# Troubleshoot L3Out Subnet Classification in ACI

## Contents



## Introduction

This document describes the classification of external subnets within Cisco ACI's L3Out EPGs,

## Abbreviations

- BD: Bridge domain
- EPG: Endpoint group
- ExEPG: External Endpoint Group
- RIB: Routing Information Base
- VRF: Virtual Routing and Forwarding
- Class ID: Tag that identifies an EPG

## External EPG Classification

An External EPG in Cisco ACI represents external routed networks connected through L3Outs. Similar to how a regular EPG classifies endpoints, an External EPG classifies external subnets on a per-VRF basis, meaning each subnet must be unique within its VRF context.

A common misconception is that an External EPG subnet only includes prefixes accepted via the dynamic routing protocol. However, when an L3Out is created, there is a default route-map filtering incoming advertisements; thus, all prefixes advertised by the dynamic routing protocol are accepted by default. The primary purpose of defining subnets under an ExEPG is classification only to assign a unique pcTag to subnets enclosed in the ExEPG for contract enforcement and policy application.

This classification enables granular policy control. For example, a single external neighbor can advertise a supernet to ACI, which can then be segmented into multiple ExEPGs. This allows different contract actions to be applied to distinct subnets, such as permitting specific internal EPGs to communicate only with designated external subnets or redirecting traffic destined for certain prefixes to a PBR-node before reaching

its final destination.

This diagram illustrates how Cisco ACI classifies external subnets based on External EPGs, enabling precise traffic segmentation and contract enforcement.

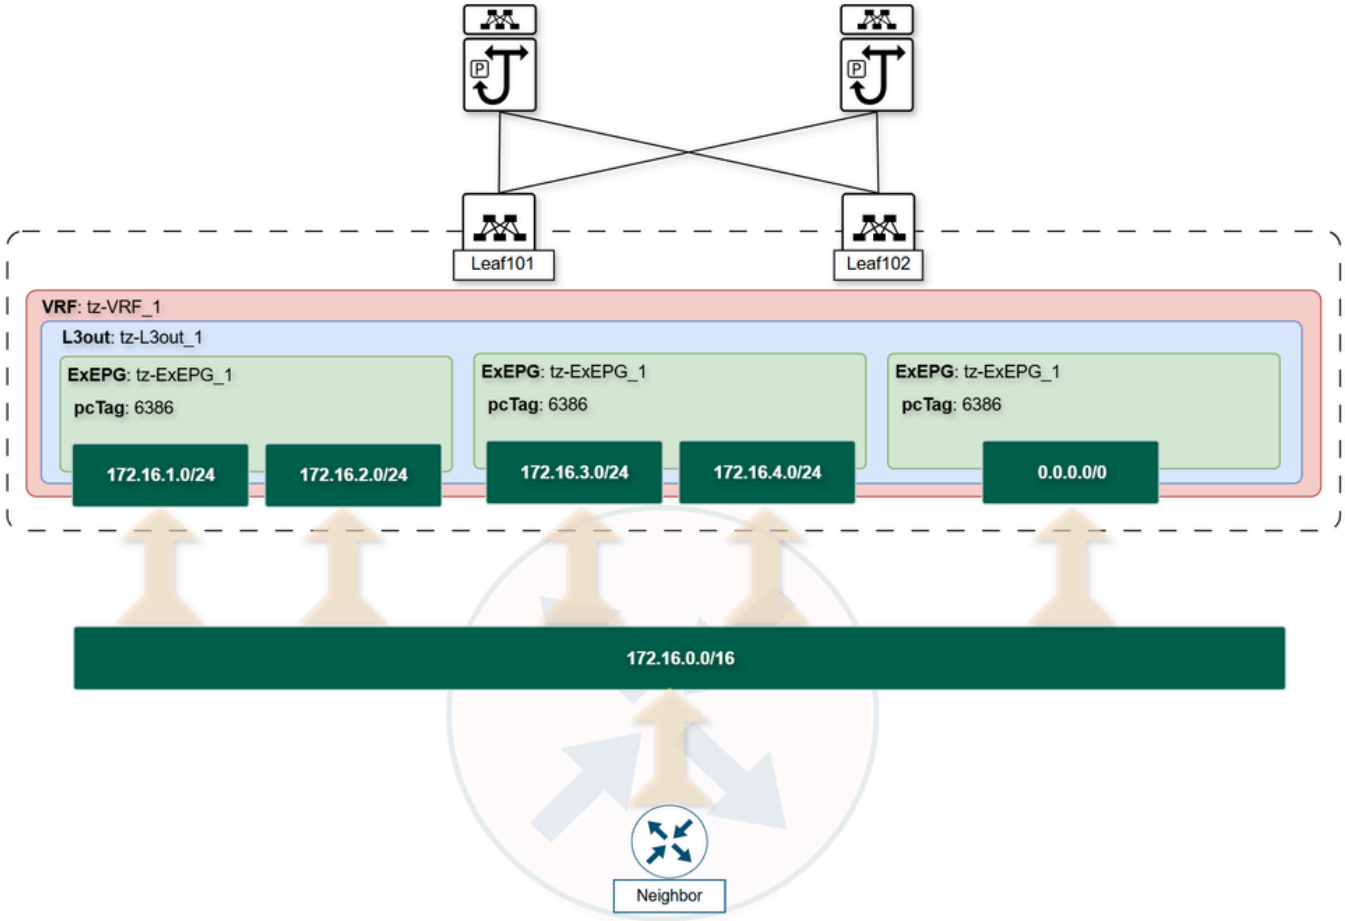


## External EPG Subnets Flags

To classify and manage external prefixes within an ExEPG in ACI, specific subnet flags are configured when creating a Subnet Prefix under an ExEPG. This section details each flag and its intended usage:

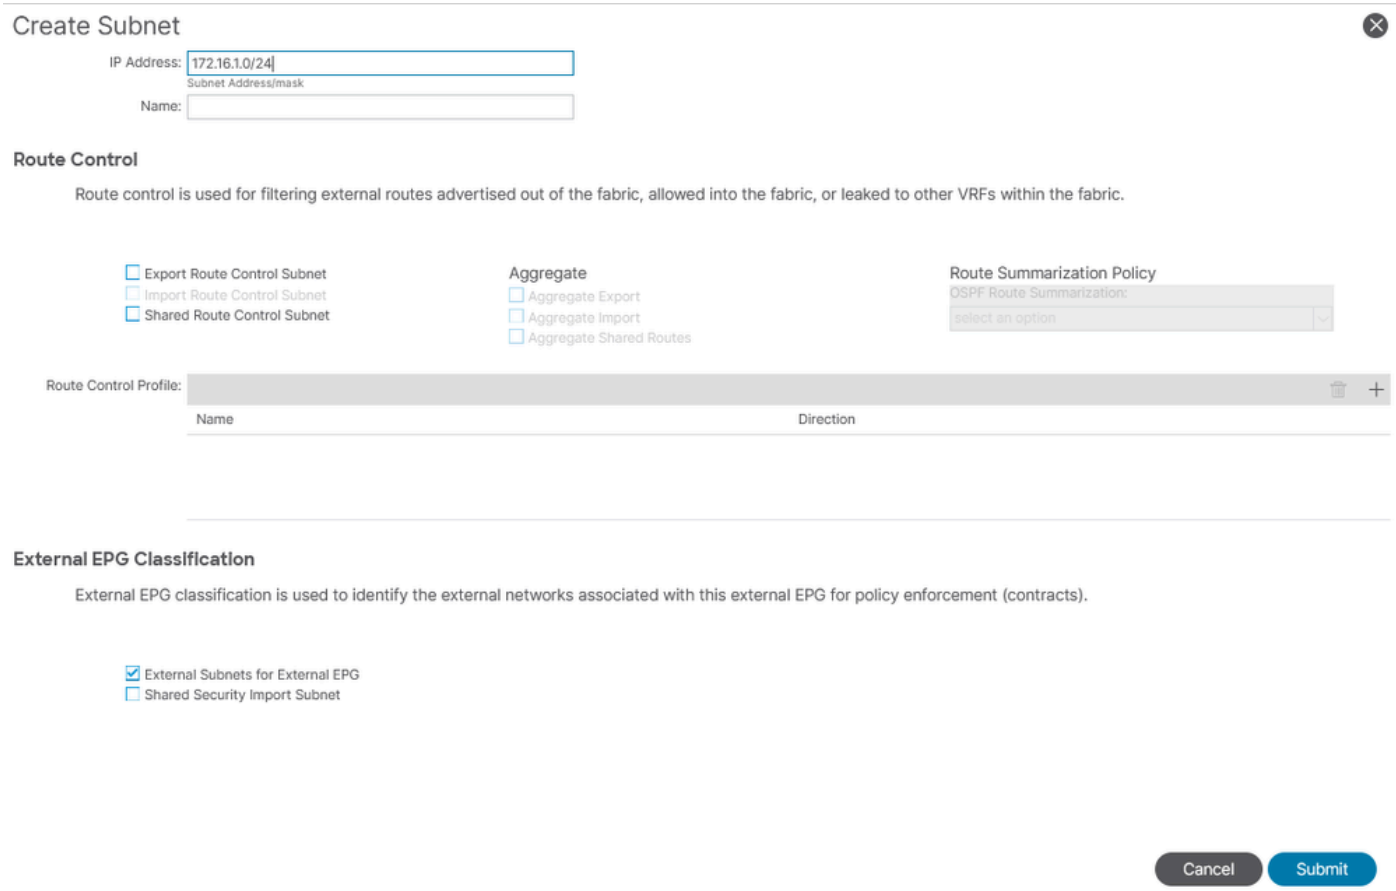


- External Subnet for External EPG:
  This flag indicates that the subnet resides outside the ACI fabric and is not configured within any Bridge Domain or EPG. It must be used only when the prefix is either advertised by a routing neighbor or statically injected into the RIB. This flag is enabled by default.

- Export Route Control Subnet:
  This flag designates that the subnet is advertised from ACI to the routing neighbor via the dynamic routing protocol. It must not be enabled simultaneously with the External Subnet for External EPG flag, as doing so can cause Layer 3 routing loops. Since ACI classifies the subnet as external and also advertises it back, this can lead to routing inconsistencies despite loop avoidance mechanisms in routing protocols.

- Shared Route Control Subnet:
  This flag is set when the subnet prefix is intended to be shared across multiple VRFs, enabling route leaking between the contexts.

- Shared Security Import Subnet:
  Used in conjunction with the Shared Route Control Subnet flag, this enables sharing of security pcTags for external subnets across different VRFs, facilitating consistent policy enforcement.

- Import Route Control Subnet:
  This flag allows granular control over the prefixes received from routing neighbors. By default, ACI accepts all incoming route advertisements; enabling this flag requires activating route control enforcement to filter incoming prefixes.

- Aggregate Section:
  Applicable only to the quad-0 (0.0.0.0/0) subnet, this section summarizes all prefixes in the RIB for aggregate export or import. When subnets are leaked to other VRFs, they are summarized as aggregated shared routes to optimize routing tables.

# Verification and Troubleshooting Commands

## Routing

To begin, the route must be present in the routing table of the VRF on the Border Leaf switches. For example, this command shows a BGP route in the VRF **tz:tz-VRF_1**:

```
<#root>

Leaf101#

show ip route bgp vrf tz:tz-VRF_1


IP Route Table for VRF "tz:tz-VRF_1"


'*' denotes best ucast next-hop
'**' denotes best mcast next-hop
'[x/y]' denotes [preference/metric]
'%<string>' in via output denotes VRF <string>


172.16.1.0/24

, ubest/mbest: 1/0

*via 10.10.1.2

%tz:tz-VRF_1, [20/0], 00:00:04, bgp-65002, external, tag 65003
Leaf101#
```

This confirms that the route is installed in the VRF routing table and is available for forwarding decisions.

## Classification

After the route is present in the routing table, classification determines how traffic is handled based on policy. In ACI, classification is tied to the ExEPG and its associated subnets.

To validate subnet classification under an ExEPG, the APIC can be queried for the l3extInstP class, which represents the External EPG instance. Its child class l3extSubnet lists the subnets configured under that ExEPG. For example:

```
<#root>

moquery -c l3extInstP -f 'l3ext.InstP.dn*"[ tenant name ].*[ l3out name ]"' -x rsp-subtree=children rsp-
```

```
<#root>

APIC#

moquery -c l3extInstP -f 'l3ext.InstP.dn*"tz.*l3out"' -x rsp-subtree=children rsp-subtree-class=l3extSub
```

```
Total Objects shown: 1

# l3ext.InstP

name : tz-ExEPG_1


!-- cut for brevity --!
configSt : applied
descr :
dn : uni/tn-tz/out-l3out/instP-tz-ExEPG_1
!-- cut for brevity --!
floodOnEncap : disabled
isSharedSrvMsiteEPg : no
lcOwn : local
matchT : AtleastOne
mcast : no
modTs : 2025-09-10T00:36:49.239+00:00
monPolDn : uni/tn-common/monepg-default
nameAlias :
pcEnfPref : unenforced

pcTag : 32771


pcTagAllocSrc : idmanager
prefGrMemb : exclude
prio : unspecified
rn : instP-tz-ExEPG_1
scope : 3047430
status : modified
targetDscp : unspecified
triggerSt : triggerable
txId : 1152921504612318828
uid : 15374
userdom : :all:

# l3ext.Subnet

ip : 172.16.1.0/24


!-- cut for brevity --!
dn : uni/tn-tz/out-l3out/instP-tz-ExEPG_1/extsubnet-[172.16.1.0/24]
extMngdBy :
lcOwn : local
modTs : 2025-09-10T01:05:13.249+00:00
monPolDn : uni/tn-common/monepg-default
!-- cut for brevity --!
rn : extsubnet-[172.16.1.0/24]

scope : import-security


status :
uid : 15374
userdom : :all:

APIC#
```

If no output is returned for the l3extSubnet class, it indicates that no subnets are configured under the External EPG. Without configured subnets, ACI cannot associate a pcTag to the incoming traffic subnet,

resulting in traffic being dropped despite the route existing in the routing table.

Another important aspect to note, is the scope of the subnet, this represents the flags set for the subnet in question:

- Import-security

The subnnet has been flagged with External Subnet for External EPG.

- export-rtctrl

The subnnet has been flagged with Export Route Control.

- import-rtctrl

The subnnet has been flagged with Import Route Control.

- shared-security

The subnnet has been flagged with Shared Security Import Subnet.

- shared-rtctrl

The subnnet has been flagged with Shared Route Control.

The routing protocols and control plane processes update the routing tables upon receiving a prefix from a mentioned neighbor, which are then programmed into the HAL L3 forwarding tables. The HAL L3 routes represent the actual Layer 3 routes programmed into the hardware forwarding tables (ASICs) on the leaf switches. These routes are derived from the routing protocols and routing table computations and are used for forwarding decisions.

<#root>

```
<-- When the prefix is not configured under the External EPG, a classification of 0xf is seen -->
Leaf101#
vsh_lc -c 'show platform internal hal l3 routes vrf tz:tz-VRF_1' | egrep "Prefix/Len|172.16.1.0" | cut -

VRF |   Prefix/Len  | RT|CLSS| Flags
4675| 172.16.1.0/ 24| UC| f|spi,dpi

Leaf101#
<-- When the prefix is configured under the External EPG, a classification of the pcTag in hexadecimal
Leaf101#
vsh_lc -c 'show platform internal hal l3 routes vrf tz:tz-VRF_1' | egrep "Prefix/Len|172.16.1.0" | cut -

VRF |   Prefix/Len  | RT|CLSS| Flags
4675| 172.16.1.0/ 24| UC|8003|spi,dpi

Leaf101#
Leaf101#
```

```
vsh_lc -c '

dec 0x8003'


32771
Leaf101#
```

Subsequently, when a subnet is configured with the External Subnet for External EPG flag in an ExEPG, an internal process called the Policy Manager (policy-mgr) updates its prefix-to-pcTag mapping table with this subnet entry and the associated pcTag. The Policy Manager serves as the fabric centralized policy orchestration engine, translating high-level policy definitions into actionable configurations across the ACI fabric. This ensures consistent and secure application connectivity and network behavior by enforcing the correct pcTags for traffic classification and forwarding decisions based on the configured external subnets.

```
<#root>

Leaf101#

vsh -c 'show system internal policy-mgr prefix' | egrep "tz:tz-VRF_1"


3047430 36 0x80000024 Up tz:tz-VRF_1 ::/0 15 True True False False
3047430 36 0x24 Up tz:tz-VRF_1 0.0.0.0/0 15 True True False False

3047430 36 0x24 Up tz:tz-VRF_1 172.16.1.0/24 32771 True True False False


Leaf101#
```

This confirms the prefix 172.16.1.0/24 is getting advertised by the neighbor to ACI border leaf switch, and ACI has classified the prefix bellow pcTag **32771**

## Contracts

A zoning-rule is the underlying process that enforces contract policies between EPGs (including ExEPGs) within the fabric. The VRF VNID (scope) and the pcTag of the External EPG can be used to define and validate the communication rules applied between source and destination EPGs. Essentially, zoning-rules translate high-level contract relationships into specific, enforceable rules programmed on the leaf switches.

An important aspect to consider is where the contract is installed in the fabric. By default, the VRF is configured with the Policy Control Enforcement Direction set to ingress. This setting determines that the zoning-rule for a given contract is installed on the leaf switch where the source endpoint resides.


For this exercise, traffic is incoming from an L3Out, the zoning-rule is installed on the border leaf that

connects to that L3Out, as this leaf acts as the source leaf for the traffic entering the fabric.

```
<#root>

Leaf101#

show zoning-rule scope 3047430 | egrep "Rule|---|32771"


+---------+--------+--------+----------+----------------+---------+---------+------------------+------
| Rule ID | SrcEPG | DstEPG | FilterID |      Dir       | operSt  |  Scope  |        Name      |
+---------+--------+--------+----------+----------------+---------+---------+------------------+------
|  4441   | 49153  | 32771  |     5    |     bi-dir     | enabled | 3047430 |     tz:Contract  |
|  4500   | 32771  | 49153  |     5    | uni-dir-ignore | enabled | 3047430 |     tz:Contract  |
+---------+--------+--------+----------+----------------+---------+---------+------------------+------
Leaf101#
```

### Transit Routing

Transit routing enables the fabric to act as a transit network by advertising external routes learned from one L3Out to another. To properly configure transit routing, the incoming subnet must be marked with the External Subnet for External EPG flag.

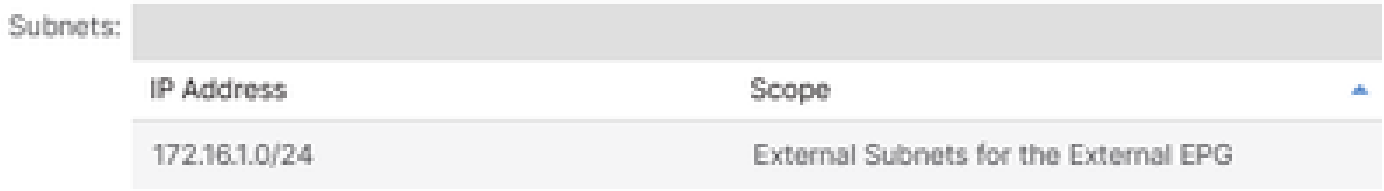


Simultaneously, the L3Out that advertises this subnet to other external peers must have the Export Route Control Subnet flag enabled on the corresponding subnet. This flag allows the subnet to be redistributed and advertised out of the fabric through the routing protocol configured on that L3Out.

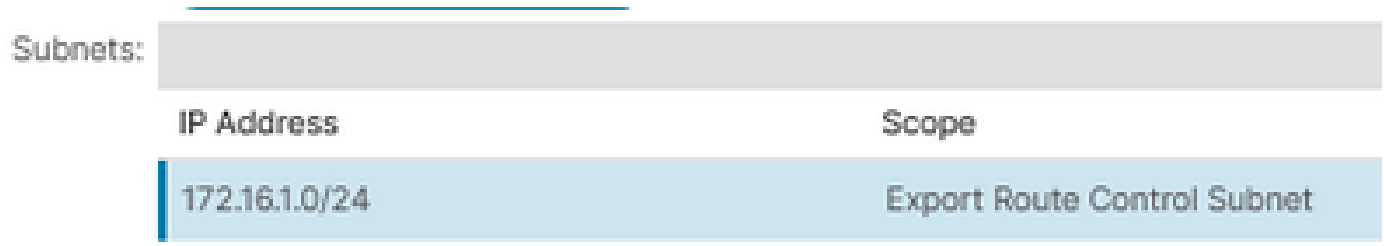


Lastly, a contract between the received and the exporting L3out needs to be configured to complete the process of having the route distributed.

## Common Issues in Subnet External EPG Classification

### pcTag 15

Previously, in this document, it was stated the ExEPG subnet helps you to classify the subnets in the correct pcTag for policy enforcement reasons. An important exception to this classification is the quad-0 subnet (0.0.0.0/0) when configured with the External Subnet for External EPG flag. This subnet is always assigned

the reserved pcTag 15, effectively acting as a wildcard for all external traffic within a VRF.

This diagram represents the issue facing configuring quad-0 with External Subnet for External EPG on multiple ExEPGs within the same VRF:

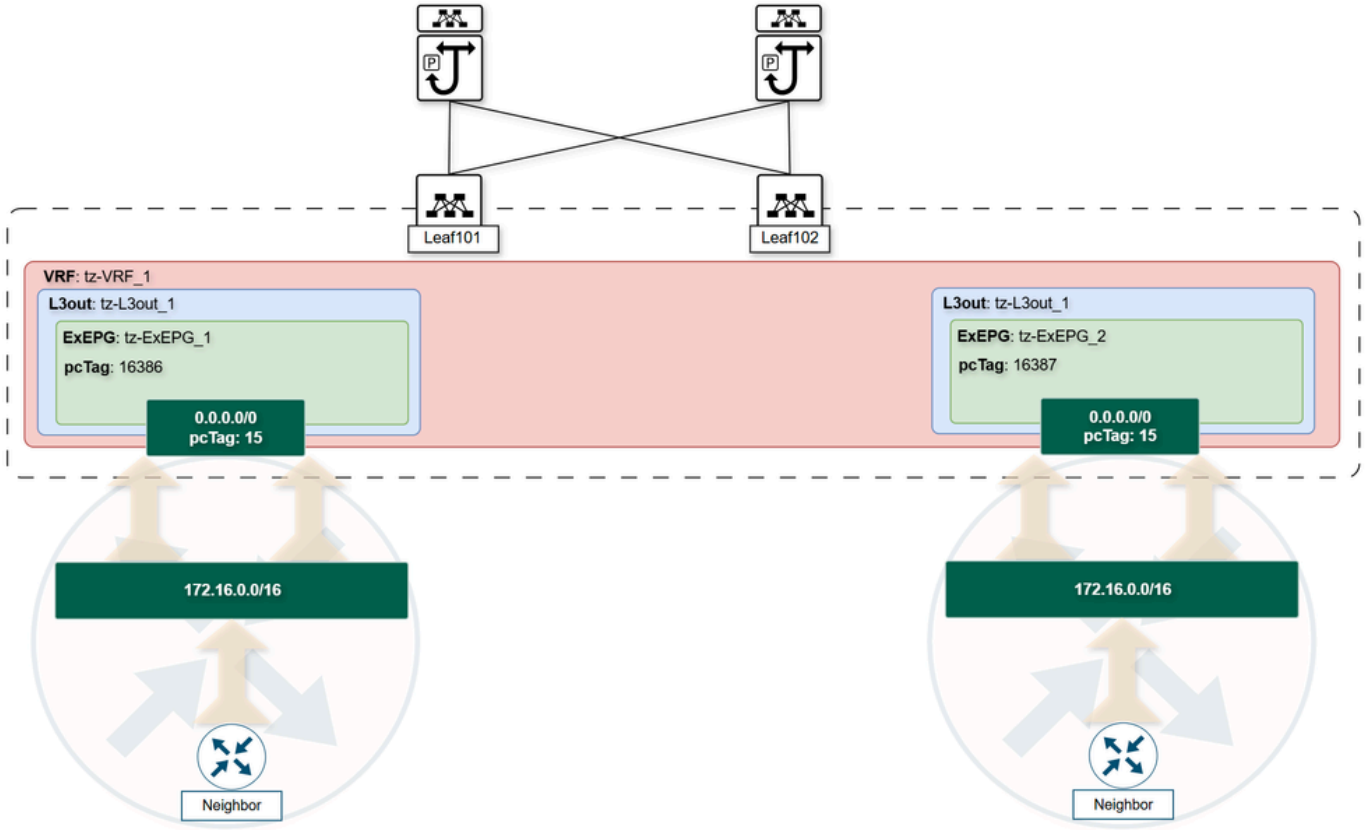


- The quad-0 subnet is often mistaken for the default route. While this is sometimes true—such as when a dynamic routing neighbor advertises only the default route to the ACI L3Out—the quad-0 subnet's role in ACI is broader as a catch-all classification.
- It is common practice to configure multiple ExEPGs with the quad-0 subnet to accept all prefixes advertised by a neighbor. Although this achieves the goal of broad acceptance, it can lead to unexpected asymmetric routing when multiple ExEPGs with quad-0 are configured within the same VRF. When multiple ExEPGs within the same VRF are configured with quad-0 as an external subnet, ACI cannot deterministically select which L3Out to use for a specific destination subnet. Instead, it selects one L3Out arbitrarily.
- This behavior can cause asymmetric routing, traffic intermittent, or even traffic drops if the randomly selected L3Out does not have the necessary contracts to permit communication.

## Overlapping Subnets

Configuration of identical subnets across different ExEPGs is not permitted. Attempting to do so triggers fault "F0467: Prefix Entry Already Used in Another EPG", preventing subnet duplication within a VRF.

However, overlapping subnets can exist across different VRFs because each VRF maintains an independent routing table context. This separation allows the same subnet to be configured in ExEPGs belonging to different VRFs. Despite this, caution is critical when performing VRF route leaking involving these overlapping subnets, as it can lead to asymmetric forwarding decisions due to conflicts in subnet classification (pcTag) versus routing information (RIB).

Key scenarios include:

- Route Leaking from Two VRFs into a Third VRF:
  When two VRFs leak the same subnet into a third VRF, the receiving VRF installs the first subnet it receives based on shared policy from the APIC. If the leaf switch handling this VRF reboots or the routing election changes, the routing table could update to a different L3Out, causing inconsistent forwarding behavior.

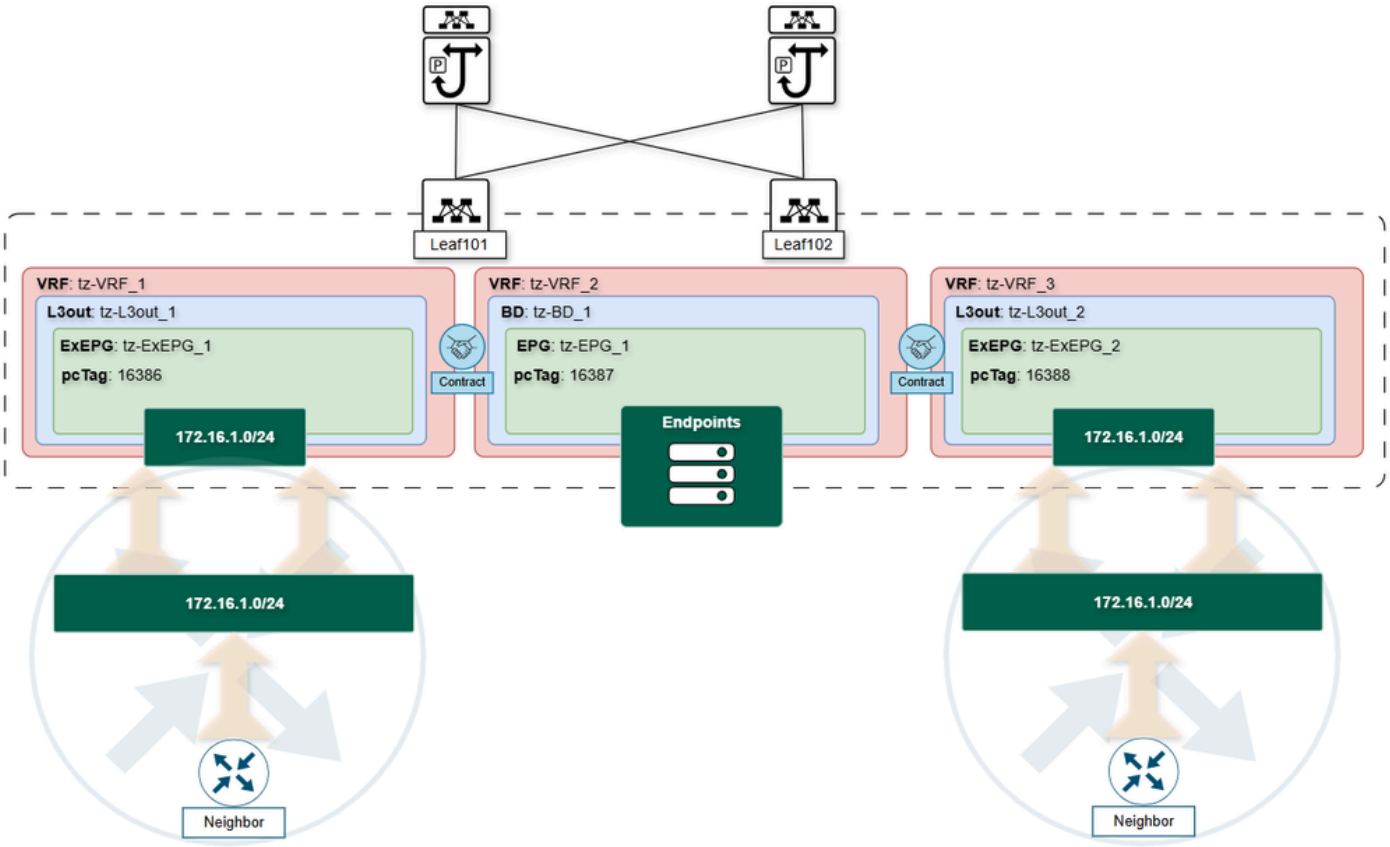


- Local-to-VRF L3Out ExEPG Overlapping with Leaked Subnets:
  In designs where route leaking is used, if a local L3Out ExEPG is configured with the same subnet as a leaked subnet, the local routing entry always take precedence over the leaked routes.

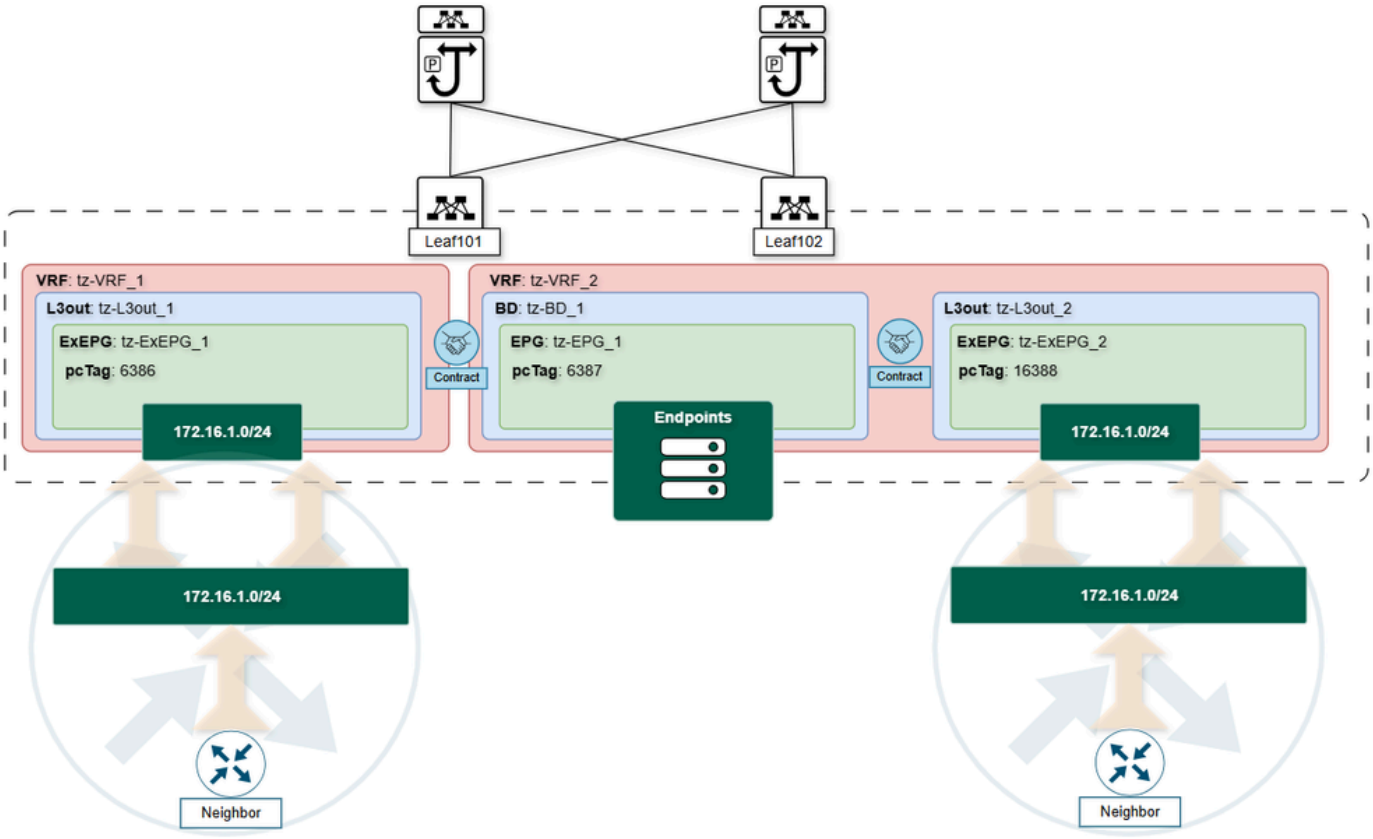


These situations highlight that the asymmetric forwarding issues arise from the classification and forwarding decision layer, not from the routing table itself. While subnet classification associates a subnet with a specific L3Out and ExEPG for policy enforcement, the routing table can point to a different L3Out destination. This mismatch can cause traffic to be forwarded inconsistently, leading to potential connectivity problems or policy enforcement gaps.

# Import Route Control Default Behavior Change

By default, ACI accepts all incoming route advertisements from neighbors. To control which prefixes are accepted, you must enable **Route Control Enforcement: inbound** on the L3Out root object:

Navigate to **Tenants > [ tenant name ] > Networking > L3outs > [ L3out name ]**.


This action creates a route-map under the selected routing protocol.

<#root>

```
Border Leaf#
```

```
show ip bgp neighbors vrf tz:tz-VRF1 | egrep route-map
```

```
Outbound route-map configured is exp-l3out-ExEPG-peer-2981888, handle obtained
Inbound route-map configured is imp-l3out-ExEPG-peer-2981888, handle obtained


Border Leaf#
show route-map imp-l3out-ExEPG-peer-2981888


route-map imp-l3out-ExEPG-peer-2981888,
permit
, sequence 15801
Match clauses:
ip address prefix-lists: IPv4-peer49155-2981888-exc-ext-inferred-import-dst


ipv6 address prefix-lists: IPv6-deny-all
Set clauses:

Border Leaf#
show ip prefix-list IPv4-peer49155-2981888-exc-ext-inferred-import-dst


ip prefix-list IPv4-peer49155-2981888-exc-ext-inferred-import-dst: 1 entries
seq 1 permit 172.16.1.0/24
Border Leaf#
```

By default, this import route-map permits all incoming prefixes. To modify this behavior:

Navigate to **Tenants > [ tenant name ] > Networking > L3outs > [ L3out name ] > Route map for import and export route control**

Select the default **import route-map** or create a new one using the **gear** icon at the top right.

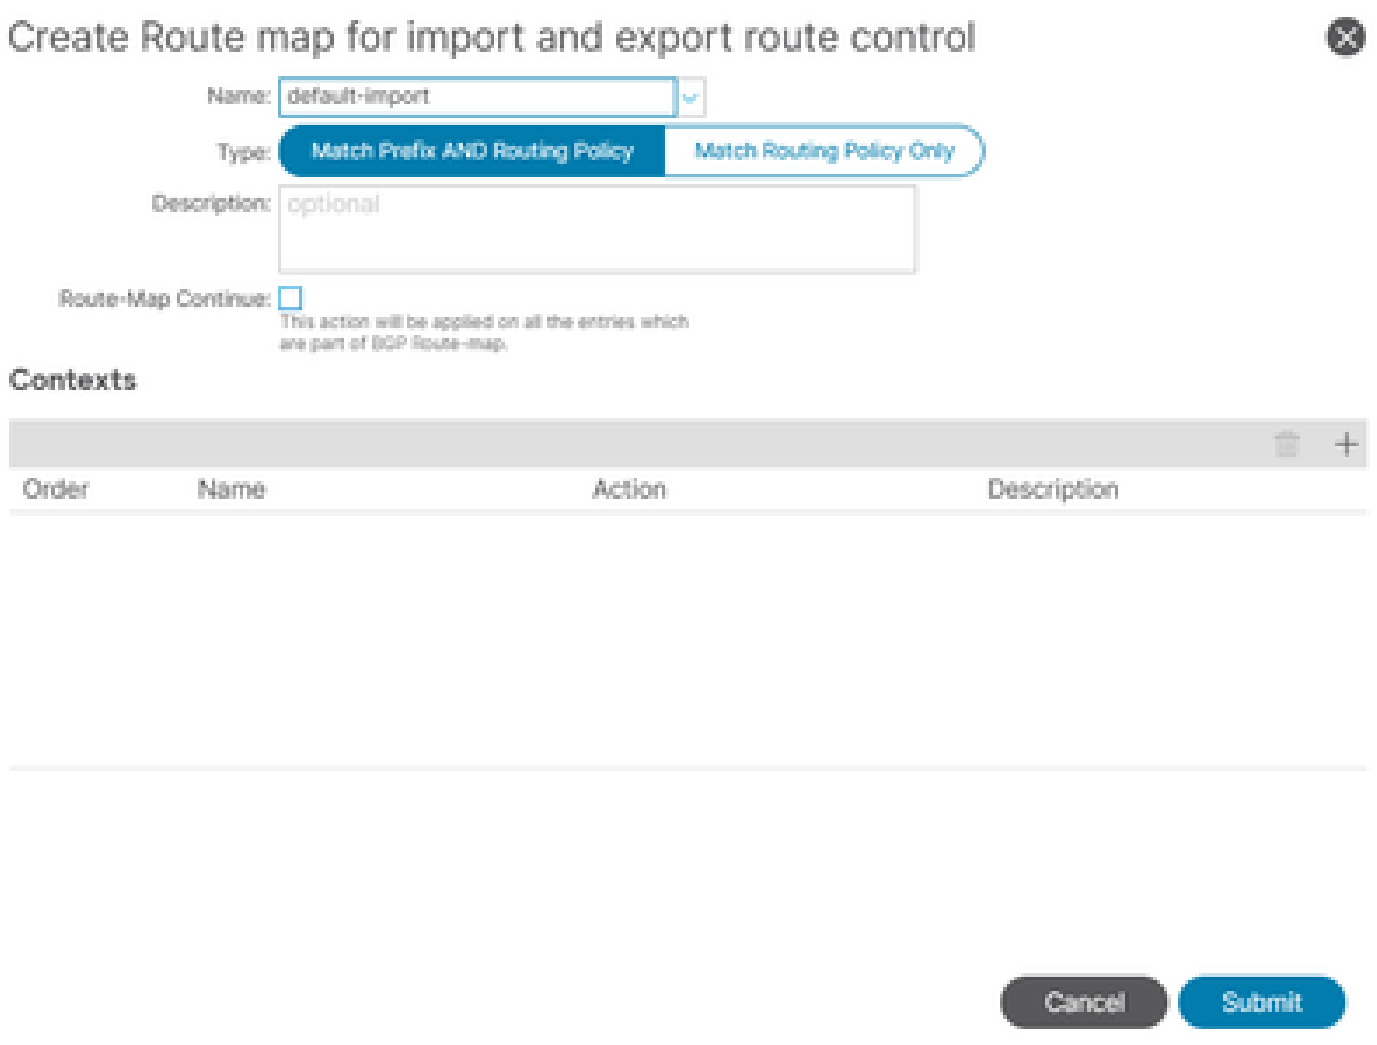


In the **Context** section, create a new **Associated Matched Rule**.

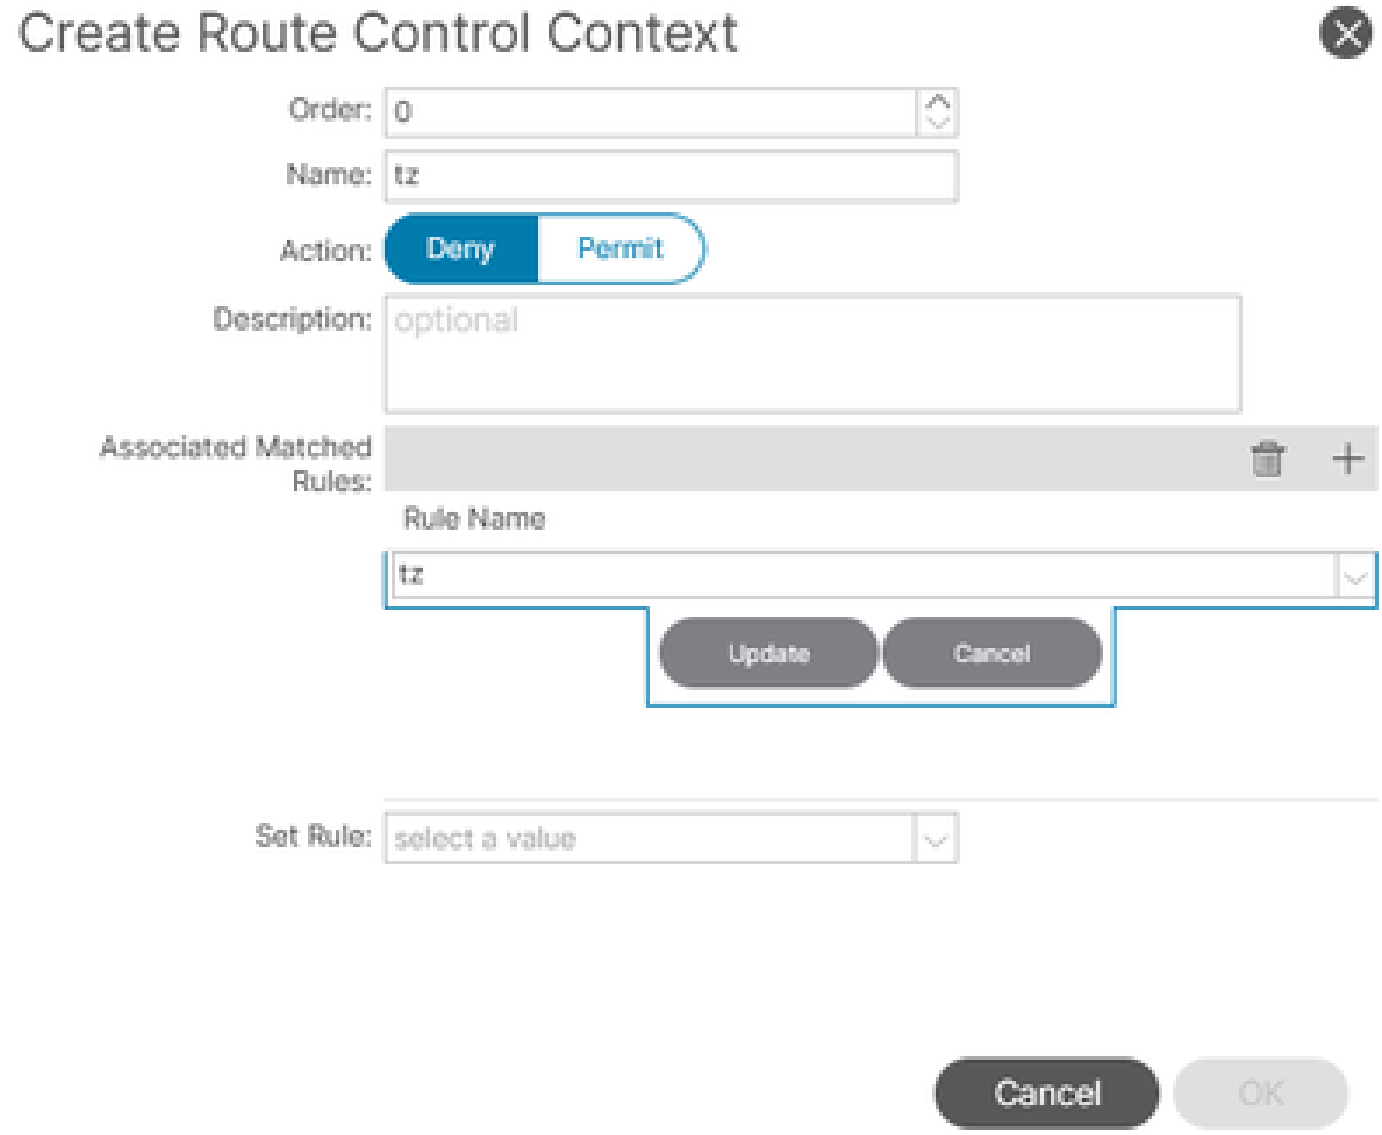


In the **Match Rules** section, scroll to **Match Prefix** and add the specific **subnet(s)** you want to control.

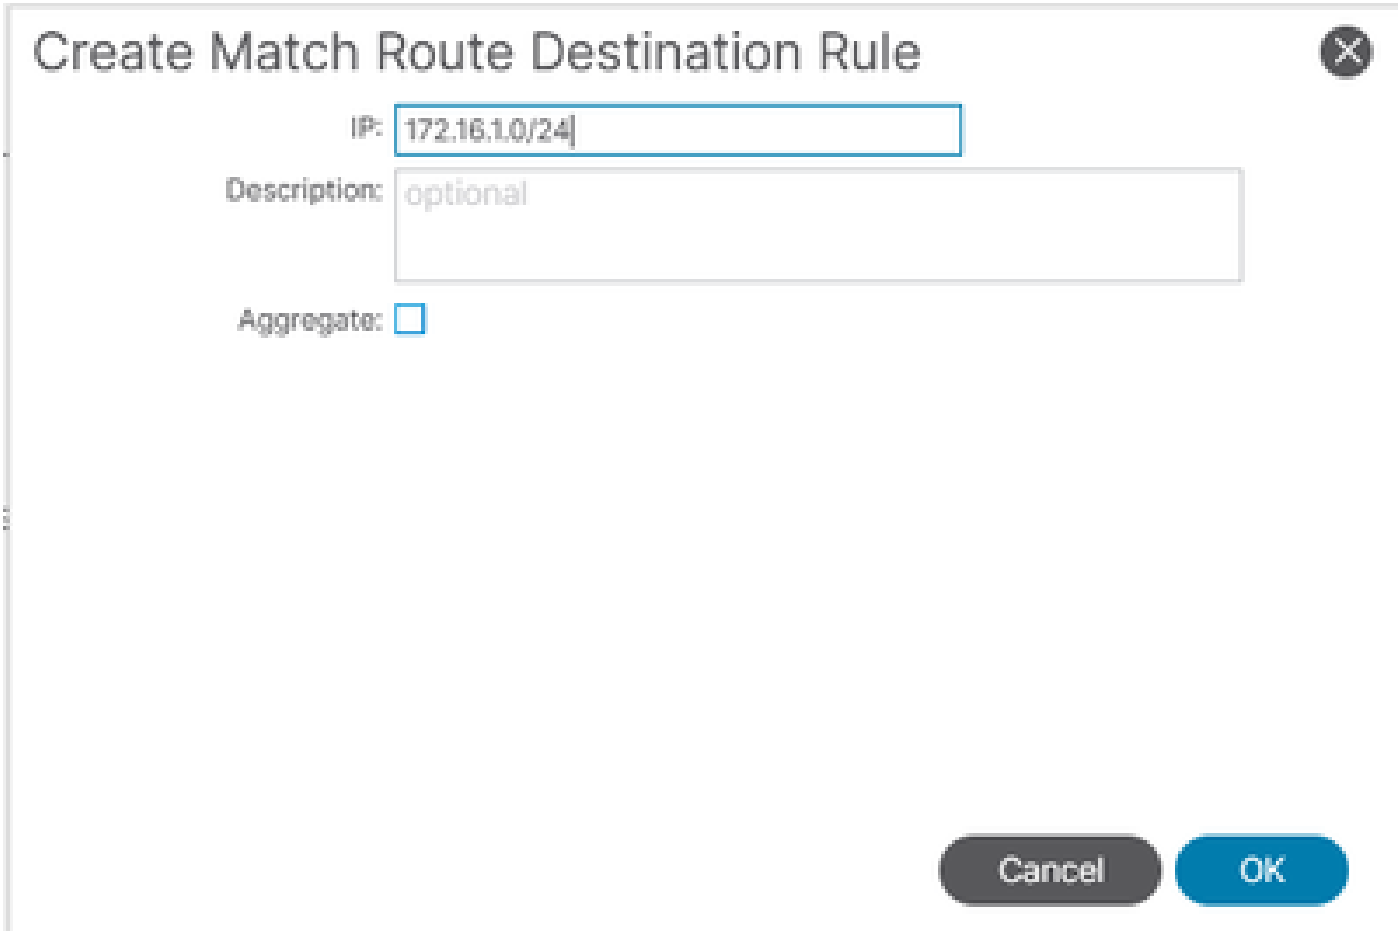


After submitting the policies, the import route-map action changes accordingly, enforcing the desired prefix filtering.

<#root>

```
Border Leaf#
```

**`show route-map imp-l3out-ExEPG-peer-2981888`**

```
route-map imp-l3out-ExEPG-peer-2981888,
```

**`deny`**

```
, sequence 8001
Match clauses:
ip address prefix-lists: IPv4-peer49155-2981888-exc-ext-in-default-import2tz0tz-dst
ipv6 address prefix-lists: IPv6-deny-all
Set clauses:

Border Leaf#
```