created. See Create a Standalone Metrics Forwarding Profile on page 208 for more information.

Procedure

1. In the Multicloud Defense Controller interface navigate to **Infrastructure** > **Profiles** > **Metrics Forwarding**.
2. Click **Create**.
3. Enter a unique **Profile Name**
4. (Optional) Enter a **Description**. This may help differentiate between profiles with a similar name.
5. Expand the **Type** drop-down menu and select **Group**.
6. Under **Group Details**, click **Add** for every new row you need to add to the profile.
7. Expand the drop-down menus for each row to select a profile to add to the group. If you want to remove a profile at any point prior to saving, select the profile's checkbox so it is highlighted and select **Remove**.
8. Click **Save**.

What to do next

Attach the profile to a policy rule set. See Rule Sets and Rule Set Groups on page 133 for more information.

Edit a Group Profile

Change existing grouped profiles by modifying, adding, or removing profile settings.

Use the following procedure to edit a set of grouped profiles that has already been created:

Procedure

1. Navigate to **Infrastructure** > **Profiles** > **Metrics Forwarding** and select the appropriate profile **Type**.
2. Check the box next to the profile you want to edit.
3. Click **Edit**.
4. Modify, add or remove group profiles.
5. Click **Save**.

Delete a Profile

Remove a monitoring profile from your dashboard to clean up outdated or unnecessary entries.

Use the following procedure to delete the profile from your dashboard:

Before you begin

You **must** remove the association between the profile and the gateway before you delete the profile from your dashboard. See Remove a Gateway Association for more information.

Procedure

1. Navigate to **Infrastructure** > **Profiles** > **Metrics Forwarding** and select the appropriate profile **Type**.
2. Check the box next to the profile you want to delete.

3. Click **Delete**.
4. Click **Yes** or **No** to either confirm or cancel the delete action.

Add an Event, Traffic Log Forwarding Profile, or Metrics Forward Profile to a Gateway

Associate an event, traffic log forwarding, or metrics forward profile with a Multicloud Defense gateway so the gateway sends its data to the selected destination.

Procedure

1. Navigate to **Infrastructure** > **Gateways** > **Gateways**.
2. Check the box next the gateway you want to associate the *Profile*.
3. Click **Edit**.
4. For the *Log Profile* parameter, select the desired *Profile* from the menu.
5. Click **Save**.

Remove an Event, Traffic Log Forwarding Profile, or Metrics Forward Profile from a Gateway

Remove an event, traffic log forwarding, or metrics forward profile from a gateway when you no longer want the gateway to send data through that profile.

Procedure

1. Navigate to **Infrastructure** > **Gateways** > **Gateways**.
2. Check the box next the gateway you want to de-associate the *Profile*.
3. Click **Edit**.
4. For the *Log Profile* parameter, click the 'X' next to the *Profile* to remove it.
5. Click **Save**.

> **Note**
>
> A Log Forwarding Profile can also be associated with a gateway at time of gateway creation. The *Log Profile* parameter is available during the gateway creation process, where the desired *Profile* can be selected from the menu.

33 Log Forwarding Destinations / SIEMs

Topics:

- AWS S3 Bucket
- Datadog
- GCP Logging
- Microsoft Sentinel
- Splunk
- Sumo Logic
- Syslogs
- Webhook

Explains how log destinations interface with SIEMs to manage and analyze security events effectively.

AWS S3 Bucket

Learn about forwarding Multicloud Defense security events and traffic logs to an AWS S3 bucket for storage and analysis. Find the bucket access requirements and profile settings needed for this destination.

Multicloud Defense supports forwarding Security Events and Traffic Logs to an AWS S3 Bucket to send Security Events and Traffic Log information for processing, storage, access and correlation. The information sent is in a semi-structured JSON format where the attribute-value pairs can be accessed and processed.

Requirements

1. Create a new or use an existing AWS S3 Bucket.
2. Apply the following policy to the AWS S3 Bucket to permit the Multicloud Defense Controller to access and write to the bucket:

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Principal": {
        "AWS": "<controller-role-arn>"
      },
      "Action": "s3:*",
      "Resource": [
        "arn:aws:s3:::<s3bucketname>/*",
        "arn:aws:s3:::<s3bucketname>"
      ]
    }
  ]
}
```

Profile Parameters

Parameter	Requirement	Default	Description
Profile Name	Required		A unique name to reference the profile.
Description	Optional		A description for the Profile.
Destination	Required	AWS S3	AWS S3 Bucket.
CSP Account	Required		The CSP Account where the AWS S3 Bucket resides.
S3 Bucket	Required		The AWS S3 Bucket name where events/logs will be forwarded.

Datadog

Learn about forwarding Multicloud Defense security events and traffic logs to Datadog for monitoring and analysis. Find the account, endpoint, and API key information needed to configure this SIEM destination.

Datadog is a very common and powerful SIEM that is used by many companies. Multicloud Defense supports Log Forwarding to Datadog to send Security Events and Traffic Log information for processing, storage, access and correlation. The information sent is in a semi-structured JSON format where the attribute-value pairs can be accessed and processed.

Requirements

In order to forward logs to Datadog, the following information is required:

- Datadog account
- Endpoint URL
- API Key

> **Tip**
>
> - To sign up for a Datadog account, refer to **Datadog Account** (https://www.datadoghq.com/).
> - To create a Datadog API Key, refer to **Datadog API Key** (https://app.datadoghq.com/account/login?next=%2Faccount%2Fsettings#api).

Profile Parameters

Parameter	Requirement	Default	Description
Profile Name	Required		A unique name to use to reference the Profile.
Description	Optional		A description for the Profile.
Destination	Required	Datadog	The SIEM used for the Profile.
Skip Verify Certificate	Optional	Unchecked	Whether to skip verifying the authenticity of the certificate.
API Key	Required		The Datadog API Key to authenticate the communication.
Endpoint	Required	https://http-intake.logs.datadoghq.com/	The URL endpoint used to receive the forwarded Events/Logs.

GCP Logging

Learn about forwarding Multicloud Defense security events and traffic logs to GCP Logging for storage and analysis. Find the service account role, log name, and event field details needed to configure this destination.

GCP Stackdriver Logging is a service offering by Google Cloud Provider (GCP) for collecting and storing logs from applications and services. Multicloud Defense supports Log Forwarding to GCP Stackdriver Logging to send Security Events and Traffic Log information for processing, storage, access and correlation. The information sent is in a semi-structured JSON format where the attribute-value pairs can be accessed and processed.

Requirements

The GCP *Multicloud Defense-firewall* Service Account must be assigned **Logs Writer** role in order for the Gateway to write events to the GCP Stackdriver Log.

Profile Parameters

Parameter	Requirement	Default	Description
Profile Name	Required		A unique name to use to reference the profile.
Description	Optional		A description for the profile.
Destination	Required	GCP Logging (From Gateway)	The SIEM used for the profile.
Log Name	Required	ciscomcd -gateway-logs	The name of the Stackdriver Log used to store events.

Field Integer to String Mappings

When events are forwarded from the Controller, the Controller introduces mappings of event field values to friendly names. When events are forwarded directly from the Gateway (e.g., GCP Logging), the Controller is not involved and thus the event field values are not mapped to more relatable names. In order to interpret these fields, you must ensure that the field value is more relatable and understandable.

Field	Integer	String
action	0	DUMMY_ACTION
	1	ALLOW
	2	DENY
	3	DROP
	4	REDIRECT
	5	PROXY
	6	LOG
	7	OTHER
	8	DELAY
	9	DETECT_SIG

Field	Integer	String
gatewaySecurityType	1	INGRESS_FIREWALL
	2	EAST_WEST_AND_EGRESS_FIREWALL

Field	Integer	String
level	1	DEBUG
	2	INFO
	3	NOTICE
	4	WARNING
	5	ERROR
	6	CRITICAL
	7	ALERT
	8	EMERGENCY

Field	Integer	String
policyMatchInfo.serviceType	0	UNKNOWN
	1	PROXY
	2	FORWARDING
	3	REVERSE_PROXY
	4	FORWARD_PROXY

Field	Integer	String
protocol	0	DUMMY
sessionSummaryInfo.egressConnection.protocol	1	ICMP
sessionSummaryInfo.ingressConnect.protocol	6	TCP
	17	UDP
	252	HTTP

Field	Integer	String
rule.type	0	DUMMY_RULE_TYPE
	1	THIRD_PARTY
	2	USER_DEFINED

Field	Integer	String
statusText	0	CLOSED
ingressConnectionStates.state	1	SYN_SENT
	2	SYN_RECV
	3	ESTABLISHED
	4	FIN_WAIT
	5	CLOSE_WAIT
	6	LAST_ACK
	7	TIME_WAIT
	8	CLOSE

Field	Integer	String
type	1	WAF
	2	DPI
	3	HTTP_REQUEST
	4	L4_FW
	5	FLOW_LOG
	6	MALICIOUS_IP
	7	TLS_ERROR
	8	TLS_LOG
	9	L7DOS
	10	SNI
	11	APPID
	12	URLFILTER
	13	SESSION_SUMMARY
	14	DLP
	15	FQDNFILTER
	16	AV

Microsoft Sentinel

Learn about forwarding Multicloud Defense security events and traffic logs to Microsoft Sentinel for monitoring and analysis. Find the Azure Log Analytics workspace, shared key, and log table information needed to configure this SIEM destination.

Microsoft Sentinel is a powerful SIEM that is used by many companies. Multicloud Defense supports Log Forwarding to Microsoft Sentinel to send Security Events and Traffic Log information for processing, storage, access and correlation. The information sent is in a semi-structured JSON format where the attribute-value pairs can be accessed and processed.

Requirements

In order to forward logs to Microsoft Sentinel, the following information is required:

- Create an Azure Log Analytics Workspace.
- Define an Azure Log Table.

Profile Parameters

Parameter	Requirement	Default	Description
Profile Name	Required		A unique name to use to reference the Profile.
Description	Optional		A description for the Profile.
Destination	Required	Microsoft Sentinel	The SIEM used for the Profile.
Azure Log Analytics Workspace ID	Required		The ID of the Azure Log Analytics Workspace.
Shared Key	Required		The Shared Key used to authenticate the communication.
Azure Log Table Name	Required		Name of the Azure Log Table where the logs/events will be stored.

Splunk

Learn about forwarding Multicloud Defense security events and traffic logs to Splunk for monitoring and analysis. Find the HTTP Event Collector endpoint, token, and index information needed to configure this SIEM destination.

Splunk is a very common and powerful SIEM that is used by many companies. Multicloud Defense supports Log Forwarding to Splunk to send Security Events and Traffic Log information for processing, storage, access and correlation. The information sent is in a semi-structured JSON format where the attribute-value pairs can be accessed and processed.

Requirements

In order to forward logs to Splunk, the following information is required:

- Splunk account
- Splunk Collector URL
- Event Collector Key
- Index Name

> **Tip**
>
> For information on the Splunk Event Collector, refer to **Splunk HTTP Event Collector** (https://docs.splunk.com/Documentation/Splunk/8.1.0/Data/UsetheHTTPEventCollector).

Profile Parameters

Parameter	Requirement	Default	Description
Profile Name	Required		A unique name to use to reference the Profile.
Description	Optional		A description for the Profile.
Destination	Required	Datadog	The SIEM used for the Profile.
Skip Verify Certificate	Optional	Unchecked	Whether to skip verifying the authenticity of the certificate.
Endpoint	Required		The URL used to access the HTTP Event Collector.
Token	Required		The Splunk Token to allow Multicloud Defense to communicate with Splunk.
Index	Required	main	The name of the Splunk index used to store events.

Sumo Logic

Learn about forwarding Multicloud Defense security events and traffic logs to Sumo Logic for monitoring and analysis. Find the account and collector endpoint information needed to configure this SIEM destination.

Sumo Logic is a very common and powerful SIEM that is used by many companies. Multicloud Defense supports Log Forwarding to Sumo Logic to send Security Events and Traffic Log information for processing, storage, access and correlation. The information sent is in a semi-structured JSON format where the attribute-value pairs can be accessed and processed.

Requirements

In order to forward logs to Sumo Logic, the following information is required:

- Sumo Logic account
- Sumo Logic collector endpoint

> **Tip**
>
> For information on how to setup Sumo Logic Collector, refer to **Sumo Logic Setup Guide** (https://help.sumologic.com/docs/send-data/setup-wizard/).

Profile Parameters

Parameter	Requirement	Default	Description
Profile Name	Required		A unique name to use to reference the Profile
Description	Optional		A description for the Profile
Destination	Required	Sumo Logic	The SIEM used for the Profile
Endpoint	Required		The URL endpoint used to receive the forwarded Events/Logs

Syslogs

Learn about forwarding Multicloud Defense security events and traffic logs to a syslog server. Find the supported log categories and profile settings needed to send data to this destination.

A syslog server is a common log collector that accepts a standard formatted syslog message. Each syslog message contains fields for facility, severity and message. Almost any SIEM can accept syslog formatted messages, although most SIEMs support other message formats. Multicloud Defense supports sending security events and traffic logs to a syslog server. The following are a list of events and logs that are forwarded:

- Flow Logs (Traffic Summary)
- Firewall Events (AppID, L4FW, GeoIP, MaliciousIP, SNI)
- HTTPS Logs (HTTP, TLS)
- Network Threats (AV, DLP, IDS/IPS)
- Web Protection (WAF, L7 DoS)

At this time this list of included events and logs is mandatory and cannot be altered. If you configure syslogs to be forwarded then **all** of these logs are included in the report.

Note

Flow logs are deprecated in gateway version 2.10 and later releases. The information contained within each flow log is made available as part of the session information available in **Investigate** > **Flow Analytics** > **Traffic Summary**.

Events can be forwarded to a syslog server using a log forwarding profile. Once created, the profile needs to be associated with a new or existing gateway in order for the events to be sent to the syslog Server. To create, modify or change the gateway association of a log forwarding profile, refer to Log Forwarding - Security Events and Traffic Logs.

Profile Parameters

Parameter	Requirement	Default	Description
Profile Name	Required		A unique name to use to reference the Profile.
Description	Optional		A description for the Profile.
SIEM Vendor	Required	Syslog	The SIEM used for the profile.
Server IP	Required		The IP address of the syslog server.
Protocol	Required	UDP	The protocol to use when sending messages (TCP / UDP).
Port	Required		The port to use when sending messages.
Format	Required	IETF	The format of the messages (only IETF is supported).
Flow Logs	Required	No	Whether to send flow logs (Yes / No).
Firewall Events	Required	No	Whether to send firewall events (Yes / No).
HTTPS Logs	Required	No	Whether to send HTTPS logs (Yes / No).
Network Threats	Required	Emergency	The lowest severity level to send network threats.
Web Attacks	Required	Emergency	The lowest severity level to send web attacks.

> **Note**
>
> The following levels of severity (highest to lowest) are available:

- Emergency
- Alert
- Criticial
- Error
- Warning
- Notice
- Info
- Debug

All events for the category that contain the severity level specified or higher will be sent to the syslog server.

Webhook

Learn about forwarding Multicloud Defense logs to a webhook for integration with external services. Find the endpoint, message, batch size, and authentication settings needed to configure the forwarding profile.

Logging forward with a webhook can be a useful practice in various scenarios, particularly when you need to integrate different systems or services in real-time. Webhooks can enable real-time data transfer, are ideal for an event-driven architecture or centralized logging, as well as support automation and integration with other third-party services. Customize the log forwarding profile to support your specific environment.

Requirements

If you are forwarding logs to a service that requires an authetnication token or password, have that configured and readily available when creating this profile.

Profile Parameters

Parameter	Requirement	Default	Description
Profile Name	Required		A unique name to use to reference the Profile.
Description	Optional		A description for the Profile.
Type	Required		Select **Standalone**.
Destination	Required	Webhook	
Endpoint	Required		Enter the URL for the webhook.
Message Field Name	Required		This value is defines the structure and meaning of the data being transmitted.

Parameter	Requirement	Default	Description
Batch size	Required	100	This value determines the number of log entries that are grouped together and sent in a single transmission.
Authentication Type	Required		Select **Basic** for username and password or **Bearer token**.

34 Management

Topics:

- Management
- Alert Profiles

Learn about using Multicloud Defense administration settings to review and manage the cloud service providers and resources associated with your account.

View and manage the overall status of the cloud service providers affiliated with your account.

Management

Learn about using Multicloud Defense administration settings to review and manage the cloud service providers and resources associated with your account.

View and manage the overall status of the cloud service providers affiliated with your account.

API Keys

Learn about viewing and managing Multicloud Defense API keys created for cloud service provider components. Search key records, review role and expiration details, or access actions to create and delete keys.

Navigate to **Systems and Accounts** > **Account** > **API Keys** to view this page.

Search

Use the search bar to seek or filter the list of API keys with key words. You must use at least three characters for the search to qualify.

API Key Table and Actions

This table lists all the API keys that are created by Multicloud Defense components for your cloud service providers. View the role, key ID, the date the key was added to Multicloud Defense, and the date the key expires.

From here you can create or delete API keys. Note that these keys are generated by Multicloud Defense and not related to the keys your cloud service provider might create to maintain communication. Continue reading for more information.

Create an API Key in Multicloud Defense

Create a new API key for authorization and management tasks within Multicloud.

Use the following procedure to create an API Key:

Procedure

1. Navigate to **Systems and Accounts** > **Account** > **API Keys**.
2. Click **Create API Key**.
3. Enter a unique **Name**.
4. Confirm the **Email Address** that Multicloud Defense automatically generates. You cannot change this option.
5. Use the drop-down menu to select one of the key roles:
   - **admin_read_only** - This role restricts interactions so you cannot modify or action anything, and can only “view” the available data.
   - **admin_read_rw** - This role allows you to read and modify available data.
6. Enter an appropriate value for **API Key Lifetime (days)**. The deafult value is 365 days.
7. Click **Save**.

Delete an API Key from Multicloud Defense

Delete an API key from Multicloud Defense when the key is no longer needed for a cloud service provider component.

Use the following procedure to delete a API Key:

Procedure

1. Navigate to **Systems and Accounts** > **Account** > **API Keys**.
2. Select the API Key from the table and check the box so it is highlighted.
3. Click **Delete**.
4. Confirm you want to delete the key and click **Yes**. The key is immediately removed from Multicloud Defense.

Account Level Settings

Learn about account-level settings for managing application tags and custom tags in Multicloud Defense. Use these settings to organize and identify applications and other resources.

This page displays some of the tags used in Multicloud Defense, including application tags and custom tags. Continue reading for more information.

Application Tags

Learn about using application tags in Multicloud Defense to classify and group applications for easier searching and vulnerability review. Understand the availability and usage considerations for these tags.

The application tag is a string of characters and is used as one of the classification criteria for the automatic classification of processes or threads. Tagging allows you to group apps based on your unique requirements so that you can search for apps and find vulnerabilities. Note that not all cloud service providers support the use of application tags.

> **Note**
>
> You can only create one appication tag at a time. If you need to create a new tag, you **must** delete the existing tag and then create a new application tag.

Create an Application Tag

Create an application tag in Multicloud Defense to classify and group applications for easier searching and vulnerability review across your environment.

Use the following procedure to create an application tag. Note that these tags are for internal use only and may not be recognized or available from your cloud ervice provider's interface.

Procedure

1. Navigate to **Systems and Accounts** > **Account** > **Account**.
2. In the **Application Tag** table, click **Create**.
3. The type of application tag is `APPLICATION_TAG_KEYS` by default.
4. Enter a brief **Description** of the tag. This can help identify or differentiate between other tags that might have a similar name or concept.
5. Enter at least one **Value**. Hit `Enter` after each value to create more than one. Note that these values are case sensitive.
6. Click **Save**. The tag is created and available in the table.

Edit an Application Tag

Edit an application tag created in Multicloud Defense to update its description or add and remove tag values for classified applications.

Use the following procedure to edit an existing application tag that has been created in Multicloud Defense. You cannot use this procedure to modify tags that were created in your cloud service provider's interface.

Procedure

1. Navigate to **Systems and Accounts** > **Account** > **Account**.
2. In the **Application Tag** table, locate the application tag you want to edit and check its box on the left so it is highlighted.
3. Click **Edit**.
4. Modify the following parameters:
   - **Description** - You can edit or delete the description.
   - **Tag Values** - You can add or remove tags here.
5. Click **Save**. Alternatively, you can cancel at any time without saving changes.

Delete an Application Tag

Delete an existing application tag from Multicloud Defense when it is no longer needed to classify or organize applications in your environment.

Use the following procedure to delete an existing application tag:

Procedure

1. Navigate to **Systems and Accounts** > **Account** > **Account**.
2. In the **Application Tag** table, locate the application tag you want to edit and check its box on the left so it is highlighted.
3. Click **Delete**.
4. Confirm you want to delete the application tag and click **Yes**.

Custom Tags

Learn about using custom tags in Multicloud Defense to identify, differentiate, and locate related objects, policies, rules, and other managed items.

Custom tags are simple pieces of data that provide details about an item and make it easy to locate related items that have the same tag. You can use a tag to easily identify or differentiate an object, policy, rule, and more.

Create a Custom Tag

Create a custom tag in Multicloud Defense to identify and distinguish related objects, policies, rules, or other managed items within your environment.

Use the following procedure to create a custom tag in Multicloud Defense. Note that these tags are for internal use only and may not be recognized or available from your cloud ervice provider's interface.

Procedure

1. Navigate to **Systems and Accounts** > **Account** > **Account**.
2. In the **Custom Tag** table, click **Create**.
3. Enter the **Value** of the tag. This can help identify or differentiate between other tags that might have a similar name or concept
4. Enter at least one **Value**.
5. Click **Save**. The tag is created and available in the table.

Edit a Custom Tag

Edit an existing custom tag in Multicloud Defense to update its key and values for managed items without replacing the tag.

Use the following procedure to modify an existing custom tag:

Procedure

1. Navigate to **Systems and Accounts** > **Account** > **Account**.
2. In the **Custom Tag** table, locate the application tag you want to edit and check its box on the left so it is highlighted.
3. Click **Edit**.
4. Modify the following parameters:
   - Key.
   - Values.
5. Click **Save**. Alternatively, you can cancel at any time without saving changes.

Delete a Custom Tag

Delete an existing custom tag from Multicloud Defense when it is no longer needed to identify or organize managed items.

Use the following procedure to delete an existing custom tag:

Procedure

1. Navigate to **Systems and Accounts** > **Account** > **Account**.
2. In the **Custom Tag** table, locate the application tag you want to edit and check its box on the left so it is highlighted.
3. Click **Delete**.
4. Confirm you want to delete the application tag and click **Yes**.

System

Learn about reviewing Multicloud Defense component versions and update history on the System page. Find gateway image, network intrusion, and web protection information for release research or Cisco Support.

The **System** page is a historical document that catalogues at least a year's worth of updates. You can use these details for general knowledge, locating the correct Release Notes version, and when you contact Cisco Support for product help. The following information collections are displayed here:

Component

This section displays the current versions for both the Multicloud Defense Controller and the user interface. Note that you cannot force an update or rollback to a previous version from this page.

Gateway Images

The gateway images table denotes when your Multicloud Defense Gateway was upgraded, which version of the gateway was in place and for how long, and what time zone the gateway is established in.

Talos/Network Intrusion

This table displays all the updates from Cisco's Talos Intelligence Group. These updates are pushed to Cisco products separate from a normal product software release.

Web Protection

This table displays all the Web Application Firewall (WAF) core and trustwave rule updates against the latest Web application vulnerabilities and threats.

Metering

Learn about reviewing Multicloud Defense usage and gateway instance data on the Metering page. Filter records by time period and download usage or instance reports for further analysis.

The **Metering** page displays graphs of usage, both for the overall useage of Multicloud Defense and the gateway instances created for your cloud service providers.

Filters

Use the filters located at the top of the page to determine the data displayed in the page. You can change this view by selecting the month and year. You can use these filter settings to generate a usage report.

Generate a Usage Report

You can generate a usage report for either of the two options from this page. Navigate to **System and Accounts** > **System Settings** > **Metering** and expand the **Download** drop-down option in the **Filter** section of the page to select either usage or instances. The file is downloaded locally as an .csv file. Use the filtering options to determine the timespan the report should generate from.

Usage Records

The **Usage Records** table details the number of accounts associated with your tenant, how many hours the accounts were interacted with, and on what days of the month selected in the Filter section. You can determine from the usage/month ratio what days were the most active.

Instance Records

The **instance Records** table displays the following instance statistics:

- Account Name.
- Account type by cloud service provider.
- Instance ID.
- Instance Type.
- Availablilty zone.
- Gateway.

- Started - When the gateway instance was created.
- Ended - When the gateway instance expired or was terminated.

Alert Profiles

Learn about using Alert Profiles in Multicloud Defense to define which events create alerts and where notifications are sent. Configure both alerts and services to support complete notification coverage.

Access the following Management views by navigating to **System and Accounts** > **Service Alerts** > **Alert Rules**.

Both the **Services** and **Alerts** page focus on alerts from Multicloud Defense. The **Alerts** page focuses on *where* alerts are sent to and the **Alerts** page details *what* alerts are sent to the endpoints configured. For ideal configuration, spend time setting up entries in both pages to successfully and wholly optimze the alert opportunity within the dashboard.

Services

Learn about configuring services that receive Multicloud Defense alerts at third-party endpoints. Search existing services, review their types and update dates, or access actions to create and delete destinations.

Navigate to **System and Accounts** > **Service Alerts** > **Services** to view this page.

Services focuses on **where** you want to send alerts to. Note that you must provide criteria from the third-party application in order to successfully configure any options on this page.

Search

Use the search bar to seek or filter the list of services with key words. You must use at least three characters for the search to qualify.

Services Table and Actions

This table lists all the services that are created by Multicloud Defense components for your cloud service providers. View the name, type of service, the date the service was updated.

From here you can create or delete services. Note that these services are generated by Multicloud Defense and not related to the services your cloud service provider might provide.

Create a Service

Create an alert service to send Multicloud Defense notifications to a supported third-party messaging, monitoring, or incident management destination for your organization.

Use the following procedure to create a service:

Before you begin

You must have service notifications or integrations enabled or allowed on your third party messaging application.

Procedure

1. Navigate to **System and Accounts** > **Service Alerts** > **Services**.
2. Click **Create**.
3. Enter a unique **Name**.
4. (Optional) Enter a **Description**. This may help differentiate between other services that may have a similar name.
5. Use the drop-down menu to select the service **Type**:

- Pager Duty.
- ServiceNow.
- Slack.
- Datadog.
- Microsoft Sentinel.
- Microsoft Teams.
- Webex.
- Splunk.

6. Depending on the service type, complete the following entries when prompted:
   - API Key.
   - API URL.
   - Azure Log Table Name.
   - Azure Log Analytics Workspace ID
   - (Optional for Splunk) Index.
7. Click **Save**.

Edit a Service

Edit an existing alert service to update its name, description, destination type, or connection settings for third-party notifications as needed.

Use the following procedure to edit an existing service:

Procedure

---

1. Navigate to Navigate to **System and Accounts** > **Service Alerts** > **Services**.
2. Locate and select the service within the table so it is highlighted.
3. Expand the Actions drop-down menu and click **Edit**.
4. Modify the following aspects of the service:
   - Name.
   - Description.
   - Type.
   - Type-specific configuration criteria.
5. Click **Save** to confirm the changes. At any point, click **Cancel** to close the window and cancel the changes.

What to do next

You may have to **Refresh** the page to see any changes.

Clone a Service

Clone an existing alert service to reuse its destination type and configuration while creating a separately named service for another notification need.

Use the following procedure to clone an existing service:

Procedure

1. Navigate to **System and Accounts** > **Service Alerts** > **Services**.
2. Locate and select the service within the table so it is highlighted.
3. Expand the Actions drop-down menu and click **Clone**.
4. A clone of the service is generated. By default, only the service **Type** and any service-specific configuration criteria is retained.
5. Enter a unique **Name**.
6. (Optional) Enter a description.
7. Click **Save** to confirm the changes. At any point, click **Cancel** to close the window and cancel the changes.

What to do next

You may have to **Refresh** the page to see changes or additions to the table.

Export a Service

Export an existing Multicloud Defense alert service for Terraform so you can download or copy its resource definition and import the service into your Terraform state.

Use the following procedure to export an existing service:

Procedure

1. Navigate to **System and Accounts** > **Service Alerts** > **Services**.
2. Locate and select the service within the table so it is highlighted.
3. Expand the Actions drop-down menu and click **Export**.
4. Multicloud Defense generates an export wizard.
5. Either click **Download** to download the terrform locally or click **Copy Code** to copy the JSON resource to manually paste into the terroform script.
6. Within the terrform prompt, execute the command provided in the lower half of the window: `terraform import "ciscomcd_alert_profile". "servicename" <number in table>`
7. Follow the prompts within terraform to complete the task. There are no more steps in the dashboard.

Delete a Service

Delete an existing alert service from Multicloud Defense when its third-party notification destination is no longer needed by your organization.

Use the following procedure to delete an existing service:

Procedure

1. Navigate to **System and Accounts** > **Service Alerts** > **Services**.
2. Locate and select the service within the table so it is highlighted.
3. Expand the Actions drop-down menu and click **Delete**.
4. Confirm you want to delete the service and click **Yes**.
5. The service is removed from Multicloud Defense.

Alerts

Learn about managing which Multicloud Defense events generate notifications for configured third-party endpoints. Pair alert definitions with services to deliver the notifications your organization needs.

The Alerts page focuses on **what** alerts are sent to the third-party endpoints.We strongly recommend configuring both alerts and services to take advantage of the alerts opportunity.

Create an Alert

Create an alert in Multicloud Defense to define the event type, severity, status, and alert profile used for notifications to configured services.

Use the following procedure to create an alert:

Procedure

1. Navigate to **System and Accounts** > **Service Alerts** > **Alert Rules**.
2. Click **Create**.
3. Enter a unique **Name**.
4. (Optional) Enter a **Description**. This may help differentiate between other services that may have a similar name.
5. Select the **Alert Profile**. At this time, `Pagerduty` is the only option available.
6. Use the drop-down menu to select the alert **Type**.
   - System Logs.
   - Audit Logs.
   - Discovery.
7. (Optional) Use the drop-down menu to select the **Sub Type**. Note that these options may change or may not be available depending on the Type you selected in step 6:
   - Gateway.
   - Account.
   - Controller.
   - Insights Rule.
8. Use the drop-down menu and select the level of **Severity**:
   - Info.
   - Warning.
   - Medium.

- High.
- Critical.

9. The **Enabled** checkbox is checked by default. This option designates whether the alert profile is active and usable or not. If it is disabled, Multicloud Defense does not include it when issuing alerts.

What to do next

Create a Service profile to designate where these alerts are sent to.

Edit an Alert

Edit an existing Multicloud Defense alert to update its profile, event criteria, severity, enabled status, or other notification settings as needed.

Use the following procedure to edit an existing alert:

Procedure

---

1. Navigate to **System and Accounts** > **Service Alerts** > **Alert Rules**.
2. Locate and select the alert within the table so it is highlighted.
3. Expand the Actions drop-down menu and click **Edit**.
4. Edit any of the fields and selections of the alert proile. Note that some of the available fields may change depending on the selections you make.
5. Click **Save** to confirm the changes. At any time, click **Cancel** to cancel the changes and close out the edit window.

Clone an Alert

Clone an existing alert to reuse its profile and event type while creating a separately configured alert for another notification need.

Use the following procedure to clone an existing alert:

Procedure

---

1. Navigate to **System and Accounts** > **Service Alerts** > **Services**.
2. Locate and select the alert within the table so it is highlighted.
3. Expand the Actions drop-down menu and click **Edit**.
4. A clone of the alert is generated. By default, only the **Alert Profile** and **Type** is retained.
5. Edit any of the remaining fields and selections of the alert. Note that some of the available fields may change depending on the selections you make.
6. Click **Save** to confirm the changes. At any time, click **Cancel** to cancel the changes and close out the edit window.

Export an Alert

Export an existing Multicloud Defense alert for Terraform so you can download or copy its resource definition and import the alert into your Terraform state.

Use the following procedure to export an existing alert:

Procedure

1. Navigate to **System and Accounts** > **Service Alerts** > **Alert Rules**.
2. Locate and select the alert within the table so it is highlighted.
3. Expand the Actions drop-down menu and click **Export**.
4. Multicloud Defense generates an export wizard.
5. Either click **Download** to download the terraform locally or click **Copy Code** to copy the JSON resource.
6. Manually paste into the terraform script.
7. Within the terraform prompt, execute the command provided in the lower half of the window: `terraform import "ciscomcd_alert_rule"."alertname" <number in table>`
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 an Alert

Delete an existing alert from Multicloud Defense when you no longer want its matching events to generate notifications for configured services.

Use the following procedure to delete an existing alert:

Procedure

1. Navigate to **System and Accounts** > **Service Alerts** > **Alert Rules**.
2. Locate and select the alert within the table so it is highlighted.
3. Expand the Actions drop-down menu and click **Delete**.
4. Confirm you want to delete the service and click **Yes**.
5. The alert is removed from Multicloud Defense.

35 Manage Your Multicloud Defense Account

Topics:

- Account (Multicloud Defense Tenant)
- Understand user roles in Security Cloud Control Firewall Management

Manage your Cisco account settings to keep them up-to-date and secure.

Account (Multicloud Defense Tenant)

Learn about the Multicloud Defense tenant Account page, where administrators access settings for managing account-level configuration information and related functions.

The Account information is used by the Administrator to create and edit the following functions.

Navigate to **Systems and Accounts** > **Account** > **Account**.

Understand user roles in Security Cloud Control Firewall Management

Learn how user roles define permissions for each Security Cloud Control Firewall Management organization.

Security Cloud Control Firewall Management assigns roles per user and per organization. If a user has access to more than one tenant, that user can have the same user ID but different roles on different organizations. For example, a user can have the `Read-Only` role on one organization and the `Super Admin` role on another.

When the interface or documentation refers to a `Read-Only` user, `Admin` user, or `Super Admin` user, it describes that user's permission level on a specific organization.

Roles in Multicloud Defense

Describes different roles and their permissions within a Multi-Cloud environment.

Roles play an important part of what a user is allowed to do when accessing the Multicloud Defense tenant through the Multicloud Defense portal. A role is a privilege that grants the user a set of permissions.

There are three available roles:

- Super Admin (admin_super) .
- Edit-only Admin (admin_rw).
- Read-only Admin (admin_read-only) .

There are two permission definitions:

- Modify - Read, write, edit, and delete.
- Read - Read-only.

The permissions for each setting associated with each role are outlined in the following table:

Setting	Super Admin(admin_super)	Edit-Only (admin_rw)	Read-Only (admin_read-only)
Management			
Users	Modify	Modify (except Super Admin)	Read
MFA Enable / Disable	Modify	Modify (except Super Admin)	Read
Reset MFA	Modify	Modify (except Super Admin)	Read
API Keys	Modify	Modify	Read
Roles	Read	Read	Read
Account > Application Tags	Modify	Modify	Read

Setting	Super Admin(admin_super)	Edit-Only (admin_rw)	Read-Only (admin_read-only)
Account > Email Domains	Modify	Read	Read
System	Read	Read	Read
Metering	Read	Read	Read
Alert Profiles			
Services	Modify	Modify	Read
Alert	Modify	Modify	Read

Only one (1) user within a Multicloud Defense tenant can be assigned the super admin role. This user is seen as the **owner** of the account and is synonymous with the owner of an AWS account or a linux root account. All other users should be assigned a read/write admin or read-only admin role.

The super admin role is assigned by Multicloud Defense and is granted to the first user created when the Multicloud Defense tenant is created. If any changes are required to a super admin user, please contact Multicloud Defense Support.

36 Cloud Accounts

Topics:

- Cloud Accounts
- Inventory

Explains the role, attributes, and importance of cloud accounts within a cloud services infrastructure.

Cloud Accounts

Learn about the Cloud Accounts page in Multicloud Defense and how it helps you review connected cloud service provider accounts before managing their configurations.

This is an overview of the cloud service providers currently connected to Multicloud Defense. You can do the following:

Add Account

Add a cloud service provider account to Multicloud Defense by selecting the account type and completing the connection wizard for that provider.

Procedure

1. Log in to the Multicloud Defense Controller and navigate to **Systems and Accounts** > **Account** > **CSP Accounts**.
2. Click **Add Account**.
3. For Account Type, use the drop-down menu to select the cloud service provider you want to connect.
4. Continue through the connection wizard to connect your cloud service provider. See the respective cloud service provider sections for more information about prerequisites and itemized values that might be specific to the type of cloud service provider.

Manage Inventory

Manage the inventory for a connected cloud account by adding or removing monitored regions, changing the refresh interval, or refreshing the account data.

Procedure

1. Log into the Multicloud Defense Controller and navigate to **Systems and Accounts** > **Account** > **CSP Accounts**.
2. In the Account table select **one** cloud service provider account.
3. From the options listed above the table, click **Manage Inventory**.
4. The generated window displays general information about the account, cloud service provider, and any currently monitored regions.
5. Modify the selection of regions with the following actions:
   - Add an individual region to an existing row.
   - Delete individual regions from an existing row.
   - Add a new row of monitored regions. Click the add button.
   - Delete an entire row of monitored regions. Click the minus button.
   - Change the **Refresh Interval** value. Default value is to refresh every 60 minutes.
   - Manually refresh the window with the **Refresh** icon.
6. Once all your changes are made, click **Save**. Alternatively, if you do not want to save the changes made, click **Cancel** to exit the window.

Edit a Cloud Account

Edit a connected cloud service provider account to update its configurable information and review the regions monitored in Multicloud Defense.

Use the following proceudre to edit the basic cloud service provider account information:

Procedure

1. Log into the Multicloud Defense Controller and navigate to **Systems and Accounts** > **Account** > **Account**.
2. In the Account table select **one** cloud service provider account.
3. From the options listed above the table, expand the **Actions** drop-down menu and click **Edit**. A window generates basic information about the account.
4. Edit any of the populated fields that are configurable. Note that not all cloud service providers have the same fields.
5. Click **Save & Continue**. Alternatively, if you do not want to save the changes, click **Cancel** to close the window.
6. The changes are successfully saved.
7. The window automatically switches to the **Manage Inventory** window for you to review and modify any currently monitored regions. Click **Save & Continue**. Alternatively, if you do not want to save the changes, click **Cancel** to close the window. This action closes the **Manage Inventory** window only. It does not revert changes made to the account from the prior window.

Update Log Profile for a Cloud Account

Update the log profile for a cloud account to select the service that receives forwarded logs from the connected cloud service provider.

Use the following procedure to modify the log fowarding service your cloud service provider is currently configured to send logs to:

Procedure

1. Log into the Multicloud Defense Controller and navigate to **Systems and Accounts** > **Account** > **Account**.
2. In the Account table select **one** cloud service provider account.
3. From the options listed above the table, expand the **Actions** drop-down menu and click **Update Log Profile**. A window generates basic information about the account.
4. Use the drop-down menu and select a service to forward logs to. To see which logs will be sent, hover your cursor over the information tag to the left of the drop-down.
5. Click **Save & Continue**. Alternatively, if you do not want to save the changes, click **Cancel** to close the window.
6. The changes are saved successfully and you are returned to the **Accounts** page.

Export a Cloud Account

Export a Multicloud Defense cloud account configuration for Terraform so you can download or copy the resource definition and import the account into your Terraform state.

Procedure

1. Log into the Multicloud Defense Controller and navigate to **Systems and Accounts** > **Account** > **Account**.
2. In the Account table select **one** cloud service provider account.

3. From the options listed above the table, expand the **Actions** drop-down menu and click **Export**. Multicloud Defense Controller generates an export wizard.
4. Either click Downloadto download the terraform locally orscroll down and click Copy Code to copy the JSON resource.
5. Manually paste into the terraform script.
6. Within the terraform prompt, execute the command provided in the lower half of the window:

```
terraform import "ciscomcd_cloud_account"."</cloud service provider>" </cloud
 service provider>
```

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

Delete a Cloud Account

Delete a cloud service provider account from Multicloud Defense when you no longer want the account available for cloud management and monitoring.

Procedure

1. Log into the Multicloud Defense Controller and navigate to **Systems and Accounts** > **Account** > **Account**.
2. In the Account table select **one** cloud service provider account.
3. From the options listed above the table, expand the **Actions** drop-down menu and click **Delete**. A window generates basic information about the account.
4. Confirm you want to delete the account and click **Yes**. If you do not want to delete the account, click **No** to exist the confirmation window.

Inventory

Learn about reviewing and managing cloud assets on the Multicloud Defense Inventory page. Filter or search discovered assets, view their history and configuration, and copy identifiers for use in related tasks.

The **Inventory** page allows you to review and manage the discovered assets affiliated with your cloud service provider account.

Use the filter option at the top of the page to organize the account page by cloud service provder type; as an alternative, click **Switch to Advanced Search** to create a search with multiple fields and values.

By clicking on any of the assets displayed in the page, you can view the history of the asset and its basic configuration. Note that you cannot edit the asset from this page. If you need to, you can copy both the VPC/VNet value and the asset ID from the table.

If you opt to **Manage Inventory on page 326**, you are redirrected to the appropriate page within Multicloud Defense.

37 Compliance Certificates and Awards

Topics:

- Compliance Certificates

Describes the significance of compliance certificates and associated awards.

Compliance Certificates

Learn about compliance certificates available for Multicloud Defense and why they matter for security, quality, and regulatory assurance. Find the current certification and access supporting documents in the Cisco Trust Portal.

Compliance certificates are official documents issued by a regulatory body, organization, or certified professional that confirms a product, service, or process meets specific regulatory standards, requirements, or guidelines. These certificates are crucial to ensure safety, quality, and adherence to laws and regulations.

See the Multicloud Defense Trust Portal to view the certificates.

Multicloud Defense has the following certificate:

MSECB ISO/IEC 27001:2022

The MSECB (Management Systems Evaluation Certification Board) awarded the **Information Security Management Systems (ISMS)** certificate. Certification according to ISO/IEC 27001 ensures that an organization has established, implemented, and maintains an effective information security management system.

38 Troubleshoot Connecting Your Account

Topics:

- Manually Onboard an Account
- Graceful Termination of Connections
- Terraform Onboarding Scripts for Cloud Accounts

Troubleshoot common issues when connecting a cloud account to Multicloud Defense.

Manually Onboard an Account

Learn about manual onboarding options for connecting a cloud service provider account to Multicloud Defense when the standard onboarding methods do not meet your needs.

In cases where onboarding a cloud service provider account to Multicloud Defense with the methods provided in onboarding section of cloud service providers, you may need to onboard your account manually. Use the following options as an alternative.

Manually Onboard a GCP Project

Learn about manually onboarding a GCP project to Multicloud Defense when standard onboarding methods are not suitable for specific environment requirements.

GCP Overview

An overview of onboarding Google Cloud projects and folders to Cisco Multicloud Defense, including supported use cases and configuration options.

GCP Project and GCP Folders

Multicloud Defense currently supports both GCP projects and GCP folders although these components are supported separately. Note the following limitations and exceptions for both of these options.

A GCP project has the potential to contain GCP resources like virtual machines, storage buckets, databases, and more. It can be used to create, enable, and use all Google Cloud services.

- Projects can be onboarded with Terraform, manual onboarding, and scripted onboarding.
- Projects are ideal for environments requiring orchestration, including discovery and investigation.
- You can interact with each project individually through the Multicloud Defense dashboard.

As of Version 23.10, you can connect a GCP folder with Terraform. A GCP folder contains projects, other folders, or a combination of both. Organization resources can use folders to group projects under the organization resource node in a hierarchy.

- Consider folders without the `roles/compute.admin` permission enabled as empty, and do not use them.
- Projects associated with onboarded folders are used for asset and traffic discovery only.
- Projects associated with onboarded folders do not accommodate orchestrating service VPC or gateway creation.
- Permissions made to folders from the GCP console must be made at the folder level. Therefore, perform, Multicloud Defense actions at the folder level.
- When creating a new gateway, do not use 172.16.0.0/16 subnet.

If you want to onboard a GCP folder, see refer to the Terraform repository.

Overview Procedure

The following is an overview of how to connect your GCP project. A shell **script** is provided by Multicloud Defense and facilitates an easy connective process as part of a wizard. The script automates the following steps so you don't have to:

1. Create two service accounts.
2. Enable the following APIs (Compute Engine, Secret Manager).
3. Create the two following VPCs (management, datapath).
4. Create firewall rules to allow traffic to the Multicloud Defense Gateway (app traffic) in the datapath VPC.

5. Create firewall rules to allow management traffic from Multicloud Defense Gateway to the Multicloud Defense Controller in the management VPC.

If you find that the script does not work, or if you need to manually change your settings, these actions can be executed using the GCP cloud console web UI, or using the gcloud CLI. See the alternative method of connecting your project here.

Service Accounts

Explains the service accounts needed for your GCP project.

Multicloud Defense requires two service accounts created in your GCP project:

- **Multicloud Defense-controller**: This account is used by theMulticloud Defense Controller to access your GCP project to create resources (Multicloud Defense Gateways), load balancers for Multicloud Defense Gateways, and read information about the VPCs, Subnets, Security Group tags etc.
- **Multicloud Defense-gateway**: This account is assigned to the Multicloud Defense Gateways (Compute VM instances). The account provides access to the secret manager (private keys for TLS decryption) and storage.

You can create these service accounts in one of two ways: by using the service available in the UI or by using the the cloud service provider's CLI.

Create Multicloud Defense Controller Service Account Using GCP Cloud Console

Create a Multicloud Defense controller service account in Google Cloud Console, assign the required roles, and download its JSON key for GCP project onboarding.

The Multicloud Defense Controller service account is used by the Multicloud Defense Controller to access and manage resources in your GCP project. You must create the account and generate a key. The key is added to the Controller as part of account onboarding to the Controller.

Procedure

1. Open **IAM** in your GCP project.
2. Click **Service Accounts**.
3. Create **Service Account**.
4. Provide a name and ID (e.g Multicloud Defense-controller) and click **Create**.
5. Add **Compute Admin** and **Service Account User roles**.
6. Click **Continue**.
7. Click **Done**.

> **Note**
>
> There is no requirement to add any users.

8. Click on the newly created account, scroll down to **Keys** and in the dropdown for **Add Key** and select **Create New Key**.
9. Choose JSON (default option) and click **Create**.
10. A file is downloaded to your computer. Save this file.

Create a Multicloud Defense Firewall Service Account Using the GCP Cloud Console

Create a Multicloud Defense firewall service account in Google Cloud Console and assign permissions for secret access and log forwarding from GCP gateways.

The Multicloud Defense firewall service account is used by the Multicloud Defense Gateway instances running inside your GCP project. The Gateways may need to access the private keys stored in the SecretManager for TLS decryption and access storage to store PCAP files etc. (if configured by the user). Also, the Gateways many need Log Writer permissions to send logs from Multicloud Defense Gateway to the GCP logging instance (if configured by the user).

Below are two (2) methods of creating this service account.

Procedure

1. Open **IAM** in your GCP project.
2. Click **Service Accounts**.
3. Create **Service Account**.
4. Provide a name and ID (e.g Multicloud Defense-firewall) and click **Create**.
5. Add **Secret Manager**, **Secret Accessor** and **Logs Writer roles**.
6. Click **Continue**.
7. Click **Done**.

> **Note**
>
> There is no requirement to add any users.

Enable API

Learn how to enable the GCP APIs required for communication with Multicloud Defense by using either Google Cloud Console or the gcloud CLI.

You can enable the API for communication between Multicloud Defense Controller and your GCP account with either GCP console or the cloud service provider's CLI.

Enable API-Using the GCP Cloud Console

Enable the GCP APIs required for Multicloud Defense to create and manage cloud resources by using Google Cloud Console during onboarding.

Enable the APIs in your project/account so that the Multicloud Defense Controller can create Multicloud Defense Gateways (Virtual Machines, Load Balancers).

Procedure

1. Search for **Compute Engine API** in the searchbar.
2. Click **Enable**.
3. Search for **Secret Manager API** in the searchbar.
4. Click **Enable**.
5. Search for **Identity and Access Management (IAM) API** in the searchbar.

6. Click **Enable**.
7. Search for **Cloud Resource Manager API** in the searchbar.
8. Click **Enable**.

VPC Setup

Learn about preparing GCP VPCs for a Multicloud Defense gateway deployment in edge mode, where gateway instances run with the applications they protect.

Multicloud Defense Gateway instances can be deployed in edge or hub mode. In edge mode, the gateway instances run in the same VPC as your applications. This document focuses on preparing you to deploy the Multicloud Defense Gateway deployment in edge mode.

VPC and Subnets

Learn about the management and datapath VPCs and subnets required for Multicloud Defense gateways in GCP. Understand where gateway interfaces connect and which traffic each network carries.

When deploying the Multicloud Defense Gateway, the Multicloud Defense Controller will prompt for the **management** and **datapath** VPC information. Multicloud Defense Gateway instances require two network interfaces. In GCP, the network interfaces of a VM instance need to be in different VPCs unlike other cloud providers where they can be in just different subnets. If you already have a VPC where the application is running, you have the **datapath** VPC and the subnet. You must create another VPC (or use an another existing VPC) for management purposes. You can either use the auto-created subnets or create them manually.

The datapath vpc is the VPC where your applications are running and will be referred to as such in the following sections

In each of the VPCs, Multicloud Defense requires one subnet for datapath and one subnet for management.

The **management** subnet is a public subnet that must be associated with the route table that has a default route to the Internet. The Multicloud Defense Gateway instance has an interface attached to this subnet that it uses to communicate with the Multicloud Defense Controller. This interface is used for policy pushes and other management and telemetry activities between the Multicloud Defense Controller and the Multicloud Defense Gateway instances. Customer application traffic **does not** flow through this interface and subnet. The interface is associated with the **Multicloud Defense-management** network tag (or any tag based on your team requirements), which is described in the network tags section below.

The **datapath** subnet is a public subnet that must be associated with the route table that has a default route to the Internet. The Multicloud Defense Controller creates a network load balancer (NLB) in this subnet. In addition, a Multicloud Defense Gateway instance has an interface attached to this subnet. The customer applications traffic **flows** through this interface. A security policy is applied to the traffic ingressing through this interface. The interface is associated with the **Multicloud Defense-datapath** network tag (or any tag based on your team requirements), which is described in the network tags section below.

Choosing a service VPC CIDR

Describes the service VPC CIDR options and their implications for subnetting and instance capacity.

For each CIDR prefix, the table shows the maximum number of Availability Zones (AZs) supported and the maximum safe instances per AZ (that is, the highest `max_instances` value that allows a rolling upgrade to complete without IP exhaustion).

Note

If the service VPC is shared among gateways, which is generally not recommended, the CIDR block must be adjusted accordingly to avoid overlaps. Additionally, there may be restrictions affecting whether gateways sharing the same VPC can upgrade, auto-scale, or recover from failure scenarios simultaneously.

AWS — Without NAT Gateway (5 subnets/AZ)

VPC CIDR	Subnet Slots	Subnets Used	Maximum AZs	Max Instances/AZ
/25	8	5 * AZ	1 [1]	5
/24	16	5 * AZ	3	5
/23	16	5 * AZ	3	13
/22	32	5 * AZ	6	13
/20	64	5 * AZ	12	29
/16	256	5 * AZ	51	125

AWS — With NAT Gateway (6 subnets/AZ)

VPC CIDR	Subnet Slots	Subnets Used	Maximum AZs	Max Instances/AZ
/25	8	6 * AZ	1 [2]	5
/24	16	6 * AZ	2 [3]	5
/23	16	6 * AZ	2	13
/22	32	6 * AZ	5	13
/20	64	6 * AZ	10	29
/16	256	6 * AZ	42	125

[1] **/25 supports 1 AZ only.** *There are 8 slots. 8 slots ÷ 5 subnets/AZ = 1. A 2-AZ request needs 10 subnets but only 8 slots exist. The creation of SVPC fails.*

Number of subnets (10) exceeds the maximum supported subnets (8) for VPC CIDR X.X.X.X/25

[2] **/25 supports 1 AZ only.** There are 8 slots. 8 slots ÷ 6 subnets/AZ = 1. Same hard limit as no NAT.

[3] **/24 with NAT — 2 AZs only** (vs 3 without NAT). The extra NAT-GW subnet per AZ reduces the AZ ceiling from 3 to 2 for a /24. A 3-AZ request needs 18 subnets but only 16 slots exists.The creation of SVPC fails.

Azure MCD (3 subnets/AZ, 2 IPs per instance)

Azure allocates **2 IPs per instance** from the datapath subnet (primary + secondary NIC config), which halves the effective instance capacity compared to AWS.

VPC CIDR	Subnet Slots	Subnets Used	Maximum AZs	Max Instances/AZ
/25	8	3 * AZ	2	2
/24	16	3 * AZ	5	2
/23	16	3 * AZ	5	6
/22	32	3 * AZ	10	6
/20	64	3 * AZ	21	14
/16	256	3 * AZ	85	62

When several gateways share one service VPC, **Max instances/AZ** is a **combined** limit across all gateways in that AZ – not per gateway. Two gateways on a /24 get max=1/AZ each; use /22 for max=3/AZ each.

GCP (1 subnet/VPC, no per-AZ subdivision)

GCP creates three separate VPCs (Management, Datapath, Second-Datapath), each with a single subnet spanning the full CIDR. IP addresses are not separated by AZ. All instances in all AZs share one pool per VPC. The limit is total instances across all AZs.

VPC CIDR	Subnets Used	Max Instances (All AZs)
/25	1 per VPC	62
/24	1 per VPC	126
/23	1 per VPC	254
/22	1 per VPC	510
/20	1 per VPC	2046
/16	1 per VPC	32766

GCP is not practically constrained by IP exhaustion at typical gateway instance counts.

Note

- **Subnet slots** = `2 ^ floor((32 – P) / 2)` where P is the CIDR prefix.
- **Max AZs** = `floor(subnet_slots / subnets_per_AZ)`. Fixed at SVPC creation – cannot be changed later.
- **Max instances/AZ** = `floor(usable / 2)` for AWS; `floor(usable / 4)` for Azure MCD (due to 2 IPs/instance).

Where usable = `2 ^ ceil((32 – P) / 2) – 5` (AWS/Azure) or `2 ^ ceil((32 – P) / 2) – 4` (GCP). This is the safe ceiling for rolling upgrades – higher values work at steady state but cause IP exhaustion when an upgrade runs.

For formulas, multi-gateway sizing, and upgrade peak analysis, refer to https://eygupta-doc.dev.valtixinc.com/valtix-sw/svpc-cidr-guide/svpc-cidr-guide.html.

Sample VPC and Subnets using CLI

Create sample management and datapath VPCs, subnets, and firewall rules for a Multicloud Defense GCP gateway deployment by using gcloud CLI commands.

Use the following commands as an example when executing your own commands to create VPCs for your GCP account. Open the Google Cloud Shell windows for these particular commands:

Procedure

1. Create VPC **apps** and subnet **apps-us-east1**

2. Create VPC `Multicloud Defense-mgmt` and subnet `Multicloud Defense-mgmt-us-east1`:

3. Create at least two Firewall rules for VPC `Multicloud Defense-mgmt` with target-tags as `Multicloud Defense-mgmt`:
   a. Egress rule to allow all the outbound traffic:
   b. Ingress rule to allow SSH into the firewall instances:

4. Create at least three Firewall rules for VPC **apps**. Use the following as examples:
   a. One egress rule to allow all the outbound traffic with target-tags as `Multicloud Defense-datapath`:
   b. One ingress rule to allow HTTP and HTTPS into the gateway instances through the non-load balancerwith target-tags as `Multicloud Defense-datapath`:

c. Once egress rule to allow all the outbound traffic with target-tags as `app-instance`:

d. One ingress rule to allow `tcp:8000` with `target-tags` as `app-instance`:

```
gcloud config set project <project-name> # incase the project is not set in
the gcloud cli shell
gcloud compute networks create apps --subnet-mode custom
gcloud compute networks subnets create apps-us-east1 --network apps --range
10.0.0.0/24 --region us-east1
gcloud compute networks create ciscomcd-mgmt --subnet-mode custom
gcloud compute networks subnets create ciscomcd-mgmt-us-east1 --network
ciscomcd-mgmt --range 172.16.0.0/24 --region us-east1
gcloud compute firewall-rules create ciscomcd-mgmt-out --direction EGRESS
--network ciscomcd-mgmt \
    --target-tags ciscomcd-mgmt --allow tcp,udp
gcloud compute firewall-rules create ciscomcd-mgmt-in --direction INGRESS
--network ciscomcd-mgmt \
    --target-tags ciscomcd-mgmt --allow tcp:22
gcloud compute firewall-rules create ciscomcd-datapath-out --direction EGRESS
 --network apps \
    --target-tags ciscomcd-datapath --allow tcp,udp
gcloud compute firewall-rules create ciscomcd-datapath-in --direction INGRESS
 --network apps \
    --target-tags ciscomcd-datapath --allow tcp:80,tcp:443
gcloud compute firewall-rules create app-instance-out --direction EGRESS
--network apps \
    --target-tags app-instance --allow tcp,udp
gcloud compute firewall-rules create app-instance-in --direction INGRESS
--network apps \
    --target-tags app-instance --allow tcp:8000,tcp:22
```

Once you run the above commands, you can create a VM instance in the **apps** VPC and launch a test web application on port 8000.

```
gcloud compute instances create app-instance1 \
  --zone=us-east1-b \
  --image-project=ubuntu-os-cloud \
  --image-family=ubuntu-2004-lts \
  --network apps \
  --subnet=apps-us-east1 \
  --tags=app-instance
gcloud compute ssh app-instance1 --zone us-east1-b
echo hello world > index.html
python3 -m http.server 8000
```

Network Tags (for GCP Gateways)

Learn about configuring management and datapath network tags for GCP gateways. Identify the firewall access required for controller communication, administrative access, and protected application traffic.

The management and datapath network tags are associated with the respective interfaces on the Multicloud Defense Gateway instance, as described in the subnets section above.

Create a gateway rule in the **management** VPC and associate that with **Multicloud Defense-management** network tag. This must allow all outbound traffic that makes the gateway instance communicate with the controller.

Optionally, for inbound rules, enable port 22 (SSH) to allow SSH access to the gateway instance. SSH is **not required** for the Multicloud Defense firewall to function properly.

Create a gateway rule in the **datapath** VPC and associate that with **Multicloud Defense-datapath** network tag. This must allow the traffic to the Multicloud Defense Gateway for all the services that you enable (are going to enable).

For example, if an application is running on port 3000 and is proxied by the Multicloud Defense Gateway on port 443, port 443 must be opened on the Multicloud Defense-datapath network security tag.

Gateway Creation

Learn about the GCP VPC, subnet, network tag, and region values used when creating a Multicloud Defense gateway for an edge-mode deployment.

Using the Multicloud Defense Gateway creation page use the following parameters:

1. Datapath VPC: **apps**.
2. Datapath Network Tag: **Multicloud Defense-datapath**.
3. Management VPC: **Multicloud Defense-mgmt**.
4. Management Network Tag: **Multicloud Defense-mgmt**.
5. Use **us-east1-b** zone.
6. Management Subnet: **Multicloud Defense-mgmt-us-east1**.
7. Datapath Subnet: **apps-us-east1**.

You can create subnets in other regions to test the Multicloud Defense Gateway in multi-availability zone mode.

Manually Onboard an Azure Subscription

Learn about manually onboarding an Azure subscription to Multicloud Defense when the dashboard script cannot establish the connection. Use the advanced procedures to prepare and connect the subscription.

If you cannot directly connect an Azure subscription with the script provided in the Multicloud Defense Controller dashboard, use the procedures below to manually connect your subscription

(Optional) User-assigned Managed Identity for Key Vault and Blob Storage access

Learn about using an optional Azure user-assigned managed identity to give Multicloud Defense gateways access to Key Vault certificates and Blob Storage for packet captures.

Multicloud Defense Gateways can optionally integrate with Azure Key Vault to retrieve TLS certificates and with Blob Storage for saving PCAP (packet capture) files. User-assigned managed identities are used to grant access to these services.

In the Azure portal, navigate to **Managed Identities** to create an identity.

Alternatively in Azure Cloud Shell, run the following command:

```
az identity create -g <RESOURCE GROUP> -n <USER ASSIGNED IDENTITY NAME>
```

For information on creating TLS certificate secrets in Azure Key Vault, see Azure Key Vault on page 170.

Register Application in Microsoft Entra ID

Register the Cisco Multicloud Defense application in Microsoft Entra ID and collect the identifiers and credentials required to connect your Azure subscription.

Use the following procedure to register the Multicloud Defense application in your Entra ID.

Procedure

1. From your Azure portal, navigate to **Microsoft Entra ID**.
2. Select **App registrations**.
3. Click **New registration**.
4. Provide a name to reference the new app registration e.g. Multicloud Defense Controller In the *Supported account types* choose the second option *Accounts in any organizational directory*.
5. Choose the option appropriate to your organization. Note that the **Redirect URI** is not needed for the creation of the App registration.
6. Click **Register**.
7. Navigate to **App registrations** section and click on the newly created application.
8. Navigate to **Manage** section and click **Certificates & secrets**.
9. Click **+ New client secret**, and then enter the required information in the *Add a client secret* dialog box
   - **Description** - Add a description (for example, Multicloud Defense-controller-secret1)
   - **Expires** - Choose **24 months**. You can also make this selection at your convenience. You will need to create new secrets when the current one expires.
10. Click **Add**. The client secret is populated under the **Value** column.
11. Copy the **Client secret** into a notepad, as this is shown only once and is never displayed again.
12. In the left navigation bar click **Overview**.
13. Copy the **Application (client) ID** and **Directory (tenant) ID** into a notepad.

Create a Custom Role to Assign to the Application

Create and assign an Azure custom role that grants the Cisco Multicloud Defense application the required access to your subscription.

The CloudFormation template creates a **custom role** that will be assigned to the application created for the Multicloud Defense Controller. The custom role gives the application permissions to read inventory information and create resources, for example, VMs, load balancers, etc.

There are multiple ways to create a custom role but we recommend the following procedure:

Procedure

1. Navigate to **Subscription** and click **Access Control (IAM)**.
2. Click on **Roles** and on the top menu bar navigate to click **+Add** > **Add Custom Role**.
3. Give a name to the custom role, for example, `Multicloud Defense-controller-role`.
4. Keep clicking **Next** until you get to the JSON editing screen.
5. Click **Edit** on the screen and in the JSON text, under the **permissions > actions** section, copy and paste the provided code content in between the square brackets (adjust the indentation if needed). Note that you are also adding permissions in the Azure role to read role assignments and role definitions, as indicated in the code content.

```
"Microsoft.ApiManagement/service/*",
"Microsoft.Compute/disks/*",
"Microsoft.Compute/images/read",
```

```
"Microsoft.Compute/sshPublicKeys/read",
"Microsoft.Compute/virtualMachines/*",
"Microsoft.ManagedIdentity/userAssignedIdentities/read",
"Microsoft.ManagedIdentity/userAssignedIdentities/assign/action",
"Microsoft.Network/loadBalancers/*",
"Microsoft.Network/natGateways/*",
"Microsoft.Network/networkinterfaces/*",
"Microsoft.Network/networkSecurityGroups/*",
"Microsoft.Network/publicIPAddresses/*",
"Microsoft.Network/routeTables/*",
"Microsoft.Network/virtualNetworks/*",
"Microsoft.Network/virtualNetworks/subnets/*",
"Microsoft.Resources/subscriptions/resourcegroups/*",
"Microsoft.Storage/storageAccounts/blobServices/*",
"Microsoft.Storage/storageAccounts/listkeys/action",
"Microsoft.Network/networkWatchers/*",
"Microsoft.Network/applicationSecurityGroups/*",
"Microsoft.Compute/diskEncryptionSets/read",
"Microsoft.Insights/Metrics/Read"
"Microsoft.Network/locations/serviceTagDetails/read",
"Microsoft.Network/locations/serviceTags/read",
"Microsoft.CognitiveServices/*/read",
"Microsoft.CognitiveServices/accounts/listkeys/action",
"Microsoft.Network/virtualHubs/*",
"Microsoft.Network/virtualHubs/hubRouteTables/*",
"Microsoft.MarketplaceOrdering/offertypes/publishers/offers/plans/agreements/*"
"Microsoft.Authorization/roleAssignments/read",
"Microsoft.Authorization/roleDefinitions/read"
```

6. **Optional** - If you plan to use multiple subscriptions with Multicloud Defense, you must edit the JSON at `assignableScopes` to add another subscription line or change it to ***** (star) so the custom role can be used with all subscriptions.
7. Click **Save** at the top of the text box.
8. Click **Review + Create** and create the role.
9. Once the custom role is created, return to **Access Control (IAM).**
10. On the top menu bar, click **Add** > **Add role assignment**.
11. In the **Role** dropdown, select the custom role created above.
12. In the **Assign access to** radio button, leave it as the default (User, group, or service principal).
13. Click **Select members**. Select the required users, groups, or service principals and click **Save**.
14. Navigate to **Review + Create** tab and copy the subscription ID from the **Scope** field to the notepad.

Required Values For Multicloud Defense Controller Onboarding

Find the Azure subscription, tenant, application, and client secret values required to manually onboard an Azure subscription to Multicloud Defense.

Make sure you have the following information before proceeding further:

- Subscription ID (*from subscription overview page*)
- Directory (Tenant) ID (*from the Azure AD app overview page*)
- Application (client) ID (*from the Azure AD app overview page*)
- Client Secret (*Copied when the Client secret was created*)

Accept Marketplace Terms

Accept the Cisco Multicloud Defense virtual machine image terms for each Azure subscription before deploying gateways from Azure Marketplace.

Multicloud Defense Controller creates gateway instances using a Multicloud Defense virtual machine (VM) image from the Azure marketplace. The Terms and Conditions must be accepted for each subscription. Open the Azure cloud shell from the Azure portal website (on the top menubar towards the right side). Choose or switch to bash shell and execute the following command (replace the subscription-id with your subscription id copied in the previous section):

```
az vm image terms accept --subscription $sub_id --publisher valtix --offer
datapath --plan valtix_dp_image
```

Graceful Termination of Connections

Learn about graceful termination of TCP connections on Multicloud Defense gateways and how it helps applications recover from interrupted flows. Understand when connections may end and what behavior to expect.

Multicloud Defense Gateway can choose to terminate an established flow for multiple reasons such as:

- Termination based on the policy. For example, FQDN filtering can only be applied after the flow is established.
- IDS/IPS can deem any packet in the flow that is sent by either the client or the server to be unsafe and can choose to terminate an established flow.
- Proxy service on the Multicloud Defense Gateway decides to terminate the flow after the flow is established.
- When one of the timers in the Multicloud Defense Gateway TCP stack decides that the flow is no longer active or alive.
- Flow termination during certain configuration changes such as PRS updates, gateway setting changes and so on.
- Flow termination when the gateway is decommissioned (controller initiated - disable/upgrade/scale-in).

Currently, when a Multicloud Defense Gateway chooses to terminate an established flow for any of the above reasons, it does so without informing the client and the server about the termination (except if there is FQDN Filtering with Reset on Deny turned on). This causes the client and server to rely on TCP or application timeouts to detect the loss of connection, causing application outages.

For TCP flows, Multicloud Defense Gateway introduces a graceful termination mechanism which causes the gateway to send a TCP Reset to the client (initiator) when the gateway stops the flow. This should enable the client TCP stack to terminate the connection quickly, enabling the applications to attempt to re-establish the interrupted flow, thereby minimizing traffic disruption. This applies to all kinds of flows - forwarded, forward proxied, and reverse proxied, that are handled by the Multicloud Defense Gateway.

Also, when a Multicloud Defense Gateway data plane goes down unexpectedly (due to a software issue), this reset mechanism does not apply. Clients will continue to rely on application timeouts to recover.

Troubleshooting

To find flows that are terminated with a TCP Reset by the Multicloud Defense Gateway, download the traffic summary (from the controller) as a CSV and search for *RESET*. It will be the last connection state for the ingress flow. Connections that are terminated naturally will not have this state as the last state. For non-TCP flows, the last connection state is always *AGED OUT*.

Terraform Onboarding Scripts for Cloud Accounts

Onboard a cloud service provider account to Multicloud Defense with a Terraform script as an alternative to the onboarding wizard or manual setup.

You an use the terraform script to onboard your cloud service provider account instead of using the onboarding wizard or the manual process.

About Terraform

Learn about using the verified Multicloud Defense Terraform provider to onboard cloud accounts, discover assets, deploy gateways, and apply security policies across AWS, Azure, GCP, and OCI environments.

Multicloud Defense customers can use the **Terraform Provider** to: **discover** - onboard public cloud accounts, gain continuous asset visibility and detect indicators of compromise (IoC); **deploy** - Multicloud Defense Gateways to protect ingress, egress and east-west traffic; and **defend** - with multicloud (AWS, Azure, GCP, OCI) dynamic policies with continuously discovered cloud assets.

> **Attention**
>
> As of Multicloud Defense Controller version 23.10, you can connect a GCP folder as well as a GCP project using the Terraform provider. See Terraform Repository on page 344 for more information.

The Multicloud Defense terraform provider is a “Verified” provider available from the Terraform registry. Customers can now use the Terraform provider for Multicloud Defense to bake security into their operations, that is, onboard their cloud accounts into Multicloud Defense, deploy Multicloud Defense Gateways and specify security policies to protect against ingress attacks from the Internet (WAF, IDS/IPS, Geo-IP), stop exfiltration on egress traffic (TLS decryption, IDS/IPS, AV, DLP, FQDN/URL filtering), and prevent east-west attacks between VPCs/VNets. The security policies can be specified based on cloud asset tags (for example, “dev”, “test”, “prod”, “pci”, “web”, “app1” and more).

For more information, refer to:

- Download the Terraform Provider for Multicloud Defense.
- Examples in GitHub.

Terraform Repository

Find Terraform repositories for onboarding AWS accounts, Azure subscriptions, and GCP projects or folders, including supported discovery-only deployment options in Multicloud Defense.

Use case	Description	Github Repository
AWS onboarding	This is for onboarding AWS account using Terraform.	AWS Github Repo
AWS discovery CFT	This CFT deployment will include all necessary privileges needed to use Multicloud Defense's discovery feature. For full feature set, please use the native product CFT.	AWS Discovery Github Repo
AWS discovery	This is for onboarding AWS account for discovery only mode using Terraform.	AWS Github Repo

Use case	Description	Github Repository
Azure onboarding	This is for onboarding Azure subscription using Terraform.	Azure Github Repo
GCP Project onboarding	This is for onboarding GCP project using Terraform.	GCP Github Repo
GCP Folder onboarding	This is for onboarding GCP folder using Terraform.	GCP Github Repo

Exporting Configuration as Terraform Block

Learn how to export a Multicloud Defense security profile as a Terraform resource block for reuse in infrastructure-as-code workflows. Understand which objects and profiles do not support Terraform export.

Customers can export security profiles into terraform resource blocks from Multicloud Defense Controller. To export configuration into Terraform block, navigate and select the intended security profile and click on **Export** button. This will download a file that has the terraform block for the selected object/security profile.

All objects and profiles support terraform export with the exception of:

- Gateways
- Service VPCs/VNets
- Diagnostics