

# NSX-T Federation Adoption Content

NSX-T 3.1.x

NSX-T Federation

You can find the most up-to-date technical documentation on the VMware website at: <https://docs.vmware.com/>

NSX-T Federation Adoption

3401 Hillview  
Ave. Palo Alto,  
CA 94304  
[www.vmware.com](http://www.vmware.com)

Copyright © 2022 VMware, Inc. All rights reserved. Copyright and trademark information.

# Contents

NSX-T 3.1.x.....	1
Introduction .....	5
Federation Key Concepts .....	6
Features and Configurations Supported in NSX-T Federation .....	7
Understanding NSX-T Federation .....	11
Using the Global and Local Manager Web Interfaces .....	16
Overriding Global Manager Configurations on Local Manager .....	18
Getting Started with NSX-T Federation .....	22
Install the Active and Standby Global Manager .....	22
Make the Global Manager Active and Add Standby Global Manager .....	24
Add a Location .....	26
Networking .....	30
Tier-0 Gateway Configurations in Federation .....	32
Tier-1 Gateway Configurations in NSX-T Federation .....	35
Configure Edge Nodes for Stretched Networking .....	38
Add a Tier-0 Gateway from Global Manager .....	40
Add a Tier-1 Gateway from Global Manager .....	48
Add a Segment from Global Manager .....	54
Security .....	58
Security in Federation .....	58
Create a Region from Global Manager .....	65
Create Groups from Global Manager .....	67
Create DFW Policies and Rules from Global Manager .....	71
Create Gateway Policies and Rules from Global Manager .....	77
Backup and Restore .....	83
Backup and Restore in NSX-T Federation .....	83
Disaster Recovery .....	85
Disaster Recovery for Global Manager .....	86
Network Recovery .....	88
Network Recovery for Local Managers .....	88

# Maintenance Activities

This section contains standard maintenance procedures for the capabilities being deployed as a part of the service.

This chapter includes the following topics:

- Introduction
- Getting Started with NSX-T Federation
- Networking
- Security
- Backup and Restore
- Disaster Recovery
- Network Recovery

## Introduction

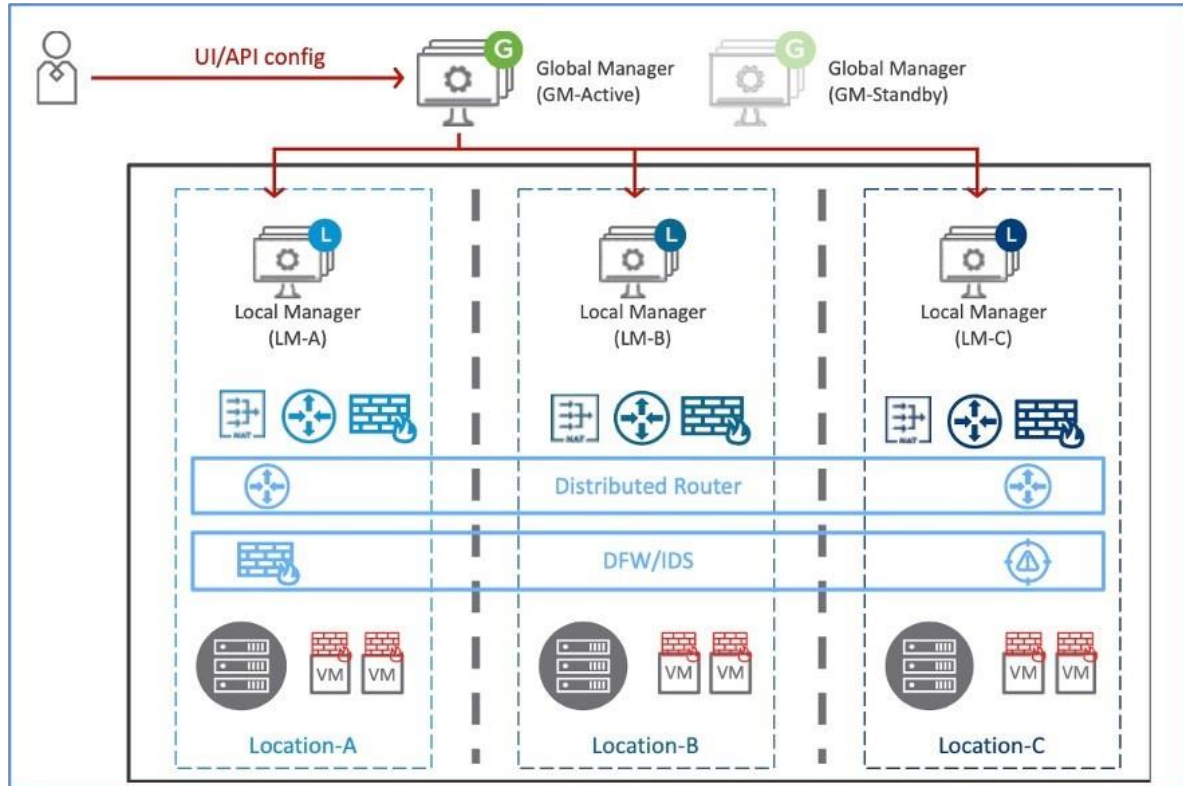
With NSX-T Federation, you can manage multiple NSX-T Data Center environments using an intuitive user interface, with a single pane of glass view.

You can create gateways and segments that span one or more locations and configure and enforce firewall rules consistently across locations.

NSX-T uses one central NSX-T Global Manager Cluster (GM) that offers central network and security services configuration for all locations:

- There is one NSX-T Manager Cluster per location, which we call the Local Manager (LM) that manages Transport Nodes (hypervisor and Edge nodes) for that location.
- The GM pushes the network and security configuration to the different LMs to implement locally.

Figure 1-1. NSX-T Federation



## Federation Key Concepts

Federation introduces some new terms and concepts, such as remote tunnel endpoint (RTEP), span, and region.

### Federation Systems: Global Manager and Local Manager

An Federation environment includes two types of management systems:

- **Global Manager:** a system similar to NSX Manager that federates multiple Local Managers.
- **Local Manager:** an NSX Manager system in charge of network and security services for a location.

### Federation Span: Local and Stretched

When you create a networking object from Global Manager, it can span one or more locations.

- **Local:** the object spans only one location.
- **Stretched:** the object spans more than one location.

You do not directly configure the span of a segment. A segment has the same span as the gateway it is attached to.

## Federation Regions

Security objects have a region. The region can be one of the following:

- **Location:** a region is automatically created for each location. This region has the span of that location.
- **Global:** a region that has the span of all available locations.
- **Custom Region:** you can create regions that include a subset of the available locations.

## Federation Tunnel Endpoints

In an Federation environment, there are two types of tunnel endpoints.

- **Tunnel End Point (TEP):** the IP address of a transport node (Edge node or Host) used for Geneve encapsulation within a location.
- **Remote Tunnel End Points (RTEP):** the IP address of a transport node (Edge node only) used for Geneve encapsulation across locations.

## Features and Configurations Supported in NSX-T Federation

All configurations made from the Global Manager are made in Policy mode. Manager mode is not available in NSX-T Federation.

See [#unique\\_5](#) for more information about the two modes.

## Configuration Maximums

An NSX-T Federation environment has the following configuration maximums:

- For most configurations, the Local Manager cluster has the same configuration maximums as an NSX Manager cluster. Go to [VMware Configuration Maximums tool](#) and select NSX-T Data Center.

Select the NSX-T Federation category for NSX-T Data Center in the [VMware Configuration Maximums tool](#) for exceptions and other NSX-T Federation-specific values.

- For a given location, the following configurations contribute to the configuration maximum:
  - Objects that were created on the Local Manager.
  - Objects that were created on the Global Manager and include the location in its span.

You can view the capacity and usage on each Local Manager. See [View the Usage and Capacity of Categories of Objects](#).

You can view the capacity and usage on each Local Manager. See *View the Usage and Capacity of Categories of Objects* in the *NSX-T Data Center Administration Guide*.



## Feature Support

Table 1-1. Features Supported in NSX-T Federation

Feature	Details	Related Links
Tier-0 Gateway	<ul style="list-style-type: none"> <li>■ 3.0.1 and later: Active Active and Active Standby</li> <li>■ 3.0.0: Active Active only</li> </ul>	<a href="#">Add a Tier-0 Gateway from Global Manager</a>
Tier-1 Gateway		<a href="#">Add a Tier-1 Gateway from Global Manager</a>
Segments	Layer 2 Bridge is not supported.	<a href="#">Add a Segment from Global Manager</a>
Groups	Some limitations. See <a href="#">Security in Federation</a> .	<a href="#">Create Groups from Global Manager</a>
Distributed Firewall		<a href="#">Create DFW Policies and Rules from Global Manager</a>
Gateway Firewall		<a href="#">Create Gateway Policies and Rules from Global Manager</a>
Network Address Translation (NAT)	<p>Tier-0 gateway:</p> <ul style="list-style-type: none"> <li>■ Active Active: You can configure stateless NAT only, that is, with action type Reflexive.</li> <li>■ Active Standby: You can create stateful or stateless NAT rules.</li> </ul> <p>Tier-1 gateway:</p> <ul style="list-style-type: none"> <li>■ You can create stateful or stateless NAT rules.</li> </ul> <p>Stateless NAT rules are pushed to all locations in the gateway's span unless scoped to one or more locations specifically.</p> <p>Stateful NAT rules are also pushed to all locations in the gateway's span or to the specific location selected. However, stateful NAT rules are realized and enforced only on the primary location.</p>	<a href="#">#unique_13</a>
DNS		<a href="#">See #unique_14</a>

DHCP and SLAAC

- DHCP Relay is supported on segments and gateways.
  - DHCPv4 server is supported on gateways with DHCP static bindings configured on segments.
  - IPv6 addresses can be assigned using SLAAC with DNS Through RA (DAD detects duplicates within a location only).
  - DHCP Relay: #unique\_15
  - DHCP Server (supported on gateway only):
    - #unique\_16
    - #unique\_17
    - #unique\_18
  - IPv6 address assignment: #unique\_19
-

Feature	Details	Related Links
Using objects created on Global Manager in a Local Manager configuration	<p>Most configurations are supported. For example:</p> <ul style="list-style-type: none"> <li>■ Connecting a Local Manager tier-1 gateway to a Global Manager tier-0 gateway.</li> <li>■ Using a Global Manager group in a Local Manager distributed firewall rule.</li> </ul> <p>These configurations are not supported:</p> <ul style="list-style-type: none"> <li>■ Connecting a Local Manager segment to a Global Manager tier-0 or tier-1 gateway.</li> <li>■ Connecting a load balancer to a Global Manager tier-1 gateway.</li> </ul>	
Backup and Restore	<ul style="list-style-type: none"> <li>■ 3.0.1 and later: backup with FQDN or IP is supported.</li> <li>■ 3.0.0: backup with FQDN is not supported.</li> </ul>	<a href="#">Backup and Restore in NSX-T Federation</a>
vMotion between locations	Cold migration between locations is not supported.	

## Understanding NSX-T Federation

In NSX-T Federation, you make configuration changes on the active Global Manager. The changes are synced with the relevant Local Managers and the standby Global Manager, if you have one.

Local Managers also sync some information with each other.

### Making Changes on Global Manager

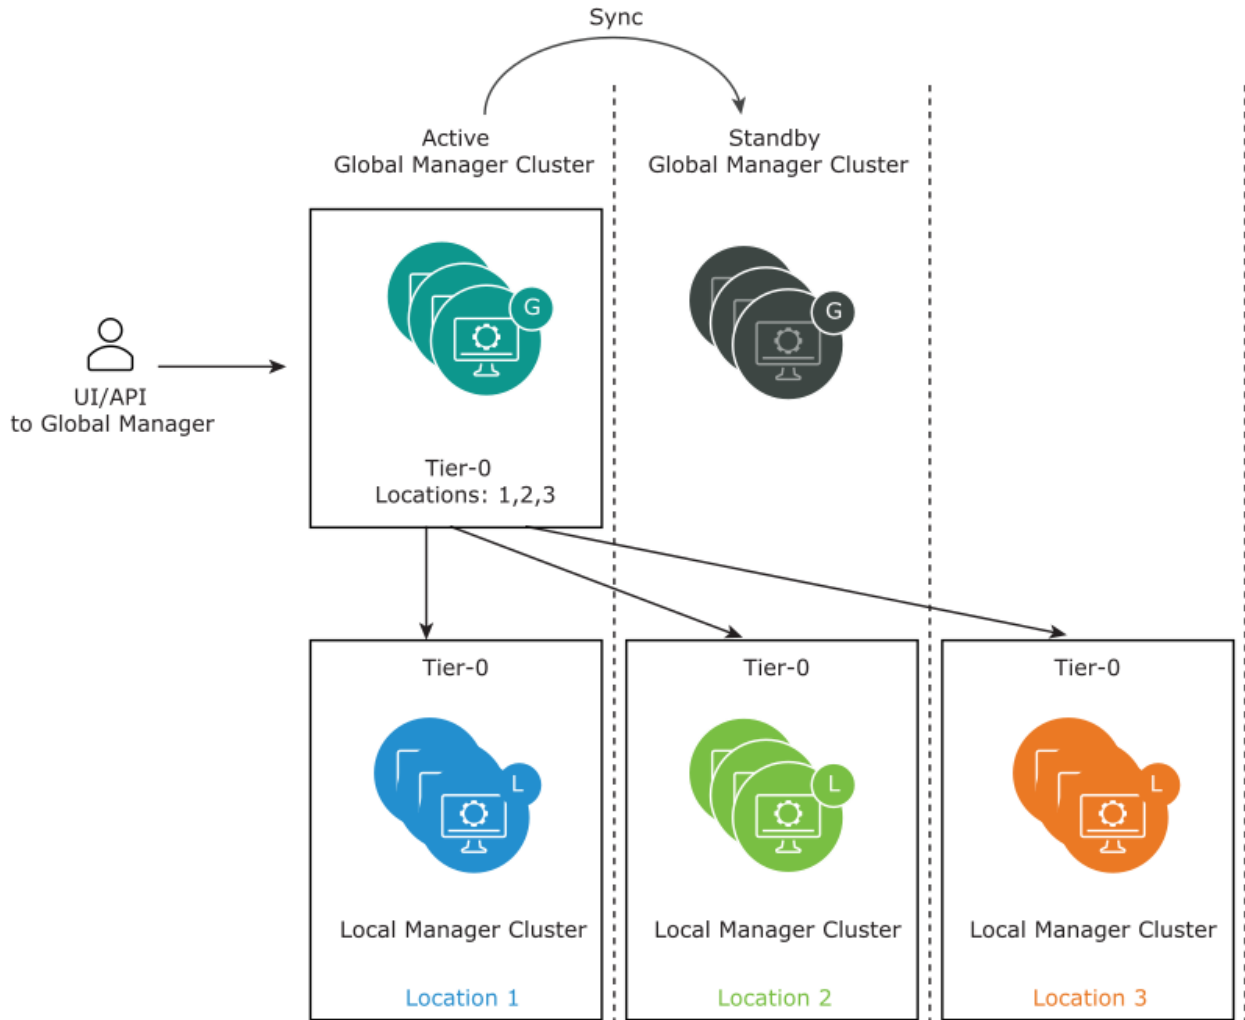
The Global Manager provides a user interface similar to the NSX Manager interface.

Configurations that are created on the Global Manager are read-only on the Local Managers. Configurations on the Local Managers are not synced with the Global Manager.

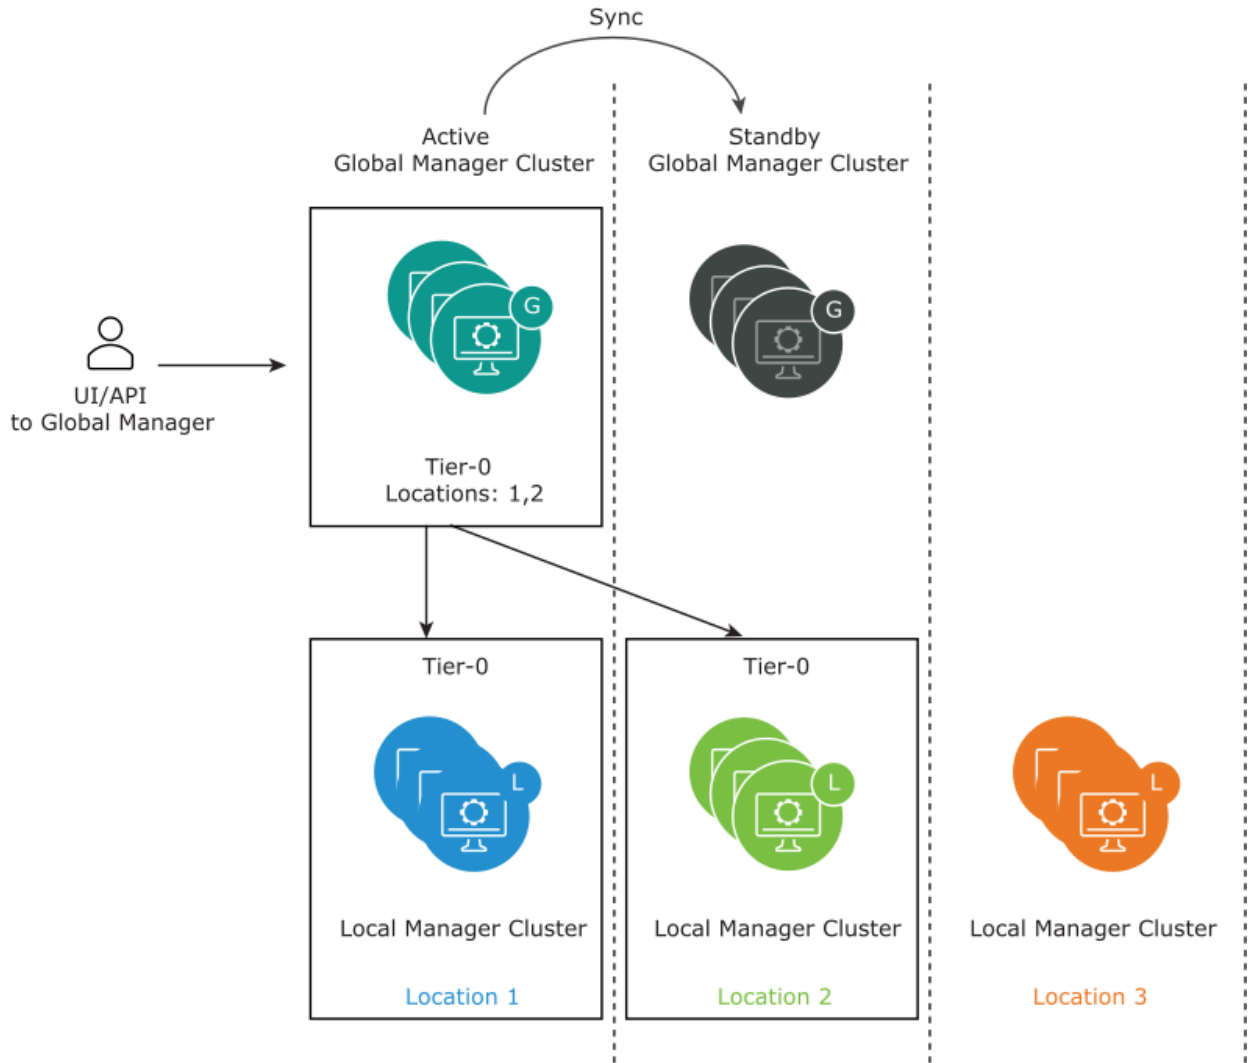
The Global Manager syncs a configuration with a Local Manager only if the configuration is relevant to that location. For example, if you create a tier-0 gateway and add it to Location 1, Location 2, and Location 3, the VMware,

configuration is synced with all three Local Managers.

If you have a standby Global Manager, configurations are also synced between the active Global Manager and the standby Global Manager.

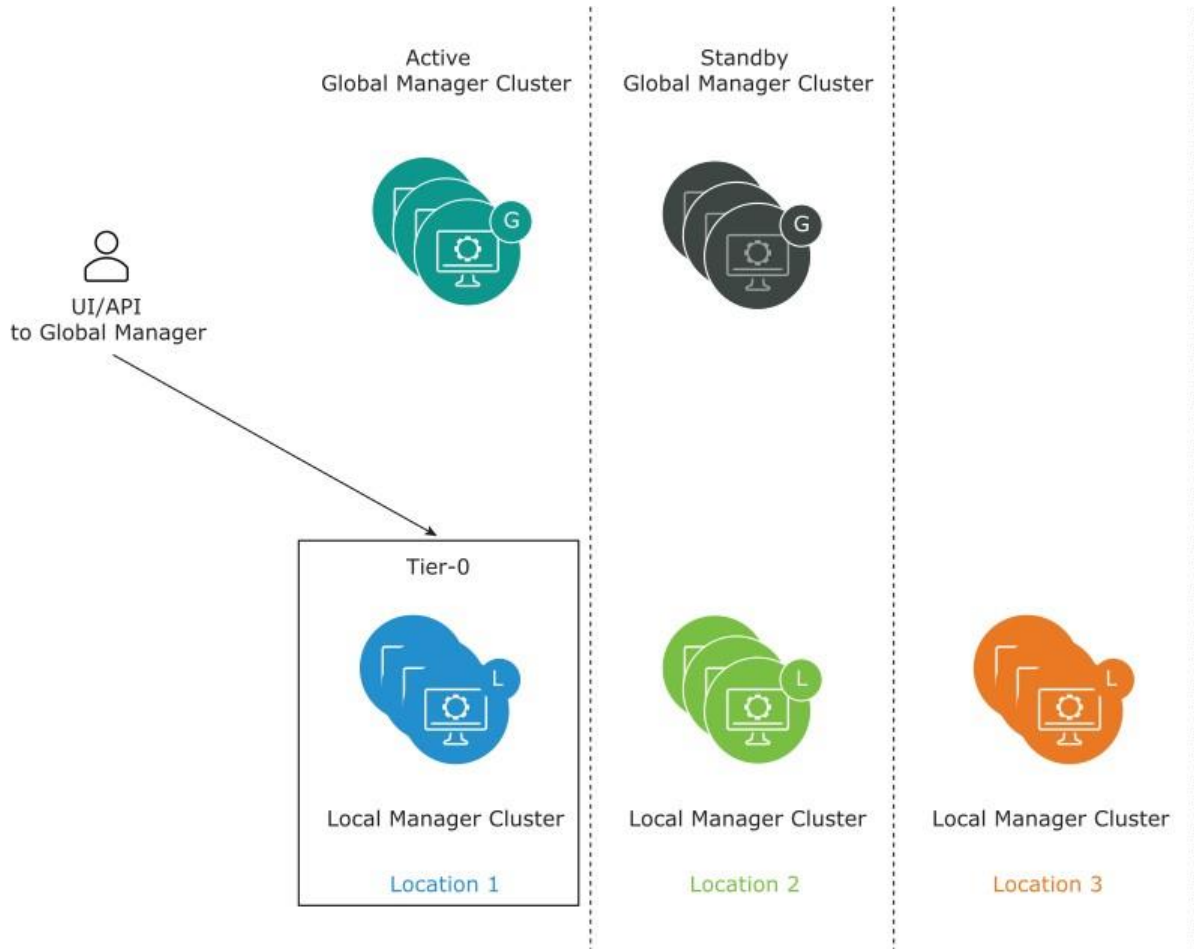


If the tier-0 gateway is added only to Location 1 and Location 2, the configuration is not synced with Location 3.



### Making Changes on Local Managers

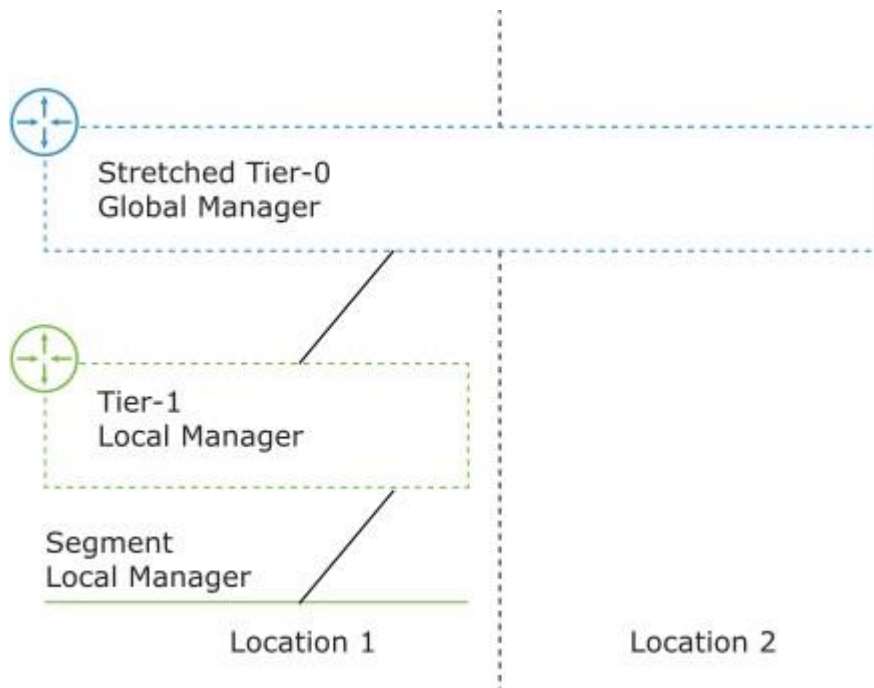
You can use the Local Manager to create objects on that specific Local Manager. These objects are not synced with the Global Manager or any other Local Manager.



### Realizing Global Manager Changes on Local Managers

The Global Manager validates changes against the Global Manager configuration only. When a Local Manager receives a configuration from the Global Manager, the configuration is realized in the fabric nodes of that Local Manager. During this realization, errors or conflicts might be detected.

For example, you can create a tier-0 gateway from Global Manager, and then from a Local Manager you can create and attach a tier-1 gateway to the tier-0 gateway.



Because Local Managers do not sync their configurations to the Global Manager, from the Global Manager context the tier-0 gateway does not appear to be connected to anything. You can delete the tier-0 gateway from the Global Manager, and this change is synced to the Local Managers.

When the changes are realized in each location, the following occurs:

- The tier-0 gateway can be deleted from the Local Manager in Location 2.
- The tier-0 gateway cannot be deleted from the Local Manager in Location 1.
- The tier-0 gateway is marked for deletion on the Global Manager.

When the tier-0 is disconnected from the tier-1 in Location 1, the tier-0 is deleted from Global Manager.

Most problems are displayed on the user interface. Additional problems can be displayed using these API calls.

- On Global Manager:

```
GET /global-manager/api/v1/global-infra/realized-state/alarms
```

- On Local Manager:

```
GET /policy/api/v1/infra/realized-state/alarms
```

## Using the Global and Local Manager Web Interfaces

You can use the Global Manager to create objects that are limited to one location, or span multiple locations.



## Location Drop-Down Menu on Global Manager


When you log into the Global Manager web interface, you see a Location drop-down menu in the top navigation bar. Using this menu, you can switch between the Global Manager and any associated Local Managers.





## Local and Global Objects


Objects created on a Local Manager are local objects. They are specific to that Local Manager and are not viewable from the Global Manager web interface.

Objects created from the Global Manager are global objects, though their span might not include all available locations.

On a Local Manager, you can see local objects, and any global objects that apply to that location. The global objects have an icon next to them: .

This screenshot from the Local Manager web interface shows two segments. The segment

segment-01 has the  icon next to it, which indicates that it was created on the Global Manager. The segment segment-02 has no  icon, which indicates that it was created on the Local Manager.



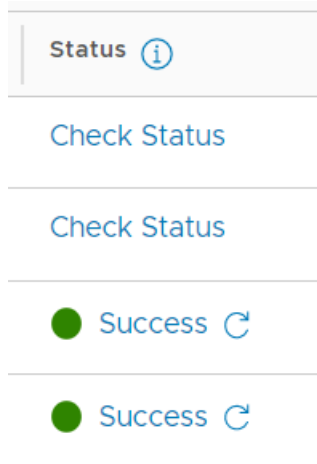
Because all objects on the Global Manager are global, there is no icon displayed when you are logged into the Global Manager.

## Status of Local and Global Objects

Local Managers display the status of both global and local objects.

The Global Manager displays only global objects, but does not automatically receive the status of the objects.


To retrieve the latest status from the Local Managers, click **Check Status** for the object. To refresh the status, click the **Refresh** icon.



## Overriding Global Manager Configurations on Local Manager

When you create an object from Global Manager, the same configuration is propagated to all relevant locations. Starting in NSX-T Data Center 3.0.1 you can override some Global Manager configurations on a Local Manager.

To override a configuration, click the three dots menu (⋮) next to the configuration, and click **Edit**. If the **Edit** menu item is dimmed, you cannot override this configuration.

If a configuration is overridden, you see this icon in the status column on both Global Manager and Local Manager: .

To remove an override, click the three dots menu (⋮) next to the configuration, and click **Revert**. The configuration from Global Manager is restored.

If you override a configuration from Global Manager on a Local Manager, and then you delete the configuration from the Global Manager, the configuration persists on the Local Manager. When you revert the configuration, the configuration is deleted from Local Manager.

You can get a list of all configurations that have been overridden. Make this API request to the Global Manager: GET `https://<global-mgr>/global-manager/api/v1/global-infra/overridden-resources`.

## Gateway Configurations

Gateway configurations are found in **Networking > Tier-0 Gateways** and **Networking > Tier-1 Gateways**.

You can modify the following gateway configurations:

- Tier-0 Gateway BGP Configuration
- Tier-0 Gateway Interfaces

## Profile Configurations

Profile configurations on Global Manager are used in all Local Managers. There is no span setting for a profile configuration.

You can override the following global profile configurations from Local Manager:

- **Segment Profiles: `Networking > Segments > Segment Profiles`**
  - IP Discovery Profiles
  - MAC Discovery Profiles
  - Segment Security Profiles
  - SpoofGuard Profiles
- **Networking Profiles: `Networking > Networking Settings`**
  - IPv6 DAD Profiles
  - IPv6 ND Profiles
  - Gateway QoS Profiles
  - BFD Profiles
- **Context Profiles: `Inventory > Context Profiles`**
- **Security Profiles: `Security > Security Profiles`**
  - Firewall Session Timer Profile
  - Edge Gateway Flood Protection Profiles
  - Firewall Flood Protection Profiles
  - DNS Security Profiles
  - CPU and Memory Threshold Profiles are API only:
    - Override with PUT/PATCH `https://<local-manager>/policy/api/v1/global-infra/settings/firewall/cpu-mem-thresholds-profiles/<id>?action=override`.
    - Revert with DELETE `https://<local-manager>/policy/api/v1/global-infra/settings/firewall/cpu-mem-thresholds-profiles/<id>`.
- **Troubleshooting Profiles: `Plan & Troubleshoot`**
  - Firewall IPFIX Profiles
  - Switch IPFIX Profiles
  - IPFIX Firewall Collector

- IPFIX Switch Collector
- Remote L3 Span Port Mirroring Profile
- Logical Span Port Mirroring Profile

- QoS Profile

## Getting Started with NSX-T Federation

To get started with NSX-T Federation, you install the Global Manager, configure the Global Manager as active, and add locations.

Task	Details
Check the requirements for Federation.	See <a href="#">#unique_25</a> .
Install the Global Manager.	See <a href="#">Install the Active and Standby Global Manager</a> .
Make the Global Manager cluster active.	See <a href="#">Make the Global Manager Active and Add Standby Global Manager</a> .
Add Locations to the active Global Manager.	See <a href="#">Add a Location</a> .

For further configuration tasks, such as preparing Edge clusters for stretched networking, and creating objects from the Global Manager, see *Federation* in the *NSX-T Data Center Administration Guide*.

## Install the Active and Standby Global Manager

To use NSX-T Federation, you must install the Global Manager.

Installing a Global Manager appliance is similar to installing an NSX Manager appliance. The only difference is that when you deploy the appliance, you select *NSX Global Manager* for the role.

Install a standby Global Manager appliance for high availability and disaster recovery. The standby Global Manager appliance must be installed in a different location with a latency of 150ms or less.

### Prerequisites

- Verify that your environment meets the requirements for NSX Manager. See [#unique\\_29](#).
- Decide which locations will contain the active and standby Global Manager appliances.
- Verify that you are installing the Global Manager appliance with NSX-T Data Center 3.1.0 or later.

---

**Important** All Global Manager and Local Manager appliances in an NSX-T Federation environment must have the same version of NSX-T Data Center installed.

### Procedure

1 Install the first Global Manager appliance.

- On vSphere: `#unique_30`.
  - Select Medium or Large for the deployment configuration size. Do not use Small.
  - Select NSX Global Manager for the **Rolename**.

- On KVM: [#unique\\_31](#).
    - Deploy a medium appliance. For example: `virt-install --import --ram 16000-- vcpus 6`.
    - Select NSX Global Manager for the `nsx_role`.
- 2 Log in to the NSX Manager appliance.
- See [#unique\\_32](#).
- 3 (Optional) If you are installing Global Manager on vSphere, configure a compute manager.
- See [#unique\\_33](#).

**Note** If you are at this step while installing the standby Global Manager, you must configure a separate compute manager. Do not use the same compute manager that you configured for the active Global Manager.

- 4 Create a Global Manager cluster. See [#unique\\_34](#) for design recommendations.
- On vSphere with a compute manager configured: See [#unique\\_35](#).
  - On vSphere without a compute manager configured: Repeat the NSX Manager install on vSphere steps to install three appliances, then form the cluster. See [#unique\\_36](#).
  - On KVM: Repeat the NSX Manager install on KVM steps to install three appliances, then form the cluster. See [#unique\\_36](#).
- 5 Configure a VIP for the Global Manager cluster.
- See [#unique\\_37](#).
- 6 In a different location, install a standby Global Manager appliance and form a cluster by repeating these steps.

#### What to do next

Select the designated Global Manager appliance as active and connect it with the standby Global Manager.

## Make the Global Manager Active and Add Standby Global Manager

After you have deployed a Global Manager appliance, you can make the Global Manager active.

Adding a standby Global Manager is optional but recommended for high availability of the Global Manager.



Procedure

- 1 Log in to the Global Manager appliance at <https://global-manager-ip-or-fqdn/>.
- 2 Select **System > Location Manager**. In the **Global Manager** tile, click **Make Active**. Provide a descriptive name for the active Global Manager and click **Save**.

3 (Optional) Add a standby Global Manager cluster.

- a Install a new Global Manager appliance in a secondary location and start it. Follow the same instructions as for installing the primary Global Manager, see [Install the Active and Standby Global Manager](#) .
- b From the active Global Manager, add this newly installed Global Manager appliance as standby.

Navigate back to your active Global Manager and click **Add Standby** and provide the following information:

Option	Description
<b>Global Manager Name</b>	Provide a name for the standby Global Manager.
<b>FQDN/IP</b>	Enter the FQDN or IP address of the Global Manager cluster VIP at the secondary location. Do not enter an individual Global Manager FQDN or IP.
<b>Username and Password</b>	Provide the admin user's credentials for the Global Manager at the secondary location.
<b>SHA-256 Thumbprint</b>	<p>Log in to any Global Manager node at the secondary location and run this command:</p> <pre>get certificate cluster thumbprint</pre> <p>The result is the cluster VIP certificate: bfaela0a...</p>
<b>Check Compatibility</b>	Click <b>Check Compatibility</b> to ensure that the Global Manager can be added as standby. This checks that the NSX-T Data Center version is compatible.

- c Click **Save**.
- d Click **Make Standby**.

## Add a Location

After you add a location to Global Manager, you can create objects from Global Manager that span that location.

You can find the number of supported locations in the [VMware Configuration Maximums tool](#). Select the appropriate version of NSX-T Data Center, select the NSX-T Federation category, and click **View Limits**.

After you add a location to the Global Manager, the NSX Manager is called a Local Manager.

### Prerequisites

- Verify that you have an NSX-T Data Center environment installed in the

location you want to add.

You can add a new NSX-T Data Center environment or an NSX-T Data Center environment with an existing configuration.

- The NSX-T Data Center environment in the new location must have three NSX Manager nodes deployed and a cluster VIP configured. See [#unique\\_37](#).

For a proof-of-concept environment, you can add a location that has only one NSX Manager node, but you must still configure a cluster VIP.

- Verify that the latency between the Global Manager and the location is 150 ms or less.
- Verify that the environment you are adding has NSX-T Data Center 3.0 installed.
- If you are using VMware Tanzu Kubernetes Grid Integrated Edition (formerly VMware Enterprise PKS), you have to install the same certificate on all your NSX Manager nodes. Currently you cannot change the certificates on NSX Manager nodes after you add the NSX Manager to the Global Manager. Therefore, update the certificates on your NSX Manager nodes before adding this location to the Global Manager, to ensure that you can use VMware Tanzu Kubernetes Grid Integrated Edition with your NSX-T Data Center deployment, while also using NSX-T Federation. See "Certificates" in the *NSX-T Data Center Administration Guide* for more information on certificates used in NSX-T Data Center and how to replace them.

Procedure

- 1 Log in to the Global Manager at <https://global-manager-ip-or-fqdn/>.
- 2 Select and click **Add On-Prem Location**.
- 3 In the **Add New Location** dialog box, enter the Location details.

Option	Description
<b>Location Name</b>	Provide a name for the location.
<b>FQDN/IP</b>	Enter the FQDN or IP address of the NSX Manager cluster VIP. Do not enter an individual NSX Manager FQDN or IP.
<b>Username and Password</b>	Provide the admin user's credentials for the NSX Manager at the location.
<b>SHA-256 Thumbprint</b>	Log in to any NSX Manager node in the cluster and run this command:  <pre>get certificate cluster thumbprint</pre> <p>The result is the cluster VIP certificate: bfae1a0a...</p>

**Check Compatibility**

Click **Check Compatibility** to ensure that the location can be added. This checks that the NSX-T Data Center version is compatible.

**What to do next**

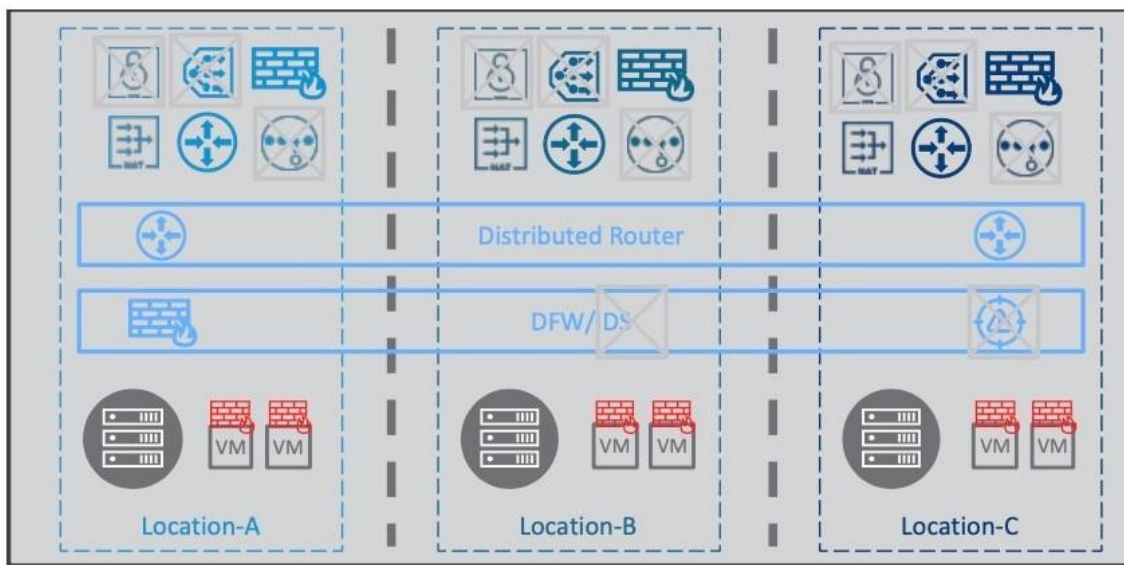
If you want to create gateways and segments that span more than one location, you must configure a remote tunnel endpoint (RTEP) on Edge nodes in each location to handle the cross- location traffic. See [Configure Edge Nodes for Stretched Networking](#).

## Networking

This section provides networking set up procedures for Federation. NSX-T Federation currently supports the following network services:

- Switching: Overlay and VLAN
- IPAM: DHCP Relay, static binding, and DNS
- Routing: NAT and route redistribution
- Routing protocols: BGP and Static

Figure 1-2. NSX-T Federation Network Services



Tier-0 gateways, tier-1 gateways, and segments can span one or more locations in the NSX Federation environment.

When you plan your network topology, keep these requirements in mind:

- Tier-0 and tier-1 gateways can have a span of one or more locations.
- The span of a tier-1 gateway must be equal to, or a subset of, the span of the tier-0 gateway it is attached to.
- A segment has the same span as the tier-0 or tier-1 gateway it is attached to. Isolated segments are not realized until they are connected to a gateway.
- NSX Edge nodes in the Edge Cluster selected on the Global Manager for tier-0 and tier-1 gateways must be configured with the Default TZ Overlay.

You can create different topologies to achieve different goals. You can create segments and gateways that are specific to a given location. Each site has its own configuration, but you can manage everything from the Global Manager interface. You can create segments and gateways that span locations. These stretched networks provide consistent networking across sites.

## Tier-0 Gateway Configurations in Federation

With Federation, you can deploy a tier-0 gateway that is limited to a single location, or you can stretch it across multiple locations.

Tier-0 gateways can have one of the following configurations:

- Non-stretched tier-0 gateway.
- Stretched active-active with primary and secondary locations.
- Stretched active-active with all primary locations.
- Stretched active-standby with primary and secondary locations.

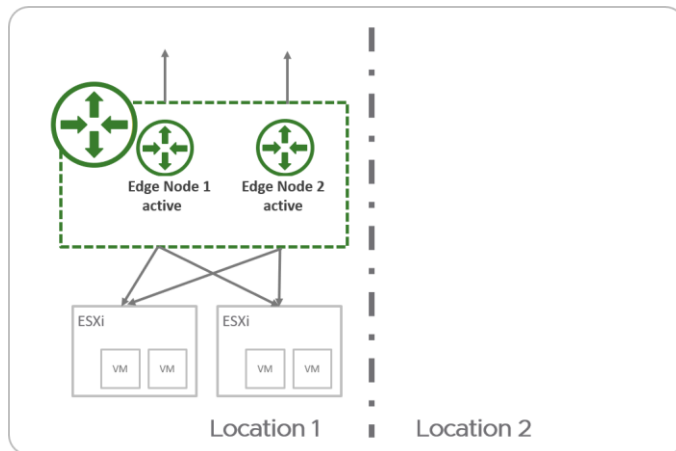
---

**Note** Active-standby tier-0 gateways are supported starting in NSX-T Data Center 3.0.1.

---

### Non-Stretched Tier-0 Gateway

You can create a tier-0 gateway from Global Manager that spans only one location. This is similar to creating the tier-0 gateway on the Local Manager directly, but has the advantage that you can manage it from Global Manager.



### Stretched Active-Active Tier-0 Gateway with Primary and Secondary Locations

In an active-active tier-0 gateway with primary and secondary locations, the following applies:

- All Edge nodes are active at the same time, therefore the tier-0 cannot run stateful services.
- All traffic enters and leaves through the Edge nodes in the primary location.



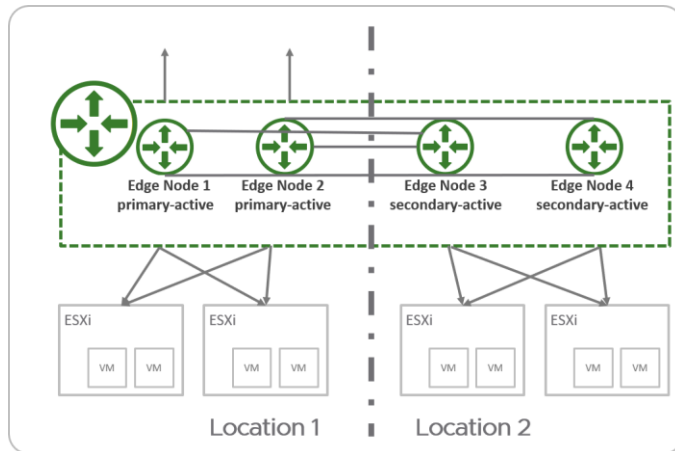
If both the tier-0 gateway and the linked tier-1 gateway have primary and secondary locations, configure the same location to be primary for both gateways to reduce cross-location traffic.

---

**Important** In this topology, NSX-T Data Center ensures that all egress traffic leaves through the primary location.

If your environment has stateful services, such as external firewall, on the physical network, you must ensure that the return traffic enters through the primary location. For example, you can add AS path prepending on the BGP peers in your secondary locations.

If you do not have stateful services on your physical network, and you choose to have asymmetric routing, you must disable Unicast Reverse Path Forwarding (uRPF) on all externally tier-0 interfaces.




### Stretched Active-Active Tier-0 Gateway with All Primary Locations

In an active-active tier-0 gateway with all primary locations, the following applies:

- All Edge nodes are active at the same time, therefore the tier-0 cannot run stateful services.
- All traffic enters and leaves through Edge nodes in the same location as the workloads.

---


**Important** This topology allows traffic to egress locally from each location. You must ensure that return traffic enters the same location to allow stateful services such as firewall. For example, you can configure a location-specific NAT IP so that return traffic is always routed back to the same location that it left.



### Stretched Active-Standby Tier-0 Gateway with Primary and Secondary Locations

In an active-standby tier-0 gateway with primary and secondary locations, the following applies:

- Only one Edge node is active at a time, therefore the tier-0 can run stateful services.
- All traffic enters and leaves through the active Edge node in the primary location.



For Active Standby tier-0 gateways, the following services are supported:

- Network Address Translation (NAT)
- Gateway Firewall
- DNS
- DHCP

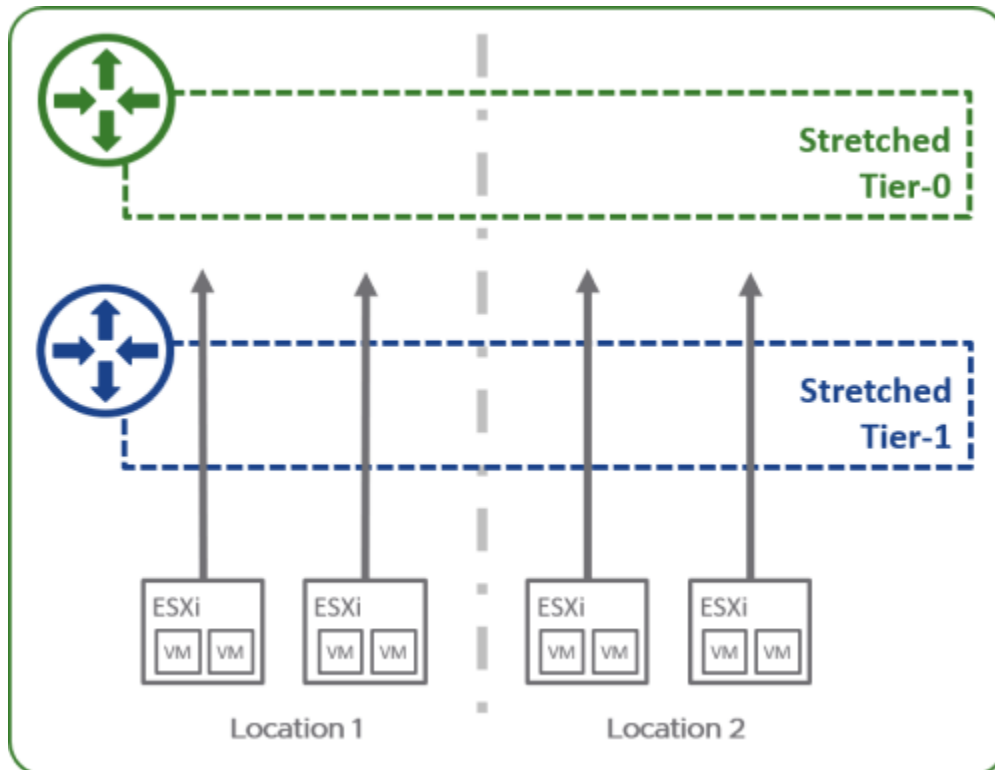
See [Features and Configurations Supported in NSX-T Federation](#) for more information.

## Tier-1 Gateway Configurations in NSX-T Federation

With Federation, you can deploy a tier-1 gateway to provide distributed routing only, or you can configure services on it.

## Tier-1 Gateway for Distributed Routing Only

You can create a tier-1 gateway in NSX-T Federation for distributed routing only. This gateway has the same span as the tier-0 gateway it is linked to. The tier-1 does not use Edge nodes for routing. All traffic is routed from host transport nodes to the tier-0 gateway. However, to enable cross-location forwarding, the tier-1 allocates two Edge nodes from the Edge cluster configured on the linked tier-0 to use for that traffic.



## Tier-1 Gateway with Services or Custom Span

You configure the tier-1 gateway with Edge clusters if you need one of the following configurations:

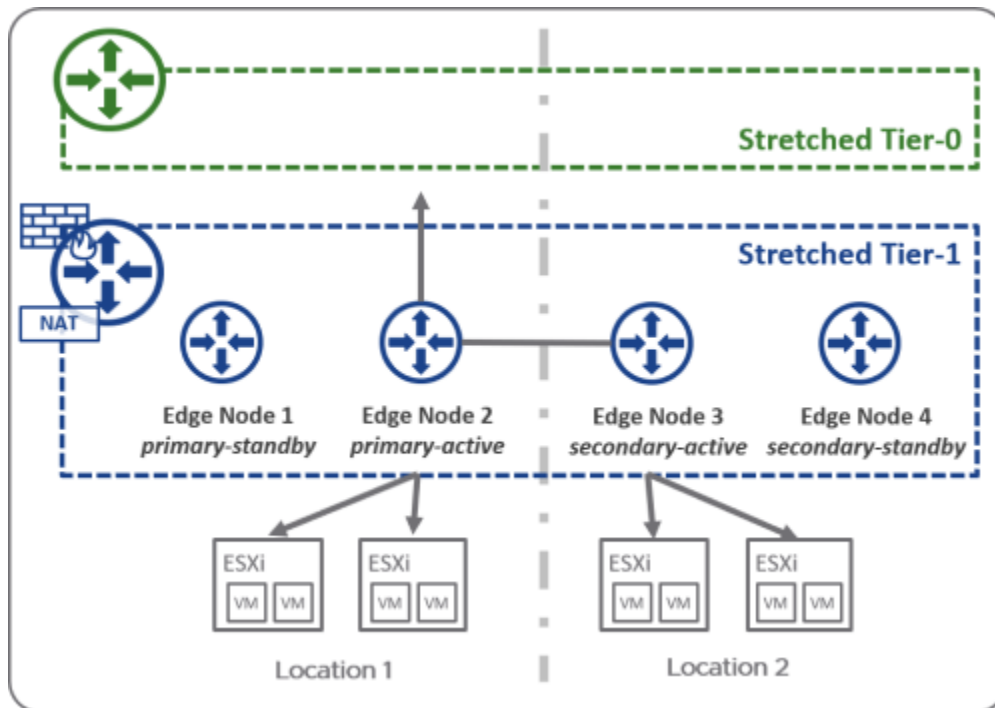
- You want to run services on the tier-1 gateway.
- You want to deploy a tier-1 gateway that has a different span than the linked tier-0 gateway.

You can remove locations, but you cannot add locations that are not already included the span of the tier-0 gateway.

You select one of the locations to be the primary location. All other locations are secondary. The HA mode for the tier-1 gateway is Active Standby. All traffic passing through this tier-1 gateway passes through the active edge node in the primary location.

If both the tier-1 gateway and the linked tier-0 gateway have primary and VMware,

secondary locations, configure the same location to be primary for both gateways to reduce cross-location traffic.



## Configure Edge Nodes for Stretched Networking

If you want to create gateways and segments that span more than one location, you must configure a remote tunnel endpoint (RTEP) on Edge nodes in each location to handle the cross- location traffic.

When you configure an RTEP, do it on an Edge cluster basis. All Edge nodes in the cluster must have an RTEP configured. You do not need to configure all Edge clusters with RTEP. RTEPs are required only if the Edge cluster is used to configure a gateway that spans more than one location.

You can configure the TEP and RTEP to use the same physical NIC on the Edge node or use separate physical NICs.

You can also configure RTEPs from each Local Manager. Select **System > Get Started > Configure Remote Tunnel Endpoint**.

You can edit RTEPs on an Edge node. Log into the Local Manager and select **System > Fabric**

**> Nodes > Edge Transport Nodes**. Select an Edge node, and click **Tunnels**. If an RTEP is configured, it is displayed in the **Remote Tunnel Endpoint** section. Click **Edit** to modify the RTEP configuration.

### Prerequisites

- Verify that each location participating in the stretched network has at least one Edge cluster.
- Determine which layer 3 networks and VLANs to use for RTEP networks.

- Intra-location tunnel endpoints (TEP) and inter-location tunnel endpoints (RTEP) must use separate VLANs and layer 3 subnets.

- Verify that all RTEP networks used in a given NSX-T Federation environment have IP connectivity to each other.
- Verify that external firewalls allow cross-location RTEP tunnels, and BGP sessions between Edges. See VMware Ports and Protocols at <https://ports.vmware.com/home/NSX-T-Data-Center>.
- Configure the MTU for RTEP on each Local Manager. The default is 1500. Set the RTEP MTU to be as high as your physical network supports. On each Local Manager, select **System > Fabric > Settings**. Click **Edit** next to **Remote Tunnel Endpoint**.

Procedure

- 1 From your browser, log in with admin privileges to the active Global Manager at `https://<global-manager-ip-address>`.
- 2 Go to **System > Location Manager** and click **Networking** from the location you want to configure for stretched networking.
- 3 Click **Configure** next to the Edge cluster for which you want to set up the RTEP.

The **Configure Edge Nodes for Stretched Networking** screen opens in the Local Manager with that Edge cluster selected.

- 4 You can select all Edge Nodes in this cluster or one node at a time. Provide the following details for the RTEP configuration:

Option	Description
<b>Host Switch</b>	Select a host switch from the drop-down menu.
<b>Teaming Policy</b>	Select a teaming policy if you have one configured.
<b>RTEP VLAN</b>	Enter the VLAN ID for the RTEP network. Valid values are between 1 and 4094.
<b>IP Pool for all nodes</b>	Select an IP pool for all nodes in this Edge Cluster. If you want to assign an IP address to an individual node, you can edit the RTEP configuration later.
<b>Inter Location MTU</b>	The default is 1500.

- 5 Click **Save**.

You can click each of the Edge Nodes that are marked as Configured to see the Edge node configuration details. Select the **Tunnels** tab to view and edit the RTEP configuration.

## Add a Tier-0 Gateway from Global Manager



You can add a tier-0 gateway from the Global Manager. This gateway can have a span of one or more locations. This span affects the span of the tier-1 gateways and segments attached to it.

See [Tier-0 Gateway Configurations in Federation](#) for details about tier-0 gateway configurations in Federation.

The following settings must be kept consistent across locations. If you change these settings from the Global Manager web interface, those changes are automatically applied on all locations. However, if you change these settings using the API, you must manually make the same changes in each location.

- Local AS
- ECMP settings
- Multipath Relax settings
- Graceful Restart

---

**Important** When you create a tier-0 gateway from Global Manager, you must configure an external interface in each location that the tier-0 is stretched to. Each external interface must be connected to a segment that was created from Global Manager, with the **Connectivity** set to None and the **Traffic type** set to VLAN. See [Add a Segment from Global Manager](#). The Edge nodes configured with those external interfaces are used for inter-location communication, even if northbound communication is not needed.

#### Prerequisites

- If you are creating a tier-0 gateway that spans more than one location, verify that each location has Edge nodes configured with RTEPs for stretched networking. See [Configure Edge Nodes for Stretched Networking](#).

#### Procedure

- 1 From your browser, log in with admin privileges to the active Global Manager at `https://<global-manager-ip-address>`.
- 2 Select **Networking > Tier-0 Gateways**.
- 3 Enter a name for the gateway.
- 4 Select an HA (high availability) mode to configure within each location.

The default mode is active-active. In the active-active mode, traffic is load balanced across edge nodes in all locations. In the active-standby mode, an elected Edge node processes traffic in each location. If the active node fails, the standby node becomes active.

---

**Note** Active-standby tier-0 gateways are supported starting in NSX-T Data Center 3.0.1.

---

- 5 If the HA mode is active-standby, select a failover mode.

Option	Description
Preemptive	If the preferred node fails and recovers, it will preempt its peer and become the active node. The peer will change its state to standby.
Non-preemptive	If the preferred node fails and recovers, it will check if its peer is the active node. If so, the preferred node will not preempt its peer and will be the standby node.

- 6 Specify the span of this tier-0 gateway by providing the following details for each location. To add additional locations, click **Add Location**.

Option	Description
<b>Location</b>	Select the location from the drop-down menu.
<b>Edge Cluster</b>	Select an Edge cluster from this location. If you are configuring a stretched tier-0, you must select an Edge cluster that contains Edge nodes that are configured with an RTEP.
<b>Mode</b>	Each location of the tier-0 gateway can have a mode of <b>Primary</b> or <b>Secondary</b> . <ul style="list-style-type: none"> <li>■ If the HA mode is <b>Active Active</b>, you can configure the tier-0 gateway with all locations mode set to primary. <ol style="list-style-type: none"> <li>1 Select the <b>Mark all locations as Primary</b> toggle to mark all locations as primary.</li> </ol> </li> <li>■ If the HA mode is <b>Active Active</b> or <b>Active Standby</b>, you can configure the tier-0 gateway with one location set to <b>Primary</b>, and all others set to <b>Secondary</b>. <ol style="list-style-type: none"> <li>1 Select <b>Primary</b> mode for one location. In all other locations, set mode to <b>Secondary</b>.</li> <li>2 For secondary locations, you must select a fallback preference.</li> </ol> </li> </ul>

- 7 Click **Additional Settings**.

- a In the **Internal Transit Subnet** field, enter a subnet.

This is the subnet used for communication between components within this gateway. The default is 169.254.0.0/24.

- b In the **T0-T1 Transit Subnets** field, enter one or more subnets.

These subnets are used for communication between this gateway and all tier-1 gateways that are linked to it. After you create this gateway and link a tier-1 gateway to it, you will see the actual IP address assigned to the link on the tier-0 gateway side and on the tier-1 gateway side. The address is displayed in **Additional Settings > Router Links** on the tier-0 gateway page and the tier-1 gateway page. The default is 100.64.0.0/16.

- c In the **Intersite Transit Subnet** field, enter a subnet. This subnet is used for cross-location communication between gateway components. The default is 169.254.32.0/20.

- 8 Click **Save**.

- 9 To configure interfaces, click **Interfaces** and **Set**. Configure an

external interface for each location that the tier-0 gateway spans.

a Click **Add**

**Interface.** b

Enter a name.

c Select a location.

d Select a type.

If the HA mode is active-standby, the choices are **External**, **Service**, and **Loopback**. If the HA mode is active-active, the choices are **External** and **Loopback**.

Service interfaces are supported only on gateways that span one location. If the gateway is stretched, service interfaces are not supported.

e Enter an IP address in CIDR

format. f Select a segment.

The segment must be created from the Global Manager, with the **Connectivity** set to None and the **Traffic type** set to VLAN. See [Add a Segment from Global Manager](#).

g If the interface type is not **Service**, select an NSX Edge node.

h (Optional) If the interface type is not **Loopback**,

enter an MTU value. i Skip **PIM** configuration.

Multicast is not supported in NSX-T

Federation. j (Optional) Add tags

and select an ND profile.

k (Optional) If the interface type is **External**, for **URPF Mode**, you can select **Strict** or **None**.

URPF (Unicast Reverse Path Forwarding) is a security feature.

l After you create an interface, you can download the ARP table by clicking the menu icon (three dots) for the interface and selecting **Download ARP table**.

10 Click **Routing** to add IP prefix lists, community lists, static routes, and route maps.

When you add a static route on a tier-0 gateway, the default behavior is that the static routes are pushed to all locations configured on the gateway. However, the routes are enabled only on the primary locations. This ensures that on the secondary locations, the routes that are learned from the primary location are preferred.

If you want to change this behavior, you can use the **Enabled on Secondary** setting and the **Scope** setting.

If you select **Enabled on Secondary**, the static route is also

enabled on the secondary locations.

When you add a next hop for a static route, you can set the **Scope**. The scope can be an interface, a gateway, or a segment. On a tier-0 gateway created from Global Manager, the scope can also be a location. You can use the scope setting to configure different next hops for each location.

**11** Click **BGP** to configure BGP.

When you configure BGP on a tier-0 gateway from the Global Manager, most settings apply to all locations.

Some of the settings within the BGP configuration, such as **Route Aggregation** and **BGP Neighbors** prompt you to provide separate values for each location.

See [#unique\\_42](#) for more information about configuring BGP.

12 To configure route redistribution, click **Route Redistribution**, and for each location, click **Set**.

Select one or more of the sources:

- Tier-0 subnets: **Static Routes, NAT IP, IPSec Local IP, DNS Forwarder IP, EVPN TEP IP, Connected Interfaces & Segments.**

Under **Connected Interfaces & Segments**, you can select one or more of the following: **Service Interface Subnet, External Interface Subnet, Loopback Interface Subnet, Connected Segment.**

- Advertised tier-1 subnets: **DNS Forwarder IP, Static Routes, LB VIP, NAT IP, LB SNAT IP, IPSec Local Endpoint, Connected Interfaces & Segments.**

Under **Connected Interfaces & Segments**, you can select **Service Interface Subnet** and/or **Connected Segment.**

What to do next

Set up a tier-1 gateway from Global Manager.

## Add a Tier-1 Gateway from Global Manager

A gateway can be configured in one or more locations. These locations are the span of the gateway. A tier-1 gateway cannot have a greater span than the tier-0 gateway it is connected to.

See [Tier-1 Gateway Configurations in NSX-T Federation](#) for details about tier-1 gateway configuration options in Federation.

Prerequisites

Verify you have a tier-0 gateway configured.

Procedure

- 1 From your browser, log in with admin privileges to an NSX Manager at `https://<global-manager-ip-address>`.
- 2 Select **Networking > Tier-1 Gateways**.
- 3 Click **Add Tier-1 Gateway**.
- 4 Enter a name for the gateway.
- 5 Select a tier-0 gateway to connect to this tier-1 gateway to create a multi-tier topology.



- If you select a tier-0 gateway, the Locations configuration is populated with the same locations that are configured on the tier-0. If needed, you can modify the locations configuration in the Locations section.
- If you do not select a tier-0 gateway, you can select locations. However, if you later connect the tier-1 gateway to a tier-0 gateway, you might need to update the locations to create a valid configuration.

6 In **Locations**, you can change the **Enable Edge Clusters for Services or Custom Span** setting. It is disabled by default.

- Leave **Enable Edge Clusters for Services or Custom Span** disabled if you want the tier-1 gateway to have the same span as the tier-0 gateway, and you do not need to enable services on the tier-1 gateway. The tier-1 gateway will perform distributed routing only.
- Enable **Enable Edge Clusters for Services or Custom Span** if you want to choose a subset of locations for the tier-1 gateway, or if you want to enable services on the tier-1 gateway.

If you enable **Enable Edge Clusters for Services or Custom Span**, enter the location, cluster, and mode information.

- a Select a location from the drop-down menu. If you linked this tier-1 gateway to a tier-0 gateway, the locations of that tier-0 gateway are automatically listed. If needed, you can delete a location.
- b Select an NSX Edge cluster for each location. If the tier-1 gateway spans more than one location, the Edge clusters must already be configured with an RTEP for each of its Edge Nodes.
- c (Optional) To select specific Edge nodes, click **Set** next the Edge cluster. Edge nodes are automatically allocated if you do not select Edge nodes.
- d Select a mode for each location. Mode can be Primary or Secondary. Only one location can be configured with Primary mode. All northbound traffic from this tier-1 gateway is sent through this location.

7 If you have enabled Edge clusters, select a failover mode.

Option	Description
Preemptive	If the preferred NSX Edge node fails and recovers, it will preempt its peer and become the active node. The peer will change its state to standby. This is the default option.
Non-preemptive	If the preferred NSX Edge node fails and recovers, it will check if its peer is the active node. If so, the preferred node will not preempt its peer and will be the standby node.

8 Skip selecting a size from the **Edge Pool Allocation Size** drop-down menu.

9 If you have enabled Edge clusters, select a setting for **Enable StandBy Relocation**.

Standby relocation means that if the Edge node where the active or

standby logical router is running fails, a new standby logical router is created on another Edge node to maintain high availability. If the Edge node that fails is running the active logical router, the original standby logical router becomes the active logical router and a new standby logical router is created. If the Edge node that fails is running the standby logical router, the new standby logical router replaces it.

10 (Optional) Click **Route Advertisement**.

Select one or more of the following:

- **All Static Routes**
- **All NAT IP's**
- **All DNS Forwarder Routes**
- **All LB VIP Routes**
- **All Connected Segments and Service Ports**
- **All LB SNAT IP Routes**
- **All IPsec Local Endpoints**

11 Click **Save**.

12 (Optional) Click **Route Advertisement**.

a In the **Set Route Advertisement Rules** field, click **Set** to add route advertisement rules.

13 (Optional) Click **Additional Settings**.

a For IPv6, you can select or create an **ND Profile** and a **DAD Profile**.

These profiles are used to configure Stateless Address Autoconfiguration (SLAAC) and Duplicate Address Detection (DAD) for IPv6 addresses.

b Select an **Ingress QoS Profile** and an **Egress QoS Profile** for traffic limitations.

These profiles are used to set information rate and burst size for permitted traffic. See [#unique\\_43](#) for more information on creating QoS profiles.

If this gateway is linked to a tier-0 gateway, the **Router Links** field shows the link addresses.

14 (Optional) Click **Service Interfaces** and **Set** to configure connections to segments. Required in some topologies such as VLAN-backed segments or one-arm load balancing.

Service interfaces are supported only on gateways that span one location. If the gateway is stretched, service interfaces are not supported.

a Click **Add Interface**.

b Enter a name and IP address in CIDR format.

If you configure multicast on this gateway, you must not configure tier-1 addresses as static RP address in the PIM profile.

c Select a segment.

d In the **MTU** field, enter a value between 64

and 9000. e For **URPF Mode**, you can select

**Strict** or **None**.

URPF (Unicast Reverse Path Forwarding) is a security feature. f Add one or more tags.

g In the **ND Profile** field, select or create a profile. h Click **Save**.

15 (Optional) Click **Static Routes** and **Set** to

configure static routes. a Click **Add Static Route**.

b Enter a name and a network address in the CIDR or IPv6 CIDR format. c Click **Set Next Hops** to add next hop information.

d Click **Save**.

What to do next

If you created a tier-1 gateway with Edge clusters for services, you can configure services now.

- For more information about NAT, see [#unique\\_13](#).
- For more information about Gateway Firewall, see [Create Gateway Policies and Rules from Global Manager](#).
- For more information about DNS, see [#unique\\_14](#).
- For more information about DHCP, see [#unique\\_17](#).

## Add a Segment from Global Manager

You can add two kinds of segments: overlay-backed segments and VLAN-backed segments. When you create segments from Global Manager, only overlay-backed segments can span multiple locations.

You can view segments ports from Global Manager, but you cannot create or modify them. If you need to create or modify a segment port, you must do it from the Local Manager.

---

**Important** Do not change the gateway connectivity of a segment in NSX-T Federation. Changing the gateway affects the span of the segment. If the span changes in such a way that it excludes a location, the segment is deleted on the excluded location. You must disconnect all VMs before you shrink the span of a segment.

Prerequisites

Verify that each location has a default overlay transport zone

configured. The default overlay transport zone is used to create global overlay segments. From each Local Manager, select **System > Fabric > Transport Zones**. Select an overlay transport zone, and click **Actions > Set as Default Transport Zone**.

#### Procedure

- 1 From your browser, log in with admin privileges to a Global Manager at <https://<global-manager-ip-address>>.
- 2 Select **Networking > Segments**.

- 3 Click **Add Segment**.
- 4 Enter a name for the segment.
- 5 Select the Connectivity, Traffic Type, and Locations for this segment.

Table 1-2. Segment Configurations

Connectivity	Traffic Type	Location and Transport Zone	Details
A global tier-0 or tier-1 gateway	Overlay	The <b>Location</b> section is populated with the following configurations: <ul style="list-style-type: none"> <li>■ the same locations that are configured on the attached gateway.</li> <li>■ the default overlay transport zone for each location.</li> </ul>	Use this configuration to create a global overlay-backed segment connected to the selected global gateway.
None	VLAN	You must select one location for this segment. You must also select a transport zone from that location.	Use this configuration to create a global VLAN-backed segment to use for a tier-0 external interface.
None	Overlay	No locations or transport zones can be selected.	This segment is created on the Global Manager but is not realized in any Local Managers. You can attach it to a gateway later.

Creating a VLAN-backed segment that is attached to a gateway is not supported.

- 6 Enter the Gateway IP address of the subnet in a CIDR format. A segment can contain an IPv4 subnet, or an IPv6 subnet, or both.
  - If a segment is not connected to a gateway, subnet is optional.
  - If a segment is connected either to a tier-1 or tier-0 gateway, subnet is required.

Subnets of one segment must not overlap with the subnets of other segments



in your network. A segment is always associated with a single virtual network identifier (VNI) regardless of whether it is configured with one subnet, two subnets, or no subnet.

**7 Skip Set DHCP Config.**

Only static bindings are supported on a segment created from Global Manager. See [Features and Configurations Supported in NSX-T Federation](#).

- 8** If the transport zone is of type VLAN, specify a list of VLAN IDs. If the transport zone is of type Overlay, and you want to support layer 2 bridging or guest VLAN tagging, specify a list of VLAN IDs or VLAN ranges

9 (Optional) Select an uplink teaming policy for the segment.

This drop-down menu displays the named teaming policies, if you have added them in the VLAN transport zone. If no uplink teaming policy is selected, the default teaming policy is used.

- Named teaming policies are not applicable to overlay segments. Overlay segments always follow the default teaming policy.
- For VLAN-backed segments, you have the flexibility to override the default teaming policy with a selected named teaming policy. This capability is provided so that you can steer the infrastructure traffic from the host to specific VLAN segments in the VLAN transport zone. Before adding the VLAN segment, ensure that the named teaming policy names are added in the VLAN transport zone.

10 Click **Save**.

11 To continue configuring the segment, click **Yes** when prompted.

12 To select segment profiles, click **Segment Profiles** .

13 To bind a static IP address to the MAC address of a VM on the segment, expand **DHCP Static Bindings**, and then click **Set**.

14 Click **Save**.

## Security

This section provides security set up procedures for Federation.


### Security in Federation

You can create distributed and gateway firewall rules from the Global Manager with global, regional or local spans.

Federation security provides the following benefits:

- Consistent security policy across your deployments managed using Federation.
- Effective disaster recovery ensuring continuity of established security framework.
- Extension of network and security framework to another location if you are running out of compute resources in one location.

Distributed and gateway firewall policies and rules created from the Global Manager are synced to Local Managers and appear in the Local Managers with

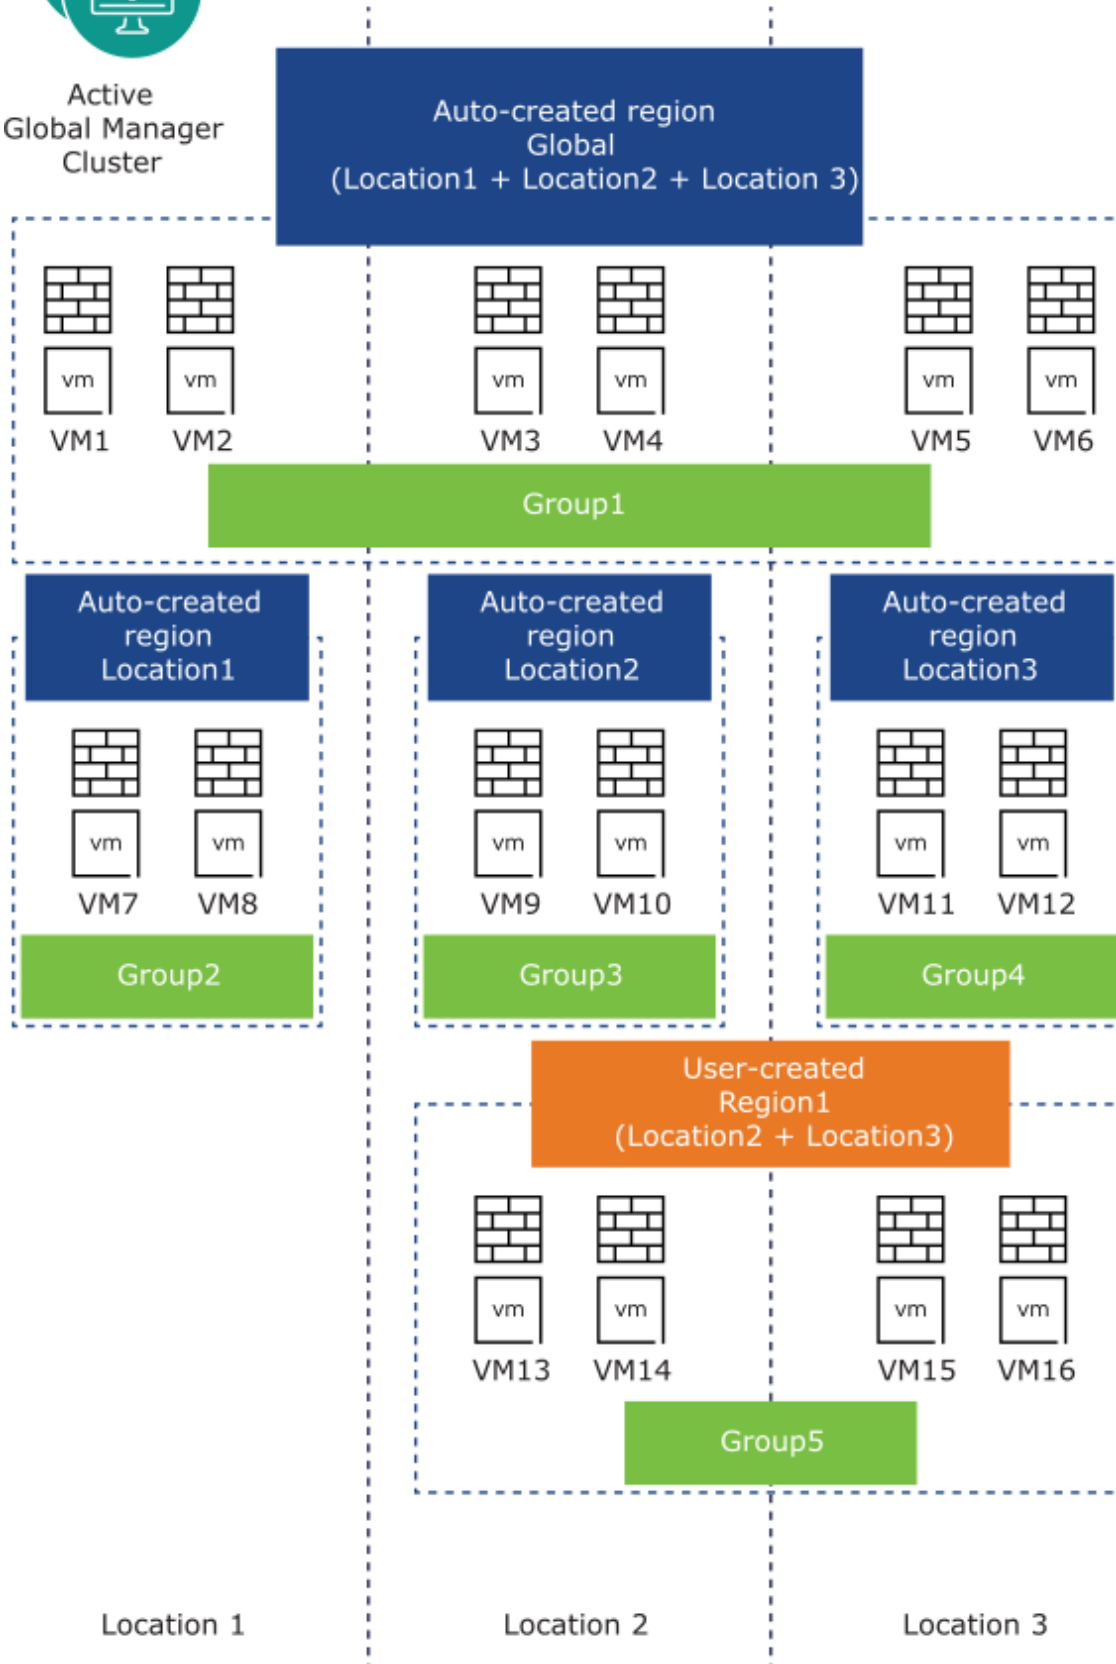
a  icon. You can edit rules created from the Global Manager only from the Global Manager. They cannot be edited from Local Managers.

## NSX-T Federation of Distributed Firewall (DFW) Policies and Rules

Use this example to understand the supported firewall workflows:



Active  
Global Manager  
Cluster



- In the example, the Global Manager has three Local Managers registered with it, named:  
*Location1*, *Location2* and *Location3*.
- The Global Manager auto-creates the following regions:
  - *Global*
  - *Location1*
  - *Location2*
  - *Location3*
- You create a customized region named: **Region1** that includes Local Managers *Location2* and *Location3*.
- You create the following groups:
  - **Group1**: Region *Global*.
  - **Group2**: Region *Location1*.
  - **Group3**: Region *Location2*.
  - **Group4**: Region *Location3*.
  - **Group5**: Region **Region1**.

## DFW Policies and Rules

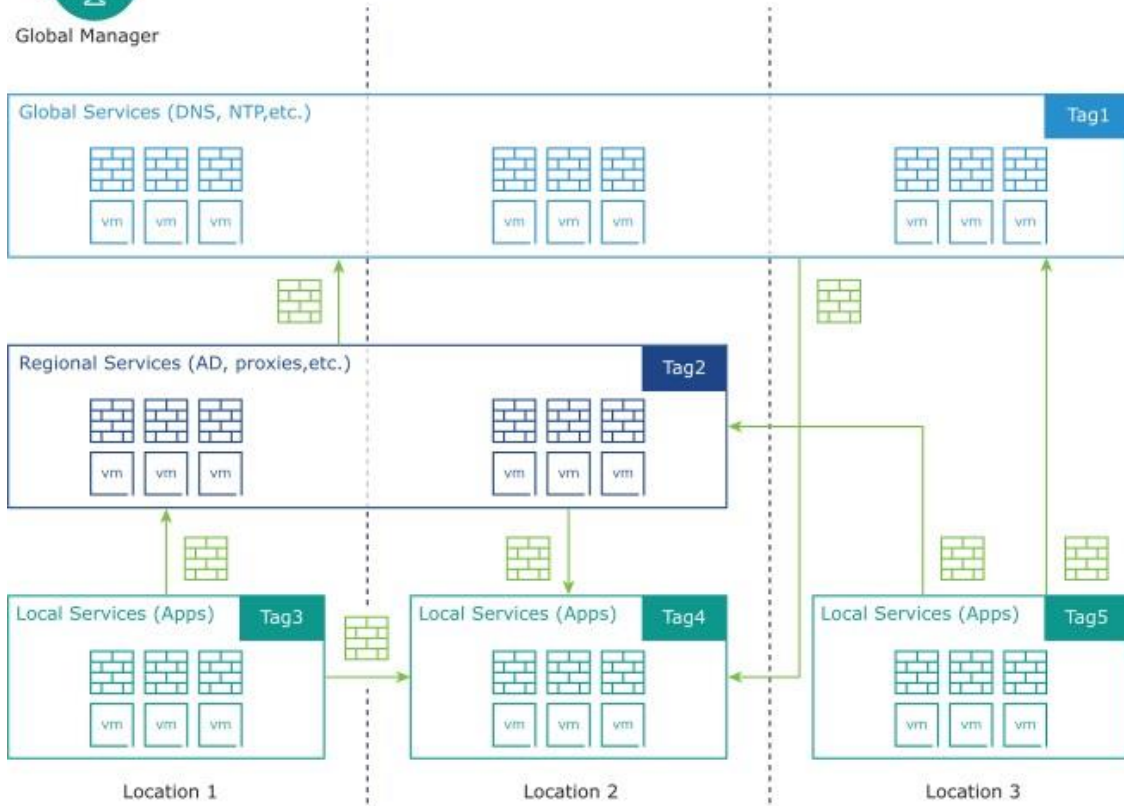
The following use cases are supported:

- **Group Span**: You can create groups in the Global Manager with a global, local or regional span. See [Create Groups from Global Manager](#) .
- **Dynamic Groups**: You can create groups based on dynamic criteria, such as tags.
- **DFW Policy Span**: DFW policies can be applied to a global, regional or local span.
- **DFW Rule's Source and Destination Groups**: Either all the groups in the source field or all the groups in the destination field must match the DFW policy's span. The system auto-creates groups in locations that are outside the policy's span.



Global Manager

- Global Manager can create groups with spans that are global, regional or local.
- Groups can be dynamic, based on tags.
- Firewall rules can be applied to a mixed span of groups.



Refer to the table for examples of valid and invalid source and destination groups in DFW rules:

Table 1-3. Valid Source and Destination for a DFW rule based on the DFW Policy's Span

DFW Policy Span (Applied To)	Scenarios supported in DFW rules
<p><i>Global</i></p> <p>From the example, this region contains the following groups:</p> <ul style="list-style-type: none"> <li>■ <b>Group1</b></li> </ul>	<p>For a DFW policy with the span of <i>Global</i> region, all groups are allowed in the DFW rule's source and destination.</p> <p>Following are some typical scenarios that are supported, using our example:</p> <ul style="list-style-type: none"> <li>■ <b>Source: Group2; Destination Group3</b></li> <li>■ <b>Source: Group3; Destination Group4</b></li> <li>■ <b>Source: Group4; Destination: Any</b></li> <li>■ <b>Source: Group1; Destination Group2.</b></li> </ul>
<p><i>Location1</i> : auto-created region for the Local Manager in location 1.</p> <p>From the example, this region contains the following groups:</p> <ul style="list-style-type: none"> <li>■ <b>Group2</b></li> </ul>	<p>For a DFW policy with the span of one location: <i>Location1</i> in this example, either the source or the destination group for the DFW rule must belong to <i>Location1</i>.</p> <p>The following scenarios are supported:</p> <ul style="list-style-type: none"> <li>■ <b>Source: Group2; Destination Group2</b></li> <li>■ <b>Source: Group3; Destination Group2.</b></li> <li>■ <b>Source: Group2; Destination Group4.</b></li> <li>■ <b>Source Group1; Destination Group2.</b></li> </ul> <p>The following is an example of unsupported group selections for this policy span. Both the source and the destination groups are outside the policy's span:</p> <ul style="list-style-type: none"> <li>■ <b>Source Group5; Destination Group3.</b></li> <li>■ <b>Source Group1; Destination Group3.</b></li> </ul>
<p><b>Region1</b> : user-created region that spans <i>Location2</i> and <i>Location3</i>.</p> <p>From the example, this region contains the following groups:</p> <ul style="list-style-type: none"> <li>■ <b>Group5</b></li> </ul>	<p>For a DFW policy with the span of a user-created region: <b>Region1</b> in this example, either the source or the destination group for the DFW rule must contain locations that belong to <b>Region1</b>.</p> <p>The following scenarios are supported:</p> <ul style="list-style-type: none"> <li>■ <b>Source: Group5; Destination Group2.</b></li> <li>■ <b>Source: Group2; Destination Group5.</b></li> <li>■ <b>Source: Group2; Destination Group3.</b></li> <li>■ <b>Source: Group3; Destination Group4.</b></li> <li>■ <b>Source: Any ;Destination: Group5</b></li> <li>■ <b>Source Group4; Destination Any</b></li> </ul> <p>The following is an example of unsupported group selections for this policy span. Both the source and the destination groups are outside the policy's span:</p> <ul style="list-style-type: none"> <li>■ <b>Source Group2; Destination Group2.</b></li> </ul>





- If a group contains segments, the span of the DFW policy must be greater than or equal to the span of the segment. For example, if you have a group containing a segment whose span is *Location1*, the DFW policy cannot be applied to region **Region1** because it only contains *Location2* and *Location3*.

## NSX-T Federation of Gateway Firewall Policies and Rules

Gateway firewall rules can be applied to all the locations included in the gateway's span, or all interfaces of a particular location, or specific interfaces of one or more locations.

---

**Note** The span of the source and destination groups for gateway firewall rules must be the same as or a subset of the gateway's span on which you are creating the rule.

Table 1-4. Span Options for Gateway Firewall Rules

Gateway Firewall Rule's Span (Applied To)	Applies to
Apply rule to gateway	The rule applies to all interfaces attached to this gateway, in all locations that this gateway is stretched to.
Select a location and then select Apply rule to all Entities.	The rule applies only to the selected location.
Select a location and then select interfaces from that location. Repeat for other locations, selecting interfaces for each location that you want to apply the rule to.	The rule applies only to the selected interfaces.

## Create a Region from Global Manager

Each location added to the Global Manager automatically becomes a region. You can also create customized regions.

Use regions to create focused groups for security and networking policies. Some regions are created automatically after you onboard locations in Global Manager. You can add more regions as necessary.

---

**Note** Each location can be a part of only one customized region.

---

The following regions are added by default:

- A Global region including all the locations added to the Global Manager.
- One region for each location added to the Global Manager.

For existing regions, you can view the following information:

- Name of the region.

- Locations included in the region.
- Groups the region belongs to.
- Security/Network policies the region is a part of.

## Prerequisites

Refer to [Security in Federation](#) for details on the implication of the span of regions and groups in creating and maintaining security policies and rules.

## Procedure

- 1 Select **Inventory > Regions**.
- 2 Click **Add Region**.
- 3 Provide the following information:

Option	Description
<b>Name</b>	Provide a name for the region, for example, EMEA, or APAC.
<b>Locations</b>	Select the locations that you want to include in this region.

- 4 Click **Save**.

The region with the specified locations is created.

## What to do next

[Create Groups from Global Manager](#) .

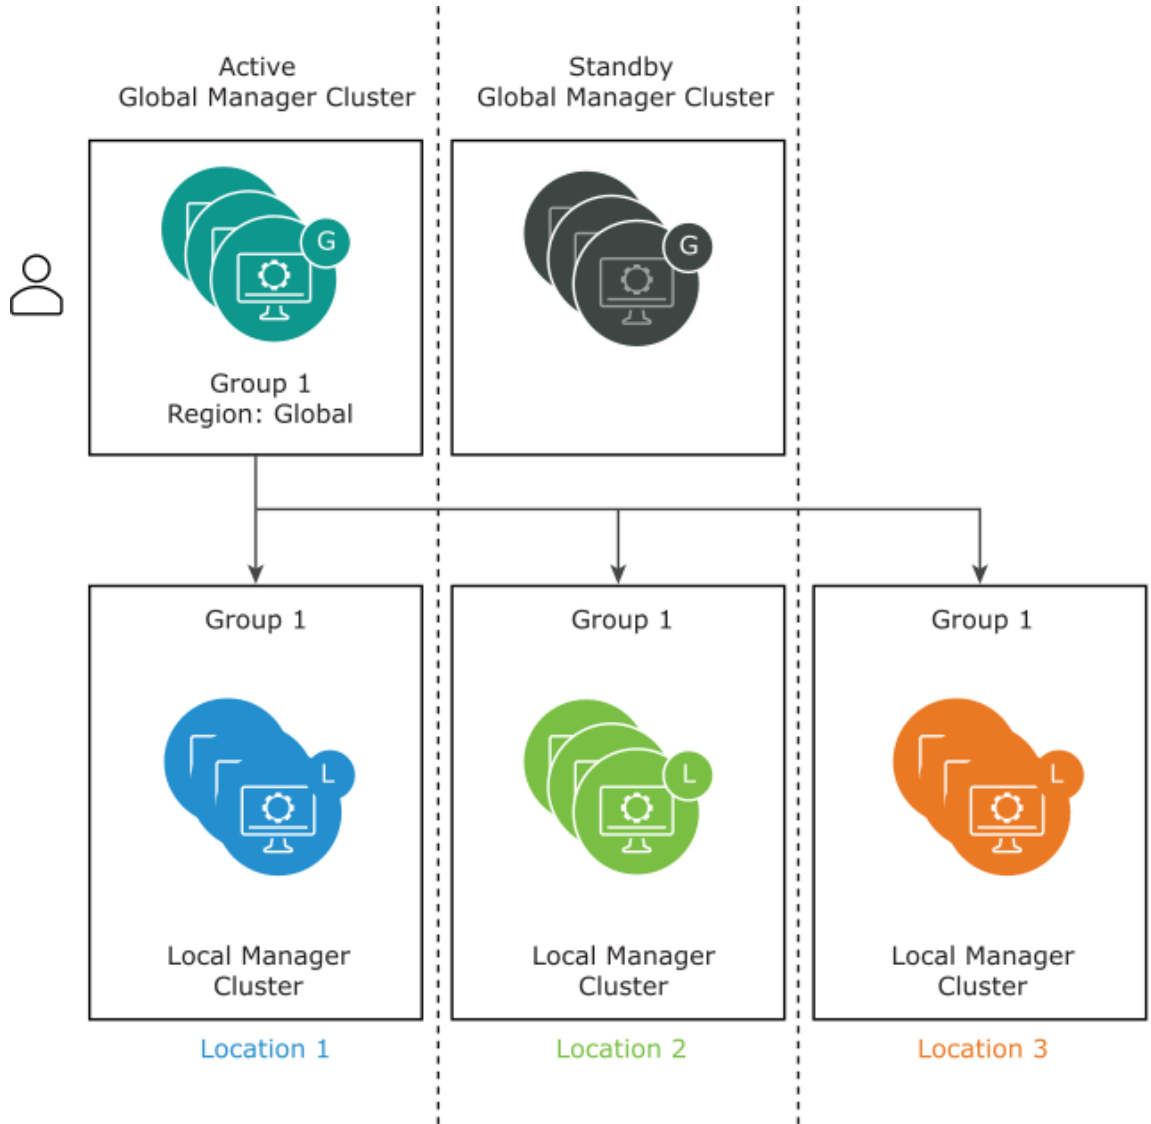
## Create Groups from Global Manager

Create Groups from Global Manager that apply globally across your NSX-T Data Center deployments or cover selected locations or regions.

### Group Span

When you create a group from the Global Manager, you select a region for the group. The group is synced with all locations in that region. A global region containing all locations, and a region for each location that has been added to the Global Manager are available automatically as regions you can select for a group's span. You can create customized regions before you create groups. See [Create a Region from Global Manager](#).

In this example, **Group1** is created in the Global region, and is therefore synced with all Local Managers.



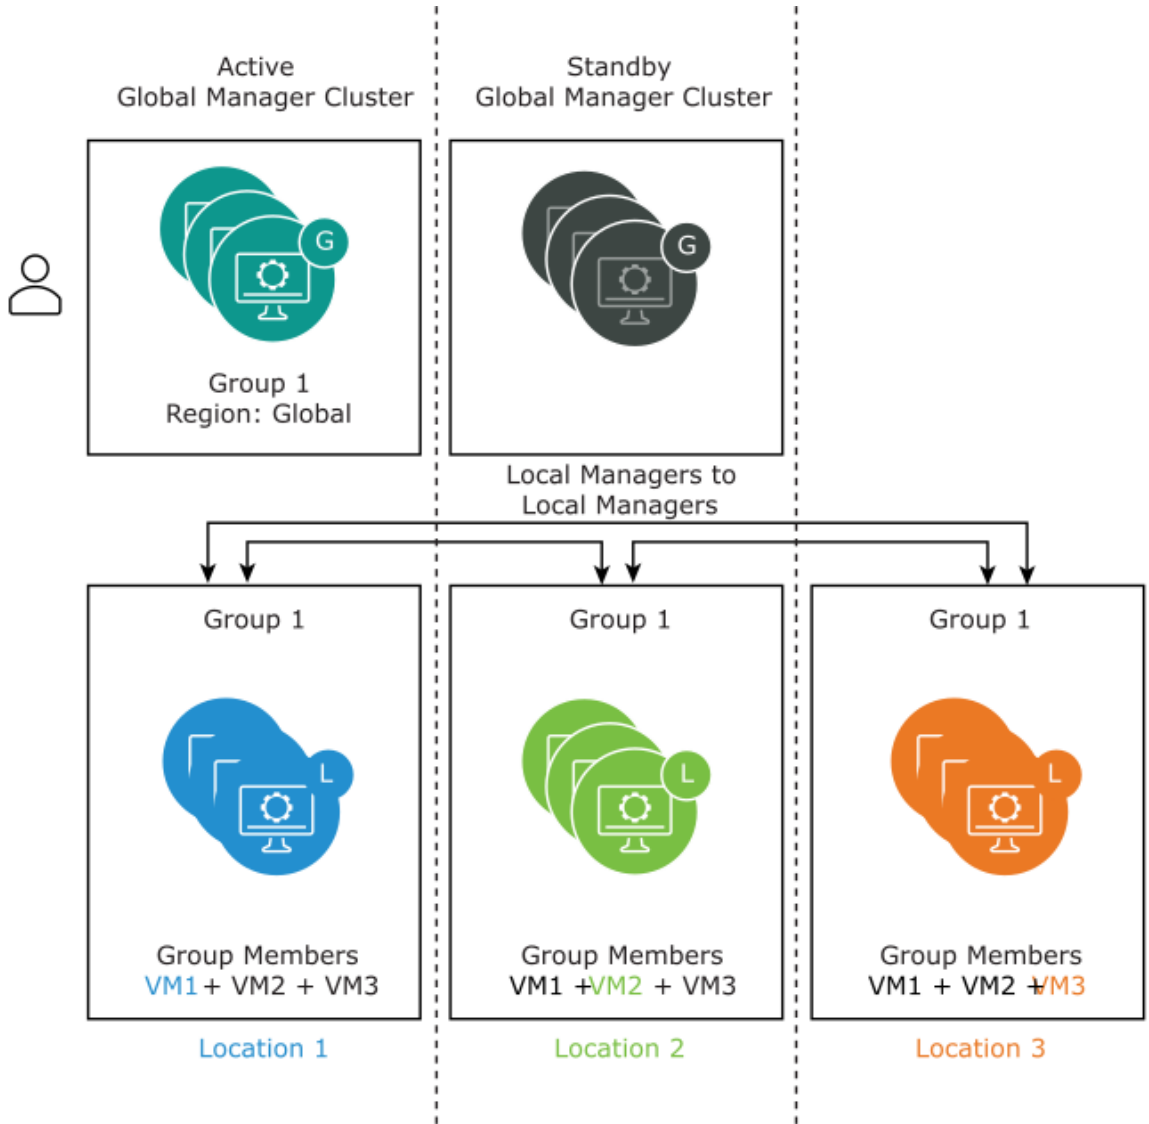
## Dynamic Groups

If a group that spans more than one location has dynamic membership, you need information from each location to list the group membership.

In this example, **Group1** has the following members:

- **VM1** in *Location1*
- **VM2** in *Location2*
- **VM3** in *Location3*

Each Local Manager syncs its dynamic group membership with the other Local Managers. As a result, each Local Manager has a complete list of group members.



## Nested Groups

For groups created from the Global Manager, you can add another group as a member if it has an equal or smaller span than the group's region.

**Note** If you are using NSX-T Data Center version 3.0.0, you can add a group as a member of another group only if the span of both the groups is exactly the same.

Extending the example using **Region1** that contains *Location2* and *Location3*, note the following additional configurations:

Task	Effect
------	--------

---

From Global Manager, create **Group-Loc2** with region *Location2*.

- **Group-Loc2** is created in Global Manager.
  - **Group-Loc2** is created in the Local Manager *Location2*.
-

<p>From Global Manager, create group <b>Group-Region1</b> with region <b>Region1</b>. Add <b>Group-Loc2</b> as a member. This is a nested group.</p>	<ul style="list-style-type: none"> <li>■ <b>Group-Region1</b> is created in Global Manager.</li> <li>■ <b>Group-Region1</b> is created in <i>Location2</i> and <i>Location3</i>.</li> <li>■ <b>Group-Loc2</b> is created in Local Manager <i>Location3</i>.</li> </ul>
<p>From Global Manager, navigate to <b>Inventory &gt; Regions</b> and edit <b>Region1</b> to remove <i>Location2</i>.</p>	<p>This action is not allowed because of the nested group <b>Group-Region1</b>.</p>

See [#unique\\_46](#) for detailed steps for creating groups.

## Create DFW Policies and Rules from Global Manager

You can create security policies and DFW rules to apply to multiple locations registered with the Global Manager.

### Prerequisites

Ensure that you have already created any customized regions that you want to use for firewall rules. See [Create a Region from Global Manager](#).

### Procedure

- 1 From your browser, log in with Enterprise Admin or Security Admin privileges to a Global Manager at `https://<global-manager-ip-address>`.
- 2 Select **Security > Distributed Firewall**
- 3 Ensure that you are in the correct pre-defined category, and click **Add Policy**. For more about categories, see [#unique\\_47](#).

---

**Note** Ethernet, Emergency categories and Default Policy are not supported on Global Manager.

- 4 Click **Add Policy**.
- 5 Enter a **Name** for the new policy section.
- 6 Click the pencil icon next to **Applied To** to set the span of this policy.
- 7 In the **Set Applied To** dialog box, you can make the following selections:
  - **Region:** select which Local Managers to apply the policy to. Each Local Manager is automatically added as a region. You can also create customized regions. See [Create a Region from Global Manager](#).
  - **Select Applied To:** By default, policy is applied to **DFW**, that is, the policy is applied to all the workloads on the Local Managers based on the selected region for this policy. You can also apply a policy to selected groups. Applied to defines the scope of enforcement per policy, and is used mainly for resource optimization on ESXi and KVM

hosts. It helps  
in defining a targeted policy for specific zones, tenants and  
application without interfering with other policy defined for other  
tenants, zones & applications.

See [DFW Policies and Rules](#) to understand how the span of the policy  
determines whether your DFW rule is valid or invalid.



8 To configure the following policy settings, click the gear icon:

Option	Description
TCP Strict	<p>A TCP connection begins with a three-way handshake (SYN, SYN-ACK, ACK) and typically ends with a two-way exchange (FIN, ACK). In certain circumstances, the distributed firewall (DFW) might not see the three-way handshake for a particular flow ( due to asymmetric traffic or the distributed firewall being enabled while a flow exists). By default, the distributed firewall does not enforce the need to see a three-way handshake, and picks up sessions that are already established. TCP strict can be enabled on a per section basis to turn off mid-session pick-up and enforce the requirement for a three-way handshake. When enabling TCP strict mode for a particular DFW policy, and using a default ANY-ANY Block rule, packets that do not complete the three-way handshake connection requirements and that match a TCP-based rule in this section are dropped. Strict is only applied to stateful TCP rules, and is enabled at the distributed firewall policy level. TCP strict is not enforced for packets that match a default ANY-ANY Allow which has no TCP service specified.</p>
Stateful	<p>A stateful firewall monitors the state of active connections and uses this information to determine which packets to allow through the firewall.</p>
Locked	<p>The policy can be locked to prevent multiple users from editing the same sections. When locking a section, you must include a comment.</p> <p>Some roles such as enterprise administrator have full access credentials, and cannot be locked out. See <a href="#">#unique_48</a>.</p>

9 Click **Publish**. Multiple policies can be added, and then published together at one time.

The new policy is shown on the screen.

10 Select a policy section and click **Add Rule**.

- 11 Enter a name for the rule.
- 12 The Source and Destination are validated based on the DFW policy's span. See [DFW Policies and Rules](#) for more information.
  - If the DFW policy is applied to a location, for example, **Loc1**, source or destination can be either the keyword **ANY** or a group that belongs to **Loc1**.
  - If DFW policy is applied to a user-created region, for example, **Region1** source or destination can be either the keyword **ANY** or a group that has the same span as **Region1** or spans a location in **Region1**.

- If DFW policy is applied to **Global**, source or destination can be anything.

**Note** Active Directory and IDFW are not supported for NSX-T Federation, that is, you cannot use these features from the Global Manager.

- a In the **Sources** column, click the pencil icon and select the source of the rule.
  - b In the **Destinations** column, click the pencil icon and select the destination of the rule. If not defined, the destination matches any.
- 13 In the **Services** column, click the pencil icon and select services. The service matches any if not defined.
- 14 In the **Profiles** column, click the edit icon and select a context profile, or click **Add New Context Profile**. See [#unique\\_49](#).
- 15 Click **Apply** to apply the context profile to the rule.
- 16 By default, the **Applied to** column is set to DFW, and the rule is applied to all workloads. You can also apply the rule or policy to a selected group. **Applied to** defines the scope of enforcement per rule, and is used mainly for optimization of resources on ESXi and KVM hosts. It helps in defining a targeted policy for specific zones, tenants, and applications without interfering with other policy defined for other tenants and zones and applications.

**Note** You cannot select the following types of groups in **Applied-to**:

- a group with IP or MAC addresses
- an Active Directory user group

- 17 In the **Action** column, select an action.

Option	Description
<b>Allow</b>	Allows all L3 or L2 traffic with the specified source, destination, and protocol to pass through the current firewall context. Packets that match the rule, and are accepted, traverse the system as if the firewall is not present.
<b>Drop</b>	Drops packets with the specified source, destination, and protocol. Dropping a packet is a silent action with no notification to the source or destination systems. Dropping the packet causes the connection to be retried until the retry threshold is reached.

---

**Reject**

Rejects packets with the specified source, destination, and protocol. Rejecting a packet is a more graceful way to deny a packet, as it sends a destination unreachable message to the sender. If the protocol is TCP, a TCP RST message is sent. ICMP messages with administratively prohibited code are sent for UDP, ICMP, and other IP connections. One benefit of using Reject is that the sending application is notified after only one attempt that the connection cannot be established.

---

**18** Click the toggle button to enable or disable the rule.

19 Click the gear icon to configure the following rule options:

Option	Description
<b>Logging</b>	Logging is turned off by default. Logs are stored at /var/log/ dfwpktlogs.log on ESXi and KVM hosts.
<b>Direction</b>	Refers to the direction of traffic from the point of view of the destination object. IN means that only traffic to the object is checked, OUT means that only traffic from the object is checked, and In/Out, means that traffic in both directions is checked.
<b>IP Protocol</b>	Enforce the rule based on IPv4, IPv6, or both IPv4-IPv6.
<b>Log Label</b>	Log Label appears in the Firewall Log when logging is enabled.

20 Click **Publish**. Multiple rules can be added and then published together at one time.

21 On each policy, click **Check Status** to view the status of rules it contains, per location. You can click **Success** or **Failed** to open the policy status window.

22 Click **Check Status** to check the realization status of policies that are applied to Transport Nodes on different locations.

## Create Gateway Policies and Rules from Global Manager

You can create gateway firewall policies and rules to be applied to multiple locations or selected interfaces for particular locations, from the Global Manager.

Tier-0 or tier-1 gateways created from the Global Manager span all or a set of locations. You have a few options when applying gateway firewall rules created from the Global Manager: Gateway firewall rules can be applied to all the locations included in the gateway's span, or all interfaces of a particular location, or specific interfaces of one or more locations.

On the Local Manager rules are enforced in the following order:

- 1 Any rules you create from the Global Manager, that get successfully realized on the Local Manager, are enforced first.
- 2 Any rules that you create from the Local Manager are enforced next.
- 3 The last rule enforced is the default gateway firewall rule. This is the allow-all or deny-all rule applicable to all locations and all workloads. You can edit the behavior for this default rule from the Global Manager.

Procedure

- 1 From your browser, log in with Enterprise Admin or Security Admin privileges to the Global Manager at <https://<global-manager-ip-address>>.
- 2 Select **Security > Gateway Firewall**.

- 3 Ensure that you are in the correct pre-defined category. Only **Pre Rules**, **Local Gateway** and **Default** categories are supported on Global Manager. To define policy under the **Local Gateway** category, click the category name from the **All Shared Rules** tab or directly click the **Gateway Specific Rules** tab.

Select a tier-0 or tier-1 gateway from the drop-down menu next to **Gateway**. The span of the tier-0 or tier-1 gateway you selected becomes the default span of the Gateway Firewall policy and rule. You can reduce the span but not expand it.

- 4 Click **Add Policy**.
- 5 Enter a **Name** for the new policy section.
- 6 (Optional) Click the gear icon to configure the following policy settings:

Settings	Description
TCP Strict	A TCP connection begins with a three-way handshake (SYN, SYN-ACK, ACK), and typically ends with a two-way exchange (FIN, ACK). In certain circumstances, the firewall may not see the three-way handshake for a particular flow (i.e. due to asymmetric traffic). By default, the firewall does not enforce the need to see a three-way handshake, and will pick up sessions that are already established. TCP strict can be enabled on a per section basis to turn off mid-session pick-up, and enforce the requirement for a three-way handshake. When enabling TCP strict mode for a particular firewall policy and using a default ANY-ANY Block rule, packets that do not complete the three-way handshake connection requirements and that match a TCP-based rule in this policy section are dropped. Strict is only applied to stateful TCP rules, and is enabled at the gateway firewall policy level. TCP strict is not enforced for packets that match a default ANY-ANY Allow which has no TCP service specified.
Stateful	A stateful firewall monitors the state of active connections, and uses this information to determine which packets to allow through the firewall.

---

Locked	The policy can be locked to prevent multiple users from making changes to the same sections. When locking a section, you must include a comment.
--------	--

---

- 7 Click **Publish**. Multiple Policies can be added, and then published together at one time.

The new policy is shown on the screen.

- 8 Select a policy section and click **Add Rule**.

- 9 Enter a name for the rule.

- 10 In the **Sources** column, click the edit icon and select the source of the rule. The source group must have the same or a subset of the gateway's span.



- 11 In the **Destinations** column, click the edit icon and select the destination of the rule. If not defined, the destination matches any. The destination group must have the same or a subset of the gateway's span.
- 12 In the **Services** column, click the pencil icon and select services. The service matches any if not defined. Click **Apply** to save.
- 13 In the **Profiles** column, click the edit icon and select a context profile, or click **Add New Context Profile**. See [#unique\\_49](#).

---

**Note** Context profiles are not supported for tier-0 gateways. You can apply L7 context profiles to tier-1 gateways.

- 14 Click the pencil icon in the **Applied to** column. In the **Applied To** dialog box:

Applied To Selection	Result
Select Apply rule to gateway	The gateway firewall rule is applied to all locations covered by the gateway's span. If you add another location to the gateway, this gateway firewall rule automatically gets applied to the location.
Select a location and then select Apply rules to all Entities	Apply this rule to all interfaces in the selected location.
Select a location and then select interfaces for that location	Apply the rule only to selected interfaces in one or more locations.

---

**Note** There is no default selection for **Applied To**. You must make a selection to be able to publish this rule.

- 15 In the **Action** column, select an action.

Option	Description
<b>Allow</b>	Allows all traffic with the specified source, destination, and protocol to pass through the current firewall context. Packets that match the rule, and are accepted, traverse the system as if the firewall is not present.
<b>Drop</b>	Drops packets with the specified source, destination, and protocol. Dropping a packet is a silent action with no notification to the source or destination systems. Dropping the packet causes the connection to be retried until the retry threshold is reached.
<b>Reject</b>	Rejects packets with the specified source, destination, and protocol. Rejecting a packet sends a destination unreachable message to the sender. If the protocol is TCP, a TCP RST message is sent. ICMP messages with administratively prohibited code are sent for UDP, ICMP, and other IP connections. The sending application is

---

notified after one attempt that the connection cannot be established.

---

**16** Click the status toggle button to enable or disable the rule.

17 Click the gear icon to set logging, direction, IP protocol, tag, and notes.

Option	Description
<b>Logging</b>	<p>Logging can be turned off or on. You can access logs using the following NSX CLI command on NSX Edge:</p> <pre>get log-file syslog   find datapathd.firewallpkt</pre> <p>Logs can also be sent to an external syslog server.</p>
<b>Direction</b>	<p>The options are <b>In</b>, <b>Out</b>, and <b>In/Out</b>. The default is <b>In/Out</b>. This field refers to the direction of traffic from the point of view of the destination object. <b>In</b> means that only traffic to the object is checked, <b>Out</b> means that only traffic from the object is checked, and <b>In/Out</b> means that traffic in both directions is checked.</p>
<b>IP Protocol</b>	<p>The options are <b>IPv4</b>, <b>IPv6</b>, and <b>IPv4_IPv6</b>. The default is <b>IPv4_IPv6</b>.</p>
<b>Log Label</b>	<p>Log label that has been added to the rule.</p>

**Note** Click the graph icon to view the flow statistics of the firewall rule. You can see information such as the byte, packet count, and sessions.

18 Click **Publish**. Multiple rules can be added and then published together at one time.

19 Click **Check Status** to view the realization status of policy applied to gateways through edge nodes in different locations. You can click **Success** or **Failed** to open the policy status window.

## Backup and Restore

### Backup and Restore in NSX-T Federation

You can configure and start backups for Global Manager and each Local Manager from within the Global Manager.

**Important** Starting in NSX-T Data Center 3.0.1, restoring Global Manager to an FQDN is supported. If you are using NSX-T Data Center 3.0.0, do not use FQDN for the Global Manager. Only IP address backups are supported for the Global Manager appliance in NSX-T Data Center 3.0.0.

- Log in to the active Global Manager and select . Each Global Manager and Local Manager in the environment is listed. See [#unique\\_51](#) for instructions.

- You cannot restore Local Manager from within the Global Manager. To restore a Local Manager backup, log in to the Local Manager to restore. See [#unique\\_52](#) for instructions.
- The system treats backup and restore operations as specific to each appliance, whether it is the Global Manager or the Local Manager you are backing up or restoring. The Global Manager's backup contains a backup of the database of that appliance only. The Local Manager contains a backup of the database and inventory of that appliance only.

- If you are restoring a Global Manager and a Local Manager, select backup timestamps of each appliance as close to each other as possible.
- After each appliance is restored, the async replicator service restores communication between the Global Manager and each Local Manager.

## Backup Scenarios in NSX-T Federation

Scenario	Backup Workflow
Global Manager has any of the following changes: <ul style="list-style-type: none"> <li>■ networking configuration</li> <li>■ security configuration</li> </ul>	Back up only the Global Manager.
A Local Manager has any of the following changes: <ul style="list-style-type: none"> <li>■ networking configuration</li> <li>■ security configuration</li> <li>■ fabric changes, such as:                             <ul style="list-style-type: none"> <li>■ host transport node added or removed (ESXi or KVM).</li> <li>■ edge transport nodes added or removed (VM or bare metal).</li> </ul> </li> </ul>	If you have any local configurations, back up the Local Manager.

## Restore Scenarios in NSX-T Federation

Scenario	Restore Workflow
Global Manager is lost.	Restore the Global Manager. When restored, the Global Manager pushes configurations to the Local Managers registered with it.
A Local Manager is lost.	Restore the Local Manager. When restored, configurations from the Global Manager are synchronized with the Local Manager.
Both the Global Manager and the Local Manager are lost.	If you are restoring both the Global Manager and the Local Manager, use the latest backups of each appliance. When the Global Manager and the Local Manager are restored, the Global Manager pushes the configurations to the Local Manager.  You must manually resolve any discrepancies in inventory and fabric related changes between the Local Manager and the Global Manager.

## Disaster Recovery

## Disaster Recovery for Global Manager

In an NSX-T Federation environment, if you lose your active Global Manager, you can switch to the standby Global Manager.

The workflows described here use the following scenario where GM denotes the Global Manager appliance:

- You have a GM cluster in location Loc1. You name this GM **GM-Loc1** and set it as the **Active** GM.
- You have another GM cluster in location Loc2. You name this GM **GM-Loc2** and set it as the **Standby** GM.

### Planned switchover to Standby GM

If you want to set the standby GM - **GM-Loc2** - as active, while the active GM - **GM-Loc1** - is running, do the following:

- 1 Log in to the standby GM - **GM-Loc2**.
- 2 Select **Make Active** from the **Actions** drop-down menu.

The system starts the process of making **GM-Loc2** active. After the process completes, **GM-Loc2** gets the status of **Active** and **GM-Loc1** gets the status of **Standby**.

### Unplanned switchover to the Standby GM

If you lose the active GM - **GM-Loc1** - do the following:

- 1 Log in to the standby GM - **GM-Loc2** - and select **Make Active** from the **Actions** drop-down menu.
- 2 (Optional) If you also lost the Local Manager at this site - **Loc1**:
  - a Follow the network recovery workflow to move stretched tier-0 and tier-1 gateways to the secondary site. See instructions at [Network Recovery for Local Managers](#).
  - b Recover the compute VMs using your preferred method, for example, VMware Site Recovery Manager.

The system starts the process of making **GM-Loc2** active. After the process completes, **GM-Loc2** gets the status of **Active**.

If **GM-Loc1** is back online after **GM-Loc2** has become active, the status of **GM-Loc1** is set to **NONE**. You can make **GM-Loc1** standby by following the steps below:

- 1 Log in to the active GM - **GM-Loc2**.

VMware,

- 2 From the tile for **GM-Loc1** showing the status of **None**, select **Make Standby** from the **Actions** drop-down menu.

See section 4.4 titled *Disaster Recovery* in the [NSX-T Data Center Multi-location Design Guide](#) for more details.

# Network Recovery

## Network Recovery for Local Managers

If a Local Manager is lost, you can recover networking configurations from it using the auto-detected Network Recovery option in the Global Manager.

You must have at least one stretched tier-0 or tier-1 gateway set up designating a Location Manager as primary. The loss of this primary Location Manager for the tier-0 or tier-1 gateway triggers the option of network recovery in the Global Manager.

- The Global Manager detects the loss of connection and prompts you to perform **Network Recovery**.
- In the first step of recovery, you recover the tier-0 gateway. You can change the preferred primary location if you want it to be different from the one you set in the fallback preference.
- In the second step, you select a preferred primary location for tier-1 gateways that have a subset of the span of the locations covered by the tier-0 network. The preferred primary location for such tier-1 gateways would be different from tier-0 gateways and you must either accept the fallback preference established by the tier-0 gateway, or elect not to move the gateway.
- In the final step, you can view the list of networking constructs that cannot be recovered because they do not have a secondary location configured.

---

**Note** If you have a tier-0 and tier-1 gateway set up using a Location Manager as primary, but the tier-0 and tier-1 gateway do not have any services attached to them, for example, tier-0 and tier-1 without NAT and firewall, then the data plane traffic still works after the loss of the primary Location Manager. For tier-0/tier-1 configuration without service, Network Recovery is not mandatory for the recovery of data plane, even though the Network Recovery option appears in the Global Manager.

### Procedure

- 1 From your browser, log in with admin privileges to the active Global Manager at `https://<global-manager-ip-address>`.



- 2 Select **System > Location Manager**.
- 3 A banner appears on this page noting the location that is down. Click **Network Recovery** on the banner and start the workflow for **Location Disaster Recovery** in the following steps.
- 4 **Tier-0 Gateways:** For each tier-0 gateway that has the failed location set as primary, you have the option to select a new primary location. This new primary location can be different from the fallback preference you elected when creating the tier-0 gateway. You can also elect to not move the tier-0 gateway. Click **Apply Configuration** for each tier-0 gateway after selecting a new primary location or retaining the priority set earlier.
- 5 Click **Next**.

- 6 Tier-1 gateways are listed for recovery only if their span differs from the span of the tier-0 gateway. If tier-1 gateways follow the same span as the tier-0 gateway, the same locations are selected to be primary as for tier-0 gateways. For a different span, you can either select a different location as primary or elect to not move the tier-1 gateway at all.
- 7 After you make your selections for each tier-1 gateway, click **Accept** and **Next** to proceed.
- 8 Under **Single Location Entities** you can see a list of tier-0 and tier-1 gateways that cannot be moved to a new primary location because they exist only in the failed location. Click **Next** to proceed.

#### Results

The stretched tier-0 and tier-1 gateways are moved to the new location which that you designated as primary.

See section 4.4.2 titled *Data Plane Recovery* in the [NSX-T Data Center Multi-location Design Guide](#) for more details.