

Configure IP SLA Feature with L3out to Track Static Route

Contents

[Introduction](#)

[Prerequisites](#)

[Requirements](#)

[Components Used](#)

[Background Information](#)

[Configure](#)

[Network Diagram](#)

[Configurations](#)

[Verify](#)

[Troubleshoot](#)

[Related Information](#)

Introduction

This document describes how to configure the Internet Protocol Service Level Agreement (IPSLA) in Cisco Application Centric Infrastructure (ACI) to track static route learn from one L3out and advertise to another L3out only if the subnet is reachable from the first L3out.

Prerequisites

Requirements

Cisco recommends that you have knowledge of these topics:

- ACI software version 4.1 and later
- L3out toward External device or Server
- EX and -FX chassis
- Track the route to use Internet Control Message Protocol (ICMP) and TCP probes (in this example ICMP probe is used)

Note: ACI image IP SLA is supported in all Cisco Nexus second-generation switches, which includes -EX and -FX chassis. Please read [Guidelines and Limitations for IP SLA](#).

Components Used

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

- ACI version 5.2(2f)
- N9K-C93180YC-FX

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.

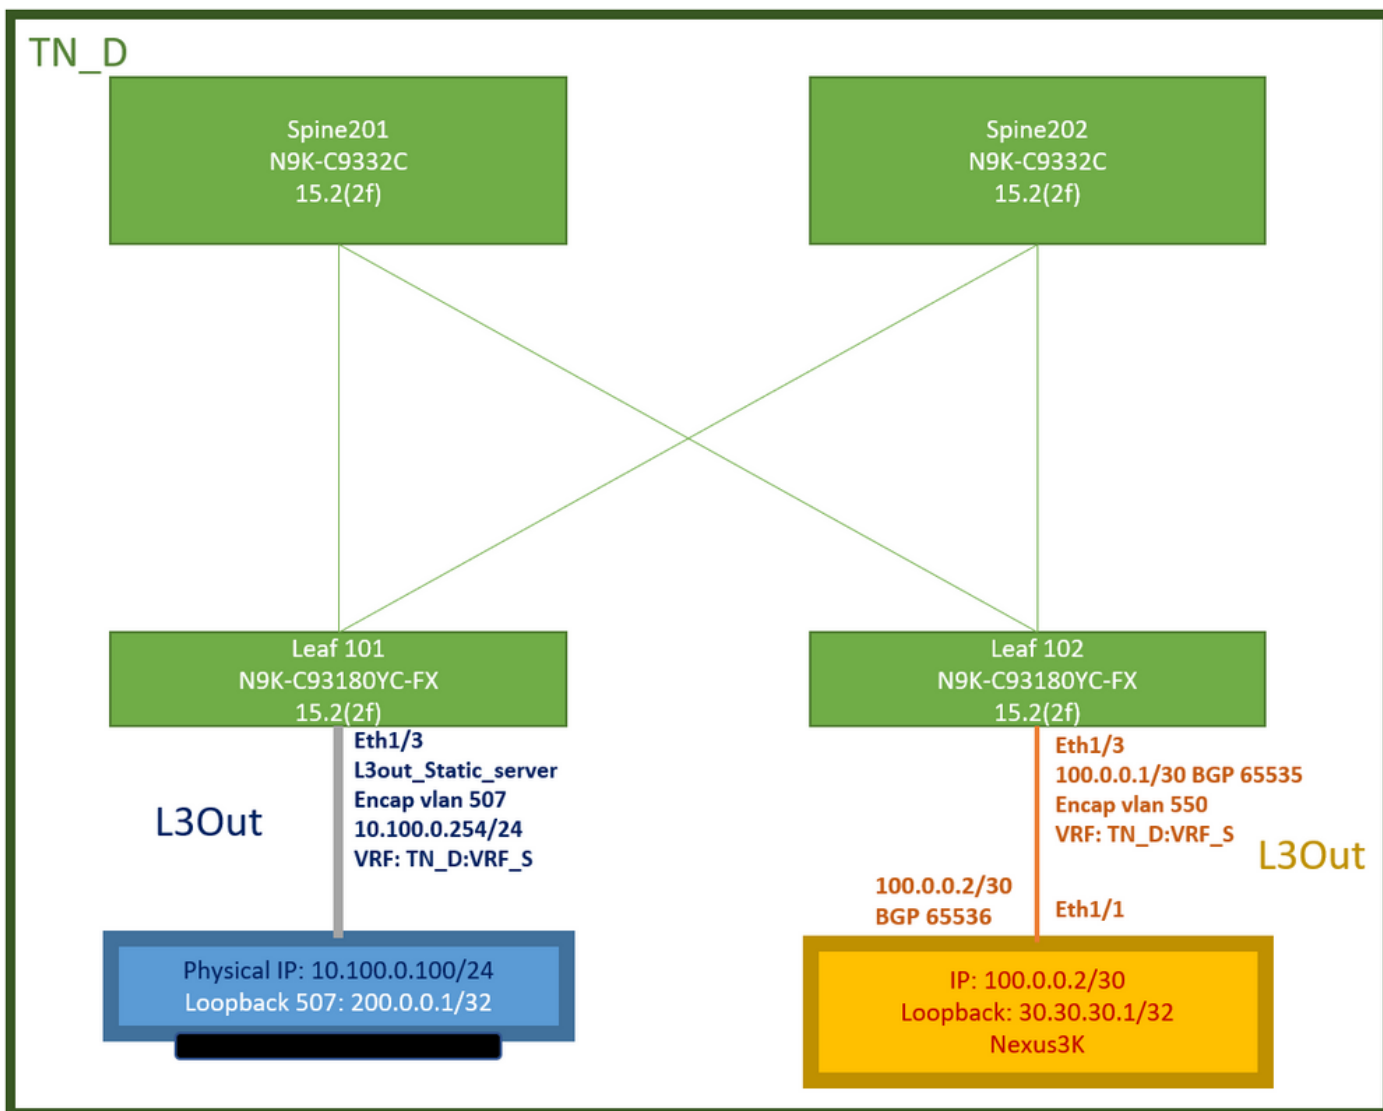
Background Information

Some servers have multiple interfaces (like a loopback) which is reachable from ACI via the physical IP address of the server. In such a case you can have a requirement to add a static route and advertise externally but only if the physical IP of the server is reachable. Hence IP SLA track feature is an unavoidable configuration that can only achieve through L3out configuration toward those servers. At this moment IP SLA track features are not supported for the [static route on a Bridge Domain](#). In this document, we will look for server examples and transit route configurations that use IP SLA.

Configure

- L3out toward Server and toward N3K devices.
- Configure IP SLA track for the physical IP address of the server.
- Configure static route under L3out toward server which use IP SLA track and advertise from another L3out toward N3K.

Network Diagram



ACI Lab Topology

Configurations

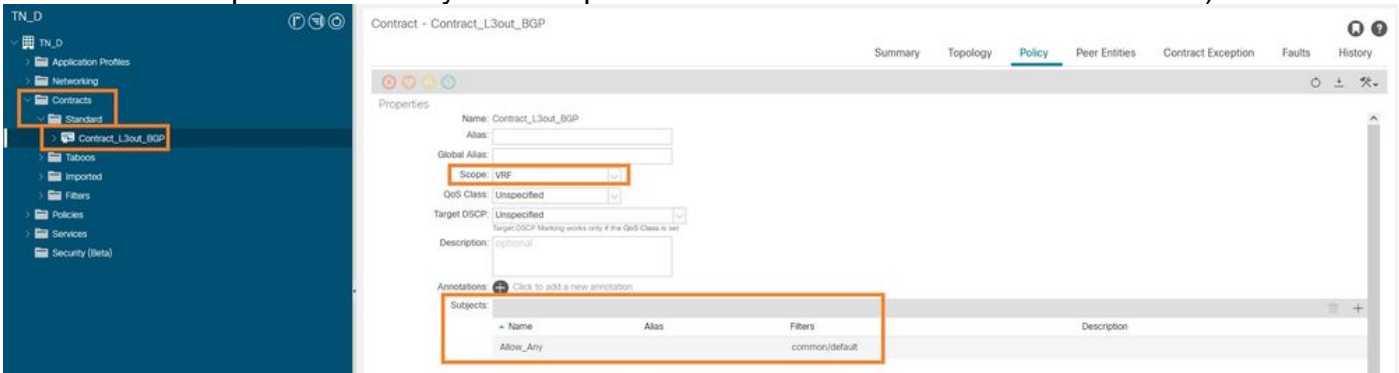
Summary Steps:

[ACI fabric policies:](#)

- Create Contract (for this example, a common default filter that allows all traffic to be used, but you can use a specific filter locally created in the same Tenant to allow specific traffic. in such a case please ensure you allow protocol that we are used for IP SLA track).
- Create new L3out toward server 10.100.0.100/24 (ACI side SVI 550 with IP address 10.100.0.254)
- Create IP SLA Track Policies (IP SLA monitor policy, Track Members policy, Track List policy)
- Add static route under L3out toward server with IP SLA tracklist.
- Create a new L3out toward the N3K device which use BGP. (EBGP) ACI AS 65535 and N3K AS 65536
- Export static route from L3out toward N3K.
- Verify configuration and reachability.

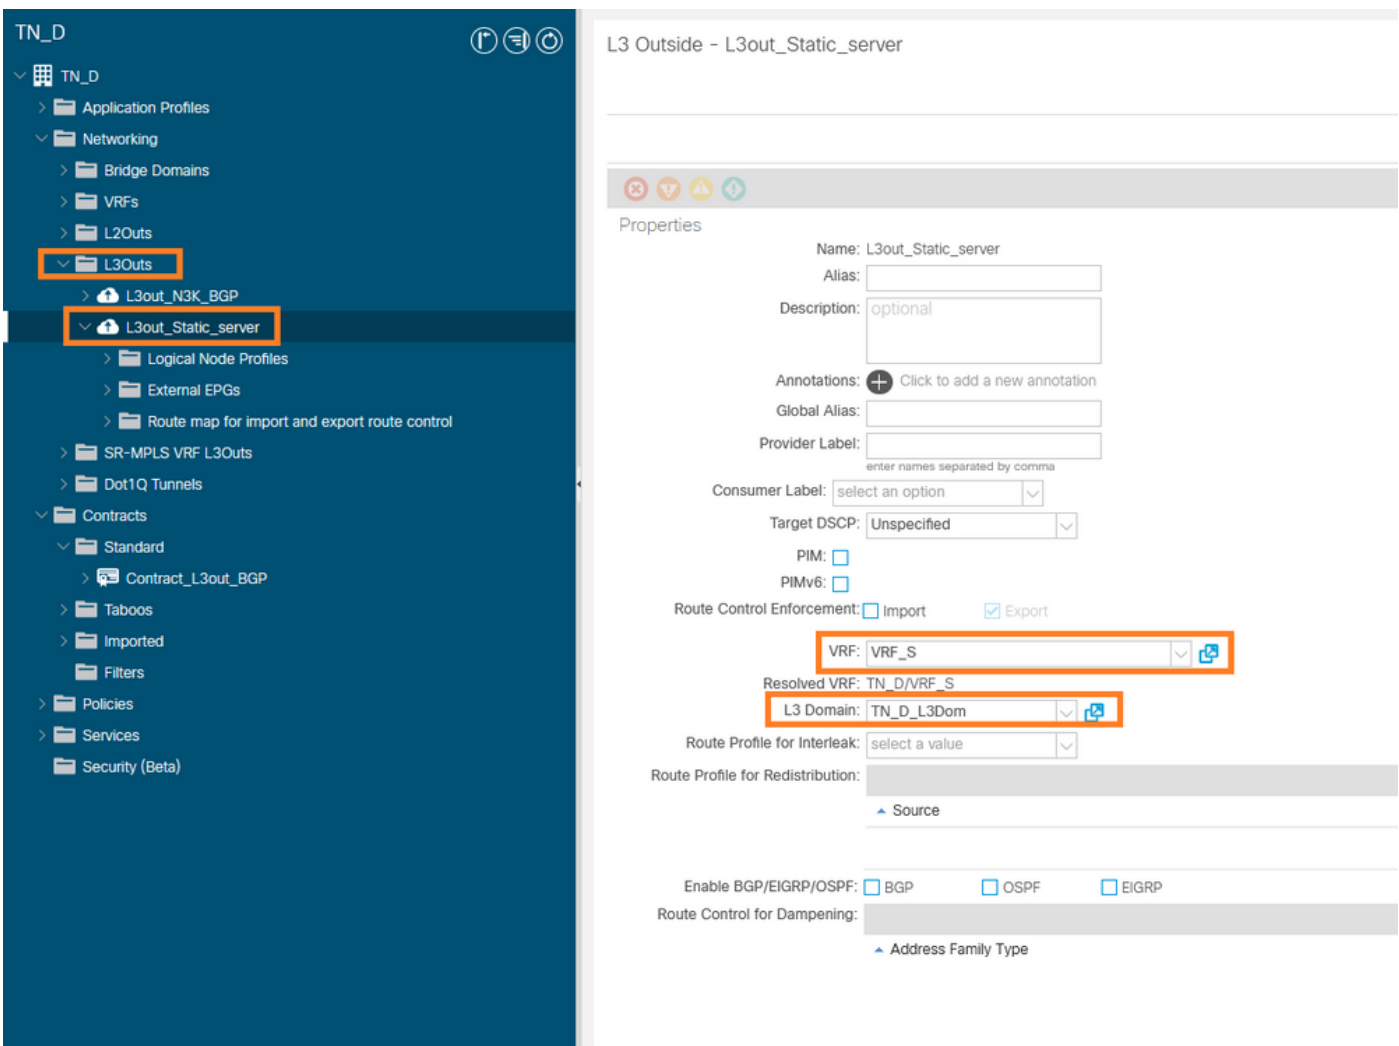
1. Create Contract (for this example, use a common default filter that allows all traffic, however, you can use a specific filter locally created in the same tenant to allow specific traffic but in

such case please ensure you allow protocol which we are used for IP SLA track).



Create Contract

2. Create a new L3out toward server 10.100.0.100/24 (ACI side SVI 550 with IP address 10.100.0.254).



Create L3out

Logical Node Profile - L3out_Static_server_nodeProfile

Properties

Name: L3out_Static_server_nodeProfile

Description: optional

Alias:

Target DSCP: Unspecified

Node ID	Router ID	Loopback Address
topology/pod-1/node-101	101.101.101.101	101.101.101.101

Create BGP Protocol Profile:

Create BFD Multihop Protocol Profile:

Attaching node to L3out

Logical Interface Profile - L3out_Static_server_interfaceProfile

Policy

General | Routed Sub-Interfaces | Routed Interfaces | SVI | Floating SVI

Path	Side A IP	Side B IP	Secondary IP Address	IP Address	MAC Address	MTU (bytes)	Encap	Encap Scope
Pod-1/Node-101/eth1/3				10.100.0.254/24	00:22:BD:FB:19:FF	inherit	vlan-507	Local

Attaching interface to L3out

External EPG - EXT_static_EPG

Policy

General | Contracts | Inherited Contracts | Subject Labels | EPG Labels

Properties

Name: EXT_static_EPG

Alias:

Annotations: Click to add a new annotation

Global Alias:

Description: optional

pcTag: 32771

Contract Exception Tag:

Configured VRF Name: VRF_S

Resolved VRF: un/tn-TN_D/ctx-VRF_S

QoS Class: Unspecified

Target DSCP: Unspecified

Configuration Status: applied

Configuration Issues:

Preferred Group Member: Exclude | Include

Intra-Ext-EPG Isolation: Enforced | Unenforced

IP Address	Scope	Name	Aggregate	Route Control Profile	Route Summarization Policy
0.0.0.0/0		External Subnets for the Extern...			

Show Usage | Reset | Submit

Configure External EPG

External EPG - EXT_static_EPG

Policy

General | Contracts | Inherited Contracts | Subject Labels | EPG Labels

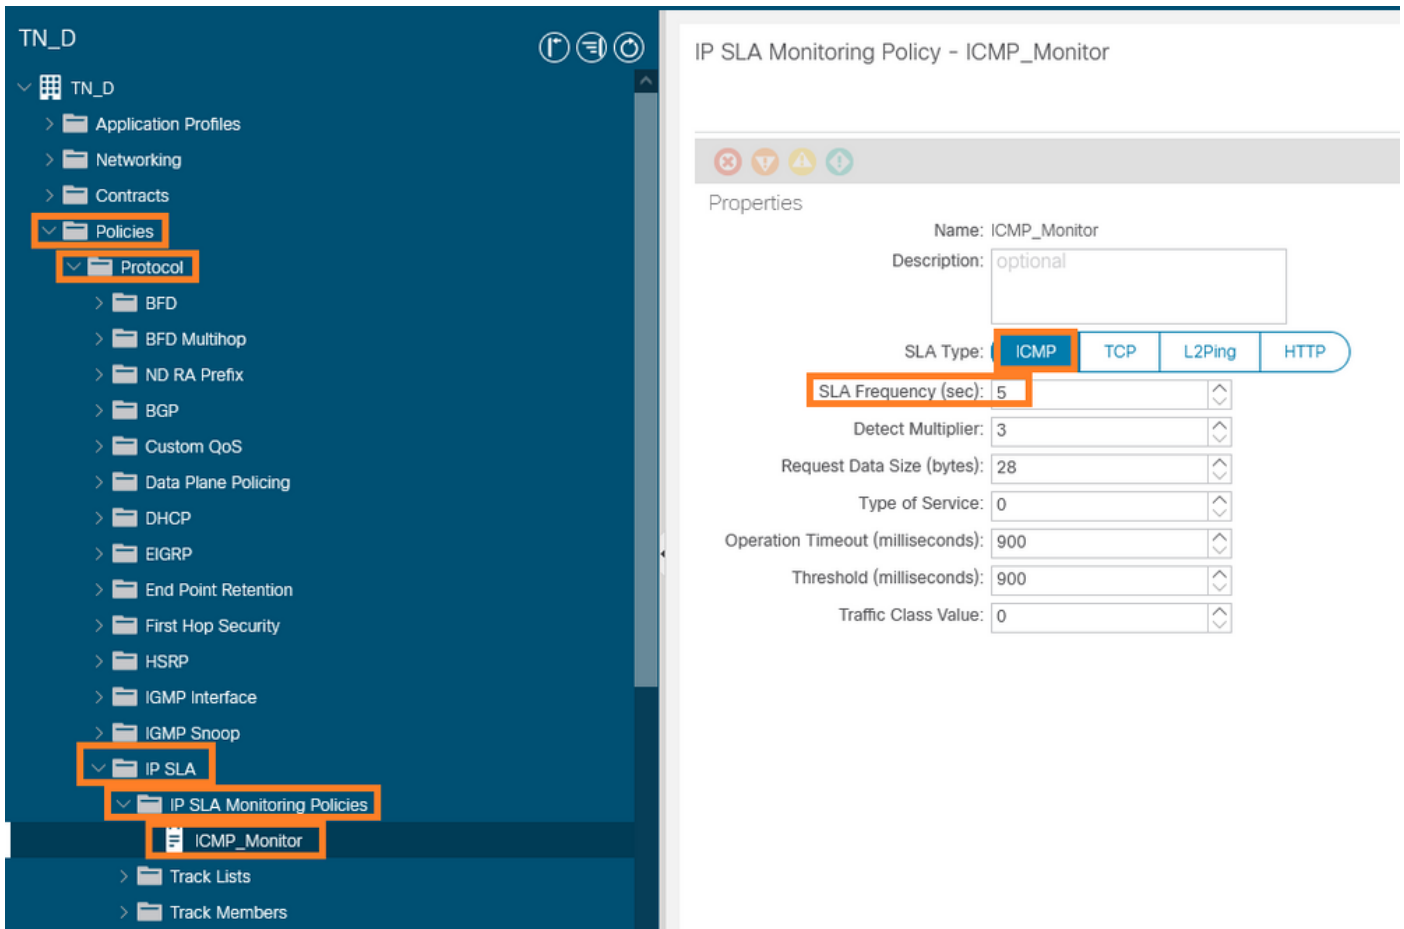
Healthy

Name	Tenant	Tenant Alias	Contract Type	Provided / Consumed	QoS Class	State	Label	Subject Label
Contract Type: Contract								
Contract_L3out_BGP	TN_D		Contract	Provided	Unspecified	formed		

Attaching Contract to L3out

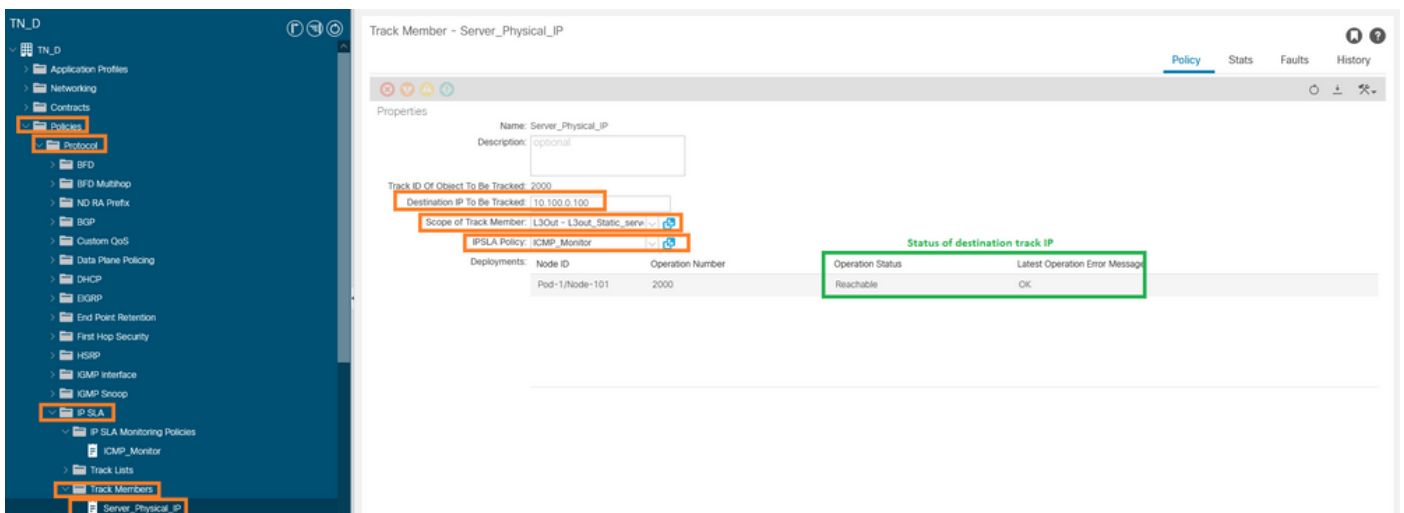
3. Create IP SLA Track Policies (IP SLA monitor policy, Track Members policy, Track List policy).

IP SLA Monitor Policy:



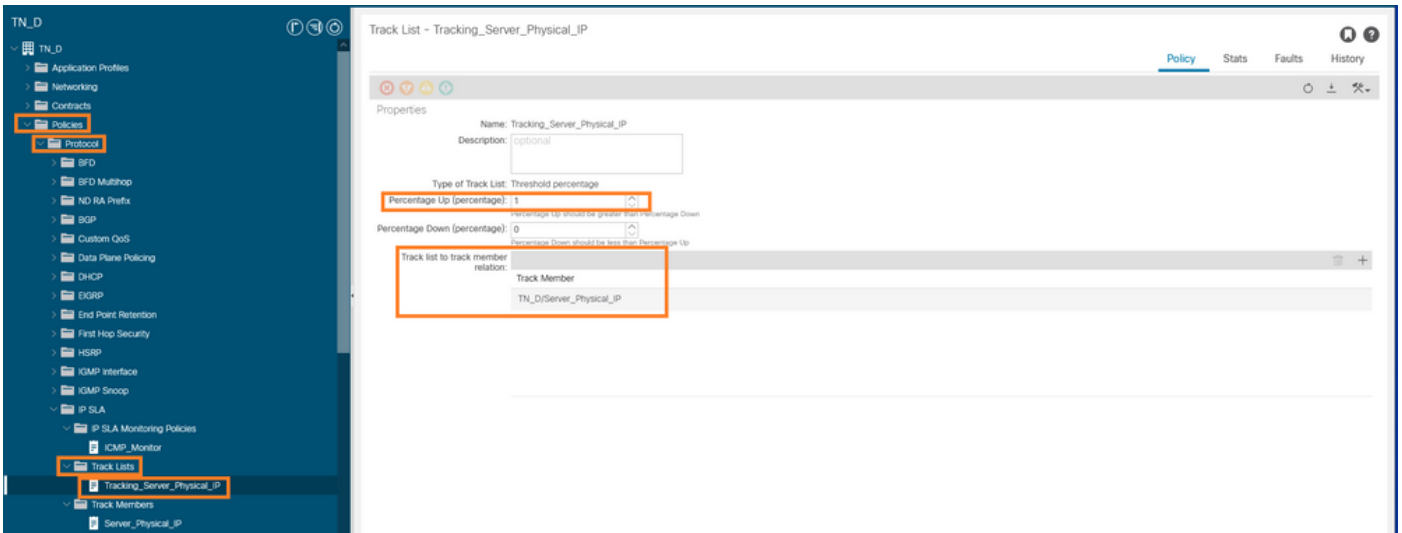
Configure IP SLA monitor policy

IP SLA Track Members:



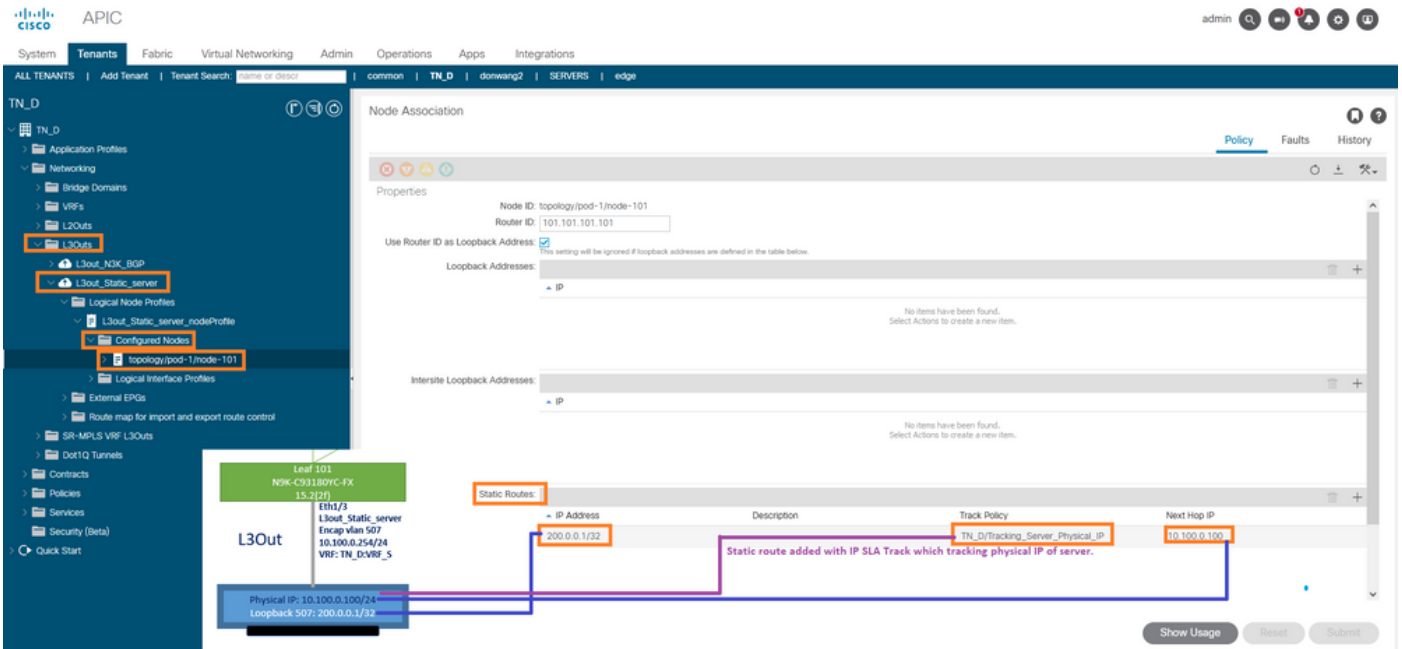
Adding IP to monitor policy

Track List Policy:



Configure Track List

4. Configure static route under L3out toward server with newly created IP SLA tracklist policy.



Configure Static route under L3out

5. Create a new L3out toward the N3K device which uses Border Gateway Protocol (BGP). (EBGP) ACI AS 65535 and N3K AS 65536.

TN_D

- Application Profiles
- Networking
 - Bridge Domains
 - VRFs
 - L2Outs
 - L3Outs**
 - L3out_N3K_BGP**
 - Logical Node Profiles
 - L3out_BGP_nodeProfile
 - Configured Nodes
 - Logical Interface Profiles
 - L3out_N3K_BGP_interfaceProfile
 - BGP Peer 100.0.0.2- Node-102/1/3
 - External EPGs
 - EXT_N3K_BGP_EPG
 - Route map for import and export route control
 - L3out_Static_server
 - SR-MPLS VRF L3Outs
 - Dot1Q Tunnels
 - Contracts
 - Policies
 - Services
 - Security (Beta)
 - Quick Start

L3 Outside - L3out_N3K_BGP

Properties

Name: L3out_N3K_BGP

Alias:

Description: optional

Annotations: Click to add a new annotation

Global Alias:

Provider Label:

Consumer Label:

Target DSCP: Unspecified

PIM:

PIMv6:

Route Control Enforcement: Import Export

VRF: VRF_S

Resolved VRF: TN_D/VRF_S

L3 Domain: TN_D L3Dom

Route Profile for Interleak:

Route Profile for Redistribution:

Enable BGP/EIGRP/OSPF: BGP OSPF EIGRP

Route Control for Dampening:

Configure BGP protocol

TN_D

- Application Profiles
- Networking
 - Bridge Domains
 - VRFs
 - L2Outs
 - L3Outs**
 - L3out_N3K_BGP**
 - Logical Node Profiles**
 - L3out_BGP_nodeProfile**
 - Configured Nodes
 - Logical Interface Profiles
 - L3out_N3K_BGP_interfaceProfile
 - BGP Peer 100.0.0.2- Node-102/1/3
 - External EPGs
 - EXT_N3K_BGP_EPG
 - Route map for import and export route control
 - L3out_Static_server
 - SR-MPLS VRF L3Outs
 - Dot1Q Tunnels
 - Contracts
 - Policies
 - Services
 - Security (Beta)
 - Quick Start

Logical Node Profile - L3out_BGP_nodeProfile

Properties

Name: L3out_BGP_nodeProfile

Description: optional

Alias:

Target DSCP: Unspecified

Nodes:

Node ID	Node	Loopback Address
1	Topology/Spw-1/node-102	102.102.102.102

BGP Peer Connectivity:

Peer IP Address: 100.0.0.2

Interface: Pod-1/Node-102/1/3

Create BGP Protocol Profile:

Create EFD Multihop Protocol Profile:

BGP Peer Profile

The screenshot displays the configuration of a BGP Peer Connectivity Profile named "100.0.0.2- Node-102/1/3". The left sidebar shows the navigation tree with the following path highlighted: TN_D > Networking > L3Outs > L3out_N3K_BGP > Logical Node Profiles > L3out_BGP_nodeProfile > Configured Nodes > Logical Interface Profiles > L3out_N3K_BGP_interfaceProfile > BGP Peer 100.0.0.2- Node-102/1/3.

The main configuration area is titled "BGP Peer Connectivity Profile 100.0.0.2- Node-102/1/3" and includes the following properties:

- Address:** 100.0.0.2
- Description:** optional
- BGP Controls:**
 - Allow Self AS
 - AS override
 - Disable Peer AS Check
 - Next-hop Self
 - Send Community
 - Send Extended Community
 - Send Domain Path
- Password:** [Empty field]
- Confirm Password:** [Empty field]
- Allowed Self AS Count:** 3
- Peer Controls:**
 - Bidirectional Forwarding Detection
 - Disable Connected Check
- Address Type Controls:**
 - AF Mcast
 - AF Ucast
- Routing Domain ID:** EBGPMultihop TTL: 3
- Weight for routes from this neighbor:** 0
- Private AS Control:**
 - Remove all private AS
 - Remove private AS
 - Replace private AS with local AS
- BGP Peer Prefix Policy:** select a value
- Site of Origin:** [Empty field]
- Remote Autonomous System Number:** 65536
- Local-AS Number Config:** [Empty field]
- Local-AS Number:** [Empty field]
- Admin State:** Disabled (Enabled button is visible)
- Route Control Profile:** [Empty field]


Configure BGP peer policy

The screenshot displays the configuration of a Logical Interface Profile named "L3out_N3K_BGP_interfaceProfile". The left sidebar shows the navigation tree with the following path highlighted: TN_D > Networking > L3Outs > L3out_N3K_BGP > Logical Node Profiles > L3out_BGP_nodeProfile > Configured Nodes > Logical Interface Profiles > L3out_N3K_BGP_interfaceProfile > BGP Peer 100.0.0.2- Node-102/1/3.


The main configuration area is titled "Logical Interface Profile - L3out_N3K_BGP_interfaceProfile" and includes the following table:

Path	Side A IP	Side B IP	Secondary IP Address	IP Address	MAC Address	MTU (bytes)	Encap	Encap Scope
Node-102/1/1/1				100.0.0.1/30	00:22:80:F8:19:0F	inherit	vlan-500	Local

Configure Logical interface profile under L3out



External EPG export subnet in transit L3out



Attaching Contract to External EPG

6. Export static route from L3out toward N3K.

```


switchname N3K
feature bgp
feature interface-vlan
interface Vlan550
  no shutdown
  vrf member BGP_L3out
  ip address 100.0.0.2/30
interface loopback200
  vrf member BGP_L3out
  ip address 30.30.30.1/32
interface Ethernet1/1
  switchport mode trunk
router bgp 65536
  address-family ipv4 unicast
  neighbor 100.0.0.1
  vrf BGP_L3out
  router-id 3.3.3.3
  address-family ipv4 unicast
  network 30.30.30.1/32
  neighbor 100.0.0.1
  remote-as 65535
  update-source Vlan550
  address-family ipv4 unicast

```

Verify

Use this section in order to confirm that your configuration works properly.

Nexus3K.



Transit route advertisement explained by topology

N3K# routing vrf BGP_L3out

```
N3K%BGP_L3out# show ip route IP Route Table for VRF "BGP_L3out" '*' denotes best ucast next-hop
'***' denotes best mcast next-hop '[x/y]' denotes [preference/metric] '%' in via output denotes
VRF 30.30.30.1/32, ubest/mbest: 2/0, attached *via 30.30.30.1, Lo200, [0/0], 02:35:27, local
*via 30.30.30.1, Lo200, [0/0], 02:35:27, direct 100.0.0.0/30, ubest/mbest: 1/0, attached *via
100.0.0.2, Vlan550, [0/0], 05:52:18, direct 100.0.0.2/32, ubest/mbest: 1/0, attached *via
100.0.0.2, Vlan550, [0/0], 05:52:18, local 200.0.0.1/32, ubest/mbest: 1/0 *via 100.0.0.1,
[20/0], 02:32:36, bgp-65536, external, tag 65535
```

Server Loopback is reachable with source as N3K loopback address.

N3K

```
interface loopback200
  vrf member BGP_L3out
  ip address 30.30.30.1/32
```

N3K# ping 200.0.0.1 vrf BGP_L3out source 30.30.30.1

```
PING 200.0.0.1 (200.0.0.1): 56 data bytes
64 bytes from 200.0.0.1: icmp_seq=0 ttl=252 time=0.94 ms
64 bytes from 200.0.0.1: icmp_seq=1 ttl=252 time=0.729 ms
64 bytes from 200.0.0.1: icmp_seq=2 ttl=252 time=0.658 ms
64 bytes from 200.0.0.1: icmp_seq=3 ttl=252 time=0.706 ms
64 bytes from 200.0.0.1: icmp_seq=4 ttl=252 time=0.655 ms
--- 200.0.0.1 ping statistics ---
5 packets transmitted, 5 packets received, 0.00% packet loss
round-trip min/avg/max = 0.655/0.737/0.94 ms
```

ACI Leaf 102 route table (which has L3out toward Nexus 3K).

Leaf102# show ip route vrf TN_D:VRF_S

```
IP Route Table for VRF "TN_D:VRF_S"
'*' denotes best ucast next-hop
'***' denotes best mcast next-hop
'[x/y]' denotes [preference/metric]
'%' in via output denotes VRF
10.100.0.0/24, ubest/mbest: 1/0
    *via 10.0.96.64%overlay-1, [200/0], 02:56:36, bgp-65535, internal, tag 65535
30.30.30.1/32, ubest/mbest: 1/0
of N3K.
```

<<address

```

    *via 100.0.0.2%TN_D:VRF_S, [20/0], 02:44:34, bgp-65535, external, tag 65536
100.0.0.0/30, ubest/mbest: 1/0, attached, direct
    *via 100.0.0.1, vlan19, [0/0], 05:09:37, direct
100.0.0.1/32, ubest/mbest: 1/0, attached
    *via 100.0.0.1, vlan19, [0/0], 05:09:37, local, local
101.101.101.101/32, ubest/mbest: 1/0
    *via 10.0.96.64%overlay-1, [1/0], 02:56:36, bgp-65535, internal, tag 65535
102.102.102.102/32, ubest/mbest: 2/0, attached, direct
    *via 102.102.102.102, lo5, [0/0], 16:49:13, local, local
    *via 102.102.102.102, lo5, [0/0], 16:49:13, direct
200.0.0.1/32, ubest/mbest: 1/0
    *via 10.0.96.64%overlay-1, [1/0], 02:42:15, bgp-65535, internal, tag 65535

```

Leaf 101 IP SLA configuration verification from CLI.

Leaf101# show ip sla configuration

```

IP SLAs Infrastructure Engine-III
Entry number: 2000
Owner: owner-icmp-echo-dme
Tag:
Operation timeout (milliseconds): 900
Type of operation to perform: icmp-echo
Target address/Source address: 10.100.0.100/0.0.0.0
Traffic-Class parameter: 0x0
Type Of Service parameter: 0x0
Request size (ARR data portion): 28
Verify data: No
Vrf Name: TN_D:VRF_S
Schedule:
  Operation frequency (seconds): 5 (not considered if randomly scheduled)
  Next Scheduled Start Time: Start Time already passed
  Group Scheduled : FALSE
  Randomly Scheduled : FALSE
  Life (seconds): Forever
  Entry Ageout (seconds): 3600
  Recurring (Starting Everyday): FALSE
  Status of entry (SNMP RowStatus): Active
Threshold (milliseconds): 900
Distribution Statistics:
  Number of statistic hours kept: 2
  Number of statistic distribution buckets kept: 1
  Statistic distribution interval (milliseconds): 20
History Statistics:
  Number of history Lives kept: 0
  Number of history Buckets kept: 15
  History Filter Type: None

```

Leaf101# show track brief

TrackId	Type	Instance	Parameter	State	Last Change
4	IP SLA	2000	reachability	up	2021-09-16T18:08:42.364+00:00
3	List	---	percentage	up	2021-09-16T18:08:42.365+00:00

Leaf101# show track

```

Track 1
  List Threshold percentage
  Threshold percentage is up
  6 changes, last change 2021-09-16T00:01:50.339+00:00
  Threshold percentage up 1% down 0%
  Tracked List Members:
    Object 2 (100)% up
  Attached to:
    Route prefix 200.0.0.1/32

```

Track 2

IP SLA 2000
reachability is up
6 changes, last change 2021-09-16T00:01:50.338+00:00
Tracked by:
Track List 1

Verification with Managed Object Query (Moquery) command:

```
apic1# moquery -c fvIPSLAMonitoringPol -f 'fv.IPSLAMonitoringPol.name=="ICMP_Monitor"'
```

Total Objects shown: 1

```
# fv.IPSLAMonitoringPol
name           : ICMP_Monitor
annotation     :
childAction    :
descr         :
dn            : uni/tn-TN_D/ipslaMonitoringPol-ICMP_Monitor
extMngdBy     :
httpMethod    : get
httpUri       : /
httpVersion   : HTTP10
ipv4Tos       : 0
ipv6TrfClass  : 0
lcOwn        : local
modTs        : 2021-09-15T21:18:48.195+00:00
monPolDn     : uni/tn-common/monepg-default
nameAlias     :
ownerKey      :
ownerTag      :
reqDataSize   : 28
rn           : ipslaMonitoringPol-ICMP_Monitor
slaDetectMultiplier : 3
slaFrequency  : 5
slaPort      : 0
slaType      : icmp
status       :
threshold    : 900
timeout      : 900
uid         : 15374
userdom     : :all:
```

```
apic1# moquery -c fvTrackMember -f 'fv.TrackMember.name=="Server_Physical_IP"'
```

Total Objects shown: 1

```
# fv.TrackMember
name          : Server_Physical_IP
annotation    :
childAction   :
descr        :
dn           : uni/tn-TN_D/trackmember-Server_Physical_IP
dstIpAddr    : 10.100.0.100
extMngdBy    :
id           : 2000
lcOwn       : local
modTs       : 2021-09-15T21:16:22.992+00:00
monPolDn    : uni/tn-common/monepg-default
nameAlias    :
ownerKey     :
ownerTag     :
rn          : trackmember-Server_Physical_IP
```

```
scopeDn      : uni/tn-TN_D/out-L3out_Static_server
status       :
uid          : 15374
userdom      : :all:
```


```
apic1# moquery -c fvTrackList -f 'fv.TrackList.name=="Tracking_Server_Physical_IP"'
Total Objects shown: 1
```

```
# fv.TrackList
name          : Tracking_Server_Physical_IP
annotation    :
childAction   :
descr         :
dn            : uni/tn-TN_D/tracklist-Tracking_Server_Physical_IP
extMngdBy     :
lcOwn         : local
modTs         : 2021-09-15T07:41:15.958+00:00
monPolDn      : uni/tn-common/monepg-default
nameAlias     :
ownerKey      :
ownerTag      :
percentageDown : 0
percentageUp  : 1
rn            : tracklist-Tracking_Server_Physical_IP
status        :
type          : percentage
uid           : 15374
userdom       : :all:
weightDown    : 0
weightUp      : 1
```

Troubleshoot

There is currently no specific troubleshooting information available for this configuration.

In event of link disconnection or physical IP address is unreachable, ACI IP SLA shows destination IP 'timeout' after configured threshold reaches.



L3out interface down

Track Member - Server_Physical_IP

Properties

Name: Server_Physical_IP
Description: optional

Track ID Of Object To Be Tracked: 2000
Destination IP To Be Tracked: 10.100.0.100
Scope of Track Member: L3Out - L3out_Static_serv
IPSLA Policy: ICMP_Monitor

Deployments:

Node ID	Operation Number	Operation Status	Latest Operation Error Message
Pod-1/Node-101	2000	Unreachable	Timeout

IP SLA monitor link status after link down

Leaf 101 CLI verification (You can see timeout for "Last Operation return code").

Leaf101# show ip sla statistics

IPSLAs Latest Operation Statistics

IPSLA operation id: 2000

Latest RTT: NoConnection/Busy/Timeout

Latest operation start time: 23:54:30 UTC Wed Sep 15 2021

Latest operation return code: Timeout

Number of successes: 658

Number of failures: 61

Operation time to live: forever

As soon as the server is reachable, it shows status **OK**.

Track Member - Server_Physical_IP

Properties

Name: Server_Physical_IP
Description: optional

Track ID Of Object To Be Tracked: 2000
Destination IP To Be Tracked: 10.100.0.100
Scope of Track Member: L3Out - L3out_Static_serv
IPSLA Policy: ICMP_Monitor

Deployments:

Node ID	Operation Number	Operation Status	Latest Operation Error Message
Pod-1/Node-101	2000	Reachable	OK

IP SLA monitor status after link brought up

Leaf101# show ip sla statistics

IPSLAs Latest Operation Statistics

IPSLA operation id: 2000

Latest RTT: 1 milliseconds

Latest operation start time: 00:03:15 UTC Thu Sep 16 2021

Latest operation return code: OK

Number of successes: 18
Number of failures: 86
Operation time to live: forever

Related Information

- [Cisco APIC Layer 3 Networking Configuration Guide, Release 5.2\(x\)](#)
- [Technical Support & Documentation - Cisco Systems](#)