



Cisco Cloud APIC Policy Model

- [About the ACI Policy Model, on page 1](#)
- [Policy Model Key Characteristics, on page 1](#)
- [Logical Constructs, on page 2](#)
- [The Cisco ACI Policy Management Information Model, on page 3](#)
- [Tenants, on page 4](#)
- [VRFs, on page 5](#)
- [Cloud Application Profiles, on page 6](#)
- [Cloud Endpoint Groups, on page 7](#)
- [Contracts, on page 9](#)
- [About the Cloud Template, on page 12](#)
- [Managed Object Relations and Policy Resolution, on page 15](#)
- [Default Policies, on page 16](#)
- [Shared Services, on page 17](#)

About the ACI Policy Model

The ACI policy model enables the specification of application requirements policies. The Cisco Cloud APIC automatically renders policies in the cloud infrastructure. When you or a process initiates an administrative change to an object in the cloud infrastructure, the Cisco Cloud APIC first applies that change to the policy model. This policy model change then triggers a change to the actual managed item. This approach is called a model-driven framework.

Policy Model Key Characteristics

Key characteristics of the policy model include the following:

- As a model-driven architecture, the software maintains a complete representation of the administrative and operational state of the system (the model). The model applies uniformly to cloud infrastructure, services, system behaviors, and virtual devices attached to the network.
- The logical and concrete domains are separated; the logical configurations are rendered into concrete configurations by applying the policies in relation to the available resources. No configuration is carried out against concrete entities. Concrete entities are configured implicitly as a side effect of the changes to the Cisco Cloud policy model.

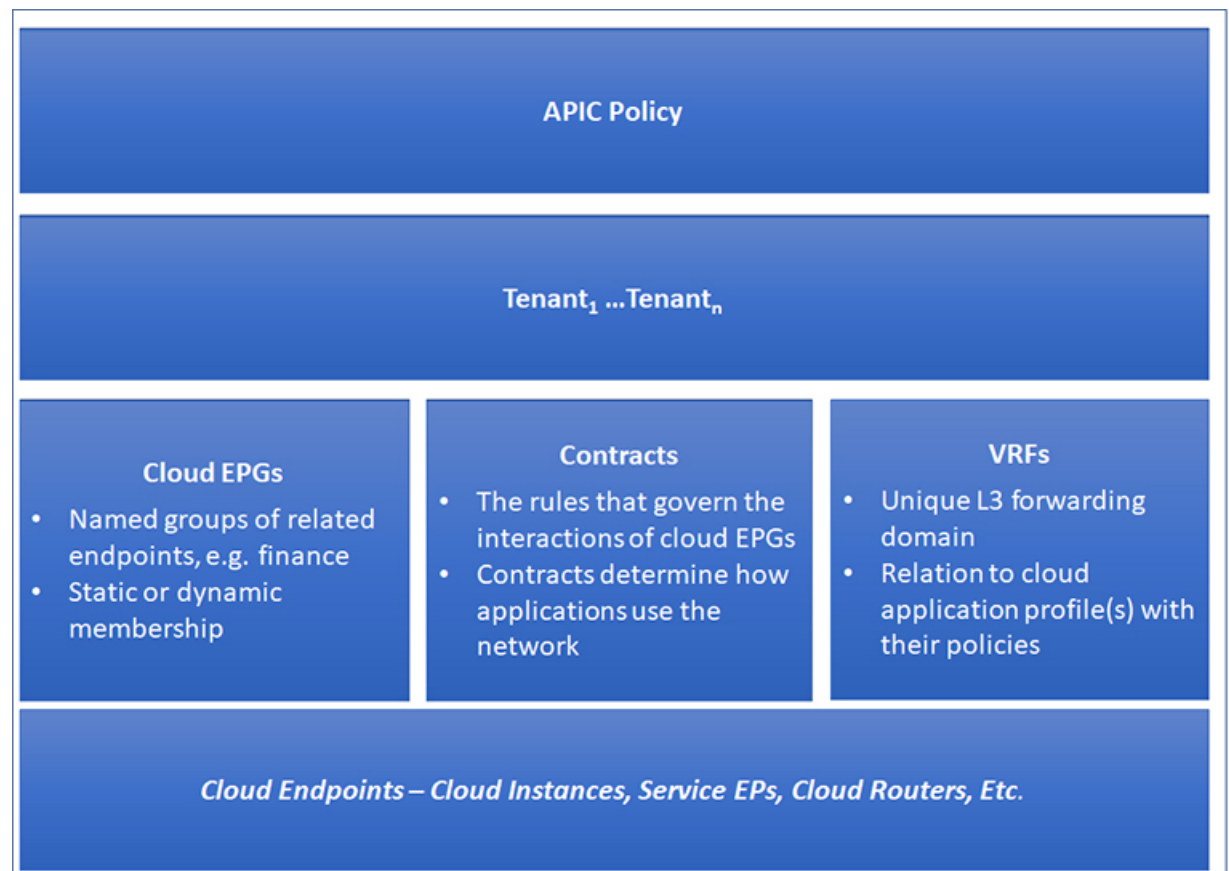
- The system prohibits communications with newly connected endpoints until the policy model is updated to include the new endpoint.
- Network administrators do not configure logical system resources directly. Instead, they define logical (hardware-independent) configurations and the Cisco Cloud APIC policies that control different aspects of the system behavior.

Managed object manipulation in the model relieves engineers from the task of administering isolated, individual component configurations. These characteristics enable automation and flexible workload provisioning that can locate any workload anywhere in the infrastructure. Network-attached services can be easily deployed, and the Cisco Cloud APIC provides an automation framework to manage the lifecycle of those network-attached services.

Logical Constructs

The policy model manages the entire cloud infrastructure, including the infrastructure, authentication, security, services, applications, cloud infrastructure, and diagnostics. Logical constructs in the policy model define how the cloud infrastructure meets the needs of any of the functions of the cloud infrastructure. The following figure provides an overview of the ACI policy model logical constructs.

Figure 1: ACI Policy Model Logical Constructs Overview




cloud infrastructure-wide or tenant administrators create predefined policies that contain application or shared resource requirements. These policies automate the provisioning of applications, network-attached services, security policies, and tenant subnets, which puts administrators in the position of approaching the resource pool in terms of applications rather than infrastructure building blocks. The application needs to drive the networking behavior, not the other way around.

The Cisco ACI Policy Management Information Model

The cloud infrastructure comprises the logical components as recorded in the Management Information Model (MIM), which can be represented in a hierarchical management information tree (MIT). The Cisco Cloud APIC runs processes that store and manage the information model. Similar to the OSI Common Management Information Protocol (CMIP) and other X.500 variants, the Cisco Cloud APIC enables the control of managed resources by presenting their manageable characteristics as object properties that can be inherited according to the location of the object within the hierarchical structure of the MIT.

Each node in the tree represents a managed object (MO) or group of objects. MOs are abstractions of cloud infrastructure resources. An MO can represent a concrete object, such as a cloud router, adapter, or a logical object, such as an application profile, cloud endpoint group, or fault. The following figure provides an overview of the MIT.

Figure 2: Cisco ACI Policy Management Information Model Overview



The hierarchical structure starts with the policy universe at the top (Root) and contains parent and child nodes. Each node in the tree is an MO and each object in the cloud infrastructure has a unique distinguished name (DN) that describes the object and locates its place in the tree.

The following managed objects contain the policies that govern the operation of the system:

- A tenant is a container for policies that enable an administrator to exercise role-based access control. The system provides the following four kinds of tenants:
 - The administrator defines user tenants according to the needs of users. They contain policies that govern the operation of resources such as applications, databases, web servers, network-attached storage, virtual machines, and so on.
 - Although the system provides the common tenant, it can be configured by the cloud infrastructure administrator. It contains policies that govern the operation of resources accessible to all tenants, such as firewalls, load balancers, Layer 4 to Layer 7 services, intrusion detection appliances, and so on.



Note As of the Cisco Application Policy Infrastructure Controller (APIC) Release 4.1(1), the Cisco Cloud APIC only supports load balancers as a Layer 4 to Layer 7 service.


- The infrastructure tenant is provided by the system but can be configured by the cloud infrastructure administrator. It contains policies that govern the operation of infrastructure resources. It also enables a cloud infrastructure provider to selectively deploy resources to one or more user tenants. Infrastructure tenant policies are configurable by the cloud infrastructure administrator.
- The cloud infra policies enable you to manage on-premises and inter-region connectivity when setting up the Cisco Cloud APIC. For more information, see the *Cisco Cloud APIC Installation Guide*.
- Cloud inventory is a service that enables you to view different aspects of the system using the GUI. For example, you can view the regions that are deployed from the aspect of an application or the applications that are deployed from the aspect of a region. You can use this information for cloud resource planning and troubleshooting.
- Layer 4 to Layer 7 service integration lifecycle automation framework enables the system to dynamically respond when a service comes online or goes offline. For more information, see [Deploying Layer 4 to Layer 7 Services](#)
- Access, authentication, and accounting (AAA) policies govern user privileges, roles, and security domains of the Cisco Cloud ACI cloud infrastructure. For more information, see [Cisco Cloud APIC Security](#)

The hierarchical policy model fits well with the REST API interface. When invoked, the API reads from or writes to objects in the MIT. URLs map directly into distinguished names that identify objects in the MIT. Any data in the MIT can be described as a self-contained structured tree text document encoded in XML or JSON.

Tenants

A tenant (`fvTenant`) is a logical container for application policies that enable an administrator to exercise domain-based access control. A tenant represents a unit of isolation from a policy perspective, but it does not represent a private network. Tenants can represent a customer in a service provider setting, an organization or domain in an enterprise setting, or just a convenient grouping of policies. The following figure provides an overview of the tenant portion of the management information tree (MIT).

Figure 3: Tenants



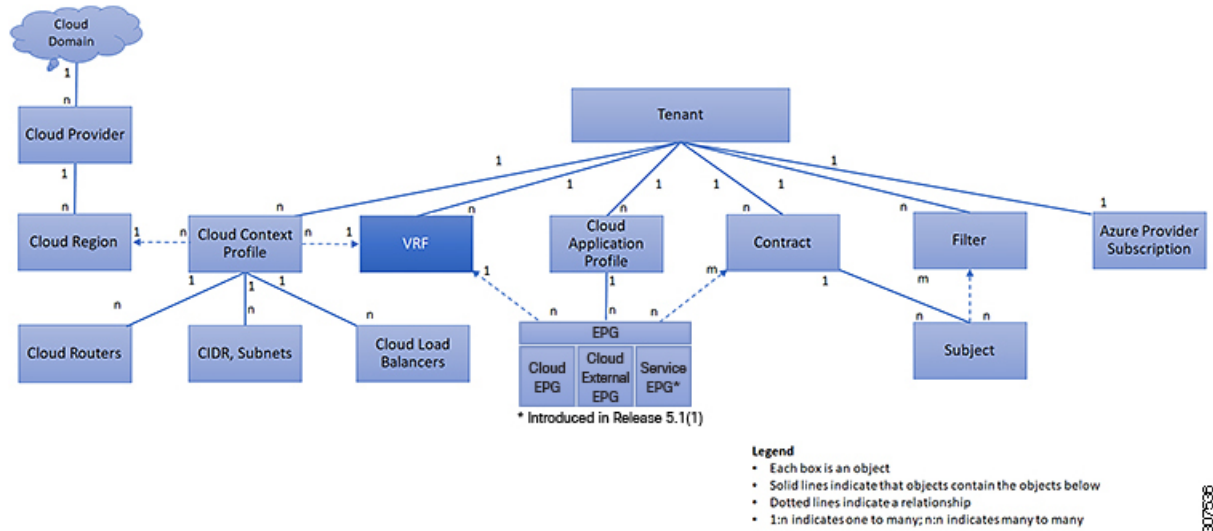
Tenants can be isolated from one another or can share resources. The primary elements that the tenant contains are filters, contracts, Virtual Routing and Forwarding (VRF) instances, cloud context profiles, Azure provider configurations, and cloud application profiles that contain cloud endpoint groups (cloud EPGs). Entities in the tenant inherit its policies. VRFs are also known as contexts; each VRF can be associated with multiple cloud context profiles. A cloud context profile, in conjunction with a VRF, tenant and region, represents a resource group in Azure. A VNET is created inside the resource group based on the VRF name.

Tenants are logical containers for application policies. The cloud infrastructure can contain multiple tenants. You must configure a tenant before you can deploy any Layer 4 to Layer 7 services. The ACI cloud infrastructure supports IPv4 and dual-stack configurations for tenant networking.

VRFs

A Virtual Routing and Forwarding (VRF) object (`fVctx`) or context is a tenant network (called a VRF in the Cisco Cloud APIC GUI). A tenant can have multiple VRFs. A VRF is a unique Layer 3 forwarding and application policy domain. The following figure shows the location of VRFs in the management information tree (MIT) and their relation to other objects in the tenant.

Figure 4: VRFs



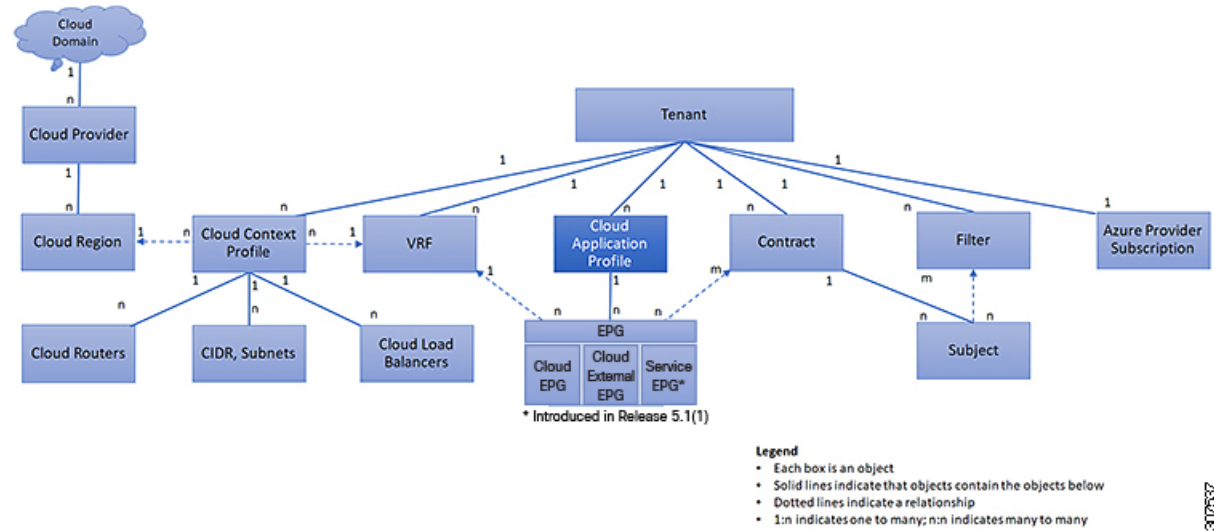
A VRF defines a Layer 3 address domain. One or more cloud context profiles are associated with a VRF. You can only associate one cloud context profile with a VRF in a given region. All the endpoints within the Layer 3 domain must have unique IP addresses because it is possible to forward packets directly between these devices if the policy allows it. A tenant can contain multiple VRFs. After an administrator creates a logical device, the administrator can create a VRF for the logical device, which provides a selection criteria policy for a device cluster. A logical device can be selected based on a contract name, a graph name, or the function node name inside the graph.

Beginning with Release 5.0(2), you can have a hub VNet (a cloudCtxProfile in the infra tenant) that can be carved out into multiple VRFs.

Cloud Application Profiles

A cloud application profile (`cloudAp`) defines the policies, services and relationships between cloud EPGs. The following figure shows the location of cloud application profiles in the management information tree (MIT) and their relation to other objects in the tenant.

Figure 5: Cloud Application Profiles



Cloud application profiles contain one or more cloud EPGs. Modern applications contain multiple components. For example, an e-commerce application could require a web server, a database server, data located in a storage service, and access to outside resources that enable financial transactions. The cloud application profile contains as many (or as few) cloud EPGs as necessary that are logically related to providing the capabilities of an application.

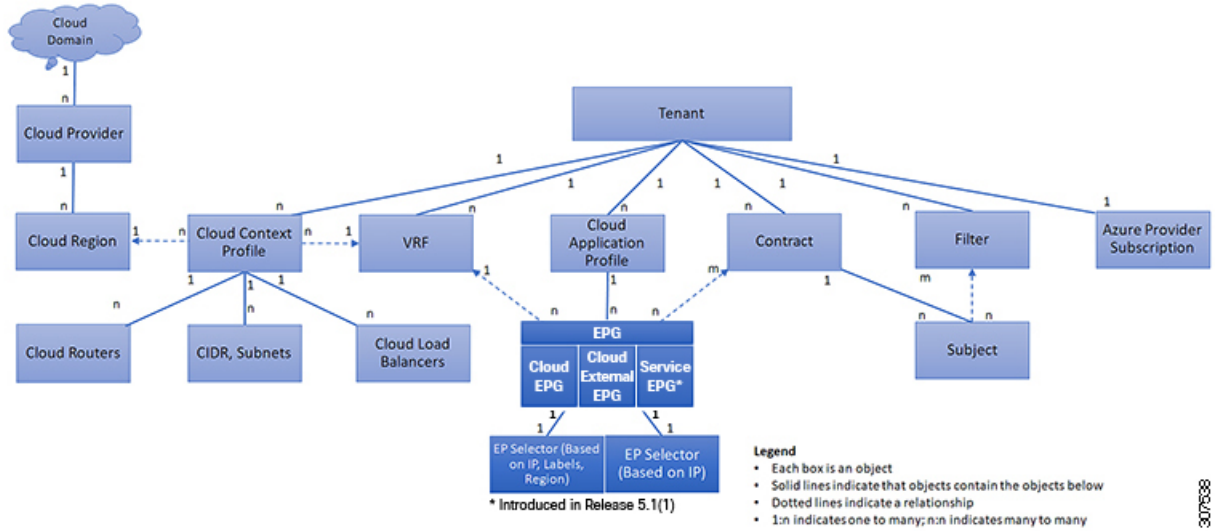
Cloud EPGs can be organized according to one of the following:

- The application they provide, such as a DNS server or SAP application (see *Tenant Policy Example* in *Cisco APIC REST API Configuration Guide*).
- The function they provide (such as infrastructure)
- Where they are in the structure of the data center (such as DMZ)
- Whatever organizing principle that a cloud infrastructure or tenant administrator chooses to use

Cloud Endpoint Groups

The cloud endpoint group (cloud EPG) is the most important object in the policy model. The following figure shows where application cloud EPGs are located in the management information tree (MIT) and their relation to other objects in the tenant.

Figure 6: Cloud Endpoint Groups



A cloud EPG is a managed object that is a named logical entity that contains a collection of endpoints. Endpoints are devices that are connected to the network. They have an address (identity), a location, attributes (such as version or patch level), and are virtual. Knowing the address of an endpoint also enables access to all its other identity details. Cloud EPGs are fully decoupled from the physical and logical topology. Endpoint examples include servers, virtual machines, storage services, or clients on the Internet. Endpoint membership in a cloud EPG can be dynamic or static.

The ACI cloud infrastructure can contain the following types of cloud EPGs:

- Cloud endpoint group (`cloudEPg`)
- Cloud external endpoint group (`cloudExtEPg`)

Cloud EPGs contain endpoints that have common policy requirements such as security or Layer 4 to Layer 7 services. Rather than configure and manage endpoints individually, they are placed in a cloud EPG and are managed as a group.

Policies apply to cloud EPGs, never to individual endpoints.

Regardless of how a cloud EPG is configured, cloud EPG policies are applied to the endpoints they contain.


WAN router connectivity to the cloud infrastructure is an example of a configuration that uses a static cloud EPG. To configure WAN router connectivity to the cloud infrastructure, an administrator configures a `cloudExtEPg` cloud EPG that includes any endpoints within an associated WAN subnet. The cloud infrastructure learns of the cloud EPG endpoints through a discovery process as the endpoints progress through their connectivity life cycle. Upon learning of the endpoint, the cloud infrastructure applies the `cloudExtEPg` cloud EPG policies accordingly. For example, when a WAN connected client initiates a TCP session with a server within an application (`cloudEPg`) cloud EPG, the `cloudExtEPg` cloud EPG applies its policies to that client endpoint before the communication with the (`cloudEPg`) cloud EPG web server begins. When the client server TCP session ends, and communication between the client and server terminates, the WAN endpoint no longer exists in the cloud infrastructure.

The Cisco Cloud APIC uses endpoint selectors to assign endpoints to Cloud EPGs. The endpoint selector is essentially a set of rules that are run against the cloud instances that are assigned to the Azure VNET managed by Cisco ACI. Any endpoint selector rules that match endpoint instances assign that endpoint to the Cloud EPG. The endpoint selector is similar to the attribute-based microsegmentation available in Cisco ACI.

Contracts

In addition to cloud EPGs, contracts (vzBrCP) are key objects in the policy model. Cloud EPGs can only communicate with other cloud EPGs according to contract rules. The following figure shows the location of contracts in the management information tree (MIT) and their relation to other objects in the tenant.

Figure 7: Contracts



An administrator uses a contract to select one or more types of traffic that can pass between cloud EPGs, including the protocols and ports allowed. If there is no contract, inter-EPG communication is disabled by default. There is no contract required for intra-EPG communication; intra-EPG communication is always implicitly allowed.

Contracts govern the following types of cloud EPG communications:

- Between cloud EPGs (cloudEPg), both intra-tenant and inter-tenant



Note In the case of a shared service mode, a contract is required for inter-tenant communication. A contract is used to specify static routes across VRFs, although the tenant VRF does not enforce a policy.

- Between cloud EPGs and cloud external EPGs (cloudExtEPg)

Contracts govern the communication between cloud EPGs that are labeled providers, consumers, or both. The relationship between a cloud EPG and a contract can be either a provider or consumer. When a cloud EPG provides a contract, communication with the cloud endpoints in that cloud EPG can be initiated from cloud endpoints in other cloud EPGs as long as the communication complies with the provided contract. When a cloud EPG consumes a contract, the cloud endpoints in the consuming cloud EPG may initiate communication with any cloud endpoint in a cloud EPG that is providing that contract.



Note A cloud EPG can both provide and consume the same contract. A cloud EPG can also provide and consume multiple contracts simultaneously.

Comma-separated Filters Support for Contract Rule Consolidation

After a contract is created, some of the rules defined in the contract are consolidated and displayed in Azure based on certain criteria. You can combine multiple ports and multiple IP addresses and ranges into a single, easy-to-understand rule. The criteria for consolidation of rules are:

- Rules are consolidated only within a contract. Two rules resulting from two different contracts are not consolidated in Azure.
- The source/ destination address prefixes and destination port(s) are consolidated.
- The conditions for multiple rules to get consolidated together in an NSG are:
 - Same contract
 - Same protocol (UDP, TCP, ICMP)
 - Same direction (inbound , outbound)
 - Same type (SG, IP)
- Overlapping port ranges for same protocol (TCP/UDP) in the same contract are consolidated to one range.

For example, TCP ports 100-200, 150-250 are consolidated to 100-250.
- If 1.2.3.4/32 (any address prefixes) is allowed, and an ext EPG with 0.0.0.0/0 is added, then the allowed Source/Destination IP would be *Any*, not [1.2.3.4/32, 0.0.0.0/0].

Example below shows the EPG1 outbound rules and the consolidated EPG1 outbound rules, based on contracts C1 and C2.

```
Contract C1:
Consumer: EPG1 , Provider: EPG2
Filter: TCP (ports 53)
Filter: UDP (port 53, 5000)
```

```
Contract C2:
Consumer: EPG1 , Provider: EPG2
Filter: TCP (ports 80, 8080)
```

```
EPG1 outbound rules:
EPG1 -> EPG2   TCP   80
EPG1 -> EPG2   TCP  8080
EPG1 -> EPG2   TCP           53
EPG1 -> EPG2   UDP   53
EPG1 -> EPG2   UDP  5000
EPG1 -> 1.1.1.1/32 TCP   80
EPG1 -> 1.1.1.1/32 TCP  8080
EPG1 -> 1.1.1.1/32 TCP   53
EPG1 -> 1.1.1.1/32 UDP   53
EPG1 -> 1.1.1.1/32 UDP  5000
```

```

EPG1 -> 2.2.2.2/32 TCP 80
EPG1 -> 2.2.2.2/32 TCP 8080
EPG1 -> 2.2.2.2/32 TCP 53
EPG1 -> 2.2.2.2/32 UDP 53
EPG1 -> 2.2.2.2/32 UDP 5000

```

Rules are consolidated by comma-separated filters (consolidated based on C1 and C2):

```

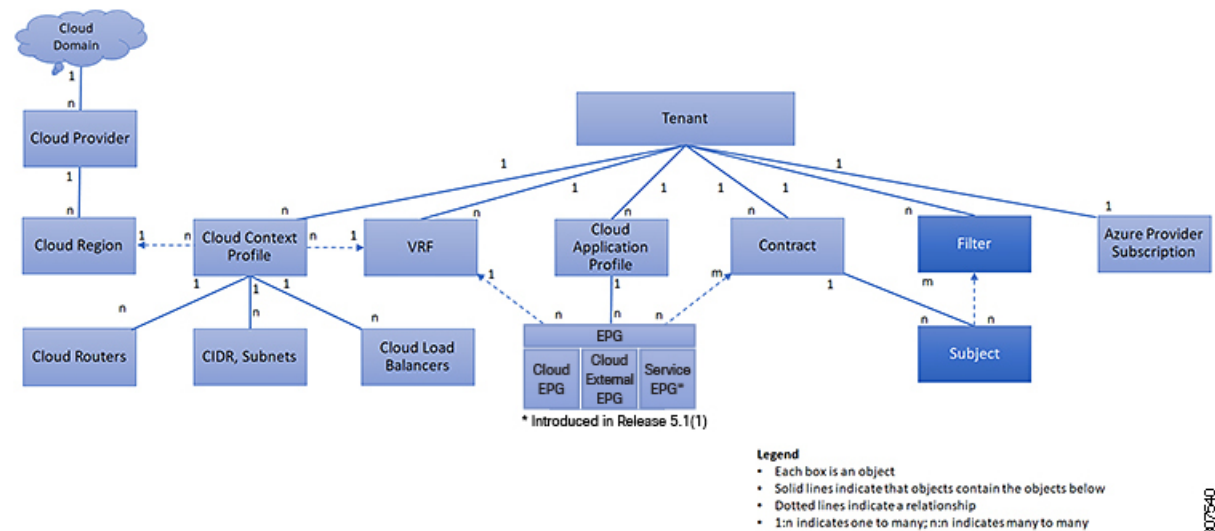
EPG1 -> EPG2 TCP 80,8080
EPG1 -> EPG2 UDP 53,5000
EPG1 -> EPG2 TCP 53
EPG1 -> 1.1.1.1/32, 2.2.2.2/32 TCP 80,8080
EPG1 -> 1.1.1.1/32, 2.2.2.2/32 UDP 53,5000
EPG1 -> 1.1.1.1/32, 2.2.2.2/32 TCP 53

```

Filters and Subjects Govern Cloud EPG Communications

Subject and filter managed-objects enable mixing and matching among cloud EPGs and contracts so as to satisfy various applications or service delivery requirements. The following figure shows the location of application subjects and filters in the management information tree (MIT) and their relation to other objects in the tenant.

Figure 8: Subjects and Filters



Contracts can contain multiple communication rules and multiple cloud EPGs can both consume and provide multiple contracts. A policy designer can compactly represent complex communication policies and re-use these policies across multiple instances of an application.



Note Subjects are hidden in Cisco Cloud APIC and not configurable. For rules installed in Azure, source port provided in the filter entry is not taken into account.

Subjects and filters define cloud EPG communications according to the following options:

- Filters are Layer 3 to Layer 4 fields, TCP/IP header fields such as Layer 3 protocol type, Layer 4 ports, and so forth. According to its related contract, a cloud EPG provider dictates the protocols and ports in

both the in and out directions. Contract subjects contain associations to the filters (and their directions) that are applied between cloud EPGs that produce and consume the contract.

- Subjects are contained in contracts. A subject within a contract uses filters to specify the type of traffic that can be communicated and how it occurs. For example, for HTTPS messages, the subject specifies the direction and the filters that specify the IP address type (for example, IPv4), the HTTP protocol, and the ports allowed. Subjects determine if filters are unidirectional or bidirectional. A unidirectional filter is used in one direction. Unidirectional filters define in or out communications but not the same for both. Bidirectional filters are the same for both; they define both in and out communications.
- ACI contracts rendered in Azure constructs are always stateful, allowing return traffic.

About the Cloud Template

The cloud template provides a template that configures and manages the Cisco Cloud APIC infra network. The template requires only the most essential elements for the configuration. From these elements, the cloud template generates a detailed configuration necessary for setting up the Cisco Cloud APIC infra network. However, it is not a one-time configuration generation—it is possible to add, modify, or remove elements of the template input. The cloud template updates the resulting configuration accordingly.

One of the central things in the Azure network configuration is the Virtual Private Cloud (VNET). Azure supports many regions worldwide and one VNET is specific to one region.

The cloud template accepts one or more region names and generates the entire configuration for the infra VNETs in those regions. They are the infra VNETs. The Cisco Cloud APIC-managed object (MO) corresponding to the Azure VNET is `cloudCtxProfile`. For every region specified in the cloud template, it generates the `cloudCtxProfile` configuration. A `cloudCtxProfile` is the topmost MO for all the configuration corresponding to a region. Underneath, it has many of other MOs organized as a tree to capture a specific configuration. The `cloudCtxProfile` MO for the infra VNet is generated by the cloud template. It carries `ctxProfileOwner == SYSTEM`, which means that this MO is generated by the system. For the non-infra network, it is possible to configure `cloudCtxProfile` directly; in this case, `cloudCtxProfile` carries `ctxProfileOwner == USER`.

A primary property of an Azure VNet is the CIDR. In Cisco Cloud APIC, you can choose and deploy CIDRs in the user VNETs. The CIDRs for the infra VNet are provided by users to the cloud template during the initial setup of the cloud site, and are deployed to the Azure cloud by the cloud template.

Beginning with Release 5.0(2), a new property called `createdBy` is added for the CIDR. The default value for this `createdBy` property is `USER`.

- For all user-created CIDRs, the value for the `createdBy` property is set to `USER`.
- For cloud template-created CIDRs, the value for the `createdBy` property is set to `SYSTEM`.

In releases prior to Release 5.0(2), you are not allowed to add more CIDRs to the infra VNet. Beginning with Release 5.0(2), multiple CIDR and subnet blocks can now be configured on the infra VNet. You can create CIDRs and associate subnets in the infra VNet. The cloud template subnets will be mapped to the overlay-1 VRF, but the user-created subnets will be implicitly mapped to the overlay-2 VRF in the same infra VNet. All subnets in the respective VRFs will have separate route tables in the cloud for VRF segregation.

In addition, beginning with Release 5.0(2), you can create cloud EPGs and cloud external EPGs in the infra tenant, where all the cloud EPGs and cloud external EPGs will be associated with the overlay-2 VRF in the infra tenant. A cloud EPG in the overlay-2 VRF can communicate with other cloud EPGs and cloud external

EPGs in the overlay-2 VRF, and can also communicate with cloud EPGs in other user tenant VRFs. We recommend that you do not use existing "cloud-infra" application profiles, and instead create a new application profile in the infra tenant and associate that new application profile to the cloud EPGs and cloud external EPGs in the overlay-2 VRF.

For more information, see [Creating an EPG Using the Cisco Cloud APIC GUI](#) and [About the Overlay-1 and Overlay-2 VRFs](#).

The cloud template generates and manages a huge number of MOs in the `cloudCtxProfile` subtree including, but not limited to, the following:

- Subnets
- Cloud routers
- IP address allocation for the cloud router interfaces
- IP address allocation and configuration for tunnels
- IP address allocation and configuration for loopbacks

Without the cloud template, you would be responsible for configuring and managing these.

The *Cisco Cloud Template MO* table contains a brief summary of the inputs (MOs) to the cloud template.

Table 1: Cloud Template MOs

MO	Purpose
<code>cloudtemplateInfraNetwork</code>	The root of the cloud template configuration. Attributes include: <code>numRoutersPerRegion</code> —The number of cloud routers for each <code>cloudRegionName</code> specified under <code>cloudtemplateIntNetwork</code> .
<code>cloudtemplateProfile</code>	Configuration profile for all the cloud routers. Attributes include: <ul style="list-style-type: none"> • <code>routerUsername</code> <p>Note</p> <ul style="list-style-type: none"> • The username cannot be "admin." • Any username restrictions from Azure applies. • <code>routerPassword</code> • <code>routerThroughput</code> • <code>routerLicenseToken</code>
<code>cloudtemplateIntNetwork</code>	Contains a list of regions, which specify where you deploy the cloud routers. Each region is captured through a <code>cloudRegionName</code> child MO


MO	Purpose
cloudtemplateExtNetwork	Contains infra network configuration input that is external of the cloud. Contains a list of regions where cloud routers are configured for external networking. Each region is captured through a <code>cloudRegionName</code> child MO
cloudtemplateVpnNetwork	Contains information for setting up a VPN with an ACI on-premises site or another Cisco Cloud APIC site.
cloudtemplateIpSecTunnel	Captures the IP address of the IPsec peer in the ACI on-premises site.
cloudtemplateOspf	Captures the OSPF area to be used for the VPN connections.
cloudtemplateBgpEvpn	Captures the peer IP address, ASN, and so forth, for setting up the BGP session with the on-premises site.

In Cisco Cloud APIC, the layering of MOs is slightly different from a regular Cisco APIC due to the cloud template. In a regular Cisco APIC, you post logical MOs that go through two layers of translation:

1. Logical MO to resolved MO
2. Resolved MO to concrete MO

In Cisco Cloud APIC, there is an additional layer of translation for the infra network. This additional layer is where the cloud template translates logical MOs in the `cloudtemplate` namespace to logical MOs in the `cloud` namespace. For configurations outside of the infra network, you post logical MOs in the `cloud` namespace. In this case, the MOs go through the usual two-layer translation as in the regular Cisco APIC.

Figure 9: Cloud and Cloud Template MO Conversion





Note For information about configuring the cloud template, see [Configuring Cisco Cloud APIC Components](#)

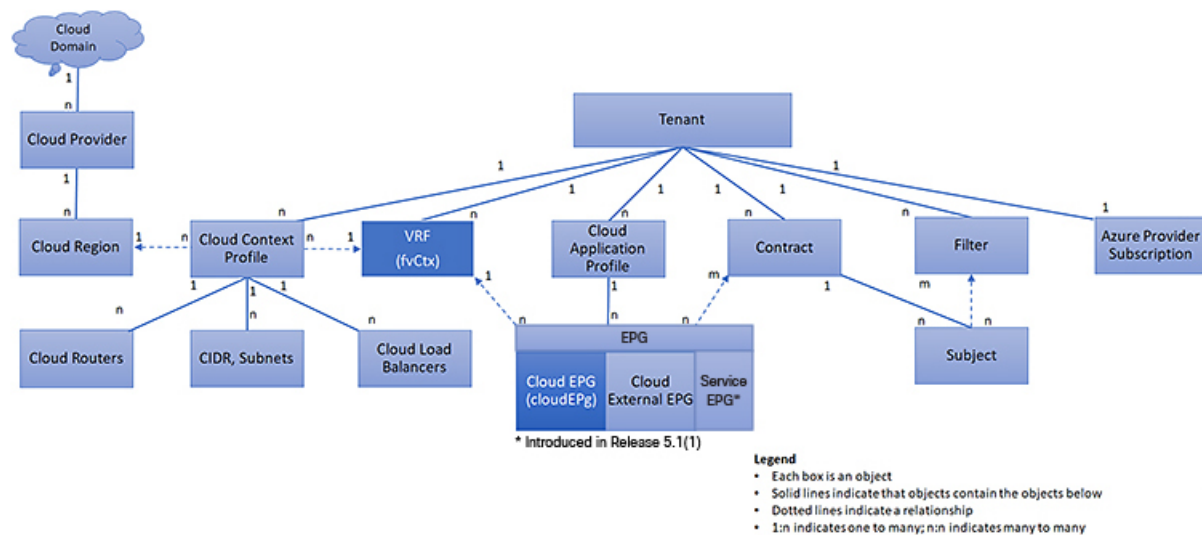
Managed Object Relations and Policy Resolution

Relationship-managed objects express the relation between managed object instances that do not share containment (parent-child) relations. MO relations are established between the source MO and a target MO in one of the following two ways:

- An explicit relation, such as with `cloudRsCloudEPgCtx`, defines a relationship that is based on the target MO distinguished name (DN).
- A named relation defines a relationship that is based on the target MO name.

The dotted lines in the following figure show several common MO relations.

Figure 10: MO Relations



For example, the dotted line between the cloud EPG and the VRF defines the relation between those two MOs. In this figure, the cloud EPG (`cloudEPg`) contains a relationship MO (`cloudRsCloudEPgCtx`) that is named with the name of the target VRF MO (`fvCtx`). For example, if production is the VRF name (`fvCtx.name=production`), then the relation name is production (`cloudRsCloudEPgCtx.tnFvCtxName=production`).

In the case of policy resolution based on named relations, if a target MO with a matching name is not found in the current tenant, the ACI cloud infrastructure tries to resolve in the common tenant. For example, if the user tenant cloud EPG contained a relationship MO targeted to a VRF that did not exist in the tenant, the system tries to resolve the relationship in the common tenant. If a named relation cannot be resolved in either the current tenant or the common tenant, the ACI cloud infrastructure attempts to resolve to a default policy. If a default policy exists in the current tenant, it is used. If it does not exist, the ACI cloud infrastructure looks for a default policy in the common tenant. Cloud context profile, VRF, and contract (security policy) named relations do not resolve to a default.

Default Policies

**Warning**

Default policies can be modified or deleted. Deleting a default policy can result in a policy resolution process to complete abnormally.

The ACI cloud infrastructure includes default policies for many of its core functions. Examples of default policies include the following:

- Cloud Azure provider (for the infra tenant)
- Monitoring and statistics

**Note**

To avoid confusion when implementing configurations that use default policies, document changes made to default policies. Be sure that there are no current or future configurations that rely on a default policy before deleting a default policy. For example, deleting a default firmware update policy could result in a problematic future firmware update.

A default policy serves multiple purposes:

- Allows a cloud infrastructure administrator to override the default values in the model.
- If an administrator does not provide an explicit policy, the Cisco CloudAPIC applies the default policy. An administrator can create a default policy and the Cisco Cloud APIC uses that unless the administrator provides any explicit policy.

The following scenarios describe common policy resolution behavior:

- A configuration explicitly refers to the default policy: if a default policy exists in the current tenant, it is used. Otherwise, the default policy in tenant **common** is used.
- A configuration refers to a named policy (not default) that does not exist in the current tenant or in tenant **common**: if the current tenant has a default policy, it is used. Otherwise, the default policy in tenant **common** is used.



Note The scenario above does not apply to a VRF in a tenant.

- A configuration does not refer to any policy name: if a default policy exists in the current tenant, it is used. Otherwise, the default policy in tenant **common** is used.

The policy model specifies that an object is using another policy by having a relation-managed object (MO) under that object and that relation MO refers to the target policy by name. If this relation does not explicitly refer to a policy by name, then the system tries to resolve a policy that is called default. Cloud context profiles and VRFs are exceptions to this rule.

Shared Services

Cloud EPGs in one tenant can communicate with cloud EPGs in another tenant through a contract interface that is contained in a shared tenant. Within the same tenant, a cloud EPG in one VRF can communicate with another cloud EPG in another VRF through a contract defined in the tenant. The contract interface is an MO that can be used as a contract consumption interface by the cloud EPGs that are contained in different tenants. By associating to an interface, a cloud EPG consumes the subjects that are represented by the interface to a contract contained in the shared tenant. Tenants can participate in a single contract, which is defined at some third place. More strict security requirements can be satisfied by defining the tenants, contract, subjects, and filter directions so that tenants remain isolated from one another.

Follow these guidelines when configuring shared services contracts:

- A shared service is supported only with non-overlapping and non-duplicate CIDR subnets. When configuring CIDR subnets for shared services, follow these guidelines:
 - CIDR subnets leaked from one VRF to another must be disjointed and must not overlap.
 - CIDR subnets advertised from multiple consumer networks into a VRF or vice versa must be disjointed and must not overlap.
 - Inter-tenant contracts require a global scope.

