Clustering for Threat Defense Virtual in a Public Cloud

Clustering lets you group multiple Threat Defense Virtuals together as a single logical device. A cluster provides all the convenience of a single device (management, integration into a network) while achieving the increased throughput and redundancy of multiple devices. You can deploy Threat Defense Virtual clusters in a public cloud using the following public cloud platforms:

- Amazon Web Services (AWS)
- Microsoft Azure
- Google Cloud Platform (GCP)

Currently, only routed firewall mode is supported.

Note Some features are not supported when using clustering. See Unsupported Features and Clustering, on page 75.

About Threat Defense Virtual Clustering in the Public Cloud

This section describes the clustering architecture and how it works.

How the Cluster Fits into Your Network

The cluster consists of multiple firewalls acting as a single device. To act as a cluster, the firewalls need the following infrastructure:

- Isolated network for intra-cluster communication, known as the *cluster control link*, using VXLAN interfaces. VXLANs, which act as Layer 2 virtual networks over Layer 3 physical networks, let the Threat Defense Virtual send broadcast/multicast messages over the cluster control link.
- Load Balancer(s)—For external load balancing, you have the following options depending on your public cloud:
  - AWS Gateway Load Balancer

    The AWS Gateway Load Balancer combines a transparent network gateway and a load balancer that distributes traffic and scales virtual appliances on demand. The Threat Defense Virtual supports the Gateway Load Balancer centralized control plane with a distributed data plane (Gateway Load Balancer endpoint) using a Geneve interface single-arm proxy.
  - Azure Gateway Load Balancer

    In an Azure service chain, Threat Defense Virtuals act as a transparent gateway that can intercept packets between the internet and the customer service. The Threat Defense Virtual defines an external interface and an internal interface on a single NIC by utilizing VXLAN segments in a paired proxy.
  - Native GCP load balancers, internal and external
  - Equal-Cost Multi-Path Routing (ECMP) using inside and outside routers such as Cisco Cloud Services Router

    ECMP routing can forward packets over multiple “best paths” that tie for top place in the routing metric. Like EtherChannel, a hash of source and destination IP addresses and/or source and destination ports can be used to send a packet to one of the next hops. If you use static routes for ECMP routing, then the Threat Defense failure can cause problems; the route continues to be used, and traffic to the failed Threat Defense will be lost. If you use static routes, be sure to use a static route monitoring feature such as Object Tracking. We recommend using dynamic routing protocols to add and remove routes, in which case, you must configure each Threat Defense to participate in dynamic routing.

Note Layer 2 Spanned EtherChannels are not supported for load balancing.

Individual Interfaces

You can configure cluster interfaces as *Individual interfaces*.

Individual interfaces are normal routed interfaces, each with their own local IP address. Interface configuration must be configured only on the control node, and each interface uses DHCP.

Note Layer 2 Spanned EtherChannels are not supported.

Control and Data Node Roles

One member of the cluster is the control node. If multiple cluster nodes come online at the same time, the control node is determined by the priority setting; the priority is set between 1 and 100, where 1 is the highest priority. All other members are data nodes. When you first create the cluster, you specify which node you want to be the control node, and it will become the control node simply because it is the first node added to the cluster.

All nodes in the cluster share the same configuration. The node that you initially specify as the control node will overwrite the configuration on the data nodes when they join the cluster, so you only need to perform initial configuration on the control node before you form the cluster.

Some features do not scale in a cluster, and the control node handles all traffic for those features.

Cluster Control Link

Each node must dedicate one interface as a VXLAN (VTEP) interface for the cluster control link. For more information about VXLAN, see Configure VXLAN Interfaces.

VXLAN Tunnel Endpoint

VXLAN tunnel endpoint (VTEP) devices perform VXLAN encapsulation and decapsulation. Each VTEP has two interface types: one or more virtual interfaces called VXLAN Network Identifier (VNI) interfaces, and a regular interface called the VTEP source interface that tunnels the VNI interfaces between VTEPs. The VTEP source interface is attached to the transport IP network for VTEP-to-VTEP communication.

VTEP Source Interface

The VTEP source interface is a regular threat defense virtual interface with which you plan to associate the VNI interface. You can configure one VTEP source interface to act as the cluster control link. The source interface is reserved for cluster control link use only. Each VTEP source interface has an IP address on the same subnet. This subnet should be isolated from all other traffic, and should include only the cluster control link interfaces.

VNI Interface

A VNI interface is similar to a VLAN interface: it is a virtual interface that keeps network traffic separated on a given physical interface by using tagging. You can only configure one VNI interface. Each VNI interface has an IP address on the same subnet.

Peer VTEPs

Unlike regular VXLAN for data interfaces, which allows a single VTEP peer, The threat defense virtual clustering allows you to configure multiple peers.

Cluster Control Link Traffic Overview

Cluster control link traffic includes both control and data traffic.

Control traffic includes:

- Control node election.
- Configuration replication.
- Health monitoring.

Data traffic includes:

- State replication.
- Connection ownership queries and data packet forwarding.

Configuration Replication

All nodes in the cluster share a single configuration. You can only make configuration changes on the control node (with the exception of the bootstrap configuration), and changes are automatically synced to all other nodes in the cluster.

Management Network

You must manage each node using the Management interface; management from a data interface is not supported with clustering.

Licenses for Threat Defense Virtual Clustering

Each threat defense virtual cluster node requires the same performance tier license. We recommend using the same number of CPUs and memory for all members, or else performance will be limited on all nodes to match the least capable member. The throughput level will be replicated from the control node to each data node so they match.

You assign feature licenses to the cluster as a whole, not to individual nodes. However, each node of the cluster consumes a separate license for each feature. The clustering feature itself does not require any licenses.

When you add the control node to the Management Center, you can specify the feature licenses you want to use for the cluster. You can modify licenses for the cluster in the **Devices** > **Device Management** > **Cluster** > **License** area.

Note If you add the cluster before the Management Center is licensed (and running in Evaluation mode), then when you license the Management Center, you can experience traffic disruption when you deploy policy changes to the cluster. Changing to licensed mode causes all data units to leave the cluster and then rejoin.

Requirements and Prerequisites for Threat Defense Virtual Clustering

Model Requirements

- FTDv5, FTDv10, FTDv20, FTDv30, FTDv50, FTDv100

> **Note** FTDv5 and FTDv10 do not support Amazon Web Services (AWS) Gateway Load Balancer (GWLB) and Azure GWLB.

- The following public cloud services:
    - Amazon Web Services (AWS)
    - Microsoft Azure
    - Google Cloud Platform (GCP)
- Maximum 16 nodes

See also the general requirements for the Threat Defense Virtual in the Cisco Secure Firewall Threat Defense Virtual Getting Started Guide.

User Roles

- Admin
- Access Admin
- Network Admin

Hardware and Software Requirements

All units in a cluster:

- Must be in the same performance tier. We recommend using the same number of CPUs and memory for all nodes, or else performance will be limited on all nodes to match the least capable node.
- The Management Center access must be from the Management interface; data interface management is not supported.
- Must run the identical software except at the time of an image upgrade. Hitless upgrade is supported.
- All units in a cluster must be deployed in the same availability zone.
- Cluster control link interfaces of all units must be in the same subnet.

MTU

Make sure the ports connected to the cluster control link have the correct (higher) MTU configured. If there is an MTU mismatch, the cluster formation will fail. The cluster control link MTU should be 154 bytes higher than the data interfaces. Because the cluster control link traffic includes data packet forwarding, the cluster control link needs to accommodate the entire size of a data packet plus cluster traffic overhead (100 bytes) plus VXLAN overhead (54 bytes).

For AWS with GWLB, the data interface uses Geneve encapsulation. In this case, the entire Ethernet datagram is being encapsulated, so the new packet is larger and requires a larger MTU. You should set the source interface MTU to be the network MTU + 306 bytes. So for the standard 1500 MTU network path, the source interface MTU should be 1806, and the cluster control link MTU should be +154, 1960.

For Azure with GWLB, the data interface uses VXLAN encapsulation. In this case, the entire Ethernet datagram is being encapsulated, so the new packet is larger and requires a larger MTU. You should set the cluster control link MTU to be the source interface MTU + 80 bytes.

The following table shows the default values for the cluster control link MTU and the data interface MTU.

Note We do not recommend setting the cluster control link MTU between 2561 and 8362; due to block pool handling, this MTU size is not optimal for system operation.

Table 1: Default MTU

Public Cloud	Cluster Control Link MTU	Data Interface MTU
AWS with GWLB	1980	1826
AWS	1654	1500
Azure with GWLB	1454	1374
Azure	1454	1300
GCP	1554	1400

Guidelines for Threat Defense Virtual Clustering

High Availability

High Availability is not supported with clustering.

IPv6

The cluster control link is only supported using IPv4.

Additional Guidelines

- When significant topology changes occur (such as adding or removing an EtherChannel interface, enabling or disabling an interface on the Threat Defense or the switch, adding an additional switch to form a VSS or vPC) you should disable the health check feature and also disable interface monitoring for the disabled interfaces. When the topology change is complete, and the configuration change is synced to all units, you can re-enable the interface health check feature.

- When adding a node to an existing cluster, or when reloading a node, there will be a temporary, limited packet/connection drop; this is expected behavior. In some cases, the dropped packets can hang your connection; for example, dropping a FIN/ACK packet for an FTP connection will make the FTP client hang. In this case, you need to reestablish the FTP connection.
- Do not power off a node without first disabling clustering on the node.
- For decrypted TLS/SSL connections, the decryption states are not synchronized, and if the connection owner fails, then decrypted connections will be reset. New connections will need to be established to a new node. Connections that are not decrypted (they match a do-not-decrypt rule) are not affected and are replicated correctly.
- Dynamic scaling is not supported.
- If you are using Secure Firewall versions 7.2 or 7.3, Stateful Target Failover is not supported when you deploy the cluster on AWS.
- Perform a global deployment after the completion of each maintenance window.
- Ensure that you do not remove more than one device at a time from the auto scale group (AWS) / instance group (GCP) / scale set (Azure). We also recommend that you run the **cluster disable** command on the device before removing the device from the auto scale group (AWS) / instance group (GCP) / scale set (Azure).
- If you want to disable data nodes and the control node in a cluster, we recommend that you disable the data nodes before disabling the control node. If a control node is disabled while there are other data nodes in the cluster, one of the data nodes has to be promoted to be the control node. Note that the role change could disturb the cluster.
- In the customized day 0 configuration scripts given in this guide, you can change the IP addresses as per your requirement, provide custom interface names, and change the sequence of the CCL-Link interface.
- If you experience CCL instability issues, such as intermittent ping failures, after deploying a Threat Defense Virtual cluster on a cloud platform, we recommend that you address the reasons that are causing CCL instability. Also, you can increase the hold time as a temporary workaround to mitigate CCL instability issues to a certain extent. For more information on how to change the hold time, see Edit Cluster Health Monitor Settings.
- When you are configuring your security firewall rule or security group for the Management Center virtual, you must include both Private and Public IP addresses of the Threat Defense Virtual in the Source IP address range. Also, ensure to specify the Private and Public IP addresses of the Management Center Virtual in the security firewall rule or security group of the Threat Defense Virtual. This is important to ensure proper registration of nodes during clustering deployment.

Defaults for Clustering

- The cLACP system ID is auto-generated, and the system priority is 1 by default.
- The cluster health check feature is enabled by default with the holdtime of 3 seconds. Interface health monitoring is enabled on all interfaces by default.
- The cluster auto-rejoin feature for a failed cluster control link is unlimited attempts every 5 minutes.

- The cluster auto-rejoin feature for a failed data interface is 3 attempts every 5 minutes, with the increasing interval set to 2.
- Connection replication delay of 5 seconds is enabled by default for HTTP traffic.

Deploy the Cluster in AWS

To deploy a cluster in AWS, you can either manually deploy or use CloudFormation templates to deploy a stack. You can use the cluster with AWS Gateway Load Balancer, or with a non-native load-balancer such as the Cisco Cloud Services Router.

AWS Gateway Load Balancer and Geneve Single-Arm Proxy

Note This use case is the only currently supported use case for Geneve interfaces.

The AWS Gateway Load Balancer combines a transparent network gateway and a load balancer that distributes traffic and scales virtual appliances on demand. The Threat Defense Virtual supports the Gateway Load Balancer centralized control plane with a distributed data plane (Gateway Load Balancer endpoint). The following figure shows traffic forwarded to the Gateway Load Balancer from the Gateway Load Balancer endpoint. The Gateway Load Balancer balances traffic among multiple Threat Defense Virtuals, which inspect the traffic before either dropping it or sending it back to the Gateway Load Balancer (U-turn traffic). The Gateway Load Balancer then sends the traffic back to the Gateway Load Balancer endpoint and to the destination.

Note Transport Layer Security (TLS) Server Identity Discovery is not supported with Geneve single-arm setup on AWS.

Figure 1: Geneve Single-Arm Proxy

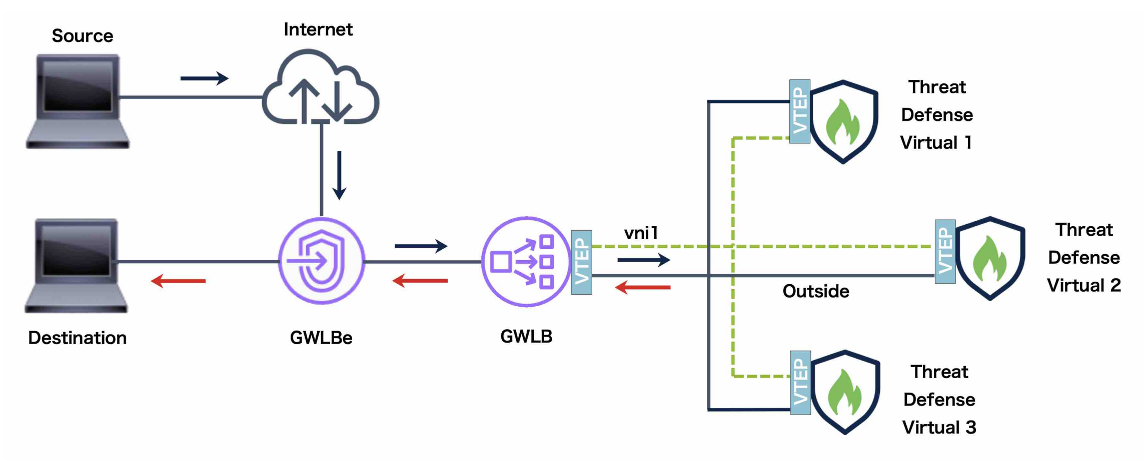

Sample Topology

The topology given below depicts both inbound and outbound traffic flow. There are three Threat Defense Virtual instances in the cluster that is connected to a GWLB. A Management Center Virtual instance is used to manage the cluster.

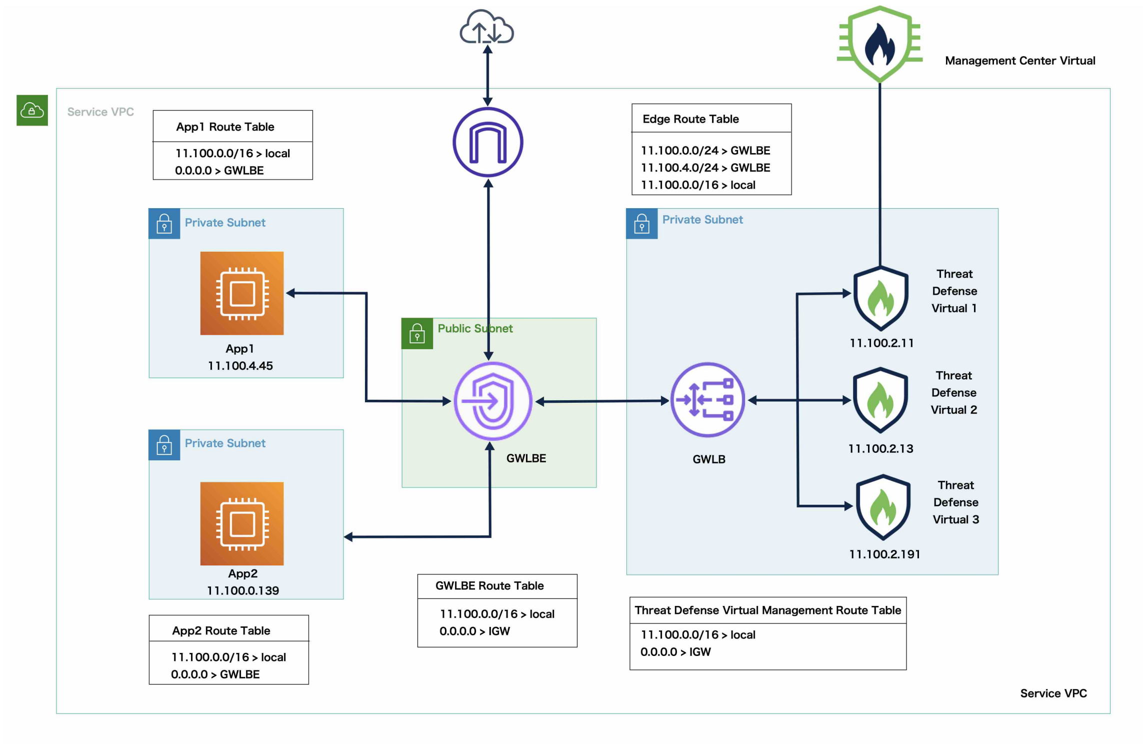

Inbound traffic from the internet goes to the GWLB endpoint which then transmits the traffic to the GWLB. Traffic is then forwarded to the Threat Defense Virtual cluster. After the traffic has been inspected by a Threat Defense Virtual instance in the cluster, it is forwarded to the application VM, App1 /App2.

Outbound traffic from App1/App2 is transmitted to the GWLB endpoint which then sends it out to the internet.

End-to-End Process for Deploying Threat Defense Virtual Cluster on AWS

Template-based Deployment

The following flowchart illustrates the workflow for template-based deployment of the Threat Defense Virtual cluster on AWS.

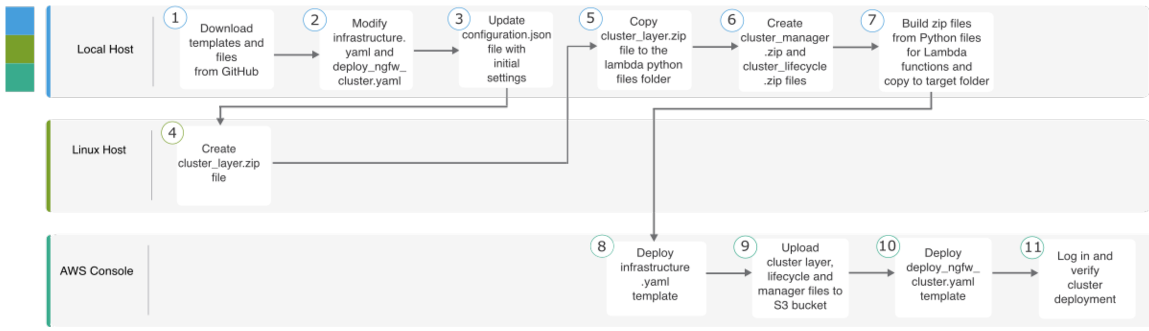

	Workspace	Steps
1	Local Host	Download templates and files from GitHub.
2	Local Host	Modify *infrastructure.yaml* and *deploy_ngfw_cluster.yaml* templates.
3	Local Host	Update the *Configuration.json* file with initial settings.
4	Linux Host	Create *cluster_layer.zip* file.
5	Local Host	Copy *cluster_layer.zip* file to the Lambda python files folder.
6	Local Host	Create *cluster_manager.zip* and *cluster_lifecycle.zip* files.
7	Local Host	Build zip files from Python files for Lambda functions and copy to target folder.
8	AWS Console	Deploy *infrastructure.yaml* template.
9	AWS Console	Upload cluster_layer.zip, cluster_lifecycle.zip, and cluster_manager.zip, to the S3 bucket.
10	AWS Console	Deploy *deploy_ngfw_cluster.yaml* template.
11	AWS Console	Log in and verify cluster deployment.

Manual Deployment

The following flowchart illustrates the workflow for manual deployment of the Threat Defense Virtual cluster on AWS.

	Workspace	Steps
1	Local Host	Create the Day0 Configuration for AWS
2	AWS Console	Deploy Threat Defense Virtual instance.
3	AWS Console	Attach interfaces to instance.
4	AWS Console	Verify if nodes have joined cluster.
5	AWS Console	Create target group and GWLB; attach target group to the GWLB.
6	AWS Console	Register instances with the target group using data interface IP.
7	Management Center	Register control node.

Templates

The templates given below are available in GitHub. The parameter values are self-explanatory with the parameter names, default values, allowed values, and description, given in the template.

- infrastructure.yaml – Template for infrastructure deployment.
- deploy_ngfw_cluster.yaml – Template for cluster deployment.

Note Ensure that you check the list of supported AWS instance types before deploying cluster nodes. This list is found in the *deploy_ngfw_cluster.yaml* template, under allowed values for the parameter InstanceType.

Deploy the Stack in AWS Using a CloudFormation Template

Deploy the stack in AWS using the customized CloudFormation template.

Before you begin

- You need a Linux computer with Python 3.
- To allow the cluster to auto-register with the management center, you need to create a user with administrative privileges on the management center that can use the REST API. See the Cisco Secure Firewall Management Center Administration Guide.
- Add an access policy in the management center that matches the name of the policy that you specified in Configuration.JSON.

Procedure

Step 1 Prepare the template.

a) Clone the github repository to your local folder. See https://github.com/CiscoDevNet/cisco-ftdv/tree/master/cluster/aws.

b) Modify **infrastructure.yaml** and **deploy_ngfw_cluster.yaml** with the required parameters.

c) Modify **cloud-clustering/ftdv-cluster/lambda-python-files/Configuration.json** with initial settings.

For example:

```
{
  "licenseCaps": ["BASE", "MALWARE", "THREAT"],
  "performanceTier": "FTDv50",
  "fmcIpforDeviceReg": "DONTRESOLVE",
  "RegistrationId": "cisco",
  "NatId": "cisco",
  "fmcAccessPolicyName": "AWS-ACL"
}
```

- Keep the fmcIpforDeviceReg setting as DONTRESOLVE.
- The fmcAccessPolicyName needs to match an access policy on the management center.

Note FTDv5 and FTDv10 tiers are not supported.

d) Create a file named **cluster_layer.zip** to provide essential Python libraries to Lambda functions.

We recommend to use the Amazon Linux with Python 3.9 installed to create the **cluster_layer.zip** file.

Note If you need an Amazon Linux environment, you can create an EC2 instance using Amazon Linux 2023 AMI or use AWS Cloudshell, which runs the latest version of Amazon Linux.

For creating the cluster-layer.zip file, you need to first create **requirements.txt** file that consists of the python library package details and then run the shell script.

1. Create the **requirements.txt** file by specifying the python package details.

   The following is the sample package details that you provide in the **requirements.txt** file:

```
$ cat requirements.txt
pycryptodome
paramiko
requests
scp
jsonschema
cffi
zipp
importlib-metadata
```

2. Run the following shell script to create **cluster_layer.zip** file.

```
$ pip3 install --platform manylinux2014_x86_64
--target=./python/lib/python3.9/site-packages
--implementation cp --python-version 3.9 --only-binary=:all:
--upgrade -r requirements.txt
$ zip -r cluster_layer.zip ./python
```

Note If you encounter a dependency conflict error during installation, such as urllib3 or cryptography, it is recommended that you include the conflicting packages along with their recommended versions in the **requirements.txt** file. After that, you can run the installation again to resolve the conflict.

e) Copy the resulting **cluster_layer.zip** file to the lambda python files folder.

f) Create the **cluster_manager.zip** and **cluster_lifecycle.zip** files.

A **make.py** file can be found in the cloned repository. This will zip the python files into a Zip file and copy to a target folder.

python3 make.py build

Step 2 Deploy **infrastructure.yaml** and note the output values for cluster deployment.

a) On the AWS Console, go to **CloudFormation** and click **Create stack**; select **With new resources(standard)**.

b) Select **Upload a template file**, click **Choose file**, and select **infrastructure.yaml** from the target folder.

c) Click **Next** and provide the required information.

d) Click **Next**, then **Create stack**.

e) After the deployment is complete, go to **Outputs** and note the S3 **BucketName**.

Figure 2: Output of infrastructure.yaml

Outputs (16)

Search outputs

Key	Value	Description	Export name
AZ	me-south-1a	Availability zone	-
AppInstanceSGId	sg-02b07af19c3e746d9	Security Group ID for Application Instances	-
ApplicationSubnetIds	subnet-03217efc6049e5fee	Application subnet ID	-
BucketName	neo-cls-infra-s3bucketcluster-13pgzkjjrx66	Name of the sample Amazon S3 bucket with Private Static Web hosting Configuration	-
BucketUrl	http://neo-cls-infra-s3bucketcluster-13pgzkjjrx66.s3-website.me-south-1.amazonaws.com	URL of S3 Bucket Static Website	-
CCLSubnetId	subnet-0caf6c4801922d8b1	CCL subnet ID	-
EIPforNATgw	15.184.208.231	EIP reserved for NAT GW	-
FmcInstanceSGID	sg-0a0d3797b04370aa3	Security Group ID for FMC if user would like to launch in this VPC itself	-
InInterfaceSGId	sg-0522ebe5acb8a2827	Security Group ID for Instances Inside Interface	-
InsideSubnetIds	subnet-056fdc9fe5389bf88	Inside subnet ID	-
InstanceSGId	sg-0be5b62647eb53dec	Security Group ID for Instances Management Interface	-
LambdaSecurityGroupId	sg-0347d191d724b2574	Security Group ID for Lambda Functions	-
LambdaSubnetIds	subnet-0989fbaeb522a906c,subnet-0c7a9b649d506f930	List of lambda subnet IDs (comma seperated)	-
MgmtSubnetIds	subnet-08c386d4b06890532	Mangement subnet ID	-
UseGWLB	Yes	Use Gateway Load Balancer	-
VpcName	vpc-0d94d3eaaa1f1354d	Name of the VPC created	-

Step 3 Upload **cluster_layer.zip**, **cluster_lifecycle.zip**, and **cluster_manager.zip** to the S3 bucket created by infrastructure.yaml.

Figure 3: S3 Bucket

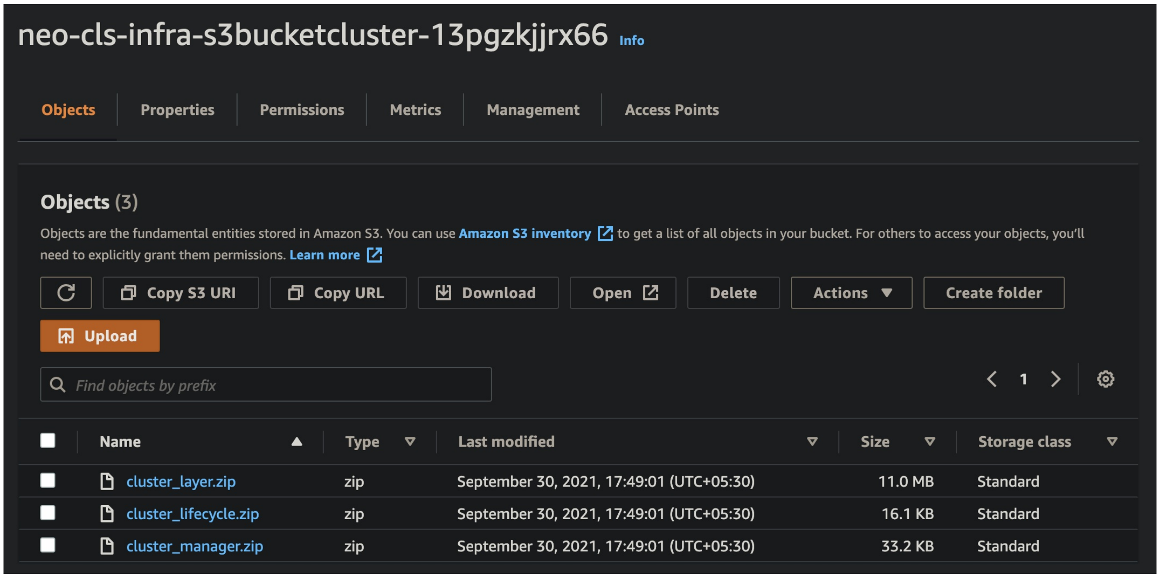

Step 4 Deploy **deploy_ngfw_cluster.yaml**.

a) Go to **CloudFormation** and click on **Create stack**; select **With new resources(standard)**.

b) Select **Upload a template file**, click **Choose file**, and select **deploy_ngfw_cluster.yaml** from the target folder.

c) Click **Next** and provide the required information.

d) Click **Next**, then **Create stack**.

The Lambda functions manage the rest of the process, and the threat defense virtuals will automatically register with the management center.

Figure 4: Deployed Resources

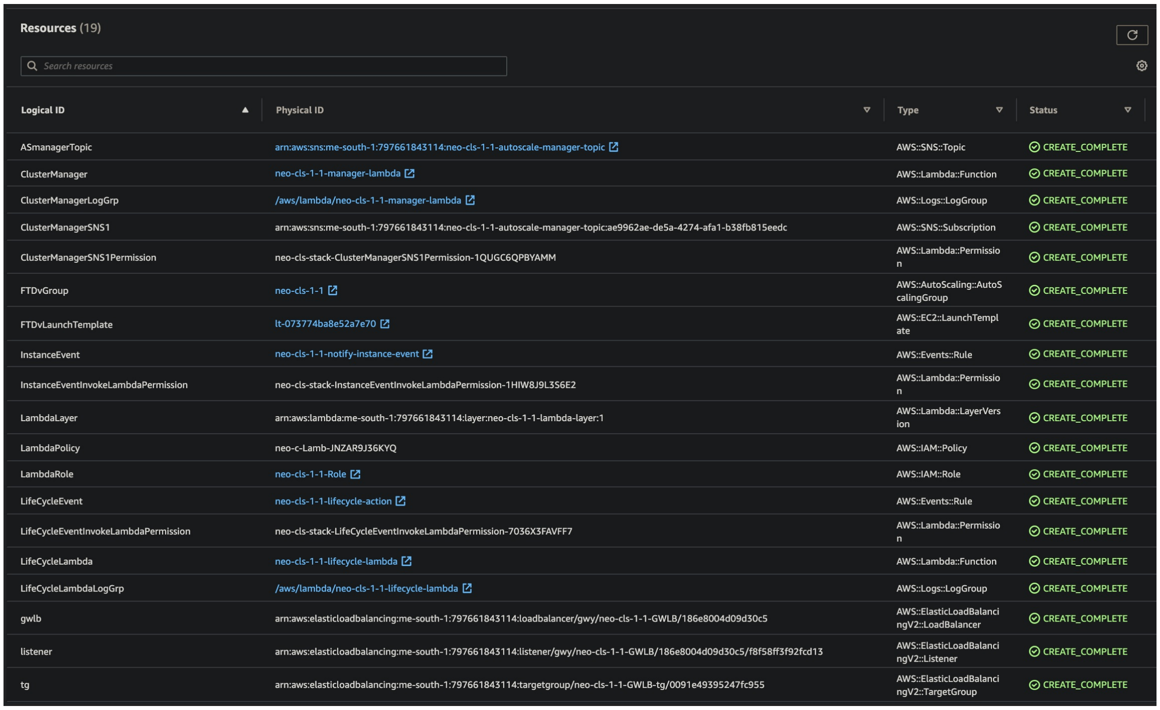

Step 5 Verify the cluster deployment by logging into any one of the nodes and using the **show cluster info** command.

Figure 5: Cluster Nodes

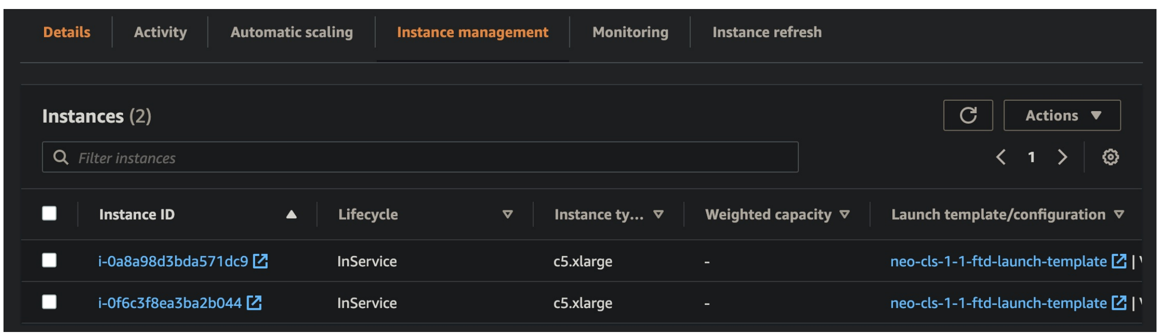

Figure 6: show cluster info

Deploy the Cluster in AWS Manually

To deploy the cluster manually, prepare the day 0 configuration, deploy each node, and then add the control node to the management center.

Create the Day0 Configuration for AWS

You can use either a fixed configuration or a customized configuration. We recommend using the fixed configuration.

Create the Day0 Configuration With a Fixed Configuration for AWS

The fixed configuration will auto-generate the cluster bootstrap configuration.

```
{
    "AdminPassword": "password",
    "Hostname": "hostname",
    "FirewallMode": "Routed",
    "ManageLocally": "No",
    "Cluster": {
        "CclSubnetRange": "ip_address_start ip_address_end",
        "ClusterGroupName": "cluster_name",
```

```
        [For Gateway Load Balancer] "Geneve": "{Yes | No}",
        [For Gateway Load Balancer] "HealthProbePort": "port"
    }
}
```

For example:

```
{
 "AdminPassword": "Sup3rnatural",
 "Hostname": "ciscoftdv",
 "FirewallMode": "Routed",
 "ManageLocally": "No",
 "Cluster": {
  "CclSubnetRange": "10.10.55.4 10.10.55.30",    //mandatory user input
  "ClusterGroupName": "ftdv-cluster",            //mandatory user input
  "Geneve": "Yes",
  "HealthProbePort": "7777"
 }
}
```

Note If you are copying and pasting the configuration given above, ensure that you remove **//mandatory user input** from the configuration.

For the **CclSubnetRange** variable, specify a range of IP addresses starting from x.x.x.4. Ensure that you have at least 16 available IP addresses for clustering. Some examples of start (*ip_address_start*) and end (*ip_address_end*) IP addresses given below.

Table 2: Examples of Start and End IP addresses

CIDR	Start IP Address	End IP Address
10.1.1.0/27	10.1.1.4	10.1.1.30
10.1.1.32/27	10.1.1.36	10.1.1.62
10.1.1.64/27	10.1.1.68	10.1.1.94
10.1.1.96/27	10.1.1.100	10.1.1.126
10.1.1.128/27	10.1.1.132	10.1.1.158
10.1.1.160/27	10.1.1.164	10.1.1.190
10.1.1.192/27	10.1.1.196	10.1.1.222
10.1.1.224/27	10.1.1.228	10.1.1.254
10.1.1.0/24	10.1.1.4	10.1.1.254

Create the Day0 Configuration With a Customized Configuration for AWS

You can enter the entire cluster bootstrap configuration using commands.

```
{
    "AdminPassword": "password",
```

```
"Hostname": "hostname",
"FirewallMode": "Routed",
"ManageLocally": "No",
"run_config": [comma_separated_threat_defense_configuration]
}
```

Gateway Load Balancer Example

The following example creates a configuration for a Gateway Load Balancer with one Geneve interface for U-turn traffic and one VXLAN interface for the cluster control link. Note the values in bold that need to be unique per node.

A sample day 0 configuration for **version 7.4 and later** is given below.

```
{
 "AdminPassword": "Sam&Dean",
 "Hostname": "ftdv1",
 "FirewallMode": "Routed",
 "ManageLocally": "No",
 "run_config": [
  "cluster interface-mode individual force",
  "interface TenGigabitEthernet0/0",
  "nameif geneve-vtep-ifc",
  "ip address dhcp",
  "no shutdown",
  "interface TenGigabitEthernet0/1",
  "nve-only cluster",
  "nameif ccl_link",
  "ip address dhcp",
  "no shutdown",
  "interface vni1",
  "description Clustering Interface",
  "segment-id 1",
  "vtep-nve 1",
  "interface vni2",
  "proxy single-arm",
  "nameif uturn-ifc",
  "vtep-nve 2",
  "object network ccl#link",
  "range 10.1.90.4 10.1.90.19",
  "object-group network cluster#group",
  "network-object object ccl#link",
  "nve 2",
  "encapsulation geneve",
  "source-interface geneve-vtep-ifc",
  "nve 1",
  "encapsulation vxlan",
  "source-interface ccl_link",
  "peer-group cluster#group",
  "jumbo-frame reservation",
  "mtu geneve-vtep-ifc 1826",
  "mtu ccl_link 1980",
  "cluster group ftdv-cluster",
  "local-unit 1",
  "cluster-interface vni1 ip 10.1.1.1 255.255.255.0",
  "priority 1",
  "enable",
  "aaa authentication listener http geneve-vtep-ifc port 7777"
 ]
}
```

A sample day 0 configuration for **version 7.3 and earlier** is given below.

```
{
 "AdminPassword": "Sam&Dean",
 "Hostname": "ftdv1",
 "FirewallMode": "Routed",
 "ManageLocally": "No",
 "run_config": [
  "cluster interface-mode individual force",
  "interface TenGigabitEthernet0/0",
  "nameif geneve-vtep-ifc",
  "ip address dhcp",
  "no shutdown",
  "interface TenGigabitEthernet0/1",
  "nve-only cluster",
  "nameif ccl_link",
  "ip address dhcp",
  "no shutdown",
  "interface vni1",
  "description Clustering Interface",
  "segment-id 1",
  "vtep-nve 1",
  "interface vni2",
  "proxy single-arm",
  "nameif uturn-ifc",
  "vtep-nve 2",
  "object network ccl#link",
  "range 10.1.90.4 10.1.90.19",
  "object-group network cluster#group",
  "network-object object ccl#link",
  "nve 2",
  "encapsulation geneve",
  "source-interface geneve-vtep-ifc",
  "nve 1",
  "encapsulation vxlan",
  "source-interface ccl_link",
  "peer-group cluster#group",
  "jumbo-frame reservation",
  "mtu geneve-vtep-ifc 1806",
  "mtu ccl_link 1960",
  "cluster group ftdv-cluster",
  "local-unit 1",
  "cluster-interface vni1 ip 10.1.1.1 255.255.255.0",
  "priority 1",
  "enable",
  "aaa authentication listener http geneve-vtep-ifc port 7777"
 ]
}
```

Note For the CCL subnet range, specify IP addresses from the CCL subnet CIDR, excluding reserved IP addresses. Refer the Table 2: Examples of Start and End IP addresses given above for some examples.

For the AWS health check settings, ensure that you specify the **aaa authentication listener http** port you set here.

Non-Native Load Balancer Example

The following example creates a configuration for use with non-native load balancers with Management, Inside, and Outside interfaces, and a VXLAN interface for the cluster control link. Note the values in bold that need to be unique per node.

```
{
 "AdminPassword": "W1nch3sterBr0s",
 "Hostname": "ftdv1",
 "FirewallMode": "Routed",
 "ManageLocally": "No",
 "run_config": [
  "cluster interface-mode individual force",
  "interface Management0/0",
  "management-only",
  "nameif management",
  "ip address dhcp",
  "interface GigabitEthernet0/0",
  "no shutdown",
  "nameif outside",
  "ip address dhcp",
  "interface GigabitEthernet0/1",
  "no shutdown",
  "nameif inside",
  "ip address dhcp",
  "interface GigabitEthernet0/2",
  "nve-only cluster",
  "nameif ccl_link",
  "ip address dhcp",
  "no shutdown",
  "interface vni1",
  "description Clustering Interface",
  "segment-id 1",
  "vtep-nve 1",
  "jumbo-frame reservation",
  "mtu ccl_link 1654",
  "object network ccl#link",
  "range 10.1.90.4 10.1.90.19",          //mandatory user input
  "object-group network cluster#group",
  "network-object object ccl#link",
  "nve 1",
  "encapsulation vxlan",
  "source-interface ccl_link",
  "peer-group cluster#group",
  "cluster group ftdv-cluster",          //mandatory user input
  "local-unit 1",
  "cluster-interface vni1 ip 10.1.1.1 255.255.255.0",
  "priority 1",
  "enable"
 ]
}
```

For the cluster control link network object, specify only as many addresses as you need (up to 16). A larger range can affect performance.

Note If you are copying and pasting the configuration given above, ensure that you remove **//mandatory user input** from the configuration.

Deploy Cluster Nodes

Deploy the cluster nodes so they form a cluster.

Procedure

Step 1 Deploy the Threat Defense Virtual instance by using the cluster day 0 configuration with the required number of interfaces - four interfaces if you are using Gateway Load Balancer (GWLB), or five interfaces if you are using non-native load balancer. To do this, in the **Configure Instance Details** > **Advanced Details** section, paste the cluster day 0 configuration.

Note Ensure that you attach interfaces to the instances in the order given below.

- AWS Gateway Load Balancer - four interfaces - management, diagnostic, inside, and cluster control link.
- Non-native load balancers - five interfaces - management, diagnostic, inside, outside, and cluster control link.

For more information on deploying Threat Defense Virtual on AWS, see Deploy the Threat Defense Virtual on AWS.

Step 2 Repeat Step 1 to deploy the required number of additional nodes.

Step 3 Use the **show cluster info** command on the Threat Defense Virtual console to verify if all nodes have successfully joined the cluster.

Step 4 Configure the AWS Gateway Load Balancer.

a) Create a target group and GWLB.

b) Attach the target group to the GWLB.

Note Ensure that you configure the GWLB to use the correct security group, listener configuration, and health check settings.

c) Register the data interface (inside interface) with the Target Group using IP addresses.

For more information, see Create a Gateway Load Balancer.

Step 5 Add the control node to the Management Center. See Add the Cluster to the Management Center (Manual Deployment), on page 52.

Configure Target Failover for Secure Firewall Threat Defense Virtual Clustering with GWLB in AWS

Threat Defense Virtual clustering in AWS utilizes the Gateway Load Balancer (GWLB) to balance and forward network packets for inspection to a designated Threat Defense Virtual node. The GWLB is designed to continue sending network packets to the target node in the event of a failover or deregistration of that node.

The Target Failover feature in AWS enables GWLB to redirect network packets to a healthy target node in the event of node deregistration during planned maintenance or a target node failure. It takes advantage of the cluster's stateful failover.

In AWS, you can configure Target Failover through the AWS Elastic Load Balancing (ELB) API or AWS console.

> **Note** If a target node fails while the GWLB routes traffic using certain protocols such as SSH, SCP, CURL, and so on, then there may be a delay in redirecting traffic to a healthy target. This delay is due to rebalancing and rerouting of traffic flow.

In AWS, you can configure Target Failover through the AWS ELB API or AWS console.

- AWS API - In the AWS ELB API - *modify-target-group-attributes* you can define the flow handling behavior by modifying the following two new parameters.
  - target_failover.on_unhealthy - It defines how the GWLB handles the network flow when the target becomes unhealthy.
  - target_failover.on_deregistration - It defines how the GWLB handles the network flow when the target is deregistered.

  The following command shows the sample API parameter configuration of defining these two parameters.

  ```
  aws elbv2 modify-target-group-attributes \
  --target-group-arn arn:aws:elasticloadbalancing:…/my-targets/73e2d6bc24d8a067 \
  --attributes \
  Key=target_failover.on_unhealthy, Value=rebalance[no_rebalance] \
  Key=target_failover.on_deregistration, Value=rebalance[no_rebalance]
  ```

  For more information, refer TargetGroupAttribute in the AWS documentation.

- AWS Console – In the EC2 console, you can enable the Target Failover option on the Target Group page by configuring the following options.
  - Edit Target Groups Attributes
  - Enable Target Failover
  - Verify Rebalance Flows

  For more information about how to enable Target Failover, see Enable Target Failover for Secure Firewall Threat Defense Virtual Clustering in AWS, on page 23.

Enable Target Failover for Secure Firewall Threat Defense Virtual Clustering in AWS

The data interface of threat defense virtual is registered to a target group of GWLB in AWS. In the threat defense virtual clustering, each instance is associated with a Target Group. The GWLB load balances and sends the traffic to this healthy instance identified or registered as a target node in the target group.

Before you begin

You must have deployed the cluster in AWS either by manual method or using CloudFormation templates.

If you are deploying a cluster using a CloudFormation template, you can also enable the **Target Failover** parameter by assigning the **rebalance** attribute that is available under **GWLB Configuration** section of the cluster deployment file, `deploy_ftdv_clustering.yaml`. In the template, by default, the value is set to **rebalance** for this parameter. However, the default value for this parameter is set to **no_rebalance** on the AWS console.

Where,

- **no_rebalance** - GWLB continues to send the network flow to the failed or deregistered target.
- **rebalance** - GWLB sends the network flow to another healthy target when the existing target is failed or deregistered.

For information on deploying stack in AWS, see:

- Deploy the Cluster in AWS Manually
- Deploy the Stack in AWS Using a CloudFormation Template

Procedure

Step 1 On the AWS Console, go to **Services** > **EC2**

Step 2 Click **Target Groups** to view the target groups page.

Step 3 Select the target group to which the threat defense virtual data interface IPs are registered. The target group details page is displayed, where you can enable the Target failover attributes.

Step 4 Go to the **Attributes** menu.

Step 5 Click **Edit** to edit the attributes.

Step 6 Toggle the **Rebalance flows** slider button to the right to enable target failover to configure GWLB to rebalance and forward the existing network packets to a healthy target node in the event of target failover or deregistration.

Deploy the Cluster in Azure

You can use the cluster with the Azure Gateway Load Balancer (GWLB), or with a non-native load-balancer. To deploy a cluster in Azure, use Azure Resource Manager (ARM) templates to deploy a Virtual Machine Scale Set.

Sample Topology for GWLB-based Cluster Deployment

Figure 7: Inbound Traffic Use Case and Topology with GWLB

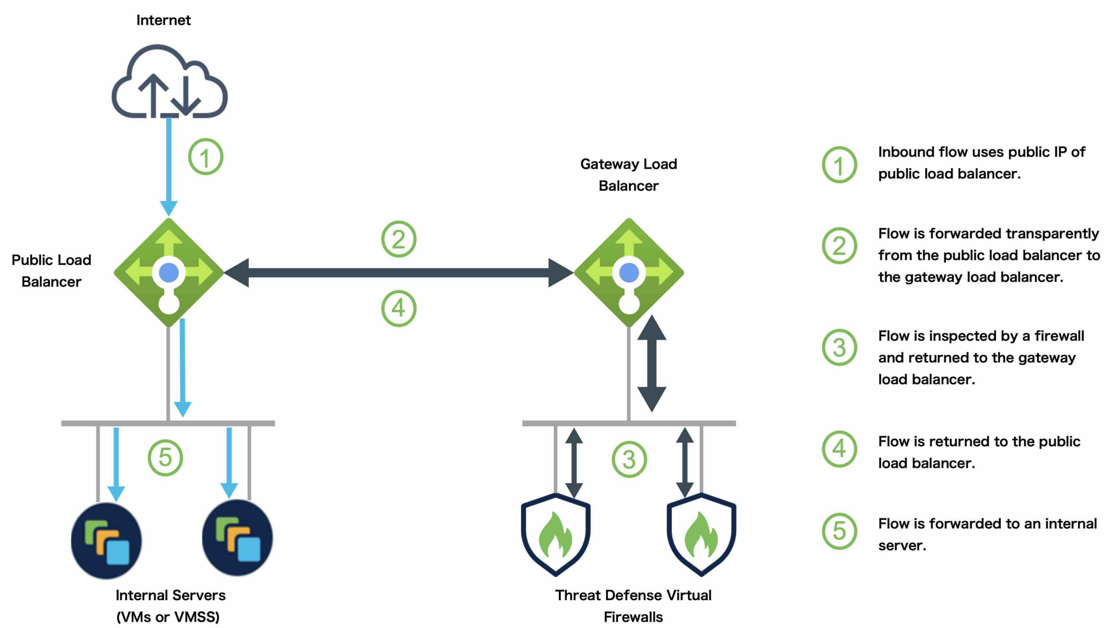

Figure 8: Outbound Traffic Use Case and Topology with GWLB

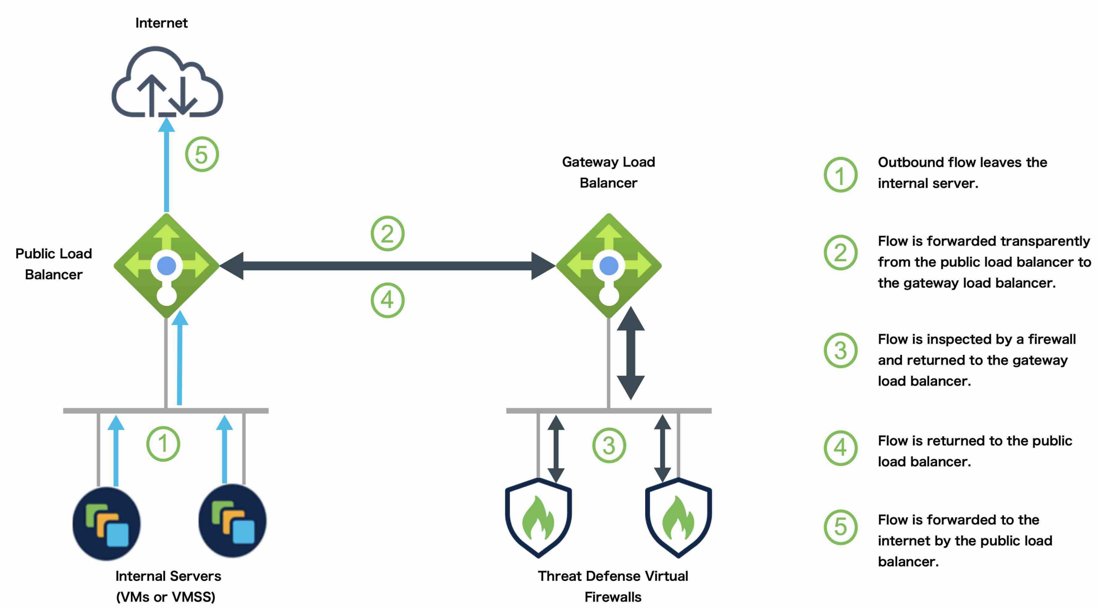

Azure Gateway Load Balancer and Paired Proxy

In an Azure service chain, Threat Defense Virtuals act as a transparent gateway that can intercept packets between the internet and the customer service. The Threat Defense Virtual defines an external interface and an internal interface on a single NIC by utilizing VXLAN segments in a paired proxy.

The following figure shows traffic forwarded to the Azure Gateway Load Balancer from the Public Gateway Load Balancer on the external VXLAN segment. The Gateway Load Balancer balances traffic among multiple Threat Defense Virtuals, which inspect the traffic before either dropping it or sending it back to the Gateway Load Balancer on the internal VXLAN segment. The Azure Gateway Load Balancer then sends the traffic back to the Public Gateway Load Balancer and to the destination.

Figure 9: Azure Gateway Load Balancer with Paired Proxy

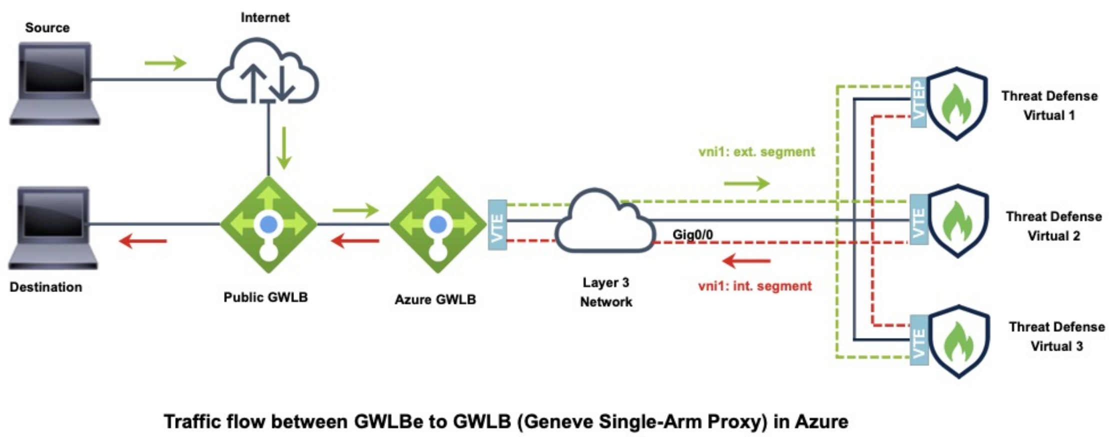

End-to-End Process for Deploying Threat Defense Virtual Cluster in Azure with GWLB

Template-based Deployment

The following flowchart illustrates the workflow for template-based deployment of the Threat Defense Virtual cluster in Azure with GWLB.

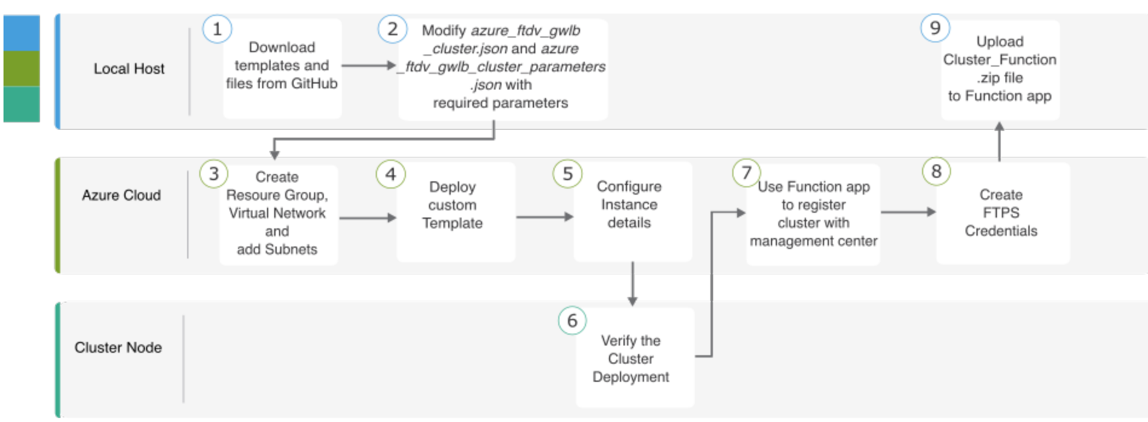

	Workspace	Steps
1	Local Host	Download templates and files from GitHub.
2	Local Host	Modify *azure_ftdv_gwlb_cluster.json* and *azure_ftdv_gwlb_cluster_parameters.json* with the required parameters.
3	Azure Cloud	Create the resource group, virtual network, and subnets.
4	Azure Cloud	Deploy custom template.
5	Azure Cloud	Configure instance details.
6	Cluster Node	Verify cluster deployment.
7	Azure Cloud	Use the Function app to register the cluster with the Management Center.
8	Azure Cloud	Create FTPS credentials.
9	Local Host	Upload *Cluster_Function.zip* file to the Function app.

Manual Deployment

The following flowchart illustrates the workflow of manual deployment of Threat Defense Virtual cluster in Azure with GWLB.

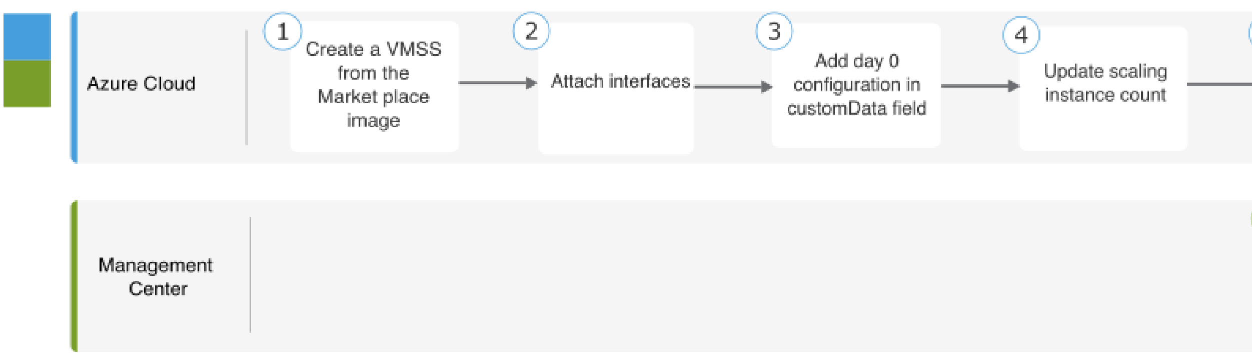

	Workspace	Steps
1	Local Host	Create a VMSS from the Marketplace image.
2	Local Host	Attach interfaces.
3	Local Host	Add day 0 configuration in the customData field.

	Workspace	Steps
4	Local Host	Update scaling instance count.
5	Local Host	Configure GWLB.
6	Management Center	Add control node.

Templates

The templates given below are available in GitHub. The parameter values are self-explanatory with the parameter names, and values, given in the template.

From Secure Firewall version 7.4.1, you can deploy the cluster without the diagnostic interface. To deploy the cluster with only the Outside, Inside, Management, and CCL interfaces, use the withoutDiagnostic templates - azure_withoutDiagnostic_ftdv_gwlb_cluster_parameters.json and azure_withoutDiagnostic_ftdv_gwlb_cluster.json files.

Templates to deploy **with Diagnostic** Interface:

- azure_ftdv_gwlb_cluster_parameters.json – Template to enter parameters for the Threat Defense Virtual cluster with GWLB.
- azure_ftdv_gwlb_cluster.json – Template to deploy Threat Defense Virtual cluster with GWLB.

Templates to deploy **without Diagnostic** Interface:

- azure_withoutDiagnostic_ftdv_gwlb_cluster_parameters.json – Template to enter parameters for the Threat Defense Virtual cluster with GWLB deployment without Diagnostic interface.
- azure_withoutDiagnostic_ftdv_gwlb_cluster.json – Template to deploy Threat Defense Virtual cluster with GWLB without Diagnostic interface.

Prerequisites

- To allow the cluster to auto-register to the management center, create a user with Network Admin & Maintenance User privileges on the management center. Users with these privileges can use REST API. See the Cisco Secure Firewall Management Center Administration Guide.
- Add an access policy in the management center that matches the name of the policy that you will specify during template deployment.
- Ensure that the Management Center Virtual is licensed appropriately.
- Perform the steps given below after the cluster is added to the Management Center Virtual:
  1. Configure platform settings with the health check port number in the Management Center. For more information on configuring this, see Platform Settings.
  2. Create a static route for data traffic. For more information on creating a static route, see Add a Static Route.

     Sample static route configuration:

```
Network: any-ipv4
Interface: vxlan_tunnel
Leaked from Virtual Router: Global
Gateway: vxlan_tunnel_gw
Tunneled: false
Metric: 2
```

Note *vxlan_tunnel_gw* is the data subnet's gateway IP address.

Deploy Cluster on Azure with GWLB Using an Azure Resource Manager Template

Deploy the Virtual Machine Scale Set for Azure GWLB using the customized Azure Resource Manager (ARM) template.

Procedure

Step 1 Prepare the template.

a) Clone the github repository to your local folder. See https://github.com/CiscoDevNet/cisco-ftdv/tree/master/cluster/azure.

b) Modify *azure_ftdv_gwlb_cluster.json* and *azure_ftdv_gwlb_cluster_parameters.json* with the required parameters.

OR

Modify withoutDiagnostic templates, *azure_withoutDiagnostic_ftdv_gwlb_cluster_parameters.json* and *azure_withoutDiagnostic_ftdv_gwlb_cluster.json*, with the required parameter for deploying cluster without the diagnostic interface.

Step 2 Log into the Azure Portal: https://portal.azure.com.

Step 3 Create a Resource Group.

a) In the **Basics** tab, choose the **Subscription** and **Resource Group** from the drop-down lists.

b) Choose the required **Region**.

Step 4 Create a virtual network with 4 subnets: Management, Diagnostic, Outside, and Cluster Control Link (CCL).

From Secure Firewall version 7.4.1, you can deploy the cluster without the diagnostic interface. To deploy the cluster with only the Outside, Inside, Management, and CCL interfaces, use the withoutDiagnostic templates - azure_withoutDiagnostic_ftdv_gwlb_cluster_parameters.json and azure_withoutDiagnostic_ftdv_gwlb_cluster.json files.

a) Create the virtual network.

1. In the **Basics** tab, choose the **Subscription** and **Resource Group** from the drop-down lists.

2. Choose the required **Region**. Click **Next: IP addresses**.

In the **IP Addresses** tab, click **Add subnet** and add the following subnets – Management, Diagnostic, Data, and Cluster Control Link.

If you are deploying the Threat Defense Virtual 7.4.1 cluster without a Diagnostic interface, then you must skip the Diagnostic subnet creation.

b) Add the subnets.

Step 5 Deploy the custom template.

a) Click **Create** > **Template deployment (deploy using custom templates)**.
b) Click **Build your own template in the editor**.
c) Click **Load File**, and upload **azure_ftdv_gwlb_cluster.json** or **azure_withoutDiagnostic_ftdv_gwlb_cluster.json**, if you have opted for without diagnostic interface deployment.
d) Click **Save**.

Step 6 Configure the Instance details.

a) Enter the required values and then click **Review + create**.
b) Click **Create** after the validation is passed.

Step 7 After the instance is running, verify the cluster deployment by logging into any one of the nodes and entering the **show cluster info** command.

Figure 10: show cluster info

```
> show cluster info
Cluster gwlb-cluster-template-with-AN: On
    Interface mode: individual
Cluster Member Limit : 16
    This is "12" in state CONTROL_NODE
        ID        : 0
        Version   : 99.19(1)180
        Serial No.: 9AKGFV8VH4G
        CCL IP    : 10.1.1.12
        CCL MAC   : 000d.3a55.5470
        Module    : NGFWv
        Resource  : 8 cores / 28160 MB RAM
        Last join : 11:13:24 UTC Sep 5 2022
        Last leave: N/A
```

Step 8 In the Azure Portal, click the Function app to register the cluster with the Management Center.

Note If you do not want to use the Function app, you can alternatively register the control node to the management center directly by using **Add** > **Device** (not **Add** > **Cluster**). The rest of the cluster nodes will register automatically.

Step 9 Create FTPS Credentials by clicking **Deployment Center** > **FTPS credentials** > **User scope** > **Configure Username and Password**, and then click **Save**.

Step 10 Upload the Cluster_Function.zip file to the Function app by executing the following **curl** command in the local terminal.

curl -X POST -u *username* **--data-binary @"Cluster_Function.zip" https://***Function_App_Name***.scm.azurewebsites.net/api/zipdeploy**

Note The **curl** command might take few minutes (~2 to 3 minutes) to complete command execution.

The function will be uploaded to the Function app. The function will start, and you can see the logs in the storage account's outqueue. The device registration with the Management Center will be initiated.

Figure 11: Functions

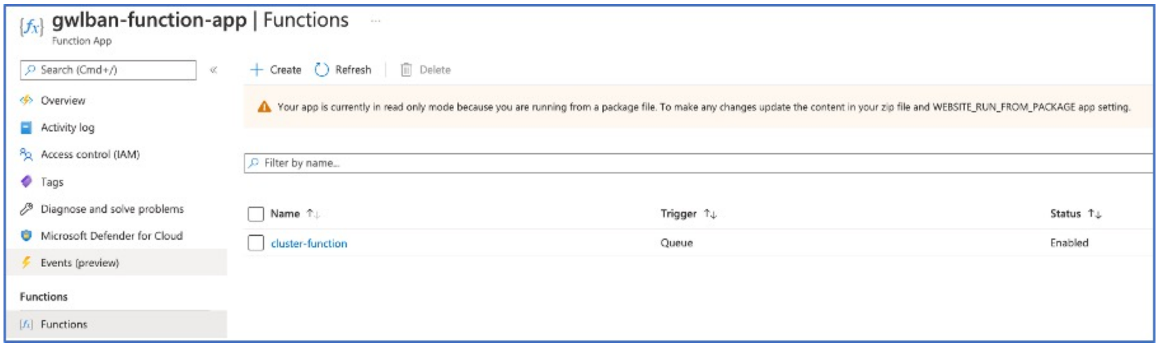

Figure 12: Queues

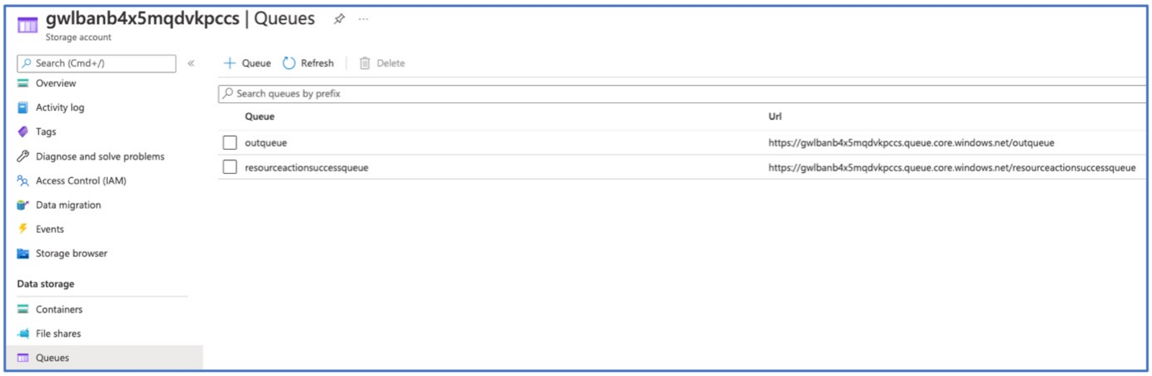

Figure 13: Outqueue

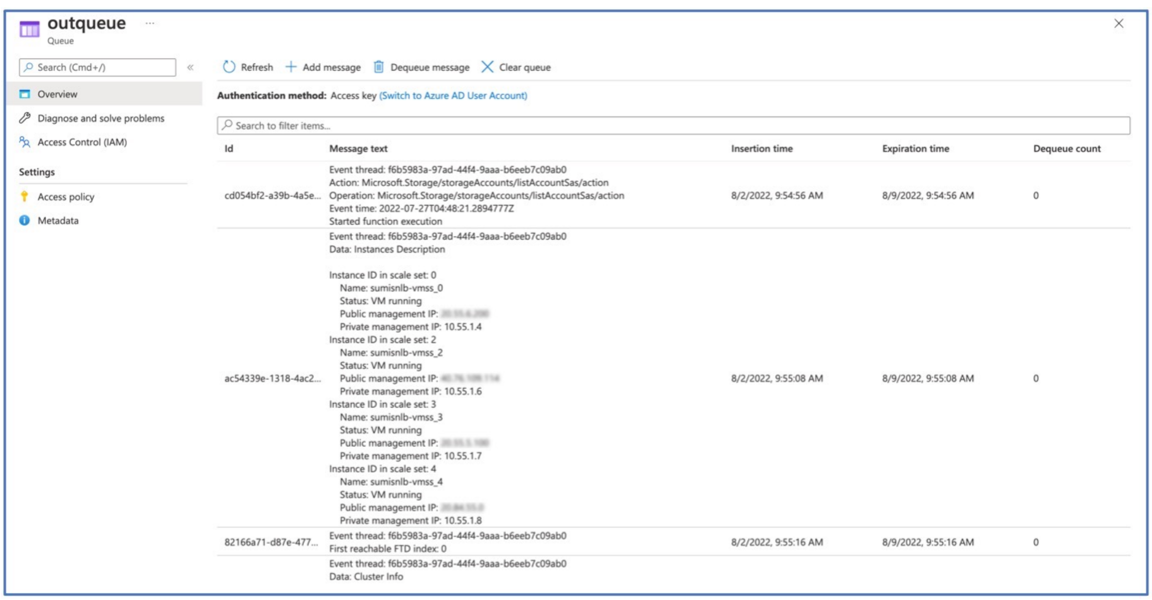

Sample Topology for NLB-based Cluster Deployment

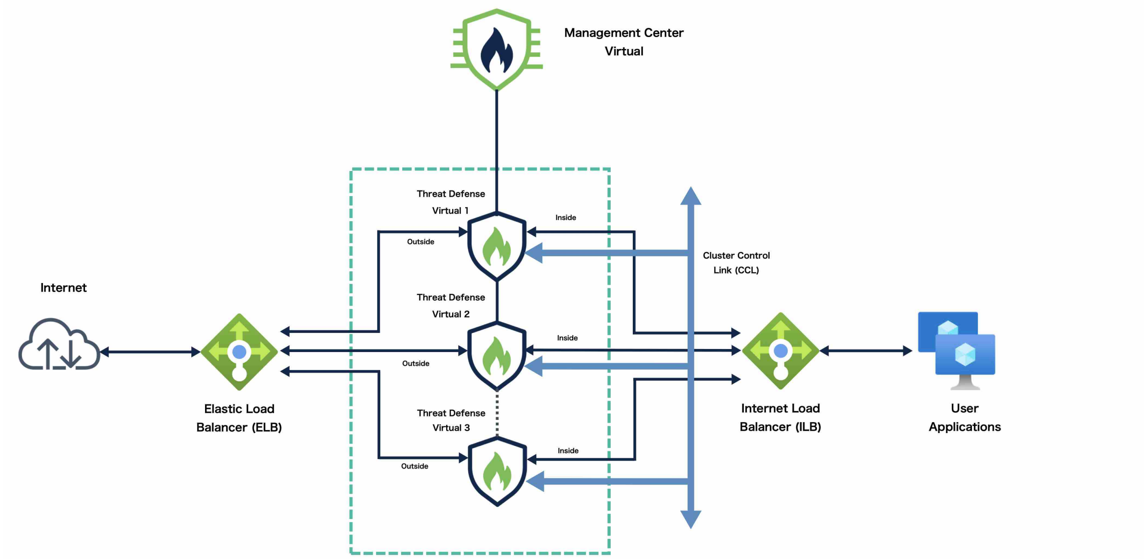

This topology depicts both inbound and outbound traffic flow. The Threat Defense Virtual cluster is sandwiched between the internal and external load balancers. A Management Center Virtual instance is used to manage the cluster.

Inbound traffic from the internet goes to the external load balancer which then transmits the traffic to the Threat Defense Virtual cluster. After the traffic has been inspected by a Threat Defense Virtual instance in the cluster, it is forwarded to the application VM.

Outbound traffic from the application VM is transmitted to the internal load balancer. Traffic is then forwarded to the Threat Defense Virtual cluster and then sent out to the internet.

End-to-End Process for Deploying Threat Defense Virtual Cluster in Azure with NLB

Template-based Deployment

The following flowchart illustrates the workflow of template-based deployment of Threat Defense Virtual cluster in Azure with NLB.

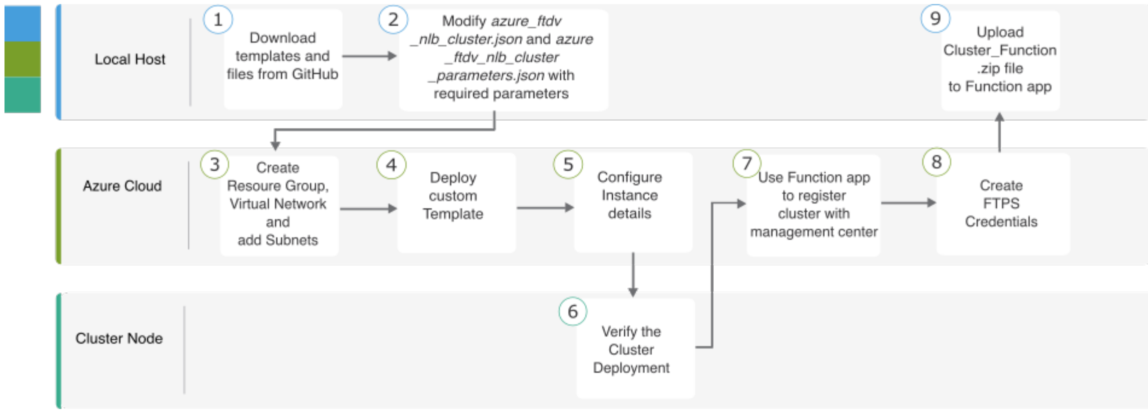

	Workspace	Steps
1	Local Host	Download templates and files from GitHub.
2	Local Host	Modify *azure_ftdv_nlb_cluster.json* and *azure_ftdv_nlb_cluster_parameters.json* with the required parameters.
3	Azure Cloud	Create the resource group, virtual network, and subnets.
4	Azure Cloud	Deploy custom template.
5	Azure Cloud	Configure instance details.
6	Cluster Node	Verify cluster deployment.
7	Azure Cloud	Use the Function app to register the cluster with the Management Center.
8	Azure Cloud	Create FTPS credentials.
9	Local Host	Upload *Cluster_Function.zip* file to the Function app.

Manual Deployment

The following flowchart illustrates the workflow of manual deployment of Threat Defense Virtual cluster in Azure with NLB.

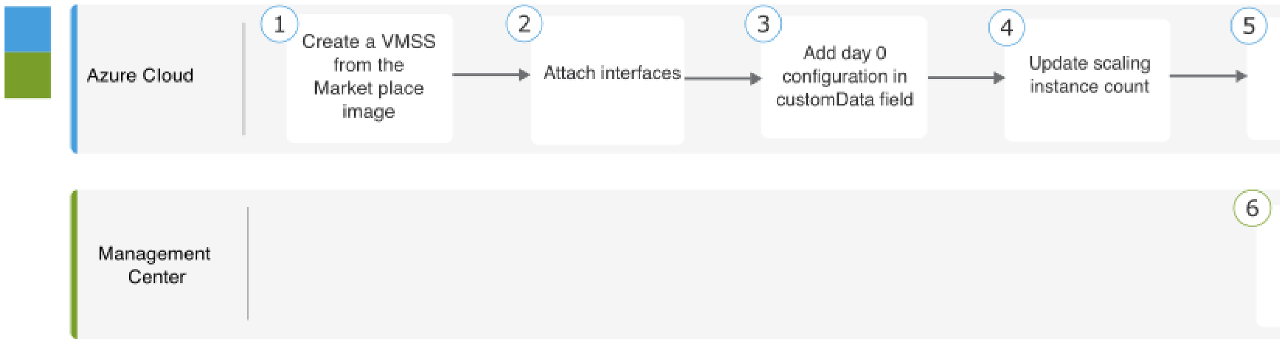

	Workspace	Steps
1	Local Host	Create a VMSS from the Marketplace image.
2	Local Host	Attach interfaces.
3	Local Host	Add day 0 configuration in the customData field.
4	Local Host	Update scaling instance count.
5	Local Host	Configure NLB.
6	Management Center	Add control node.

Templates

The templates given below are available in GitHub. The parameter values are self-explanatory with the parameter names, and values, given in the template.

From Secure Firewall version 7.4.1, you can deploy the cluster without the diagnostic interface. To deploy the cluster with only the Outside, Inside, Management, and CCL interfaces, use the withoutDiagnostic templates - azure_withoutDiagnostic_ftdv_nlb_cluster_parameters.json and azure_withoutDiagnostic_ftdv_nlb_cluster.json files.

Templates to deploy **with Diagnostic** Interface:

- azure_ftdv_nlb_cluster_parameters.json – Template to enter parameters for the Threat Defense Virtual cluster with NLB.
- azure_ftdv_nlb_cluster.json – Template to deploy Threat Defense Virtual cluster with NLB.

Templates to deploy **without Diagnostic** Interface:

- azure_withoutDiagnostic_ftdv_nlb_cluster_parameters.json – Template to enter parameters for the Threat Defense Virtual cluster with NLB deployment without Diagnostic interface.
- azure_withoutDiagnostic_ftdv_nlb_cluster.json – Template to deploy Threat Defense Virtual cluster with NLB without Diagnostic interface.

Prerequisites

- To allow the cluster to auto-register with the Management Center, create a user with Network Admin & Maintenance User privileges on the Management Center. Users with these privileges can use REST API. See the Cisco Secure Firewall Management Center Administration Guide.
- Add an access policy in the Management Center that matches the name of the policy that you will specify during template deployment.
- Ensure that the Management Center Virtual is licensed appropriately.
- After the cluster is added to the Management Center Virtual:
  1. Configure platform settings with the health check port number in the Management Center. For more information on configuring this, see Platform Settings.
  2. Create static routes for traffic from outside and inside interfaces. For more information on creating a static route, see Add a Static Route.

     Sample static route configuration for the outside interface:

     ```
     Network: any-ipv4
     Interface: outside
     Leaked from Virtual Router: Global
     Gateway: ftdv-cluster-outside
     Tunneled: false
     Metric: 10
     ```

     Note *ftdv-cluster-outside* is the outside subnet's gateway IP address.

     Sample static route configuration for the inside interface:

     ```
     Network: any-ipv4
     Interface: inside
     Leaked from Virtual Router: Global
     Gateway: ftdv-cluster-inside-gw
     Tunneled: false
     Metric: 11
     ```

     Note *ftdv-cluster-inside-gw* is the inside subnet's gateway IP address.

  3. Configure NAT rule for data traffic. For more information on configuring NAT rules, see Network Address Translation.

Deploy Cluster on Azure with NLB Using an Azure Resource Manager Template

Deploy the cluster for Azure NLB using the customized Azure Resource Manager (ARM) template.

Procedure

Step 1 Prepare the template.

a) Clone the github repository to your local folder. See https://github.com/CiscoDevNet/cisco-ftdv/tree/master/cluster/azure.

b) Modify *azure_ftdv_nlb_cluster.json* and *azure_ftdv_nlb_cluster_parameters.json* with the required parameters.

Modify withoutDiagnostic templates, *azure_withoutDiagnostic_ftdv_nlb_cluster_parameters.json* and *azure_withoutDiagnostic_ftdv_nlb_cluster.json*, with the required parameter for deploying cluster without the diagnostic interface.

Step 2 Log into the Azure Portal: https://portal.azure.com.

Step 3 Create a Resource Group.

a) In the **Basics** tab, choose the **Subscription** and **Resource Group** from the drop-down lists.

b) Choose the required **Region**.

Step 4 Create a virtual network with 5 subnets: Management, Diagnostic, Inside, Outside, and Cluster Control Link.

From Secure Firewall version 7.4.1, you can deploy the cluster without the diagnostic interface. To deploy the cluster with only the Outside, Inside, Management, and Cluster Control Link interfaces, use the withoutDiagnostic templates - azure_withoutDiagnostic_ftdv_nlb_cluster_parameters.json and azure_withoutDiagnostic_ftdv_nlb_cluster.json files.

a) Create the virtual network.

1. In the **Basics** tab, choose the **Subscription** and **Resource Group** from the drop-down lists.
2. b) Choose the required **Region**. Click **Next: IP addresses**.

b) Add the subnets.

In the **IP Addresses** tab, click **Add subnet** and add the following subnets – Management, Diagnostic, Inside, Outside, and Cluster Control Link.

If you are deploying the Threat Defense Virtual 7.4.1 cluster without a Diagnostic interface, then you must skip the Diagnostic subnet creation.

Step 5 Deploy the custom template.

a) Click **Create** > **Template deployment (deploy using custom templates)**.

b) Click **Build your own template in the editor**.

c) Click **Load File**, and upload **azure_ftdv_nlb_cluster.json** or **azure_withoutDiagnostic_ftdv_nlb_cluster.json**, if you have opted for without diagnostic interface deployment.

d) Click **Save**.

Step 6 Configure the instance details.

a) Enter the required values and then click **Review + create**.

Note For the cluster control link starting and ending addresses, specify only as many addresses as you need (up to 16). A larger range can affect performance.

b) Click **Create** after the validation is passed.

Step 7 After the instance is running, verify the cluster deployment by logging into any one of the nodes and using the **show cluster info** command.

Figure 14: show cluster info

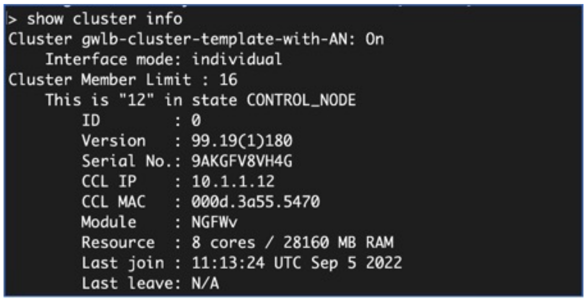

Step 8 In the Azure Portal, click the Function app to register the cluster to the management center.

Note If you do not want to use the Function app, you can alternatively register the control node with the Management Center directly by using **Add** > **Device** (not **Add** > **Cluster**). The rest of the cluster nodes will register automatically.

Step 9 Create FTPS Credentials by clicking **Deployment Center** > **FTPS credentials** > **User scope** > **Configure Username and Password**, and then click **Save**.

Step 10 Upload the Cluster_Function.zip file to the Function app by executing the following **curl** command in the local terminal.

curl -X POST -u *username* **--data-binary @"Cluster_Function.zip" https://** *Function_App_Name***.scm.azurewebsites.net/api/zipdeploy**

Note The **curl** command might take a few minutes (~2 to 3 minutes) to complete command execution.

The function will be uploaded to the Function app. The function will start, and you can see the logs in the storage account's outqueue. The device registration with the Management Center will be initiated.

Deploy the Cluster in Azure Manually

To deploy the cluster manually, prepare the day0 configuration, deploy each node, and then add the control node to the management center.

Create the Day0 Configuration for Azure

You can use either a fixed configuration or a customized configuration.

Create the Day0 Configuration With a Fixed Configuration for Azure

The fixed configuration will auto-generate the cluster bootstrap configuration.

```
{
 "AdminPassword": "password",
 "FirewallMode": "Routed",
 "ManageLocally": "No",
 "Diagnostic": "OFF",    //For deployment of version 7.4.1 and later without Diagnostics
```

```
template, set this parameter to OFF.
. "FmcIp": "<FMC_IP>",
 "FmcRegKey": "<REGISTRATION_KEY>",
 "FmcNatId": "<NAT_ID>",
 "Cluster": {
  "CclSubnetRange": "ip_address_start ip_address_end",
  "ClusterGroupName": "cluster_name",
  "HealthProbePort": "port_number",
  "GatewayLoadBalancerIP": "ip_address",
  "EncapsulationType": "vxlan",
  "InternalPort": "internal_port_number",
  "ExternalPort": "external_port_number",
  "InternalSegId": "internal_segment_id",
  "ExternalSegId": "external_segment_id"
 }
}
```

Example

A sample day 0 configuration is given below.

```
{
 "AdminPassword": "password",
 "FirewallMode": "routed",
 "ManageLocally": "No",
 "Diagnostic": "OFF",    //For deployment of version 7.4.1 and later without Diagnostics
template, set this parameter to OFF.
 "FmcIp":"<FMC_IP>",
 "FmcRegKey":"<REGISTRATION_KEY>",
 "FmcNatId":"<NAT_ID>",
 "Cluster": {
  "CclSubnetRange": "10.45.3.4 10.45.3.30",    //mandatory user input
  "ClusterGroupName": "ngfwv-cluster",         //mandatory user input
  "HealthProbePort": "7777",                   //mandatory user input
  "GatewayLoadBalanceIP": "10.45.2.4",         //mandatory user input
  "EncapsulationType": "vxlan",
  "InternalPort": "2000",
  "ExternalPort": "2001",
  "InternalSegId": "800",
  "ExternalSegId": "801"
 }
}
```

Note If you are copying and pasting the configuration given above, ensure that you remove **//mandatory user input** from the configuration

For the Azure health check settings, be sure to specify the **HealthProbePort** you set here.

For the **CclSubnetRange** variable, specify a range of IP addresses starting from x.x.x.4. Ensure that you have at least 16 available IP addresses for clustering. Some examples of start and end IP addresses are given below.

Table 3: Examples of Start and End IP addresses

CIDR	Start IP Address	End IP Address
10.1.1.0/27	10.1.1.4	10.1.1.30
10.1.1.32/27	10.1.1.36	10.1.1.62

CIDR	Start IP Address	End IP Address
10.1.1.64/27	10.1.1.68	10.1.1.94
10.1.1.96/27	10.1.1.100	10.1.1.126
10.1.1.128/27	10.1.1.132	10.1.1.158
10.1.1.160/27	10.1.1.164	10.1.1.190
10.1.1.192/27	10.1.1.196	10.1.1.222
10.1.1.224/27	10.1.1.228	10.1.1.254

Create the Day0 Configuration With a Customized Configuration for Azure

You can enter the entire cluster bootstrap configuration using commands.

```
{
 "AdminPassword": "password",
 "FirewallMode": "Routed",
 "ManageLocally": "No",
 "Diagnostic": "OFF",    //For deployment of version 7.4.1 and later without Diagnostics
template, set this parameter to OFF.
 "FmcIp": "<FMC_IP>",
 "FmcRegKey": "<REGISTRATION_KEY>",
 "FmcNatId": "<NAT_ID>",
 "Cluster": {
  "CclSubnetRange": "ip_address_start ip_address_end",
  "ClusterGroupName": "cluster_name",
  "HealthProbePort": "port_number",
  "GatewayLoadBalancerIP": "ip_address",
  "EncapsulationType": "vxlan",
  "InternalPort": "internal_port_number",
  "ExternalPort": "external_port_number",
  "InternalSegId": "internal_segment_id",
  "ExternalSegId": "external_segment_id"
 }
}
```

Example

A sample day 0 configuration for **version 7.4 and later** is given below.

```
{
 "AdminPassword": "Sup3rnatural",
 "Hostname": "clusterftdv",
 "FirewallMode": "routed",
 "ManageLocally": "No",
 "Diagnostic": "OFF",    //For deployment of version 7.4.1 and later without Diagnostics
template, set this parameter to OFF.
 "FmcIp": "<FMC_IP>",
 "FmcRegKey": "<REGISTRATION_KEY>",
 "FmcNatId": "<NAT_ID>",
 "run_config": [
  "cluster interface-mode individual force",
  "policy-map global_policy",
  "class inspection_default",
  "no inspect h323 h225",
  "no inspect h323 ras",
  "no inspect rtsp",
```

```
  "no inspect skinny",
  "interface Management0/0",
  "management-only",
  "nameif management",
  "security-level 0",
  "ip address dhcp",
  "interface GigabitEthernet0/0",
  "no shutdown",
  "nameif vxlan_tunnel",
  "security-level 0",
  "ip address dhcp",
  "interface GigabitEthernet0/1",
  "no shutdown",
  "nve-only cluster",
  "nameif ccl_link",
  "security-level 0",
  "ip address dhcp",
  "interface vni1",
  "description Clustering Interface",
  "segment-id 1",
  "vtep-nve 1",
  "interface vni2",
  "proxy paired",
  "nameif GWLB-backend-pool",
  "internal-segment-id 800",
  "external-segment-id 801",
  "internal-port 2000",
  "external-port 2001",
  "security-level 0",
  "vtep-nve 2",
  "object network ccl#link",
  "range 10.45.3.4 10.45.3.30",                          //mandatory user input
  "object-group network cluster#group",
  "network-object object ccl#link",
  "nve 1 ",
  "encapsulation vxlan",
  "source-interface ccl_link",
  "peer-group cluster#group",
  "nve 2 ",
  "encapsulation vxlan",
  "source-interface vxlan_tunnel",
  "peer ip <GatewayLoadbalancerIP>",
  "cluster group ftdv-cluster",                          //mandatory user input
  "local-unit 1",
  "cluster-interface vni1 ip 1.1.1.1 255.255.255.0",
  "priority 1",
  "enable",
  "mtu vxlan_tunnel 1454",
  "mtu ccl_link 1454"
 ]
}
```

A sample day 0 configuration for **version 7.3 and earlier** is given below.

```
{
 "AdminPassword": "Sup3rnatural",
 "Hostname": "clusterftdv",
 "FirewallMode": "routed",
 "ManageLocally": "No",
 "FmcIp": "<FMC_IP>",
 "FmcRegKey": "<REGISTRATION_KEY>",
 "FmcNatId": "<NAT_ID>",
 "run_config": [
  "cluster interface-mode individual force",
```

```
"policy-map global_policy",
"class inspection_default",
"no inspect h323 h225",
"no inspect h323 ras",
"no inspect rtsp",
"no inspect skinny",
"interface Management0/0",
"management-only",
"nameif management",
"security-level 0",
"ip address dhcp",
"interface GigabitEthernet0/0",
"no shutdown",
"nameif vxlan_tunnel",
"security-level 0",
"ip address dhcp",
"interface GigabitEthernet0/1",
"no shutdown",
"nve-only cluster",
"nameif ccl_link",
"security-level 0",
"ip address dhcp",
"interface vni1",
"description Clustering Interface",
"segment-id 1",
"vtep-nve 1",
"interface vni2",
"proxy paired",
"nameif GWLB-backend-pool",
"internal-segment-id 800",
"external-segment-id 801",
"internal-port 2000",
"external-port 2001",
"security-level 0",
"vtep-nve 2",
"object network ccl#link",
"range 10.45.3.4 10.45.3.30",                    //mandatory user input
"object-group network cluster#group",
"network-object object ccl#link",
"nve 1 ",
"encapsulation vxlan",
"source-interface ccl_link",
"peer-group cluster#group",
"nve 2 ",
"encapsulation vxlan",
"source-interface vxlan_tunnel",
"peer ip <GatewayLoadbalancerIP>",
"cluster group ftdv-cluster",                    //mandatory user input
"local-unit 1",
"cluster-interface vni1 ip 1.1.1.1 255.255.255.0",
"priority 1",
"enable",
"mtu vxlan_tunnel 1454",
"mtu ccl_link 1554"
 ]
}
```

Note If you are copying and pasting the configuration given above, ensure that you remove **//mandatory user input** from the configuration.

Deploy Cluster Nodes Manually - GWLB-based Deployment

Deploy the cluster nodes so they form a cluster.

Procedure

Step 1 Create a Virtual Machine Scale Set from the Marketplace image with 0 instance count using the **az vmss create** CLI.

az vmss create --resource-group <*ResourceGroupName*> **--name** <*VMSSName*> **--vm-sku** <*InstanceSize*> **--image** <*FTDvImage*> **--instance-count 0 --admin-username** <*AdminUserName*> **--admin-password** <*AdminPassword*> **--plan-name** <*ftdv-azure-byol/ftdv-azure-payg*> **--plan-publisher cisco --plan-product cisco-ftdv --plan-promotion-code** <*ftdv-azure-byol/ftdv-azure-payg*> **--vnet-name** <*VirtualNetworkName*> **--subnet** <*MgmtSubnetName*>

Step 2 Attach three interfaces—Diagnostic, Data, and Cluster Control Link.

Step 3 Go to the virtual machine scale set you have created and perform the following steps:

a) Under the **Operating system** section, add the day 0 configuration in the **customData** field.
b) Click **Save**.
c) Under the **Scaling** section, update the instance count with the required cluster node. You can set the instance count range of minimum 1 and maximum 16.

Step 4 Configure the Azure Gateway Load Balancer. See Auto Scale with Azure Gateway Load Balancer Use Case for more information.

Step 5 Add the control node to the management center. See Add the Cluster to the Management Center (Manual Deployment), on page 52.

Deploy Cluster Nodes Manually - NLB-based Deployment

Deploy the cluster nodes so they form a cluster.

Procedure

Step 1 Create a Virtual Machine Scale Set from the Marketplace image with 0 instance count using the **az vmss create** CLI.

az vmss create --resource-group <*ResourceGroupName*> **--name** <*VMSSName*> **--vm-sku** <*InstanceSize*> **--image** <*FTDvImage*> **--instance-count 0 --admin-username** <*AdminUserName*> **--admin-password** <*AdminPassword*> **--plan-name** <*ftdv-azure-byol/ftdv-azure-payg*> **--plan-publisher cisco --plan-product cisco-ftdv --plan-promotion-code** <*ftdv-azure-byol/ftdv-azure-payg*> **--vnet-name** <*VirtualNetworkName*> **--subnet** <*MgmtSubnetName*>

Step 2 Attach 4 interfaces—Diagnostic, Inside, Outside, and Cluster Control Link.

Step 3 Go to the virtual machine scale set you have created and perform the following:

a) Under the **Operating system** section, add the day0 configuration in the **customData** field.
b) Click **Save**.

c) Under the **Scaling** section, update the instance count with the required cluster node. You can set the instance count range of minimum 1 and maximum 16.

Step 4 Add the control node to the Management Center. See Add the Cluster to the Management Center (Manual Deployment), on page 52.

Troubleshooting Cluster Deployment in Azure

- Issue: No traffic flow

  Troubleshooting:

  - Check if the health probe status of the Threat Defense Virtual instances deployed with a GWLB is healthy.
  - If the Threat Defense Virtual instance's health probe status is unhealthy-
    - Check if the static route is configured in the Management Center Virtual.
    - Check if the default gateway is the data subnet's gateway IP.
    - Check if the Threat Defense Virtual instance is receiving health probe traffic.
    - Check if the access list configured in the Management Center Virtual allows health probe traffic.

- Issue: Cluster is not formed

  Troubleshooting:

  - Check the IP address of the nve-only cluster interface. Ensure that you can ping the nve-only cluster interface of other nodes.
  - Check the IP address of the nve-only cluster interfaces are part of the object group.
  - Ensure that the NVE interface is configured with the object group .
  - Ensure that the cluster interface in the cluster group has the right VNI interface. This VNI interface has the NVE with the corresponding object group.
  - Ensure that the nodes are pingable from each other. Since each node has its own cluster interface IP, these should be pingable from each other.
  - Check if the CCL Subnet's Start and End Address mentioned during template deployment is correct. The start address should begin with the first available IP address in the subnet. For example, if the subnet is 192.168.1.0/24. The start address should be 192.168.1.4 (the three IP addresses at the start are reserved by Azure).
  - Check if the Management Center Virtual has a valid license.

- Issue: Role-related error while deploying resources again in the same resource group.

  Troubleshooting: Remove the roles given below by using the following commands on the terminal.

  Error message:

```
"error": {
"code": "RoleAssignmentUpdateNotPermitted",
"message": "Tenant ID, application ID, principal ID, and scope are not allowed to be
updated."}
```

- **az role assignment delete --resource-group** <*Resource Group Name*> **--role "Storage Queue Data Contributor"**
- **az role assignment delete --resource-group** <*Resource Group Name*> **--role "Contributor"**

Deploy the Cluster in GCP

To deploy a cluster in GCP, you can either manually deploy or use an instance template to deploy an instance group. You can use the cluster with native GCP load-balancers, or non-native load balancers such as the Cisco Cloud Services Router.

Note Outbound traffic requires interface NAT and is limited to 64K connections.

Sample Topology

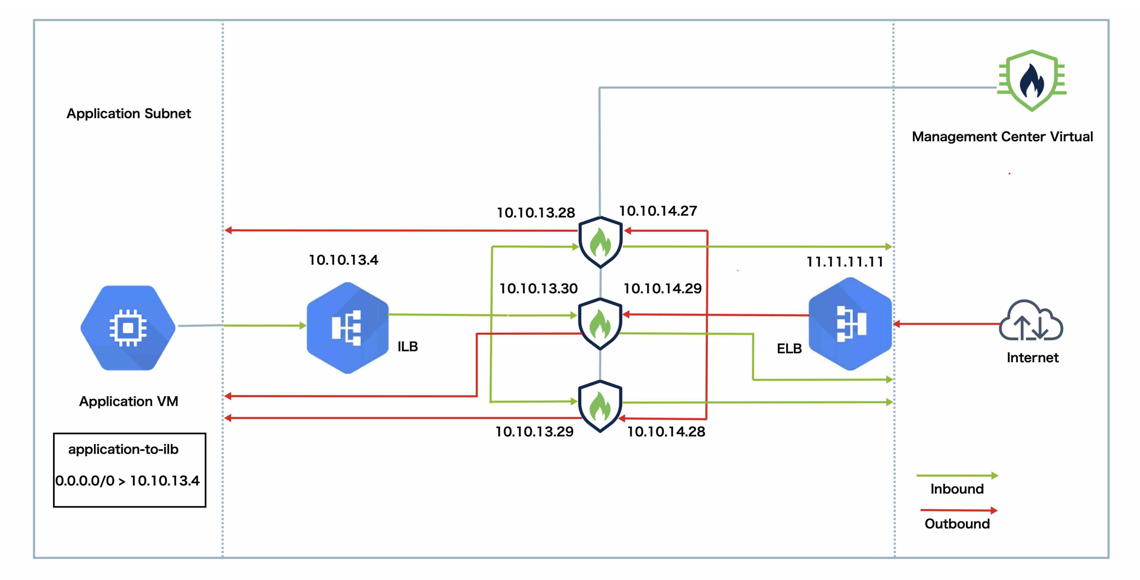

This topology depicts both inbound and outbound traffic flow. The Threat Defense Virtual cluster is sandwiched between the internal and external load balancers. A Management Center Virtual instance is used to manage the cluster.

Inbound traffic from the internet goes to the external load balancer which then transmits the traffic to the Threat Defense Virtual cluster. After the traffic has been inspected by a Threat Defense Virtual instance in the cluster, it is forwarded to the application VM.

Outbound traffic from the application VM is transmitted to the internal load balancer. Traffic is then forwarded to the Threat Defense Virtual cluster and then sent out to the internet.

End-to-End Process for Deploying Threat Defense Virtual Cluster in GCP

Template-based Deployment

The following flowchart illustrates the workflow for template-based deployment of the Threat Defense Virtual cluster on GCP.

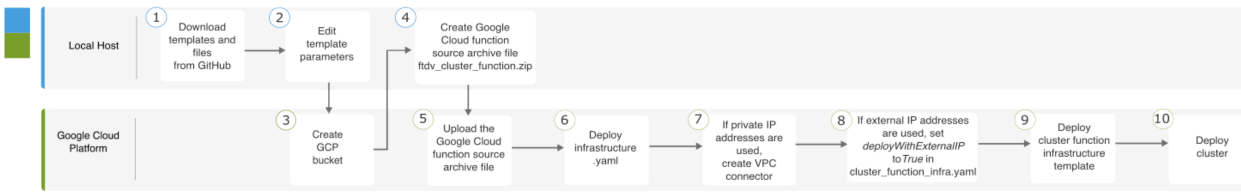

	Workspace	Steps
1	Local Host	Download templates and files from GitHub.
2	Local Host	Edit template parameters.
3	Google Cloud Platform	Create GCP bucket.
4	Local Host	Create Google Cloud function source archive file *ftdv_cluster_function.zip*.
5	Google Cloud Platform	Upload the Google function source archive file.
6	Google Cloud Platform	Deploy infrastructure.yaml.
7	Google Cloud Platform	If private IP addresses are used, create VPC connector.
8	Google Cloud Platform	If external IP addresses are used, set *deployWithExternalIP* to *True* in *cluster_function_infra.yaml*.
9	Google Cloud Platform	Deploy cluster function infrastructure template.
10	Google Cloud Platform	Deploy cluster.

Manual Deployment

The following flowchart illustrates the workflow for manual deployment of the Threat Defense Virtual cluster on GCP.

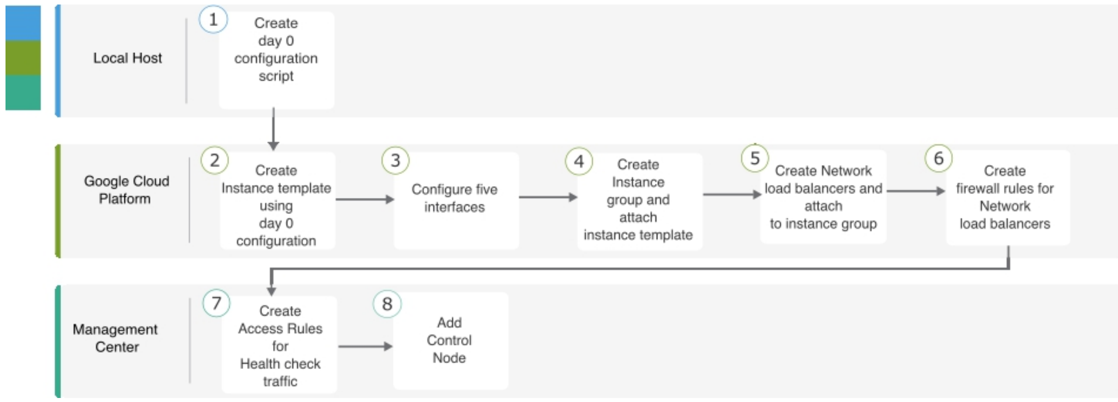

	Workspace	Steps
1	Local Host	Create the Day0 Configuration for GCP
2	Google Cloud Platform	Create instance template using day 0 configuration.
3	Google Cloud Platform	Configure the interfaces.
4	Google Cloud Platform	Create instance group and attach instance template.
5	Google Cloud Platform	Create NLB and attach to instance group.
6	Google Cloud Platform	Create firewall rules for NLB.
7	Management Center	Create access rules for health check traffic.
8	Management Center	Add control node.

Templates

The templates given below are available in GitHub. The parameter values are self-explanatory with the parameter names, and values, given in the template.

- Cluster deployment template for East-West traffic - *deploy_ngfw_cluster_yaml*
- Cluster deployment template for North-South traffic - deploy_ngfw_cluster.yaml

Deploy the Instance Group in GCP Using an Instance Template

Deploy the instance group in GCP using an instance template.

Before you begin

- Use Google Cloud Shell for deployment. Alternatively, you can use Google SDK on any macOS/Linux/Windows machine.
- To allow the cluster to auto-register with the Management Center, you need to create a user with administrative privileges on the Management Center that can use the REST API. See the Cisco Secure Firewall Management Center Administration Guide.
- Add an access policy in the Management Center that matches the name of the policy that you specified in *cluster_function_infra.yaml*.

Procedure

Step 1 Download the templates from GitHub to your local folder.

Step 2 Edit **infrastructure.yaml** , **cluster_function_infra.yaml** and **deploy_ngfw_cluster.yaml** with the required *resourceNamePrefix* parameter (for example, ngfwvcls) and other required user inputs.

From Secure Firewall version 7.4.1, you can deploy the cluster without the diagnostic interface. To deploy the cluster with only the Outside, Inside, Management, and CCL interfaces, set the *withDiagnostic* variable to **False** in both the **infrastructure.yaml** and the **deploy_ngfw_cluster.yaml** files.

Note that there is a **deploy_ngfw_cluster.yaml** file in both the **east-west** and **north-south** folders in GitHub. Download the appropriate template as per your traffic flow requirement.

Step 3 Create a bucket using Google Cloud Shell to upload the Google cloud function source archive file *ftdv_cluster_function.zip*.

gsutil mb --pap enforced gs://*resourceNamePrefix***-ftdv-cluster-bucket/**

Ensure that the *resourceNamePrefix* variable here matches the *resourceNamePrefix* variable that you specified in **cluster_function_infra.yaml**.

Step 4 Create an archive file for the cluster infrastructure.

Example:

```
zip -j ftdv_cluster_function.zip ./cluster-function/*
```

Step 5 Upload the Google source archive that you created earlier.

gsutil cp ftdv_cluster_function.zip gs://*resourceNamePrefix***-ftdv-cluster-bucket/**

Step 6 Deploy infrastructure for the cluster.

gcloud deployment-manager deployments create *cluster_name* **--config infrastructure.yaml**

Step 7 If you are using private IP addresses, perform the steps given below:

a) Launch and set up the Management Center Virtual with a Threat Defense Virtual management VPC.
b) Create a VPC connector to connect the Google Cloud functions with the Threat Defense Virtual management VPC.

   gcloud compute networks vpc-access connectors create *vpc-connector-name* **--region us-central1 --subnet** *resourceNamePrefix***-ftdv-mgmt-subnet28**

Step 8 If the Management Center is remote from the Threat Defense Virtual, and the Threat Defense Virtual needs an external IP address, ensure that you set **deployWithExternalIP** to **True** in **cluster_function_infra.yaml**.

Step 9 Deploy the cluster function infrastructure.

gcloud deployment-manager deployments create *cluster_name* **--config cluster_function_infra.yaml**

Step 10 Deploy the cluster.

a. For North-South topology deployment:

   gcloud deployment-manager deployments create *cluster_name* **--config north-south/deploy_ngfw_cluster.yaml**

b. For East-West topology deployment:

   gcloud deployment-manager deployments create *cluster_name* **--config east-west/deploy_ngfw_cluster.yaml**

Deploy the Cluster in GCP Manually

To deploy the cluster manually, prepare the day0 configuration, deploy each node, and then add the control node to the management center.

Create the Day0 Configuration for GCP

You can use either a fixed configuration or a customized configuration.

Create the Day0 Configuration With a Fixed Configuration for GCP

The fixed configuration will auto-generate the cluster bootstrap configuration.

```
{
    "AdminPassword": "password",
    "Hostname": "hostname",
    "FirewallMode": "Routed",
    "ManageLocally": "No",
                "Diagnostic": "OFF",       //Optional user input from version 7.4.1 - use
to deploy cluster without Diagnostic interface
    "Cluster": {
        "CclSubnetRange": "ip_address_start ip_address_end",
        "ClusterGroupName": "cluster_name"
    }
}
```

For example:

```
{
    "AdminPassword": "DeanW1nche$ter",
    "Hostname": "ciscoftdv",
    "FirewallMode": "Routed",
    "ManageLocally": "No",
    "Cluster": {
        "CclSubnetRange": "10.10.55.2 10.10.55.253",     //mandatory user input
        "ClusterGroupName": "ftdv-cluster"               //mandatory user input
```

```
        }
}
```

> **Note** If you are copying and pasting the configuration given above, ensure that you remove **//mandatory user input** from the configuration.

For the **CclSubnetRange** variable, note that you cannot use the first two IP addresses and the last two IP addresses in the subnet. See Reserved IP addresses in IPv4 subnets for more information. Ensure that you have at least 16 available IP addresses for clustering. Some examples of start and end IP addresses are given below.

Table 4: Examples of Start and End IP addresses

CIDR	Start IP Address	End IP Address
10.1.1.0/27	10.1.1.2	10.1.1.29
10.1.1.32/27	10.1.1.34	10.1.1.61
10.1.1.64/27	10.1.1.66	10.1.1.93
10.1.1.96/27	10.1.1.98	10.1.1.125
10.1.1.128/27	10.1.1.130	10.1.1.157
10.1.1.160/27	10.1.1.162	10.1.1.189
10.1.1.192/27	10.1.1.194	10.1.1.221
10.1.1.224/27	10.1.1.226	10.1.1.253
10.1.1.0/24	10.1.1.2	10.1.1.253

Create the Day0 Configuration With a Customized Configuration for GCP

You can enter the entire cluster bootstrap configuration using commands.

```
{
    "AdminPassword": "password",
    "Hostname": "hostname",
    "FirewallMode": "Routed",
    "ManageLocally": "No",
    "run_config": [comma_separated_threat_defense_configuration]
    }
```

The following example creates a configuration with Management, Inside, and Outside interfaces, and a VXLAN interface for the cluster control link. Note the values in bold that need to be unique per node.

```
{
 "AdminPassword": "W1nch3sterBr0s",
 "Hostname": "ftdv1",
 "FirewallMode": "Routed",
 "ManageLocally": "No",
 "run_config": [
  "cluster interface-mode individual force",
  "interface Management0/0",
  "management-only",
  "nameif management",
```

```
  "ip address dhcp",
  "interface GigabitEthernet0/0",
  "no shutdown",
  "nameif outside",
  "ip address dhcp",
  "interface GigabitEthernet0/1",
  "no shutdown",
  "nameif inside",
  "ip address dhcp",
  "interface GigabitEthernet0/2",
  "nve-only cluster",
  "nameif ccl_link",
  "ip address dhcp",
  "no shutdown",
  "interface vni1",
  "description Clustering Interface",
  "segment-id 1",
  "vtep-nve 1",
  "object network ccl#link",
  "range 10.1.90.2 10.1.90.17",
  "object-group network cluster#group",
  "network-object object ccl#link",
  "nve 1",
  "encapsulation vxlan",
  "source-interface ccl_link",
  "peer-group cluster#group",
  "cluster group ftdv-cluster",
  "local-unit 1",
  "cluster-interface vni1 ip 10.1.1.1 255.255.255.0",
  "priority 1",
  "enable",
  "mtu outside 1400",
  "mtu inside 1400"
 ]
}
```

Note For the cluster control link network object, specify only as many addresses as you need (up to 16). A larger range can affect performance.

Deploy Cluster Nodes Manually

Deploy the cluster nodes so they form a cluster. For clustering on GCP, you cannot use the 4 vCPU machine type. The 4 vCPU machine type only supports four interfaces, and five are needed. Use a machine type that supports five interfaces, such as c2-standard-8.

Procedure

Step 1 Create an instance template using the day 0 configuration (in the **Metadata** > **Startup Script** section) with 5 interfaces: outside, inside, management, diagnostic, and cluster control link.

See Cisco Secure Firewall Threat Defense Virtual Getting Started Guide.

Step 2 Create an instance group, and attach the instance template.

Step 3 Create GCP network load balancers (internal and external), and attach the instance group.

Step 4 For GCP network load balancers, allow health checks in your security policy on the Management Center. See Allow Health Checks for GCP Network Load Balancers, on page 51.

Step 5 Add the control node to the Management Center. See Add the Cluster to the Management Center (Manual Deployment), on page 52.

Allow Health Checks for GCP Network Load Balancers

Google Cloud provides health checks to determine if backends respond to traffic.

See https://cloud.google.com/load-balancing/docs/health-checks to create firewall rules for network load balancers. Then in the management center, create access rules to allow the health check traffic. See https://cloud.google.com/load-balancing/docs/health-check-concepts for the required network ranges. See Access Control Rules.

You also need to configure dynamic manual NAT rules to redirect the health check traffic to the Google metadata server at 169.254.169.254. See Configure Dynamic Manual NAT.

North-South NAT Rules Sample Configuration

```
nat (inside,outside) source dynamic GCP-HC ILB-SOUTH destination static ILB-SOUTH METADATA
nat (outside,outside) source dynamic GCP-HC ELB-NORTH destination static ELB-NORTH METADATA

nat (outside,inside) source static any interface destination static ELB-NORTH Ubuntu-App-VM
nat (inside,outside) source dynamic any interface destination static obj-any obj-any

object network Metadata
 host 169.254.169.254

object network ILB-SOUTH
host <ILB_IP>
object network ELB-NORTH
host <ELB_IP>

object-group network GCP-HC
 network-object 35.191.0.0 255.255.0.0
 network-object 130.211.0.0 255.255.252.0
 network-object 209.85.204.0 255.255.252.0
 network-object 209.85.152.0 255.255.252.0
```

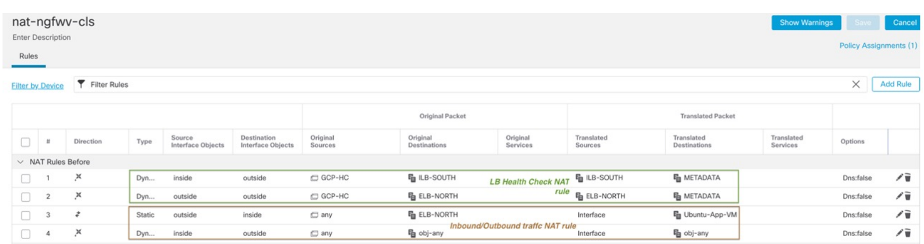

East-West NAT Rules Sample Configuration

```
nat (inside,outside) source dynamic GCP-HC ILB-East destination static ILB-East Metadata
nat (outside,outside) source dynamic GCP-HC ILB-West destination static ILB-West Metadata

object network Metadata
```

```
 host 169.254.169.254

object network ILB-East
host <ILB_East_IP>
object network ILB-West
host <ILB_West_IP>

object-group network GCP-HC
 network-object 35.191.0.0 255.255.0.0
 network-object 130.211.0.0 255.255.252.0
 network-object 209.85.204.0 255.255.252.0
 network-object 209.85.152.0 255.255.252.0
```

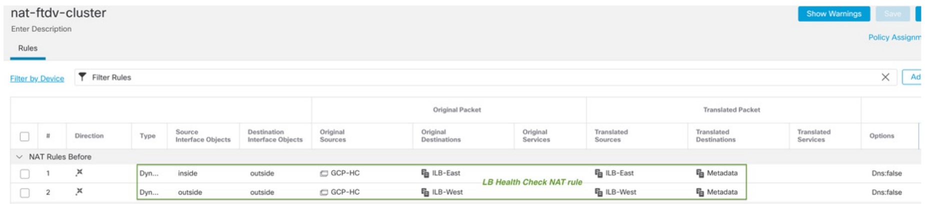

Add the Cluster to the Management Center (Manual Deployment)

Use this procedure to add the cluster to the management center if you manually deployed the cluster. If you used a template, the cluster will auto-register on the management center.

Add one of the cluster units as a new device to the management center; the management center auto-detects all other cluster members.

Before you begin

- All cluster units must be in a successfully-formed cluster prior to adding the cluster to the management center. You should also check which unit is the control unit. Use the threat defense **show cluster info** command.

Procedure

Step 1 In the management center, choose **Devices** > **Device Management**, and then choose **Add** > **Add Device** to add the control unit using the unit's management IP address.

Figure 15: Add Device

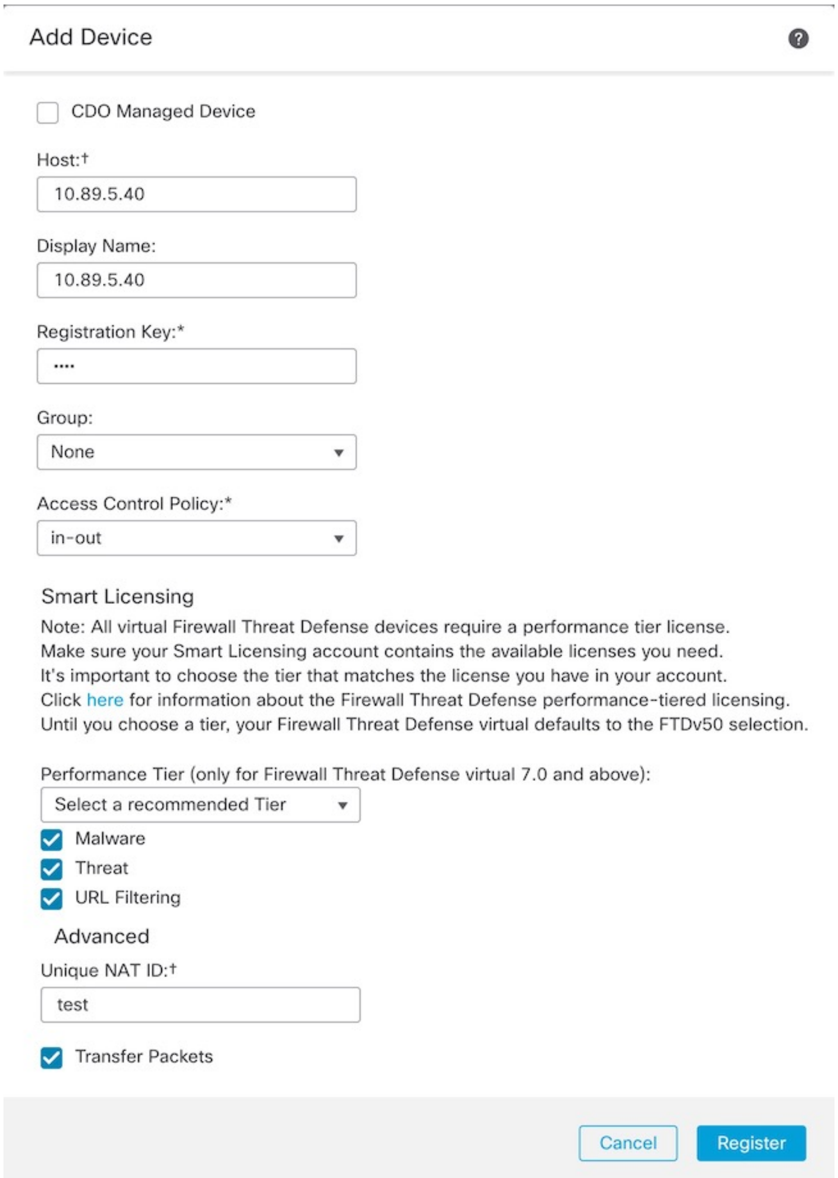

a) In the **Host** field, enter the IP address or hostname of the control unit.

   We recommend adding the control unit for the best performance, but you can add any unit of the cluster.

   If you used a NAT ID during device setup, you may not need to enter this field.

b) In the **Display Name** field, enter a name for the control unit as you want it to display in the management center.

   This display name is not for the cluster; it is only for the control unit you are adding. You can later change the name of other cluster members and the cluster display name.

c) In the **Registration Key** field, enter the same registration key that you used during device setup. The registration key is a one-time-use shared secret.

d) (Optional) Add the device to a device **Group**.

e) Choose an initial **Access Control Policy** to deploy to the device upon registration, or create a new policy.

If you create a new policy, you create a basic policy only. You can later customize the policy as needed.

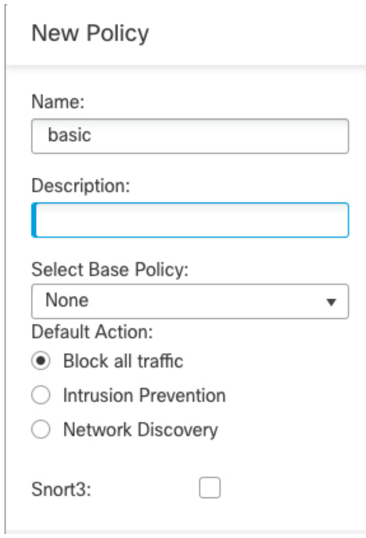

f) Choose licenses to apply to the device.

g) If you used a NAT ID during device setup, expand the **Advanced** section and enter the same NAT ID in the **Unique NAT ID** field.

h) Check the **Transfer Packets** check box to allow the device to transfer packets to the management center.

This option is enabled by default. When events like IPS or Snort are triggered with this option enabled, the device sends event metadata information and packet data to the management center for inspection. If you disable it, only event information will be sent to the management center but packet data is not sent.

i) Click **Register**.

The management center identifies and registers the control unit, and then registers all data units. If the control unit does not successfully register, then the cluster is not added. A registration failure can occur if the cluster was not up, or because of other connectivity issues. In this case, we recommend that you try re-adding the cluster unit.

The cluster name shows on the **Devices** > **Device Management** page; expand the cluster to see the cluster units.

Figure 16: Cluster Management

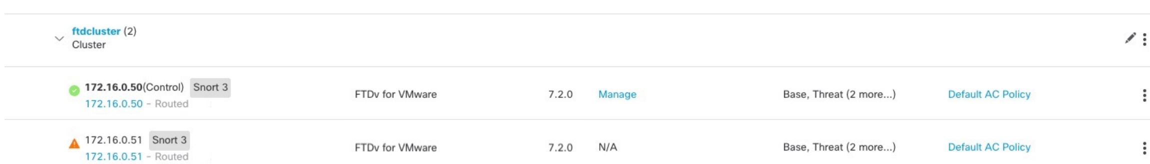

A unit that is currently registering shows the loading icon.

Figure 17: Node Registration

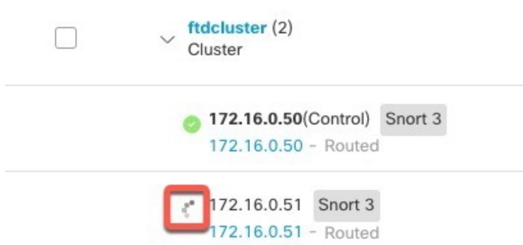

You can monitor cluster unit registration by clicking the **Notifications** icon and choosing **Tasks**. The management center updates the Cluster Registration task as each unit registers. If any units fail to register, see Reconcile Cluster Nodes, on page 64.

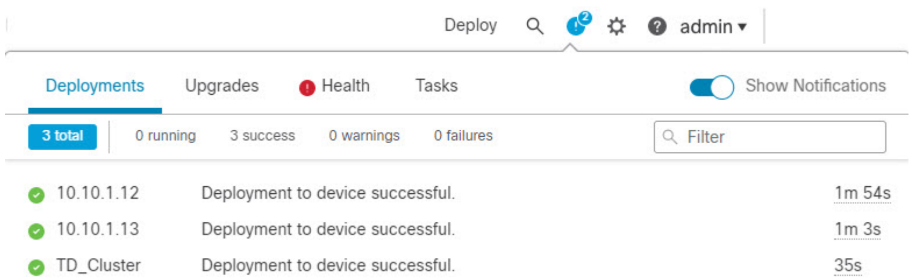

Step 2 Configure device-specific settings by clicking the **Edit** (✎) for the cluster.

Most configuration can be applied to the cluster as a whole, and not nodes in the cluster. For example, you can change the display name per node, but you can only configure interfaces for the whole cluster.

Step 3 On the **Devices** > **Device Management** > **Cluster** screen, you see **General**, **License**, **System**, and **Health** settings.

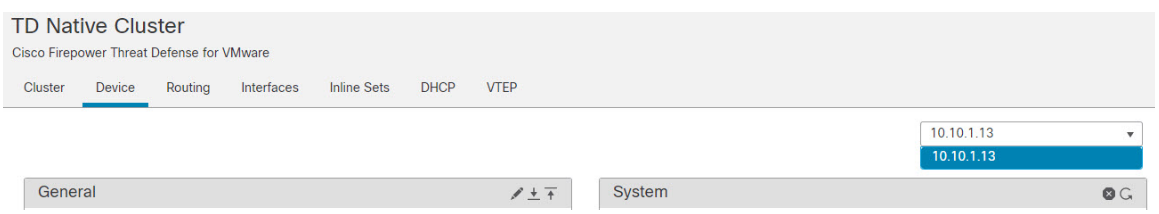

See the following cluster-specific items:

- **General** > **Name**—Change the cluster display name by clicking the **Edit** (✎).

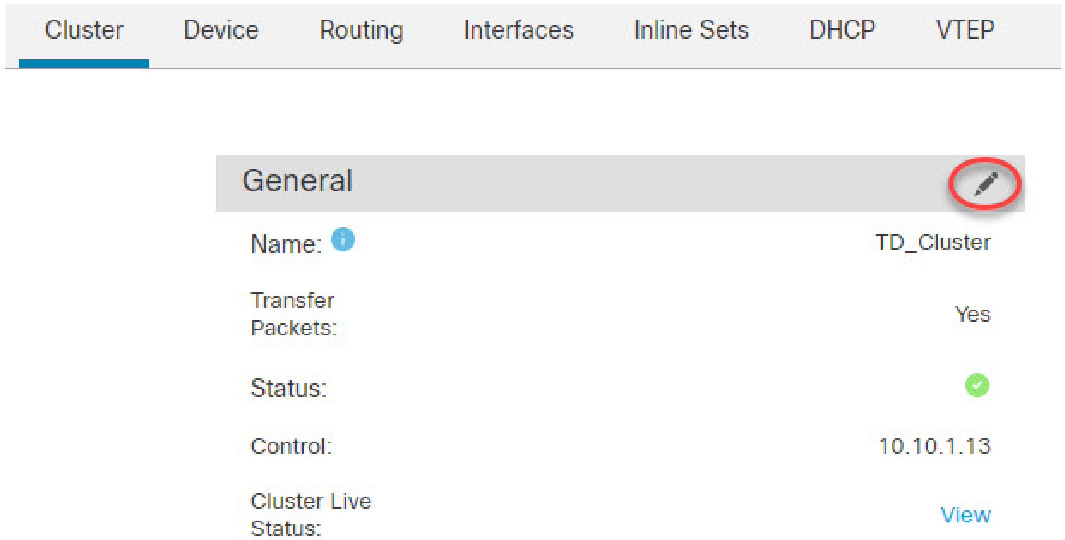

Then set the **Name** field.

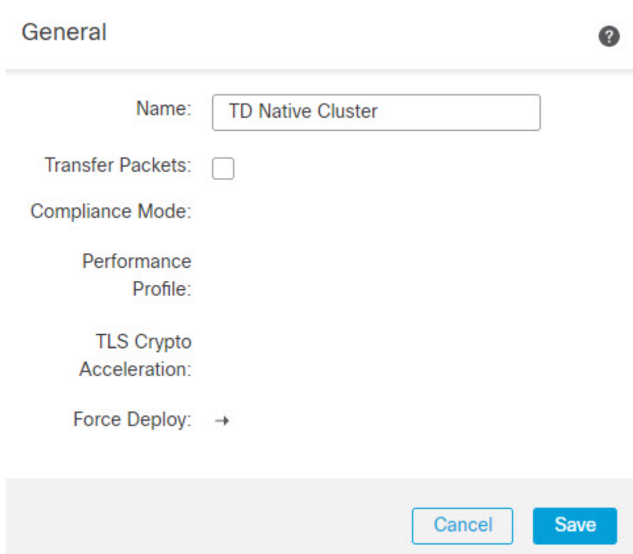

- **General** > **Cluster Live Status**—Click the **View** link to open the **Cluster Status** dialog box.

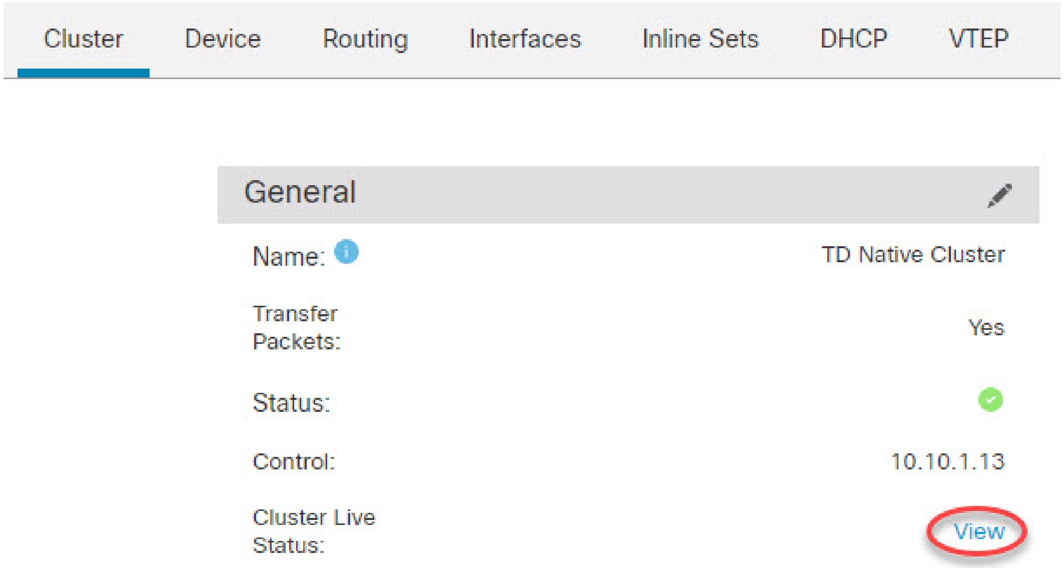

The **Cluster Status** dialog box also lets you retry data unit registration by clicking **Reconcile**.You can also ping the cluster control link from a node. See Perform a Ping on the Cluster Control Link, on page 72.

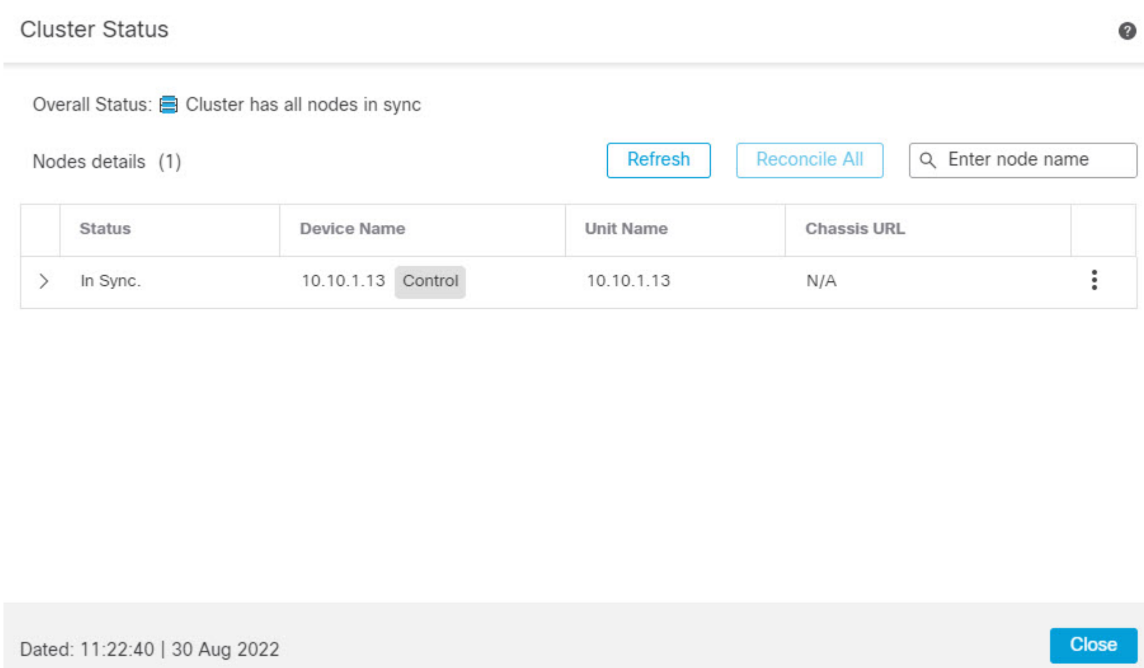

- **General** > **Troubleshoot**—You can generate and download troubleshooting logs, and you can view cluster CLIs. See Troubleshooting the Cluster, on page 72.

Figure 18: Troubleshoot

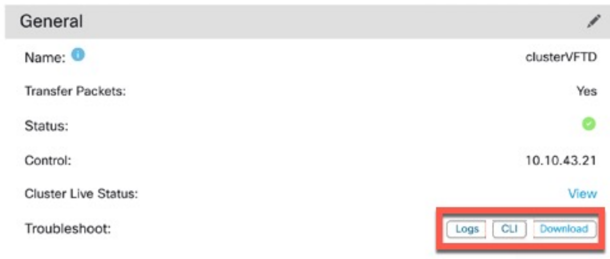

- **License**—Click **Edit** (✎) to set license entitlements.

Step 4 On the **Devices** > **Device Management** > **Devices**, you can choose each member in the cluster from the top right drop-down menu and configure the following settings.

- **General** > **Name**—Change the cluster member display name by clicking the **Edit** (✎).

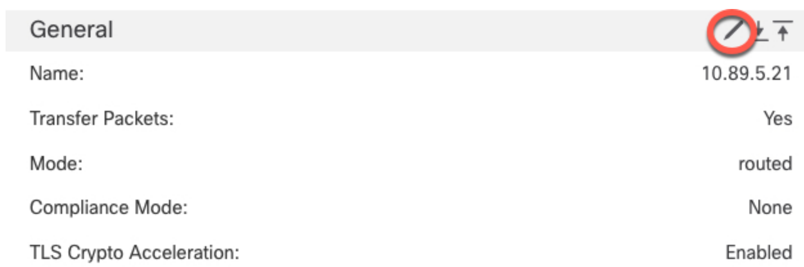

Then set the **Name** field.

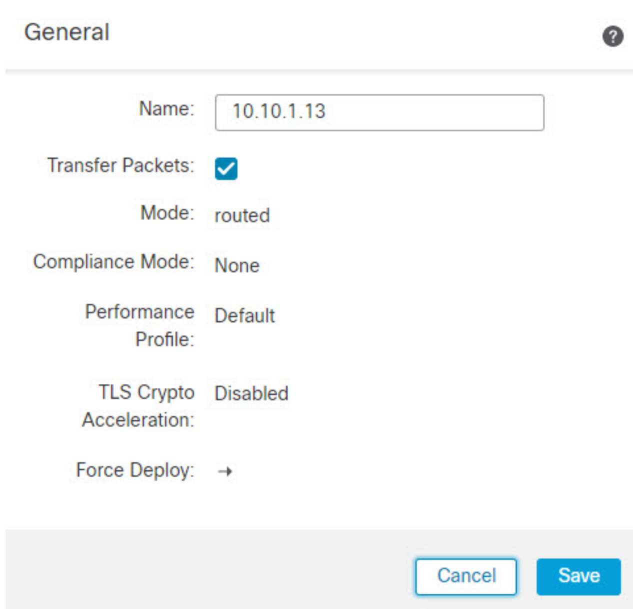

- **Management** > **Host**—If you change the management IP address in the device configuration, you must match the new address in the management center so that it can reach the device on the network; edit the **Host** address in the **Management** area.

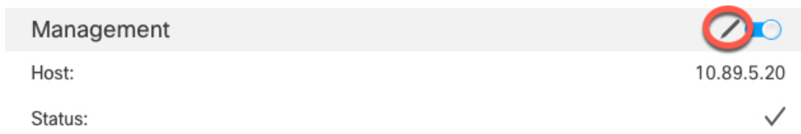

Configure Cluster Health Monitor Settings

The **Cluster Health Monitor Settings** section of the **Cluster** page displays the settings described in the table below.

Figure 19: Cluster Health Monitor Settings

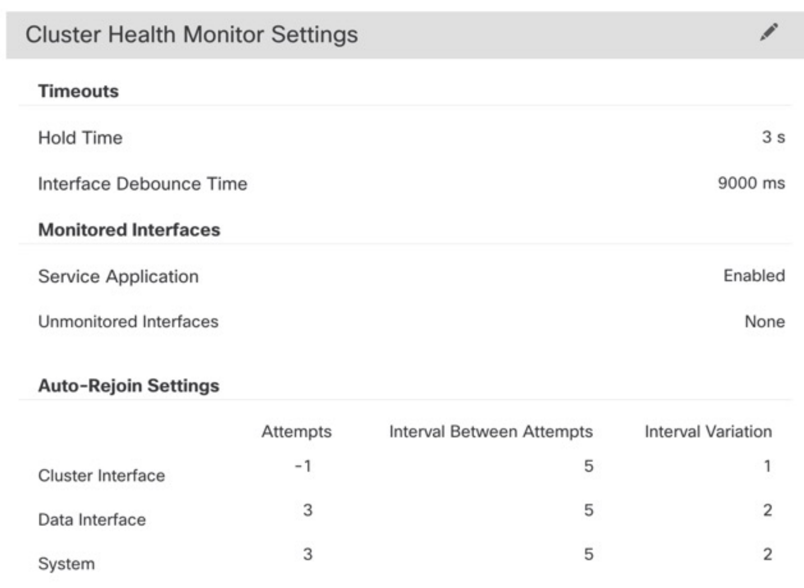

Table 5: Cluster Health Monitor Settings Section Table Fields

Field	Description
Timeouts	
Hold Time	Between .3 and 45 seconds; The default is 3 seconds. To determine node system health, the cluster nodes send heartbeat messages on the cluster control link to other nodes. If a node does not receive any heartbeat messages from a peer node within the hold time period, the peer node is considered unresponsive or dead.

Field	Description
Interface Debounce Time	Between 300 and 9000 ms. The default is 500 ms. The interface debounce time is the amount of time before the node considers an interface to be failed, and the node is removed from the cluster.
Monitored Interfaces	The interface health check monitors for link failures. If all physical ports for a given logical interface fail on a particular node, but there are active ports under the same logical interface on other nodes, then the node is removed from the cluster. The amount of time before the node removes a member from the cluster depends on the type of interface and whether the node is an established node or is joining the cluster.
Service Application	Shows whether the Snort and disk-full processes are monitored.
Unmonitored Interfaces	Shows unmonitored interfaces.
Auto-Rejoin Settings	
Cluster Interface	Shows the auto-rejoin settings after a cluster control link failure.
Attempts	Between -1 and 65535. The default is -1 (unlimited). Sets the number of rejoin attempts.
Interval Between Attempts	Between 2 and 60. The default is 5 minutes. Defines the interval duration in minutes between rejoin attempts.
Interval Variation	Between 1 and 3. The default is 1x the interval duration. Defines if the interval duration increases at each attempt.
Data Interfaces	Shows the auto-rejoin settings after a data interface failure.
Attempts	Between -1 and 65535. The default is 3. Sets the number of rejoin attempts.
Interval Between Attempts	Between 2 and 60. The default is 5 minutes. Defines the interval duration in minutes between rejoin attempts.
Interval Variation	Between 1 and 3. The default is 2x the interval duration. Defines if the interval duration increases at each attempt.
System	Shows the auto-rejoin settings after internal errors. Internal failures include: application sync timeout; inconsistent application statuses; and so on.
Attempts	Between -1 and 65535. The default is 3. Sets the number of rejoin attempts.
Interval Between Attempts	Between 2 and 60. The default is 5 minutes. Defines the interval duration in minutes between rejoin attempts.
Interval Variation	Between 1 and 3. The default is 2x the interval duration. Defines if the interval duration increases at each attempt.

> **Note** If you disable the system health check, fields that do not apply when the system health check is disabled will not show.

You can change these settings from this section.

You can monitor any port-channel ID, single physical interface ID, as well as the Snort and disk-full processes. Health monitoring is not performed on VLAN subinterfaces or virtual interfaces such as VNIs or BVIs. You cannot configure monitoring for the cluster control link; it is always monitored.

Procedure

Step 1 Choose **Devices** > **Device Management**.

Step 2 Next to the cluster you want to modify, click **Edit** (✎).

Step 3 Click **Cluster**.

Step 4 In the **Cluster Health Monitor Settings** section, click **Edit** (✎).

Step 5 Disable the system health check by clicking the **Health Check** slider .

Figure 20: Disable the System Health Check

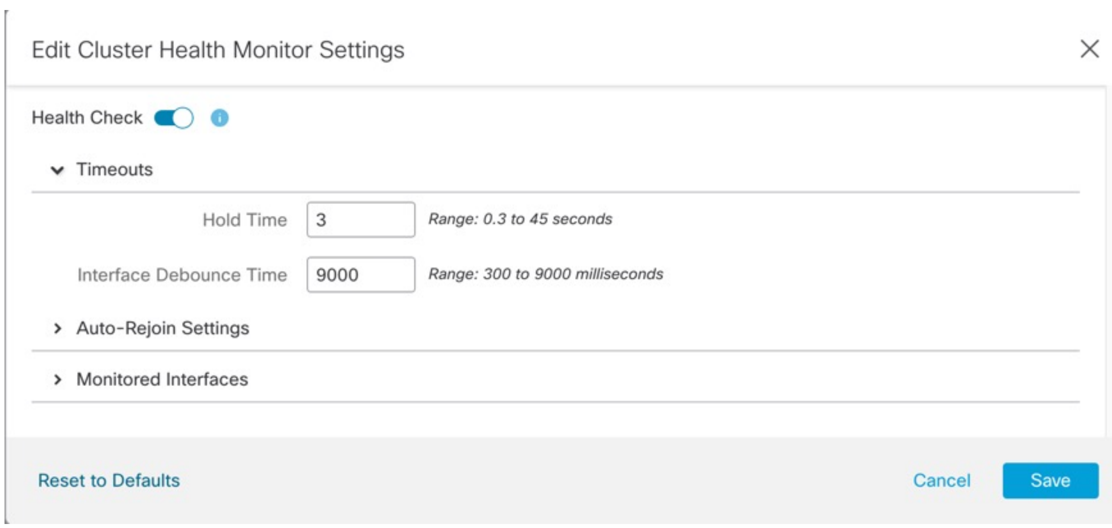

When any topology changes occur (such as adding or removing a data interface, enabling or disabling an interface on the node or the switch, or adding an additional switch to form a VSS or vPC) you should disable the system health check feature and also disable interface monitoring for the disabled interfaces. When the topology change is complete, and the configuration change is synced to all nodes, you can re-enable the system health check feature and monitored interfaces.

Step 6 Configure the hold time and interface debounce time.

- **Hold Time**—Set the hold time to determine the amount of time between node heartbeat status messages, between .3 and 45 seconds; The default is 3 seconds.
- **Interface Debounce Time**—Set the debounce time between 300 and 9000 ms. The default is 500 ms. Lower values allow for faster detection of interface failures. Note that configuring a lower debounce

time increases the chances of false-positives. When an interface status update occurs, the node waits the number of milliseconds specified before marking the interface as failed, and the node is removed from the cluster. In the case of an EtherChannel that transitions from a down state to an up state (for example, the switch reloaded, or the switch enabled an EtherChannel), a longer debounce time can prevent the interface from appearing to be failed on a cluster node just because another cluster node was faster at bundling the ports.

Step 7 Customize the auto-rejoin cluster settings after a health check failure.

Figure 21: Configure Auto-Rejoin Settings

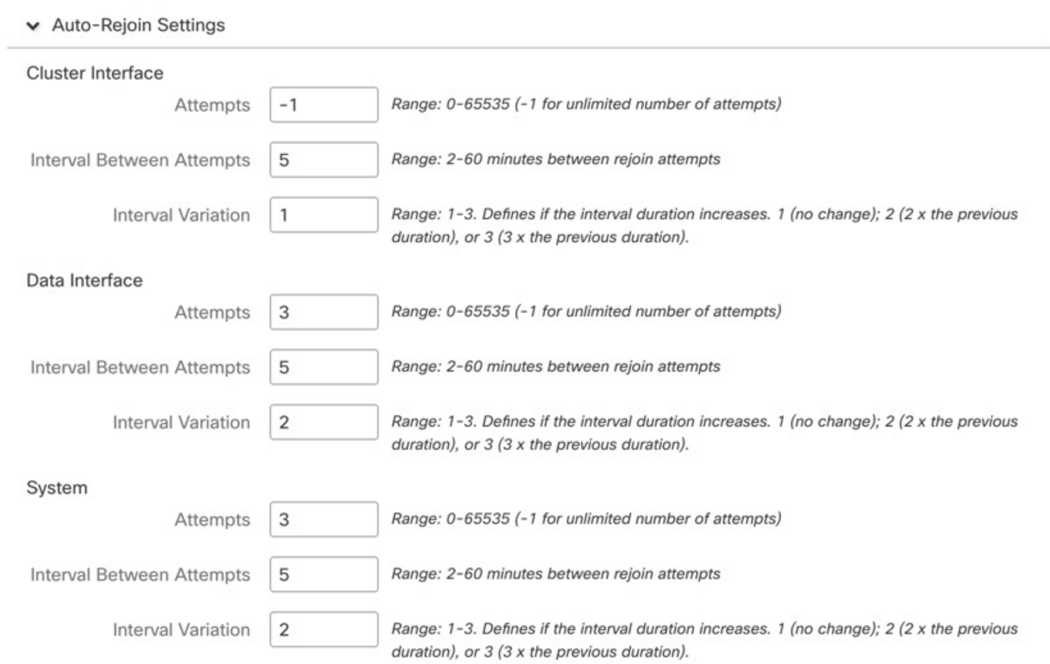

Set the following values for the **Cluster Interface**, **Data Interface**, and **System** (internal failures include: application sync timeout; inconsistent application statuses; and so on):

- **Attempts**—Sets the number of rejoin attempts, between -1 and 65535. **0** disables auto-rejoining. The default for the **Cluster Interface** is -1 (unlimited). The default for the **Data Interface** and **System** is 3.
- **Interval Between Attempts**—Defines the interval duration in minutes between rejoin attempts, between 2 and 60. The default value is 5 minutes. The maximum total time that the node attempts to rejoin the cluster is limited to 14400 minutes (10 days) from the time of last failure.
- **Interval Variation**—Defines if the interval duration increases. Set the value between 1 and 3: **1** (no change); **2** (2 x the previous duration), or **3** (3 x the previous duration). For example, if you set the interval duration to 5 minutes, and set the variation to 2, then the first attempt is after 5 minutes; the 2nd attempt is 10 minutes (2 x 5); the 3rd attempt 20 minutes (2 x 10), and so on. The default value is **1** for the **Cluster Interface** and **2** for the **Data Interface** and **System**.

Step 8 Configure monitored interfaces by moving interfaces in the **Monitored Interfaces** or **Unmonitored Interfaces** window. You can also check or uncheck **Enable Service Application Monitoring** to enable or disable monitoring of the Snort and disk-full processes.

Figure 22: Configure Monitored Interfaces

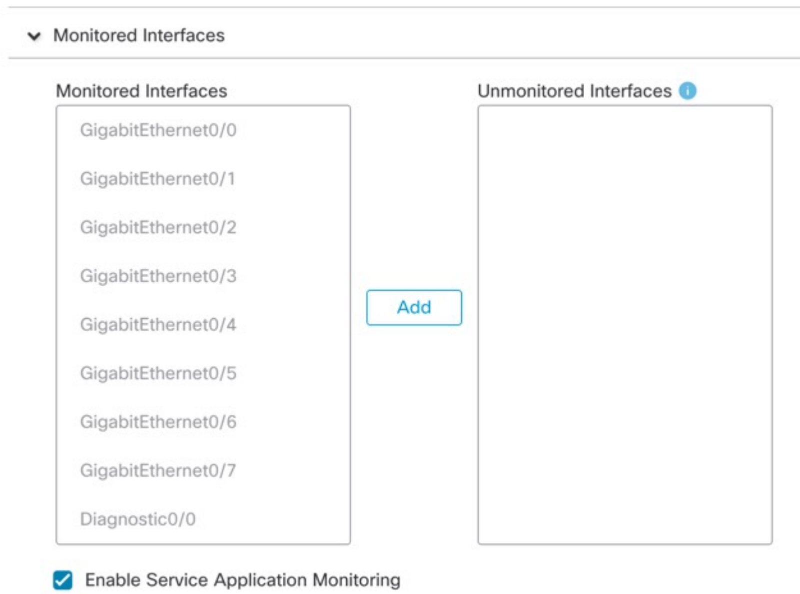

The interface health check monitors for link failures. If all physical ports for a given logical interface fail on a particular node, but there are active ports under the same logical interface on other nodes, then the node is removed from the cluster. The amount of time before the node removes a member from the cluster depends on the type of interface and whether the node is an established node or is joining the cluster. Health check is enabled by default for all interfaces and for the Snort and disk-full processes.

You might want to disable health monitoring of non-essential interfaces.

When any topology changes occur (such as adding or removing a data interface, enabling or disabling an interface on the node or the switch, or adding an additional switch to form a VSS or vPC) you should disable the system health check feature and also disable interface monitoring for the disabled interfaces. When the topology change is complete, and the configuration change is synced to all nodes, you can re-enable the system health check feature and monitored interfaces.

Step 9 Click **Save**.

Step 10 Deploy configuration changes.

Manage Cluster Nodes

•

Disable Clustering

You may want to deactivate a node in preparation for deleting the node, or temporarily for maintenance. This procedure is meant to temporarily deactivate a node; the node will still appear in the management center device list. When a node becomes inactive, all data interfaces are shut down.

> **Note** Do not power off the node without first disabling clustering.

Procedure

Step 1 For the unit you want to disable, choose **Devices** > **Device Management**, click the **More** (⋮), and choose **Disable Node Clustering**.

Step 2 Confirm that you want to disable clustering on the node.

The node will show **(Disabled)** next to its name in the **Devices** > **Device Management** list.

Step 3 To reenable clustering, see Rejoin the Cluster, on page 64.

Rejoin the Cluster

If a node was removed from the cluster, for example for a failed interface or if you manually disabled clustering, you must manually rejoin the cluster. Make sure the failure is resolved before you try to rejoin the cluster.

Procedure

Step 1 For the unit you want to reactivate, choose **Devices** > **Device Management**, click the **More** (⋮), and choose **Enable Node Clustering**.

Step 2 Confirm that you want to enable clustering on the node.

Reconcile Cluster Nodes

If a cluster node fails to register, you can reconcile the cluster membership from the device to the management center. For example, a data node might fail to register if the management center is occupied with certain processes, or if there is a network issue.

Procedure

Step 1 Choose **Devices** > **Device Management > More** (⋮) for the cluster, and then choose **Cluster Live Status** to open the **Cluster Status** dialog box.

Step 2 Click **Reconcile All**.

Figure 23: Reconcile All

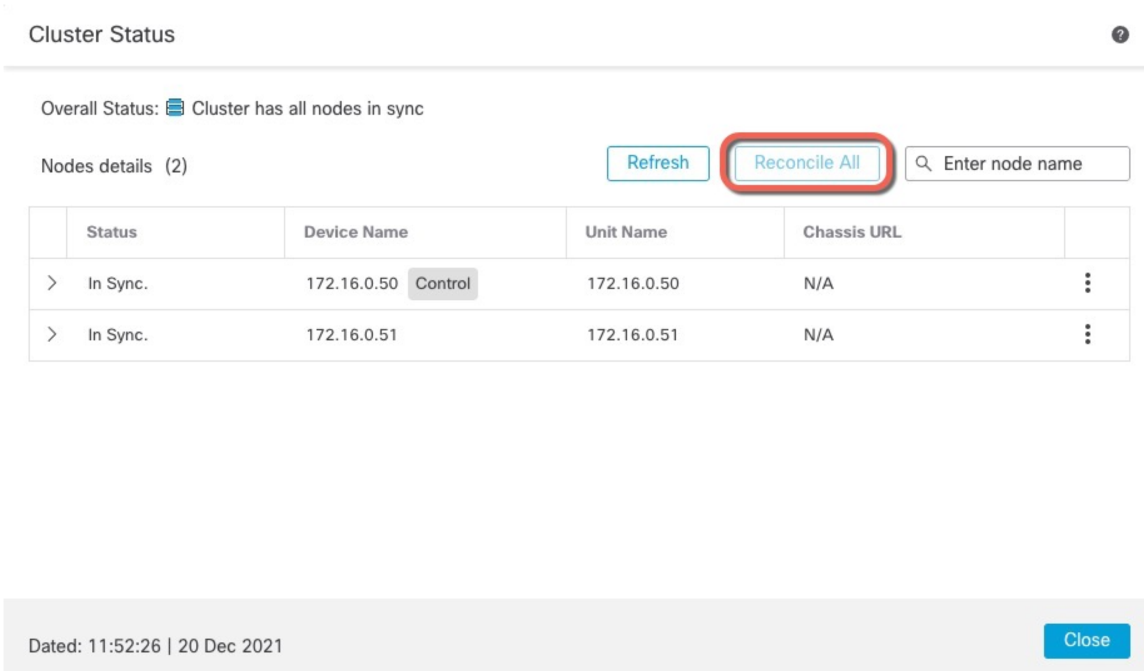

For more information about the cluster status, see Monitoring the Cluster, on page 66.

Unregister the Cluster or Nodes and Register to a New Management Center

You can unregister the cluster from the management center, which keeps the cluster intact. You might want to unregister the cluster if you want to add the cluster to a new management center.

You can also unregister a node from the management center without breaking the node from the cluster. Although the node is not visible in the management center, it is still part of the cluster, and it will continue to pass traffic and could even become the control node. You cannot unregister the current control node. You might want to unregister the node if it is no longer reachable from the management center, but you still want to keep it as part of the cluster while you troubleshoot management connectivity.

Unregistering a cluster:

- Severs all communication between the management center and the cluster.
- Removes the cluster from the **Device Management** page.
- Returns the cluster to local time management if the cluster's platform settings policy is configured to receive time from the management center using NTP.
- Leaves the configuration intact, so the cluster continues to process traffic.

  Policies, such as NAT and VPN, ACLs, and the interface configurations remain intact.

Registering the cluster again to the same or a different management center causes the configuration to be removed, so the cluster will stop processing traffic at that point; the cluster configuration remains intact so you can add the cluster as a whole. You can choose an access control policy at registration, but you will have to re-apply other policies after registration and then deploy the configuration before it will process traffic again.

Before you begin

This procedure requires CLI access to one of the nodes.

Procedure

Step 1 Choose **Devices** > **Device Management**, click **More** (⋮) for the cluster or node, and choose **Unregister**.

Step 2 You are prompted to unregister the cluster or node; click **Yes**.

Step 3 You can register the cluster to a new (or the same) management center by adding one of the cluster members as a new device.

You only need to add one of the cluster nodes as a device, and the rest of the cluster nodes will be discovered.

Step 4 To re-add a deleted node, see Reconcile Cluster Nodes, on page 64.

Monitoring the Cluster

You can monitor the cluster in the management center and at the threat defense CLI.

- **Cluster Status** dialog box, which is available from the **Devices** > **Device Management > More** (⋮) icon or from the **Devices** > **Device Management** > **Cluster** page **> General** area **> Cluster Live Status** link.

Figure 24: Cluster Status

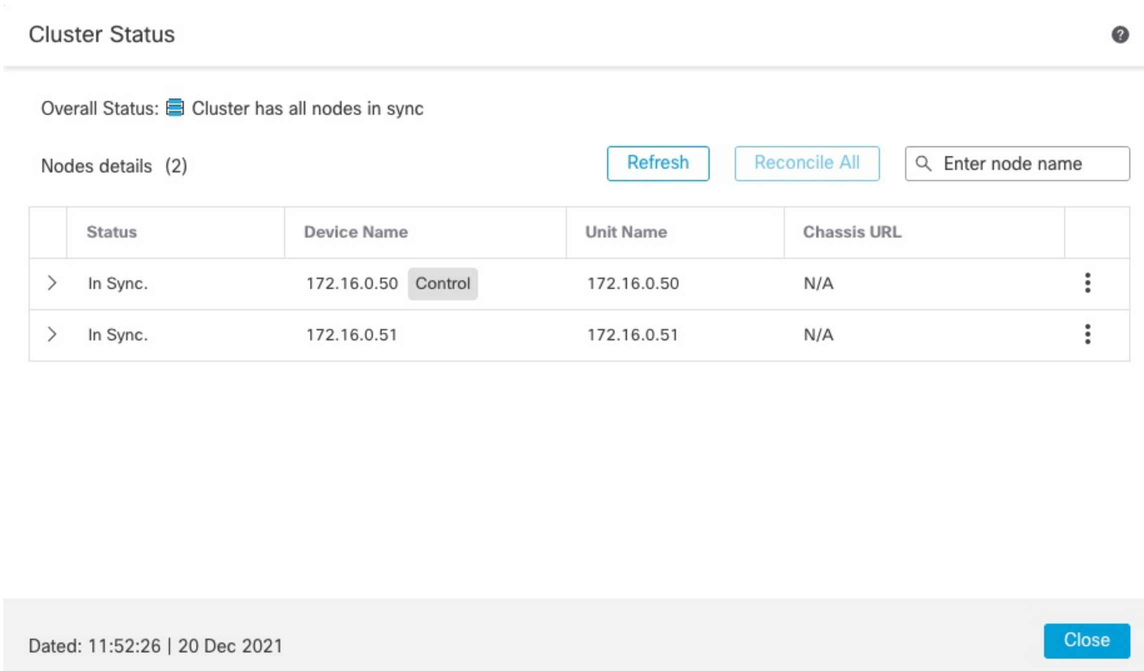

The Control node has a graphic indicator identifying its role.

Cluster member **Status** includes the following states:

- In Sync.—The node is registered with the management center.
- Pending Registration—The node is part of the cluster, but has not yet registered with the management center. If a node fails to register, you can retry registration by clicking **Reconcile All**.
- Clustering is disabled—The node is registered with the management center, but is an inactive member of the cluster. The clustering configuration remains intact if you intend to later re-enable it, or you can delete the node from the cluster.
- Joining cluster...—The node is joining the cluster on the chassis, but has not completed joining. After it joins, it will register with the management center.

For each node, you can view the **Summary** or the **History**.

Figure 25: Node Summary

Figure 26: Node History

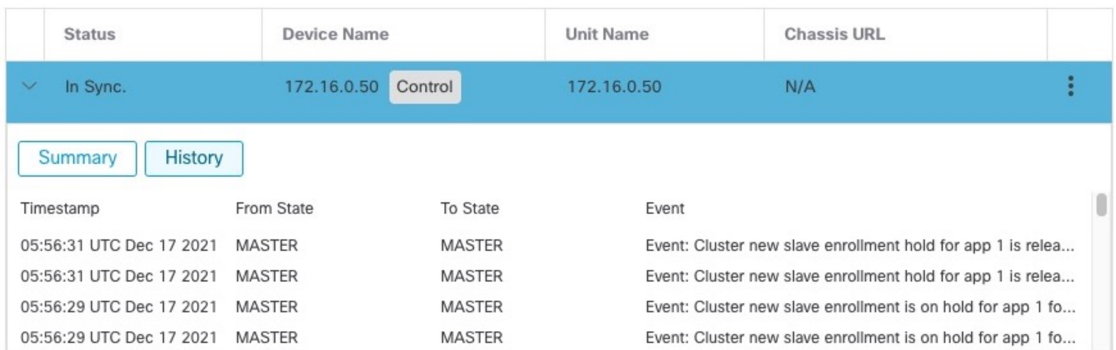

- **System** ( ) > **Tasks** page.

  The **Tasks** page shows updates of the Cluster Registration task as each node registers.

- **Devices** > **Device Management** > *cluster_name*.

  When you expand the cluster on the devices listing page, you can see all member nodes, including the control node shown with its role next to the IP address. For nodes that are still registering, you can see the loading icon.

- **show cluster** {**access-list** [*acl_name*] | **conn** [**count**] | **cpu** [**usage**] | **history** | **interface-mode** | **memory** | **resource usage** | **service-policy** | **traffic** | **xlate count**}

  To view aggregated data for the entire cluster or other information, use the **show cluster** command.

- **show cluster info** [**auto-join** | **clients** | **conn-distribution** | **flow-mobility counters** | **goid** [*options*] | **health** | **incompatible-config** | **loadbalance** | **old-members** | **packet-distribution** | **trace** [*options*] | **transport** { **asp** | **cp**}]

  To view cluster information, use the **show cluster info** command.

Cluster Health Monitor Dashboard

Cluster Health Monitor

When a threat defense is the control node of a cluster, the management center collects various metrics periodically from the device metric data collector. The cluster health monitor is comprised of the following components:

- Overview dashboard—Displays information about the cluster topology, cluster statistics, and metric charts:
  - The topology section displays a cluster's live status, the health of individual threat defense, threat defense node type (control node or data node), and the status of the device. The status of the device could be *Disabled* (when the device leaves the cluster), *Added out of box* (in a public cloud cluster, the additional nodes that do not belong to the management center), or *Normal* (ideal state of the node).
  - The cluster statistics section displays current metrics of the cluster with respect to the CPU usage, memory usage, input rate, output rate, active connections, and NAT translations.

Note The CPU and memory metrics display the individual average of the data plane and snort usage.

  - The metric charts, namely, CPU Usage, Memory Usage, Throughput, and Connections, diagrammatically display the statistics of the cluster over the specified time period.
- Load Distribution dashboard—Displays load distribution across the cluster nodes in two widgets:
  - The Distribution widget displays the average packet and connection distribution over the time range across the cluster nodes. This data depicts how the load is being distributed by the nodes. Using this widget, you can easily identify any abnormalities in the load distribution and rectify it.
  - The Node Statistics widget displays the node level metrics in table format. It displays metric data on CPU usage, memory usage, input rate, output rate, active connections, and NAT translations across the cluster nodes. This table view enables you to correlate data and easily identify any discrepancies.
- Member Performance dashboard—Displays current metrics of the cluster nodes. You can use the selector to filter the nodes and view the details of a specific node. The metric data include CPU usage, memory usage, input rate, output rate, active connections, and NAT translations.
- CCL dashboard—Displays, graphically, the cluster control link data namely, the input, and output rate.
- Troubleshooting and Links — Provides convenient links to frequently used troubleshooting topics and procedures.
- Time range—An adjustable time window to constrain the information that appears in the various cluster metrics dashboards and widgets.
- Custom Dashboard—Displays data on both cluster-wide metrics and node-level metrics. However, node selection only applies for the threat defense metrics and not for the entire cluster to which the node belongs.

Viewing Cluster Health

You must be an Admin, Maintenance, or Security Analyst user to perform this procedure.

The cluster health monitor provides a detailed view of the health status of a cluster and its nodes. This cluster health monitor provides health status and trends of the cluster in an array of dashboards.

Before you begin

- Ensure you have created a cluster from one or more devices in the management center.

Procedure

Step 1 Choose **System** (⚙) > **Health** > **Monitor**.

Use the Monitoring navigation pane to access node-specific health monitors.

Step 2 In the device list, click **Expand** (>) and **Collapse** (∨) to expand and collapse the list of managed cluster devices.

Step 3 To view the cluster health statistics, click on the cluster name. The cluster monitor reports health and performance metrics in several predefined dashboards by default. The metrics dashboards include:

- Overview — Highlights key metrics from the other predefined dashboards, including its nodes, CPU, memory, input and output rates, connection statistics, and NAT translation information.
- Load Distribution — Traffic and packet distribution across the cluster nodes.
- Member Performance — Node-level statistics on CPU usage, memory usage, input throughput, output throughput, active connection, and NAT translation.
- CCL — Interface status and aggregate traffic statistics.

You can navigate through the various metrics dashboards by clicking on the labels. For a comprehensive list of the supported cluster metrics, see Cisco Secure Firewall Threat Defense Health Metrics.

Step 4 You can configure the time range from the drop-down in the upper-right corner. The time range can reflect a period as short as the last hour (the default) or as long as two weeks. Select **Custom** from the drop-down to configure a custom start and end date.

Click the refresh icon to set auto refresh to 5 minutes or to toggle off auto refresh.

Step 5 Click on deployment icon for a deployment overlay on the trend graph, with respect to the selected time range.

The deployment icon indicates the number of deployments during the selected time-range. A vertical band indicates the deployment start and end time. For multiple deployments, multiple bands/lines appear. Click on the icon on top of the dotted line to view the deployment details.

Step 6 (For node-specific health monitor) View the **Health Alerts** for the node in the alert notification at the top of page, directly to the right of the device name.

Hover your pointer over the **Health Alerts** to view the health summary of the node. The popup window shows a truncated summary of the top five health alerts. Click on the popup to open a detailed view of the health alert summary.

Step 7 (For node-specific health monitor) The device monitor reports health and performance metrics in several predefined dashboards by default. The metrics dashboards include:

- Overview — Highlights key metrics from the other predefined dashboards, including CPU, memory, interfaces, connection statistics; plus disk usage and critical process information.
- CPU — CPU utilization, including the CPU usage by process and by physical cores.
- Memory — Device memory utilization, including data plane and Snort memory usage.
- Interfaces — Interface status and aggregate traffic statistics.
- Connections — Connection statistics (such as elephant flows, active connections, peak connections, and so on) and NAT translation counts.
- Snort — Statistics that are related to the Snort process.
- ASP drops — Statistics related to the dropped packets against various reasons.

You can navigate through the various metrics dashboards by clicking on the labels. See Cisco Secure Firewall Threat Defense Health Metrics for a comprehensive list of the supported device metrics.

Step 8 Click the plus sign (+) in the upper right corner of the health monitor to create a custom dashboard by building your own variable set from the available metric groups.

For cluster-wide dashboard, choose Cluster metric group, and then choose the metric.

Cluster Metrics

The cluster health monitor tracks statistics that are related to a cluster and its nodes, and aggregate of load distribution, performance, and CCL traffic statistics.

Table 6: Cluster Metrics

Metric	Description	Format
CPU	Average of CPU metrics on the nodes of a cluster (individually for data plane and snort).	percentage
Memory	Average of memory metrics on the nodes of a cluster (individually for data plane and snort).	percentage
Data Throughput	Incoming and outgoing data traffic statistics for a cluster.	bytes
CCL Throughput	Incoming and outgoing CCL traffic statistics for a cluster.	bytes
Connections	Count of active connections in a cluster.	number
NAT Translations	Count of NAT translations for a cluster.	number
Distribution	Connection distribution count in the cluster for every second.	number

Metric	Description	Format
Packets	Packet distribution count in the cluster for every second.	number

Troubleshooting the Cluster

You can use the **CCL Ping** tool to make sure the cluster control link is operating correctly. You can also use the following tools that are available for devices and clusters:

- Troubleshooting files—If a node fails to join the cluster, a troubleshooting file is automatically generated. You can also generate and download troubleshooting files from the **Devices** > **Device Management** > **Cluster** > **General** area. See Generate Troubleshooting Files.

  You can also generate files from the **Device Management** page by clicking **More** (⋮) and choosing **Troubleshoot Files**.

- CLI output—From the **Devices** > **Device Management** > **Cluster** > **General** area, you can view a set of pre-defined CLI outputs that can help you troubleshoot the cluster. The following commands are automatically run for the cluster:

  - **show running-config cluster**
  - **show cluster info**
  - **show cluster info health**
  - **show cluster info transport cp**
  - **show version**
  - **show asp drop**
  - **show counters**
  - **show arp**
  - **show int ip brief**
  - **show blocks**
  - **show cpu detailed**
  - **show interface** *ccl_interface*
  - **ping** *ccl_ip* **size** *ccl_mtu* **repeat 2**

  You can also enter any **show** command in the Command field. See View CLI Output for more information.

Perform a Ping on the Cluster Control Link

You can check to make sure all the cluster nodes can reach each other over the cluster control link by performing a ping. One major cause for the failure of a node to join the cluster is an incorrect cluster control link configuration; for example, the cluster control link MTU may be set higher than the connecting switch MTUs.

Procedure

Step 1 Choose **Devices** > **Device Management**, click the **More** (⋮) icon next to the cluster, and choose **> Cluster Live Status**.

Figure 27: Cluster Status

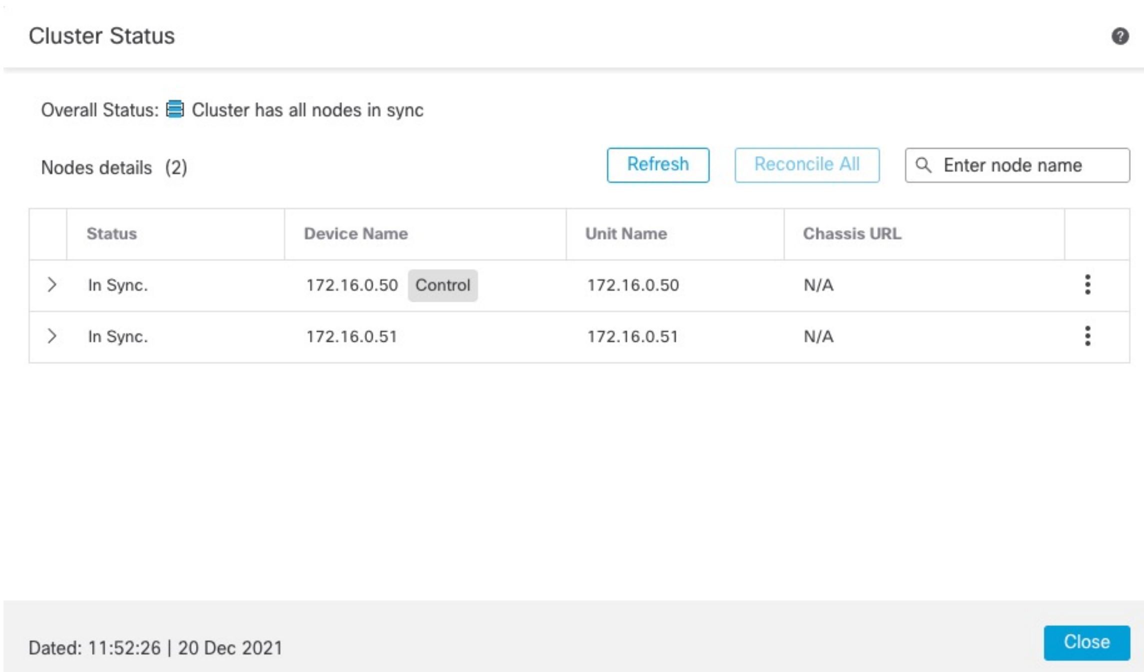

Step 2 Expand one of the nodes, and click **CCL Ping**.

Figure 28: CCL Ping

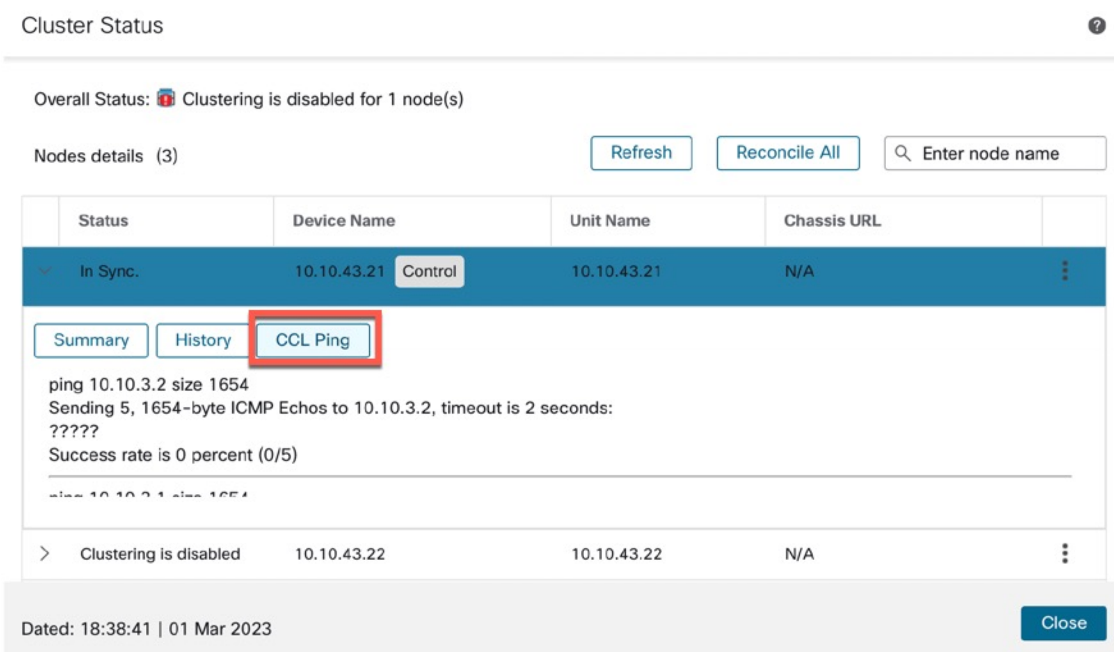

The node sends a ping on the cluster control link to every other node using a packet size that matches the maximum MTU.

Upgrading the Cluster

Perform the following steps to upgrade a threat defense virtual cluster:

Procedure

Step 1 Upload the target image version to the cloud image storage.

Step 2 Update the cloud instance template of the cluster with the updated target image version.

a) Create a copy of the instance template with the target image version.

b) Attach the newly created template to cluster instance group.

Step 3 Upload the target image version upgrade package to the management center.

Step 4 Perform readiness check on the cluster that you want to upgrade.

Step 5 After successful readiness check, initiate installation of upgrade package.

Step 6 The management center upgrades the cluster nodes one at a time.

Step 7 The management center displays a notification after successful upgrade of the cluster.

There is no change in the serial number and UUID of the instance after the upgrade.

Note

- If you initiate the cluster upgrade from the management center, ensure that no threat defense virtual device is accidentally terminated or replaced by the auto scaling group during the post-upgrade reboot process. To prevent this, go to the AWS console, click **Auto scaling group -> Advanced configurations**, and suspend the processes - Health Check and Replace Unhealthy. After the upgrade is completed, go to **Advanced configurations** again and remove any suspended processes to detect unhealthy instances.
- If you upgrade a cluster deployed on AWS from a major release to a patch release and then scale up the cluster, the new nodes will come up with the major release version instead of the patch release. You have to then manually upgrade each node to the patch release from the management center.

  Alternatively, you can also create an Amazon Machine Image (AMI) from a snapshot of a standalone threat defense virtual instance on which the patch has been applied and which does not have a day 0 configuration. Use this AMI in the cluster deployment template. Any new nodes that come up when you scale up the cluster will have the patch release.

Reference for Clustering

This section includes more information about how clustering operates.

Threat Defense Features and Clustering

Some threat defense features are not supported with clustering, and some are only supported on the control unit. Other features might have caveats for proper usage.

Unsupported Features and Clustering

These features cannot be configured with clustering enabled, and the commands will be rejected.

Note To view FlexConfig features that are also not supported with clustering, for example WCCP inspection, see the ASA general operations configuration guide. FlexConfig lets you configure many ASA features that are not present in the management center GUI. See FlexConfig Policies.

- Remote access VPN (SSL VPN and IPsec VPN)
- DHCP client, server, and proxy. DHCP relay is supported.
- Virtual Tunnel Interfaces (VTIs)
- High Availability
- Integrated Routing and Bridging
- Management Center UCAPL/CC mode

Centralized Features for Clustering

The following features are only supported on the control node, and are not scaled for the cluster.

Note Traffic for centralized features is forwarded from member nodes to the control node over the cluster control link.

If you use the rebalancing feature, traffic for centralized features may be rebalanced to non-control nodes before the traffic is classified as a centralized feature; if this occurs, the traffic is then sent back to the control node.

For centralized features, if the control node fails, all connections are dropped, and you have to re-establish the connections on the new control node.

Note To view FlexConfig features that are also centralized with clustering, for example RADIUS inspection, see the ASA general operations configuration guide. FlexConfig lets you configure many ASA features that are not present in the management center GUI. See FlexConfig Policies.

- The following application inspections:
  - DCERPC
  - ESMTP
  - NetBIOS
  - PPTP
  - RSH
  - SQLNET
  - SUNRPC
  - TFTP
  - XDMCP
- Static route monitoring

Cisco Trustsec and Clustering

Only the control node learns security group tag (SGT) information. The control node then populates the SGT to data nodes, and data nodes can make a match decision for SGT based on the security policy.

Connection Settings and Clustering

Connection limits are enforced cluster-wide. Each node has an estimate of the cluster-wide counter values based on broadcast messages. Due to efficiency considerations, the configured connection limit across the cluster might not be enforced exactly at the limit number. Each node may overestimate or underestimate the cluster-wide counter value at any given time. However, the information will get updated over time in a load-balanced cluster.

Dynamic Routing and Clustering

In Individual interface mode, each node runs the routing protocol as a standalone router, and routes are learned by each node independently.

Figure 29: Dynamic Routing in Individual Interface Mode

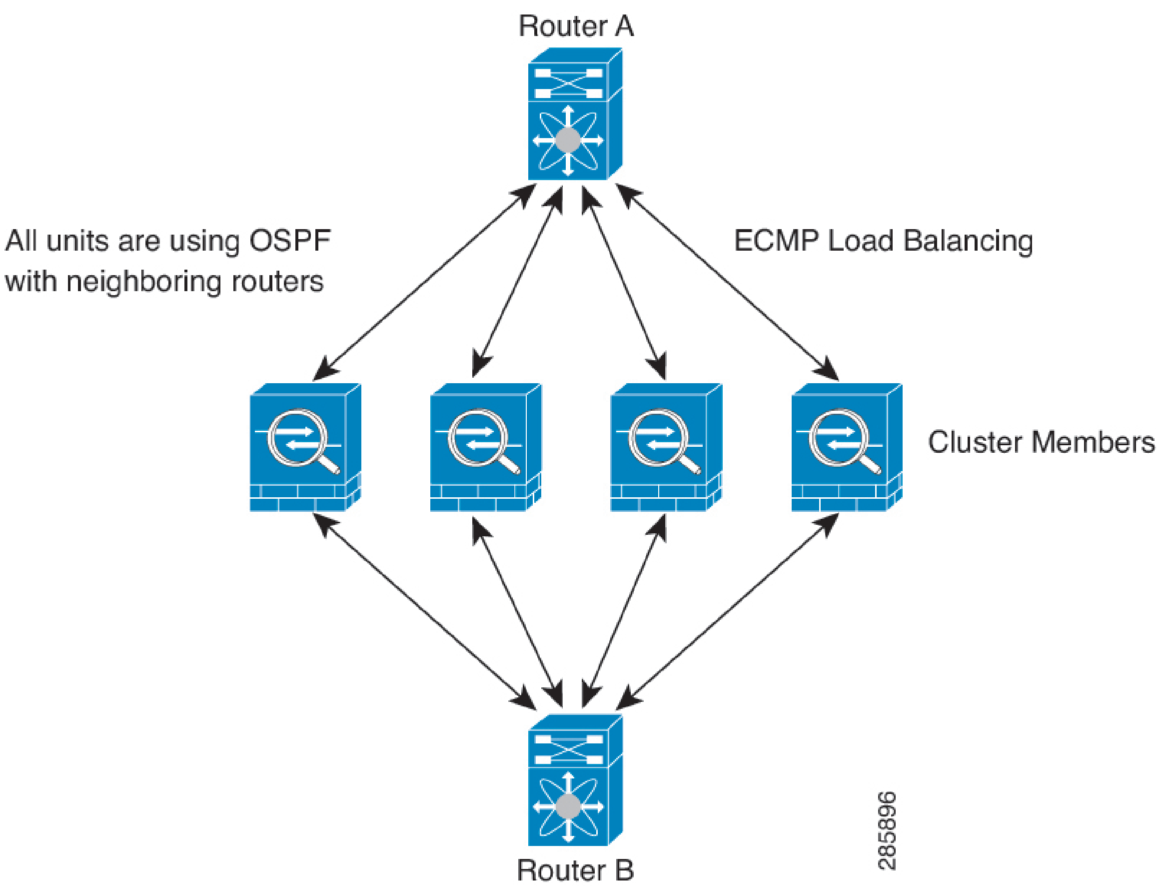

In the above diagram, Router A learns that there are 4 equal-cost paths to Router B, each through a node. ECMP is used to load balance traffic between the 4 paths. Each node picks a different router ID when talking to external routers.

You must configure a cluster pool for the router ID so that each node has a separate router ID.

FTP and Clustering

- If FTP data channel and control channel flows are owned by different cluster members, then the data channel owner will periodically send idle timeout updates to the control channel owner and update the idle timeout value. However, if the control flow owner is reloaded, and the control flow is re-hosted, the parent/child flow relationship will not longer be maintained; the control flow idle timeout will not be updated.

NAT and Clustering

For NAT usage, see the following limitations.

NAT can affect the overall throughput of the cluster. Inbound and outbound NAT packets can be sent to different threat defenses in the cluster, because the load balancing algorithm relies on IP addresses and ports, and NAT causes inbound and outbound packets to have different IP addresses and/or ports. When a packet arrives at the threat defense that is not the NAT owner, it is forwarded over the cluster control link to the owner, causing large amounts of traffic on the cluster control link. Note that the receiving node does not create a forwarding flow to the owner, because the NAT owner may not end up creating a connection for the packet depending on the results of security and policy checks.

If you still want to use NAT in clustering, then consider the following guidelines:

- No Proxy ARP—For Individual interfaces, a proxy ARP reply is never sent for mapped addresses. This prevents the adjacent router from maintaining a peer relationship with an ASA that may no longer be in the cluster. The upstream router needs a static route or PBR with Object Tracking for the mapped addresses that points to the Main cluster IP address.
- PAT with Port Block Allocation—See the following guidelines for this feature:
  - Maximum-per-host limit is not a cluster-wide limit, and is enforced on each node individually. Thus, in a 3-node cluster with the maximum-per-host limit configured as 1, if the traffic from a host is load-balanced across all 3 nodes, then it can get allocated 3 blocks with 1 in each node.
  - Port blocks created on the backup node from the backup pools are not accounted for when enforcing the maximum-per-host limit.
  - On-the-fly PAT rule modifications, where the PAT pool is modified with a completely new range of IP addresses, will result in xlate backup creation failures for the xlate backup requests that were still in transit while the new pool became effective. This behavior is not specific to the port block allocation feature, and is a transient PAT pool issue seen only in cluster deployments where the pool is distributed and traffic is load-balanced across the cluster nodes.
  - When operating in a cluster, you cannot simply change the block allocation size. The new size is effective only after you reload each device in the cluster. To avoid having to reload each device, we recommend that you delete all block allocation rules and clear all xlates related to those rules. You can then change the block size and recreate the block allocation rules.
- NAT pool address distribution for dynamic PAT—When you configure a PAT pool, the cluster divides each IP address in the pool into port blocks. By default, each block is 512 ports, but if you configure port block allocation rules, your block setting is used instead. These blocks are distributed evenly among the nodes in the cluster, so that each node has one or more blocks for each IP address in the PAT pool. Thus, you could have as few as one IP address in a PAT pool for a cluster, if that is sufficient for the number of PAT'ed connections you expect. Port blocks cover the 1024-65535 port range, unless you configure the option to include the reserved ports, 1-1023, on the PAT pool NAT rule.
- Reusing a PAT pool in multiple rules—To use the same PAT pool in multiple rules, you must be careful about the interface selection in the rules. You must either use specific interfaces in all rules, or "any" in all rules. You cannot mix specific interfaces and "any" across the rules, or the system might not be able to match return traffic to the right node in the cluster. Using unique PAT pools per rule is the most reliable option.
- No round-robin—Round-robin for a PAT pool is not supported with clustering.
- No extended PAT—Extended PAT is not supported with clustering.

- Dynamic NAT xlates managed by the control node—The control node maintains and replicates the xlate table to data nodes. When a data node receives a connection that requires dynamic NAT, and the xlate is not in the table, it requests the xlate from the control node. The data node owns the connection.
- Stale xlates—The xlate idle time on the connection owner does not get updated. Thus, the idle time might exceed the idle timeout. An idle timer value higher than the configured timeout with a refcnt of 0 is an indication of a stale xlate.
- No static PAT for the following inspections—
    - FTP
    - RSH
    - SQLNET
    - TFTP
    - XDMCP
    - SIP
- If you have an extremely large number of NAT rules, over ten thousand, you should enable the transactional commit model using the **asp rule-engine transactional-commit nat** command in the device CLI. Otherwise, the node might not be able to join the cluster.

SIP Inspection and Clustering

A control flow can be created on any node (due to load balancing); its child data flows must reside on the same node.

SNMP and Clustering

You should always use the Local address, and not the Main cluster IP address for SNMP polling. If the SNMP agent polls the Main cluster IP address, if a new control node is elected, the poll to the new control node will fail.

When using SNMPv3 with clustering, if you add a new cluster node after the initial cluster formation, then SNMPv3 users are not replicated to the new node. You must remove the users, and re-add them, and then redeploy your configuration to force the users to replicate to the new node.

Syslog and Clustering

- Each node in the cluster generates its own syslog messages. You can configure logging so that each node uses either the same or a different device ID in the syslog message header field. For example, the hostname configuration is replicated and shared by all nodes in the cluster. If you configure logging to use the hostname as the device ID, syslog messages generated by all nodes look as if they come from a single node. If you configure logging to use the local-node name that is assigned in the cluster bootstrap configuration as the device ID, syslog messages look as if they come from different nodes.

VPN and Clustering

Site-to-site VPN is a centralized feature; only the control node supports VPN connections.

Note Remote access VPN is not supported with clustering.

VPN functionality is limited to the control node and does not take advantage of the cluster high availability capabilities. If the control node fails, all existing VPN connections are lost, and VPN users will see a disruption in service. When a new control node is elected, you must reestablish the VPN connections.

For connections to an Individual interface when using PBR or ECMP, you must always connect to the Main cluster IP address, not a Local address.

VPN-related keys and certificates are replicated to all nodes.

Performance Scaling Factor

When you combine multiple units into a cluster, you can expect the total cluster performance to be approximately 80% of the maximum combined throughput.

For example, if your model can handle approximately 10 Gbps of traffic when running alone, then for a cluster of 8 units, the maximum combined throughput will be approximately 80% of 80 Gbps (8 units x 10 Gbps): 64 Gbps.

Control Node Election

Nodes of the cluster communicate over the cluster control link to elect a control node as follows:

1. When you enable clustering for a node (or when it first starts up with clustering already enabled), it broadcasts an election request every 3 seconds.
2. Any other nodes with a higher priority respond to the election request; the priority is set between 1 and 100, where 1 is the highest priority.
3. If after 45 seconds, a node does not receive a response from another node with a higher priority, then it becomes the control node.

   Note If multiple nodes tie for the highest priority, the cluster node name and then the serial number is used to determine the control node.

4. If a node later joins the cluster with a higher priority, it does not automatically become the control node; the existing control node always remains as the control node unless it stops responding, at which point a new control node is elected.
5. In a "split brain" scenario when there are temporarily multiple control nodes, then the node with highest priority retains the role while the other nodes return to data node roles.

Note You can manually force a node to become the control node. For centralized features, if you force a control node change, then all connections are dropped, and you have to re-establish the connections on the new control node.

High Availability within the Cluster

Clustering provides high availability by monitoring node and interface health and by replicating connection states between nodes.

Node Health Monitoring

Each node periodically sends a broadcast heartbeat packet over the cluster control link. If the control node does not receive any heartbeat packets or other packets from a data node within the configurable timeout period, then the control node removes the data node from the cluster. If the data nodes do not receive packets from the control node, then a new control node is elected from the remaining nodes.

If nodes cannot reach each other over the cluster control link because of a network failure and not because a node has actually failed, then the cluster may go into a "split brain" scenario where isolated data nodes will elect their own control nodes. For example, if a router fails between two cluster locations, then the original control node at location 1 will remove the location 2 data nodes from the cluster. Meanwhile, the nodes at location 2 will elect their own control node and form their own cluster. Note that asymmetric traffic may fail in this scenario. After the cluster control link is restored, then the control node that has the higher priority will keep the control node's role.

Interface Monitoring

Each node monitors the link status of all named hardware interfaces in use, and reports status changes to the control node.

All physical interfaces are monitored; only named interfaces can be monitored. You can optionally disable monitoring per interface.

A node is removed from the cluster if its monitored interfaces fail. The node is removed after 500 ms.

Status After Failure

If the control node fails, then another member of the cluster with the highest priority (lowest number) becomes the control node.

The Threat Defense automatically tries to rejoin the cluster, depending on the failure event.

Note When the Threat Defense becomes inactive and fails to automatically rejoin the cluster, all data interfaces are shut down; only the Management interface can send and receive traffic.

Rejoining the Cluster

After a cluster member is removed from the cluster, how it can rejoin the cluster depends on why it was removed:

- Failed cluster control link when initially joining—After you resolve the problem with the cluster control link, you must manually rejoin the cluster by re-enabling clustering.
- Failed cluster control link after joining the cluster—The threat defense automatically tries to rejoin every 5 minutes, indefinitely.
- Failed data interface—The threat defense automatically tries to rejoin at 5 minutes, then at 10 minutes, and finally at 20 minutes. If the join is not successful after 20 minutes, then the threat defense application

disables clustering. After you resolve the problem with the data interface, you have to manually enable clustering.

- Failed node—If the node was removed from the cluster because of a node health check failure, then rejoining the cluster depends on the source of the failure. For example, a temporary power failure means the node will rejoin the cluster when it starts up again as long as the cluster control link is up. The threat defense application attempts to rejoin the cluster every 5 seconds.
- Internal error—Internal failures include: application sync timeout; inconsistent application statuses; and so on.
- Failed configuration deployment—If you deploy a new configuration from management center, and the deployment fails on some cluster members but succeeds on others, then the nodes that failed are removed from the cluster. You must manually rejoin the cluster by re-enabling clustering. If the deployment fails on the control node, then the deployment is rolled back, and no members are removed. If the deployment fails on all data nodes, then the deployment is rolled back, and no members are removed.

Data Path Connection State Replication

Every connection has one owner and at least one backup owner in the cluster. The backup owner does not take over the connection in the event of a failure; instead, it stores TCP/UDP state information, so that the connection can be seamlessly transferred to a new owner in case of a failure. The backup owner is usually also the director.

Some traffic requires state information above the TCP or UDP layer. See the following table for clustering support or lack of support for this kind of traffic.

Table 7: Features Replicated Across the Cluster

Traffic	State Support	Notes
Up time	Yes	Keeps track of the system up time.
ARP Table	Yes	—
MAC address table	Yes	—
User Identity	Yes	—
IPv6 Neighbor database	Yes	—
Dynamic routing	Yes	—
SNMP Engine ID	**No**	—

How the Cluster Manages Connections

Connections can be load-balanced to multiple nodes of the cluster. Connection roles determine how connections are handled in both normal operation and in a high availability situation.

Connection Roles

See the following roles defined for each connection:

- Owner—Usually, the node that initially receives the connection. The owner maintains the TCP state and processes packets. A connection has only one owner. If the original owner fails, then when new nodes receive packets from the connection, the director chooses a new owner from those nodes.
- Backup owner—The node that stores TCP/UDP state information received from the owner, so that the connection can be seamlessly transferred to a new owner in case of a failure. The backup owner does not take over the connection in the event of a failure. If the owner becomes unavailable, then the first node to receive packets from the connection (based on load balancing) contacts the backup owner for the relevant state information so it can become the new owner.

  As long as the director (see below) is not the same node as the owner, then the director is also the backup owner. If the owner chooses itself as the director, then a separate backup owner is chosen.

  For clustering on the Firepower 9300, which can include up to 3 cluster nodes in one chassis, if the backup owner is on the same chassis as the owner, then an additional backup owner will be chosen from another chassis to protect flows from a chassis failure.
- Director—The node that handles owner lookup requests from forwarders. When the owner receives a new connection, it chooses a director based on a hash of the source/destination IP address and ports (see below for ICMP hash details), and sends a message to the director to register the new connection. If packets arrive at any node other than the owner, the node queries the director about which node is the owner so it can forward the packets. A connection has only one director. If a director fails, the owner chooses a new director.

  As long as the director is not the same node as the owner, then the director is also the backup owner (see above). If the owner chooses itself as the director, then a separate backup owner is chosen.

  ICMP/ICMPv6 hash details:

  - For Echo packets, the source port is the ICMP identifier, and the destination port is 0.
  - For Reply packets, the source port is 0, and the destination port is the ICMP identifier.
  - For other packets, both source and destination ports are 0.
- Forwarder—A node that forwards packets to the owner. If a forwarder receives a packet for a connection it does not own, it queries the director for the owner, and then establishes a flow to the owner for any other packets it receives for this connection. The director can also be a forwarder. Note that if a forwarder receives the SYN-ACK packet, it can derive the owner directly from a SYN cookie in the packet, so it does not need to query the director. (If you disable TCP sequence randomization, the SYN cookie is not used; a query to the director is required.) For short-lived flows such as DNS and ICMP, instead of querying, the forwarder immediately sends the packet to the director, which then sends them to the owner. A connection can have multiple forwarders; the most efficient throughput is achieved by a good load-balancing method where there are no forwarders and all packets of a connection are received by the owner.

Note We do not recommend disabling TCP sequence randomization when using clustering. There is a small chance that some TCP sessions won't be established, because the SYN/ACK packet might be dropped.

- Fragment Owner—For fragmented packets, cluster nodes that receive a fragment determine a fragment owner using a hash of the fragment source IP address, destination IP address, and the packet ID. All fragments are then forwarded to the fragment owner over the cluster control link. Fragments may be load-balanced to different cluster nodes, because only the first fragment includes the 5-tuple used in the switch load balance hash. Other fragments do not contain the source and destination ports and may be load-balanced to other cluster nodes. The fragment owner temporarily reassembles the packet so it can determine the director based on a hash of the source/destination IP address and ports. If it is a new connection, the fragment owner will register to be the connection owner. If it is an existing connection, the fragment owner forwards all fragments to the provided connection owner over the cluster control link. The connection owner will then reassemble all fragments.

New Connection Ownership

When a new connection is directed to a node of the cluster via load balancing, that node owns both directions of the connection. If any connection packets arrive at a different node, they are forwarded to the owner node over the cluster control link. If a reverse flow arrives at a different node, it is redirected back to the original node.

Sample Data Flow for TCP

The following example shows the establishment of a new connection.

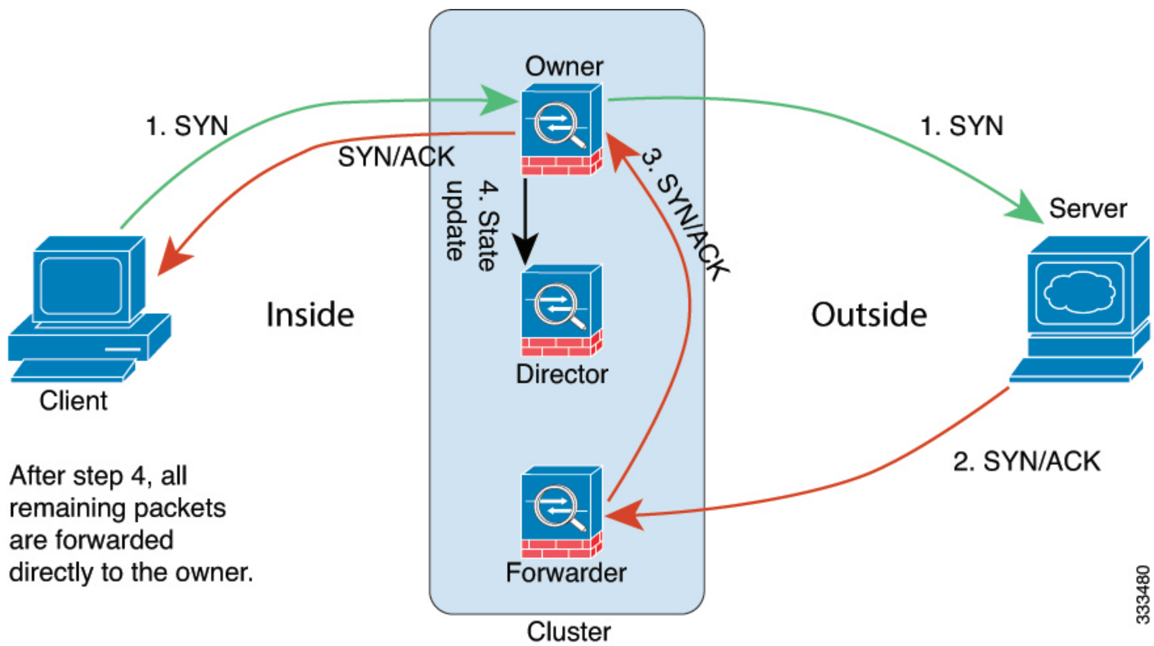

1. The SYN packet originates from the client and is delivered to one threat defense (based on the load balancing method), which becomes the owner. The owner creates a flow, encodes owner information into a SYN cookie, and forwards the packet to the server.
2. The SYN-ACK packet originates from the server and is delivered to a different threat defense (based on the load balancing method). This threat defense is the forwarder.
3. Because the forwarder does not own the connection, it decodes owner information from the SYN cookie, creates a forwarding flow to the owner, and forwards the SYN-ACK to the owner.
4. The owner sends a state update to the director, and forwards the SYN-ACK to the client.

5. The director receives the state update from the owner, creates a flow to the owner, and records the TCP state information as well as the owner. The director acts as the backup owner for the connection.
6. Any subsequent packets delivered to the forwarder will be forwarded to the owner.
7. If packets are delivered to any additional nodes, it will query the director for the owner and establish a flow.
8. Any state change for the flow results in a state update from the owner to the director.

Sample Data Flow for ICMP and UDP

The following example shows the establishment of a new connection.

1. ***Figure 30: ICMP and UDP Data Flow***

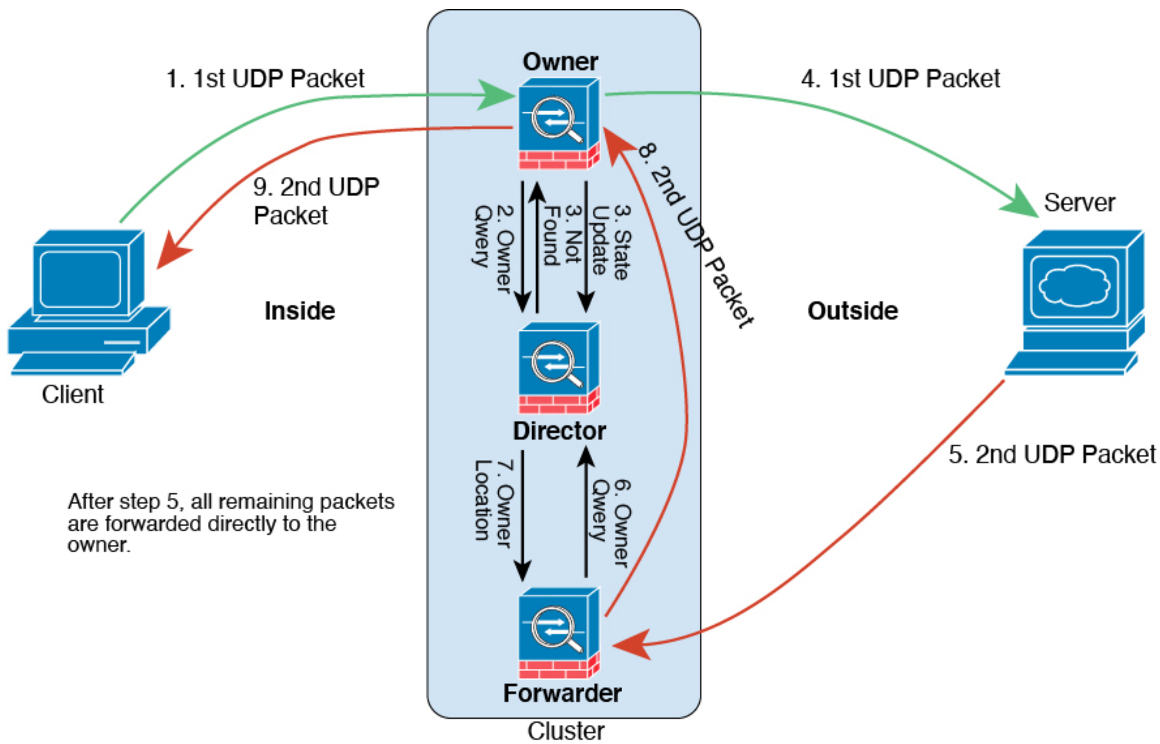

The first UDP packet originates from the client and is delivered to one threat defense (based on the load balancing method).

2. The node that received the first packet queries the director node that is chosen based on a hash of the source/destination IP address and ports.
3. The director finds no existing flow, creates a director flow and forwards the packet back to the previous node. In other words, the director has elected an owner for this flow.
4. The owner creates the flow, sends a state update to the director, and forwards the packet to the server.
5. The second UDP packet originates from the server and is delivered to the forwarder.
6. The forwarder queries the director for ownership information. For short-lived flows such as DNS, instead of querying, the forwarder immediately sends the packet to the director, which then sends it to the owner.

7. The director replies to the forwarder with ownership information.
8. The forwarder creates a forwarding flow to record owner information and forwards the packet to the owner.
9. The owner forwards the packet to the client.

History for Threat Defense Virtual Clustering in the Public Cloud

Feature	Minimum Management Center	Minimum Threat Defense	Details
Cluster control link ping tool.	20240203	Any	You can check to make sure all the cluster nodes can reach each other over the cluster control link by performing a ping. One major cause for the failure of a node to join the cluster is an incorrect cluster control link configuration; for example, the cluster control link MTU may be set higher than the connecting switch MTUs. New/modified screens: **Devices** > **Device Management > More** (⋮) **> Cluster Live Status**
Troubleshooting file generation and download available from Device and Cluster pages.	20240203	7.4.1	You can generate and download troubleshooting files for each device on the Device page and also for all cluster nodes on the Cluster page. For a cluster, you can download all files as a single compressed file. You can also include cluster logs for the cluster for cluster nodes. You can alternatively trigger file generation from the **Devices** > **Device Management > More** (⋮) **> Troubleshoot Files** menu. New/modified screens: • **Devices** > **Device Management** > **Device** > **General** • **Devices** > **Device Management** > **Cluster** > **General**
View CLI output for a device or device cluster.	20240203	Any	You can view a set of pre-defined CLI outputs that can help you troubleshoot the device or cluster. You can also enter any **show** command and see the output. New/modified screens: **Devices** > **Device Management** > **Cluster** > **General**
Cluster health monitor settings	20221213	Any	You can now edit cluster health monitor settings. New/Modified screens: **Devices** > **Device Management** > **Cluster** > **Cluster Health Monitor Settings** **Note** If you previously configured these settings using FlexConfig, be sure to remove the FlexConfig configuration before you deploy. Otherwise the FlexConfig configuration will overwrite the management center configuration.
Cluster health monitor dashboard	20221213	Any	You can now view cluster health on the cluster health monitor dashboard. New/Modified screens: **System** (⚙) > **Health** > **Monitor**

Feature	Minimum Management Center	Minimum Threat Defense	Details
Clustering for the threat defense virtual in Azure	20221213	7.3.0	You can now configure clustering for up to 16 nodes the threat defense virtual in Azure for the Azure Gateway Load Balancer or for external load balancers. New/modified screens: • **Devices** > **Device Management** > **Add Cluster** • **Devices** > **Device Management** > **More** menu • **Devices** > **Device Management** > **Cluster** Supported platforms: Threat Defense Virtual in Azure
Clustering for the Threat Defense Virtual in the Public Cloud (Amazon Web Services and Google Cloud Platform)	Any	7.2.0	The threat defense virtual supports Individual interface clustering for up to 16 nodes in the public cloud (AWS and GCP). New/Modified screens: • **Devices** > **Device Management** > **Add Device** • **Devices** > **Device Management** > **More** menu • **Devices** > **Device Management** > **Cluster** Supported platforms: Threat Defense Virtual in AWS and GCP