



Configuring Tenants and Tenant Domains, Release 4.2.1

Table of Contents

New and changed information	1
Understand tenants and tenant domains	2
Key terms	2
Understand tenants	3
Understand tenant domains	4
Understand how users work with tenants and tenant domains	5
Tenants in Nexus Dashboard and ACI	5
AAA defined users	6
NX-OS AAA users	6
Orchestration AAA users	7
Guidelines and limitations	8
General guidelines and limitations: Tenants and tenant domains	8
Guidelines and limitations: Tenants and object names	9
Guidelines and limitations: Tenants and VXLAN-ACI fabric groups	10
Upgrade considerations	11
Guidelines and limitations: Upgrades	11
Understand tenant domains and upgrades	11
Prior to the upgrade	12
After the upgrade	12
Navigate to the Multi-tenancy page	14
Configuring tenants	15
Understand the Tenants page	15
Create a tenant	15
View tenant information	16
Edit a tenant	17
Associate fabrics with a tenant	17
Associate tenants with a fabric group	18
Edit fabrics that are associated with a tenant	20
Import a tenant from ACI	20
Delete a tenant	21
Working with tenants under Segmentation and security	21
View tenants information	21
Deploying tenants	22
Import tenant policies	23
Understanding the process for importing tenant policies from ACI fabrics into VXLAN-ACI fabric groups	26
Mapping between ACI and Nexus Dashboard policies	26
Import tenant policies from ACI fabrics	27
Configuring tenant domains	35
Understand the Tenant domains page	35
Create a tenant domain	35

Edit a tenant domain	36
Delete a tenant domain	36
Information that is displayed or not displayed	37
Information that is displayed or not displayed based on tenant and tenant domain associations	37
Information displayed and hidden: Admin user role	38
Information displayed and hidden: User1 user role	38
Information displayed and hidden: User2 user role	38
Information displayed and hidden: User3 user role	38
Anomalies information that is displayed or not displayed	39
Endpoints	41
Connectivity analysis	43
L3 neighbors	44
Routes	44
Switch view: Multicast	45
Additional areas affected by tenant and tenant domain creation	46
Multi-tenancy considerations for flow telemetry and traffic analytics	48
Flow configuration	48
Ingest	48
Analysis/visualization	49
Anomalies	49
Multi-tenancy considerations for endpoints and multicast routes	49
Multi-tenancy considerations for connectivity analysis	49
GUI changes	51
Legacy Tenant CRUD APIs	52
Copyright	55

New and changed information

The following table provides an overview of the significant changes up to this current release. The table does not provide an exhaustive list of all changes or of the new features up to this release.

Release Version	Feature	Description
Nexus Dashboard release 4.2.1	Introduction of tenants and tenant domains	Beginning with Nexus Dashboard release 4.2.1, you can now use tenants and tenant domains in Nexus Dashboard to unify the configuration of networking policies that are applied to fabrics, regardless of whether that fabric is an NX-OS or an ACI fabric.

Understand tenants and tenant domains

Prior to Nexus Dashboard release 4.2.1, tenants have been available in the Nexus Dashboard Orchestrator and ACI products, but were not used in Nexus Dashboard. Beginning with Nexus Dashboard release 4.2.1, you can now use tenants and tenant domains in Nexus Dashboard to help in separating and isolating networking policy configurations that are applied to NX-OS or ACI fabrics into different administrative containers.

- Since Orchestration (Nexus Dashboard Orchestrator in previous releases) works exclusively with ACI fabrics, configuring and using tenants and tenant domains is a mandatory task for any configuration on an ACI fabric.
- However, for NX-OS fabrics, configuring and using tenants and tenant domains is an optional task, not a mandatory one. You can continue to configure users, roles and security, as described in [Configuring Users, Roles, and Security](#), without using tenants. Configure and use tenants and tenant domains for NX-OS fabrics only if you want to take advantage of the functionality as described in the following sections.

Support is also available for importing tenant policies from ACI fabrics using Nexus One, which allows you to integrate both ACI and NX-OS fabrics into a single VXLAN-ACI fabric group.

- For more information on Nexus One and VXLAN-ACI fabric groups, see [Creating Fabrics and Fabric Groups](#).
- For more information on importing tenant policies from ACI fabrics, see [Importing Tenant Policies from ACI Fabrics](#).

Key terms

These are the key terms that you should be familiar with when configuring multi-tenancy in Nexus Dashboard.

- **Tenant:** A logical entity representing a group of resources/users. It is associated with specific fabrics (ACI or NX). A tenant is a logical container that represents an independent space for managing and isolating logical network resources and their configurations.
- **Tenant Domain:** A logical construct representing a group of tenants. It simplifies user access management across multiple tenants.
- **User Mapping:** Users are assigned to a Security Domain (defines physical resource access) and a Tenant Domain (defines logical resource access).
- **Security Domain:** A Role Based Access Control collection of fabrics. When a user is assigned access to a security domain they can see and manage the physical fabrics and resources (switches, links, interfaces, etc.) that belong to that domain.
- **Fabric:** A set of devices working together as if they were a single multi-module chassis by way of a single controller and consistent configuration. Generally, fabrics share a coordinated control plane through protocols like BGP or OSPF to share state.
- **Node:** A single physical or virtual instance of Nexus Dashboard
- **Cluster:** A single instance of Nexus Dashboard made up of one or three nodes each running microservices
- **Federation:** Multiple Nexus Dashboard clusters integrated through a Federation Primary cluster

with consistent AAA, Security domain, Tenant, and Tenant domains synchronized between them to allow for scale-out above single cluster capabilities.

- **Tenant Fabric Associations:** Tenant association allows creation of tenant specific objects including VRFs and Security Groups inside that fabric.
- **User Mapping:** Users are assigned to a Security Domain (defines physical resource access) and a Tenant Domain (defines logical resource access). Ultimate access is restricted to the intersection of these two types of mappings.

Understand tenants

A tenant is a logical entity that represents a group of resources that share the same physical infrastructure but are isolated from other tenants.

These are the characteristics of a tenant:

- A tenant acts as a filter for policies and allows you to exercise tenant-based access control.
- A tenant represents a unit of isolation from a policy perspective, but it does not represent a private network.
- Tenants are associated with specific fabrics.
- A tenant can be a member of multiple tenant domains. See [Understand tenant domains](#) for description of tenant domains.

You would use tenants in Nexus Dashboard for these reasons:

- **Policy isolation:** Tenants act as a container for networking policies, ensuring that policies applied to one tenant do not interfere with another. This is particularly important for multi-customer or multi-department environments.
- **Access control:** Tenants enable tenant-based access control, allowing administrators to define who can manage resources.
- **Interconnection of ACI and NX-OS fabrics:** Stretching a policy across fabrics does not normally require tenants. However, because Nexus Dashboard now supports the interconnection of ACI and NX-OS fabrics, and since ACI has always supported tenants, it is now necessary for tenants to exist on NX-OS fabrics as well. Because of this, tenants are required in order to support the interconnection of ACI and NX-OS fabrics.
- **Organizational grouping:** Tenants can represent a customer in an organization or domain in an enterprise, or simply a convenient grouping of policies.

A special system-created tenant called **default-tenant** is also available in Nexus Dashboard. These are the characteristics of the **default-tenant**:

- The default-tenant is only applicable for NX-OS fabrics. The default-tenant has no role to play on ACI fabrics and should not be used in Orchestration configurations.
- The default-tenant is used to manage virtual resources not assigned to any specific tenant.
- All policies that were previously not associated to a tenant belong to the default-tenant on NX-OS fabrics.
- The default-tenant becomes available when upgrading brownfield NDFC deployments to Nexus Dashboard release 4.2.1.

- From the Nexus One perspective, the default-tenant, along with its NX-OS configurations, can be added to a VXLAN-ACI fabric group; however, the configurations from the default-tenant cannot be deployed to ACI fabrics. It's only applicable to NX-OS fabrics, as mentioned previously. See [Creating Fabrics and Fabric Groups](#) for more information on Nexus One and VXLAN-ACI fabric groups.

Understand tenant domains

A tenant domain is a virtual construct that contains one or more tenants, grouping those tenants for the purpose of user association. Tenant domains are primarily used by Nexus Dashboard to manage user access to tenants across different ACI or NX-OS fabrics. Tenant domains offer a convenient way to associate a user with a profile that contains the list of tenants. When you add or delete tenants from a tenant domain, all users are automatically affected. This allows an administrator to make updates in one place instead of making changes to all users individually.

These are the characteristics of a tenant domain:

- A user is associated to a tenant domain rather than tenants. Tenant domains are not directly associated with fabrics, but rather are associated with users and their av-pairs.
- A tenant domain does not have a direct role association but it does restrict resource access.

You would use tenant domains in Nexus Dashboard for these reasons:

- **Simplified user access management:** Tenant domains provide a convenient way to associate a user with a profile or template of tenants they can access. Instead of assigning multiple tenants to individual users, an administrator can assign a user to a tenant domain, which then grants access to all tenants within that domain.
- **Efficient policy changes:** If an administrator needs to add or remove tenants from a group, they can modify the tenant domain, and all users associated with that domain are automatically affected. This centralizes management and reduces administrative overhead.
- **Logical access control:** While security domains define user access to physical resources (such as fabrics or switches), tenant domains define access to logical entities (the tenants themselves). A user is assigned to a tenant domain in addition to security domains.

This special system-created tenant domain is available in Nexus Dashboard:

- **all-tenants-domain:** Grants users access to all current and future tenants. Users without an explicit tenant domain association default to this tenant domain.

In addition:

- Only users associated with the all-tenants-domain can manage physical resources. User-defined tenant domains restrict users to logical objects.
- In a federation, tenant and tenant domains are always created on the federation primary even if executed on a different cluster.
- Tenant fabric associations are always created at the individual cluster.

Understand how users work with tenants and tenant domains

The concept of a user in Nexus Dashboard, as described in [Configuring Users, Roles, and Security](#), remains unchanged with the introduction of the tenants and tenant domains feature in Nexus Dashboard release 4.2.1, where a user is simply an individual or entity that interacts with Nexus Dashboard.

However, how users interact with tenants and tenant domains is new for Nexus Dashboard release 4.2.1. Users are now assigned to tenant domains along with security domains, but with these differences:

- Assigning a user to a security domain defines user access to physical resources, such as fabrics and switches. Mapping a user to a security domain and a role determines that user's access rights.
- Assigning a user to a tenant domain defines user access to logical entities. Users can be associated with at most one tenant domain, which results in a list of tenants that they can access.
- Orchestration requires access to all fabrics in order to orchestrate the stretching of policies across interconnected fabrics. Therefore, Orchestration users must use security domain **all**.

These are additional ways that tenant domains work with users, roles, and security:

- A tenant domain does not carry a role assignment.
- If a user is not associated with a specific tenant domain, that user will be granted access to the **all-tenants** tenant domain on Controller and Telemetry pages, but will not have access to tenants on Orchestration pages.

Tenants in Nexus Dashboard and ACI

As mentioned earlier, tenants have been available in ACI previously but were not used in Nexus Dashboard. Tenants and tenant domains are now available in Nexus Dashboard beginning with Nexus Dashboard release 4.2.1.

Because tenants have been available in ACI previously, existing ACI fabrics that are brought into Nexus Dashboard might already have tenants configured. In addition, it is possible that the same tenant name is used in different ACI fabrics that are not part of the same fabric group.

When these tenants are brought into Nexus Dashboard, in order to resolve name space collisions, Nexus Dashboard supports an ACI fabric tenant name (a local name) that is used in Nexus Dashboard. This allows different fabric groups to support tenants on the ACI fabric with the same name.



An import of policies from an APIC will not be allowed if that import causes the **CLI_PREFIX:PolicyName** to exceed the maximum length of the policy name on an NX-OS switch.

This table provides feature comparison information for tenants that have historically been available in ACI and tenants that are now available in Nexus Dashboard.

Field	Nexus Dashboard	ACI
Tenant	Defined outside of the fabric	Defined at the fabric level
Tenant Fabric Association	Required to associate any logical objects with that tenant	Not Applicable
VRF	Defined inside a fabric group or fabric with an optional tenant association	Defined inside a tenant
Network/bridge domain (BD)	Defined inside a fabric group or fabric with an optional tenant association	Defined inside a tenant
Security Group	Defined inside a fabric group or fabric with an optional tenant association	Defined inside a tenant

AAA defined users

With the multi-tenancy feature in Nexus Dashboard 4.2.1, AAA users are still required to have RBAC set but the rules are a bit more complex.

- [NX-OS AAA users](#)
- [Orchestration AAA users](#)

NX-OS AAA users

- If AAA users do not have a tenant-domain specified in their Av-pair, they are automatically placed in the all-tenants-domain
- Similar to previous releases, a user might have multiple roles within a given security domain
- Av-pairs are defined for users with different profiles:
 - `shell:domains = <security-domain1>/<role>,<security-domain2>/<role>|<role> tenant-domain=<tenant-domain>`

For example:

```
shell:domains = SD1/fabric-admin,SD2/support-engineer|designer tenant-domain=all-tenants-domain
```

These are example entries for different users:

- Sales admin: `shell:domains = SD1/fabric-admin tenant-domain=sales-dmz-domain`
- Marketing admin: `shell:domains = SD1/fabric-admin tenant-domain=marketing-dmz-domain`
- Sales admin in a fabric group: `shell:domains = all/fabric-admin tenant-domain=sales-dmz-domain`
- Physical Infra admin: `shell:domains = all/fabric-admin`
- DC admin: `shell:domains = all/super-admin`
- Fabric group admin: User must use security domain `all`, such as this:

```
shell:domains = all/fabric-admin,SD2/support-engineer|designer tenant-domain=all-tenants-
```

domain *

- o Orchestration users: User must use security domain **all**, but with no tenant domain, such as this:

```
shell:domains = all/designer tenant-domain=sales-dmz-domain
```

Orchestration AAA users

Orchestration has always supported tenancy. Prior to release 4.2, users could view all tenants but could only edit policies on the tenants that they were associated with. Starting with release 4.2, views are also tenant constrained.

To continue supporting tenancy constraints and prevent privilege escalation, it is important that remote users are associated with an appropriate tenant domain. An admin (remote or local) will continue to have access to the system and the admin should define appropriate tenant domains and update all remote users to their appropriate tenant domain.

Here is an example of a AV pair for a user who is a **designer** as well as an **approver**:

```
"avpair" : " shell:domains=all/approver|designer tenant-domain=TD1" ,
```

- If the tenant domain is not specified in the AV pair, the following behavior will be seen:
 - o If it is an upgraded setup and the user previously had access to a tenant through a NDO tenant association, they will continue to have access. This cannot be modified within Nexus Dashboard; it can only be changed by association of a tenant domain in the AV pair.
 - o If a remote user does not have a pre-upgrade association, they will not have functional access to Orchestration until an AV pair specifying the tenant domain is set.
 - o If they are a super-admin, then they have access to the all-tenant-domain.
 - o If they are a fabric-admin in the security domain all, then they have access to the all-tenant-domain.

Guidelines and limitations

These are the guidelines and limitations when configuring tenants and tenant domains.

- [General guidelines and limitations: Tenants and tenant domains](#)
- [Guidelines and limitations: Tenants and object names](#)
- [Guidelines and limitations: Tenants and VXLAN-ACI fabric groups](#)

General guidelines and limitations: Tenants and tenant domains

- You are limited to 10 tenants or fewer in a tenant domain.
- To create or manage resources within a tenant's context, you must first associate a fabric with a tenant or a tenant with a fabric group. This association defines the operational scope of tenants, allowing each tenant to have distinct aliases when linked to various fabrics.
- When you associate an ACI fabric with a tenant, if you delete that ACI fabric association with the tenant using the procedures in [Edit fabrics that are associated with a tenant](#), this action only deletes the fabric association with the tenant and does not delete the tenant from the ACI fabric. You must manually delete the tenant from the ACI fabric afterward to avoid potential stale configurations on the APIC.
- A tenant can exist across multiple fabrics.
- Tenant-fabric associations are cluster specific, whereas tenants and tenant domains are considered federation-wide resources and must have unique names throughout the federation.
- These workflows require access to the **all-tenants-domain** tenant domain:
 - All configurations involving physical switches, links, and interfaces, as well as switch- and fabric-wide policies
 - Pre-change Analysis (fabric admin only, only users with all-tenant access)
 - Delta Analysis (aggregated, only users with all-tenant access)
 - Compliance rule
 - Anomaly rule
 - Explorer
 - Inter-fabric connections
 - Layer 3 neighbors that are underlay connections (L3Out-based neighbors can be filtered by their tenant association)
 - Physical (routed) interfaces (both eth and po)
 - Sustainability (including PDU)
 - Advisories
 - Integrations: VCenter (VM display), AppD
 - Hardware resources
 - Capacity dashboard

- Tenants and tenant domains are supported in these areas:
 - Anomalies
 - Endpoints
 - Routes
 - Segmentation & security
 - Flow Telemetry / Traffic Analytics
 - Multicast group & Routes
 - Interfaces (Logical)
 - L3 Neighbors (overlay)
 - Topology
 - Search
 - Connectivity analysis (Source / dest selection based on endpoints)
 - Custom dashboards (Behavior defined by underlying APIs)
- Tenants and tenant domains are *not* supported in these areas:
 - Co-located fabrics are not supported
 - You will not be able to enable change control if a fabric has tenants assigned to it
- Anomaly levels located under **Manage > Fabrics** and **Manage > Inventory** are shown only to users who belong to the **all-tenants-domain** tenant domain or are not associated with any tenant domain.
- If a VM is not associated to any tenant, then any user can view the VM details. Tenant association to a fabric does not restrict the entire fabric and integrations associated to the fabric to the users who have access to that tenant. Tenant association by way of having a specific VM connected to a network or VRF applies because the network or VRF are themselves associated to a tenant. When a VM is associated to a network or VRF in a tenant, only users who have that tenant in their tenant domain associations can view the VM.

Guidelines and limitations: Tenants and object names

- Object names use the format **tenant_prefix:resource_name**, where
 - The **tenant_prefix** portion has these characteristics:
 - Is used to provide name-spacing to all named logical objects, which are written to the Nexus switches.
 - Is limited to no more than 10 characters, even though the tenant name could be up to 63 characters in length.
 - The combination of **tenant_prefix:resource_name** should not exceed the character limit on the switches. These are the character limits on the switches:
 - VRF: 32
 - SG (Policy Map, Class Map, SG Name): 40
 - L3Out (Route Map, ACL, Community List, Prefix List) : 63
- Object creation will limit each object name length to the total allowed, as listed above:

- o Minus the **tenant_prefix** length, and
- o Minus 1 for the colon (:) in the object name

For example, a fabric with a five character **tenant_prefix** would be limited to a VRF name (**tenant_prefix:resource_name**) for that tenant of no more than 26 characters:

- o 32 character limit for VRFs, as listed above,
 - o Minus 5 (**tenant_prefix**), and
 - o Minus 1 for the colon (:) in the object name
- Object names will be limited by Nexus capabilities when used with a VXLAN-ACI fabric group.
 - Object names will be the natural name when deployed to ACI but the tenant in ACI might not be the same as the tenant name on Nexus Dashboard. This setting provides the mapping for VXLAN-ACI.

Guidelines and limitations: Tenants and VXLAN-ACI fabric groups

- Deleting a tenant in Nexus Dashboard does not delete the corresponding tenant in the Cisco Application Policy Infrastructure Controller (APIC).
- Policies defined within a user tenant are not accessible or referenceable by any other user tenant, including the **common** tenant.

Upgrade considerations

This section provides information that you should understand when upgrading from a previous release to Nexus Dashboard 4.2.1.

- [Guidelines and limitations: Upgrades](#)
- [Understand tenant domains and upgrades](#)

Guidelines and limitations: Upgrades

These guidelines and limitations apply when you upgrade from a release prior to Nexus Dashboard release 4.1.1 (prior to the unification of services under Nexus Dashboard, such as NDO release 4.3.1):

- In older NDO releases, all templates are associated with a tenant, except for fabric templates. After upgrading to Nexus Dashboard release 4.2.1, you will have visibility of those templates based on your tenant associations.
- In Orchestration, the **common** tenant is a special system-created tenant that is automatically created and associated with all of the ACI fabrics that it controls. All users have read access to the **common** tenant, but only users that have **common** tenant in their tenant domain can update the **common** tenant.
- Orchestration will uplift tenants and their fabric and user association to Nexus Dashboard. Local users will be assigned their own tenant domains so that their functionality will not be affected on upgrade. The administrator can reconfigure the user to a tenant-domain configuration at a later date based on their design.
- The administrator should also create appropriate tenant domains to accommodate remote users and update the AV pair of remote users to their appropriate tenant domains.
- To ensure that remote users are not totally shut out on upgrade due to a lack of tenant domain association, remote users with admin and fabric-admin privileges will automatically be associated with the all-tenants-domain. Remote users with the designer role will have access to their old (pre-4.2) tenant associations and will continue to function. If a remote user does not have a pre-upgrade association, they will not have functional access to Orchestration until an AV pair specifying the tenant domain is set.

Understand tenant domains and upgrades



The information in this section is only applicable for situations where you are upgrading from an NDO release prior to the unification of services under Nexus Dashboard, such as NDO release 4.3.1, to Nexus Dashboard release 4.2.1.

Because tenant domains are a new concept in the Nexus Dashboard 4.2 release, Nexus Dashboard automatically creates tenant domains for users based on their existing configurations and roles, where Nexus Dashboard creates a tenant domain for each user as part of the upgrade process and associates the tenant domain with the existing tenants that the user is associated with.

System-created tenant domains serve as a temporary solution, allowing you to maintain and modify your existing configurations after an upgrade. After the upgrade is complete, you can create new tenant domains, then transfer users to these new tenant domains and delete the system-created tenant domains afterward.

These tables provide examples of how tenant domains are created for each user based on roles.

- [Prior to the upgrade](#)
- [After the upgrade](#)

Prior to the upgrade

Assume that these user-role-tenant associations were set before the upgrade:

User	Role	Tenant association	Notes
User 1	Designer	T1	
User 2	Designer	T2	
User 3	Designer	T1, T2	
User 4	Fabric Admin	All (view and update)	Fabric admin is associated with all tenants on NDO.
User 5	Observer	All (view only)	Observer can view all tenants.
User 6	Approver	See Notes	Approver can view all tenants and approve any template belonging to any tenant; cannot update any template belonging to any tenant.
User 7	Support Engineer (SE)	See Notes	Support Engineer can view all tenants and deploy any template belonging to any tenant; cannot update any template belonging to any tenant.
User 8	Multiple roles; for example, Designer + SE (deployer)	T1	For a user with Designer + SE (deployer) roles: <ul style="list-style-type: none"> • As Designer, user is associated only with T1 • As SE (deployer), user can deploy any tenant

After the upgrade

After the upgrade, these tenant domains will be created:

Tenant domain	Associated tenants	Associated user	Roles	Notes
TD_User1	T1	User 1	Designer	
TD_User2	T2	User 2	Designer	
TD_User3	T1, T2	User 3	Designer	
All_tenant_domain	All tenants	User 4	Fabric Admin	Fabric admin could update any fabric and any tenant, so is associated with All_tenant_domain.
All_tenant_domain	All tenants	User 5	Observer	Observer could observe any tenant and any template, so is associated with All_tenant_domain.

Tenant domain	Associated tenants	Associated user	Roles	Notes
All_tenant_domain	All tenants	User 6	Approver	Approver could approve any tenant, so is associated with All_tenant_domain.
All_tenant_domain	All tenants	User 7	Support Engineer (SE)	Support Engineer could deploy any tenant, so is associated with All_tenant_domain.
TD_User8	T1	User 8	Multiple roles; for example, Designer + SE (deployer)	<ul style="list-style-type: none"> ▪ As a Designer, the user was associated with tenant T1, so is tenant-restricted for template updates ▪ As an SE (deployer), the user could deploy any template <p>Nexus Dashboard cannot associate the user with the All_tenant_domain because, as a Designer, they previously were associated only with tenant T1, so associating the user with the All_tenant_domain would result in a privilege escalation.</p> <p>In these situations, Nexus Dashboard restricts the user to the lowest common denominator. So, in this example, the user loses certain privileges as an SE and is restricted to deploying only on the tenants that they have access to (in this case, T1).</p>

Navigate to the Multi-tenancy page

Click **Admin > Multi-tenancy**.

- Click **Tenants** to configure tenants.
- Click **Tenant domains** to configure tenant domains.

Configuring tenants

These sections describe how to configure tenants.

- [Understand the Tenants page](#)
- [Create a tenant](#)
- [View tenant information](#)
- [Edit a tenant](#)
- [Associate fabrics with a tenant](#)
- [Associate tenants with a fabric group](#)
- [Edit fabrics that are associated with a tenant](#)
- [Import a tenant from ACI](#)
- [Delete a tenant](#)

Understand the Tenants page

1. [Navigate to the Multi-tenancy page.](#)
2. Click **Tenants**.

Previously-configured tenants are displayed in the **Tenants** page, with this information.

Field	Description
Name	Provides the name of the tenant. For user-created tenants, click the entry in the Name column to bring up detailed information on a tenant.
Description	Provides the description information for the tenant.
Associated fabrics	Shows the fabrics that are associated with each tenant. For user-created tenants, click the entry in the Associated fabrics column to bring up detailed information on an associated fabric.

Create a tenant

You can create a tenant on an individual cluster or in a federated cluster. Note that creating a tenant in a federated cluster might take longer as it will always be proxied to the federation primary and synced to all the other clusters.

1. [Navigate to the Multi-tenancy page.](#)
2. Click **Tenants**.

Previously-configured tenants are displayed in the **Tenants** page.

3. Click **Actions > Create tenant**.

The **Create tenant** page is displayed.

4. Enter a name and description for the tenant, then click **Save**.

A confirmation window appears, with these options:

- o Click **View tenant** to see information on this configured tenant. See [View tenant information](#).
- o Click **Associate fabrics** to associate one or more fabrics with this tenant. See [Associate fabrics with a tenant](#).
- o Click **Create tenant** to create another tenant.

View tenant information

1. [Navigate to the Multi-tenancy page](#).
2. Click **Tenants**.

Previously-configured tenants are displayed in the **Tenants** page.

3. Locate the tenant that you want to view information for and click the link in the **Name** column for that tenant.

The **View <tenant-name> details** page is displayed, with this information.

Field	Description
General	
Name	Provides the name of the tenant.
Description	Provides the descriptive text associated with this tenant.
Associated fabrics	
Fabric name	<p>Displays the fabric that is associated with this tenant.</p> <p>Note that if the tenant is associated with a fabric group, as described in Associate fabrics with a tenant, then two entries are displayed in the Fabric name column:</p> <ul style="list-style-type: none">• The fabric group that is associated with the tenant• The member fabric within that fabric group that is associated with the tenant <p>An arrow next to the fabric group name expands or collapses the displayed fabric and fabric group information.</p>
ACI tenant name	For ACI fabrics. Displays the entry that is used to create the tenant on the ACI fabric.
Tenant prefix	For NX-OS fabrics. Displays the entry that is a unique identifier that is applied to the switch CLI configuration to distinguish tenant-specific policies for NX-OS VXLAN fabrics.
Allowed VLANs	For NX-OS fabrics. Displays the VLANs that can be used by the tenant.

4. Perform additional actions on this tenant using the **Actions** drop-down list, if necessary.
 - o To edit the configuration for this tenant, click **Actions > Edit tenant**. See [Edit a tenant](#).
 - o To edit the fabric association for this tenant, click **Actions > Edit fabric association**. See [Edit](#)

[fabrics that are associated with a tenant.](#)

- o To associate a fabric with this tenant, click **Actions > Associate fabrics**. See [Associate fabrics with a tenant](#).

Edit a tenant

1. [Navigate to the Multi-tenancy page.](#)
2. Click **Tenants**.

Previously-configured tenants are displayed in the **Tenants** page.

3. Click the circle next to the tenant that you want to edit, then click **Actions > Edit tenant**.



You can also click on the configured tenant in the **Tenants** page, then click **Actions > Edit tenant** in the tenant-specific page.

The **Edit tenant** page is displayed.

4. Make the necessary changes to the tenant, then click **Save**.

Associate fabrics with a tenant

These procedures describe how to associate fabrics with a specific tenant. To associate tenants with a specific fabric group, see [Associate tenants with a fabric group](#).

- You must associate a fabric with a tenant on either the controller or the Orchestration cluster.
- VLAN restrictions and tenant prefix are applicable to VXLAN fabrics, as well as VXLAN and VXLAN-ACI fabric groups.
- Allowed VLANs may overlap with other tenants. VLAN restriction applies to network and VRF VLANs and network attachment.
- You can modify allowed VLANs for the tenant within a fabric.
- You cannot modify the ACI tenant name or tenant prefix after associating a fabric with a tenant. Because you can add VXLAN fabrics to VXLAN-ACI fabric groups, this restriction for ACI mapping in case future normalization is needed.
- When a fabric is added to a fabric group, if there is no conflict, all objects, including tenant associations, VRFs, networks, security groups, and contracts, are copied up to the fabric to the fabric group, then down to all fabrics in the fabric group.
- For Orchestration, only consistent tenant names between Nexus Dashboard and APIC are supported.
- See [\[Guidelines and limitations: Object names\]](#) for a list of guidelines and limitations around object names.

To associate fabrics with a tenant:

1. [Navigate to the Multi-tenancy page.](#)
2. Click **Tenants**.

Previously-configured tenants are displayed in the **Tenants** page.

3. Click the circle next to the tenant that you want to associate with a fabric, then click **Actions > Associate fabrics**.



You can also click on the configured tenant in the **Tenants** page, then click **Actions > Associate fabrics** in the tenant-specific page.

The **Associate fabrics for <tenant-name>** page is displayed.

4. Choose the type of fabric that you want to associate with this tenant.



For ACI fabrics, only ACI fabrics that are orchestration-enabled or are part of a VXLAN-ACI fabric group can be associated with a tenant.

- o ACI fabrics
 - o VXLAN fabrics
 - o Fabric groups
5. In the **Fabrics/Fabric groups** field, choose the fabrics or fabric groups that you want to associate with this tenant.
 - o You can associate more than one fabric or fabric group with each tenant.
 - o Supported fabric groups include multi-cluster fabric groups, including multi-cluster fabric groups with ACI and VXLAN fabrics in a single fabric group.
 6. Enter the necessary information to associate the fabrics with the tenant.

Field	Description
ACI tenant name	Accept the default entry or change the entry. This is the local ACI tenant name to create the tenant on the ACI fabric if this association is ever inherited into a VXLAN-ACI fabric group.
Tenant prefix	Accept the default entry or change the entry. This entry is a unique identifier that is applied to the switch CLI configuration to distinguish tenant-specific policies for NX-OS VXLAN fabrics. The switch software limits some policy names to just 32 characters. Because of this character limitation, you can configure a tenant prefix (essentially, a shortened tenant name) that is used on NX-OS switches. There is a character limit of 10 for this field. See [Guidelines and limitations: Object names] for more information.
Allowed VLANs	Specify the VLANs that can be used by the tenant.

7. Click **Save**.

The **View <tenant-name> details** page appears with the fabric now associated with this tenant.

Associate tenants with a fabric group

These procedures describe how to associate tenants with a specific fabric group. To associate

fabrics with a specific tenant, see [Associate fabrics with a tenant](#).

1. Navigate to **Manage > Fabric Groups**, then click the fabric group where you want to associate a tenant.

That fabric group's **Overview** page is displayed.

2. Click **Actions > Associate tenants**.

Previously-configured tenants are displayed in the **Associate tenants** page, with this information.

Field	Description
Tenant	Provides the name of the tenant. For user-created tenants, click the entry in the Tenant column to bring up detailed information on a tenant. See View tenant information .
ACI tenant name	For ACI fabrics. Displays the entry that is used to create the tenant on the ACI fabric.
Tenant prefix	For NX-OS fabrics. Displays the entry that is a unique identifier that is applied to the switch CLI configuration to distinguish tenant-specific policies for NX-OS VXLAN fabrics.
Allowed VLANs	For NX-OS fabrics. Displays the VLANs that can be used by the tenant.

3. Click **+ Associate tenants**.

4. Make the necessary choices in the **Associate tenants** page.

- a. In the **Tenant** field, choose the tenant that you want to associate with this fabric group.
- b. For ACI fabrics, in the **ACI tenant name** field, accept the default entry or change the entry.

This entry is used to create the tenant on the ACI fabric.

- c. For NX-OS fabrics:

- i. In the **Tenant prefix** field, accept the default entry or change the entry.

This entry is a unique identifier that is applied to the switch CLI configuration to distinguish tenant-specific policies for NX-OS VXLAN fabrics.

The switch software limits some policy names to just 32 characters. Because of this character limitation, you can configure a tenant prefix (essentially, a shortened tenant name) that is used on NX-OS switches. There is a character limit of 12 for this field.

- ii. In the **Allowed VLANs** field, specify the VLANs that can be used by the tenant.

- d. Click **Save** in the row where you are entering the information for the tenant that you are associating with this fabric group.

The information for the newly-associated tenant is displayed.

5. Click **+ Associate tenants** to associate additional tenants with this fabric group, or click **Save** if you have associated all of the necessary tenants with this fabric group.

Edit fabrics that are associated with a tenant

1. [Navigate to the Multi-tenancy page.](#)
2. Click **Tenants**.

Previously-configured tenants are displayed in the **Tenants** page.

3. Click the circle next to the tenant that has a fabric association that you want to edit.
4. Click **Actions > Edit associated fabrics**.

The **Edit associated fabrics** page is displayed.

5. Choose the appropriate fabric and make the necessary fabric association edits.
 - o If you want to associate a different fabric with this tenant, click **Actions > Associate fabrics**. See [Associate fabrics with a tenant](#).
 - o If you want to edit an existing fabric association, click **Actions > Edit**, then make the necessary modifications to this fabric association.
 - o If you want to delete an existing fabric association, click **Actions > Delete**.

Import a tenant from ACI

This process creates an ACI tenant in the Nexus Dashboard as a Nexus Dashboard tenant, and associates that tenant with the fabric that you provide in the **Fabric** field below and any fabric group that this fabric is a member of.

1. [Navigate to the Multi-tenancy page.](#)
2. Click **Tenants**.

Previously-configured tenants are displayed in the **Tenants** page.

3. Click **Actions > Import tenant (ACI)**.

The **Import tenant (ACI)** page is displayed.

4. In the **Fabric** field, choose the ACI fabric where you want to import the tenant.
5. In the **ACI tenants on fabric** field, choose the tenant that you want to import from the ACI fabric.

Only tenants that are importable are listed. This associates the tenant that you enter in the **ACI tenants on fabric** field with the fabric that you chose in the **Fabric** field.

6. In the **Tenant actions** field, determine if you want to:
 - o Associate the ACI tenant with an existing Nexus Dashboard tenant, or
 - o Create a new Nexus Dashboard tenant for this ACI tenant.
7. If you chose to associate the ACI tenant with an existing Nexus Dashboard tenant, in the **Nexus dashboard tenant name** field, choose the existing Nexus Dashboard tenant that you will associate the ACI tenant with.
8. In the **Tenant prefix** field, enter a prefix for the Nexus Dashboard tenant.

This entry is a unique identifier that is applied to the switch CLI configuration to distinguish tenant-specific policies for NX-OS VXLAN fabrics.

The switch software limits some policy names to just 32 characters. Because of this character limitation, you can configure a tenant prefix (essentially, a shortened tenant name) that is used on NX-OS switches. There is a character limit of 12 for this field.

9. In the **Allowed VLANs** field, specify the VLANs that can be used by the tenant.

Enter the VLANs as ranges (for example, **100-200**) or as individual IDs, separated with commas (for example, **100, 150, 200**).

10. Click **Save**.

A confirmation window appears.

Delete a tenant

1. [Navigate to the Multi-tenancy page](#).
2. Click **Tenants**.

Previously-configured tenants are displayed in the **Tenants** page.

3. Click the box next to the tenant that you want to delete.
4. Click **Actions > Delete**.

The tenant is now deleted from the list of tenants displayed in the in the **Tenants** page.

Working with tenants under Segmentation and security

1. Navigate to the main **Fabrics** page.

Manage > Fabrics

2. Locate the VXLAN fabric where you want to work with tenants under **Segmentation and security**.
3. Single-click the appropriate VXLAN fabric.

The **Summary** page for that VXLAN fabric appears with the **Overview** tab selected by default.

4. Click **Segmentation and security**.
5. On the **Segmentation and security** tab, click the **Tenants** subtab.
6. Review the information on the **Tenants** page.

View tenants information

This page shows tenant fabric associations.

- You can deploy individual tenants (including all pending logical objects for that tenant across all switches or fabrics) from this page.
- The default tenant is always present but not shown if there are no other tenant fabric associations.

- The default tenant might not be deployed from here as other VRF and network options exist for these deployments.

To view tenants information:

1. Review the information provided on the **Tenants** page.

The table provides information on tenants in the fabric.

Field	Description
Name	Provides the name of the tenant.
ACI tenant name	For ACI fabrics. Displays the entry that is used to create the tenant on the ACI fabric.
Tenant prefix	For NX-OS fabrics. Displays the entry that is a unique identifier that is applied to the switch CLI configuration to distinguish tenant-specific policies for NX-OS VXLAN fabrics.
VLANs	For NX-OS fabrics. Displays the VLANs that can be used by the tenant.
Config status	Displays the configuration status for the tenant.

2. Click the table header to sort the entries in alphabetical order for the selected parameter.
3. Perform any of the listed actions on the **Tenants** page.

The table describes the action items that are available in the **Actions** drop-down list.

Tenants actions and descriptions

Action Item	Description
Deploy	Allows you to deploy the tenant in the fabric. See Deploy tenants .
Import tenant policies	Allows you to import the tenant policies. See Import tenant policies .

Deploying tenants

- [Types of deployment](#)
- [Deployment dependencies](#)
- [Deploy tenants](#)

Types of deployment

Before deploying a tenant, it's useful to understand the different types of deployment that are available in Nexus Dashboard as related to multi-tenancy.

- **Recalculate and deploy:**
 - Primarily used for inter-fabric connectivity and for updating the infra-tenant.
 - Also deploys all user tenant's objects in pending state.
- **VRF deploy**
- **Network deploy:** Not supported on VXLAN-ACI fabric groups.

- **Tenant deploy:** Described in these procedures.
 - All objects within a specific tenant are deployed.
 - You can select and deploy multiple tenants at once.

Deployment dependencies

- VRF deployment requires the VRF to be in the attached state.
- Network deployment requires both the network and the VRF to be in the attached state.
- Security group deployment requires both the security group and the VRF to be in the attached state.
- Security association deployment requires the security group to be in the attached state.
- The network selector of a security group requires the network attachment to be in the attached state.
- If a network in a user tenant references a VRF in the **common** tenant, the VRF in the **common** tenant must be deployed first.
- If a child network in a user tenant references a normal network in the **common** tenant, the **common** tenant must be deployed first.

Deploy tenants

1. In the **Tenants** page, choose one or more tenants to deploy into the fabric.
2. Click **Actions > Deploy**.

The **Deploy configuration** page appears.

3. Verify the information in the **Deploy configuration** page.
 - For NX-OS fabrics, the entry in the **Pending config** column displays the number of lines that have changed as part of the preview.
 - For ACI deployments, the entry in the **Pending config** column displays the total number of objects in APIC that have been added, modified, or deleted.

Click the entry in the **Pending config** column to bring up the **Pending config** page.

- The **Pending config** tab on this page displays the pending configurations on the fabric.
 - The **Side-by-side comparison** tab on this page displays the running configuration and expected configuration in a side-by-side representation.
4. Click **Deploy all**.

Import tenant policies

1. In the **Tenants** page, choose a tenant that contains the policies that you want to import.
2. Click **Actions > Import tenant policies**.

The **Import tenant policies from import** page appears.

3. Navigate through the import steps.

- 1. [Select fabric](#)
- 2. [Select VRFs](#)
- 3. [Select networks/BDs](#)
- 4. [Map EPGs to security groups](#)
- 5. [Map ACI contracts to ND](#)
- 6. [Review security associations](#)
- 7. [Summary](#)

1. Select fabric

1. In the **Select fabric** step, choose the appropriate ACI fabrics to be imported into your Nexus Dashboard fabric group.
2. Click **Next**.

You advance to the [2. Select VRFs](#) step.

2. Select VRFs

The VRFs that you choose in this step influences the verification information that is presented later in this workflow. If you find issues with that verification information, you might have to return to this page and modify the choices that you made here.

1. In the **Select VRFs** step, choose the appropriate VRFs from the ACI fabric that will be imported into your Nexus Dashboard fabric group.

If a VRF belongs to the common tenant, you will see the text **common** underneath the VRF name.

2. Click **Next**.

You advance to the [3. Select networks/BDs](#) step.

3. Select networks/BDs

The networks and bridge domains that you choose in this step influences the verification information that is presented later in this workflow. If you find issues with that verification information, you might have to return to this page and modify the choices that you made here.

1. In the **Select networks/BDs** step, choose the appropriate networks and bridge domains from the ACI fabric that will be imported into your Nexus Dashboard fabric group.

- If you want to import the network or bridge domains and all EPGs that are associated with that network or bridge domain, click the box next that network or bridge domain.

All EPGs that are associated with the network or bridge domain are automatically imported when you choose the upper-level network or bridge domain.

- If you want to import only specific EPGs, click the arrow next to a bridge domain or network to show any associated children EPGs, then choose the specific EPGs that you want to import.

2. Click **Next**.

You advance to the [4. Map EPGs to security groups](#) step.

4. Map EPGs to security groups

The **4. Map EPGs to security groups** step in the process is a verification step; you cannot make changes on this page. However, if you see any issues with the information that is presented on this page, click **Back** to return to the [3. Select networks/BDs](#) or [2. Select VRFs](#) step, then correct the issue and move forward to this verification stage in the process again.

1. In the **Map EPGs to security groups** step, verify that the EPGs have been mapped to the intended ESGs (security groups) correctly.
2. Click **Next**.

You advance to the [5. Map ACI contracts to ND](#) step.

5. Map ACI contracts to ND

The **Map ACI contracts to ND** step in the process is a verification step; you cannot make changes on this page. However, if you see any issues with the information that is presented on this page, click **Back** to return to the [3. Select networks/BDs](#) or [2. Select VRFs](#) step, then correct the issue and move forward to this verification stage in the process again.

1. In the **Map ACI contracts to ND** step, verify that the ACI contracts will be mapped to the Nexus Dashboard correctly.

The contracts listed in this step will be uplifted from ACI and will be used to create contracts in Nexus Dashboard.

- o Each ACI contract consists of the following pieces of information:
 - Source and destination (also known as consumer and provider) information.
 - Filter information, including vzSubject information, which is a component of a contract that groups together one or more filters defining the types of traffic allowed or denied between EPGs.

For more information on ACI contracts, see [Cisco ACI Contract Guide White Paper](#).

- o When the ACI contracts are mapped to the Nexus Dashboard, the process essentially performs a one-to-one mapping of the vzSubject information from the ACI contract to a Nexus Dashboard contract. In addition, filters within the ACI contract become protocol definitions within the Nexus Dashboard contract.

For more information on Nexus Dashboard contracts, see [\[Working with security contracts\]](#).

2. Click **Next**.

You advance to the [6. Review security associations](#) step.

6. Review security associations

In the **Review security associations** step, verify that the source and destination pairs are associated with the uplifted contracts correctly.

- If any information on this page is incorrect, click **Back** to repeat any procedures that were not configured correctly.

- Click **Next** if all of the information on this page is correct.

You advance to the [7. Summary](#) step.

7. Summary

1. In the **Summary** step, verify that all of the information from the **Import tenant policies from import** steps are configured correctly.
 - If any information on this page is incorrect, click **Back** to repeat any procedures that were not configured correctly.
 - Click **Submit** if all of the information on this page is correct and you are ready to import the ACI tenant policies to your Nexus Dashboard.

A progress screen appears.

2. When you see the message **Successfully imported selected policies**, click **Done**.

You are returned to the **Tenants** tab under **Segmentation and security**.

- Click **VRFs** to see the new Nexus Dashboard VRFs that were created from the [2. Select VRFs](#) step.
- Click **Networks** to see the new Nexus Dashboard networks that were created from the [3. Select networks/BDs](#) step.
- Click **Security groups** to see the new Nexus Dashboard security groups that were created from the [4. Map EPGs to security groups](#) step.
- Click **Security contracts**, **Security associations**, and **Protocol definitions** to see the new Nexus Dashboard security components that were created from the [5. Map ACI contracts to ND](#) step.

Understanding the process for importing tenant policies from ACI fabrics into VXLAN-ACI fabric groups

This feature provides the ability to migrate endpoint groups (EPGs) on APIC to endpoint security groups (ESGs) using the ESG Migration Assistant script, and import those ESGs, as well as VRFs, BDs, and so on, into Nexus Dashboard.

Mapping between ACI and Nexus Dashboard policies

Use the information in this table to understand how ACI components map to Nexus Dashboard components when importing tenant policies from existing ACI fabrics into Nexus Dashboard VXLAN-ACI fabric groups. For a more extensive set of mapping information, see [Mapping between ND, NX-OS and ACI policies for VXLAN-ACI fabric groups](#).

ACI (before import)	Nexus Dashboard (after import)
Bridge domain (BD) + 1 endpoint group (EPG) (Network-centric design)	Network (normal)

ACI (before import)	Nexus Dashboard (after import)
BD + 2 or more EPGs (Application-centric design)	Normal network + child networks: <ul style="list-style-type: none"> • BD + 1st EPG = normal network • Remaining EPGs associated with this BD = child networks to the normal network
EPG contract (original EPG contract, pre-migration)	Security contract
ESG contract	Security association

Import tenant policies from ACI fabrics

This section walks you through the process of importing tenant policies from ACI fabrics into Nexus Dashboard.

- [Guidelines and limitations: ESG Migration Assistant script](#)
- [Using the ESG Migration Assistant script](#)
- [Create fabric group and associate fabrics](#)
- [Import the tenant and tenant policies from ACI into Nexus Dashboard](#)

Guidelines and limitations: ESG Migration Assistant script

- Migrating the EPGs on APIC to ESGs using the ESG Migration Assistant script will cause traffic disruptions because the pcTag associated to the classified resources will change as part of the process. If you do not want to have that traffic impact, you can temporarily "open up" the security policies in a VRF by creating a vzAny-to-vzAny permit-all contract.

Using the ESG Migration Assistant script

- [Download and extract the ESG Migration Assistant script](#)
- [Run the ESG Migration Assistant script in the ACI fabric APIC](#)

Download and extract the ESG Migration Assistant script

1. Log into your APIC as the **admin** user.

```
$ ssh -l admin apic-ip-address
admin@apic-ip-address's password:
```

2. In your APIC, change directories to **/data/techsupport**:

```
# cd /data/techsupport
```

3. Locate or download the ESG Migration Assistant script.

For APIC release 6.1(4) or later, manually download the image from this location:

```
https://github.com/datacenter/ACI-ESG-Migration-Assistant
```

4. Copy or download the `ESGMigrationAssistant-<version>.zip` file to the `/data/techsupport` directory in your APIC.
5. Extract the `ESGMigrationAssistant-<version>.zip` file.

```
# unzip ESGMigrationAssistant-<version>.zip
```

6. Change directories to the `ESGMigrationAssistant-<version>` directory.

```
# cd ESGMigrationAssistant-<version>
```

These contents become available after you extract the zip file.

- o deps (directory)
- o ESGMigrationAssistant
- o ESGMigrationAssistant.py
- o README.md
- o util.py
- o version.txt

Run the ESG Migration Assistant script in the ACI fabric APIC

There are three phases in the ESG Migration Assistant script, as described below:

1. Dry run analysis phase
2. Conversion phase
3. Cleanup phase

You have to run the ESG Migration Assistant script in the order provided above, but you do not have to run each phase within a certain timeframe from the previous phase.

In the ACI fabric:

1. Perform a dry run analysis.

You can perform a dry run analysis of the user configurations and groups similar EPGs into ESGs. The analysis is outputted to a YAML file.

Use the information provided with the `--help` option to determine how you want to perform the dry run.

```
./ESGMigrationAssistant dryrun --help
usage: ESGMigrationAssistant dryrun [-h] [--json JSON | --xml XML | --targz TARGZ |
--dbxml DBXML] [--disableNdMode] [--apic APIC]
                                     [--username USERNAME] [--password PASSWORD] [--mode
{optimized,one-to-one}]
```

```
[--tenantdns TENANTDNS] [--vrfdns VRFDNS] [--outYaml  
OUTYAML] [--prefix PREFIX] [--suffix SUFFIX]
```

options:

```
-h, --help          show this help message and exit  
--json JSON         Configuration snapshot JSON file  
--xml XML           Configuration snapshot XML file  
--targz TARGZ       Configuration snapshot TAR.GZ file  
--dbxml DBXML       ifc_policydist.db.xml file coming from DB conversion phase
```

[INTERNAL ONLY USE]

```
--disableNdMode    Disable Nexus Dashboard compatibility mode  
--apic APIC         APIC IP address or hostname to connect to  
--username USERNAME Username for APIC  
--password PASSWORD Password for APIC  
--mode {optimized,one-to-one}  
                    Select the mode of analysis: optimized (default) or one-to-one
```

```
--tenantdns TENANTDNS
```

Filter analysis to all the VRFs configured inside to the specified Tenants.

Use comma separated Tenant DN's without spaces. Example: uni/tn-T1,uni/tn-T2

```
--vrfdns VRFDNS    Filter analysis to a subset of VRFs. Use comma separated VRF  
DN's without spaces. Example: uni/tn-T1/ctx-ctx1,uni/tn-T2/ctx-ctx2
```

```
--outYaml OUTYAML  YAML file in which we report the execution plan
```

```
--prefix PREFIX    Prefix to add to cloned names (default: empty). Example: contract  
name is "web" and prefix is "e", cloned contract will be named "e_web"
```

```
--suffix SUFFIX    Suffix to add to cloned names (default: e). Example: contract name  
is "web" and suffix is "e", cloned contract will be named "web_e"
```

where the **--mode** option gives you these options:

- o **one-to-one**: Migrates a single EPG to a single ESG.
- o **optimized**: The default option. The script analyzes the contract relationships and tries to optimize the number of ESGs created. In other words, the script won't necessarily create 10 ESGs if it finds 10 EPGs; rather, it analyzes how your EPGs are providing and consuming contracts and optimizes the result. For example, if two EPGs are consuming the same contracts and providing the same contract, then the script groups those together into a single ESG. Note that if the relationships are too complex, the **optimized** option falls back to the more basic **one-to-one** mode.



- Do not use the **--disableNdMode** option. The Nexus Dashboard compatibility mode (**NdMode**) is enabled by default and should be left enabled for these procedures.
- Since you are running this script on an APIC, you will not need to use these options:
 - **--apic APIC**
 - **--username USERNAME**

- `--password PASSWORD`

You would only use those options if you are running the script outside of your APIC.

For example, if you wanted to:

- Use the tenant `import (uni/tn-import)` as the tenant of interest to analyze all the EPGs
- Use the filename `migrate.yaml` for the output

then you would enter this command:

```
./ESGMigrationAssistant dryrun --tenantdns uni/tn-import --outYaml migrate.yaml
```

Similarly, if you wanted to:

- Filter the analysis to the VRF `uni/tn-common/ctx-vrf-common`
- Use the filename `migrate.yaml` for the output

then you would enter this command:

```
./ESGMigrationAssistant dryrun --vrfdns uni/tn-common/ctx-vrf-common  
--outYaml migrate.yaml
```



Do not enter the `dryrun` command with both `--tenantdns` and `--vrfdns` options. The `ESGMigrationAssistant` script is designed so that the `--tenantdns` option will migrate the EPGs to ESGs for all VRFs in that tenant, whereas the `--vrfdns` option will migrate the EPGs to ESGs for all the VRFs, and those VRFs might be used by multiple tenants (user/common).

After you enter the command to perform the dry run analysis, the script takes a snapshot of the APIC configuration and asks you if you want to POST this configuration.

2. Review the output from the dry run analysis.
3. Locate the YAML file that was produced in the step above and edit it, if necessary.

If you see any EPG to ESG migration from the output of the dry run analysis that you want to change, such as the names of the cloned ESG contracts, you can make those changes in the YAML file.

There is a section `contractClones` at the end of the YAML file. Clones essentially means that there is an original EPG contract, which is denoted by `cloneFromDn`, and the script uses this EPG contract to create the ESG contract. The `cloneName` fields in this section provide the name of the cloned ESG contracts, along information on the original EPG that it was cloned from.

4. Run the conversion phase of the ESG Migration Assistant script using the YAML file that was produced in the step above.

The conversion phase in the process converts the EPGs to ESGs, and EPG contracts to ESG

contracts.

Use the information provided with the `--help` option to determine how you want to run the ESG Migration Assistant conversion script.

```
./ESGMigrationAssistant conversion --help
usage: ESGMigrationAssistant conversion [-h] --inYaml INYAML --apic APIC [--
username USERNAME] [--password PASSWORD] [--noConfig]
      [--configStrategy {interactive,vrf}] [--outputFile OUTPUTFILE]

options:
  -h, --help            show this help message and exit
  --inYaml INYAML       YAML file in which we report the execution plan
  --apic APIC           APIC IP address or hostname to connect to
  --username USERNAME   Username for APIC
  --password PASSWORD   Password for APIC
  --noConfig            Proposed configuration is not applied to APIC
  --configStrategy {interactive,vrf}
                        Select the configuration strategy mode: in interactive mode (default)
                        EPGs/External EPGs are migrated one by one, in vrf
                        mode all EPGs/External EPGs assigned to a single VRF are migrated in a
                        single transaction
  --outputFile OUTPUTFILE
                        Output file for generated configuration (default: output.xml). Use .xml or
                        .json extension to save in respective format
```

- o If you don't want to deploy these changes to the APIC just yet and you want to see a preview of the changes, enter the command using the `--noConfig` option. The changes are saved to an XML or JSON file in that case, where you can review the changes that would be made using this script.
- o Run the conversion phase again without the `--noConfig` option when you are comfortable with the changes.

This is an example entry for the conversion phase:

```
./ESGMigrationAssistant conversion --inYaml migrate.yaml --configStrategy vrf
--outputFile apic-config.xml
```

After you enter the command to perform the conversion phase, the script takes a snapshot of the APIC configuration and asks you if you want to POST this configuration.

The script also prompts you at each step of the conversion phase.

- o If you want to make all the conversions shown in a step, enter **A** for Yes to All.
- o If you want to split the conversion into pieces in a step, enter **Y** for Yes.

- o If you do not want to make any of the changes in this step in the conversion phase, enter **N** for No.
 - o If you want to quit out of the conversion phase entirely, enter 1 for Quit.
5. Run the cleanup phase of the ESG Migration Assistant script.

The cleanup phase in the process removes all of the unused EPG contracts.

Use the information provided with the **--help** option to determine how you want to run the cleanup phase of the ESG Migration Assistant script.

```
./ESGMigrationAssistant cleanup --help
usage: ESGMigrationAssistant cleanup [-h] --apic APIC [--username USERNAME] [--password PASSWORD] [--noConfig]
                                     [--configStrategy {interactive,vrf,global}] [--outputFile OUTPUTFILE]

options:
  -h, --help            show this help message and exit
  --apic APIC           APIC IP address or hostname to connect to
  --username USERNAME  Username for APIC
  --password PASSWORD  Password for APIC
  --noConfig           Proposed configuration is not applied to APIC
  --configStrategy {interactive,vrf,global}
                       Select the configuration strategy mode: in interactive mode (default)
                       EPGs/External EPGs are cleaned up one by one, in vrf
                       mode all EPGs/External EPGs assigned to a single VRF are cleaned up in
                       a single transaction, in global mode (not
                       recommended unless noConfig option is used) all EPGs/External EPGs are
                       cleaned up in a single transaction
  --outputFile OUTPUTFILE
                       Output file for generated configuration (default: output.xml). Use .xml or
                       .json extension to save in respective format.
```

This is an example entry for the cleanup phase:

```
./ESGMigrationAssistant conversion --inYaml migrate.yaml --configStrategy vrf
--outputFile apic-config.xml
```

Verify the migration updates in APIC

1. Log into your APIC and navigate to the tenant that you used in the dry run analysis in the EPG-to-ESG migration procedures.
2. Verify that the new Endpoint Security Groups (ESGs) were created successfully.
 - a. Navigate to:

import > Application Profiles > *app_profile* > Endpoint Security Groups

- b. Verify that the Endpoint Security Groups (ESGs) are displayed on this page.
3. Verify that the ESG contracts were cloned successfully.
 - a. Navigate to:

import > Contracts > Standard

- b. For each contract, verify that the contract and ESG relationships are configured correctly.

For each contract, in the contracts page, access the topology view to see the ESGs that are providing and consuming that contract.

Create fabric group and associate fabrics

In Nexus Dashboard:

1. If the ACI fabric is not already added to the Nexus Dashboard, onboard it to the Nexus Dashboard using multi-cluster connectivity.

See [Connecting Clusters](#) for more information.

2. Create the fabric group and associate the fabrics.

You will import tenant policies from ACI fabrics and you will integrate both ACI and NX-OS fabrics into a single VXLAN-ACI fabric group. For more information, see [\[Create fabric groups\]](#).

- a. Create a VXLAN-ACI fabric group.
- b. Add the ACI and VXLAN fabrics to the VXLAN-ACI fabric group in the correct order.

You must add the ACI and VXLAN fabrics to the VXLAN-ACI fabric group in this order:

- First, add the ACI fabric to the new VXLAN-ACI fabric group.
- Then add the VXLAN fabric to the new VXLAN-ACI fabric group.

Import the tenant and tenant policies from ACI into Nexus Dashboard

Use these procedures to:

- Create an APIC tenant into the Nexus Dashboard and associate the tenant to the ACI fabric.

Because the ACI fabric is a member of the VXLAN-ACI fabric group, these procedures associate the tenant to the VXLAN-ACI fabric group as well.

- Import tenant policies from the ACI fabric.

In Nexus Dashboard:

1. Import the first tenant from ACI into the Nexus Dashboard.
 - a. Follow the procedures provided in the "Import a tenant from ACI" section in [Configuring Tenants and Tenant Domains](#) to complete this task.

Enter these values in this procedure:

- **Fabric:** Choose the ACI fabric that you added to the VXLAN-ACI fabric group in [Create fabric group and associate fabrics](#).
 - **ACI tenants on fabric:** Choose the first tenant from ACI that you want to import into the Nexus Dashboard.
- b. Enter the necessary values for the remaining fields in the **Import tenant (ACI)** page, then click **Save**.
2. Repeat this step to import the remaining tenants from ACI into the Nexus Dashboard, if necessary.

Once you have imported all the necessary tenants from ACI into Nexus Dashboard, navigate into the VXLAN-ACI fabric group and click **Segmentation and security > Tenants** to verify that all of the imported tenants are associated with this fabric group.

3. Choose the first tenant that you imported from ACI into the Nexus Dashboard.
4. Import the tenant policies from the ACI fabric.

Follow the procedures provided in the section "Import tenant policies" in [Working with Segmentation and Security for Your Nexus Dashboard VXLAN Fabric](#) to complete this task.

5. Review or edit the policies that you've created.

Once you have completed all of the tasks in this [Import tenant policies from ACI fabrics](#) section, you can now review the policies that you've created and edit them, if necessary (for example, you might want to stretch a VRF from the policy). Refer to these sections in the [Working with Segmentation and Security for Your Nexus Dashboard VXLAN Fabric](#) for more information:

- [Working with networks](#)
- [Working with VRFs](#)
- [Working with security groups](#)
- [Working with security contracts](#)

6. Deploy the tenant, if necessary.

After importing the tenant policies, you can deploy the tenant that you used when you imported the policies. See the section "Deploy tenants" in [Working with Segmentation and Security for Your Nexus Dashboard VXLAN Fabric](#) for more information.

Configuring tenant domains

These sections describe how to configure tenant domains.

- [Understand the Tenant domains page](#)
- [Create a tenant domain](#)
- [Edit a tenant domain](#)
- [Delete a tenant domain](#)

Understand the Tenant domains page

1. [Navigate to the Multi-tenancy page.](#)
2. Click **Tenant domains**.

Previously-configured tenant domains are displayed in the **Tenant domains** page, with this information.

Field	Description
Name	Provides the name of the tenant domain. For user-created tenant domains, click the entry in the Name column to bring up detailed information on a tenant domain.
Description	Provides the description information for the tenant domain.
Member tenants	Shows the tenants that are members of each tenant domain. For user-created tenant domains, click the entry in the Member tenants column to bring up detailed information on a member tenant.

Create a tenant domain

- You can create a tenant domain on an individual cluster or federated cluster.
- In a federation, the creation may take longer as it will always be proxied to the federation primary and synced to all other clusters.
- Tenant domains reference tenants but are not exclusive containers.
- Many tenant domains may include a given tenant.

To create a tenant domain:

1. [Navigate to the Multi-tenancy page.](#)
2. Click **Tenant domains**.

Previously-configured tenant domains are displayed in the **Tenant domains** page.

3. Click **Actions > Create tenant domain**.

The **Create tenant domain** page is displayed.

4. Enter a name and description for the tenant domain.

5. In the **Tenants** area:

- Choose an already-configured tenant from the list of tenants, or
- Click **+Create tenant** to create a new tenant for this tenant domain.

See [Create a tenant](#) for more information.

6. Click **Save**.

The newly-configured tenant domain is displayed in the **Tenant domains** page.

Edit a tenant domain

1. [Navigate to the Multi-tenancy page](#).
2. Click **Tenant domains**.

Previously-configured tenant domains are displayed in the **Tenant domains** page.

3. Click the box next to the tenant domain that you want to edit.
4. Click **Actions > Edit**.

The **Edit tenant domain** page is displayed.

5. Make the necessary changes to the tenant domain, then click **Save**.

Delete a tenant domain

1. [Navigate to the Multi-tenancy page](#).
2. Click **Tenant domains**.

Previously-configured tenant domains are displayed in the **Tenant domains** page.

3. Click the box next to the tenant domain that you want to delete.
4. Click **Actions > Delete**.

The tenant domain is now deleted from the list of tenant domains displayed in the in the **Tenant domains** page.

Information that is displayed or not displayed

With multi-tenancy enabled, you might see information, and information might not be displayed to you, based on several factors.

- [Information that is displayed or not displayed based on tenant and tenant domain associations](#)
- [Anomalies information that is displayed or not displayed](#)
- [Endpoints](#)

Information that is displayed or not displayed based on tenant and tenant domain associations

When you have tenants and tenant domains configured, information may or may not be displayed to you in certain pages, depending on the associations that you set between the tenants, tenant domains, and users in your Nexus Dashboard.

The best way to describe this concept is to first provide an example Nexus Dashboard configuration:

User	Associated tenants	Associated tenant domains	Notes
admin	Tenant1-1 Tenant1-2 Tenant1-3 Tenant2-1 Tenant3-1	all-tenants-domain	Because the admin user is able to create and edit tenants and tenant domains, all of these tenants in the Nexus Dashboard are essentially associated with the admin user. However, each user is only allowed to be associated with a single tenant domain, so the admin user is associated with the all-tenants-domain in this situation.
User1	Tenant1-1 Tenant1-2 Tenant1-3	TenantDomain1	User1 is associated with TenantDomain1, and Tenant1-1, Tenant1-2, and Tenant1-3 are members of TenantDomain1.
User2	Tenant2-1	TenantDomain2	User2 is associated with TenantDomain2, and Tenant2-1 is a member of TenantDomain2.
User3	Tenant1-1 Tenant2-1 Tenant3-1	TenantDomain3	User3 is associated with TenantDomain3, and Tenant1-1, Tenant2-1, and Tenant3-1 are members of TenantDomain3.

and then to describe what tenant- and tenant domain-based information is shown or hidden when you log in as any of these users:

- [Information displayed and hidden: Admin user role](#)
- [Information displayed and hidden: User1 user role](#)
- [Information displayed and hidden: User2 user role](#)
- [Information displayed and hidden: User3 user role](#)

Information displayed and hidden: Admin user role

In this scenario, you are logged into the Nexus Dashboard as admin.

As admin, you are able to see all of the tenants and tenant domains that are configured under **Admin > Multi-tenancy**, and nothing is hidden from you. In addition, you are also able to see all tenant- and tenant domain-based information, as described in [Additional areas affected by tenant and tenant domain creation](#).

Information displayed and hidden: User1 user role

In this scenario, you are logged into the Nexus Dashboard as User1.

Because User1 is associated with TenantDomain1, you are only able to see tenant-based information for these tenants that are associated with TenantDomain1:

- Tenant1-1
- Tenant1-2
- Tenant1-3

and you are not able to see any tenant- or tenant domain-based information for these tenants or tenant domains:

- TenantDomain2, TenantDomain3
- Tenant2-1, Tenant3-1

Information displayed and hidden: User2 user role

In this scenario, you are logged into the Nexus Dashboard as User2.

Because User2 is associated with TenantDomain2, you are only able to see tenant-based information for these tenants that are associated with TenantDomain2:

- Tenant2-1

and you are not able to see any tenant- or tenant domain-based information for these tenants or tenant domains:

- TenantDomain1, TenantDomain3
- Tenant1-1, Tenant1-2, Tenant1-3, Tenant3-1

Information displayed and hidden: User3 user role

In this scenario, you are logged into the Nexus Dashboard as User3.

Because User3 is associated with TenantDomain3, you are only able to see tenant-based information for these tenants that are associated with TenantDomain3:

- Tenant1-1
- Tenant2-1
- Tenant3-1

and you are not able to see any tenant domain-based information for these tenant domains, even though Tenant1-1, Tenant2-1, and Tenant3-1 are also members of each of these tenant domains:

- TenantDomain1, TenantDomain2

This is because, while a tenant can be a member of multiple tenant domains as described in [Understand tenant domains](#), the tenant- or tenant domain-based information that is displayed is based on the tenant domain that User3 is associated with, which is TenantDomain3.

In addition, you are not able to see any tenant-based information for these tenants because they are not associated with TenantDomain3:

- Tenant1-2, Tenant1-3

Anomalies information that is displayed or not displayed

As they relate to tenants and tenant domains, anomalies are broken down into these categories:

- Infra anomalies (anomalies at the Nexus Dashboard system level)
- Tenant-level anomalies

In the Anomalies areas in the GUI, users can only view the portion of the graph that they have access to.

For example, here are two areas in the GUI where you can view anomaly information:

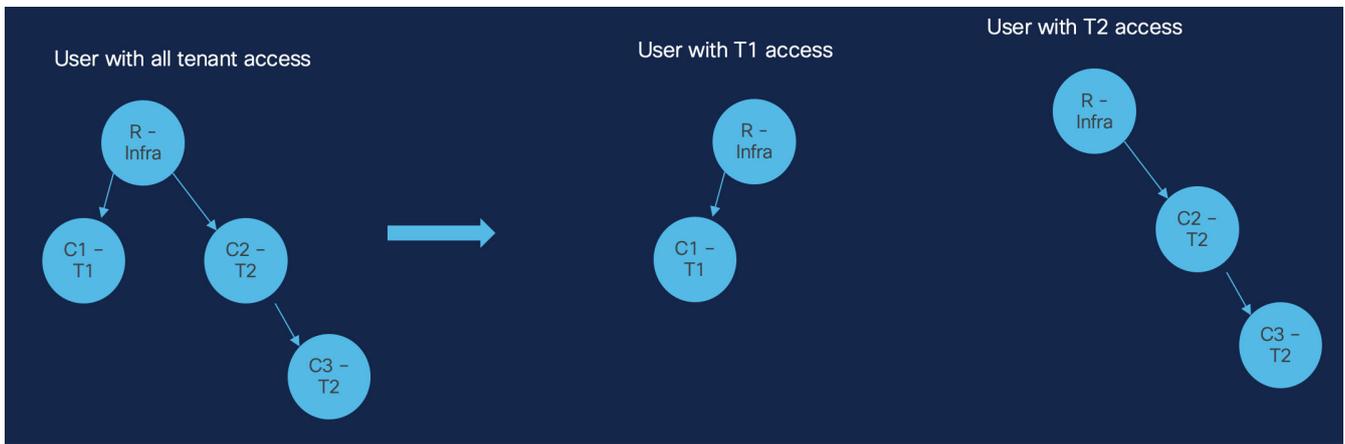
- In the main **Anomalies** page (**Analyze > Anomalies**), where you might see root (Infra) level anomalies, as well as anomalies associated with specific tenants.
- By choosing the **Root cause and uncorrelated anomalies** option in the **Anomalies** page, where you would see correlated anomalies associated with all tenants collapsed into a single row. In addition, if you were to click on a root level anomaly, you would see a graph that would display the correlated anomalies based on the tenant relations with this user.

In addition:

- In the analyze anomalies list, users can view the list of anomalies that are either classified as Infra or affecting the tenants that they have access to.
- Similarly, when looking at a graph, users can only view the portion of the graph that they have access to.
- Users can only update anomalies that they have access to. They cannot update Infra level anomalies unless they are associated with the **all-tenants-domain**.

These figures provide examples of situations where Anomalies areas in the GUI might be displayed or might not be displayed to certain users. In each of these figures:

- R signifies a root (Infra) anomaly or event
- C signifies a correlated anomaly or event



Anomalies: Two Tenants

In this figure:

- The user at the far left of the figure (User1) is associated with the **all-tenants-domain** and is able to see Infra (Nexus Dashboard-level) anomalies, as well as correlated anomalies associated with tenants T1 (C1) and T2 (C2 and C3).
- The user in the middle of the figure (User2) is associated only with the **T1** tenant, and:
 - Is able to see Infra level anomalies
 - Is able to see correlated anomalies where the impact is seen at the **T1** tenant level (C1)
 - Is not able to see correlated anomalies where the impact is seen at the **T2** tenant level (C2 or C3)
- The user in the far right of the figure (User3) is associated only with the **T2** tenant, and:
 - Is able to see Infra level anomalies
 - Is able to see correlated anomalies where the impact is seen at the **T2** tenant level (C2 and C3)
 - Is not able to see correlated anomalies where the impact is seen at the **T1** tenant level (C1)



Anomalies: Multiple Tenants

In this figure:

- The user at the far left of the figure (User1) is associated with the **all-tenants-domain** and is able to see Infra (Nexus Dashboard-level) anomalies, as well as Infra anomalies related to tenants T1

(C1), T2 (C2), and T3 (C3 and C4).

- The user in the middle of the figure (User2) is associated only with the T1 tenant, and:
 - Is able to see Infra level anomalies
 - Is able to see correlated anomalies where the impact is seen at the T1 tenant level (C1)
 - Is not able to see correlated anomalies where the impact is seen at the T2 tenant level (C2) or at the T3 tenant level (C3 and C4)
- The user in the far right of the figure (User3) is associated only with the T2 tenant, and:
 - Is able to see Infra level anomalies
 - Is able to see correlated anomalies where the impact is seen at the T2 tenant level (C2)
 - Is not able to see correlated anomalies where the impact is seen at the T1 tenant level (C1) or at the T3 tenant level (C3 and C4)

Note that the user that is associated with the T3 tenant (User4) can see tenant-level anomalies for the tenant T3 (C3 and C4), as well as the root Infra anomalies; however, User4 is not able to see anomalies associated with tenants T1 (C1) or T2 (C2). In addition, User4 is not able to see that certain Infra anomalies are correlated to C3 through C2.

With regards to the information displayed in the Anomalies tables:

- There is no change for users who are associated with the **all-tenants-domain**.
- Users with tenant-specific access will not see a count value for the Correlated anomalies/events column, and sorting will be disabled. This is applicable to both ungrouped and root events views.

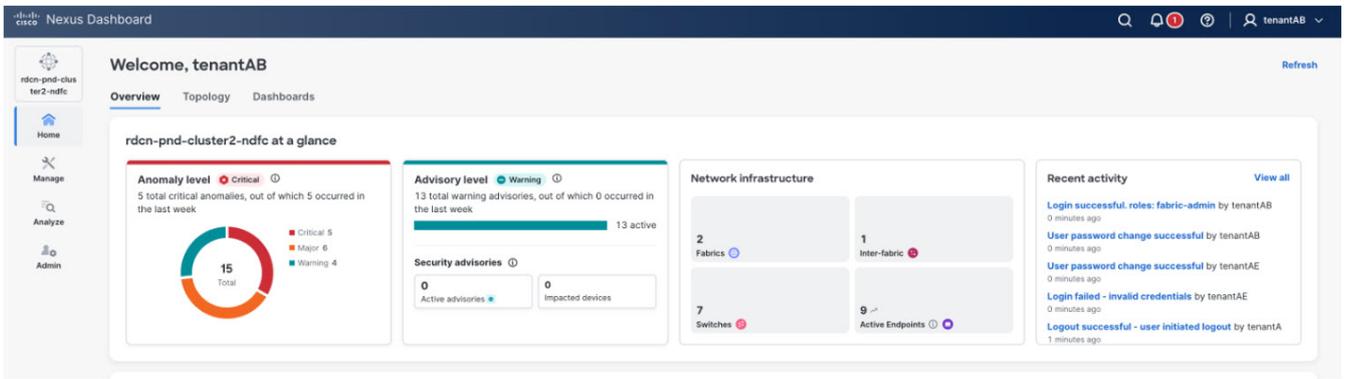
Detection time	Correlated anomalies/events ⓘ	Correlation time	
August 4, 2025 11:11:11 PM	View details	August 4, 2025 11:11:11 PM	
August 4, 2025 11:11:11 PM	-	August 4, 2025 11:11:11 PM	
August 4, 2025 11:11:11 PM	View details	August 4, 2025 11:11:11 PM	
August 4, 2025 11:11:11 PM	-	August 4, 2025 11:11:11 PM	
August 4, 2025 11:11:11 PM	View details	August 4, 2025 11:11:11 PM	

Rows per page < 2 ... 10 >

Anomalies table

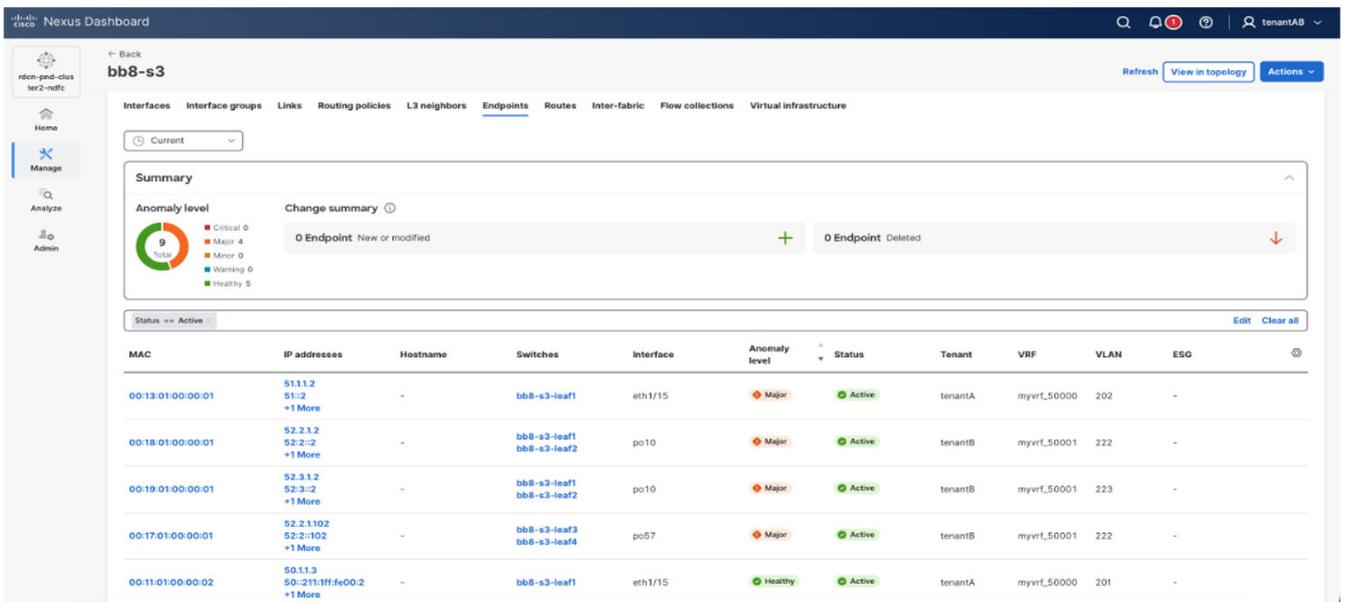
Endpoints

The active endpoints count that is displayed in the global view shows the number of endpoints based on the tenants that a user has access to across all fabrics in the cluster.

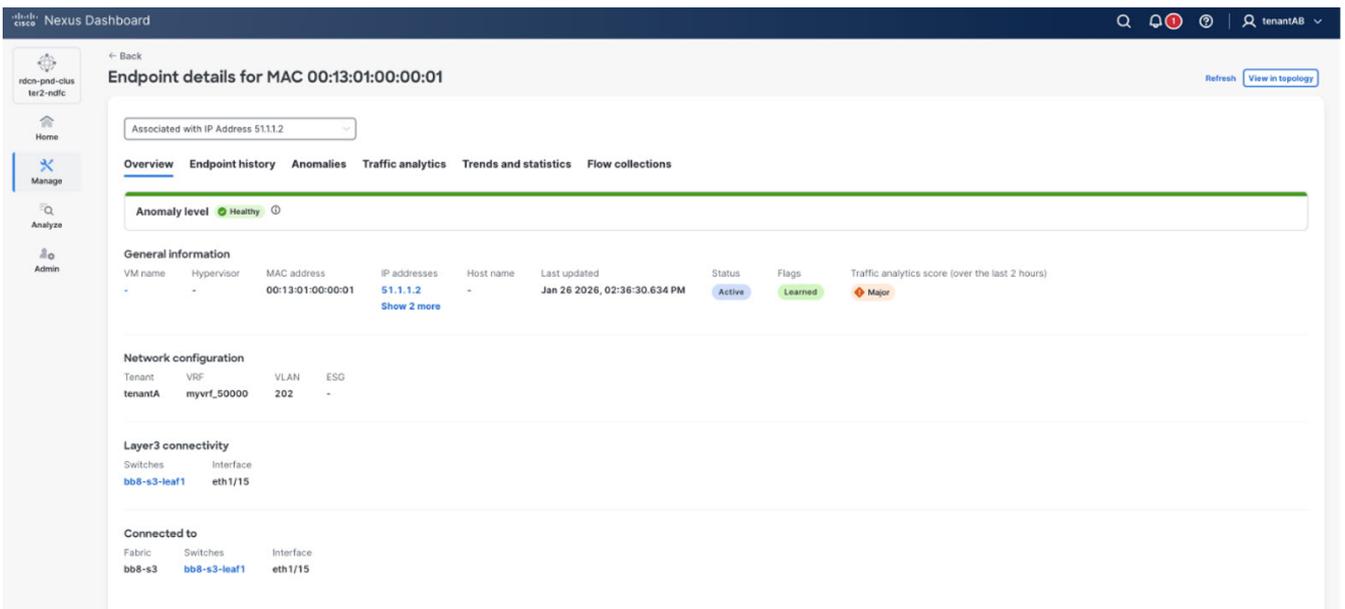


In the fabric view, endpoints are listed based on user access in the tenant-associated fabric.

In the following figure, the user **tenantAB** has access to **tenantA** and **tenantB** only.



The **Endpoint details** page shows tenant information for specific endpoints.



These endpoint anomalies are tenant-scoped and are filtered based on user access:

- ENDPOINT_DUPLICATE_MAC

- ENDPOINT_DUPLICATE_IP
- ENDPOINT_PRIVATE_SUBNET_IP
- ENDPOINT_MOVE_BETWEEN_LEAFS

The following figure shows anomalies that have endpoints in affected entities that are tenant-scoped. For example, the interface anomaly shown in the figure lists affected endpoints based on user access.

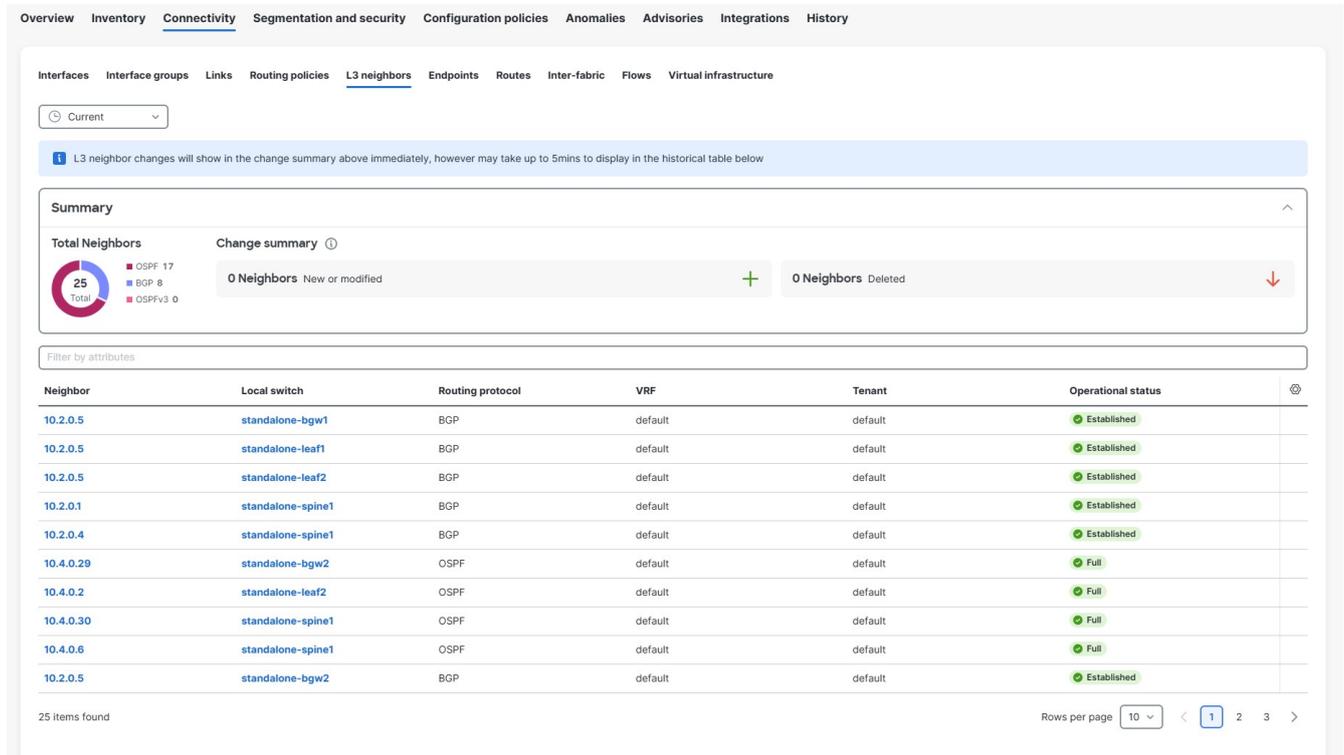
Connectivity analysis

When starting a new connectivity analysis job, only endpoints in the tenants that a user has access to are listed in the source and destination drop-down lists for both IP and MAC.

- Connectivity analysis jobs created by a user with **tenantA** access is not visible to the **tenantB** user.
- When a user starts a multi-site job across tenants, the job fails to reach a tenant that the user doesn't have access to.

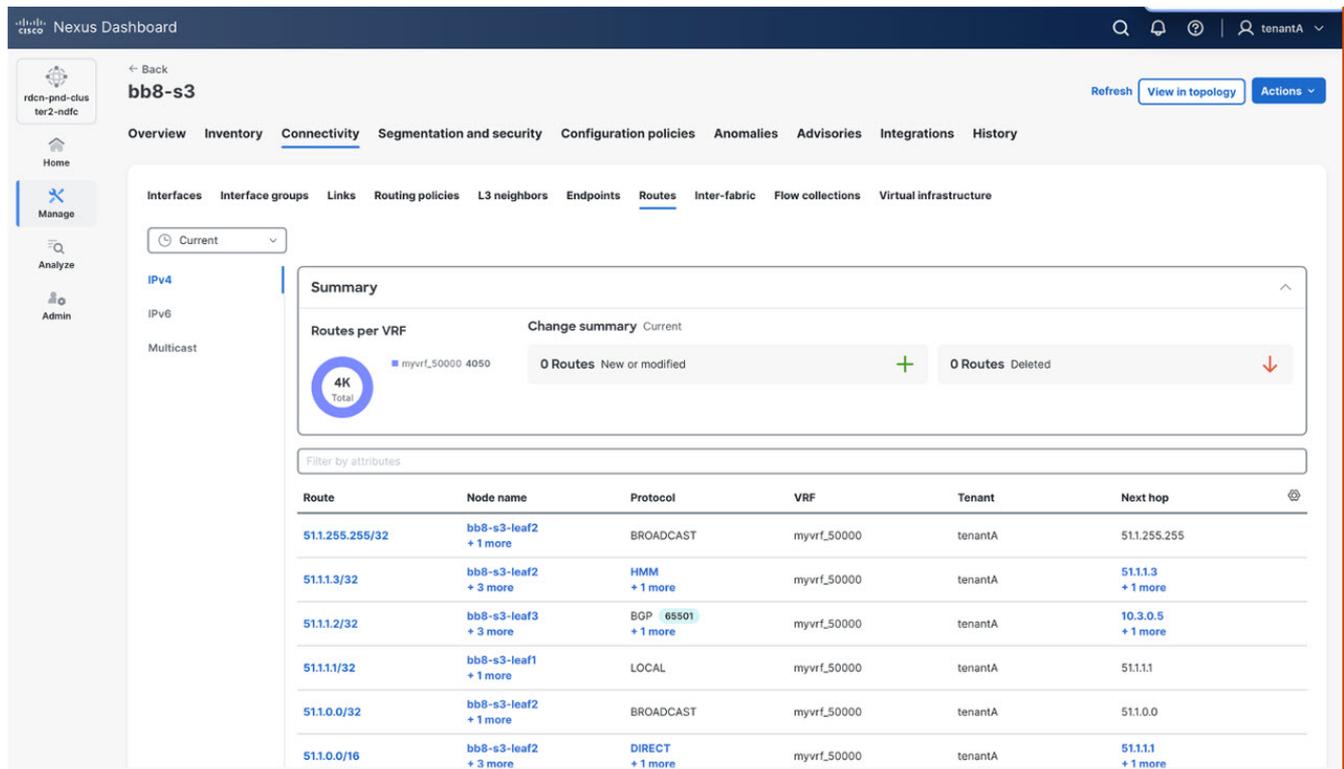
L3 neighbors

The following figure shows an example of L3 neighbors in a tenant-associated fabric. In this page, the tenant and VRF fields are added for L3 neighbors.



Routes

The following figure shows an example of IPv4/IPv6 routes in a tenant-associated fabric.



The following figure shows an example of multicast routes in a tenant-associated fabric.

- The multicast routes table view lists the routes based on a user's tenants access.
- In this page, the tenant column is added in the browse table.

Multicast group	Sources	Sender traffic rate (last 5 min)	Sender packet rate (last 5 min)	Receivers	Tenant	VRF	Status	Last updated
225.0.1.1	-	0 bps	0 packets/s	1 →	default	vrf01	Inactive	Jan 21 2026 05:16:55.569 AM
225.0.1.15	42.0.11.23	2.48 mbps	2.03 K packets/s	1 →	default	vrf01	Active	Jan 21 2026 06:31:55.578 AM
225.0.1.2	-	0 bps	0 packets/s	1 →	default	vrf01	Inactive	Jan 21 2026 05:16:55.569 AM
225.0.1.3	33.0.130.3	1.01 mbps	1.01 K packets/s	1 →	default	vrf01	Active	Jan 21 2026 06:31:55.578 AM
226.0.1.100	-	0 bps	0 packets/s	1 →	default	vrf02	Inactive	Jan 21 2026 05:16:55.569 AM
226.0.1.101	42.0.23.2	2.62 mbps	2.03 K packets/s	1 →	default	vrf02	Active	Jan 21 2026 06:31:55.578 AM
226.0.1.102	42.0.20.2	2.51 mbps	2.03 K packets/s	1 →	default	vrf02	Active	Jan 21 2026 06:31:55.578 AM
239.1.1.1	10.251.1.5 + 1 more	128.00 bps	0 packets/s	0 →	default	default	Active	Jan 21 2026 06:31:55.578 AM
239.1.1.99	10.251.1.3 + 2 more	128.00 bps	0 packets/s	0 →	default	default	Active	Jan 21 2026 06:31:55.578 AM

Switch view: Multicast

The following figures show examples of PIM/IGMP/IGMP snooping protocols tables that show the tenant field in the tenant associated NX-OS fabrics.

Interface	Admin state	Operational status	Tenant	VRF	IP address	Querier address	Membership count	Querier version	Errors
Vlan23	Enabled	Up	default	vrf02	42.0.23.1	42.0.23.1	1	v3	4191
Vlan10	Enabled	Up	default	vrf01	42.0.10.1	42.0.10.1	1	v3	4190
Ethernet1/33	Enabled	Up	default	vrf01	33.0.131.1	33.0.131.1	1	v2	0

Tenant	VLAN	Admin state	Querier address	Querier version	Multicast routing state	Fabric Querier State	Errors
default	vlan201	Enabled	42.0.201.1	v2	No	No	0
default	vlan202	Enabled	42.0.202.1	v2	No	No	0

Additional areas affected by tenant and tenant domain creation

After you create a tenant or a tenant domain, and after you perform additional tenant-related configurations, such as associating fabrics with a tenant or importing a tenant from ACI, certain tenant- and tenant domain-based information also becomes available in these areas.

In each of these areas, you will only see information based on the associations that you set between the tenants, tenant domains, and users in your Nexus Dashboard, as described in [Information that is displayed or not displayed based on tenant and tenant domain associations](#).

For example, assume you have these user, tenant, and tenant domains configured on your Nexus Dashboard:

User	Associated tenants	Associated tenant domains
User1	Tenant1-1, Tenant1-2	TenantDomain1
User2	Tenant2-1, Tenant2-2	TenantDomain2



- If you are logged in User1, you will only see tenant- and tenant domain-based information for other tenants in TenantDomain1 (Tenant1-1, Tenant1-2) and you will not see any tenant- and tenant domain-based information for users or tenants in TenantDomain2 because User1 is not associated with TenantDomain2.
- Similarly, in some configuration pages, you might see only tenant- and tenant domain-based information that are related to your tenant or tenant domain, and other tenants or tenant domains are not displayed.

For example, assume that you create a VRF and you associate Tenant1-1 (from TenantDomain1) with that VRF. Then, if you were to create a network next and you associated Tenant1-1 with that network, in the **VRF name** field, only VRFs associated with Tenant1-1 are available to choose from. You will not see any options associated with different tenants.

- **When you add local users:** You can associate a tenant domain with a local user in the **Create local user** page. See [Configuring Users, Roles, and Security](#) for more information.
- Under **Connectivity**:
 - **For tenant-based overlays:** Tenant-based overlays are attached to interfaces. Overlays that you see as a tenant-specific user are restricted to the tenants that you have access to. See [Working with Connectivity in Your Nexus Dashboard LAN Fabrics](#) for more information.
 - **To display endpoints associated with a tenant:** Within a fabric that has tenants associated with it, navigate to **Connectivity > Endpoints** to view endpoints associated with a tenant. See [Working with Connectivity in Your Nexus Dashboard LAN Fabrics](#) for more information.
 - **To display L3 neighbors associated with a tenant:** Within a fabric that has tenants associated with it, navigate to **Connectivity > L3 neighbors** to view L3 neighbors associated with a tenant. See [Working with Connectivity in Your Nexus Dashboard LAN Fabrics](#) for more information.

- **VMs shown based on tenancy:** VMs are filtered in the **Virtual infrastructure** page based on tenancy. See [Working with Connectivity in Your Nexus Dashboard LAN Fabrics](#) for more information.
- Under **Segmentation and security:**
 - **In the Segmentation and security area for your fabric:** The **Tenants** tab becomes available under the **Segmentation and security** area for your fabric, allowing you to view segmentation and security information at the tenant level. See [Working with Segmentation and Security for Your Nexus Dashboard VXLAN Fabric](#) or [Working with Segmentation and Security for Your Nexus Dashboard ACI Fabric](#) for more information.
 - **When you create or view a VRF:** You can associate a tenant with a VRF in the **Create a VRF** page if a tenant is associated with the fabric that is associated with a VRF. See [Working with Segmentation and Security for Your Nexus Dashboard VXLAN Fabric](#) or [Working with Segmentation and Security for Your Nexus Dashboard ACI Fabric](#) for more information.
 - **When you create or view a network:** You can associate a tenant with a network in the **Networks** page. See [Working with Segmentation and Security for Your Nexus Dashboard ACI Fabric](#) for more information.
 - **When you create or view an ESG/security group:** You can associate a tenant with an ESG/security group in the **Security groups** page. See [Working with Segmentation and Security for Your Nexus Dashboard VXLAN Fabric](#) or [Working with Segmentation and Security for Your Nexus Dashboard ACI Fabric](#) for more information.
 - **When you create or view a security contract:** You can associate a tenant with a security contract in the **Security contracts** page. See [Working with Segmentation and Security for Your Nexus Dashboard ACI Fabric](#) for more information.
 - **When you create or view a security association:** You can associate a tenant with a security association in the **Security associations** page. See [Working with Segmentation and Security for Your Nexus Dashboard VXLAN Fabric](#) for more information.
 - **When you create or view an L3Out:** You can associate a tenant with an L3Out in the **L3Outs** page. See [Working with Segmentation and Security for Your Nexus Dashboard VXLAN Fabric](#) or [Working with Segmentation and Security for Your Nexus Dashboard ACI Fabric](#) for more information.
- Under **Configuration policies:**
 - **Policies** and **Resources** are filtered under the **Configuration policies** based on tenancy. For example, resources are VLANs. A VLAN might be associated with a specific overlay network and that overlay network might be for a specific tenant, so VLAN information under **Resources** might be shown or might not be shown based on tenancy.
- Under **Anomalies** and **Advisories:** Anomalies and advisories that are displayed are also based on tenancy. See [Multi-tenancy considerations for flow telemetry and traffic analytics](#) for more information.
- Under **History:** For example, audit logs are filtered in the **History** page based on tenancy. See [Reviewing History and Logs in Your Nexus Dashboard](#) for more information.

Multi-tenancy considerations for flow telemetry and traffic analytics

There are four areas where flow telemetry and traffic analytics are affected by multi-tenancy.

- [Flow configuration](#)
- [Ingest](#) (including enrichment and processing)
- [Analysis/visualization](#)
- [Anomalies](#)

Flow configuration

Flow configuration can be broken down into these two levels:

- **Feature configurations**—This is where you enable and disable flow telemetry, Netflow, SFlow, and traffic analytics features. Configurations done at this level are fabric-wide configurations.
- **Flow rule configurations**—Flow rule configurations take place at these levels:
 - **Global**—These configurations take place under a tenant. The configuration of these rules are scoped to users with access to the specific tenant. However, due to the interaction with interface rules as described below, the scope is restricted to users with access to ALL tenants.
 - **Interface**—These configurations take place under an interface. While interfaces are not tenant-scoped, there is a limitation with respect to the co-existence of interface and global rules.
 - In ACI fabrics, a switch can be configured in one of two modes: **vrf-scope** or **intf-scope**. If an interface rule is configured on a switch, all VRF rules on it are deleted from the switch.
 - In NX fabrics, interface and VRF rules co-exist but interface rules take precedence over VRF rules in the programmed ACLs.

Since the configuration of interface flow rules can impact the operational functionality of existing and future configured VRF rules, Nexus Dashboard limits the configuration of all flow rules to users with access to ALL tenants.

Users that are tenant-scoped have a view-only access and will be able to view only certain pages.

Ingest

The flow record exported from a switch contains the numerical representation of the tenant/VRF, along with other fields. The mapping of the numerical values to human understandable names that you have configured happens through the enrichment stage of the flow processing pipeline.

- ACI fabrics have had support for tenants prior to Nexus Dashboard release 4.2.1. Nexus Dashboard subscribes the relevant objects from the APIC/switches to derive the necessary enrichment mappings.
- Tenant support for NX-OS fabrics is new to Nexus Dashboard and is part of the multi-tenancy

feature introduced in Nexus Dashboard release 4.2.1. Nexus Dashboard pushes the tenant configurations to the switches as part of the VRF configuration, prefixing the VRF names with the corresponding tenant name, separated by a “:”. Switches export the VRF ID to VRF name (including the tenant name) mappings to Nexus Dashboard.

The multi-tenancy core layer also maintains a mapping between the Nexus Dashboard tenant names and fabric tenant names, which can be different.

The enriched data, along with tenant info, is then queried through the Nexus Dashboard for the analysis/visualization step, described below.

Analysis/visualization

In the analysis/visualization point in the process, based on the tenant domains (which translates to tenants) that the user has access to, Nexus Dashboard filters the API/GUI response only to the user-accessible tenant list.

- For Flow Telemetry, a user can view a flow only if they have access to the tenants on both ends (source and destination) of the flow.
- For Traffic Analytics, a user is able to view all services within tenants that they have access to. This is regardless of whether they have access to the tenants on the client-side of the conversation. The data for each service (health score, client, SYN, RST counts, and so on) will only aggregate conversations for which the user has access to the tenants on the client-side.



The filter for aggregation of conversations does not apply to the Tx (total), Rx (total), Tx Rate (total), and Rx Rate (total) columns.

- The service detail page will only show clients within tenants to which a user has access. Consequently, a flow troubleshoot job for a conversation can only be triggered if the user has access to tenants on both ends of the conversations.

Anomalies

Configuration anomalies are not tenant-scoped and do not go through the tenant filters when queried by a user. Flow and traffic analytics anomalies are tagged with the tenant fields and are filtered based on a user access list.

Multi-tenancy considerations for endpoints and multicast routes

Similar to flow, any endpoint and multicast records that are received from the switch contains the numerical representation of the tenant/VRF and tenant/VLAN information, which are then mapped to the user-configured Nexus Dashboard tenant names.

Endpoint and multicast anomalies are enriched to include tenant information to facilitate the tenant filtering based on user access.

Multi-tenancy considerations for connectivity analysis

Connectivity analysis (CA) supports multi-tenancy by enforcing tenant-specific validations for both VRF and VLAN contexts, but only for VXLAN fabrics. User inputs are ingested through two primary

mechanisms:

- **Auto-population:** From the endpoint database (EP DB).
- **Manual entry:** The source VRF is mandatory, but the destination VRF is optional.

IP-based Flows

- **Known Source and Destination VRFs:** When both source and destination VRFs are provided, CA validates these VRFs against the allowed tenant list (as provided by infra). If the validation fails, CA terminates the analysis for non-compliant flows.
- **Unknown Destination VRF:** If the destination VRF is not specified, CA retrieves the VRF from the RPM response received from each switch. The retrieved VRF is then validated against the allowed tenant list before proceeding.

MAC-based Flows

- **Known Source and Destination VLANs:** CA queries the infra-provided API to obtain the list of allowed VLANs per tenant. When both source and destination VLANs are known, CA validates both against the tenant's allowed VLAN list and takes the appropriate action (continue or abort).
- **Unknown Destination VLAN:** If the destination VLAN is not provided, CA extracts the next-hop neighbor VLAN from the RPM response of each switch and performs validation against the allowed VLANs per tenant. Flow analysis proceeds only if the VLAN is permitted.



All tenant validation logic is dynamically referenced from infra-provided lists as part of the on-demand CA job, ensuring compliance with tenant-specific access policies throughout the analysis process. The "default" VRF is exempted from the interim validation check.

GUI changes

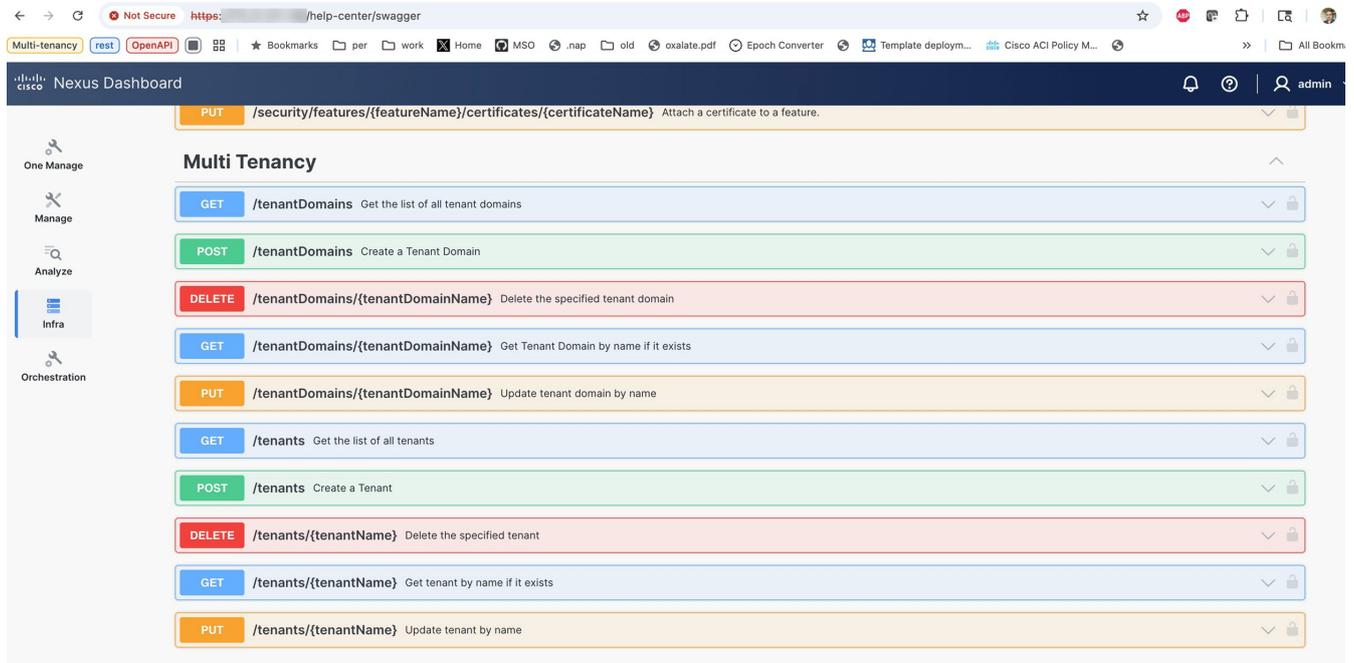
Starting with Nexus Dashboard release 4.2.1, the tenant field is shown in the GUI for NX-OS fabrics as well as ACI fabrics (which had tenant-related fields prior to Nexus Dashboard release 4.2.1), making the behavior consistent across fabric types.

These are the tenant-related changes now shown in the GUI for both NX-OS and ACI fabrics:

- CUD operations, allowed only for users with ALL-tenant access:
 - Tenant-scoped users have view-only permissions and can view the configurations.
 - Feature level configuration
 - Flow rules
- Certain tables to show the tenant field for NX-OS fabrics (was previously shown only for ACI fabrics)
- Tables where changes are required in these areas:
 - Service endpoint tables for traffic analytics.
 - Client details table for traffic analytics.
 - Flow collections table for traffic analytics flow troubleshooting. This table is also used for the Flow browse page in flow telemetry/netflow/SFlow.
- If appropriate tenant-related information is available, then that table will display the relevant tenant field; otherwise, the tenant field will display in the table with an empty value (-).
- Users are shown tenant-related information and are allowed to edit fields when they have access to view, edit, create, or delete tenant-related configurations.

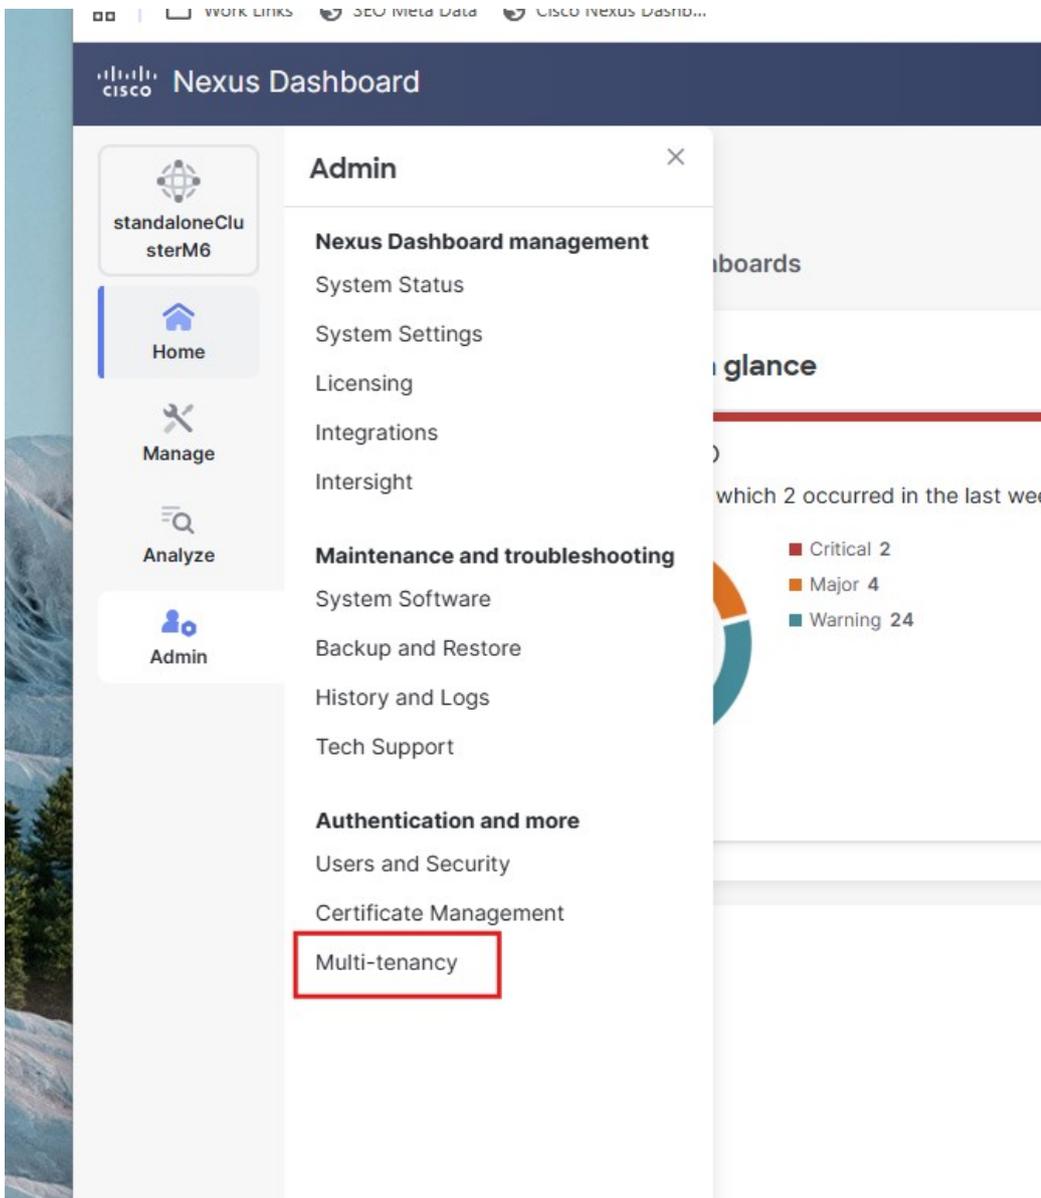
Legacy Tenant CRUD APIs

Prior to Nexus Dashboard release 4.2.1, Orchestration (NDO) was the only app that supported multi-tenancy. NDO has always supported the ability to create tenants, and to associate tenants to fabrics and users. This functionality has now been consolidated under the Nexus Dashboard platform. Starting with Nexus Dashboard release 4.2.1, the tenant CRUD (Create, Read, Update, and Delete) APIs and GUI are now available at the platform level.



Tenant CRUD in APIs

In order to be backwards compatible for automation, Nexus Dashboard release 4.2.1 continues to support the older tenant CRUD APIs that were available under NDO. However, the NDO GUI pages that are used to help with tenant management are not available in Nexus Dashboard, since those pages are now provided in the unified Nexus Dashboard platform.



Tenant CRUD in GUI

There are no changes to the tenant management APIs that were available in NDO. With regards to tenant management, the only difference between Nexus Dashboard release 4.2.1 and prior releases is that release 4.2.1 introduces the concept of a tenant domain. Users are now associated with a tenant domain that contains tenants rather than the older mechanism of associating each tenant with a user.

To accommodate this change and still be backwards compatible with the older NDO tenant CRUD APIs, Orchestrator in Nexus Dashboard performs the tenant domain management for the user association to tenants.

When upgrading from a previous release, each existing local user is assigned their own tenant domain, named `sys_<UserName>`. If the legacy tenant CRUD APIs are used for tenant management, then the user's tenant domain is updated according to the API requirements.

Alternatively, when upgrading from a previous release, we recommend that the admin reconfigures the users tenant domain according to their needs, and switches over to the Nexus Dashboard platform APIs and GUI for tenant management and its association with fabrics and users.

Legacy tenant CRUD APIs perform a one-to-one mapping of users to tenant domains. On Nexus

Dashboard, if a tenant domain is shared among multiple users and any changes are made to any user association for any of those tenants from legacy APIs, Nexus Dashboard will split them and create separate tenant domains for each of the users, named `sys_<UserName>`. So the original tenant domain is left with only one user associated with it and the remaining users are associated to their individual tenant domains with respective tenants.

You can use either the legacy NDO APIs for tenant management for the NDO application or the Nexus Dashboard platform APIs in a mixed fashion. However, there is a limitation on the mixed use of these APIs in that you cannot use them simultaneously. If you are using the NDO legacy APIs, you should wait for 30 seconds to switch over to using the Nexus Dashboard platform APIs.

For remote users, you should update the users AV pair to include the tenant domain. If the tenant domain is not specified in the AV pair, NDO will attempt to provide the user with tenant access based on the information stored in NDO for the remote user from the previous release. When upgrading from a previous release, we strongly recommend that you define new tenant domains as required and that you add tenant domains to the remote user's AV pair.

An example of the AV pair for a user is:

```
shell:domain=all/designer tenant-domain=TD_engineering
```

Copyright

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

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

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

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

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

Any Internet Protocol (IP) addresses and phone numbers used in this document are not intended to be actual addresses and phone numbers. Any examples, command display output, network topology diagrams, and other figures included in the document are shown for illustrative purposes only. Any use of actual IP addresses or phone numbers in illustrative content is unintentional and coincidental.

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

Cisco and the Cisco logo are trademarks or registered trademarks of Cisco and/or its affiliates in the U.S. and other countries. To view a list of Cisco trademarks, go to this URL: <https://www.cisco.com/go/trademarks>. Third-party trademarks mentioned are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (1110R)

© 2017–2026 Cisco Systems, Inc. All rights reserved.

Americas Headquarters

Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706

USA

<https://www.cisco.com>

Tel: 408 526-4000

800 553-NETS (6387)

Fax: 408 527-0883