



The bridge to possible

ACI Troubleshooting

Advanced L3out Features

Roland Ducombe, Technical Leader, CX-EMEAR – CCIE 3745

Cisco Webex App


Questions?

Use Cisco Webex App to chat with the speaker after the session

How

- 1 Find this session in the Cisco Live Mobile App
- 2 Click “Join the Discussion”
- 3 Install the Webex App or go directly to the Webex space
- 4 Enter messages/questions in the Webex space

Webex spaces will be moderated until February 24, 2023.





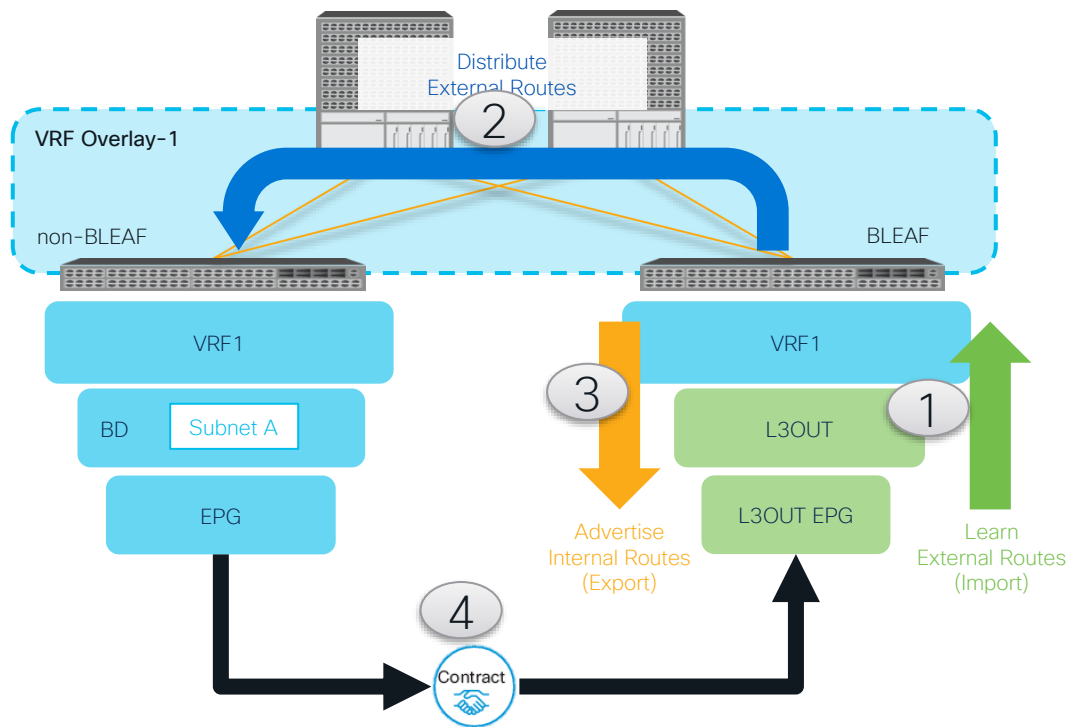
Agenda

- Introduction
- Example 1 – Simple Layer 3 out using eBGP
- Example 2 – Transit routing between OSPF and eBGP Layer 3 out
- Example 3 – Policy enforcement option with Layer 3 out
- Route-map in ACI overview
- Example 4 – Route-map example :
 - eBGP setting community ingress and OSPF matching community egress
- Summary

Introduction



L3OUT Key Components



- 1 Learn external routes
 - Routing Protocol in L3OUT
 - Import route-control (optional)
- 2 Distribute external routes to other leaves
 - MP-BGP
- 3 Advertise internal or other external routes (BD subnet or routes from other L3out) to outside
 - Redistribution – export route-control and
 - Contract
- 4 Allow traffic with contracts
 - L3OUT EPG (Prefix Based EPG)

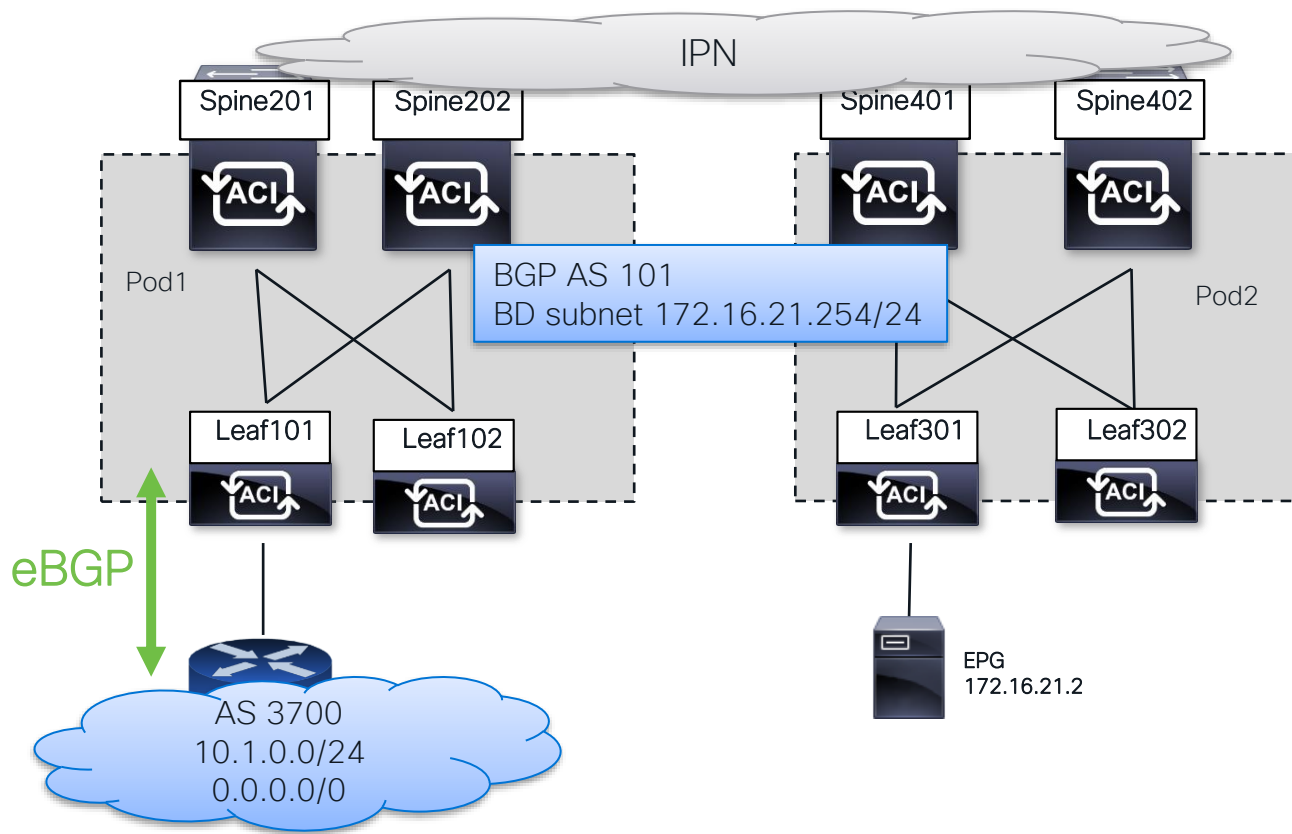
Example 1 : Simple eBGP L3 out



Setup 1 – Simple L3 out eBGP

Simple L3 out on leaf 101
Receive eBGP route

EP 172.16.21.1 in Pod2




Setup 1 – Control plane – L3 out route to Server Leaf

1 BGP routes received on BL
(default is import all)

2 Spine are Route Reflector
per pod to distribute to
local pod

3 BGP VPNv4 exchanged
Route across Pod (spine to
spine)

4 Spine egress Pod RR
reflect to leaf in pod2



Border Leaf – CLI check – BGP

1. BGP peering to external router is up and we received 2 routes

```
S1P1-Leaf101# show bgp ipv4 unicast summary vrf DC:DC
```

```
...
Neighbor      V      AS MsgRcvd MsgSent  TblVer  InQ OutQ Up/Down State/PfxRcd
192.168.101.2  4    3700    21     18     10    0   0 00:10:46 2
```

Any numeric value (even 0)
Means session is up
Idle or active means
session is flapping or down

2. Verify routes in bgp table – here received from AS 3700

```
S1P1-Leaf101# show bgp ipv4 unicast vrf DC:DC
```

```
...
Network          Next Hop          Metric      LocPrf      Weight Path
..
*>e10.1.0.0/24    192.168.101.2          0 3700 i
```

3. Routes is injected in VPNv4 address family

```
S1P1-Leaf101# show bgp vpnv4 unicast vrf DC:DC | egrep "Net|Route|10.1.0.0"
```

```
Network          Next Hop          Metric      LocPrf      Weight Path
Route Distinguisher: 101:2359302 (VRF DC:DC)
*>e10.1.0.0/24    192.168.101.2          0 3700 i
```

Border Leaf – CLI check

4. Border leaf have VPNv4 peering with 2 spine Route Reflector

```
S1P1-Leaf101# show bgp vpnv4 unicast summary vrf overlay-1 | egrep "Neig|10\.0"
BGP router identifier 10.0.0.64, local AS number 101
Neighbor      V      AS MsgRcvd MsgSent   TblVer  InQ  OutQ  Up/Down  State/PfxRcd
10.0.0.65     4     101   17769   17615     712    0     0    1w5d 72
10.0.0.66     4     101   17778   17615     712    0     0    1w5d 72
```

5. Routes are exported with Route Target of format : BGP-ASN:VRF-VNID

```
S1P1-Leaf101# show bgp process vrf DC:DC | egrep -A 5 "Export"
Export RT list:
101:2359302
Import RT list:
101:2359302
```

Server Leaf – CLI check

1. Server leaf receive BGP VPNv4 from spine

```
S1P2-Leaf301# show bgp vpnv4 unicast 10.1.0.0/24 vrf DC:DC
```

```
Path type (0xa25a1c60): internal 0xc0000018 0x40 ref 0 adv path ref 2, path is valid, is best path
Imported from (0xa25f74b4) 101:2359302:10.1.0.0/24
AS-Path: 3700 , path sourced external to AS
10.0.0.64 (metric 33) from 10.1.96.64 (172.16.2.4)
Origin IGP, MED not set, localpref 100, weight 0 tag 0, 1
Received label 0
Received path-id 2
Extcommunity:
RT:101:2359302
COST:pre-bestpath:165:2415919104
VNID:2359302.
```

BL PTEP – BGP NH

Spine that reflected that path

Route-Target that we import :

```
S1P2-Leaf301# show bgp process vrf DC:DC | egrep -A 2 Import
Import RT list:
101:2359302
```

2. Server leaf install route in RIB with NH PTEP of BL

```
S1P2-Leaf301# show ip route 10.1.0.0 vrf DC:DC
```

```
..
10.1.0.0/24, ubest/mbest: 1/0
*via 10.0.0.64%overlay-1, [200/0], 2d20h, bgp-101, internal, tag 3700
recursive next hop 10.0.0.64/32%overlay-1
```

```
bdsol-aci37-apic1# acidiag fnvread | egrep "10.0.0.64"
```

101	1	S1P1-Leaf101	FDO224702JA	10.0.0.64/32	leaf	active
-----	---	--------------	-------------	--------------	------	--------

BRKDCN-3678

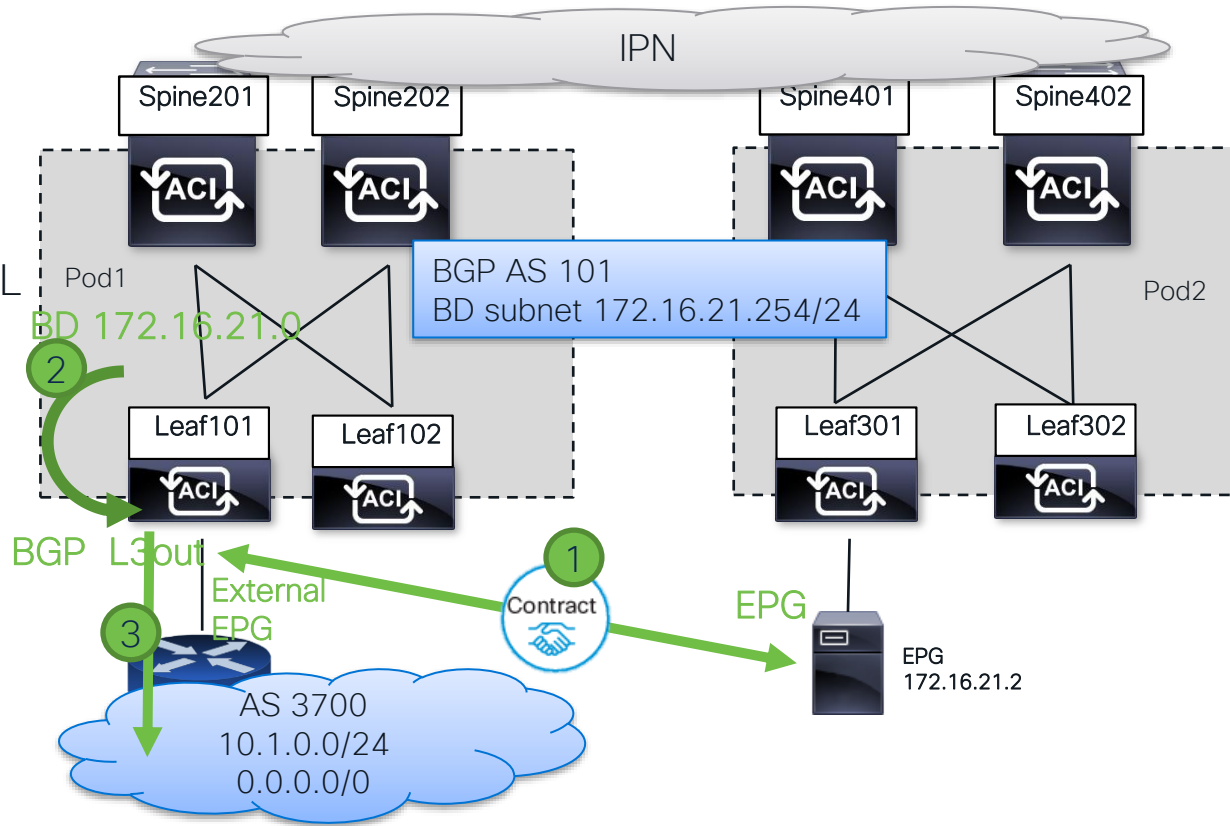
Setup 1 – Control plane BD subnet to external

No route exported by default
(deny all route-map default)

1 BD subnet must be in RIB of BL
(contract)

2 BD subnet must be
redistributed to BGP process
(VPNv4)

3 BD subnet added to outbound
route-map



Border leaf – Routing and contract

```
S1P1-Leaf101# show ip route 172.16.21.0 vrf DC:DC
..
S1P1-Leaf101#
```

```
S1P1-Leaf101# show ip route 172.16.21.0 vrf DC:DC
172.16.21.0/24, ubest/mbest: 1/0, attached, direct, pervasive
*via 10.0.72.64%overlay-1,[1/0],00:00:04, static, tag 4294967294
recursive next hop: 10.0.72.64/32%overlay-1
```

```
S1P1-Leaf101# show bgp ipv4 unicast 172.16.21.0/24 vrf DC:DC
..
S1P1-Leaf101#
```

```
S1P1-Leaf101# show bgp ipv4 unicast neighbors 192.168.101.2 vrf
DC:DC
```

```
..
Inbound route-map configured is permit-all, handle obtained
Outbound route-map configured is exp-l3out-BGP-peer-235930
```

```
S1P1-Leaf101# show route-map exp-l3out-BGP-peer-2359302
route-map exp-l3out-BGP-peer-2359302, deny, sequence 16000
Match clauses:
route-type: direct
```

1. No BD subnet in BL RIB

2. BD subnet added **when a contract is added** between Ext EPG and EPG

Default Route tag
for a private subnet

3. BD subnet not in BGP yet

4. By default outbound route-map deny all

Border leaf – Sending BD subnet

Step 1

Subnet Advertised Externally

→ Route tag of BD subnet is removed and it pushes the subnet to BGP

```
S1P1-Leaf101# show ip route 172.16.21.0 vrf DC:DC
172.16.21.0/24, ubest/mbest: 1/0, attached, direct, pervasive
  *via 10.0.72.64%overlay-1,[1/0],00:00:04, static,
    recursive next hop: 10.0.72.64/32%overlay-1
```

```
S1P1-Leaf101# show bgp process vrf DC:DC | egrep -A 2 "Redis"
Redistribution
  static, route-map imp-ctx-bgp-st-interleak-2359302
```

```
S1P1-Leaf101# show route-map imp-ctx-bgp-st-interleak-2654211
route-map imp-ctx-bgp-st-interleak-2359302, deny, sequence 1
Match clauses:
  tag: 4294967294
Set clauses:
route-map imp-ctx-bgp-st-interleak-2359302, permit, sequence 20000
Match clauses:
```

```
S1P1-Leaf101# show bgp vpnv4 unicast vrf RD-MPOD:RD | egrep "172.16.11.0"
*>r172.16.11.0/24      0.0.0.0      0      100      32768 ?
```

Properties

IP Address: 172.16.21.254/24

Description: optional

Treat as virtual IP address: ☐

Make this IP address primary: ☐

Scope: ☒ Advertised Externally
☐ Shared between VRFs

Subnet Control: ☐ No Default SVI Gateway
☐ Querier IP

L3 Out for Route Profile: select a value

Configuration Issues:

Policy Tags: + Click to add a new tag

IP Data-plane Learning:

No more route tag

Static to BGP Route-map used

Sequence 1 deny private subnet based on tag
Sequence 20000 permit all the rest → route goes to BGP VPNv4

Border leaf – Sending BD subnet –

Method 1 – Step 2

```
S1P1-Leaf101# show route-map exp-l3out-BGP-peer-2359302
```

```
..  
route-map exp-l3out-BGP-peer-2359302, permit, sequence 15801  
Match clauses:  
  ip address prefix-lists: IPv4-peer16387-2359302-exc-int-inferred-export-dst  
..
```

```
S1P1-Leaf101# show ip prefix-list IPv4-peer16387-2359302-exc-int-inferred-export-dst  
ip prefix-list IPv4-peer16387-2359302-exc-int-inferred-export-dst: 1 entries  
seq 1 permit 172.16.21.254/24
```

BD to L3 out association
→ Act on route-map

Bridge Domain - BD1

Properties

Unicast Routing: ☒

Operational Value for Unicast Routing: true

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

Virtual MAC Address: Not Configured

Subnets:

Gateway Address	Description
172.16.21.254/24	

EP Move Detection Mode: ☐ GARP based detection

Associated L3 Outs:

- L3 Out
- BGP

```
S1P1-Leaf101#show bgp ipv4 unicast neighbor 192.168.101.2 advertised-routes vrf DC:DC
```

Network	Next Hop	Metric	LocPrf	Weight	Path
*>r172.16.21.0/24	0.0.0.0	0	100	32768	101 ?

Troubleshooting : Routing protocol unexpected behavior



- RP traffic is targeted to cpu you can always use tcpdump to see what you receive (on kpm_inb)
- Easier on kpm_inb if linux interface on leaf/spine
 - `bdsol-aci32-leaf1# tcpdump -ni kpm_inb proto eigrp`
 - `bdsol-aci32-leaf1# tcpdump -ni kpm_inb proto ospf`
 - `bdsol-aci32-leaf1# tcpdump -ni kpm_inb -f port 179`
- You can add extra filter such as :
 - `bdsol-aci32-leaf1# tcpdump -ni kpm_inb -f port 179 and host 1.1.1.1`
- Or get more verbose :
 - `bdsol-aci32-leaf1# tcpdump -nxxvvi kpm_inb -f port 179 and host 1.1.1.1`
- Or write to pcap file
 - `bdsol-aci32-leaf1# tcpdump -i kpm_inb -f port 179 -w /bootflash/bgp-trace.pcap`

Tcpdump is your friend

Debug log for routing protocol



- Bgp, eigrp, isis, ospf and some other protocol traces are binary encoded.
- File end up with .bl
- All can be decoded using : `"log_trace_bl_print_tool" <file name>`
- Note in latest code a lot more process have bl trace (arp, pim, urib, acllog,...)

```
SlP1-Leaf101# ls -al /var/sysmgr/tmp_logs/*.bl
-rw-rw-rw- 1 root root 43439057 Oct 19 10:53 bgp_trace.bl
-rw-rw-rw- 1 root root 39618433 Oct 19 10:53 coop_trace.bl
-rw-rw-rw- 1 root root 59710790 Oct 19 10:53 isis_trace.bl
-rw-rw-rw- 1 root root 37771710 Oct 19 10:53 ospfv2_1_trace.bl
-rw-rw-rw- 1 root root      6671 Oct  3 14:38 ospfv2_2_trace.bl
-rw-rw-rw- 1 root root      2666 Sep 22 13:20 ospfv3_1_trace.bl
-rw-rw-rw- 1 root root   374065 Oct 19 10:53 rpm_trace.bl
...

SlP1-Leaf101# log_trace_bl_print_tool /var/sysmgr/tmp_logs/bgp_trace.bl | more
version: 1, pid: 60215
[2019 Sep 11 07:13:23.632899106:main:4257] (0) OBJ: kcache lib initialized succesfully in BGP
[2019 Sep 11 07:13:23.634434168:main:4298] BGP process bgp-132 startup, reason: configuration
```

Example 2 : Transit Layer 3 out.




Setup 2- Transit L3 out

Simple L3 out on leaf 101
Receive eBGP route 10.1.0.0

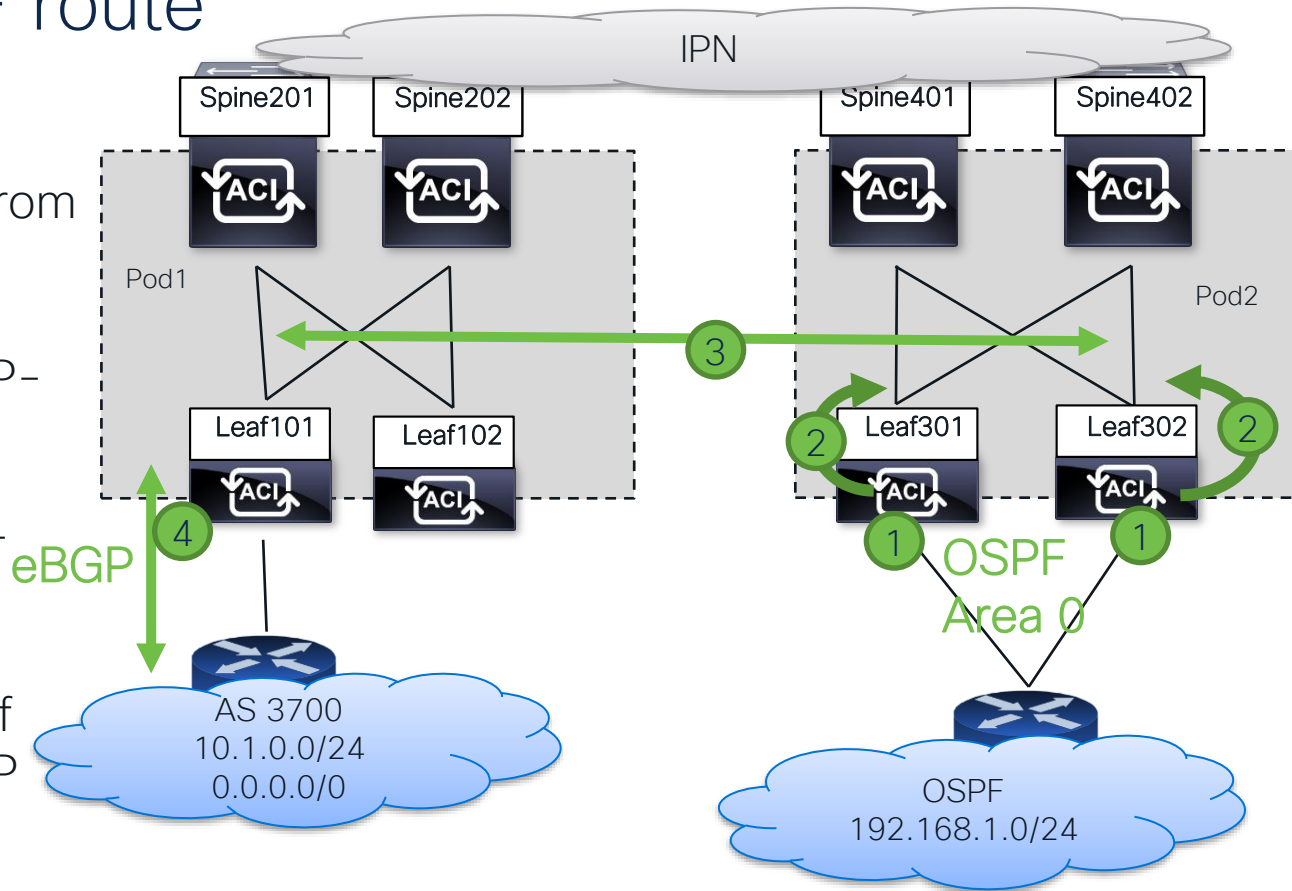
OSPF on leaf301-302
receive 192.168.1.0

Goal : Transit between to the
layer 3 out



Setup 2- OSPF route

- 1 OSPF routes is received from external router
- 2 Route is redistributed in MP-BGP
- 3 Route is propagated in MP-BGP to all leaf in VRF
- 4 Route (192.168.1.0) on leaf 101 from iMP-BGP to eBGP



OSPF CLI check

1. Verify OSPF interface parameters – matching with neighbors ?

```
S1P2-Leaf301# show ip ospf interface vrf DC:DC
Vlan101 is up, line protocol is up
  IP address 192.168.102.1/29, Process ID default VRF DC:DC, area backbone
  State DR-OTHER, Network type BROADCAST, cost 4
  Index 141, Transmit delay 1 sec, Router Priority 1
  Designated Router ID: 192.168.0.13, address: 192.168.102.3
  Backup Designated Router ID: 192.168.0.4, address: 192.168.102.2
  2 Neighbors, flooding to 2, adjacent with 2
  Timer intervals: Hello 10, Dead 40, Wait 40, Retransmit 5
S1P2-Leaf301# show interface vlan 101 | egrep MTU
  MTU 9000 bytes, BW 10000000 Kbit, DLY 1 usec
```

2. Verify OSPF neighbor is established on Broadcast network – FULL/(B)DR or TWOWAY/DROTHER

```
S1P2-Leaf301# show ip ospf neighbor vrf DC:DC
```

Neighbor ID	Pri	State	Up Time	Address	Interface
192.168.0.4	1	FULL/BDR	00:51:21	192.168.102.2	Vlan101
192.168.0.13	1	FULL/DR	00:51:23	192.168.102.3	Vlan101

3. Receiving OSPF external routes ?

```
S1P2-Leaf301# show ip route 192.168.1.0 vrf DC:DC
192.168.1.0/24, ubest/mbest: 1/0
  *via 192.168.102.3, vlan101, [110/5], 00:57:47, ospf-default, intra
```

Border Leaf – CLI check – From OSPF to MP-BGP

1. Verify the OSPF route is inject in MP-BGP (default permit-all import route-map

```
S1P2-Leaf301# show bgp process vrf DC:DC | egrep ospf  
ospf, route-map permit-all
```

Import-all route-map from OSPF
To BGP

```
S1P2-Leaf301# show bgp vpnv4 unicast 192.168.1.0/24 vrf DC:DC
```

Route Distinguisher: 301:2359302 (VRF DC:DC)

Advertised path-id 1, VPN AF advertised path-id 1

AS-Path: NONE, **path locally originated**

0.0.0.0 (metric 0) from 0.0.0.0 (10.1.208.65)

Origin incomplete, MED 5, localpref 100, weight 32768 tag 0, propagate 0

Extcommunity:

RT:101:2359302

VNID:2359302

1 Path locally imported from OSPF

VPN AF advertised path-id 2

Path type (0xa25a1e50): internal 0xc0000018 0x40000000 adv path ref 1, path is valid
reason: Weight

Imported from (0xa25f6cf4) 302:2359302:192.168.1.0/24

AS-Path: NONE, path sourced internal to AS


10.1.208.64 (metric 3) from 10.1.96.64 (172.16.2.4)

Origin incomplete, MED 5, localpref 100, weight 0 tag 0, propagate 0

Extcommunity:

RT:101:2359302

2 Path from leaf 302 (2nd OSPF BL)



Setup 2- export transit route - method 1

1 OSPF routes (now iBGP) to eBGP on leaf 101

External EPG - ExtEPG-BGP

Properties

Name: ExtEPG-BGP

Alias:

Annotations: [+ Click to add a new annotation](#)

Global Alias:

Description: optional

pcTag: 16387

Contract Exception Tag:

Configured VRF Name: DC

Resolved VRF: uni/tn-DC/ctx-DC

QoS Class: Unspecified

Target DSCP: Unspecified

Configuration Status: applied

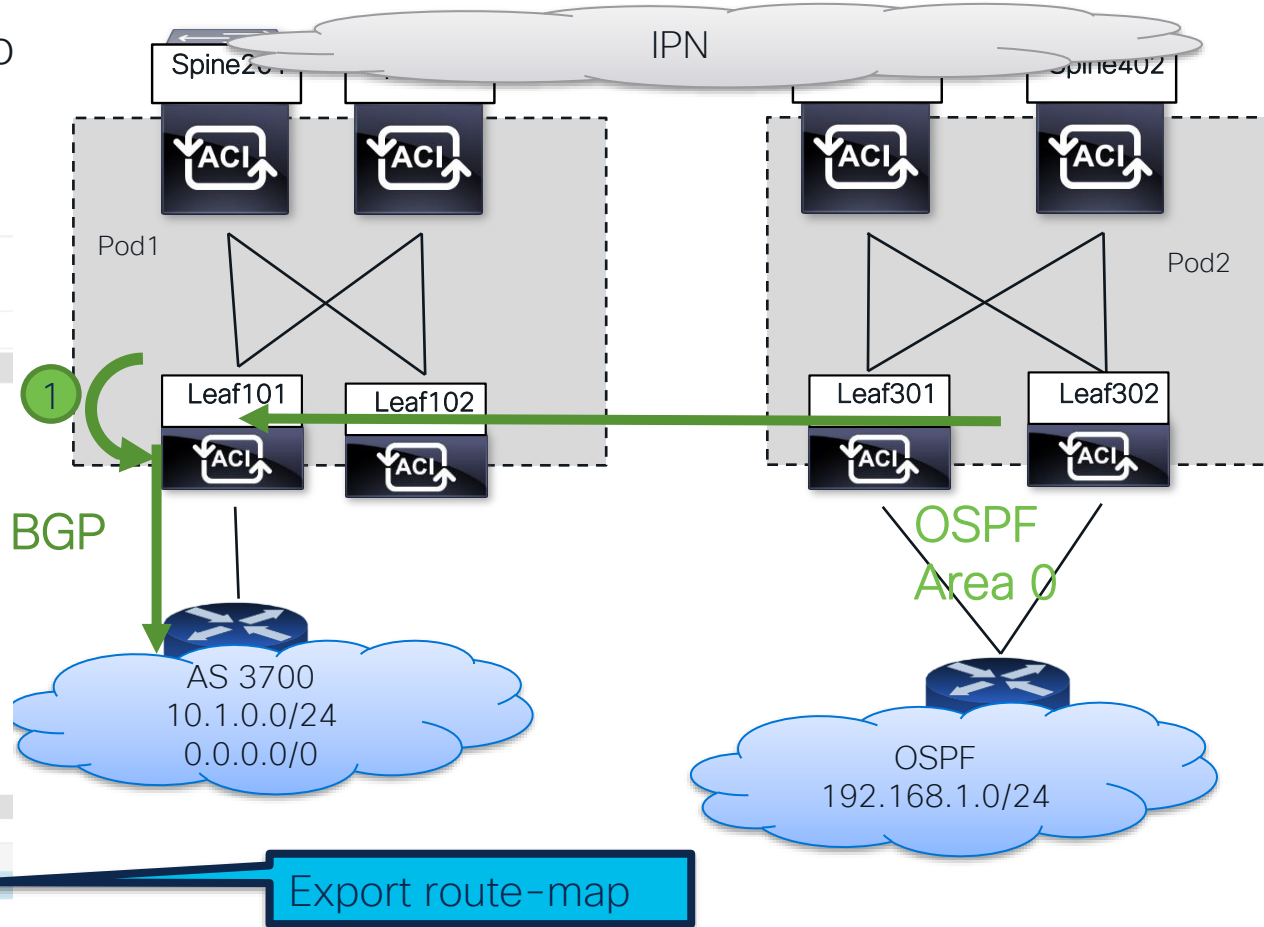
Configuration Issues:

Preferred Group Member:

Intra Ext-EPG Isolation:

Subnets:

IP Address	Scope
192.168.1.0/24	Export Route Control Subnet



Export to eBGP – Cli check

1. Find outbound route-map for BGP neighbor

```
S1P1-Leaf101# show bgp ipv4 unicast neighbors 192.168.101.2 vrf DC:DC | egrep Outbound
Outbound route-map configured is exp-l3out-BGP-peer-2359302, handle obtained
```

2. Route-map sequence for External prefix inferred export

```
S1P1-Leaf101# show route-map exp-l3out-BGP-peer-2359302
..
route-map exp-l3out-BGP-peer-2359302, permit, sequence 15802
  Match clauses:
    ip address prefix-lists: IPv4-peer16387-2359302-exc-ext-inferred-export-dst
    ipv6 address prefix-lists: IPv6-deny-all
  Set clauses:
    tag 4294967295
```

3. Prefix-list

```
S1P1-Leaf101# show ip prefix-list IPv4-peer16387-2359302-exc-ext-inferred-export-dst
ip prefix-list IPv4-peer16387-2359302-exc-ext-inferred-export-dst: 1 entries
  seq 1 permit 192.168.1.0/24
```


Export to eBGP – Cli check

4. Route-map sequence for External prefix inferred export

```
S1P1-Leaf101# show bgp ipv4 unicast neighbors 192.168.101.2 advertised-routes vrf DC:DC
```

```
Peer 192.168.101.2 routes for address family IPv4 Unicast:
```

```
BGP table version is 56, local router ID is 192.168.100.1
```

```
Status: s-suppressed, x-deleted, S-stale, d-dampened, h-history, *-valid, >-best
```

```
Path type: i-internal, e-external, c-confed, l-local, a-aggregate, r-redist, I-injected
```

```
Origin codes: i - IGP, e - EGP, ? - incomplete, | - multipath, & - backup
```

Network	Next Hop	Metric	LocPrf	Weight	Path
..					
*>i192.168.1.0/24	10.1.208.64	5	100	0	101 ?

Setup 2- Redistribute iBGP to OSPF

1 eBGP routes (now iBGP) to OSPF on leaf 301-302

External EPG - ExtEPG-OSPF

Properties

Name: ExtEPG-OSPF

Alias:

Annotations: + Click to add a new annotation

Global Alias:

Description: optional

pcTag: 32770

Contract Exception Tag:

Configured VRF Name: DC

Resolved VRF: uni/tn-DC/ctx-DC

QoS Class: Unspecified

Target DSCP: Unspecified

Configuration Status: applied

The QoS priority class identifier.

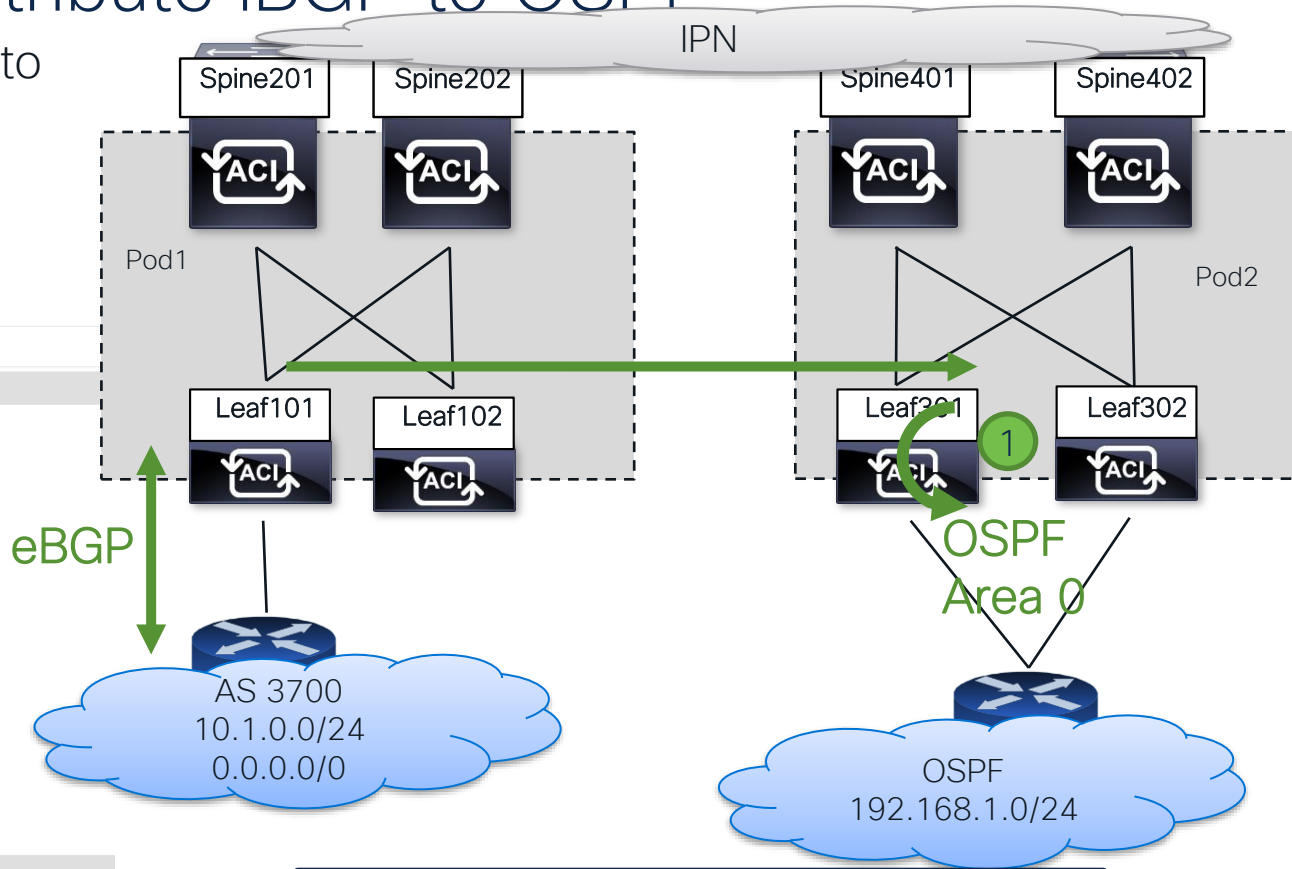
Configuration Issues:

Preferred Group Member: ☒ Exclude ☐ Include

Intra Ext-EPG Isolation: ☒ Enforced ☐ Unenforced

Subnets:

IP Address	Scope
0.0.0.0/0	External Subnets for the External EPG
10.1.0.0/24	Export Route Control Subnet



Redistribute BGP to OSPF Route-map

Export to OSPF – Cli check

1. Find outbound route-map for BGP to OSPF redistribution

```
S1P2-Leaf301# show ip ospf vrf DC:DC | egrep bgp
  bgp route-map exp-ctx-proto-2359302
```

2. Route-map sequence for External prefix inferred export

```
S1P2-Leaf301# show route-map exp-ctx-proto-2359302
..
route-map exp-ctx-proto-2359302, permit, sequence 15801
  Match clauses:
    ip address prefix-lists: IPv4-proto32770-2359302-exc-ext-inferred-export-dst
    ipv6 address prefix-lists: IPv6-deny-all
  Set clauses:
    tag 4294967295
```

3. Prefix-list

```
S1P2-Leaf301# show ip prefix-list IPv4-proto32770-2359302-exc-ext-inferred-export-dst
ip prefix-list IPv4-proto32770-2359302-exc-ext-inferred-export-dst: 1 entries
  seq 1 permit 10.1.0.0/24
```

Export to OSPF – Cli check

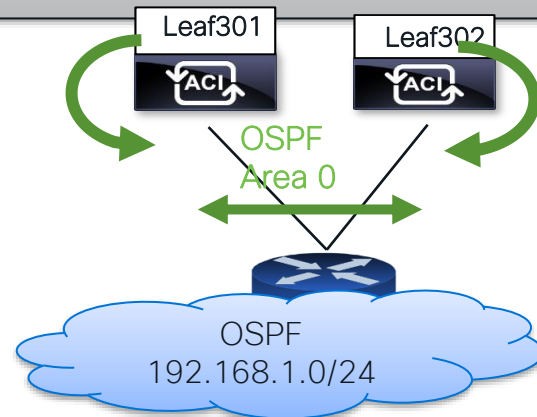
4. Verify prefix is in OSPF database as external LSA (type 5)

```
S1P2-Leaf301# show ip ospf database external 10.1.0.0 vrf DC:DC
      OSPF Router with ID (192.168.0.3) (Process ID default VRF DC:DC)
```

Type-5 AS External Link States

Link ID	ADV Router	Age	Seq#	Checksum	Tag
10.1.0.0	192.168.0.3	410	0x80000002	0x3e1b	4294967295
10.1.0.0	192.168.0.4	410	0x80000002	0x3820	4294967295

Two LSA in the OSPF DB
One redistributed on leaf 301 and
one on leaf 302




Example 3 – Policy enforcement and Layer 3 out

or how to derive pcTag when source and/or destination is an L3 out prefix ?

Policy enforcement – Where ?

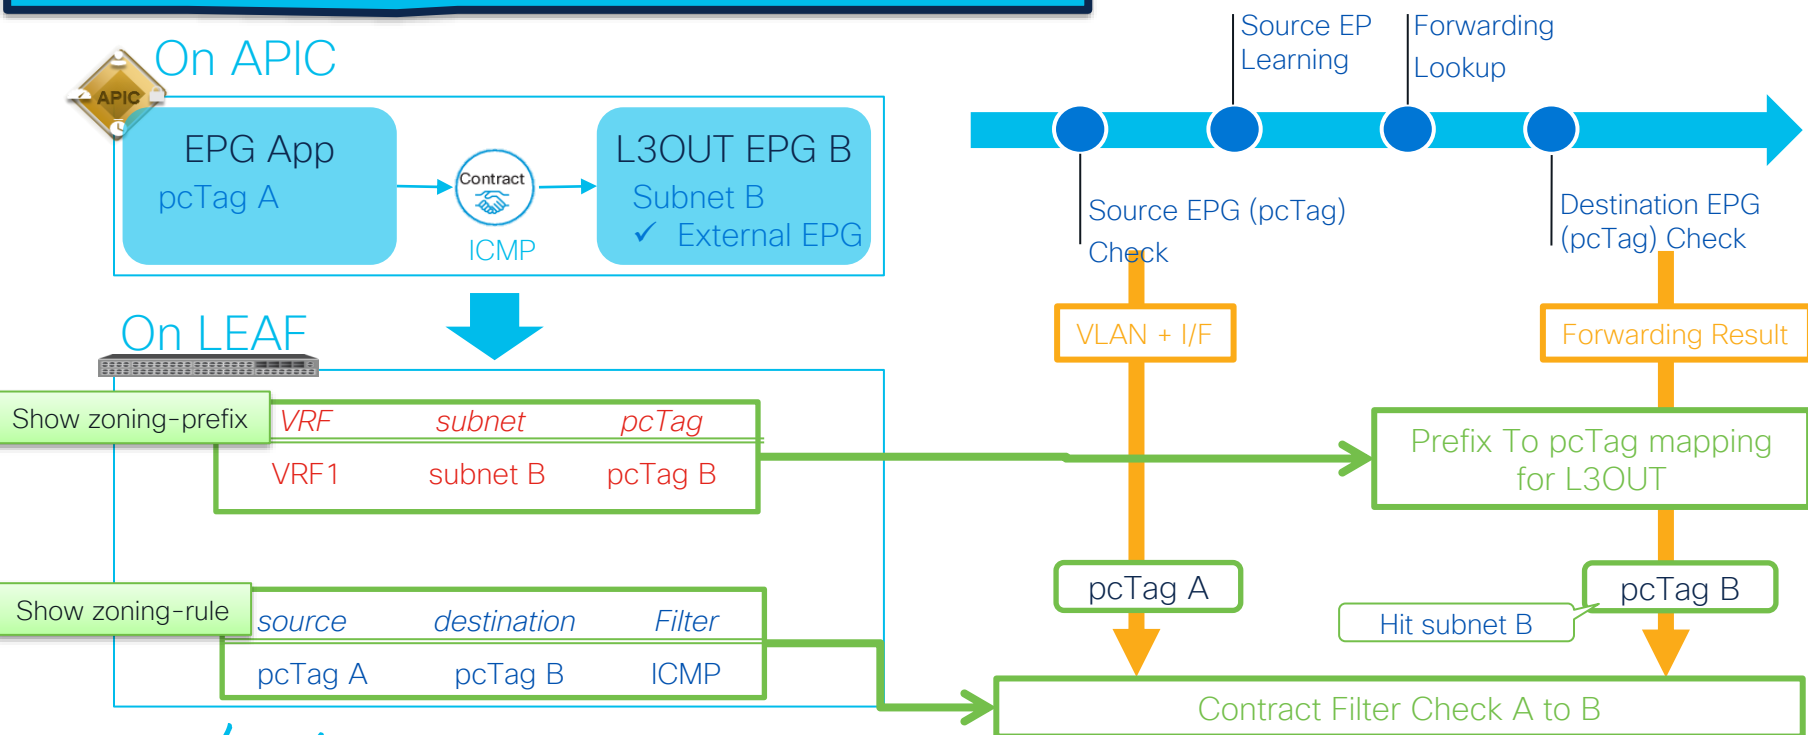
Assumption is using default vrf ingress enforcement mode



In that case policy is always enforced on server leaf in both direction

L3OUT Contract - EPG to external EPG

Prefix to pcTag table distributed to all leaf where VRF exists build from all “External subnets for External EPG” of all L3 out in the VRF



L3OUT Contract - External EPG to EPG

Prefix to pcTag table is used in both direction (ingress or egress from an L3 out)

On APIC

L3OUT EPG B
Subnet B
✓ External EPG



EPG App
pcTag A

On LEAF



VRF	subnet	pcTag
VRF1	subnet B	pcTag B

source	destination	Filter
pcTag B	pcTag A	ICMP

Hit subnet B

Src: Subnet B -> Dst: Subnet A

Forwarding
Lookup (EP or
Pervasive BD)

Source EPG (pcTag)
Check

Destination EPG (pcTag)
Check

VLAN + I/F → Ext EPG

Forwarding Result

Src Prefix To pcTag mapping
for L3OUT in vrf

pcTag B

pcTag A


Contract Filter Check B to A

L3OUT Contract – EPG to external EPG

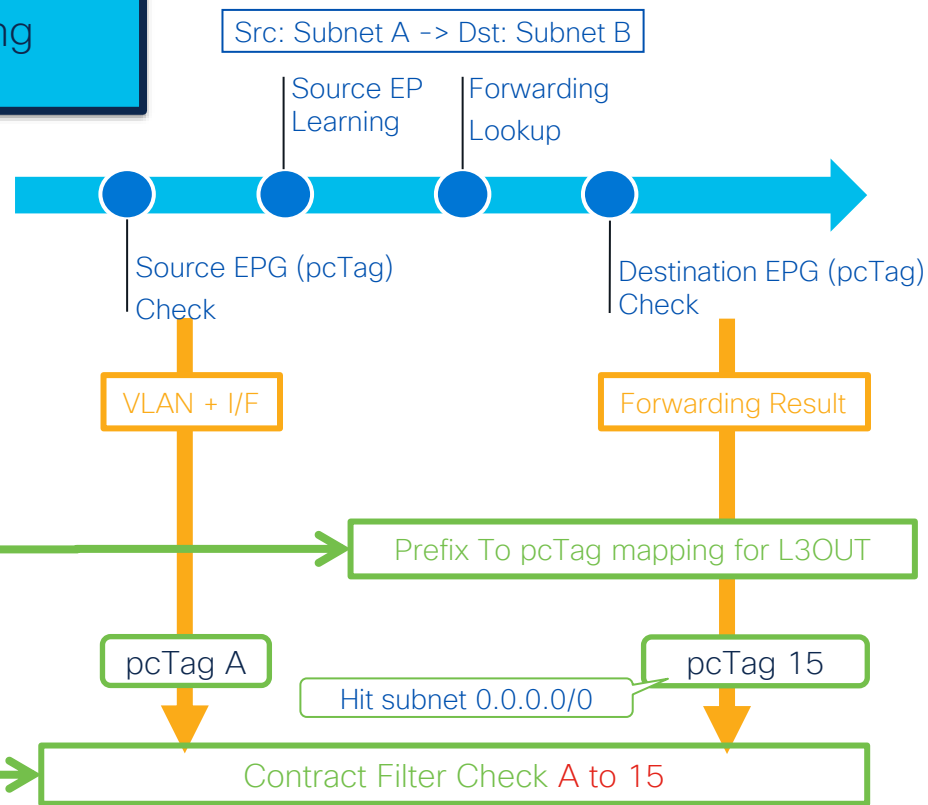
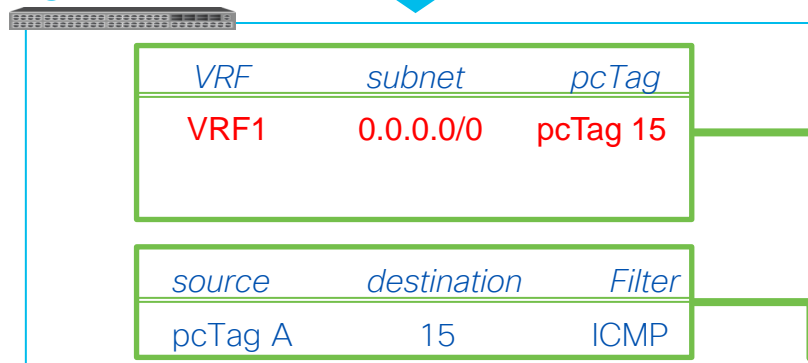
L3OUT EPG with 0.0.0.0/0

Exception is External Subnet for External EPG is 0.0.0.0/0 → considered “wildcard” always setting reserved pcTag 15 in egress

On APIC



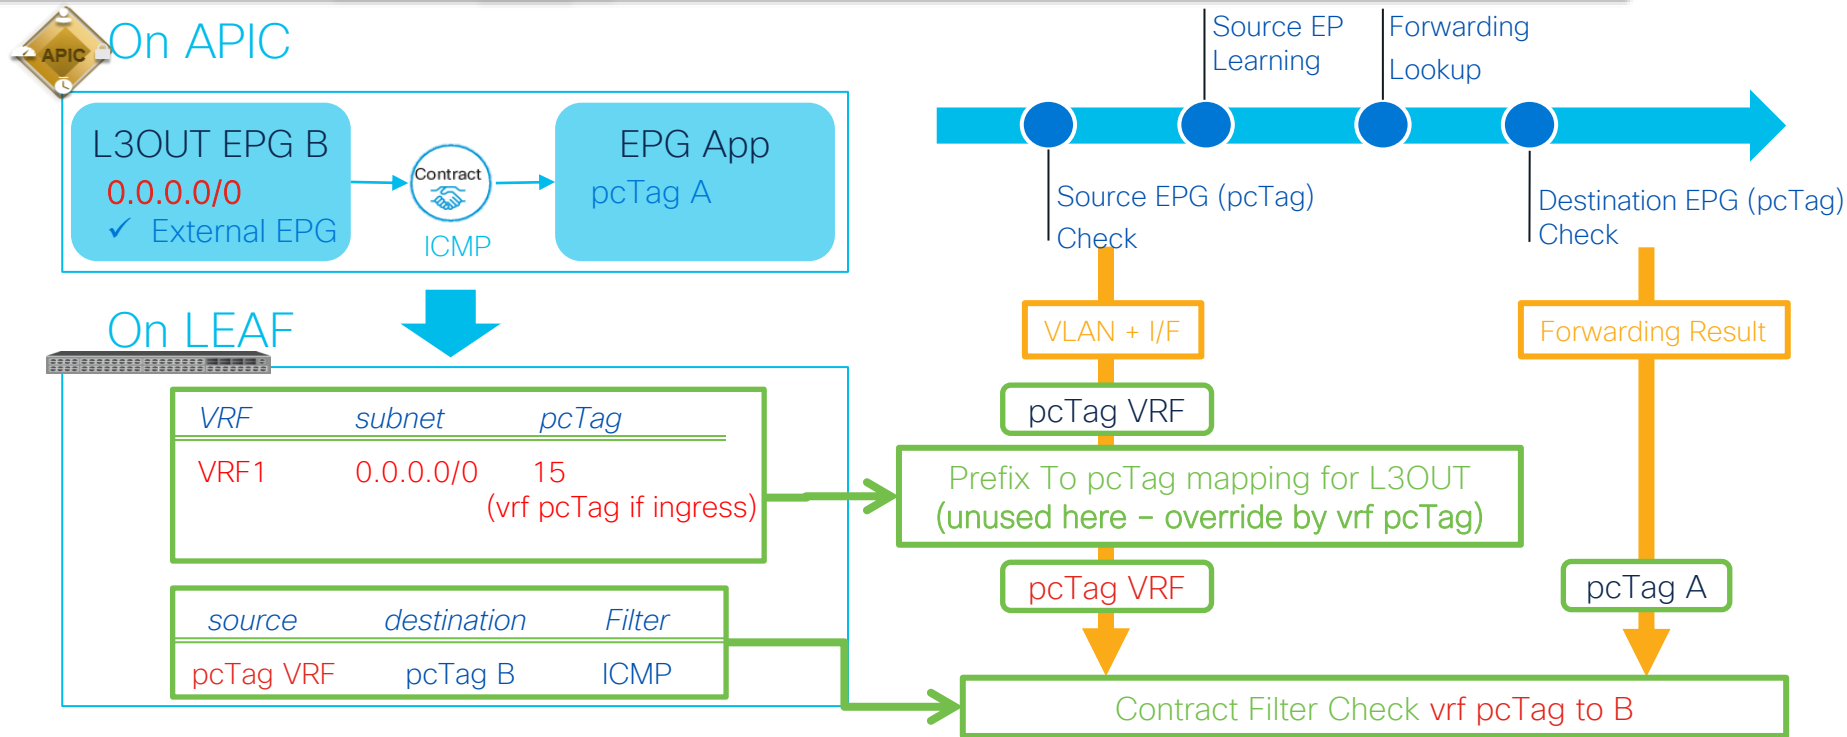
On LEAF



L3OUT Contract – External EPG to EPG

L3OUT EPG with 0.0.0.0/0


In ingress direction we use vrf pcTag if matching 0.0.0.0/0 external subnet



Setup 3a - Contract and L3 out specific L3 out subnet

Goal :
HTTPS contract between
eBGP L3 out (10.1/24) and
BD 172.16.21/24

ICMP contract between
OSPF L3 out (192.168.1/24)
and BD2 172.16.22/24



Subnet Config 3a – external subnet – cli check

External EPG - ExtEPG-BGP

Properties

Name: ExtEPG-BGP

Alias:

Annotations: [Click to add a new annotation](#)

Global Alias:

Description: optional

pcTag: 16387

Contract Exception Tag:

Configured VRF Name: DC

Resolved VRF: uni/trn-DC/ctx-DC

QoS Class: Unspecified

Target DSCP: Unspecified

Configuration Status: applied

Configuration Issues:

Preferred Group Member:

Intra-EPG Isolation:

Subnets:

IP Address	Scope
0.0.0.0/0	External Subnets for the External EPG
10.1.0.0/24	External Subnets for the External EPG
192.168.1.0/24	Export Route Control Subnet

External EPG - ExtEPG-OSPF

Properties

Name: ExtEPG-OSPF

Alias:

Annotations: [Click to add a new annotation](#)

Global Alias:

Description: optional

pcTag: 32770

Contract Exception Tag:

Configured VRF Name: DC

Resolved VRF: uni/trn-DC/ctx-DC

QoS Class: Unspecified

Target DSCP: Unspecified

Configuration Status: applied

Configuration Issues:

Preferred Group Member:

Intra-EPG Isolation:

Subnets:

IP Address	Scope
0.0.0.0/0	External Subnets for the External EPG
10.1.0.0/24	Export Route Control Subnet
192.168.1.0/24	External Subnets for the External EPG

Zoning-prefix table is distributed to all leaf where vrf exists (if vrf policy enforcement is left to default "ingress")

Note zoning-prefix is a longest prefix match table (LPM). Supernet can be used 10.0.0.0/8 and 192.168.0.0/16 for example

S1P2-Leaf301# show zoning-prefixes | egrep "DC:DC|---|Vrf"

Vrf-Vni	Vrf-Name	Address	Class	OperState
2359302	DC:DC	::/0	15	enabled
2359302	DC:DC	0.0.0.0/0	15	enabled
2359302	DC:DC	10.1.0.0/24	16387	enabled
2359302	DC:DC	192.168.1.0/24	32770	enabled

Subnet Config 3a – exact subnet – cli check



Zoning-rule check

```
S1P2-Leaf301# show zoning-rule scope 2359302 | egrep "ICMP|HTTPS|Filter|--"
```

Rule ID	SrcEPG	DstEPG	FilterID	Dir	Scope	Name	Action	Priority
4290	16387	49153	4	bi-dir	2359302	DC:HTTPS	permit	fully_qual(7)
4287	49153	16387	6	uni-dir-ignore	2359302	DC:HTTPS	permit	fully_qual(7)
4296	32770	16389	5	uni-dir-ignore	2359302	DC:ICMP	permit	fully_qual(7)
4295	16389	32770	5	bi-dir	2359302	DC:ICMP	permit	fully_qual(7)

Filter check

```
S1P2-Leaf301# show zoning-filter | egrep "FilterId|4|5|6|--"
```

FilterId	Name	EtherT	ArpOpc	Prot	SFromPort	SToPort	DFromPort	DToPort
5	5_0	ip	unspecified	icmp	unspecified	unspecified	unspecified	unspecified
6	6_0	ipv4	unspecified	tcp	https	https	unspecified	unspecified
4	4_0	ipv4	unspecified	tcp	unspecified	unspecified	https	https

Python script (embedded in leaf code)

```
S1P2-Leaf301# contract_parser.py --vrf DC:DC | egrep "ICMP|HTTPS"
```

```
[7:4290] [vrf:DC:DC] permit ipv4 tcp tn-DC/13out-BGP/instP-ExtEPG-BGP(16387) tn-DC/ap-App/epg-EPG1(49153) eq 443 [contract:uni/tn-DC/brc-HTTPS] [hit=0]  
[7:4295] [vrf:DC:DC] permit ip icmp tn-DC/ap-App/epg-EPG2(16389) tn-DC/13out-OSPF/instP-ExtEPG-OSPF(32770) [contract:uni/tn-DC/brc-ICMP] [hit=0]  
[7:4296] [vrf:DC:DC] permit ip icmp tn-DC/13out-OSPF/instP-ExtEPG-OSPF(32770) tn-DC/ap-App/epg-EPG2(16389) [contract:uni/tn-DC/brc-ICMP] [hit=0]  
[7:4287] [vrf:DC:DC] permit ipv4 tcp tn-DC/ap-App/epg-EPG1(49153) eq 443 tn-DC/13out-BGP/instP-ExtEPG-BGP(16387) [contract:uni/tn-DC/brc-HTTPS] [hit=0]
```

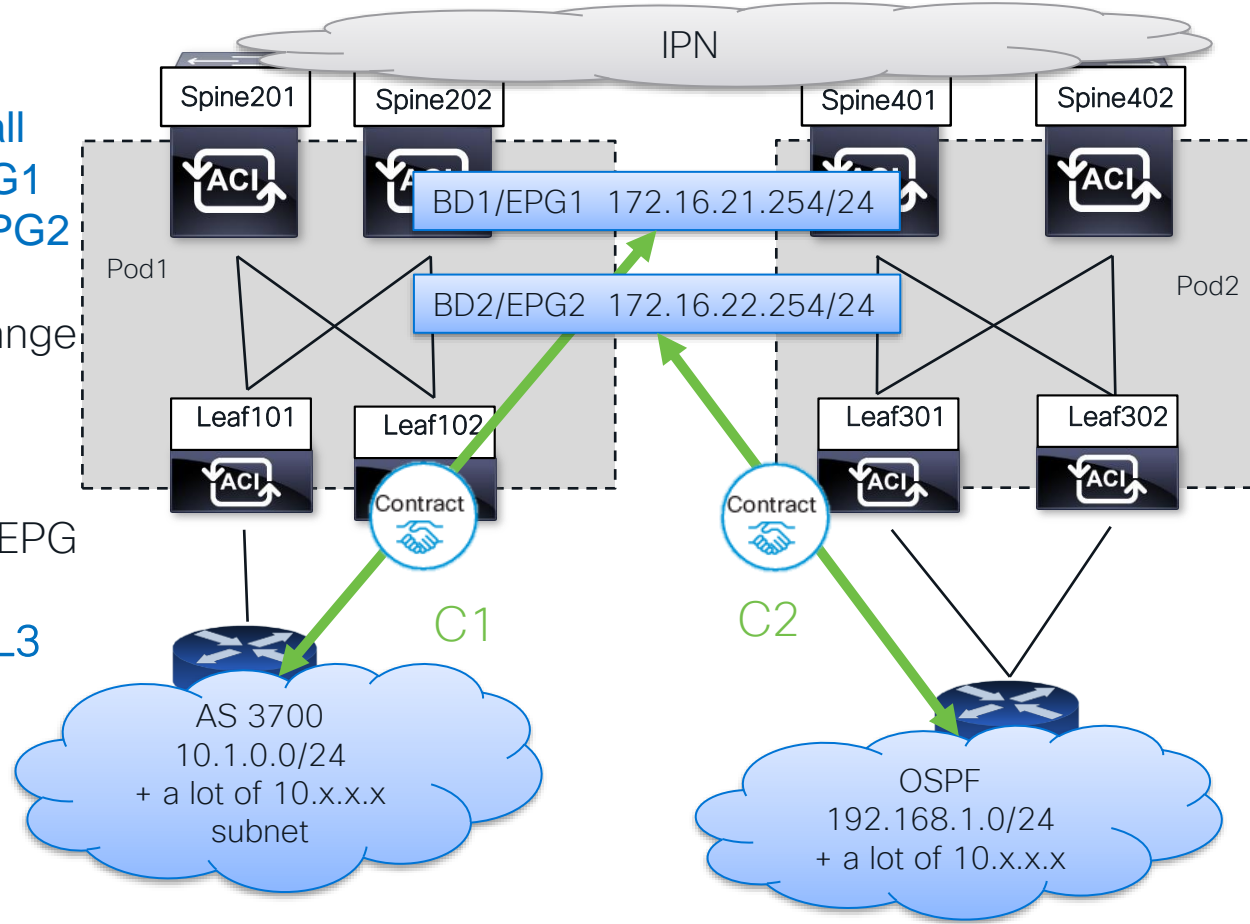
Setup 3b - Contract and L3 out - random mix of subnet across l3 out

Contract C1 and C2 both allow all
C1 from BGP L3 out to EPG1
C2 from OSPF L3 out to EPG2

Now many subnet from same range
10.x.x.x on both L3 out

Can't use the same LPM prefix
10.0.0.0/8 on multiple external EPG

Need to use 0.0.0.0/0 on both L3
out



Subnet Config 3b – 0.0.0.0/0 usage

External EPG - ExtEPG-BGP

Properties

Name: ExtEPG-BGP

Alias:

Annotations: [Click to add a new annotation](#)

Global Alias:

Description: optional

pcTag: 16387

Contract Exception Tag:

Configured VRF Name: DC

Resolved VRF: uni/tn-DC/ctx-DC

QoS Class: Unspecified

Target DSCP: Unspecified

Configuration Status: applied

Configuration Issues:

Preferred Group Member:

Intra Ext-EPG Isolation:

Subnets:

IP Address	Scope	Name
0.0.0.0/0	External Subnets for the External EPG	

External EPG - ExtEPG-OSPF

Properties

Name: ExtEPG-OSPF

Alias:

Annotations: [Click to add a new annotation](#)

Global Alias:

Description: optional

pcTag: 32770

Contract Exception Tag:

Configured VRF Name: DC

Resolved VRF: uni/tn-DC/ctx-DC

QoS Class: Unspecified

Target DSCP: Unspecified

Configuration Status: applied

Configuration Issues:

Preferred Group Member:

Intra Ext-EPG Isolation:

Subnets:

IP Address	Scope	Name
0.0.0.0/0	External Subnets for the External EPG	
10.1.0.0/24	Export Route Control Subnet	

Both L3 out use 0.0.0.0/0
external epg only 15 used
Class in zoning-prefix

External epg pcTag unused

```
S1P2-Leaf301# show zoning-prefixes | egrep "DC:DC|---|Vrf"
```

Vrf-Vni	Vrf-Name	Address	Class	OperState
2359302	DC:DC	0.0.0.0/0	15	enabled

Subnet Config 3b – “catch-all” subnet – cli check

Zoning-rule check

```
S1P2-Leaf301# show zoning-rule scope 2359302 | egrep "ICMP|HTTPS|Filter|--"
```

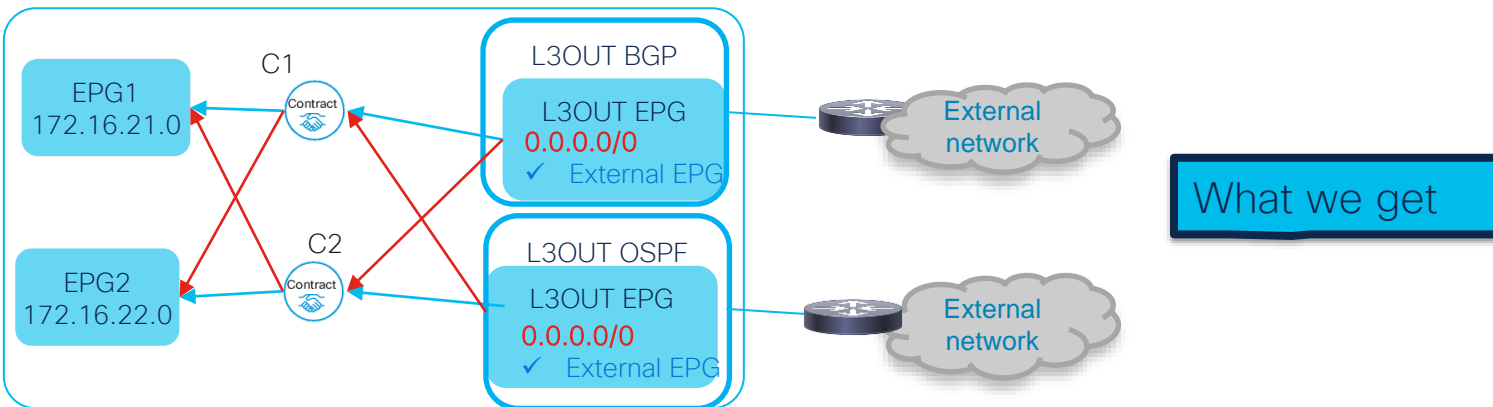
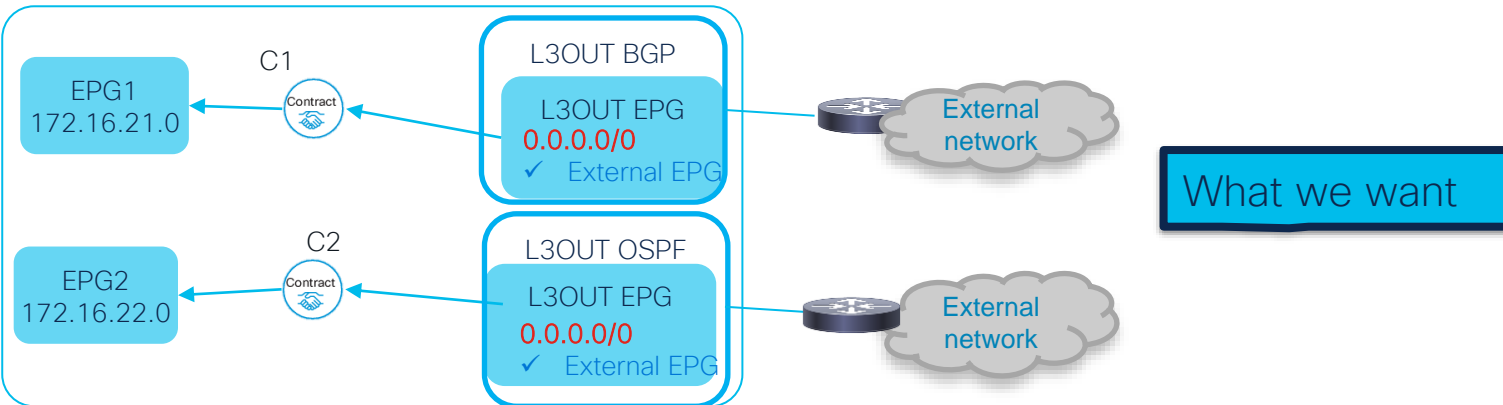
Rule ID	SrcEPG	DstEPG	FilterID	Dir	Scope	Name	Action	Priority
4283	16389	15	default	uni-dir	2359302	DC:C1	permit	fully_qual(7)
4284	49153	15	default	uni-dir	2359302	DC:C2	permit	fully_qual(7)
4285	16386	16389	default	uni-dir	2359302	DC:C1	permit	fully_qual(7)
4286	16386	49153	default	uni-dir	2359302	DC:C2	permit	fully_qual(7)

Healthy				
▲ VRF Name	VRF Alias	Class ID	Segment ID	Scope
DC		16386	2359302	2359302

Healthy			
▲ Application Profile Name	EPG Name	Class ID	Scope
App	EPG1	49153	2359302
App	EPG2	16389	2359302

Both EPG1 (49153) and EPG2 (16389) can go to both L3 out
(not expected)

Risk of using 0.0.0.0/0 on multiple L3 out



L3Out Internal Route Maps



(OSPF, EIGRP) Two types of route maps

OSPF

Route-map used to determine what is allowed To OSPF/EIGRP

```
border-leaf# show ip ospf vrf TK:VRFA | egrep 'direct|static|bgp|eigrp'  
direct route-map exp-ctx-st-2785280  
static route-map exp-ctx-st-2785280  
bgp route-map exp-ctx-proto-2785280  
eigrp route-map exp-ctx-proto-2785280
```

EIGRP

```
border-leaf# show ip eigrp vrf TK:VRFA | egrep 'direct|static|ospf|bgp'  
bgp-65002 route-map exp-ctx-proto-2785280  
direct route-map exp-ctx-st-2785280  
ospf-default route-map exp-ctx-proto-2785280  
static route-map exp-ctx-st-2785280
```

exp-ctx-st-<VRF VNID>

Route maps for direct or static routes

- L3Out association to a BD
- Export Route Control Subnet
- Route map like default-export

exp-ctx-proto-<VRF VNID>

Route maps for routing protocols

- Export Route Control Subnet
- Route map like default-export

(BGP) a route map per L3Out or per peer

Route-map used to determine what is outbound or inbound of a BGP L3 out

(when not using a per peer route map)

```
border-leaf# show bgp ipv4 unicast neighbors vrf TK:VRFA | grep Outbound
Outbound route-map configured is exp-l3out-BGP-peer-2785280, handle obtained
```

(when using a per peer route map)

```
border-leaf# show bgp ipv4 unicast neighbors vrf TK:VRFA | grep Outbound
Outbound route-map configured is TK-BGP_PEER1-BGP-out, handle obtained
```

Without per-peer route-map (default behavior)

exp-l3out-<L3Out>-peer-<VRF VNID>

- L3Out association to a BD
- Export Route Control Subnet
- Route map like default-export (best) or named route-map

[in 4.2 and after](#)

With per-peer route-map

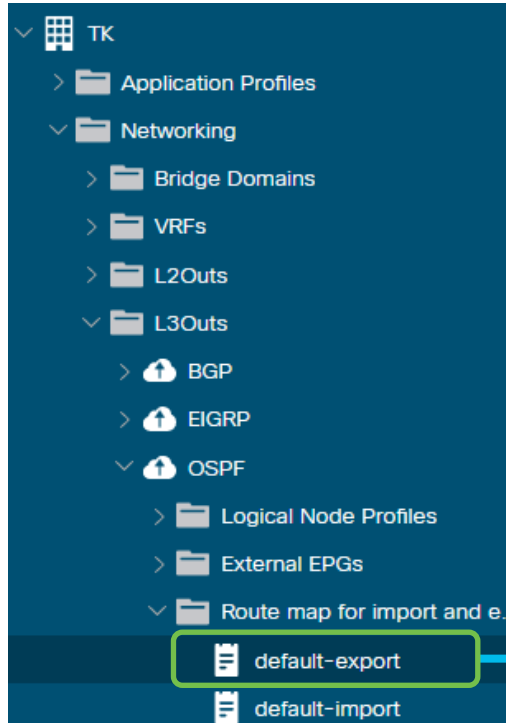
<tenant>-<route_map>-<L3Out>-out

- Non-default route map in BGP peer connectivity profile
- Override regular ACI behavior (subnet flags, BD to L3 out association)

default-export route map configuration



All route advertisement (both BD subnets and transit routing) in one single component while L3Out external EPGs are dedicated for security.



Name: default-export

Type: Match Prefix AND Routing Policy Match Routing Policy Only

Description: optional

Route-Map Continue: ☐ This action will be applied on all the entries which are part of BGP Route-map.

Contexts:

Order	Name	Action
0	BD_SUBNETS	Permit

IP	Description	Aggregate	From Prefix	To Prefix
192.168.1.0/24	BD1	False	0	0
192.168.2.0/24	BD2	False	0	0

BGP per-peer route maps (4.2+)



TK

Quick Start

TK

Application Profiles

Networking

Bridge Domains

VRFs

L2Outs

L3Outs

BGP

Logical Node Profiles

IPv4

BGP Peer 10.51.255.33

BGP Peer 10.51.255.34

Configured Nodes

Logical Interface Profiles

External EPGs

BGP Peer Connectivity Profile 10.51.255.34

Policy

Faults

History

Properties

☐ Replace private AS with local AS

BGP Peer Prefix Policy: select a value

Pre-existing BGP session must be reset to apply the Prefix policy

Site of Origin:

e.g. extended:as2-nn2:1000:65534
e.g. extended:ipv4-nn2:1.2.3.4:65515
e.g. extended:as4-nn2:1000:65505
e.g. extended:as2-nn4:1000:6554387

Local-AS Number Config:

Local-AS Number:

This value must not match the MP-BGP RR policy

Route Control Profile:

Name	Direction
BGP_PEER1	Route Export Policy

TK

Application Profiles

Networking

Contracts

Policies

Protocol

BFD

BFD Multihop

ND RA Prefix

BGP

...

Route Maps for Multicast

Route Maps for Route Control

BGP_PEER1

Route Tag

Set Rules

Ways to advertise routes from an L3Out

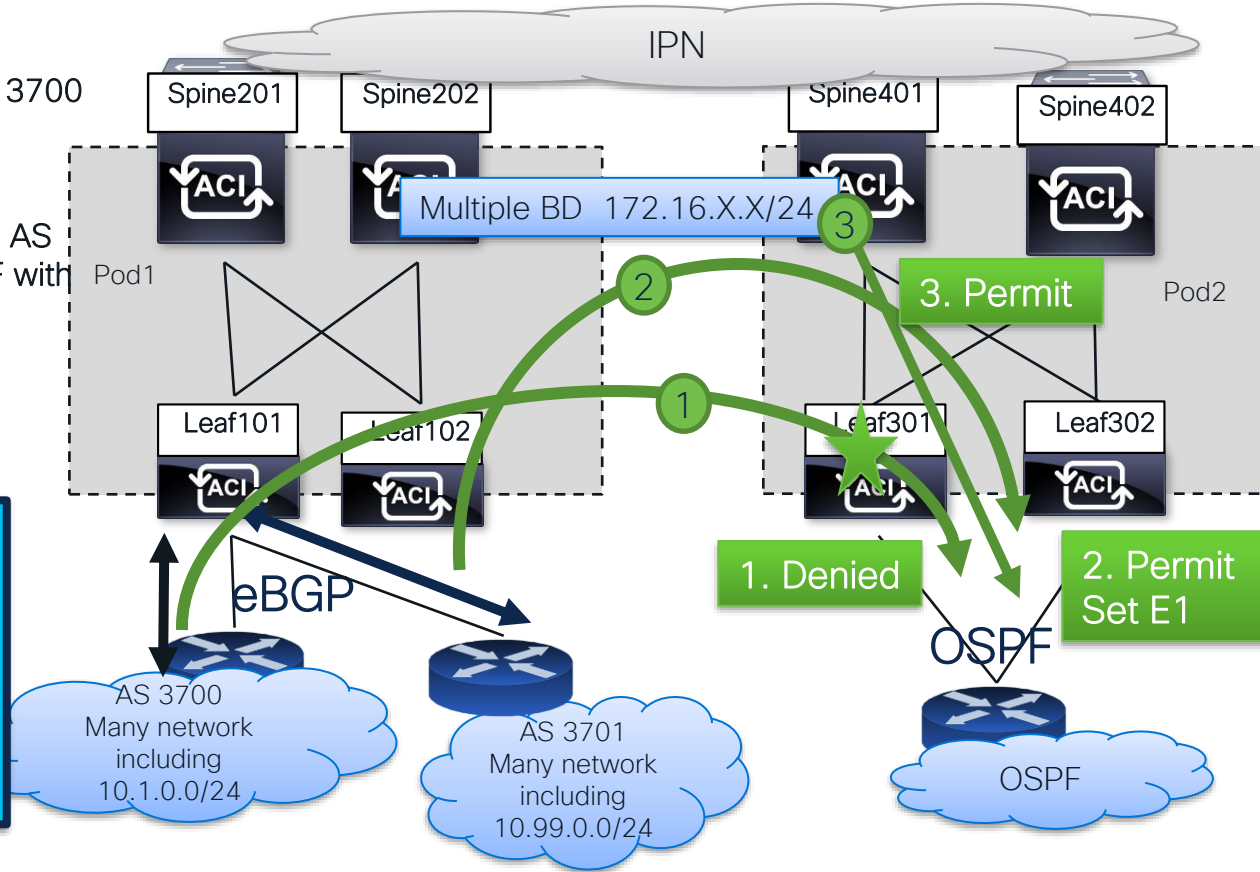
- BD association to an L3Out (see example 1)
- “Export Route Control Subnet” in L3Out EPGs (see example 2)
- Non-default route maps in L3Out EPGs/Subnets (not recommended)
- The default route map (default-export) in an L3Out (example 4)
- Non-default route maps (per-peer route maps) in BGP peer connectivity profile

Example 4 –Routing policies BGP community setting and matching

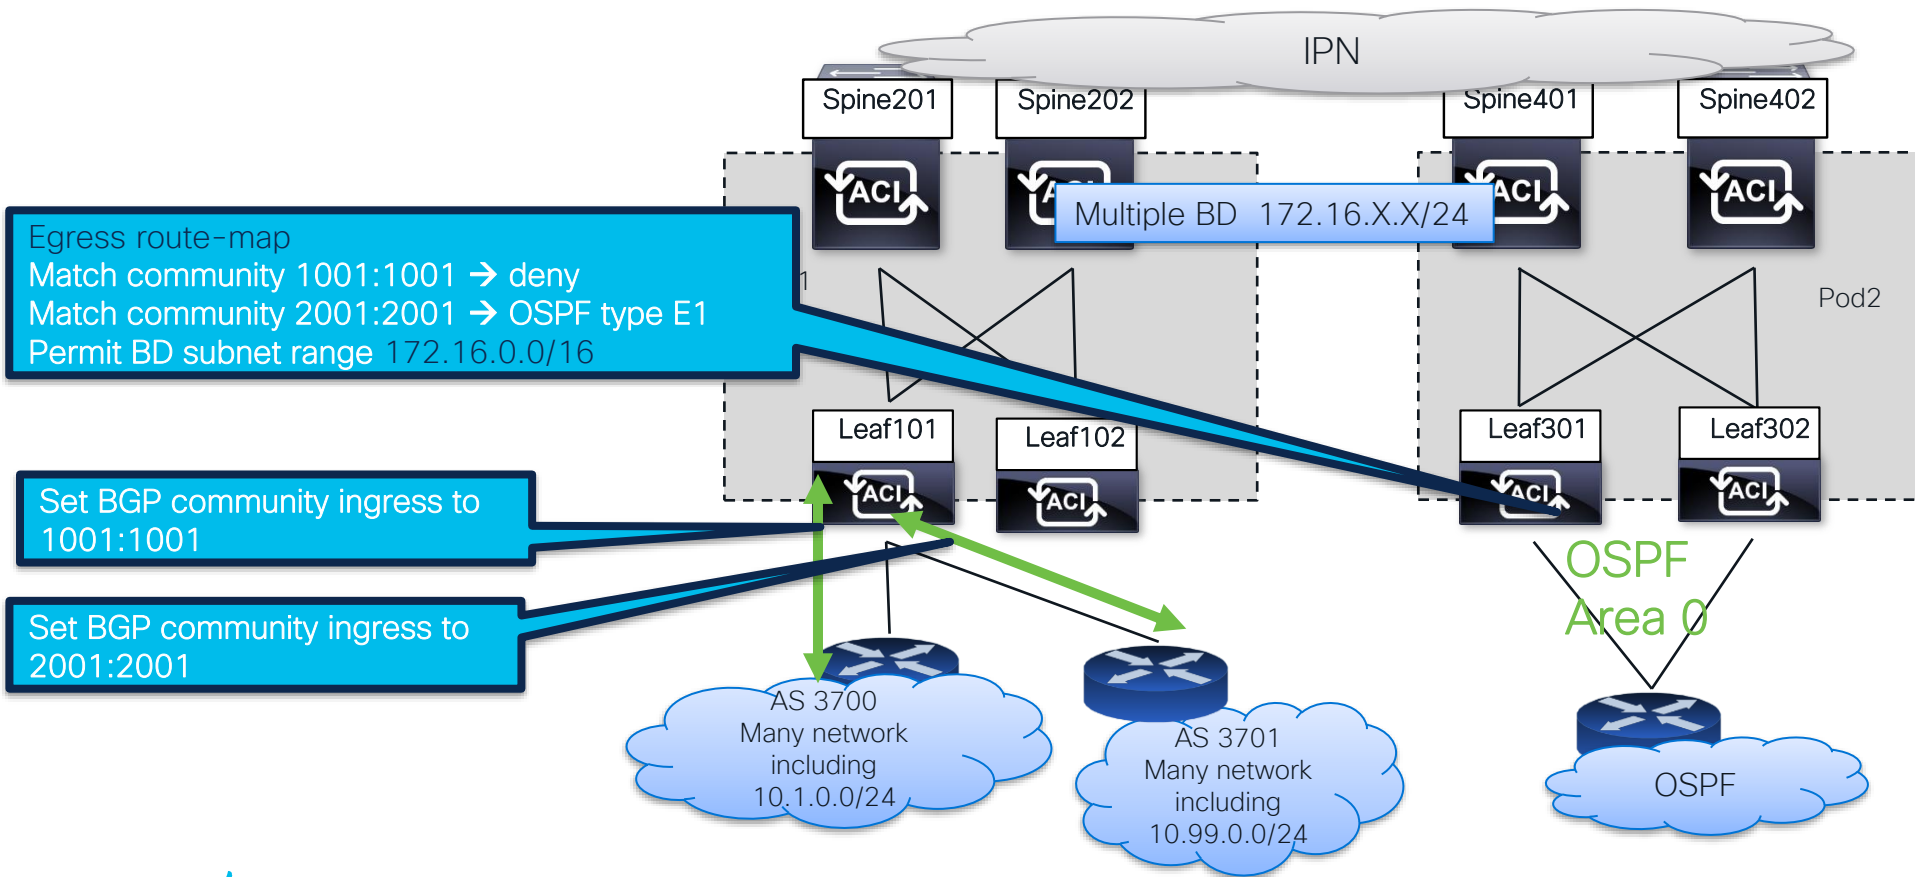
Setup 4

- 1 We receive many network from eBGP AS 3700
We do not want to send those network to OSPF.
- 2 We also have an interconnect with eBGP AS 3701 we do want to send those to OSPF with type E1
- 3 We should send to OSPF all BD subnet

We can't start to list specific prefix received from both BGP peering (too much overlapping prefix) →>
Using BGP community Set and Match



Setup 4 – Set and Match community



Prerequisite to make any ingress route-map

Before enabling import route-control

```
S1P1-Leaf101# show bgp ipv4 unicast neighbor 192.168.101.2 vrf DC:DC | egrep route-map
Inbound route-map configured is permit-all
Outbound route-map configured is exp-l3out-BGP-peer-2359302
```

PIMv6: ☐

Route Control Enforcement: ☒ Import ☐ Export

After enabling import route-control

```
S1P1-Leaf101# show bgp ipv4 unicast neighbor 192.168.101.2 vrf DC:DC | egrep route-map
Inbound route-map configured is imp-l3out-BGP-peer-2359302
Outbound route-map configured is exp-l3out-BGP-peer-2359302
S1P1-Leaf101# show route-map imp-l3out-BGP-peer-2359302
% Policy imp-l3out-BGP-peer-2359302 not found
```

Route-map config – BGP1

default-import

1 Create route-map (here default-import used)

Create Route map for import and export route control

Name:

Type: ☒ Match Prefix AND Routing Policy ☐ Match Routing Policy Only

Description: optional

Route-Map Continue: ☐ This action will be applied on all the entries which are part of BGP Route-map.

Contexts

Order	Name	Action	Description

2 Create permit context to match all and set community

Create Route Control Context

Order:

Name:

Action: ☒ Deny ☒ Permit

Description: optional

Associated Matched

Rule Name
<input checked="" type="checkbox"/> Match-All-BGP1

Set Rule:

3 Create Match Rule

Match All aggregate subnet

Name:

Description: optional

Match Regex Community Terms:

Name	Regular Expression	Community Type	Description
------	--------------------	----------------	-------------

Match Community Terms:

Name	Description
------	-------------

Match Prefix:

IP	Description	Aggregate	Greater than Mask	Less than Mask
<input checked="" type="checkbox"/> 0.0.0.0/0		<input checked="" type="checkbox"/> True	<input type="text" value="0"/>	<input type="text" value="0"/>

4 Set community

Create Set Rules for a Route Map

STEP 1 > Select

Name:

Description: optional

Set Community: ☐

Set Route Tag: ☐

Set Dampening: ☐

Set Weight: ☐

Set Next Hop: ☐

Set Preference: ☐

Set Metric: ☐

Set Metric Type: ☐

Additional Communities: ☒

Next Hop Propagation: ☐

Multipath: ☐

Set External EPG: ☐

Create Set Rules for a Route Map

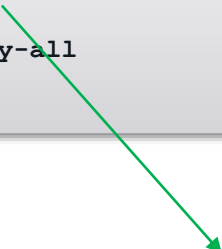
STEP 2 > Additional Communities

Community	Set Criteria	Description
<input checked="" type="checkbox"/> regular:as2-nn2:1001:1001	Append community	

BGP BL resulting route-map

- Sequence matching Prefix-list for all routes and setting community

```
S1P1-Leaf101# show route-map imp-l3out-BGP-peer-2359302
route-map imp-l3out-BGP-peer-2359302, permit, sequence 18201
  Match clauses:
    ip address prefix-lists: IPv4-peer16387-2359302-agg-ext-in-default-import4Set-Community10015Match-All-BGP1-dst
    ipv6 address prefix-lists: IPv6-deny-all
  Set clauses:
    community 1001:1001 additive
```



```
S1P1-Leaf101# show ip prefix-list IPv4-peer16387-2359302-agg-ext-in-default-import4Set-Community10015Match-All-BGP1-dst
ip prefix-list IPv4-peer16387-2359302-agg-ext-in-default-import4Set-Community10015Match-All-BGP1-dst: 1
entries
  seq 1 permit 0.0.0.0/0 le 32
```



Resulting route-map BGP2

Next apply similar config in BGP2 layer 3 out to set community to 2001:2001 for the 2nd AS connection

```
S1P1-Leaf101# show bgp ipv4 unicast neighbors 192.168.201.2 vrf DC:DC | egrep route-map
Inbound route-map configured is imp-l3out-BGP2-peer-2359302, handle obtained
Outbound route-map configured is exp-l3out-BGP2-peer-2359302, handle obtained

S1P1-Leaf101# show route-map imp-l3out-BGP2-peer-2359302
route-map imp-l3out-BGP2-peer-2359302, permit, sequence 18201
Match clauses:
  ip address prefix-lists: IPv4-peer32771-2359302-agg-ext-in-Import-BGP2-SetComm2SetComm10025Match-All-BGP1-dst
  ipv6 address prefix-lists: IPv6-deny-all
Set clauses:
  community 2001:2001 additive

S1P1-Leaf101# show ip prefix-list IPv4-peer32771-2359302-agg-ext-in-Import-BGP2-SetComm2SetComm10025Match-All-BGP1-dst
ip prefix-list IPv4-peer32771-2359302-agg-ext-in-Import-BGP2-SetComm2SetComm10025Match-All-BGP1-dst: 1 entries
seq 1 permit 0.0.0.0/0 le 32
```

BGP route on BL with Community set on both BGP L3 out

```
S1P1-Leaf101# show bgp ipv4 unicast 10.1.0.0/24 vrf DC:DC
BGP routing table information for VRF DC:DC, address family IPv4 Unicast
BGP routing table entry for 10.1.0.0/24, version 103 dest ptr 0xa259ee90
```

```
..
Multipath: eBGP iBGP
```

```
Advertised path-id 1, VPN AF advertised path-id 1
Path type (0xaab9dcf8): external 0x28 0x0 ref 0 adv path ref 2, path is valid, is best path
```


```
AS-Path: 3700, path sourced external to AS
192.168.101.2 (metric 0) from 192.168.101.2 (192.168.0.11)
Origin IGP, MED not set, localpref 100, weight 0 tag 0, propagate 0
Community: 1001:1001
Extcommunity:
RT:101:2359302
VNID:2359302
```

AS 3700

```
S1P1-Leaf101# show bgp ipv4 unicast 10.99.0.0/24 vrf DC:DC
BGP routing table information for VRF DC:DC, address family IPv4 Unicast
BGP routing table entry for 10.99.0.0/24, version 106 dest ptr 0xa25a0788
..
Multipath: eBGP iBGP
```

```
Advertised path-id 1, VPN AF advertised path-id 1
Path type (0xaab9fee0): external 0x28 0x0 ref 0 adv path ref 2, path is valid, is best path
AS-Path: 3701, path sourced external to AS
192.168.201.2 (metric 0) from 192.168.201.2 (192.168.0.111)
Origin IGP, MED not set, localpref 100, weight 0 tag 0, propagate 0
Community: 2001:2001
Extcommunity:
RT:101:2359302
VNID:2359302
```

AS 3701



BGP VPNv4 routes on OSPF BL



```
S1P2-Leaf301# show bgp vpnv4 unicast 10.1.0.0/24 vrf DC:DC
BGP routing table information for VRF overlay-1, address family VPNv4 Unicast
Route Distinguisher: 301:2359302 (VRF DC:DC)
```

```
Advertised path-id 1, VPN AF advertised path-id 1
Path type (0xa25a1c60): internal 0xc0000018 0x40 ref 0 adv path ref 2, path is
valid, is best path
```

```
Imported from (0xa25f74b4) 101:2359302:10.1.0.0/24
```

```
AS-Path: 3700 , path sourced external to AS
```

```
10.0.0.64 (metric 33) from 10.1.96.64 (172.16.2.4)
```

```
Origin IGP, MED not set, localpref 100, weight 0 tag 0, propagate 0
```

```
Received label 0
```

```
Received path-id 2
```

```
Community: 1001:1001
```


```
Extcommunity:
```

```
RT:101:2359302
```

```
COST:pre-bestpath:165:2415919104
```

```
VNID:2359302
```

```
Originator: 10.0.0.64 Cluster list: 172.16.2.
```



```
S1P2-Leaf301# show bgp vpnv4 unicast 10.99.0.0/24 vrf DC:DC
```

```
BGP routing table information for VRF overlay-1, address family VPNv4 Unicast
Route Distinguisher: 301:2359302 (VRF DC:DC)
```

```
..
```

```
Advertised path-id 1, VPN AF advertised path-id 1
```

```
Path type (0xa25a2230): internal 0xc0000018 0x40 ref 0 adv path ref 2, path
is valid, is best path
```

```
Imported from (0xa25f5f64) 101:2359302:10.99.0.0/24
```

```
AS-Path: 3701 , path sourced external to AS
```

```
10.0.0.64 (metric 33) from 10.1.96.64 (172.16.2.4)
```

```
Origin IGP, MED not set, localpref 100, weight 0 tag 0, propagate 0
```

```
Received label 0
```

```
Received path-id 2
```

```
Community: 2001:2001
```

```
Extcommunity:
```

```
RT:101:2359302
```

```
COST:pre-bestpath:165:2415919104
```

```
VNID:2359302
```

```
Originator: 10.0.0.64 Cluster list: 172.16.2.4 172.16.100.101
```


OSPF L3 out – Use default-export Route-map with 3 sequence

- Deny all prefix AND match community 1001:1001
- Permit all prefix AND match community 2001:2001 + set OSPF type E1
- Permit all the rest (matching BD subnet range say 172.16.0.0/16)

Route Control Profile - default-export

Properties

Name: default-export

Type: Match Prefix AND Routing Policy Match Routing Policy Only

Description: optional

Route-Map Continue: ☐ This action will be applied on all the entries which are part of BGP Route-map.

Contexts:

Order	Name	Action	Description
3	DenyComm1001	Deny	
5	MatchComm2001	Permit	
7	PermitBD	Permit	

Default-export is
automatically applied
outbound on the L3out

```
S1P2-Leaf301# show ip ospf vrf DC:DC | egrep bgp  
bgp route-map exp-ctx-proto-2359302
```

Sequence 1

Create Route Control Context

Order: 3

Name: DenyComm1001

Action: **Deny** Permit

Description: optional

Associated Matched Rules:

Rule Name
MatcComm1001

Set Rule: select a value

Create Match Rule

Name: MatcComm1001

Description: optional

Match Regex Community Terms:

Name	Regular Expression	Community Type	Desc
------	--------------------	----------------	------

Match Community Terms:

Name	Description
Comm1001	

Match Prefix:

IP	Description	Aggregate	Greater than Mask	Less than Mask
0.0.0.0/0		True	0	0

Match Community Term

Properties

Name: Comm1001

Description: optional

Match Community Factors:

Community	Scope
regular:as2-nn2:1001:1001	Transitive

```
route-map exp-ctx-proto-2359302, deny, sequence 17201
```

```
Match clauses:
```

```
ip address prefix-lists: IPv4-proto32770-2359302-agg-ext-out-default-export4DenyComm10013MatcComm1001-dst
```

```
ipv6 address prefix-lists: IPv6-deny-all
```

```
community (community-list filter): proto32770-2359302-agg-ext-out-default-export4DenyComm10013MatcComm1001-rgcom
```

```
Set clauses:
```

```
S1P2-Leaf301# show ip community-list proto32770-2359302-agg-ext-out-default-export4DenyComm10013MatcComm1001-rgcom  
Standard Community List proto32770-2359302-agg-ext-out-MatchCommOut2DenyComm10013MatcComm1001-rgcom
```

```
permit 1001:1001
```

```
S1P2-Leaf301# show ip prefix-list IPv4-proto32770-2359302-agg-ext-out-default-export4DenyComm10013MatcComm1001-dst  
ip prefix-list IPv4-proto32770-2359302-agg-ext-out-MatchCommOut2DenyComm10013MatcComm1001-dst: 1 entries
```

```
seq 1 permit 0.0.0.0/0 le 32
```

Sequence 2

Create Route Control Context

Order: 5

Name: MatchComm2001

Action: ☐ Deny ☒ Permit

Description: optional

Associated Matched Rules:

Rule Name
MatchComm2001

Set Rule: SetOSPF-E1

Create Match Rule

Name: MatchComm2001

Description: optional

Match Regex Community Terms:

Name	Regular Expression	Community Type	Description

Match Community Terms:

Name	Description
MatchComm2001	

Match Prefix:

IP	Description	Aggregate	Greater than Mask	Less than Mask
0.0.0.0/0		True	0	0

Create Match Community Term

Name: MatchComm2001

Description: optional

Match Community Factors:

Community	Scope
regular.as2-nn2:2001:2001	Transitive

Create Set Rules for a Route Map

STEP 1 > Select

Name: SetOSPF-E1

Description: optional

Set Community: ☐

Set Route Tag: ☐

Set Dampening: ☐

Set Weight: ☐

Set Next Hop: ☐

Set Preference: ☐

Set Metric: ☒

Set Metric Type: ☒ Metric Type: OSPF type1 metric

Additional Communities: ☐

```
route-map exp-ctx-proto-2359302, permit, sequence 18201
```

Match clauses:

```
ip address prefix-lists: IPv4-proto32770-2359302-agg-ext-out-default-export4MatchComm20015MatchComm2001-dst
```

```
ipv6 address prefix-lists: IPv6-deny-all
```

```
community (community-list filter): proto32770-2359302-agg-ext-out-default-export4MatchComm20015MatchComm2001-rgcom
```

Set clauses:

```
tag 4294967295
```

```
metric-type type-1
```

```
S1P2-Leaf301# show ip community-list proto32770-2359302-agg-ext-out-default-export4MatchComm20015MatchComm2001-rgcom
```

```
Standard Community List proto32770-2359302-agg-ext-out-default-export4MatchComm20015MatchComm2001-rgcom
```

```
permit 2001:2001
```

```
S1P2-Leaf301# show ip prefix-list IPv4-proto32770-2359302-agg-ext-out-default-export4MatchComm20015MatchComm2001-dst
```

```
ip prefix-list IPv4-proto32770-2359302-agg-ext-out-default-export4MatchComm20015MatchComm2001-dst: 1 entries
```

```
seq 1 permit 0.0.0.0/0 le 32
```

Sequence 3

Create Route Control Context

Order: 7

Name: PermitBD

Action: Deny Permit

Description: optional

Set Rule: select a value

Associated Matched

Rules:

Rule Name

MatchBd

Create Match Rule

Name: MatchBd

Description: optional

Match Regex Community Terms:

Name	Regular Expression	Community Type	Description
------	--------------------	----------------	-------------

Match Community Terms:

Name	Description
------	-------------

Match Prefix:

IP	Description	Aggregate	Greater than Mask	Less than Mask
172.16.0.0/16		True	0	0

```
route-map exp-ctx-st-2359302, permit, sequence 11001
```

```
Match clauses:
```

```
ip address prefix-lists: IPv4-st32770-2359302-exc-ext-out-default-export4PermitBD7MatchBd-dst
```

```
ipv6 address prefix-lists: IPv6-deny-all
```

```
Set clauses:
```

```
tag 4294967295
```

```
S1P2-Leaf301# show ip prefix-list IPv4-st32770-2359302-exc-ext-out-default-export4PermitBD7MatchBd-dst
```

```
ip prefix-list IPv4-st32770-2359302-exc-ext-out-default-export4PermitBD7MatchBd-dst: 1 entries
```

```
seq 1 permit 172.16.0.0/16 le 32
```

Full route-map used from BGP to OSPF

proto route-map

```
S1P2-Leaf301# show route-map exp-ctx-proto-2359302
```

Permit BD subnet

```
..
```

```
route-map exp-ctx-proto-2359302, permit, sequence 11001
```

```
Match clauses:
```

```
ip address prefix-lists: IPv4-proto32770-2359302-exc-ext-out-default-export4PermitBD7MatchBd-dst
```

```
ipv6 address prefix-lists: IPv6-deny-all
```

```
Set clauses:
```

```
tag 4294967295
```

```
route-map exp-ctx-proto-2359302, deny, sequence 17201
```

```
Match clauses:
```

```
ip address prefix-lists: IPv4-proto32770-2359302-agg-ext-out-default-export4DenyComm10013MatcComm1001-dst
```

```
ipv6 address prefix-lists: IPv6-deny-all
```

```
community (community-list filter): proto32770-2359302-agg-ext-out-default-export4DenyComm10013MatcComm1001-rgcom
```

```
Set clauses:
```

```
route-map exp-ctx-proto-2359302, permit, sequence 18201
```

```
Match clauses:
```

```
ip address prefix-lists: IPv4-proto32770-2359302-agg-ext-out-default-export4MatchComm20015MatchComm2001-dst
```

```
ipv6 address prefix-lists: IPv6-deny-all
```

```
community (community-list filter): proto32770-2359302-agg-ext-out-default-export4MatchComm20015MatchComm2001-rgcom
```

```
Set clauses:
```

```
tag 4294967295
```

```
metric-type type-1
```

Deny community 1001:1001 for all subnet

Set OSPF E1 for Community 1002:1002

Note : order of sequence in route-map have 2 rules:

1. Sequence containing 0.0.0.0/0 le 32 are always after sequence with more specific prefix-list
2. After rule 1, order in route-map adhere the sequence number used in GUI

External OSPF router RIB

```
POD2-router2# show ip route vrf DC:DC
```

```
...  
10.99.0.0/24, ubest/mbest: 2/0
```

```
  *via 192.168.102.1, Vlan942, [110/41], 00:08:03, ospf-1, type-1, tag 4294967295
```

```
  *via 192.168.102.2, Vlan942, [110/41], 00:08:03, ospf-1, type-1, tag 4294967295
```

```
..
```

```
172.16.21.0/24, ubest/mbest: 2/0
```

```
  *via 192.168.102.1, Vlan942, [110/20], 00:23:24, ospf-1, type-2, tag 4294967295
```

```
  *via 192.168.102.2, Vlan942, [110/20], 00:23:24, ospf-1, type-2, tag 4294967295
```

```
172.16.22.0/24, ubest/mbest: 2/0
```

```
  *via 192.168.102.1, Vlan942, [110/20], 00:23:24, ospf-1, type-2, tag 4294967295
```

```
  *via 192.168.102.2, Vlan942, [110/20], 00:23:24, ospf-1, type-2, tag 4294967295
```

```
..
```

```
NO 10.1.0.0/24 (filtered by outbound ospf route-map)
```

Set by the match comm statement

Bd subnet match in Route-map (no set) regular E2

10.1.0.0 is not in RIB as filtered by deny route-map matching community

Summary



Route Control Strategy – Approach 1a and 1b

ACI Day 0 implementation

Pros :


Easy to deploy subnet – All controls under BD

Cons:

On L3out itself no visibility or information on which BD subnet will be send out

Aggregation of BD subnet in prefix-list is impossible

1a – BD subnet on BD – Transit on L3out



Pros :

L3out can decide which BD is advertise

Cons:

No differentiation between internal BD subnet and transit route, neither in UI, neither in route-map on leaf

Aggregation of subnet in prefix-list possible but requires route-map

1b –Transit and BD subnet on L3out



Route Control Strategy – Approach 2a and 2b Full route-map approach

Pros :

- Closer feeling to regular router
- Very tight control on routing
- Aggregation of subnet in prefix-list very easy
- More scalable – easier to troubleshoot
- Only one configuration place for route control

Cons :

- Hard to migrate from Approach 1 to route-map
- Little more complicated
- For BGP : common route-map for all neighbors on same L3

2a – All protocols (2.1+ code)

BD

- BD Subnet A
- Advertised externally

L3 Out

Route-map (default-export)

- Match prefix (BD and transit)
- Aggregate (optional)
- Set parameters (optional)

Pros :

- Even closer to regular router
- Same Pros as 2a

Cons :

- Even Hard to migrate from Approach 1 to route-map per neighbor → best for greenfield
- None

2b – BGP L3out (4.2+ code)

BD

- BD Subnet A
- Advertised externally


L3 Out

Route-map defined per neighbor

- Match prefix (BD and transit)
- Aggregate (optional)
- Set parameters (optional)

Policy enforcement for Layer 3 summary

Assumption is using default vrf ingress enforcement mode



Apic GUI

External subnet for external Epg in L3out

- Allocate external subnet to External EPG pcTag
- Possibility to use supernet
- GUI refused duplicate subnet with same mask in same VRF
- 0.0.0.0/0 is the exception that can be reused
 - Flexibility BUT risks of unwanted traffic flow

Switch

Show zoning-prefix

- Aggregate all External subnet of all External EPG per VRF
- Distributed to all leaf of the VRF
- Policy enforcement always on non Border-leaf in both direction
- Longest prefix match behavior (LPM)
- 0.0.0.0/0 uses wildcard pcTag 15 in egress and vrf pcTag in ingress

Complete your Session Survey

- Please complete your session survey after each session. Your feedback is very important.
- Complete a minimum of 4 session surveys and the Overall Conference survey (open from Thursday) to receive your Cisco Live t-shirt.
- All surveys can be taken in the Cisco Events Mobile App or by logging in to the Session Catalog and clicking the "Attendee Dashboard" at <https://www.ciscolive.com/emea/learn/sessions/session-catalog.html>



Continue Your Education



Visit the Cisco Showcase for related demos.



Book your one-on-one Meet the Engineer meeting.



Attend any of the related sessions at the DevNet, Capture the Flag, and Walk-in Labs zones.



Visit the On-Demand Library for more sessions at ciscolive.com/on-demand.



The bridge to possible

Thank you

CISCO *Live!*

CISCO *Live!*

ALL IN