

Inter-Fabric L3Out with PBR for ACI Fabrics, 4.3.1

Table of Contents

New and changed information . . . 1
Inter-Fabric L3Out with PBR . . . 2
  Configuration Workflow . . . 2
  Supported Use Cases . . . 3
    Intra-VRF vs Inter-VRF . . . 3
    L3Out to Stretched EPG . . . 3
    L3Out to Fabric-Local EPG . . . 4
Guidelines and Limitations . . . 7
Create Service Device Template . . . 8
Add Service Chaining to Contract . . . 10
Copyright . . . 11

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 4.3.1	There were no major changes from the previous release.	

Inter-Fabric L3Out with PBR

Cisco Application Centric Infrastructure (ACI) policy-based redirect (PBR) enables traffic redirection for service appliances, such as firewalls or load balancers, and intrusion prevention system (IPS). Typical use cases include provisioning service appliances that can be pooled, tailored to application profiles, scaled easily, and have reduced exposure to service outages. PBR simplifies the insertion of service appliances by using contract between the consumer and provider endpoint groups even if they are all in the same virtual routing and forwarding (VRF) instance.

PBR deployment consists of configuring a route redirect policy and a cluster redirect policy, and creating a service graph template that uses these policies. After the service graph template is deployed, you can attach it to a contract between EPGs so that all traffic following that contract is redirected to the service graph devices based on the PBR policies you have created. Effectively, this allows you to choose which type of traffic between the same two EPGs is redirected to the L4-L7 device, and which is subject to a security policy applied at the fabric level.

More in-depth information specific to services graphs and PBR is available in the *Cisco APIC Layer 4 to Layer 7 Services Deployment Guide*

Configuration Workflow

The use cases described in the following sections are an extension of a basic inter-fabric L3Out (without PBR) use case which is in turn an extension on basic external connectivity (L3Out) configuration in each fabric. The workflow to configure the supported use cases is the same, with the only differences being whether you create the objects in the same or different VRFs (intra-VRF vs inter-VRF) and where you deploy the objects (stretched vs non-stretched).

1. Configure basic external connectivity (L3Out) for each fabric.

   The inter-fabric L3Out with PBR configuration described in the following sections is built on top of existing external connectivity (L3Out) in each fabric. If you have not configured an L3Out, create and deploy one as described in External Connectivity (L3Out) before proceeding with the following sections.

2. Configure an inter-fabric L3Out use case **without** PBR.

   We recommend configuring a simple inter-fabric L3Out use case without any policy-based redirection before adding service chaining to it. This is described in detail in Inter-Fabric L3Out.

3. Add service chaining to the L3Out contract as described in the following sections, which includes:
   - Adding an external TEP pool for each Pod in each fabric where inter-fabric L3Out is deployed.
   - Creating a Service Device template and assigning it to fabrics.

     The service device template must be assigned to the same fabrics as the L3Out and application templates that contain other configuration objects.

   - Providing fabric-level configurations for the Service Device template.

     Each fabric can have its own service device configuration including different high-availability models (such as active/active, active/standby, or independent service nodes).

- Associating the service device you defined to the contract used for the basic inter-fabric L3Out use case you deployed in the previous step.

Supported Use Cases

The following diagrams illustrate the traffic flows between an ACI internal endpoint in application EPG and an external endpoint through the L3Out in another fabric in the supported inter-fabric L3Out with PBR use cases.

Intra-VRF vs Inter-VRF

When creating and configuring the application EPG and the external EPG, you will need to provide a VRF for the application EPG's bridge domain and for the L3Out. You can choose to use the same VRF (intra-VRF) or different VRFs (inter-VRF).

When establishing a contract between the EPGs, you will need to designate one EPG as the provider and the other one as the consumer:

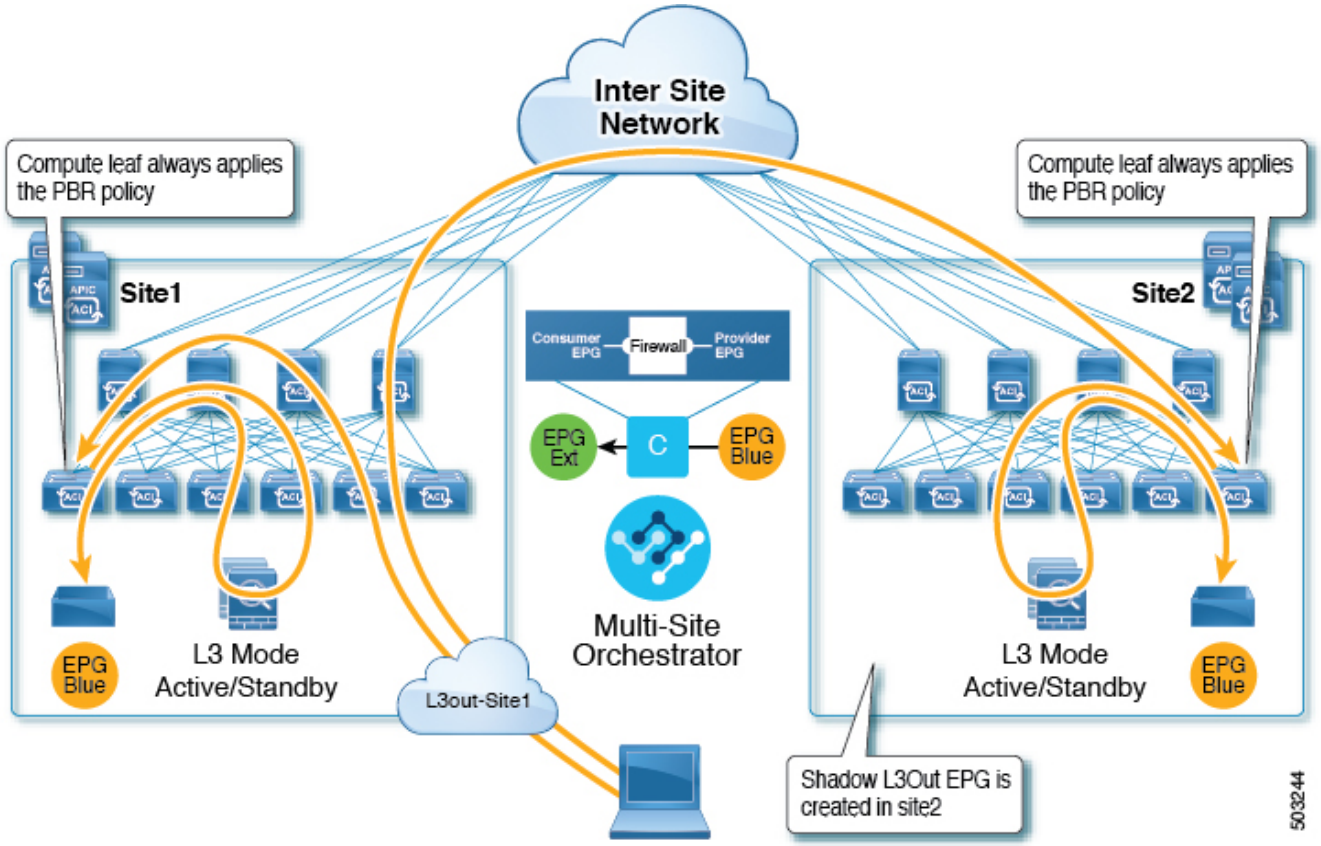
- When both EPGs are in the same VRF, either one can be the consumer or the provider.
- If the EPGs are in different VRFs, the external EPG must be the provider and the application EPG must be the consumer.

L3Out to Stretched EPG

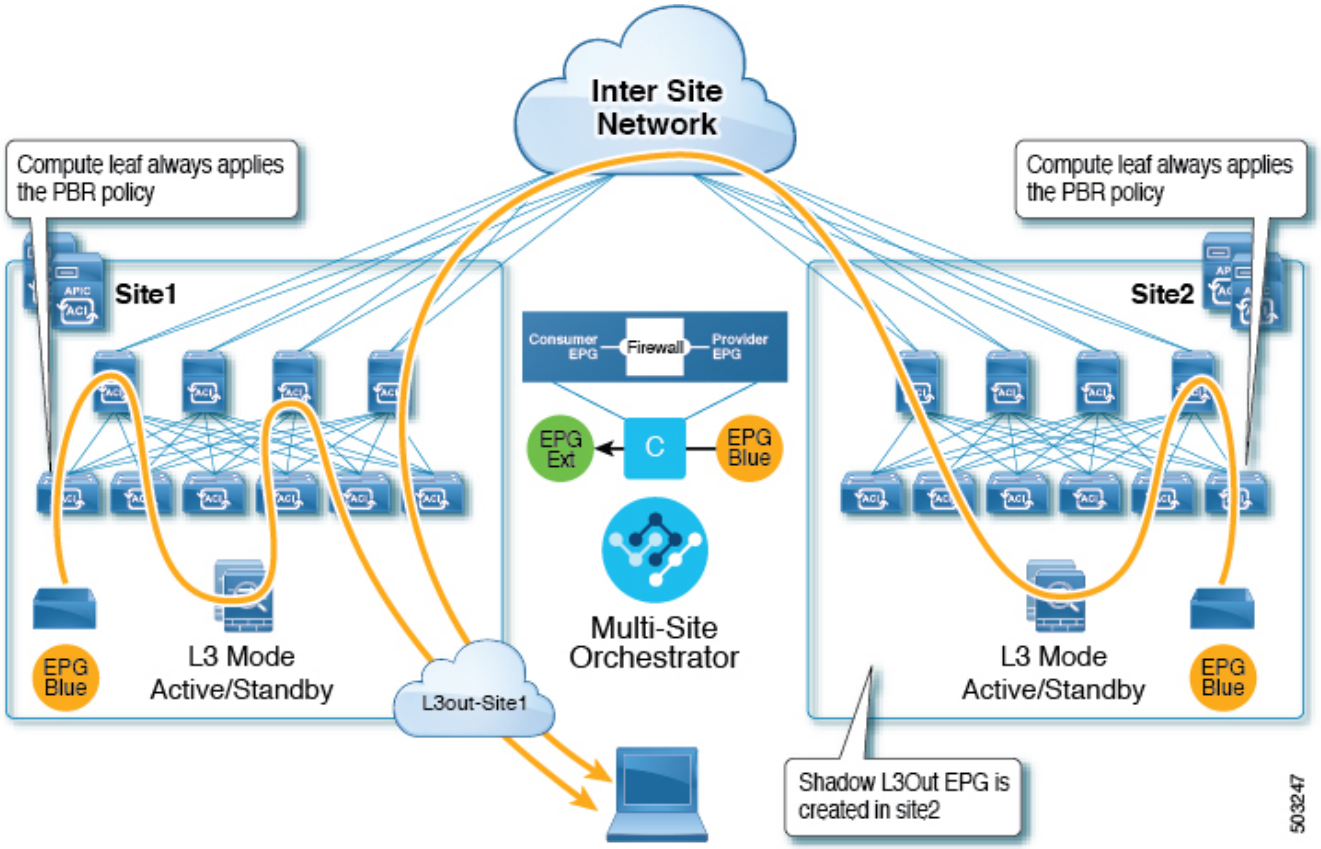
This use case illustrates a single application EPG that is stretched between two fabrics and a single L3Out created in only one of the fabrics. Regardless of whether the application EPG's endpoint is in the same fabric as the L3Out or the other fabric, traffic will go through the same L3Out. However, the traffic will always go through the service node that is local to the endpoint's fabric because for North-South traffic the PBR policy is always applied only on the compute leaf nodes (and not on the border leaf nodes).



The same flow applies in cases when the external EPG is stretched and each fabric has its own L3Out, but the L3Out in the fabric where the traffic is originating or is destined to is down.



Inbound Traffic



Outbound Traffic

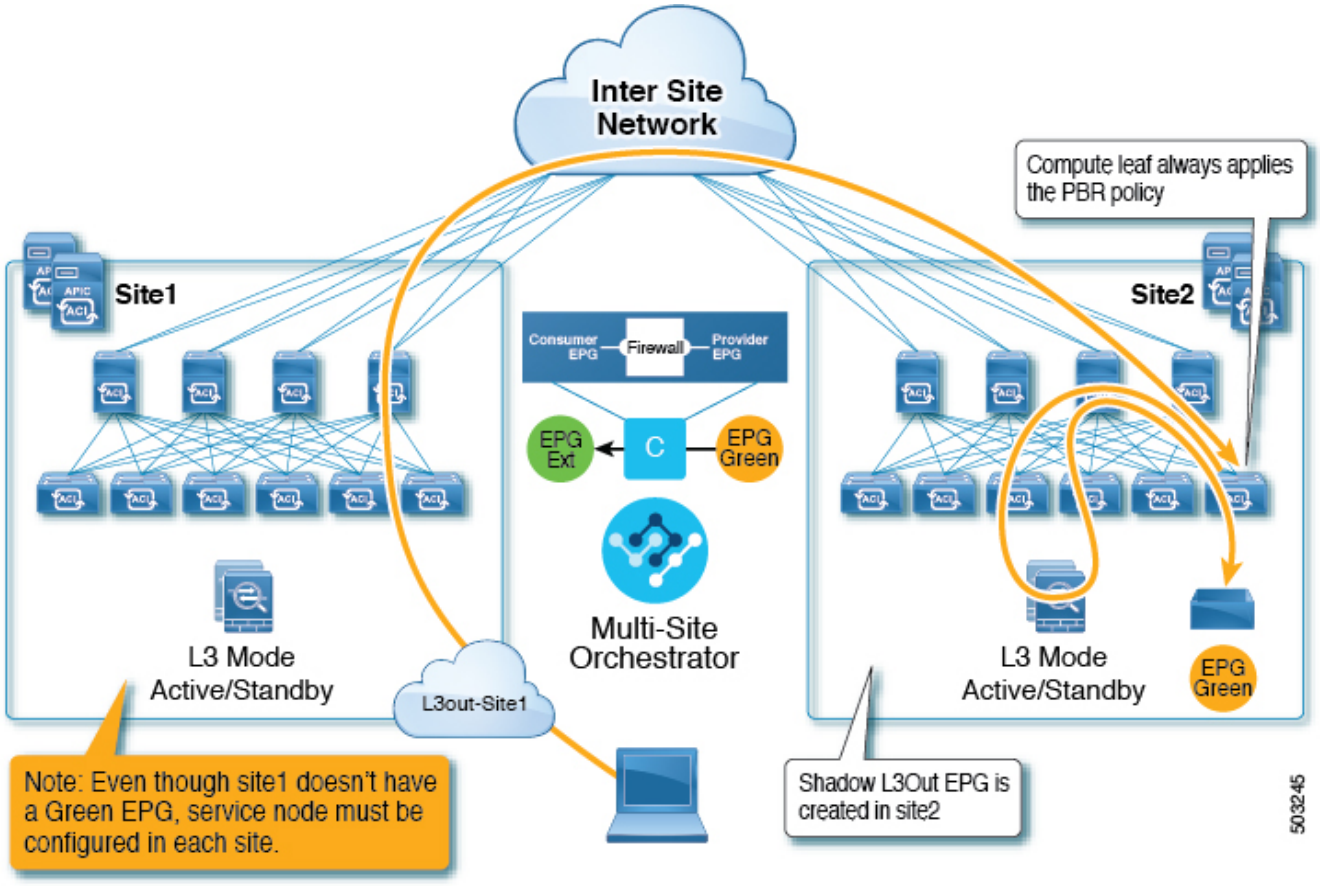
L3Out to Fabric-Local EPG

This use case illustrates a fabric-local application EPG that will use the L3Out in the other fabric for

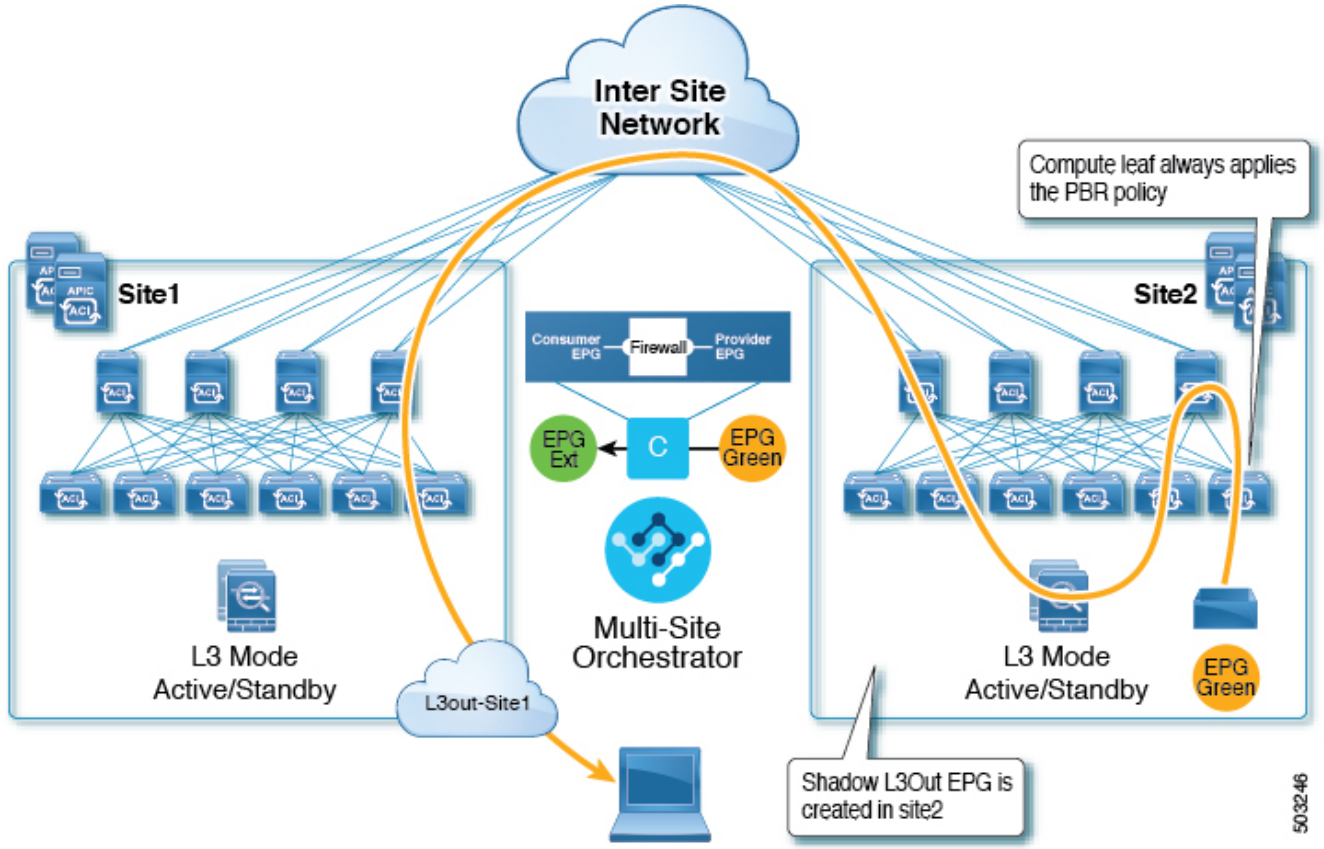
North-South traffic. Like in the previous example, all traffic will use the EPG's fabric-local service graph device.



The same flow applies in cases where the external EPG is stretched and each fabric has its own L3Out, but the EPG's local L3Out is down.



Inbound Traffic



Outbound Traffic

Guidelines and Limitations

When configuring an inter-fabric L3Out with PBR, the following restrictions apply:

- For inter-fabric L3Out with PBR, the following use cases are supported:
  - Inter-VRF inter-fabric L3Out with the application EPG as the consumer.

    For inter-VRF contracts, the External EPG associated to the L3Out must be the provider.

    This use case is supported for fabrics running Cisco APIC release 4.2(5) and later or release 5.1(x), it is not supported for APIC release 5.0(x).

  - Intra-VRF inter-fabric L3Out with the application EPG as either the provider or the consumer

    This use case is supported for fabrics running Cisco APIC release 4.2(5) and later or release 5.1(x); it is not supported for APIC release 5.0(x).

- For EPG-to-L3Out use cases, the application EPG can be stretched or fabric-local.
- For EPG-to-L3Out use cases, both one-arm and two-arm deployment models are supported; for L3Out-to-L3Out use case, only one-arm firewall devices are supported.

  In one-arm deployment, both the inside and outside interfaces of the service graph are connected to the same bridge domain. In two-arm deployments, the service graph interfaces are connected to separate BDs.

- For EPG-to-L3Out use cases, when configuring a load balancer with PBR, the load balancer and the real servers for the virtual IP (VIP) must be in the same fabric. If PBR is disabled, the load balancer and the real servers can be in different fabrics.

  L3Out-to-L3Out case does not support load balancers.

- You must have the basic use case of inter-fabric L3Out already configured before you insert a service device by enabling service chaining on the contract that is already configured between an L3Out in one fabric and an EPG in another fabric or between two L3Outs in different fabrics.

  Detailed instructions on deploying an inter-fabric L3Out without PBR are described in Inter-Fabric L3Out.

Create Service Device Template

- Ensure that you have read and completed the requirements described in Guidelines and Limitations.

This section describes how to configure one or more devices for a service graph.

1. Navigate to the **Orchestration** page.

   Manage > Orchestration

2. Create a Service Device template and associate it with the fabrics.
3. Choose **Tenant Templates > Service Device**.
   a. Click **Create Service Device Template**.
   b. In the template properties sidebar that opens, provide the **Name** for the template and **Select a Tenant**.
   c. In the **Template Properties** page, choose **Actions Add/Remove Fabrics** and associate the template with both fabrics.
   d. Click **Save** to save the template.
4. Create and configure the device cluster.
   a. In the **Template Properties** page (template-level configuration), choose **Create Object > Service Device Cluster**.

      The device cluster defines the service for which you want to redirect traffic.

   b. In the **<cluster-name>** sidebar, provide the **Name** for the cluster.

      The **Device Location** and **Device Mode** are pre-populated based on the currently supported use case.

   c. Choose the **Device Type**.
   d. For **Device Mode**, choose L3.
   e. Chose the **Connectivity Mode**.

      If you are configuring an L3Out-to-L3Out use case, you must use One Arm

   f. Provide the **Interface Name**.
   g. For the **Interface Type**, choose BD.

      For vzAny use cases, this release supports attaching the service device to a bridge domain only.

   h. Click **Select BD >** to choose the service bridge domain to which you want to attach this device.

      This is the stretched service BD you created in the previous section, for example FW-external.

   i. For the **Redirect** option, choose Yes.

You must choose to enable redirect for the PBR use case. After choosing Yes, the **IP SLA Monitoring Policy** option becomes available.

j. (Optional) Click **Select IP SLA Monitoring Policy** and choose an IP SLA policy if you had created one.

k. (Optional) In the **Advanced Settings** area, choose **Enable** if you want to provide additional settings for the service cluster.

You can configure the following advanced settings:

- **QoS Policy** - allows you assign a specific QoS level within the ACI fabrics for the redirected traffic.
- **Preferred Group** - specifies whether or not this service cluster is part of the preferred group.
- **Load Balancing Hashing** - allows you to specify the hashing algorithm for PBR load balancing.

  For additional information, see *ACI Policy-Based Redirect Service Graph Design*.

- **Pod Aware Redirection** - can be configured in Multi-Pod configuration if you want to specify the preferred PBR node. When you enable Pod-aware redirection, you can specify the Pod ID and redirection is programmed only in the leaf switches located in the specified Pod.
- **Rewrite Source MAC** - updates the source MAC address if the PBR node uses "source MAC based forwarding" instead of IP based forwarding.

  For additional information, see *ACI Policy-Based Redirect Service Graph Design*.

- **Advanced Tracking Options** - allows you to configure a number of advanced settings for the service node tracking. For additional information, see *Policy-Based Redirect and Threshold Settings for Tracking Service Nodes*

l. Click **Ok** to save.

Note that after you create the Service Device Cluster, it is highlighted in red in the **Template Properties** (template-level configuration) page. At this point, you have defined redirection to a firewall service, but you must still provide the firewall information and the redirect policy you want to use at the fabric-local level.

Add Service Chaining to Contract

After you have deployed the base inter-fabric L3Out use case and the Service Device template, you can add policy-based redirection by adding service chaining to the contract you created between the L3Out and an application EPG or another L3Out.

1. Navigate back to the application template where you defined the contract.
2. Select the contract.
3. Click **Service Chaining**.
4. In the **Service Chaining** area, click **+Service Chaining**.
5. Choose the **Device Type**.



The L3Out-to-L3Out use case supports only one-arm Firewall devices. For other inter-fabric L3Out with PBR use cases, you can chain multiple devices.

6. From the **Device** dropdown, choose the FW device cluster you created in the previous step.
7. Ensure that **Consumer Connector Type Redirect** is enabled.
8. Ensure that **Provider Connector Type Redirect** is enabled.
9. Click **Add** to continue.
10. Click **Save** to save the template.
11. Click **Deploy Template** to re-deploy it.

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