Deploying Layer 4 to Layer 7 Services

Overview

The Cisco Cloud APIC enables you to deploy Layer 4 to Layer 7 service devices to the public cloud. This initial release supports application load balancer (ALB) deployments in Amazon Web Services (AWS).

About Application Load Balancers

An application load balancer (ALB) is a Layer 7 load balancer that inspects packets and creates access points to HTTP and HTTPS headers. It also identifies the load and spreads it out to the targets with higher efficiency. You deploy an ALB using a service graph, which enables you to define how you want traffic to come into the network, the devices that the traffic passes through, and how the traffic leaves the network. You specify these actions by configuring one or more listeners.

Listeners enable you to specify the ports and protocols (HTTP or HTTPS) that the ALB accepts traffic on. When specifying HTTPS, you also choose a security policy and an SSL certificate.

Note A listener can have multiple certificates.

All listeners require you to configure at least one rule (a default rule, which does not have a condition). Rules enable you to specify the action that the load balancer takes when a condition is met. For example, you can create a rule that redirects traffic to a specified URL when a request is made to a specified hostname or path.

There are two deployment types: internet-facing and internal-facing. An internet-facing deployment inserts the ALB as a service between the consumer external EPG and the provider cloud EPG. The following figure shows the contract configuration within the VRF and the ALB as a service inserted between the consumer external EPG and the provider cloud EPG.

Figure 1: Internet-Facing Deployment

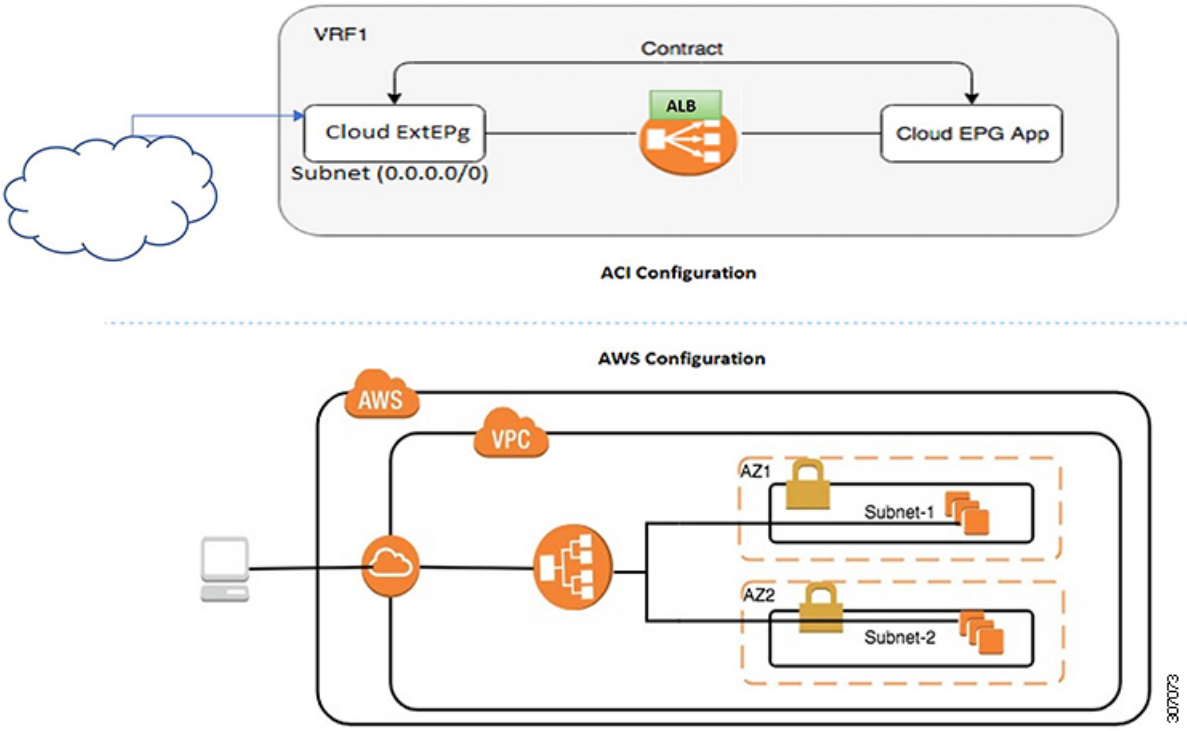

An internal-facing deployment inserts the ALB as a service between the consumer cloud EPG and the provider cloud EPG. The following figure shows the contract configuration within the VRF and the ALB as a service inserted between the consumer cloud EPG and provider cloud EPG.

Figure 2: Internal-Facing Deployment

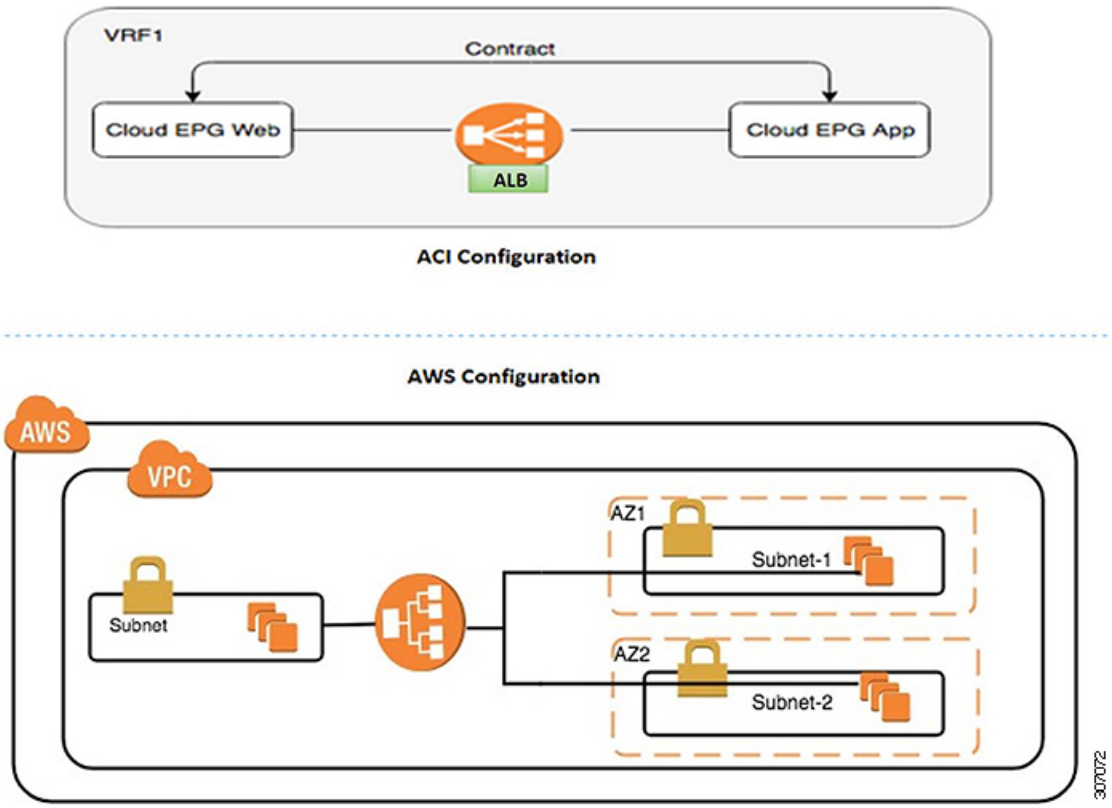

Note You can find more information about ALBs in the documentation on the AWS website.

Dynamic Server Attachment to Server Pool

Servers in the server pool or target group are dynamically added. You do not need to specify the IP addresses or instance Ids for the targets. The relation from a listener rule to a provider cloud EPG is used for the dynamic selection of endpoints. The relation is also used for adding the endpoints to the target group. By default, the endpoints are registered with the port number 80.

Based on the target group-to-security group association that is provided in the ALB, and the EPG (security group) of the endpoint, the EC2 instance (server) is associated to the target group dynamically on the target group's default port. Alternatively, instead of registering the EC2 instance on the target group port, you can attach the custom port by specifying the ports in the following table:

Table 1: Custom Port-Based Attachment

Provider EPG	Ports
EPGMap:<Epg1DN>	9090

Provider EPG	Ports
EPGMap:<Epg2DN>	9091, 9099

You can specify EPGMap:<EpgDN> as the tag and the list of ports to be registered on the target group as a list separated by commas.

About Service Graphs

The Cisco Application Centric Infrastructure (ACI) treats services as a part of an application. Any services that are required are treated as a service graph that is instantiated on the Cisco ACI fabric from the Cisco APIC. You define the service for the application while service graphs identify the set of network or service functions that the application needs.

A service graph represents the network using the following elements:

- Function node—A function node represents a function that is applied to the traffic, such as a load balancer. A function within the service graph might require one or more parameters and have one or more connectors.
- Terminal node—A terminal node enables input and output from the service graph.
- Connector—A connector enables input and output from a node.

After the graph is configured, the Cisco APIC automatically configures the services according to the service function requirements that are specified in the service graph. The Cisco APIC also automatically configures the network according to the needs of the service function that is specified in the service graph, which does not require any change in the service device.

A service graph is represented as two or more tiers of an application with the appropriate service function inserted between them.

A service appliance (device) performs a service function within the graph. One or more service appliances might be required to render the services required by a graph. A single-service device can perform one or more service functions.

Service graphs and service functions have the following characteristics:

- Traffic sent from specific endpoint groups can be redirected based on a policy.
- Service graph redirection is directional. In other words, redirection can be applied to both traffic directions or either one of them.
- Logical functions can be rendered on the appropriate device, based on the policy.
- The service graph supports splits and joins of edges, and it does not restrict the administrator to linear service chains.
- Traffic can be reclassified again in the network after a service appliance emits it.

By using a service graph, you can install a service, a load balancer, once and deploy it multiple times in different logical topologies. Each time the graph is deployed, Cisco ACI takes care of changing the configuration on the service device to enable the forwarding in the new logical topology.

About Function Nodes

A function node represents a single service function. A function node has function node connectors, which represent the network requirement of a service function.

A function node within a service graph requires the following parameters:

- A tenant
- A cloud context profile with subnets in two availability zones

Function parameters can be specified when the service graph is rendered. For example, if the function node is a load balancer, the listener and its rule can be specified for the function node at the time the graph is rendered.

About Terminal Nodes

Terminal nodes connect a service graph with the contracts. You can insert a service graph for the traffic between two application cloud EPGs by connecting the terminal node to a contract. Once connected, traffic between the consumer cloud EPG and provider cloud EPG of the contract is redirected to the service graph.

Deploying a Service Graph

The service graph enables you to define how traffic flows between devices, how the traffic comes into the network, which devices the traffic passes through, and how the traffic leaves the network.

Before you can configure a service graph, you must first configure the following:

1. A tenant
2. A cloud context profile
3. Subnets
4. An application profile
5. A consumer EPG
6. A provider EPG
7. A contract

Deploying the Service Graph Using the Cloud APIC GUI

Creating a Load Balancer Using the Cisco Cloud APIC GUI

This section explains how to create a load balancer using the Cisco Cloud APIC GUI.

Step 1 Click **Application Management** > **Services**.

The **Services** page appears.

Step 2 Click the Devices tab, then click **Actions** > **Create Device**.

The **Create Device** page appears.

Step 3 Enter the appropriate values in each field as listed in the following *Create Device Dialog Box Fields* table then continue.

Table 2: Create Device Dialog Box Fields

Properties	Description
General	
Name	Enter the name of the load balancer.
Tenant	To choose a tenant: **a.** Click **Select Tenant**. The **Select Tenant** dialog appears. **b.** From the column on the left, click to choose a tenant. **c.** Click **Select**. You return to the **Create Device** dialog box.
Settings	
Service Type	Choose **Application Load Balancer**.
Scheme	Choose **Internal** or **Internet Facing**.
Add Availability Zone	You can specify only one subnet per availabilty zone. You must specifiy subnets from at least two availability zones to increase the availability of your load balancer. To choose an availability zone: **a.** Click **Add Availability Zone**. The **Add Availability Zone** dialog box appears. **b.** Click **Select Availability Zone**. The **Select Availability Zone** dialog box appears. **c.** From the column on the left, click to choose an availability zone. **d.** Click **Select**. You return to the **Add Availability Zone** dialog box.

Properties	Description
Subnet	For Cisco Cloud APIC deployed in AWS, two subnets are required (one subnet per availability zone). To choose a subnet: **a.** From the **Add Availability Zone** dialog box, click **Select Subnet**. The **Select Subnet** dialog box appears. **b.** From the column on the left, click to choose a subnet. **c.** Click **Select**. You return to the **Add Availability Zone** dialog box. **d.** Click **Add** to add the availability zone and subnet.

Step 4 Click **Save** when finished.

Creating a Service Graph Template Using the Cisco Cloud APIC GUI

This section explains how to configure a service graph template using the Cisco Cloud APIC GUI.

Before you begin

You have already created a device.

Step 1 Click **Application Management** > **Services**.

The **Services** page appears.

Step 2 Click the **Service Graphs** tab, then click **Actions** > **Create Service Graph**.

The **Create Service Graph** page appears.

Step 3 Enter the appropriate values in each field as listed in the following *Create Service Graph Dialog Box Fields* table then continue.

Table 3: Create Service Graph Dialog Box Fields

Properties	Description
General	
Name	Enter the name of service graph template.
Tenant	To choose a tenant: **a.** Click **Select Tenant**. The **Select Tenant** dialog appears. **b.** From the column on the left, click to choose a tenant. **c.** Click **Select**. You return to the **Create Service Graph** dialog box.

Properties	Description
Description	Enter a description of the service graph template.
Settings	
Select a Device	To choose a device: **a.** Drag and drop the Application Load Balancer icon to the **Drop Device** area in the service graph. The **Service Node** dialog box appears. **b.** Click **Select Application Load Balancer**. The **Select Application Load Balancer** dialog appears. **c.** From the column on the left, click to choose a device. **d.** Click **Select**. You return to the **Service Node** dialog box. **e.** Click **Add**. You return to the **Create Service Graph** window.

Step 4 Click **Save** when finished.

Deploying Layer 4 to Layer 7 Services Using the Cisco Cloud APIC GUI

This section explains how to deploy Layer 4 to Layer 7 services.

Before you begin

- You have configured a device.
- You have configured a service graph.

Step 1 Click the **Intent** icon. The **Intent** menu appears.

Step 2 Click the drop-down arrow below the **Intent** search box and choose **Configuration**.

A list of **Configuration** options appear in the **Intent** menu.

Step 3 From the **Configuration** list in the **Intent** menu, click **EPG Communication**. The **EPG Communication** dialog box appears with the **Consumer EPGs**, **Contract**, and **Provider EPGs** information.

Step 4 To choose a contract:

a) Click **Select Contract**. The **Select Contract** dialog appears.
b) In the pane on the left side of the **Select Contract** dialog, click to choose a contract then click **Select**. The **Select Contract** dialog box closes.

Step 5 To add a consumer EPG:

a) Click **Add Consumer EPGs**. The **Select Consumer EPGs** dialog appears.

b) In the pane on the left side of the **Select Consumer EPGs** dialog, click to place a check in a check box to choose a cloud EPG (for an internal facing load balancer) or a cloud external EPG (for an internet facing load balancer) then click **Select**. The **Select Consumer EPGs** dialog box closes.

Step 6 To add a provider EPG:

a) Click **Add Provider EPGs**. The **Select Provider EPGs** dialog appears.

b) In the pane on the left side of the **Select Provider EPGs** dialog, click to place a check in a check box to choose a provider EPG then click **Select**. The **Select Provider EPGs** dialog box closes.

Step 7 To choose a service graph:

a) From the **EPG Communication Configuration** dialog, click **Select Service Graph**. The **Select Service Graph** dialog box appears.

b) In the pane on the left side of the **Select Service Graph** dialog, click to choose a service graph then click **Select**. The **Select Service Graph** dialog box closes.

Step 8 Under **Service Graph Preview**, click **Add Cloud Load Balancer Listener**. The **Add Cloud Load Balancer Listener** dialog appears that enables you to add listeners.

Listeners are the ports and protocols that the device will work on.

Step 9 Enter the appropriate values in each field as listed in the following *Add Cloud Load Balancer Listener Dialog Box Fields* table then continue.

Table 4: Add Cloud Load Balancer Listener Dialog Box Fields

Properties	Description
Name	Enter the name of the listener.
Port	Enter the port that the device will accept traffic on.
Protocol	Click to choose **HTTP** or **HTTPS**.
Security Policy	Click the drop-down list and choose a security policy (only available when **HTTPS** is chosen).

Properties	Description
SSL Certificate	To choose an SSL certificate(only available when **HTTPS** is chosen): **a.** Click **Add SSL Certificates**. **b.** Click to place a check mark in the check box of the certificates you want to add. **c.** Choose a key ring: **1.** Click **Select Key Ring**. The **Select Key Ring** dialog appears. **2.** From the **Select Key Ring** dialog, click to choose a key ring in the left column then click **Select**. The **Select Key Ring** dialog box closes. **d.** Click the **Certificate Store** drop-down list and choose a certificate. **Note** A listener can have multiple certificates.
Add Rule	To add rule settings to the device listener, click **Add Rule**. A new row appears in the **Rules** list an the **Rules Settings** fields are enabled.

Properties	Description
Rule Settings	

Properties	Description
	The **Rule Settings** pane contains the following options: • **Name**—Enter a name for the rule. • **Host**—Enter a hostname to create a host-based condition. When a request is made for this hostname, the action you specify is taken. • **Path**—Enter a path to create a path-based condition. When a request is made for this path, the action you specify is taken. • **Type**—The action type tells the device which action to take. The action type options: • **Return fixed response**—Returns a response using the following options: • **Fixed Response Body**—Enter a response message. • **Fixed Response Code**—Enter a response code. • **Fixed response Content-Type**—Choose a content type. • **Forward**—Forwards traffic using the following options: • **Port**—Enter the port that the device will accept traffic on. • **Protocol**—Click to choose **HTTP** or **HTTPS**. • **Provider EPG**—The EPG with the web server that handles the traffic. • **EPG**—To choose an EPG: **a.** Click **Select EPG**. The **Select EPG** dialog box appears. **b.** From the **Select EPG** dialog ox, click to choose an EPG in the left column then click **Select**. The **Select EPG** dialog box closes. • **Redirect**—Redirects requests to another location using the following options: • **Redirect Code**—Click the **Redirect Code** drop-down list and choose a code.

Properties	Description
	• **Redirect Hostname**—Enter a hostname for the redirect. • **Redirect Path**—Enter a redirect path. • **Redirect Port**—Enter the port that the device will accept traffic on. • **Redirect Protocol**—Click to the **Redirect Protocol** drop-down list and choose **HTTP**, **HTTPS**, or **Inherit**. • **Redirect Query**—Enter a redirect query. Click **Add Rule** when finished.

Step 10 Click **Add** when finished.
The service graph is deployed.

Deploying a Service Graph Using the REST API

Creating an Internal-Facing Load Balancer Using the REST API

This example demonstrates how to create an internal-facing load balancer using the REST API.

To create an internal-facing load balancer:

Example:

```
<polUni>
  <fvTenant name="t2" status="">
    <cloudLB name="ALB1" type="application" scheme="internal" status="">
     <cloudRsLDevToCloudSubnet tDn="uni/tn-t2/ctxprofile-c1/cidr-[10.33.0.0/16]/subnet-[10.33.7.0/24]"
 status=""/>
     <cloudRsLDevToCloudSubnet tDn="uni/tn-t2/ctxprofile-c1/cidr-[10.33.0.0/16]/subnet-[10.33.8.0/24]"
 status=""/>
    </cloudLB>
  </fvTenant>
</polUni>
```

Configuring an Internet-Facing Load Balancer Using the REST API

This example demonstrates how to create an internet-facing load balancer using the REST API.

To create an internet-facing load balancer:

Example:

```
<polUni>
  <fvTenant name="t2" status="">
    <cloudLB name="ALB1" type="application" scheme="internet" status="">
     <cloudRsLDevToCloudSubnet tDn="uni/tn-t2/ctxprofile-c1/cidr-[10.33.0.0/16]/subnet-[10.33.5.0/24]"
 status=""/>
     <cloudRsLDevToCloudSubnet tDn="uni/tn-t2/ctxprofile-c1/cidr-[10.33.0.0/16]/subnet-[10.33.6.0/24]"
 status=""/>
    </cloudLB>
  </fvTenant>
</polUni>
```

Creating a Service Graph Using the REST API

This example demonstrates how to create a service graph using the REST API.

To create a service graph:

```
<polUni>
  <fvTenant name="t2">
    <vnsAbsGraph name="CloudGraph" type="cloud" status="">
      <vnsAbsTermNodeProv name="Input1">
        <vnsAbsTermConn name="C1"/>
      </vnsAbsTermNodeProv>
      <vnsAbsTermNodeCon name="Output1">
        <vnsAbsTermConn name="C2"/>
      </vnsAbsTermNodeCon>
      <vnsAbsNode funcType="GoTo" name="N1" managed="yes">
        <vnsRsNodeToCloudLDev tDn="uni/tn-t2/clb-ALB1" status=""/>
        <vnsAbsFuncConn name="provider"/>
        <vnsAbsFuncConn name="consumer"/>
      </vnsAbsNode>
      <vnsAbsConnection connDir="consumer" connType="external" name="CON2">
       <vnsRsAbsConnectionConns tDn="uni/tn-t2/AbsGraph-CloudGraph/AbsTermNodeCon-Output1/AbsTConn"/>

        <vnsRsAbsConnectionConns tDn="uni/tn-t2/AbsGraph-CloudGraph/AbsNode-N1/AbsFConn-consumer"/>
      </vnsAbsConnection>
      <vnsAbsConnection connDir="provider" connType="internal" name="CON1">
       <vnsRsAbsConnectionConns tDn="uni/tn-t2/AbsGraph-CloudGraph/AbsTermNodeProv-Input1/AbsTConn"/>

        <vnsRsAbsConnectionConns tDn="uni/tn-t2/AbsGraph-CloudGraph/AbsNode-N1/AbsFConn-provider"/>
      </vnsAbsConnection>
    </vnsAbsGraph>
  </fvTenant>
</polUni>
```

Attaching a Service Graph Using the REST API

This example demonstrates how to attach a service graph using the REST API.

To attach a service graph:

```
<polUni>
  <fvTenant name="t2">
    <vzBrCP name="httpFamily">
      <vzSubj name="default" revFltPorts="yes" targetDscp="unspecified">
        <vzRsSubjGraphAtt tnVnsAbsGraphName="CloudGraph"/>
      </vzSubj>
    </vzBrCP>
  </fvTenant>
</polUni>
```

Configuring an HTTP Service Policy Using the REST API

This example demonstrates how to create an HTTP service policy using the REST API.

To create an HTTP service policy:

```
<polUni>
  <fvTenant name="t2">
    <vnsAbsGraph name="CloudGraph" type="cloud" status="">
      <vnsAbsNode funcType="GoTo" name="N1" managed="yes">
        <cloudSvcPolicy tenantName="t2" contractName="httpFamily" subjectName="consubj">
          <cloudListener name="http_listener1" port="80" protocol="http" status="">
            <cloudListenerRule name="rule1" priority="10" default="yes" status="">
              <cloudRuleAction type="forward" port="80" protocol="http"
epgdn="uni/tn-t2/cloudapp-ap/cloudepg-provEPG"/>
            </cloudListenerRule>
            <cloudListenerRule name="redirectRule" priority="20">
              <cloudRuleCondition type="path" value="/img/*"/>
              <cloudRuleAction type="redirect" RedirectPort="8080"/>
            </cloudListenerRule>
            <cloudListenerRule name="FixedRspRule" priority="30">
              <cloudRuleCondition type="host" value="example.com"/>
              <cloudRuleAction type="fixedResponse" FixedResponseCode="200"/>
            </cloudListenerRule>
            <cloudListenerRule name="redirectHPRule" priority="40" status="">
              <cloudRuleCondition type="host" value="example.com"/>
              <cloudRuleCondition type="path" value="/img/*"/>
              <cloudRuleAction type="forward" port="80" protocol="http"
epgdn="uni/tn-t2/cloudapp-ap/cloudepg-provEPG"/>
            </cloudListenerRule>
          </cloudListener>
        </cloudSvcPolicy>
      </vnsAbsNode>
    </vnsAbsGraph>
  </fvTenant>
</polUni>
```

Configuring a Key Ring Using the REST API

This example demonstrates how to configure a key ring using the REST API. For more information about key ring configuration, see the *Cisco APIC Basic Configuration Guide*.

To configure a key ring:

```
<polUni>
  <fvTenant name="t2">
    <cloudCertStore>
      <pkiKeyRing status="" name="lbCert" tp="lbTP" key="-----BEGIN RSA PRIVATE KEY-----
MIIEpQIBAAKCAQEA4DGxaK+RHv/nToHLnmDBq2BfLimgX/zNJQC9bGuzr8Mj7dm0
XuHfQYGv0h1PtL4Pdxf5qjB0NbHjAVB1Gw8cDiErEgAXy9Km27ySo2foKryNqCRe
Ginn/CgF75QPIed568eScNDZPt/eMeHAuRX/PykKUatWWncGanjvHqc+SOLPF6TD
gQ5nwOHHFvyM2DY8bfdYWrWmGsO7JqZzbPMptA2QWblILsSoIrdkIIgf6ZfYy/EN
bH+nYN2rJT8lzYsxz0YmR0oRQHTiN2NiDY/ZV63yxCXfLg9qpNZCuD8KOfdCZPEq
8takiWBxiR5/HRPscWAdWQsoiKgG1k4NEbFA9QIDAQABAoIBAQDQqA9IslYrdtqN
q6mZ3s2BNfF/4kgb7gn0DWs+9EJJLCJNZVhFEo2ZxxyfPp6HRnjYS50W83/E1anD
+GD1bSucTuxqFWIQVh7r1ebYZIWk+NYSjr5yNVxux8U2hCNNV8WWVqkJjKcUqICB
Bm47FKj53LV46zE0gyCaibFrYxZJ9+farGneyBdnoV+3thmez7534KCi0t3J3Eri
lgSY3ql6hPXB2ZXAP4jdAoLgWDU4I1M6OqOiWopZM/QYIE/WtPYyJ0QzNCXObtc5
FboDcvedsgd4x5GlfV2A4xTBQMCTZUZJ9fYAcFogTZXD+UVqxorh47tf/mz+1fjq
f1XphEDlAoGBAPVlvKfGW46qqRnYovfryxxz4OMlsVSgcJpQTQtBQi2koJ8OwEZJ
2s+CX0r+oDqwP23go/QEVYVkcic9RGkJBNge1+dm/bTjzgMQYtqSCNtecTsZD5JN
y1jkciizznDkjcjReSZ2kh3dGXIbRiYk7ezp2z7EKfDrHe5x5ouGMgCnAoGBAOnh
buDEohv8KJaB+DiUfhtoa3aKNPBO+zWPCHp0HFGjPXshJcIYZc1GcycmuDKVNnDd
MxhE/yOnQHowi4T9FMLpz5yh5zuCUVqOBgB1P6MzbC5t5MtLrEYr/AqFN11CqyXQ
cVcT6iCW1OAFJRw3c/OiESwLMzchsl8RnbwOi6kDAoGBANVlzmPb07zB3eGTCU0t
KGiqwFLncUkVaDZZRFZYPpNwiRkoe73j9brkNbgCqxW+NLp5UjoeFry0N6y106q/
ZA4I7FnXryLBw2HYuw41Vixl+XOZ/HeO3RmFN1z717dGmaGbv43aKIB9x+X5n8wF
6z1NtBHmBk7yNwom1IRag1sbAoGAX0p4cJ/tJNXSe7AswHDQCL68uimJdDfZ5nKG
k83nE+Qc0qQozDJAmCiSFmuSNRnSep3FiafjBFXK0X4h+mdbJCc7bagRnI92Mh0X
mOwsp4P2GdywkZwdbuHQ6UBp1Ferf9aztzTn+as6xKOUATEezy9DK9zMWzQhhtaY
m9yZTp0CgYEA1UtcpWjAzQbXODJGmxGdAAakPpeiKw/Da3MccrTdGJt88ezM1Oej
Pdoab0G2PcfgJZoTSGk7N4XArVKeq7pgZ0kwcYAshO6A2Hal+D1z/bGoZP+kmD/x
Ny82phxYOXCnEc5Vv92lU59+j7e067UFLAYJe6fu+oFImvofRnP4DIQ= -----END RSA PRIVATE KEY-----" cert="-----BEGIN
 CERTIFICATE----- MIIElTCCA32gAwIBAgIJAKWNjp//arBsMA0GCSqGSIb3DQEBCwUAMIGNMQswCQYD
VQQGEwJVUzELMAkGA1UECBMCQ0ExETAPBgNVBAcTCFNhbiBKb3NlMRIwEAYDVQQK
EwlNeUNvbXBhbnkxDjAMBgNVBAsTBU15T3JnMRgwFgYDVQQDFA8qLmFtYXpvbmF3
cy5jb20xIDAeBgkqhkiG9w0BCQEWEXJhbXNoYWhAY2lzY28uY29tMB4XDTE4MTAw
MjIwNTMwNVoXDTE5MTAwMjIwNTMwNVowgY0xCzAJBgNVBAYTAlVTMQswCQYDVQQI
EwJDQTERMA8GA1UEBxMIU2FuIEpvc2UxEjAQBgNVBAoTCU15Q29tcGFueTEOMAwG
A1UECxMFTXlPcmcxGDAWBgNVBAMUDyouYW1hem9uYXdzLmNvbTEgMB4GCSqGSIb3
DQEJARYRcmFtc2hhaEBjaXNjby5jb20wggEiMA0GCSqGSIb3DQEBAQUAA4IBDwAw
ggEKAoIBAQDgMbFor5Ee/+dOgcueYMGrYF8uKaBf/M0lAL1sa7OvwyPt2bRe4d9B
ga/SHU+0vg93F/mqMHQ1seMBUHUbDxwOISsSABfL0qbbvJKjZ+gqvI2oJF4aKef8
KAXvlA8h53nrx5Jw0Nk+394x4cC5Ff8/KQpRq1ZadwZqeO8epz5I4s8XpMOBDmfA
4ccW/IzYNjxt91hataYaw7smpnNs8ym0DZBZuUguxKgit2QgiB/pl9jL8Q1sf6dg
3aslPyXNizHPRiZHShFAdOI3Y2INj9lXrfLEJd8uD2qk1kK4Pwo590Jk8Sry1qSJ
YHGJHn8dE+xxYB1ZCyiIqAbWTg0RsUD1AgMBAAGjgfUwgfIwHQYDVR0OBBYEFBYq
K3b39+1oOr4IBSsePwcOpML7MIHCBgNVHSMEgbowgbeAFBYqK3b39+1oOr4IBSse
PwcOpML7oYGTpIGQMIGNMQswCQYDVQQGEwJVUzELMAkGA1UECBMCQ0ExETAPBgNV
BAcTCFNhbiBKb3NlMRIwEAYDVQQKEwlNeUNvbXBhbnkxDjAMBgNVBAsTBU15T3Jn
MRgwFgYDVQQDFA8qLmFtYXpvbmF3cy5jb20xIDAeBgkqhkiG9w0BCQEWEXJhbXNo
YWhAY2lzY28uY29tggkApY2On/9qsGwwDAYDVR0TBAUwAwEB/zANBgkqhkiG9w0B
AQsFAAOCAQEAe/RuzCheLIbHbrurGet6eaVx9DPYydNiKVBSAKO+5iuR84mQzhoT
nx5CN109xu5ml5baCYZZsSnn6D7usC092bPA/kRCGxt29gkjpWA74tJHqIhVWgbM
mOrLiSHoelewv+wRl0oVRChlTfKtXO68TUk6vrqpw76hKfOHIa7b2h1IIMdq6VA/
+A5FQ0xqYfqKdVd2RaINpzI8mqZiszqw+7E6j1PL5k4tftWEaYpfGPlVesFEyJEL
gHBUiPt8TIbaMYI8qUQmB/emnLXeKQ5PRxdRnleA3h8jfq3D1CQRTLjmDL3tpFwg qopM6et5ZKqShX4T87BsgZIoiquzXqsuHg==
 -----END CERTIFICATE-----">
      </pkiKeyRing>
      <pkiTP status="" name="lbTP" certChain="-----BEGIN CERTIFICATE-----
MIIElTCCA32gAwIBAgIJAKWNjp//arBsMA0GCSqGSIb3DQEBCwUAMIGNMQswCQYD
VQQGEwJVUzELMAkGA1UECBMCQ0ExETAPBgNVBAcTCFNhbiBKb3NlMRIwEAYDVQQK
EwlNeUNvbXBhbnkxDjAMBgNVBAsTBU15T3JnMRgwFgYDVQQDFA8qLmFtYXpvbmF3
cy5jb20xIDAeBgkqhkiG9w0BCQEWEXJhbXNoYWhAY2lzY28uY29tMB4XDTE4MTAw
```

```
MjIwNTMwNVoXDTE5MTAwMjIwNTMwNVowgY0xCzAJBgNVBAYTAlVTMQswCQYDVQQI
EwJDQTERMA8GA1UEBxMIU2FuIEpvc2UxEjAQBgNVBAoTCU15Q29tcGFueTEOMAwG
A1UECxMFTXlPcmcxGDAWBgNVBAMUDyouYW1hem9uYXdzLmNvbTEgMB4GCSqGSIb3
DQEJARYRcmFtc2hhaEBjaXNjby5jb20wggEiMA0GCSqGSIb3DQEBAQUAA4IBDwAw
ggEKAoIBAQDgMbFor5Ee/+dOgcueYMGrYF8uKaBf/M0lAL1sa7OvwyPt2bRe4d9B
ga/SHU+0vg93F/mqMHQ1seMBUHUbDxwOISsSABfL0qbbvJKjZ+gqvI2oJF4aKef8
KAXvlA8h53nrx5Jw0Nk+394x4cC5Ff8/KQpRq1ZadwZqeO8epz5I4s8XpMOBDmfA
4ccW/IzYNjxt91hataYaw7smpnNs8ym0DZBZuUguxKgit2QgiB/pl9jL8Q1sf6dg
3aslPyXNizHPRiZHShFAdOI3Y2INj9lXrfLEJd8uD2qk1kK4Pwo590Jk8Sry1qSJ
YHGJHn8dE+xxYB1ZCyiIqAbWTg0RsUD1AgMBAAGjgfUwgfIwHQYDVR0OBBYEFBYq
K3b39+1oOr4IBSsePwcOpML7MIHCBgNVHSMEgbowgbeAFBYqK3b39+1oOr4IBSse
PwcOpML7oYGTpIGQMIGNMQswCQYDVQQGEwJVUzELMAkGA1UECBMCQ0ExETAPBgNV
BAcTCFNhbiBKb3NlMRIwEAYDVQQKEwlNeUNvbXBhbnkxDjAMBgNVBAsTBU15T3Jn
MRgwFgYDVQQDFA8qLmFtYXpvbmF3cy5jb20xIDAeBgkqhkiG9w0BCQEWEXJhbXNo
YWhAY2lzY28uY29tggkApY2On/9qsGwwDAYDVR0TBAUwAwEB/zANBgkqhkiG9w0B
AQsFAAOCAQEAe/RuzCheLIbHbrurGet6eaVx9DPYydNiKVBSAKO+5iuR84mQzhoT
nx5CN109xu5ml5baCYZZsSnn6D7usC092bPA/kRCGxt29gkjpWA74tJHqIhVWgbM
mOrLiSHoelewv+wRl0oVRChlTfKtXO68TUk6vrqpw76hKfOHIa7b2h1IIMdq6VA/
+A5FQ0xqYfqKdVd2RaINpzI8mqZiszqw+7E6j1PL5k4tftWEaYpfGPlVesFEyJEL
gHBUiPt8TIbaMYI8qUQmB/emnLXeKQ5PRxdRnleA3h8jfq3D1CQRTLjmDL3tpFwg qopM6et5ZKqShX4T87BsgZIoiquzXqsuHg==
 -----END CERTIFICATE-----">
        </pkiTP>
    </cloudCertStore>
  </fvTenant>
</polUni>
```

Creating an HTTPS Service Policy Using the REST API

This section demonstrates how to create an HTTPS service policy using the REST API.

Note A listener can have multiple certificates. The certificate options are:

- ELBSecurityPolicy-2016-08 – The default when no security policy is chosen.
- ELBSecurityPolicy-FS-2018-06
- ELBSecurityPolicy-TLS-1-2-2017-01
- ELBSecurityPolicy-TLS-1-2-Ext-2018-06
- ELBSecurityPolicy-TLS-1-1-2017-01
- ELBSecurityPolicy-2015-05
- ELBSecurityPolicy-TLS-1-0-2015-04

If you use multiple certificates, you must specify the default certificate. The default is specified using the **defaultCert** property in **cloudRsListenerToCert**.

Before you begin

You have already configured a key ring certificate.

To create an HTTPS service policy:

```
<polUni>
  <fvTenant name="t2">
    <vnsAbsGraph name="CloudGraph" type="cloud" status="">
      <vnsAbsNode funcType="GoTo" name="N1" managed="yes">
        <cloudSvcPolicy tenantName="t2" contractName="httpFamily" subjectName="consubj">
          <cloudListener name="https_listener" port="443" protocol="https"
secPolicy="eLBSecurityPolicy-2016-08" status="">
            <cloudRsListenerToCert defaultCert="yes" certStore="iam"
tDn="uni/tn-t2/certstore/keyring-lbCert" status=""/>
            <cloudListenerRule name="defaultRule" default="yes" priority="100" status="">
              <cloudRuleAction type="forward" port="80" protocol="http"
epgdn="uni/tn-t1/cloudapp-ap/cloudepg-ep1">
                                                </cloudRuleAction>
            </cloudListenerRule>
          </cloudListener>
        </cloudSvcPolicy>
      </vnsAbsNode>
    </vnsAbsGraph>
  </fvTenant>
</polUni>
```