

ACI Remote Leaf Discovery and Configuration

Contents

[Introduction](#)

[Background Information](#)

[Prerequisites](#)

[Requirements](#)

[Components Used](#)

[Background Information](#)

[Configure](#)

[Network Diagram](#)

[Configurations](#)

[Remote WAN Configuration \(ACI Main Fabric Side\)](#)

[Remote WAN Configuration \(RLEAF side\)](#)

[ACI Configuration Step 1. Configure Pod Fabric Setup Policy](#)

[ACI Configuration Step 2. Configure Routed Outside from Spine to IPN](#)

[ACI Configuration Step 3. Discover the Remote Leaf\(s\)](#)

[ACI Configuration Step 4. Configure Routed Outside from RLEAF to IPN](#)

[ACI Configuration Step 5. Create Fabric External Connection Policy](#)

[ACI Configuration Step 6. Register Remote Leaf to ACI Fabric Pod](#)

[ACI Configuration Step 7. QoS Configuration for Remote Leaf](#)

[ACI Configuration Step 8 \(Optional\). Create Virtual Port-Channel \(vPC\) Explicit Protection Group with Remote Leaf](#)

[Verify](#)

[RLeaf TEP Review](#)

[Spine TEP Review](#)

[Remote Leaf Routable Subnets](#)

[F0467 Configuration Failed](#)

[Bootstrap Validation](#)

[Related Information](#)

Introduction

This document describes the steps to discover and configure a Remote Leaf (RLEAF) in an existing Application Centric Infrastructure (ACI) fabric with the Application Policy Infrastructure Controller (APIC) GUI.

Background Information

The ACI remote leaf switch deployment helps the customers to extend ACI fabric (ACI services and APIC management) to remote datacenters where there is no local spine switch or APIC attached. The remote leaf switches are added to an existing pod in the fabric over the Wide Area Network (WAN). All policies deployed in the main datacenter are deployed in the remote switches, which behave like local leaf switches that belong to the fabric. In the remote leaf topology, all unicast traffic is transmitted through VXLAN over Layer 3. Layer 2 Broadcast, Unknown unicast, and Multicast (BUM) traffic are sent with Head End Replication (HER) tunnels without the use of Multicast. All local traffic on the remote site is switched directly between endpoints, whether physical or virtual. Any traffic that requires the use of the spine proxy is forwarded to the main fabric. Just like a local leaf, the remote leaf can be used to connect virtual servers, physical servers, and containers. Traffic to the endpoints connected to the remote leaf is locally forwarded

through remote leaf switches.

Prerequisites

Requirements

Cisco recommends that you have knowledge of these topics:

- ACI Fabric
- ACI GUI
- ACI Leaf and Spine Switch CLI
- NXoS switch CLI

Components Used

The information in this document is based on these software and hardware versions:

- APIC Version 3.1(2q)
- Spine LC N9K-X9732C-EX
- Leaf N9K-X9732C-EX
- IP Network(IPN) - Nexus 7000, N7K-SUP2E, N7K-F348XP-25

The information in this document was created from the devices in a specific lab environment. All of the devices used in this document started with a cleared (default) configuration. If your network is live, ensure that you understand the potential impact of any command.

Background Information

The remote leaf solution is supported by ACI 3.1(1) release. The list of hardware that supports the Remote leaf solutions as of the date this document was written is listed in the table.

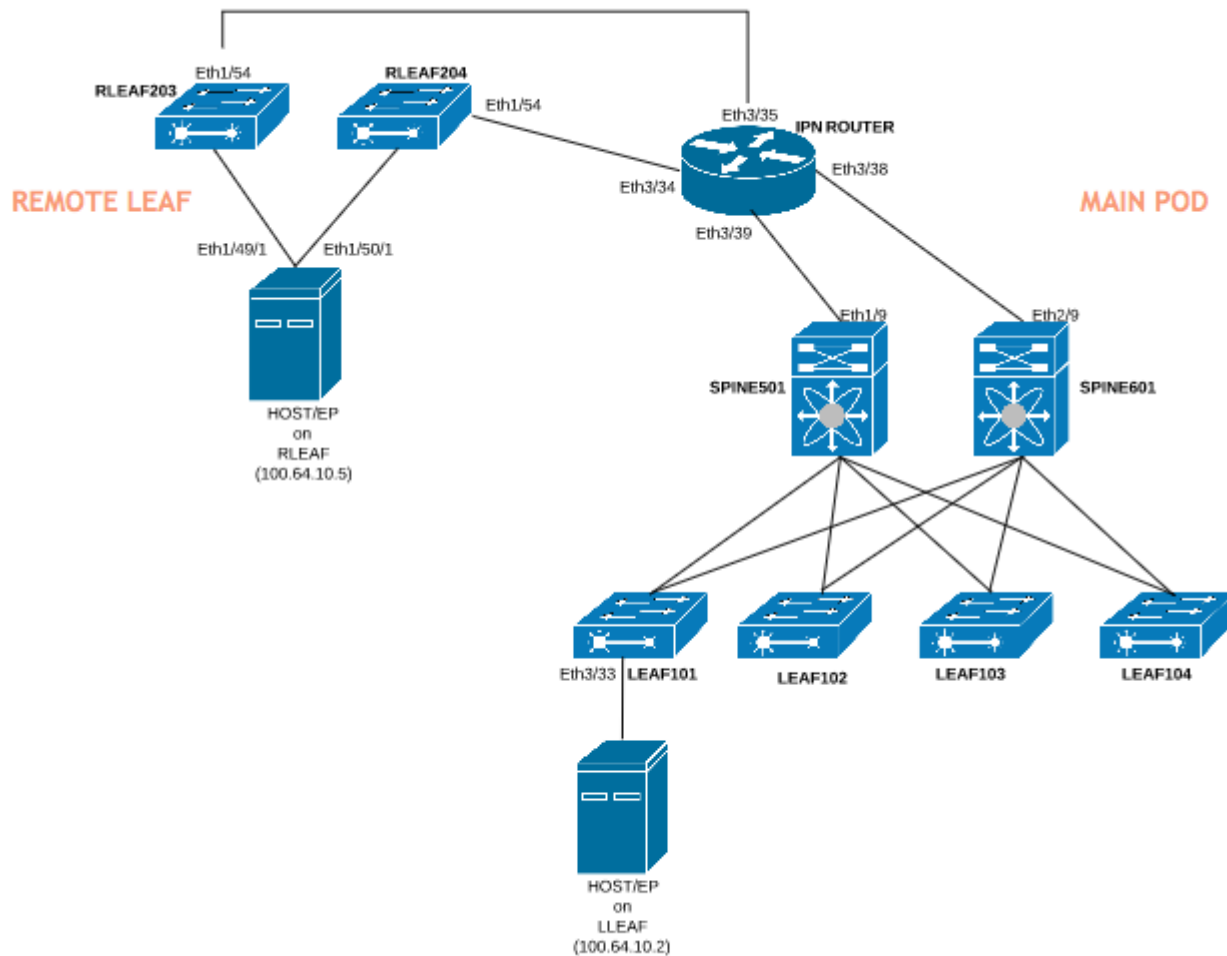
Spine/Leaf	Model
Fixed Spine	N9364C
Modular Spine LC	N9732C-EX N9736C-FX
Leaf	N93180YC-EX N93180YC-FX N93108TC-EX N93108TC-FX N93180LC-EX

	N9348GC-FXP
	N9336C-FX2

Configure

Network Diagram

This network diagram depicts the topology used in this example.



Configurations

This document mainly focuses on the ACI side of configuration for remote leaf deployment and it doesn't cover full details of the WAN side configuration of the IPN switch. However, a few important configurations from IPN are listed here for reference.

Remote WAN Configuration (ACI Main Fabric Side)

These are the configurations used in the IPN device connected to the ACI Spine(s) in the main fabric:

```
vrf context RLEAF
```

```
description VRF created for remote-leaf lab
```

```
router ospf 1

vrf RLEAF
  router-id 172.16.191.191
  area 0.0.0.1 nssa
```

In this example same IPN router is used to connect to RLEAF and SPINE

```
interface loopback191
vrf member RLEAF
ip address 172.16.191.191/32
```

Interface specific configurations on the IPN that connects to the Spine.

Towards Spine-601	Towards Spine-501
<pre>interface Ethernet3/38 mtu 9150 no shutdown interface Ethernet3/38.4 description link to spine601 mtu 9150 encapsulation dot1Q 4 vrf member RLEAF ip address 10.10.19.10/24 ip ospf network point-to-point ip router ospf 1 area 0.0.0.1 no shutdown</pre>	<pre>interface Ethernet3/39 mtu 9150 no shutdown interface Ethernet3/39.4 description link to spine501 mtu 9150 encapsulation dot1Q 4 vrf member RLEAF ip address 10.10.20.10/24 ip ospf network point-to-point ip router ospf 1 area 0.0.0.1 no shutdown</pre>

Remote WAN Configuration (RLEAF side)

These are the configurations used in the IPN device connected to the remote leaf:

```

vrf context RLEAF
description VRF created for remote-leaf lab

router ospf 1

vrf RLEAF
router-id 172.16.191.191
area 0.0.0.1 nssa

# In this example same IPN router is used to connect to RLEAF and SPINE

interface loopback191
vrf member RLEAF
ip address 172.16.191.191/32

```

Interface specific configurations on the IPN that connects to the RLEAF:

Towards RLEAF-204	Towards RLEAF-203
<pre> interface Ethernet3/34 mtu 9150 no shutdown interface Ethernet3/34.4 description link to rleaf204 mtu 915 encapsulation dot1Q 4 ip access-group 100 in vrf member RLEAF ip address 10.10.21.10/24 ip ospf network point-to-point ip router ospf 1 area 0.0.0.1 ip dhcp relay address 10.0.0.1 ip dhcp relay address 10.0.0.2 ip dhcp relay address 10.0.0.3 no shutdown </pre>	<pre> interface Ethernet3/35 mtu 9150 no shutdown interface Ethernet3/35.4 description link to rleaf203 mtu 9150 encapsulation dot1Q 4 ip access-group 100 in vrf member RLEAF ip address 10.10.22.10/24 ip ospf network point-to-point ip router ospf 1 area 0.0.0.1 ip dhcp relay address 10.0.0.1 ip dhcp relay address 10.0.0.2 ip dhcp relay address 10.0.0.3 no shutdown </pre>

Note: Ensure the dhcp-relay IP is configured with the APIC fabric IP address under the interface connected to the remote-leaf. This is required for the remote leaf to obtain the bootstrap files from APIC. In this example, 10.0.0.1, 10.0.0.2, 10.0.0.3 are the APIC TEP IP address. The remote-leaf sends DHCP DISCOVER towards the WAN to get the IP address of the subinterface connected to the WAN Router. WAN Router relays the DHCP DISCOVER message from Remote Leaf to the APICs in POD.

ACI Configuration Step 1. Configure Pod Fabric Setup Policy

1. Navigate to **Fabric > Inventory > Pod Fabric Setup Policy**.
2. Double click to open **Fabric Setup Policy** for existing Pod.
3. Add (+) Remote Pool, provide a Remote ID (in this example: 11) and Remote Pool (in this example: 172.17.0.0/20) and click **Submit**. The image snippet may show a different IP addressing scheme.

Note: The remote leaf TEP pool subnet must not overlap with the main fabric TEP pool subnet. The subnet used must be /24 or lower.

The screenshot displays the ACI GUI interface. The top navigation bar includes 'Inventory', 'Fabric Policies', and 'Access Policies'. The left sidebar shows the 'Inventory' section with a tree view containing 'Quick Start', 'Topology', 'Pod 1', 'Pod Fabric Setup Policy' (selected), 'Fabric Membership', 'Unmanaged Fabric Node', 'Unreachable Nodes', and 'Disabled Interfaces and D'. The main content area is titled 'Pod Fabric Setup Policy' and contains a table with two columns: 'Pod ID' and 'TEP Pool'. The table has one row with '1' in the 'Pod ID' column and '10.0.0.0/16' in the 'TEP Pool' column. A modal window titled 'Fabric Setup Policy For A POD - Pod 1' is open, showing the 'Properties' section with 'ID: 1' and 'TEP Pool: 10.0.0.0/16'. Below this, the 'Remote Pools' section is visible, showing a table with one row: 'Remote ID' is '11'.

ACI Configuration Step 2. Configure Routed Outside from Spine to IPN

1. Navigate to **Tenant > Infra > External Routed Networks**.
2. Right-click and create **Routed Outside**.
3. Configure OSPF Routed Outside for Spine to IPN.

4. Use OSPF as a routing protocol.
5. Use overlay-1 as VRF.

In this example, the remote leaf is attached to a single pod fabric. So, "Enable remote leaf with Multipod" is not selected. If you use remote leaf with a multipod fabric, this option must be checked.

The screenshot displays the Cisco SD-WAN configuration interface. On the left, the 'Tenant infra' tree shows the hierarchy: Tenant infra > Networking > External Routed Networks > spine2rleaf. The 'spine2rleaf' node is selected. On the right, the 'L3 Outside - spine2rleaf' configuration page is shown. The 'Properties' section includes the following settings:

- Route Control Enforcement: ☐ Import
- VRF: overlay-1 (highlighted with a red box)
- Resolved VRF: infra/overlay-1
- External Routed Domain: spine-I3
- Route Profile for Interleak: select a value
- Route Control For Dampening: ☐ Address
- Enable BGP/EIGRP/OSPF: ☐ BGP
- OSPF Area ID: 0.0.0.1
- OSPF Area Control: ☒ ☐
- ☒ Send redi
- ☒ Originate
- ☐ Suppress
- OSPF Area Type: NSSA area
- OSPF Area Cost: 1
- Enable remote leaf with Multipod: ☐

Configure the Node profile for each spine (in this example node-501 & node-601) connected to IPN. The example for node-501 is shown here. Follow the same steps for node-601. Image snippets may show different IP addressing schemes.

Tenant infra

- > Bridge Domains
- > VRFs
- > External Bridged Networks
- > External Routed Networks
 - > Route Maps/Profiles
 - > Set Rules for Route Maps
 - > Match Rules for Route Maps
 - > rleaf-wan
- > spine2rleaf
 - > Logical Node Profiles
 - > spine501
 - > Logical Interface Profiles
 - > spine501
 - > Configured Nodes
 - topology/pod-1/node-501
 - > spine601
 - > Logical Interface Profiles
 - > spine601
 - > Configured Nodes
 - topology/pod-1/node-601
 - > Networks
 - > rleaf
 - > Route Maps/Profiles



Node Association



Properties

Node ID: topology/pod-1/node-501

Router ID: 50.50.50.50

Use Router ID as Loopback Address: ☒

This setting will be ignored if loopback

Loopback Addresses:

▲ IP

Static Routes:

▲ IP Address

< < Page 0 Of 0

External Control Peering: ☒

Configure interface profile for IPN connection.

Note: Ensure to use encap vlan-4 for remote leaf integration with a single pod.

Example for path-501/pathep-[eth1/9].

Tenant infra

Quick Start

Tenant infra

Application Profiles

Networking

Bridge Domains

VRFs

External Bridged Networks

External Routed Networks

Route Maps/Profiles

Set Rules for Route Maps

Match Rules for Route Map

rlaef-wan

spine2rleaf

Logical Node Profiles

spine501

Logical Interface

spine501

Configured Node

spine601

Logical Interface

Configured Node

Networks

Route Maps/Profiles

Dot1Q Tunnels

Contracts

Logical Interface Profile - spine501

Routed Sub-Interface

Properties

Target: topology/pod-1/paths-501/pathep-[eth1/9]

Description: optional

Encap: VLAN 4 Integer Value

IPv4 Primary / IPv6 Preferred Address: 10.10.20.11/24 address/mask

IPv6 DAD: disabled enabled

IPv4 Secondary / IPv6 Additional Addresses:

Address IPv6 DAD

No items have been found. Select Actions to create a new item.

Link-Local Address: ::

ND RA Prefix: ☐

MAC Address: 00:22:BD:F8:19:FF

MTU (bytes): inherit

Target DSCP: Unspecified

Example for path-601/pathep-[eth2/9].

Tenant infra

Quick Start

Tenant infra

Application Profiles

Networking

Bridge Domains

VRFs

External Bridged Networks

External Routed Networks

Route Maps/Profiles

Set Rules for Route Maps

Match Rules for Route Maps

leaf-wan

spine2rleaf

Logical Node Profiles

spine501

Logical Interface

spine501

Configured Node

spine601

Logical Interface

spine601

Configured Node

Networks

Route Maps/Profiles

Dot1Q Tunnels

Logical Interface Profile - spine601

Routed Sub-Interface

Properties

Target: topology/pod-1/paths-601/pathep-[eth2/9]

Description: optional

Encap: VLAN 4

Integer Value

IPv4 Primary / IPv6 Preferred Address: 10.10.19.11/24

address/mask

IPv6 DAD: disabled enabled

IPv4 Secondary / IPv6 Additional Addresses:

Address IPv6 DAD

No items have been found.
Select Actions to create a new item.

Link-Local Address: ::

ND RA Prefix: ☐

Configure L3Out Network(External EPG) for IPN.

Tenant infra

- > Bridge Domains
- > VRFs
- > External Bridged Networks
- > External Routed Networks
 - > Route Maps/Profiles
 - > Set Rules for Route Maps
 - > Match Rules for Route Maps
 - > rleaf-wan
 - > spine2rleaf
 - > Logical Node Profiles
 - > spine501
 - > Logical Interface Profiles
 - > spine501
 - > Configured Nodes
 - topology/pod-1/node-501
 - > spine601
 - > Logical Interface Profiles
 - > spine601
 - > Configured Nodes
 - topology/pod-1/node-601
 - > Networks
 - > rleaf



External Network Instance

✖
⚠
✔

Properties

QoS Class:

Target DSCP:

Configuration Status: applied

Configuration Issues:

Preferred Group Member: Exclude Include

L3Out Contract Masters:

L3Out Contract

Subnets:

▲ IP Address
 0.0.0.0/0

Now that you've configured OSPF L3Out from Spine (node-501 & node-601) to the IPN device. The OSPF adjacency should be up and routes are exchanged provided the OSPF on the IPN is configured correctly. So, check the OSPF neighbourship from Spine to IPN device.

From Spine:

```
spine501# show ip ospf neighbors vrf overlay-1
OSPF Process ID default VRF overlay-1
Total number of neighbors: 1
Neighbor ID    Pri State           Up Time  Address      Interface
172.16.191.191 1 FULL/-           00:00:36 10.10.20.10  Eth1/9.9
spine501#
```

```
spine601# show ip ospf neighbors vrf overlay-1
OSPF Process ID default VRF overlay-1
Total number of neighbors: 1
Neighbor ID    Pri State           Up Time  Address      Interface
172.16.191.191 1 FULL/-           00:00:39 10.10.19.10  Eth2/9.9
```

```
spine601#
```

From IPN:

```
SPINE-IPN# show ip ospf neighbors vrf RLEAF
OSPF Process ID 1 VRF RLEAF
Total number of neighbors: 2
Neighbor ID      Pri State           Up Time  Address      Interface
172.16.60.60     1 FULL/ -         00:00:06 10.10.19.11  Eth3/38.4
172.16.50.50     1 FULL/ -         00:00:17 10.10.20.11  Eth3/39.4
SPINE-IPN#
```

Now you have OSPF neighborship between Spine and IPN and you can see that the route to ACI fabric pod infra network is learned at IPN over the OSPF.

```
SPINE-IPN# show ip route vrf RLEAF
IP Route Table for VRF "RLEAF"
'*' denotes best ucast next-hop
'**' denotes best mcast next-hop
'[x/y]' denotes [preference/metric]
 '%' in via output denotes VRF

10.0.0.0/16, ubest/mbest: 2/0
  *via 10.10.19.11, Eth3/38.4, [110/20], 00:01:21, ospf-1, nssa type-2
  *via 10.10.20.11, Eth3/39.4, [110/20], 00:01:21, ospf-1, nssa type-2
< snip >
SPINE-IPN#
```

ACI Configuration Step 3. Discover the Remote Leaf(s)

At this stage, the fabric is ready to discover a remote leaf connected to IPN across the WAN. Ensure that the IPN connected to the RLEAF has the route to the ACI pod infra network over the WAN network.

```
<#root>
```

```
RLEAF-IPN#
```

```
show lldp neighbors
```

Capability codes:

(R) Router, (B) Bridge, (T) Telephone, (C) DOCSIS Cable Device

(W) WLAN Access Point, (P) Repeater, (S) Station, (O) Other

Device ID	Local Intf	Hold-time	Capability	Port ID
switch	Eth3/34	120	BR	Eth1/54
switch	Eth3/35	120	BR	Eth1/54

Total entries displayed: 2

```
RLEAF-IPN#
```

RLEAF-IPN#

```
show ip route vrf RLEAF
```

IP Route Table for VRF "RLEAF"

'*' denotes best ucast next-hop

'**' denotes best mcast next-hop

'[x/y]' denotes [preference/metric]

'%<string>' in via output denotes VRF <string>

10.0.0.0/16, ubest/mbest: 2/0

*via 10.10.19.11, Eth3/38.4, [110/20], 00:01:21, ospf-1, nssa type-2

*via 10.10.20.11, Eth3/39.4, [110/20], 00:01:21, ospf-1, nssa type-2

< snip >

Review the IPN connected to the remote leaf and ensure that the APIC fabric IP address is configured as DHCP relay.

<#root>

RLEAF-IPN#

```
show ip dhcp relay
```

< snip >

Helper addresses are configured on the following interfaces:

Interface	Relay Address	VRF Name
-----	-----	-----
Ethernet3/34.4	10.0.0.1	
Ethernet3/34.4	10.0.0.2	
Ethernet3/34.4	10.0.0.3	
Ethernet3/35.4	10.0.0.1	
Ethernet3/35.4	10.0.0.2	
Ethernet3/35.4	10.0.0.3	

RLEAF-IPN#

Navigate to **Inventory > Fabric Membership** in ACI GUI and check for the new switches discovered.

Inventory

- Quick Start
- Topology
- Pod 1
 - leaf101 (Node-101)
 - leaf102 (Node-102)
 - leaf103 (Node-103)
 - leaf104 (Node-104)
 - leaf201 (Node-201)
 - leaf202 (Node-202)
 - spine501 (Node-501)
 - spine601 (Node-601)
- Pod Fabric Setup Policy
- Fabric Membership**
- Unmanaged Fabric Nodes
- Unreachable Nodes
- Disabled Interfaces and Decommissioned Switches

Fabric Membership

Serial Number	Pod ID	Node ID	RL TEP Pool	Node Name	Rack Name	Model
FDO20331BFQ	1	202	0	leaf202		N9K-
FDO21031WXP	1	201	0	leaf201		N9K-
FDO220810B0	1	0	0			N9K-
FOX1948G9EA	1	601	0	spine601		N9K-
FOX1949GHHM	1	501	0	spine501		N9K-
SAL1946SWJM	1	101	0	leaf101		N9K-
SAL1946SWNS	1	102	0	leaf102		N9K-
SAL1946SWNT	1	104	0	leaf104		N9K-
SAL1946SWNU	1	103	0	leaf103		N9K-
FDO22080JDA	1	0	0			N9K-

Register the newly discovered leaf to the existing fabric:

1. Identify the new leaf based on the serial number.
2. Right-click on the newly discovered leaf and click **Register**.
3. Provide the right **Pod ID** and **Node ID**.
4. Select the **RL TEP POOL**.
5. Provide a **Node Name**.
6. Check and Confirm the **Role** is selected as remote leaf.
7. Click **Update**.

Serial Number	Pod ID	Node ID	RL TEP Pool	Node Name	Rack Name	Model
FDO22080JDA	1	203	11	rleaf203	select	N9K-C9318
FDO220810B0	1	204	11	rleaf204	select	N9K-C9318
FOX1948G9EA	1	601	0	spine601		

Update

Cancel

Note: Ensure to select the correct **RL TEP Pool** you configured in Step 1. Also, check and confirm the **Role** is selected as a remote leaf automatically when you select the **RL TEP POOL** from the dropdown.

Now you can see the node type is identified as "remote leaf" and status as "Discovering". The node hasn't got a fabric IP address yet.

Inventory

- > Quick Start
- Topology
- Pod 1
 - > leaf101 (Node-101)
 - > leaf102 (Node-102)
 - > leaf103 (Node-103)
 - > leaf104 (Node-104)
 - > leaf201 (Node-201)
 - > leaf202 (Node-202)
 - > spine501 (Node-501)
 - > spine601 (Node-601)
- Pod Fabric Setup Policy
- > Fabric Membership
- > Unmanaged Fabric Nodes
- Unreachable Nodes
- Disabled Interfaces and Decommissioned Switches

Fabric Membership

Serial Number	Pod ID	Node ID	RL TEP I
FDO20331BFQ	1	202	0
FDO21031WXP	1	201	0
FDO22080JDA	1	203	11
FDO220810B0	1	204	11
FOX1948G9EA	1	601	0
FOX1949GHHM	1	501	0
SAL1946SWJM	1	101	0
SAL1946SWNS	1	102	0
SAL1946SWNT	1	104	0
SAL1946SWNU	1	103	0

ACI Configuration Step 4. Configure Routed Outside from RLEAF to IPN

1. Navigate to **Tenant > Infra > External Routed Networks** and create **Routed Outside**.

Create Routed Outside

STEP 1 > Identity

Define the Routed Outside

Description: optional

Tags:

enter tags separated by comma

PIM: ☐

Route Control Enforcement: ☐ Import ☒ Export

Target DSCP: Unspecified

VRF: overlay-1

External Routed Domain: rleaf-l3

Route Profile for Interleaf: select a value

Route Control For Dampening:

Address Family Type

Con...

☐ E

OSPF Area ID:

OSPF Area

Control:

OSPF Area Type:

OSPF Area Cost:

Enable Remote

Leaf:

Address F

Nodes and Interfaces Protocol Profiles

Name	Description	DSCP
rleaf-node-profile		Unspecified

2. Create RLEAF node profiles for node-203 and 204.
Now, you can select the **rleaf-203 (Node-203)** and **rleaf-204(Node-204)** from the **Node ID** dropdown list as they are now discovered in the fabric.

RLEAF 203 Node Profile:

Create Routed Outside

Select Node

Select Node and Configure Static Routes

Node ID: rleaf203 (Node-203)



Router ID: 203.203.203.203

Use Router ID as Loopback Address: ☒

External Control Peering: ☒

Loopback Addresses:

IP

Static Routes:

IP Address

Node: topology/pod-1/node-203

Path: topology/pod-1/paths-203/pathep-[eth1/54]

Select Routed Sub-Interface

Specify the Interface

Node:
Ex: topology/pod-1/node-1

Path:
Ex: topology/pod-1/paths-101/p

Description:

Encap:

VLAN

▼

Integer Value

IPv4 Primary / IPv6 Preferred Address:
address/mask

IPv6 DAD:

disabled

enabled

IPv4 Secondary / IPv6 Additional
Addresses:

Address

MAC Address:

MTU (bytes):

Link-local Address:

Create the interface profile for node-204. Manually enter Node and Path fields as shown.

Node: topology/pod-1/node-204

Path: topology/pod-1/paths-204/pathep-[eth1/54]

Select Routed Sub-Interface

Specify the Interface

Node: topology/pod-1/node-203

Ex: topology/pod-1/node-1

Path: topology/pod-1/paths-203/path

Ex: topology/pod-1/paths-101/path

Description: optional

Encap: VLAN

4

Integer Value

IPv4 Primary / IPv6 Preferred Address: 10.10.22.11/24

address/mask

IPv6 DAD:

disabled

enabled

IPv4 Secondary / IPv6 Additional
Addresses:

Address

MAC Address: 00:22:BD:F8:19:FF

MTU (bytes): inherit

Link-local Address:

Create the L3Out External Network (External EPG):

Create Routed Outside

STEP 2 > External EPG Networks

Configure External EPG Networks

Create Route Profiles: ☐

External EPG Networks

Name	QoS Class	Description	Target DS
rleaf-l3out	Unspecified		Unspecified

ACI Configuration Step 5. Create Fabric External Connection Policy

1. Navigate to **Tenant > Infra > Policies > Protocol > Fabric Ext Connection Policy > Fabric External Connection Policy** and create **Intrasite/Intersite Profile**.
2. Add **Fabric External Routing Profile** with an external network of RLEAF203 and RLEAF204 connected to the WAN router (IPN).
3. In this case, those are 10.10.22.0/24 and 10.10.21.0/24 respectively.

Tenant infra

Quick Start

▼ Tenant infra

> Application Profiles

> Networking

> Contracts

▼ Policies

▼ Protocol

> Route Maps

> BFD

> BGP

> OSPF

> EIGRP

> IGMP Snoop

> IGMP Interface

> Custom QOS

> End Point Retention

> DHCP

> ND Interface

> ND RA Prefix

> Route Tag

> L4-L7 Policy Based Redirect

> L4-L7 Redirect Health Groups

> Data Plane Policing

▼ Fabric Ext Connection Policies

Fabric Ext Connection Policy rleaf



Intrasite/Intersite Profile



Properties

Fabric ID: 1

Name: rleaf

Community: extend
Ex: extend

Pod Peering Profile

Peering Type: Full M

Password:

Confirm Password:

Pod Connection Profile



▲ Pod ID

Fabric External Routing P



Name

rleaf203

rleaf204

ACI Configuration Step 6. Register Remote Leaf to ACI Fabric Pod

Then the remote leaf obtains the fabric IP address from the APIC TEP pool.

Inventory | Fabric Policies | Access Policies

Inventory

Quick Start

Topology

Pod 1

Pod Fabric Setup Policy

Fabric Membership

Unmanaged Fabric Nodes

Unreachable Nodes

Fabric Membership

Serial Number	Pod ID	Node ID	RL TEP Pool	Node Name
FDO20331BFQ	1	202	0	leaf202
FDO21031WXP	1	201	0	leaf201
FDO22080JDA	1	203	11	rleaf203
FDO220810B0	1	204	11	rleaf204

Wait for a while to get the remote-leaf to active status. Now, you can see that the remote leaf rleaf-203 and rleaf-204 are registered to ACI fabric.

Inventory

Quick Start

Topology

Pod 1

Pod Fabric Setup Policy

Fabric Membership

Unmanaged Fabric Nodes

Unreachable Nodes

Disabled Interfaces and Decommissioned Switches

Fabric Membership

Serial Number	Pod ID	Node ID	RL TEP Pool
FDO20331BFQ	1	202	0
FDO21031WXP	1	201	0
FDO22080JDA	1	203	11
FDO220810B0	1	204	11
FOX1948G9EA	1	601	0

<#root>

apic3#

acidiag fmvread

ID	Pod ID	Name	Serial Number	IP Address	Role	State	LastUp
101	1	leaf101	SAL1946SWJM	10.0.232.64/32	leaf	active	0
102	1	leaf102	SAL1946SWNS	10.0.232.73/32	leaf	active	0
103	1	leaf103	SAL1946SWNU	10.0.232.69/32	leaf	active	0
104	1	leaf104	SAL1946SWNT	10.0.8.64/32	leaf	active	0
201	1	leaf201	FD021031WXP	10.0.232.72/32	leaf	active	0
202	1	leaf202	FD020331BFQ	10.0.232.68/32	leaf	active	0
203	1	rleaf203	FDO22080JDA	172.17.3.128/32	leaf	active	0

204	1	rleaf204	FDO220810B0	172.17.7.128/32	leaf	active	0
501	1	spine501	FOX1949GHHM	10.0.232.66/32	spine	active	0
601	1	spine601	FOX1948G9EA	10.0.232.65/32	spine	active	0

Total 10 nodes

apic3#

Now, you can see the OSPF neighbourhood between remote leaf and IPN.

From RLEAF 203:

```
<#root>
```

```
rleaf203#
```

```
show ip ospf neighbors vrf overlay-1
```

```

OSPF Process ID default VRF overlay-1
Total number of neighbors: 1
Neighbor ID      Pri State           Up Time  Address      Interface
172.16.191.191  1 FULL/ -         00:24:57  10.10.22.10  Eth1/54.6
rleaf203#

```

```
rleaf203#
```

```
show ip route vrf overlay-1
```

```

IP Route Table for VRF "overlay-1"
'*' denotes best ucast next-hop
 '**' denotes best mcast next-hop
 '[x/y]' denotes [preference/metric]
 '%<string>' in via output denotes VRF <string>

```

```

10.0.0.0/16, ubest/mbest: 1/0
    *via 10.10.22.10, eth1/54.6, [110/20], 00:30:24, ospf-default, nssa type-2
10.0.0.1/32, ubest/mbest: 1/0
    *via 10.10.22.10, eth1/54.6, [110/20], 00:30:24, ospf-default, nssa type-2

```

```
< snip >
```

From RLEAF 204:

```
<#root>
```

```
rleaf204#
```

```
show ip ospf neighbors vrf overlay-1
```

```

OSPF Process ID default VRF overlay-1
Total number of neighbors: 1

```

Neighbor ID	Pri	State	Up Time	Address	Interface
172.16.191.191	1	FULL/ -	00:25:36	10.10.21.10	Eth1/54.6

rleaf204#

show ip route vrf overlay-1

IP Route Table for VRF "overlay-1"

'*' denotes best ucast next-hop

'**' denotes best mcast next-hop

'[x/y]' denotes [preference/metric]

'%<string>' in via output denotes VRF <string>

10.0.0.0/16, ubest/mbest: 1/0

*via 10.10.21.10, eth1/54.6, [110/20], 00:31:37, ospf-default, nssa type-2

10.0.0.1/32, ubest/mbest: 1/0

*via 10.10.21.10, eth1/54.6, [110/20], 00:31:37, ospf-default, nssa type-2

< snip >

From IPN:

<#root>

RLEAF-IPN#

show ip ospf neighbors vrf RLEAF

OSPF Process ID 1 VRF RLEAF

Total number of neighbors: 4

Neighbor ID	Pri	State	Up Time	Address	Interface
172.16.204.204	1	FULL/ -	00:26:03	10.10.21.11	Eth3/34.4
172.16.203.203	1	FULL/ -	00:26:03	10.10.22.11	Eth3/35.4

RLEAF-IPN#

ACI Configuration Step 7. QoS Configuration for Remote Leaf

It is required to classify ACI fabric classes (QoS Levels) to a DSCP value within IPN. To achieve this requirement, ACI Fabric should be enabled with **DSCP class-cos translation policy for L3 traffic**. Use this configuration to map ACI QoS levels and default classes to DSCP values in IPN.

Navigate to **Tenant > Infra > Policies > DSCP class-cos translation policy for L3 traffic** as shown in the image.

DSCP class-cos translation policy for L3 traffic

Properties

Translation Policy State:	Disabled	Enabled
User Level 1:	CS0	▼
User Level 2:	CS1	▼
User Level 3:	CS2	▼
Control Plane Traffic:	CS3	▼
Policy Plane Traffic:	CS4	▼
Span Traffic:	CS5	▼
Traceroute Traffic:	CS6	▼

ACI Configuration Step 8 (Optional). Create Virtual Port-Channel (vPC) Explicit Protection Group with Remote Leaf

Since the remote leaf switches are registered to ACI fabric, you can create vPC Explicit Protection Group with a remote leaf. Navigate to **Fabric > Access Policies > Switch Policies > Policies > Virtual Port Channel Default** and create **Explicit VPC Protection Groups (+)**. Image snippets show different IP addressing schemes.

VPC Explicit Protection Group - VPC Protection Group



Properties

Name: rleaf-vpc

Logical Pair ID: 234

VPC Domain Policy: default

Virtual IP: 11.0.3.130/32

Switch Pairs: Node ID

203
204

<#root>

rleaf203#

show system internal epm vpc

Local TEP IP : 172.17.3.128

Peer TEP IP : 172.17.7.129

vPC configured : Yes

vPC VIP : 172.17.3.130

MCT link status : Up

Local vPC version bitmap : 0x7

Peer vPC version bitmap : 0x7

Negotiated vPC version : 3

```
Peer advertisement received    : Yes
Tunnel to vPC peer            : Up

vPC# 343
if : port-channel1, if index : 0x16000000
local vPC state : MCEC_STATE_UP, peer vPC state : MCEC_STATE_UP
current link state : LOCAL_UP_PEER_UP
vPC fast conv : Off
rleaf203#
```

Verify

RLeaf TEP Review

Remote leaf Data-Plane Tunnel End Point (RL-DP-PTEP) - This IP address is assigned to each Remote leaf switch from the TEP pool that is allocated to the Remote Location. VXLAN packets from a Remote leaf node are originated with this TEP as a source IP address when the Remote leaf nodes are not part of a vPC domain.

Remote Leaf vPC Tunnel End Point (RL-vPC) - This is an anycast IP address assigned to the vPC pair of Remote leaf nodes from the TEP pool that is allocated to the Remote Location. All the VXLAN packets sourced from both Remote leaf switches originated from this TEP address if the Remote leaf switches are part of a vPC domain.

```
<#root>
```

```
rleaf203#
```

```
show ip int vrf overlay-1
```

```
IP Interface Status for VRF "overlay-1"
eth1/54.6, Interface status: protocol-up/link-up/admin-up, iod: 64, mode: external
  IP address: 10.10.22.11, IP subnet: 10.10.22.0/24
  IP broadcast address: 255.255.255.255
  IP primary address route-preference: 1, tag: 0
lo0, Interface status: protocol-up/link-up/admin-up, iod: 4, mode: ptep
  IP address: 172.17.3.128, IP subnet: 172.17.3.128/32
  IP broadcast address: 255.255.255.255
  IP primary address route-preference: 1, tag: 0
lo1, Interface status: protocol-up/link-up/admin-up, iod: 65, mode: unspecified
  IP address: 172.16.203.203, IP subnet: 172.16.203.203/32
  IP broadcast address: 255.255.255.255
  IP primary address route-preference: 1, tag: 0
lo2, Interface status: protocol-up/link-up/admin-up, iod: 72, mode: vpc

  IP address: 172.17.3.130, IP subnet: 172.17.3.130/32

  IP broadcast address: 255.255.255.255
```

IP primary address route-preference: 1, tag: 0

lo3, Interface status: protocol-up/link-up/admin-up, iod: 75, mode: dp-ptep

IP address: 172.17.3.129, IP subnet: 172.17.3.129/32

IP broadcast address: 255.255.255.255

IP primary address route-preference: 1, tag: 0

lo1023, Interface status: protocol-up/link-up/admin-up, iod: 66, mode: ftep

IP address: 172.17.0.32, IP subnet: 172.17.0.32/32

IP broadcast address: 255.255.255.255

IP primary address route-preference: 1, tag: 0

rleaf203#

<#root>

rleaf204#

show ip int vrf overlay-1

IP Interface Status for VRF "overlay-1"

eth1/54.6, Interface status: protocol-up/link-up/admin-up, iod: 64, mode: external

IP address: 10.10.21.11, IP subnet: 10.10.21.0/24

IP broadcast address: 255.255.255.255

IP primary address route-preference: 1, tag: 0

lo0, Interface status: protocol-up/link-up/admin-up, iod: 4, mode: ptep

IP address: 172.17.7.128, IP subnet: 172.17.7.128/32

IP broadcast address: 255.255.255.255

IP primary address route-preference: 1, tag: 0

lo1, Interface status: protocol-up/link-up/admin-up, iod: 65, mode: unspecified

IP address: 172.16.204.204, IP subnet: 172.16.204.204/32

IP broadcast address: 255.255.255.255

IP primary address route-preference: 1, tag: 0

lo2, Interface status: protocol-up/link-up/admin-up, iod: 71, mode: dp-ptep

IP address: 172.17.7.129, IP subnet: 172.17.7.129/32

IP broadcast address: 255.255.255.255

IP primary address route-preference: 1, tag: 0

lo9, Interface status: protocol-up/link-up/admin-up, iod: 81, mode: vpc

```
IP address: 172.17.3.130, IP subnet: 172.17.3.130/32
```

```
IP broadcast address: 255.255.255.255
```

```
IP primary address route-preference: 1, tag: 0
```

```
lo1023, Interface status: protocol-up/link-up/admin-up, iod: 66, mode: ftep
```

```
IP address: 172.17.0.32, IP subnet: 172.17.0.32/32
```

```
IP broadcast address: 255.255.255.255
```

```
IP primary address route-preference: 1, tag: 0
```

```
rleaf204#
```

Spine TEP Review

Remote leaf Unicast Tunnel End Point (RL-UCAST) – This is an anycast IP address part of the local TEP pool automatically assigned to all the spines to which the Remote leaf switches are being associated. When unicast packets are sent from endpoints connected to the RLEAF nodes to the ACI main Pod, VXLAN encapsulated packets are sent with the destination as RL-Ucast-TEP address and source as RL-DP-TEP or RL-vPC. Any Spine in the ACI main DC Pod can hence receive the traffic, decapsulate it, perform the required L2 or L3 lookup and finally re-encapsulate it and forward it to the final destination.

Remote leaf ~~Unicast~~-Multicast Tunnel End Point (RL-MCAST-HREP) - This is another anycast IP address part of the local TEP pool automatically assigned to all the spines to which the Remote leaf switches are being associated. When BUM (Layer 2 Broadcast, Unknown Unicast, or Multicast) traffic is generated by an endpoint connected to the Remote leaf nodes, packets are VXLAN encapsulated by the RLEAF node and sent with the destination as RL-Mcast-TEP address and source as RL-DP-TEP or RL-vPC. Any of the spines in the ACI Pod can receive the BUM traffic and forward it inside the fabric.

```
<#root>
```

```
spine501#
```

```
show ip int vrf overlay-1
```

```
< snip >
```

```
lo12, Interface status: protocol-up/link-up/admin-up, iod: 88, mode: rl-mcast-hrep
```

```
IP address: 10.0.0.37, IP subnet: 10.0.0.37/32
```

```
IP broadcast address: 255.255.255.255
```

```
IP primary address route-preference: 1, tag: 0
```

lo13, Interface status: protocol-up/link-up/admin-up, iod: 91, mode: rl-ucast

IP address: 10.0.0.36, IP subnet: 10.0.0.36/32

IP broadcast address: 255.255.255.255

IP primary address route-preference: 1, tag: 0

spine501#

<#root>

spine601#

show ip int vrf overlay-1

< snip >

lo11, Interface status: protocol-up/link-up/admin-up, iod: 76, mode: rl-mcast-hrep

IP address: 10.0.0.37, IP subnet: 10.0.0.37/32

IP broadcast address: 255.255.255.255

IP primary address route-preference: 1, tag: 0

lo12, Interface status: protocol-up/link-up/admin-up, iod: 79, mode: rl-ucast

IP address: 10.0.0.36, IP subnet: 10.0.0.36/32

IP broadcast address: 255.255.255.255

IP primary address route-preference: 1, tag: 0

spine601#

Remote Leaf Routable Subnets

If the APICs are reachable via the IPN/ISN/WAN, then the remote subnet is not required for remote leaf nodes to be discovered. This subnet is used to create NAT entries for APICs on the spine switches and to support vPOD. This feature can also be used in conjunction with RL Direct.

F0467 Configuration Failed

If you use the wizard to add a remote leaf to a pod, the wizard does not configure the required access policies listed here and you see the popular **F0467** fault message. You need to manually create them.

- Leaf Switch Profile for Remote Leaf nodes and Leaf Selectors
- Leaf Interface Profile for Remote Leaf nodes and Interface Selector for the uplinks
- Access Interface Policy group

F0467 Fault delegate: Configuration failed for uni/tn-infra/out-rl-infra.l3out/instP-ipnInstP node 203 e
topology/pod-1/node-203/local/svc-policyelem-id-0/uni/epp/rtd-[uni/tn-infra/out-rl-infra.l3out/instP-ipr

Bootstrap Validation

Remote leaves can fail to be discovered due to enabling Bootscript Validation.

Basically with the Bootscript Validation enabled the leaf comes up and learns about the version it should be running via DHCP and it should then download the image from the APICS. The problem though is that for remote with RL direct enabled to work there are specific tcam rules that must be installed on the leaf to allow the natted APIC traffic. Because the bootscript validation is failing these rules and objects aren't getting installed on the RL. However, because those objects/rules aren't installed it can't successfully download the image from the APIC.

If you encounter a situation like this, try to discover the RL with BSV turned off.

Related Information

- <https://www.cisco.com/c/en/us/solutions/collateral/data-center-virtualization/application-centric-infrastructure/white-paper-c11-740861.html>
- [Technical Support & Documentation - Cisco Systems](#)