Cisco live!

6-9 March 2018 · Melbourne, Australia

ACI
Layer 4-7 Integration

Minako Higuchi

BRKACI-2016

Cisco live!

Session Objectives

- At the end of the session, the participants should be able to:
  - Understand the functionalities and specific design considerations associated to the ACI L4-L7 Service (Firewall, Load Balancer etc) Integration

- Initial assumption:
  - The audience already has a good knowledge of ACI main concepts (Tenant, BD, EPG, L2Out, L3Out, etc.)
  - The audience already has a good knowledge of L4-L7 Service design in traditional (non-ACI) network.

Agenda

- ACI Service Insertion Basics
- Design Use Cases and Considerations
  - Firewall
  - ADC (Load-Balancer)
  - IPS
  - PBR (Policy Based Redirect)
  - Multi-pod + Service
- Upcoming new features

ACI Service Insertion Basics

What is Service Graph?

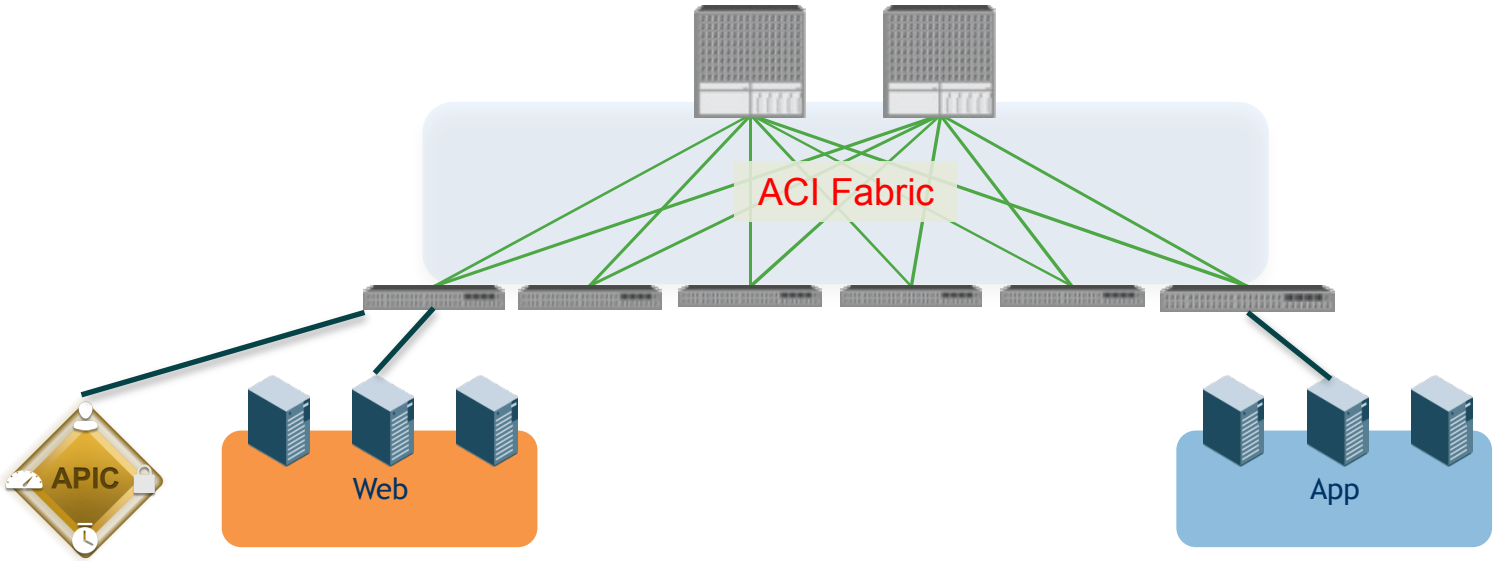

What is Service Graph?

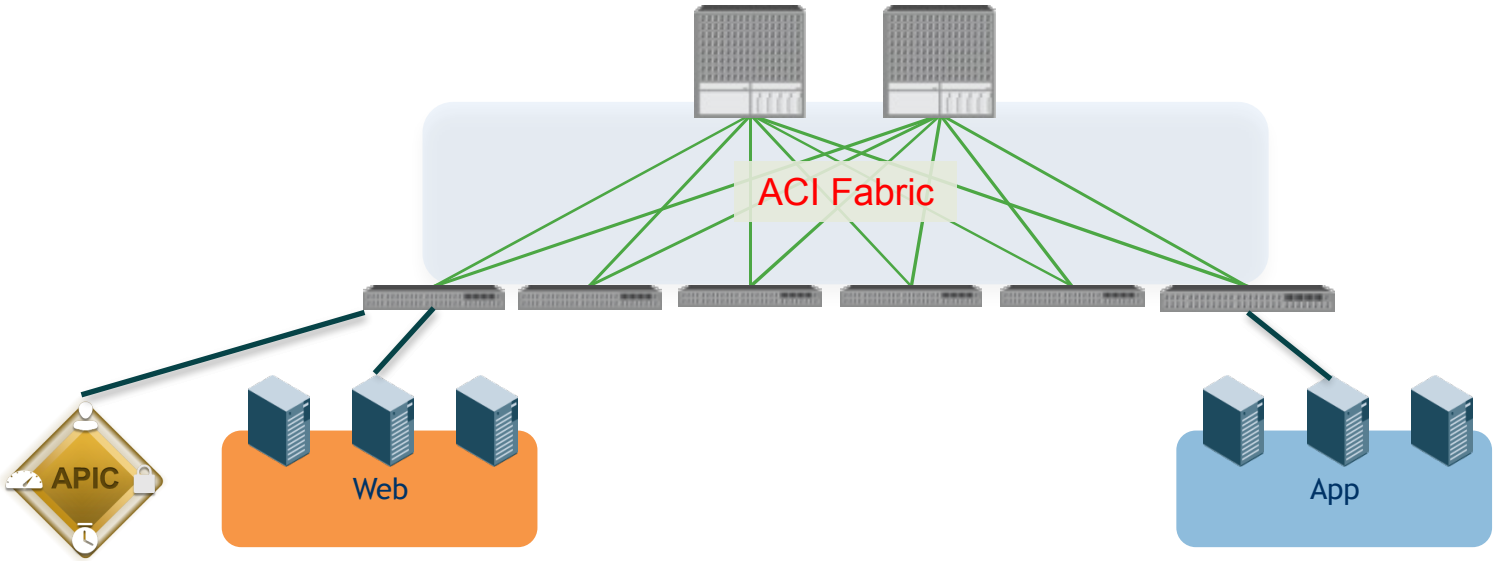

What is Service Graph?

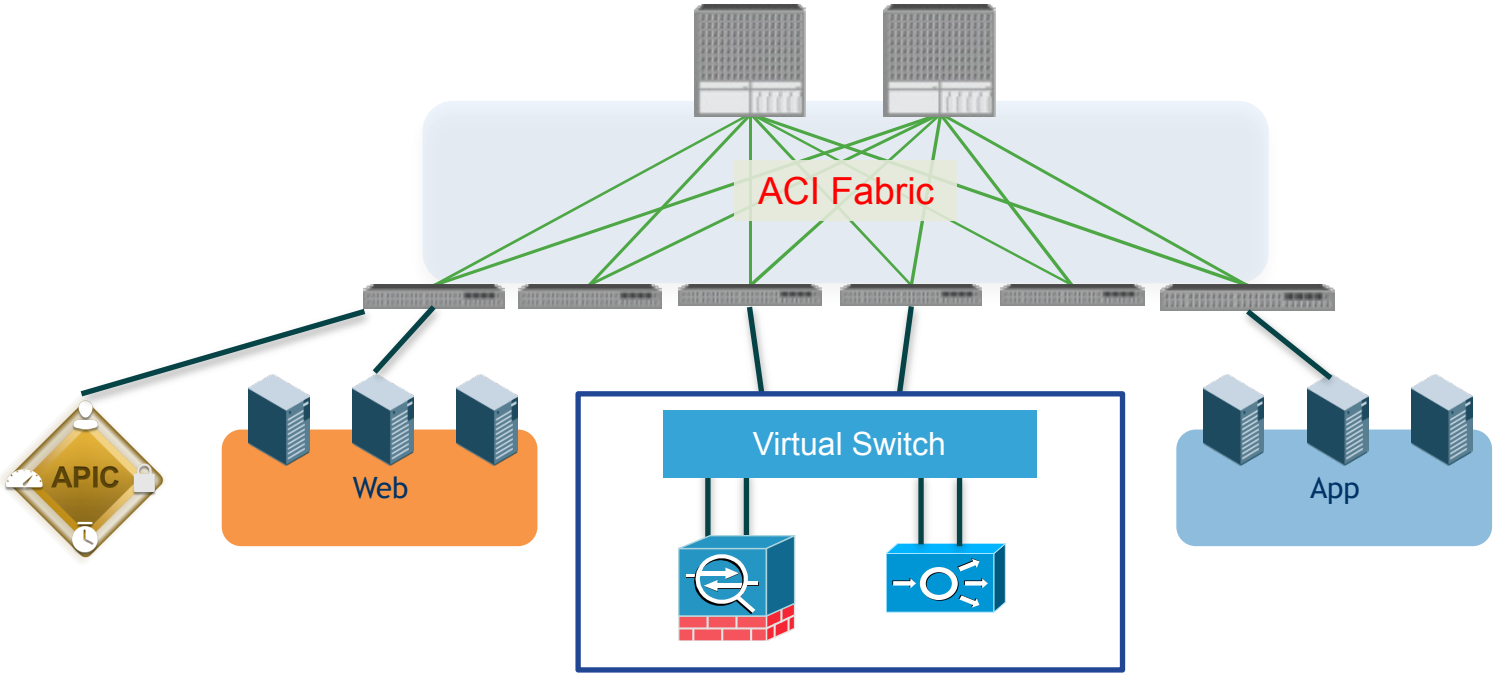

What is Service Graph?

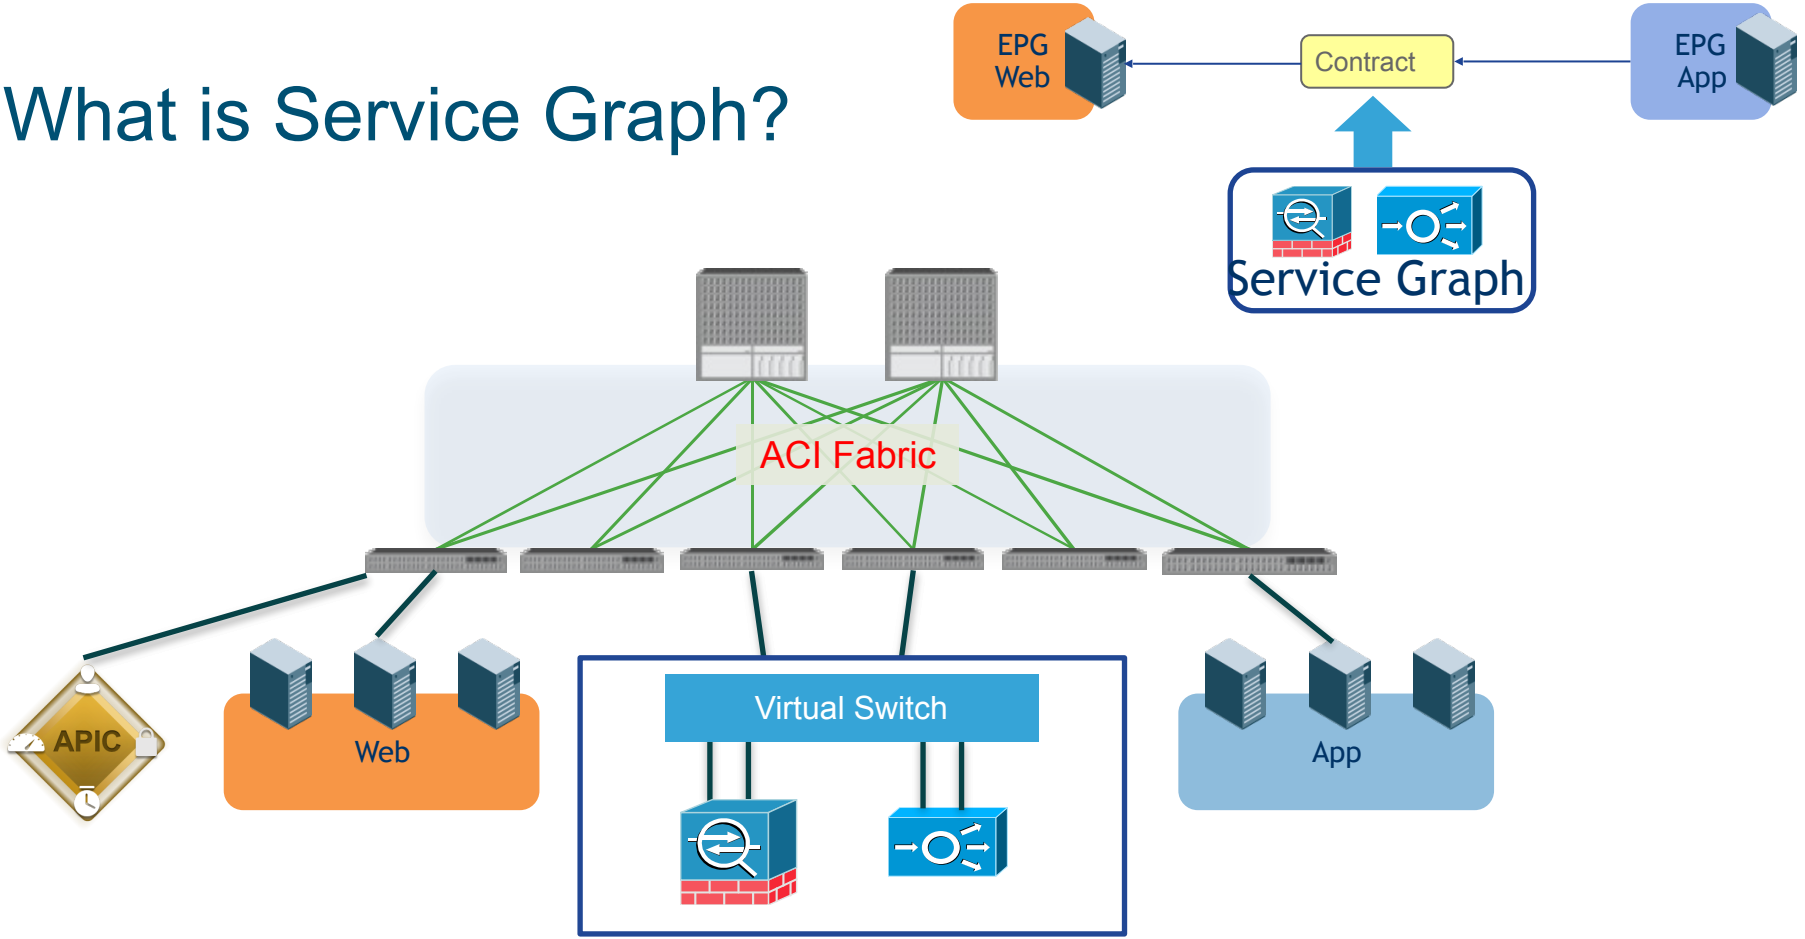

What is Service Graph?

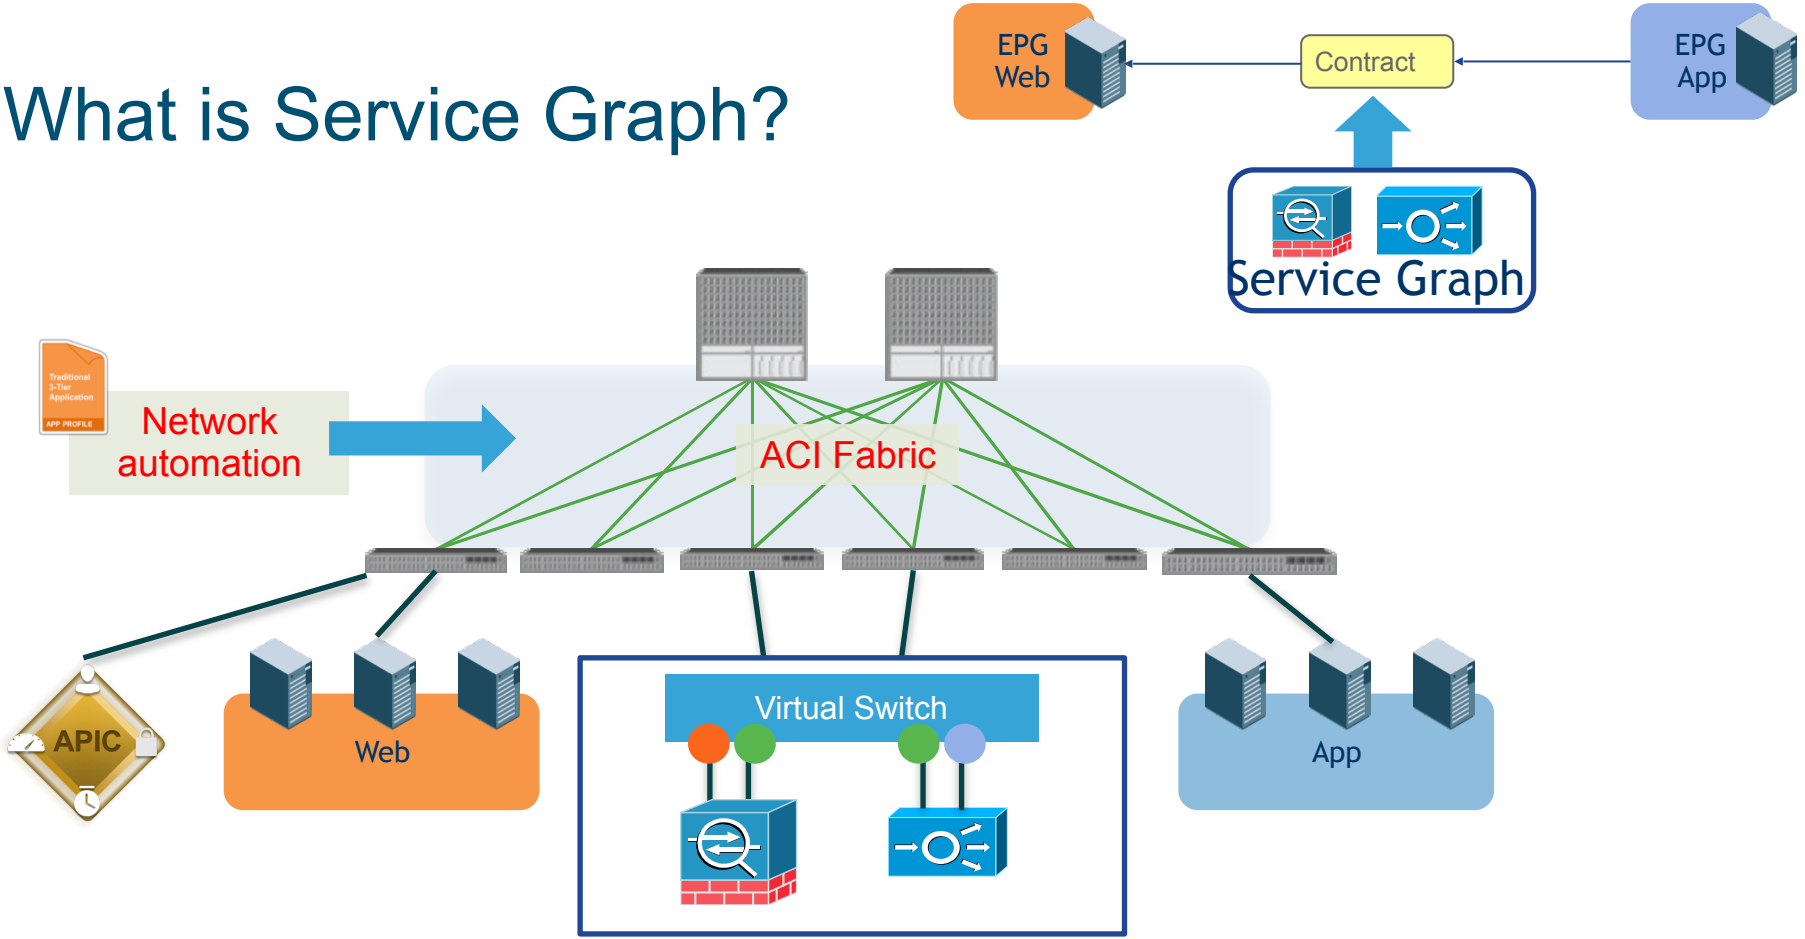

What is Service Graph?

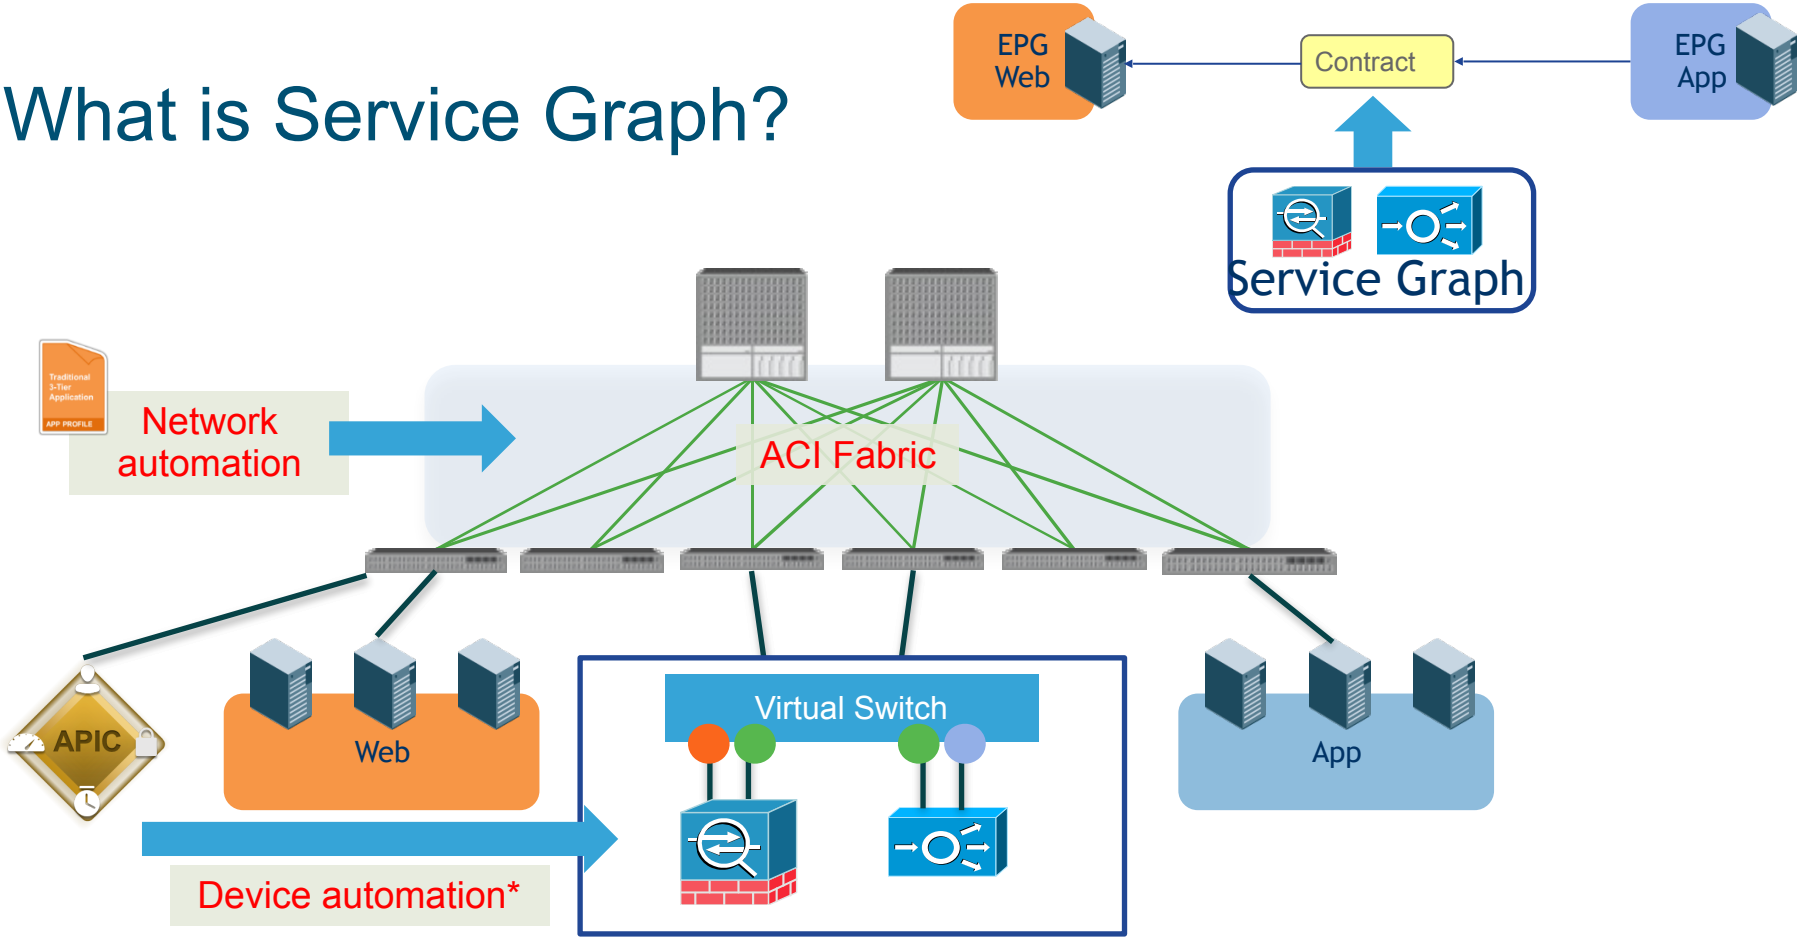

*Managed Mode/Hybrid Mode Only

What is Service Graph?

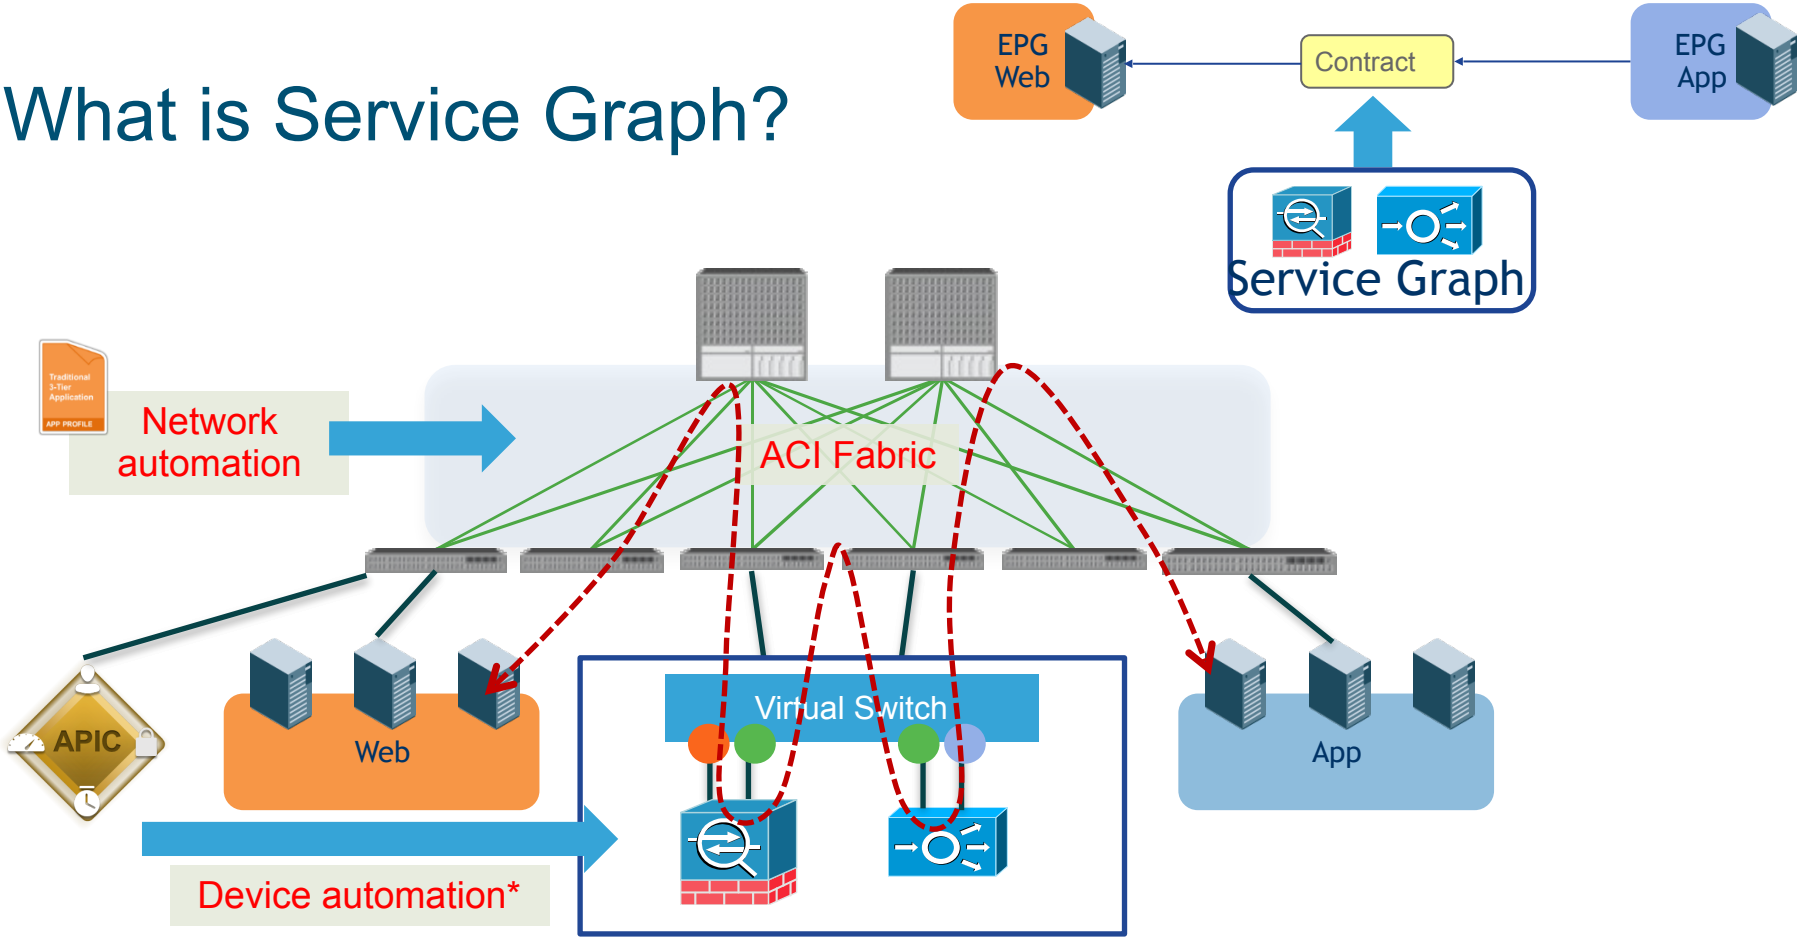

*Managed Mode/Hybrid Mode Only

Service Graph Configuration Workflow

Service Graph Configuration Workflow

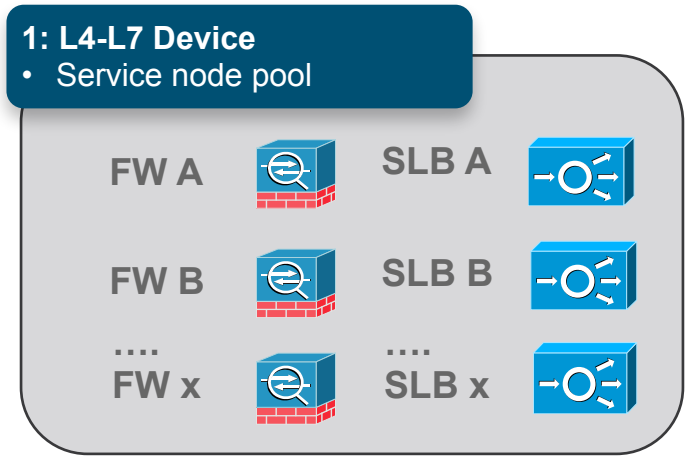

Service Graph Configuration Workflow

2: L4-L7 Service Graph Template
- How to chain function node

1: L4-L7 Device
- Service node pool

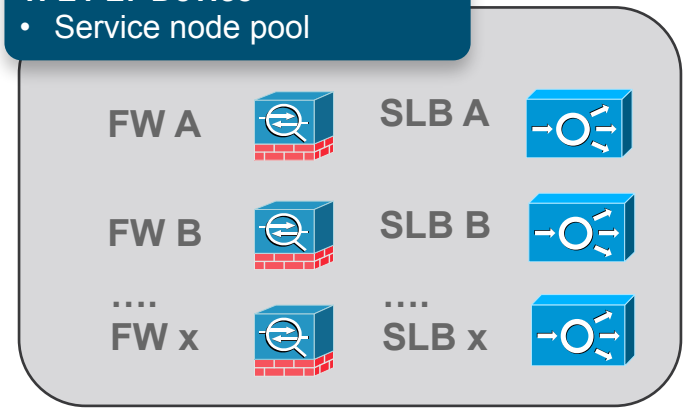

Service Graph Configuration Workflow

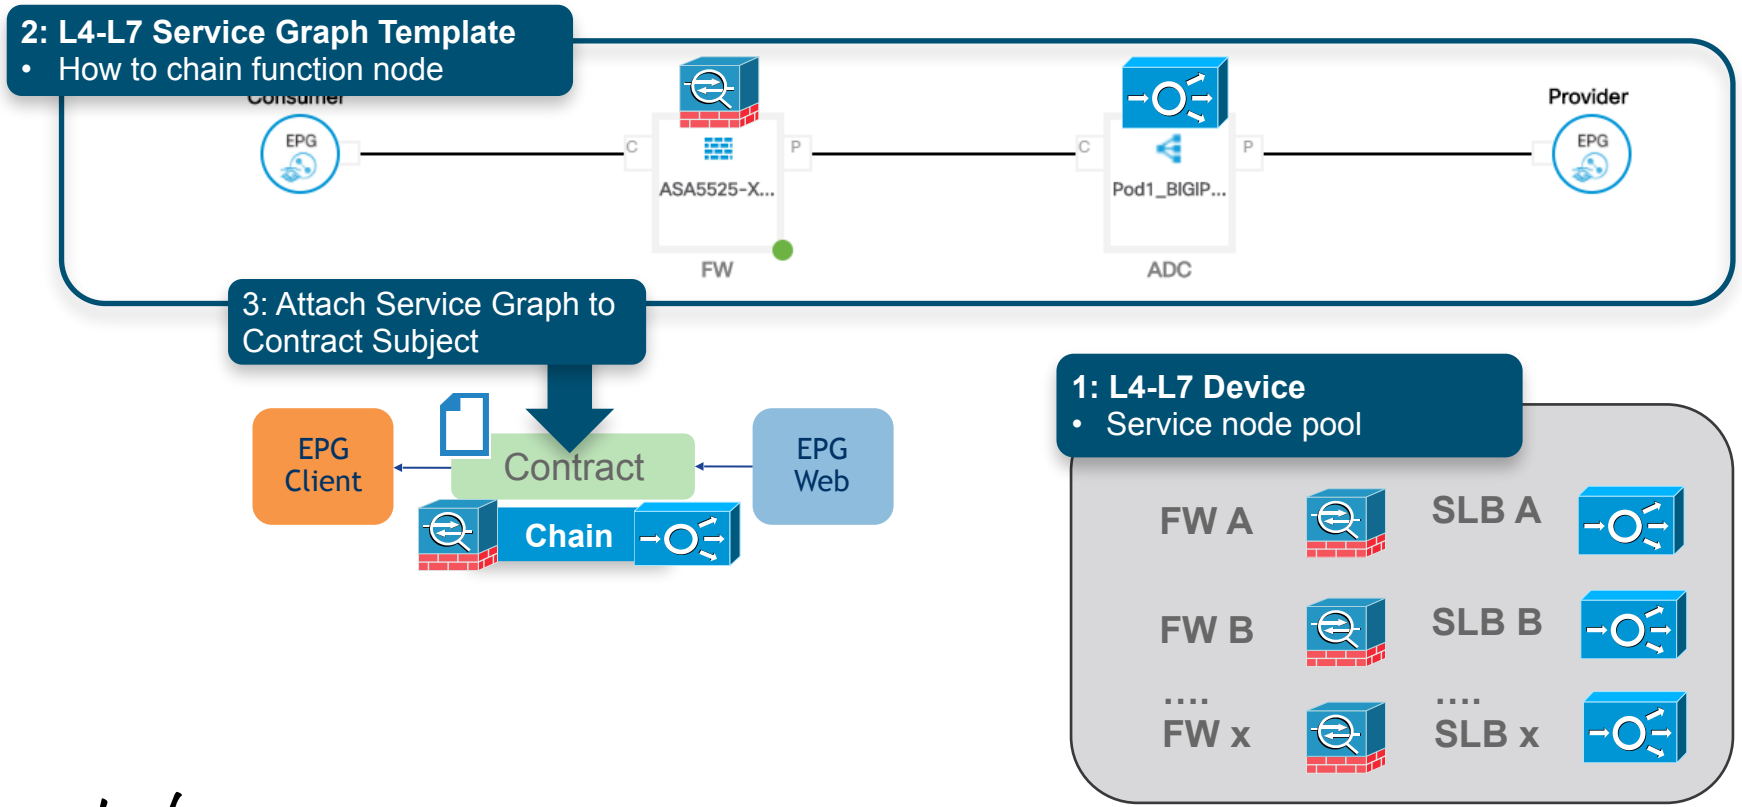

Service Graph Configuration Workflow

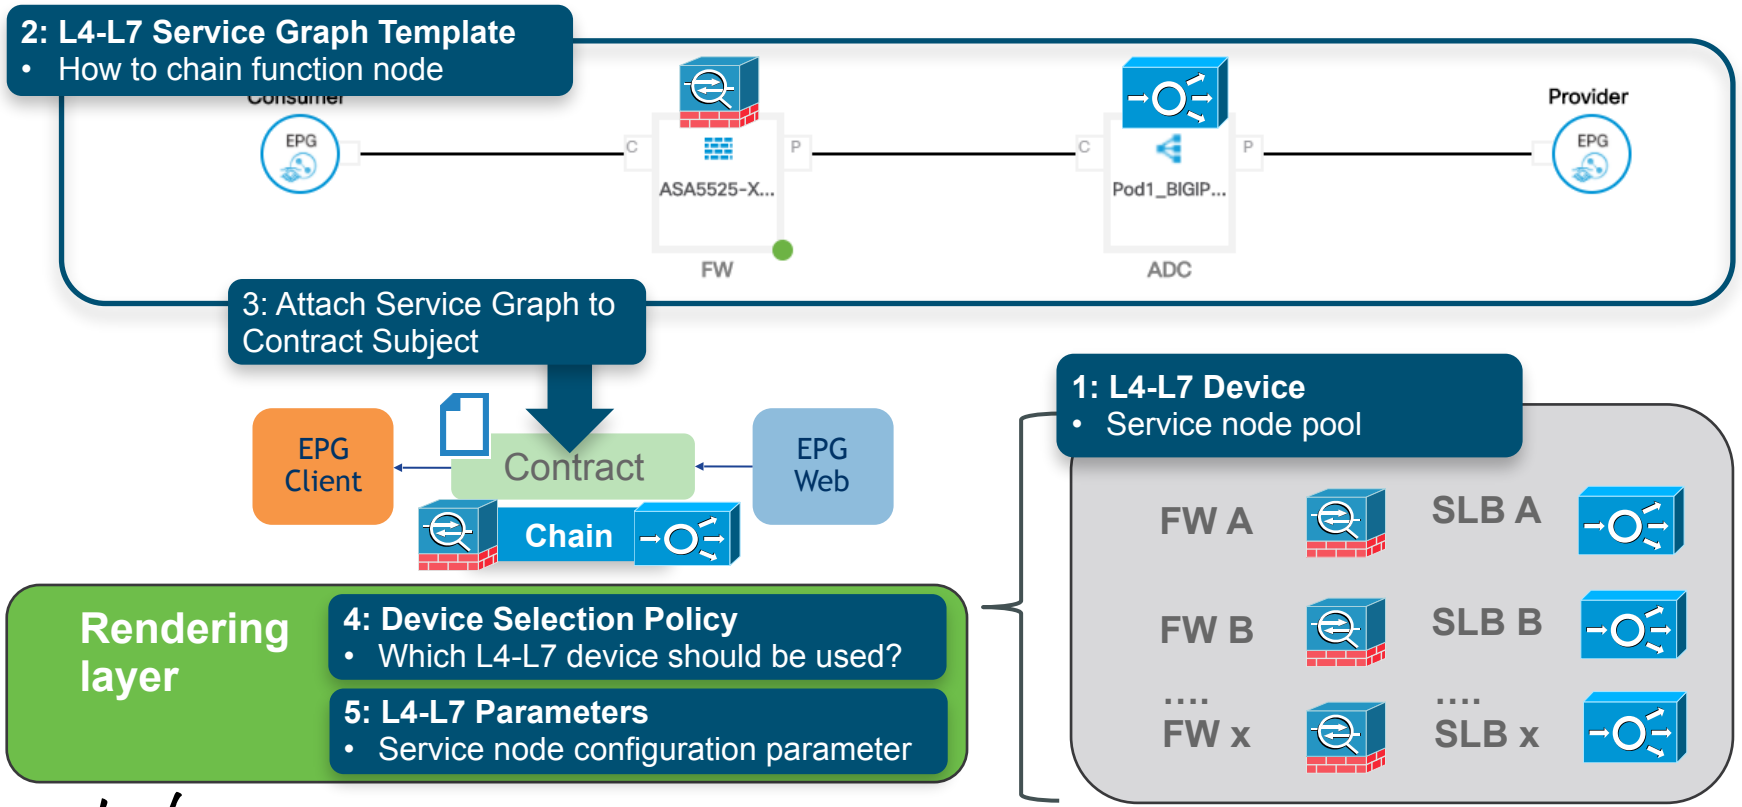

*Managed Mode/Hybrid Mode Only

Device Packages

- Service functions are added to APIC through device package
- Device Package contains a device model and device python scripts
- Device Model defines Service Function and Configuration
- Device scripts translates APIC API callouts to device specific callouts
- Script can interface with the device using REST, SSH or any mechanism

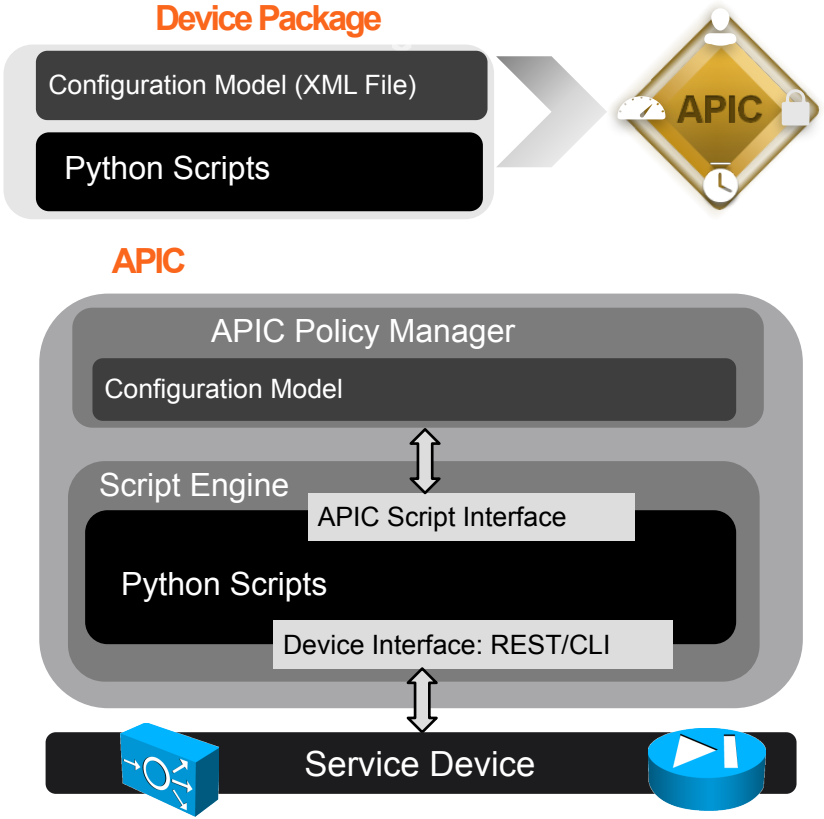

L4-L7 Devices
Logical and Concrete Devices

- Concrete Device
  Represents a service device, e.g. one load balancer, or one firewall. Can be physical or virtual
- Logical Device
  Represents a cluster of 2 devices that operate in active/standby mode for instance.
  Configure logical interfaces (defined in the device model) to be used for device selection policy
  Defines Cluster wide parameters where applicable, such as NTP, DNS, etc.

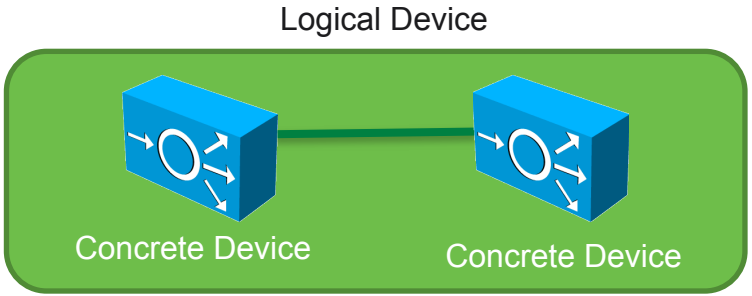

L4-L7 Service Graph Template
Overview

- "Generic" representation of the expected traffic flow
- Defines
  - Function Nodes
  - Terminal Nodes
  - Connections
- Reusable
  - Service Graph template can be used in multiple Contracts.

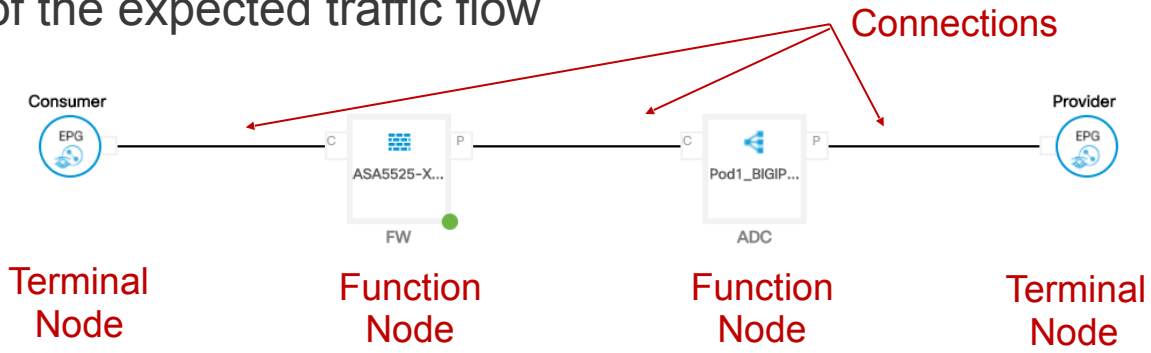

Device Selection Policy
Example

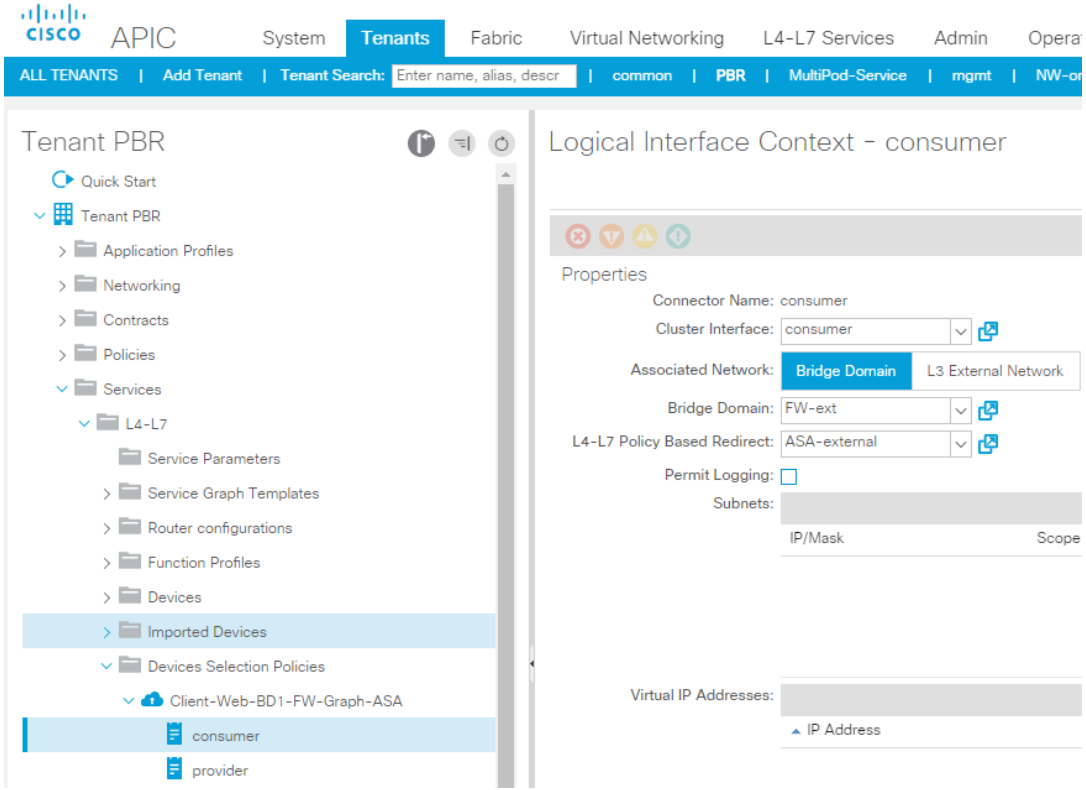

Device Selection Policy
Example

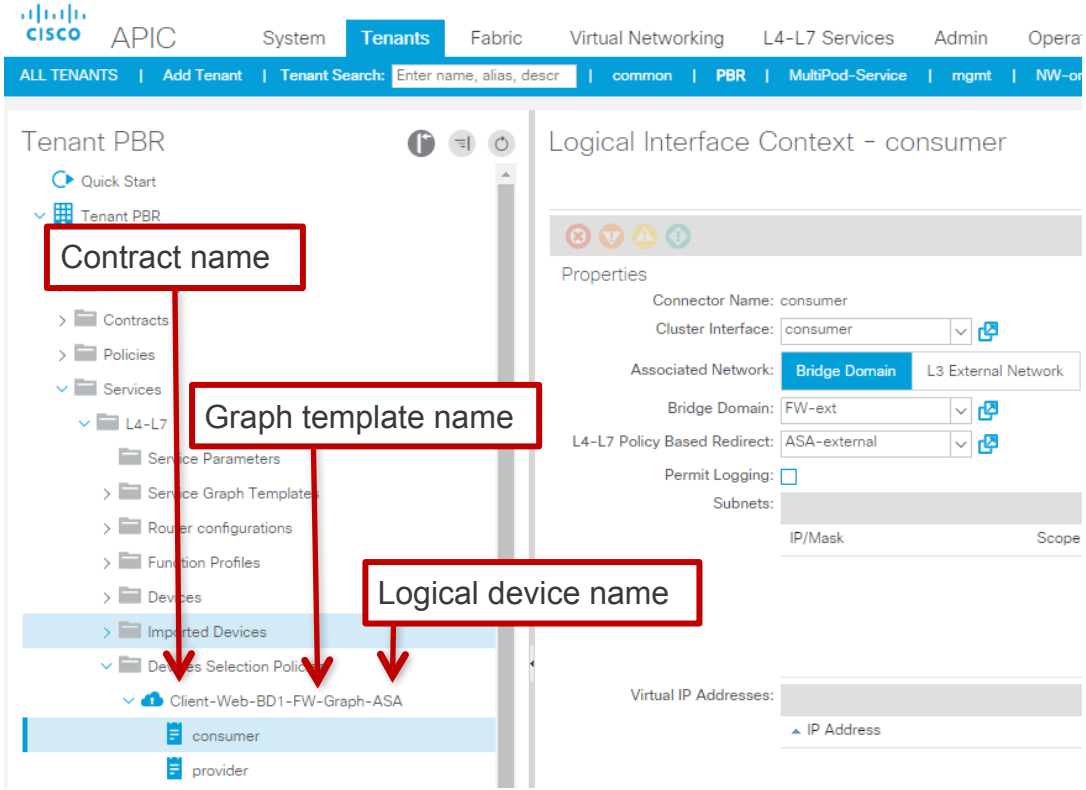

Device Selection Policy
Example

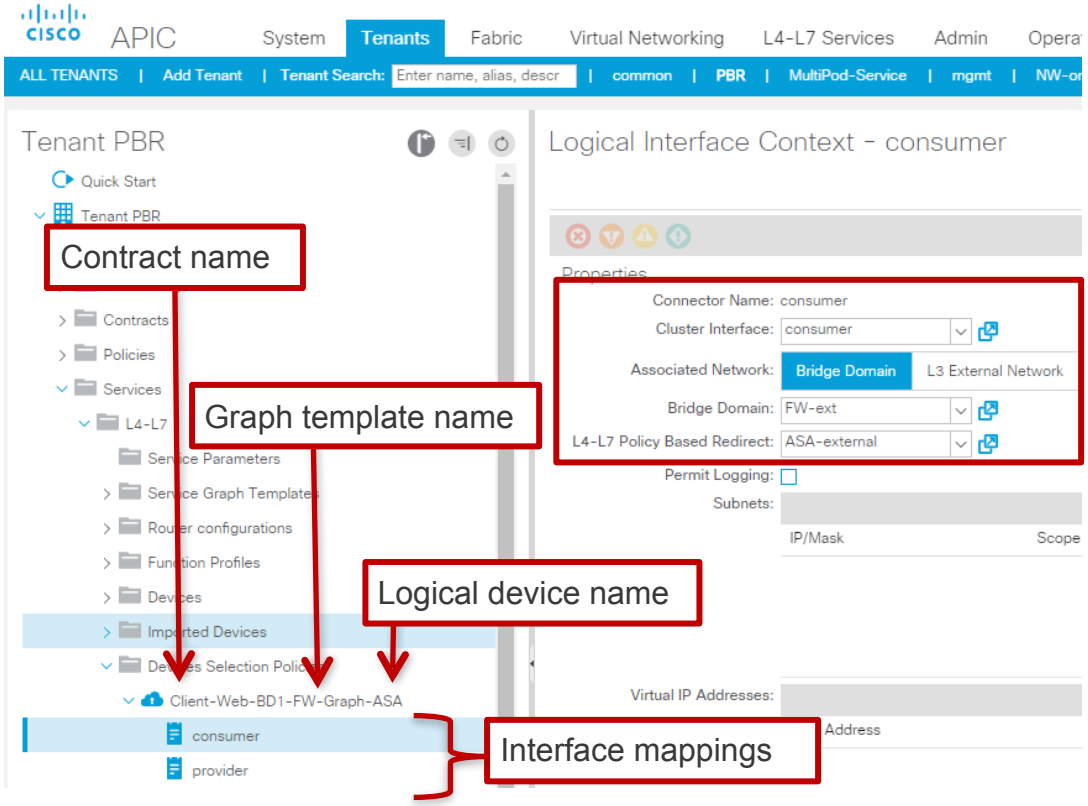

Deployed Graph Instances & Deployed Devices

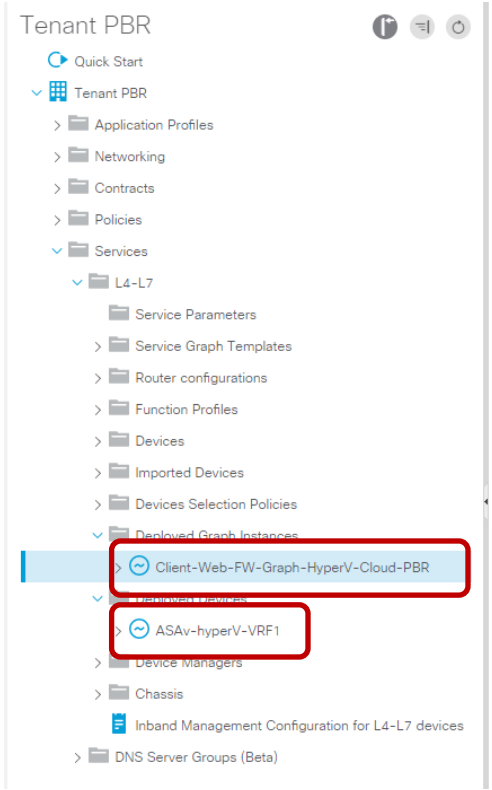

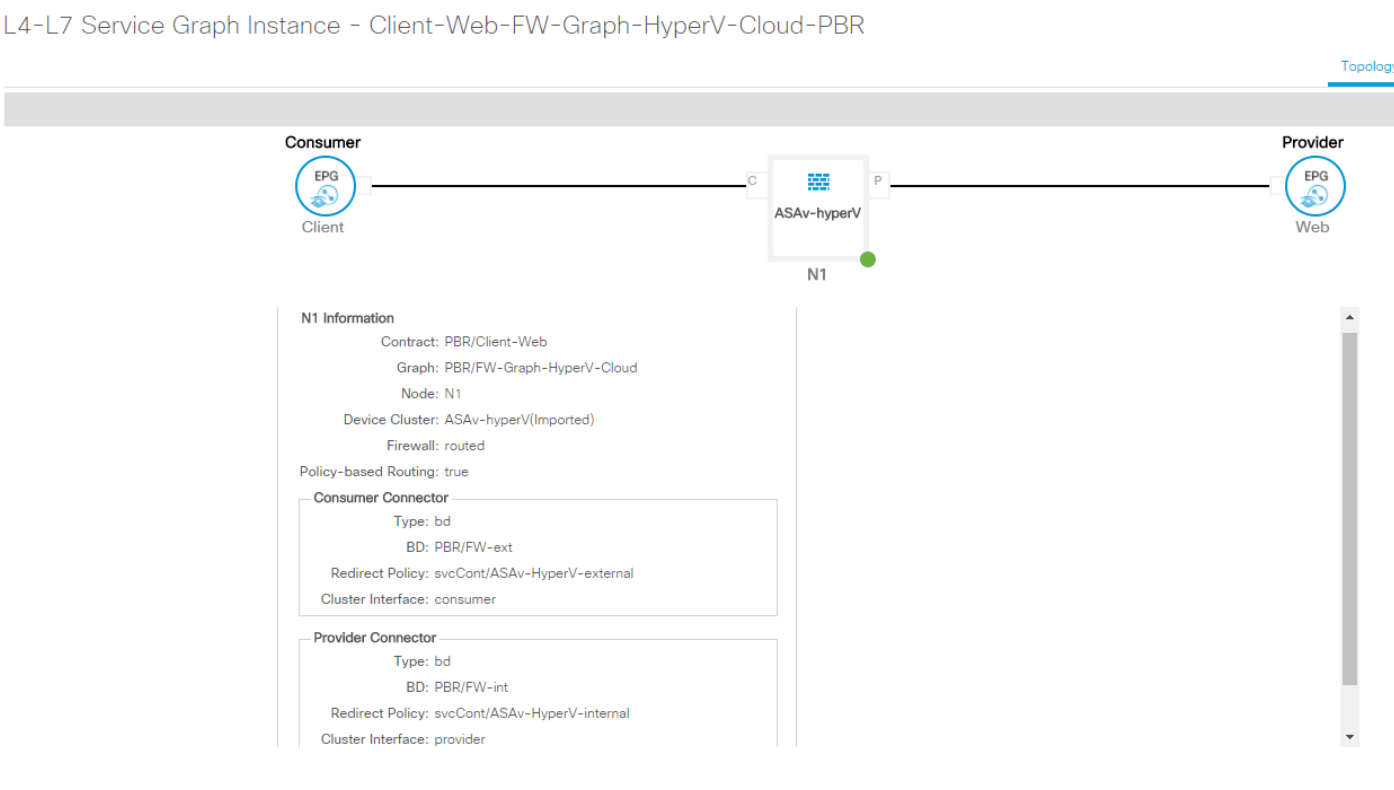

Service Graph Operations
Internal EPGs & Contracts

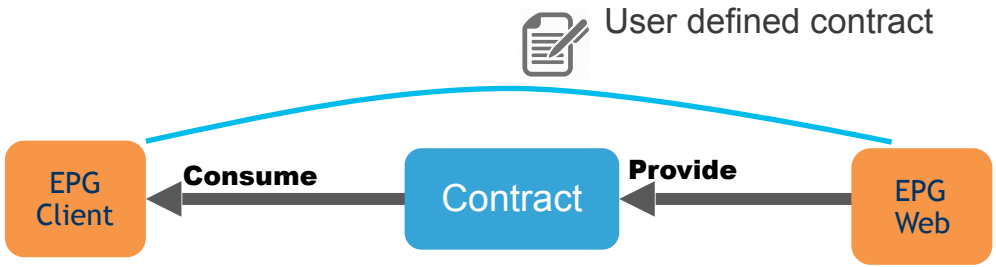

Service Graph Operations
Internal EPGs & Contracts

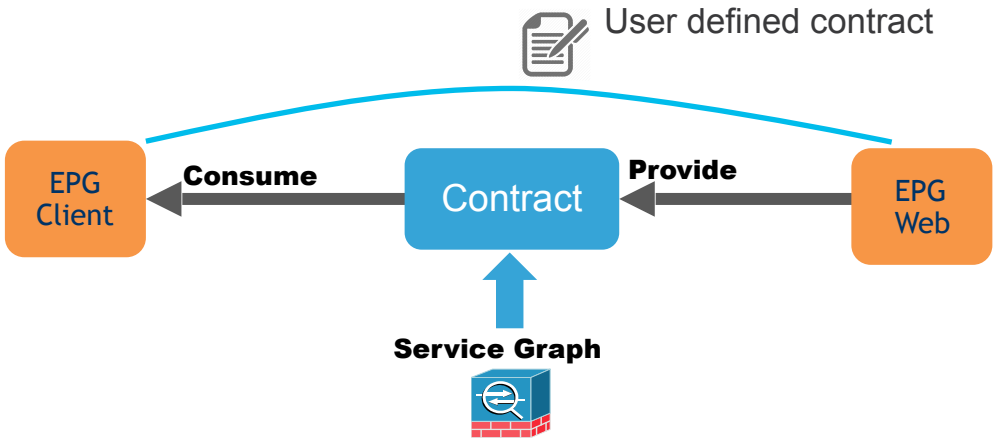

Service Graph Operations
Internal EPGs & Contracts

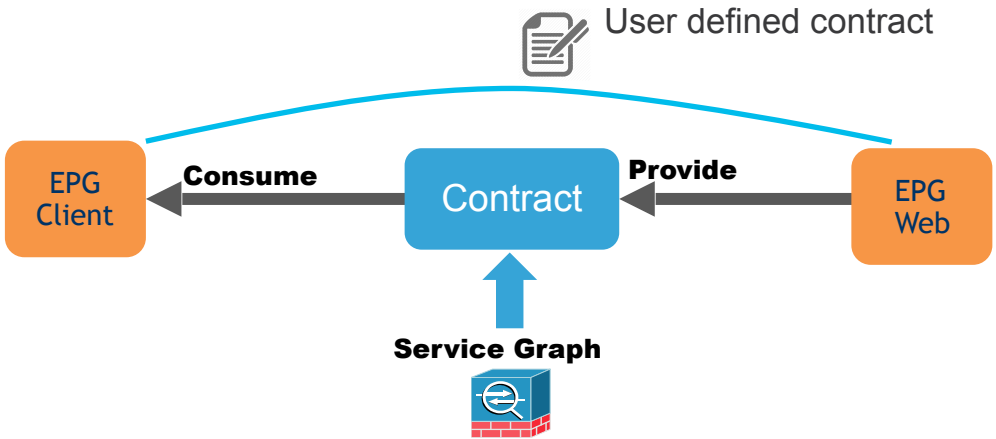

Cluster Interfaces:	Logical Interface	Encap
	ASAv-L2_consumer	vlan-1040
	ASAv-L2_provider	vlan-1006

Service Graph Operations
Internal EPGs & Contracts

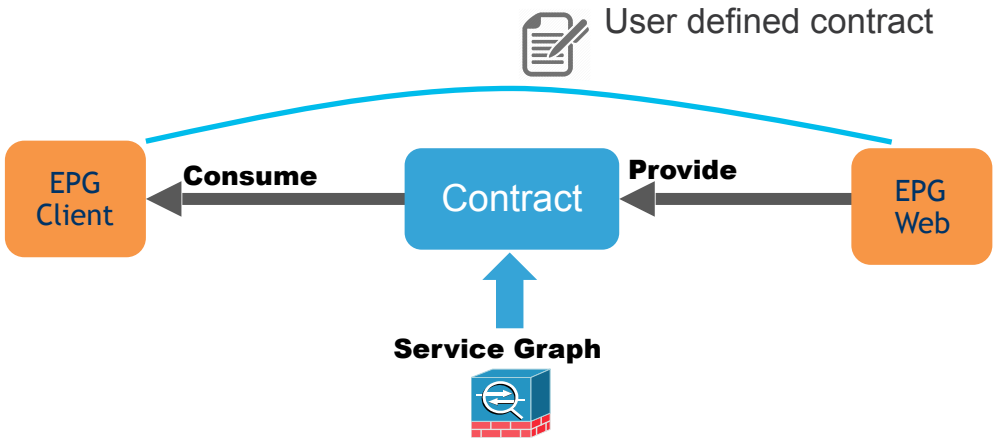

Cluster Interfaces:

Logical Interface	Encap
ASAv-L2_consumer	vlan-1040
ASAv-L2_provider	vlan-1006

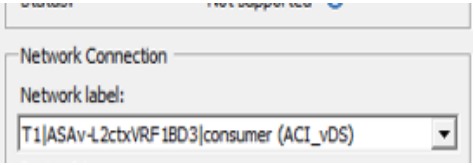

Service Graph Operations
Internal EPGs & Contracts

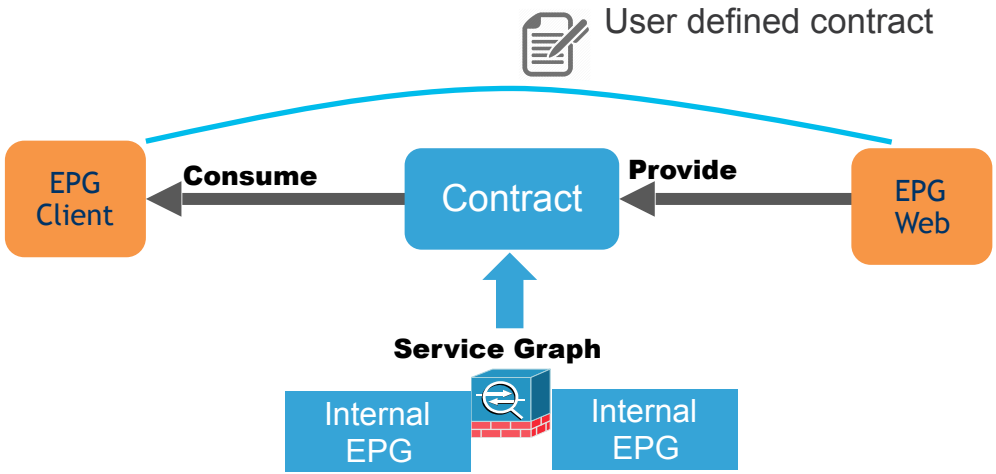

Cluster Interfaces:

Logical Interface	Encap
ASAv-L2_consumer	vlan-1040
ASAv-L2_provider	vlan-1006

Network adapter 2	T1|ASAv-L2ctxVRF1BD3...
Network adapter 3	T1|ASAv-L2ctxVRF1BD1...
Network adapter 4	Management
Network adapter 5	Management

Network Connection

Network label:

T1|ASAv-L2ctxVRF1BD3|consumer (ACI_vDS)

Service Graph Operations
Internal EPGs & Contracts

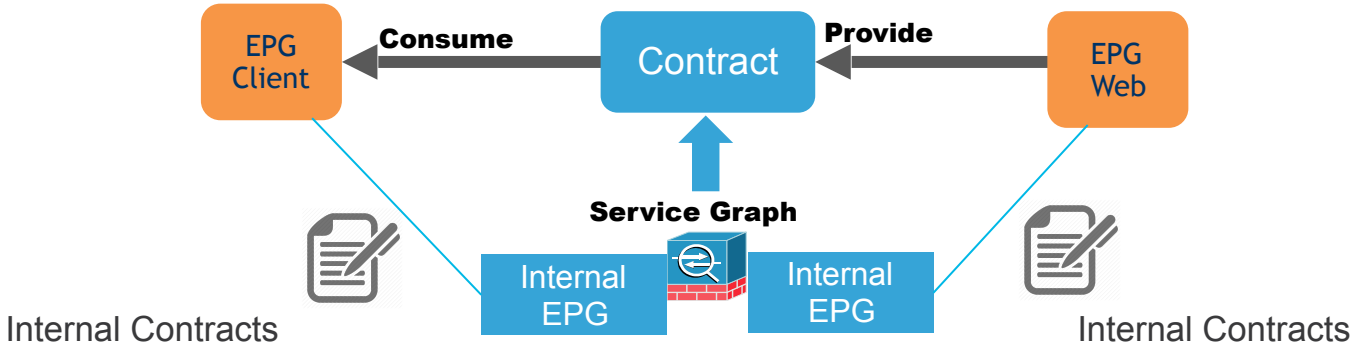

Cluster Interfaces:

Logical Interface	Encap
ASAv-L2_consumer	vlan-1040
ASAv-L2_provider	vlan-1006

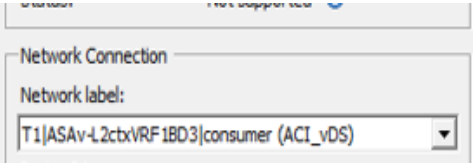

Service Graph Operations
Applied Policy View

Without Service Graph

```
Leaf1# show zoning-rule scope 2621442
Rule ID         SrcEPG          DstEPG          FilterID        operSt          Scope           Action                          Priority
=======         ======          ======          ========        ======          =====           ======                          ========
<snip>
4211            32771           32774           39              enabled         2621442         permit                          fully_qual(7)
4212            32774           32771           38              enabled         2621442         permit                          fully_qual(7)
```

Service Graph Operations
Applied Policy View

Without Service Graph

```
Leaf1# show zoning-rule scope 2621442
Rule ID         SrcEPG          DstEPG          FilterID        operSt          Scope           Action                          Priority
=======         ======          ======          ========        ======          =====           ======                          ========
<snip>
4211            32771           32774           39              enabled         2621442         permit                          fully_qual(7)
4212            32774           32771           38              enabled         2621442         permit                          fully_qual(7)
```

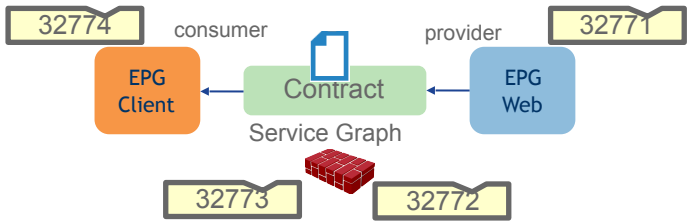

With Service Graph

Service Graph Operations
Applied Policy View

Without Service Graph

```
Leaf1# show zoning-rule scope 2621442
Rule ID         SrcEPG          DstEPG          FilterID        operSt          Scope           Action                          Priority
=======         ======          ======          ========        ======          =====           ======                          ========
<snip>
4211            32771           32774           39              enabled         2621442         permit                          fully_qual(7)
4212            32774           32771           38              enabled         2621442         permit                          fully_qual(7)
```

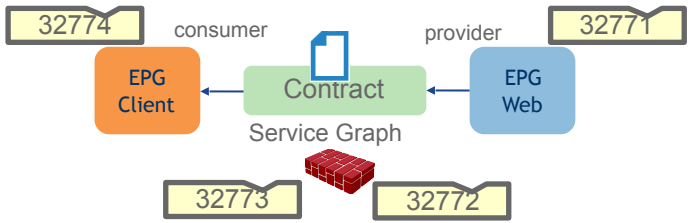

With Service Graph

```
Leaf1# show zoning-rule scope 2621442
Rule ID         SrcEPG          DstEPG          FilterID        operSt          Scope           Action                          Priority
=======         ======          ======          ========        ======          =====           ======                          ========
<snip>
4233            32772           32771           default         enabled         2621442         permit                          src_dst_any(9)
4137            32773           32774           39              enabled         2621442         permit                          fully_qual(7)
4168            32774           32773           38              enabled         2621442         permit                          fully_qual(7)
4169            32771           32772           default         enabled         2621442         permit                          src_dst_any(9)
```

*Managed Mode/Hybrid Mode Only

Dynamic Attach Endpoint
Overview

- APIC dynamically detect new endpoint, then EP is automatically added to
  - Member pool of VIP for ADC
  - EPG based network objects for FW
- Device package is required
- Function connector attachment notification

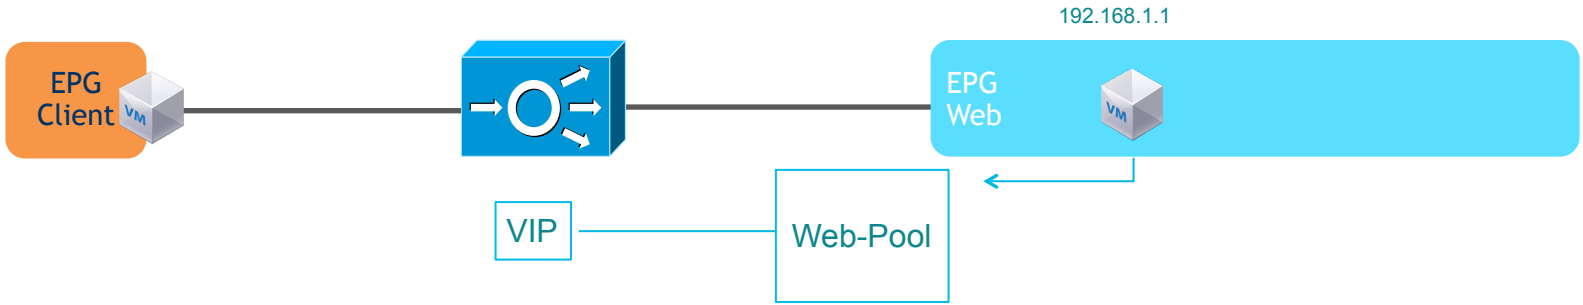

*Managed Mode/Hybrid Mode Only

Dynamic Attach Endpoint
Overview

- APIC dynamically detect new endpoint, then EP is automatically added to
  - Member pool of VIP for ADC
  - EPG based network objects for FW
- Device package is required
- Function connector attachment notification

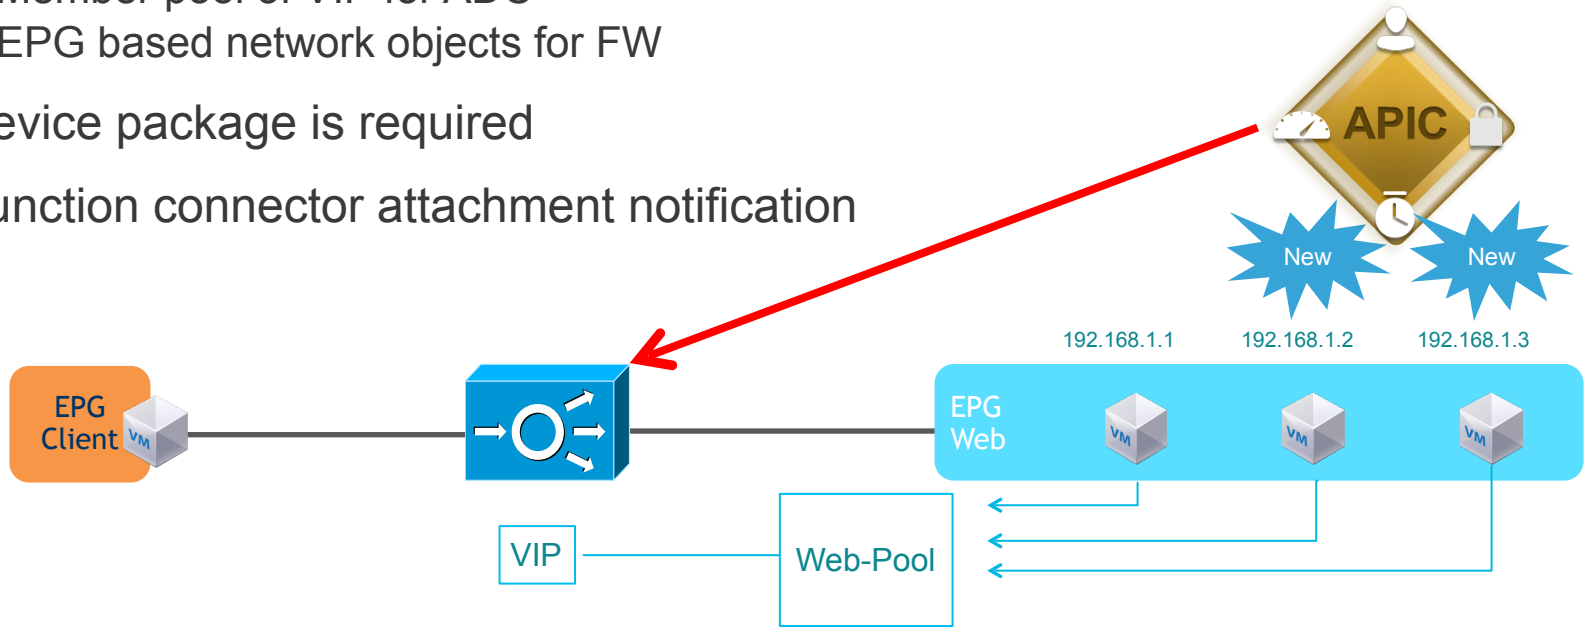

Policy Based Redirect (PBR)
Overview

- APIC 2.0 or later
- Use case 1: Inspect specific traffic by FW

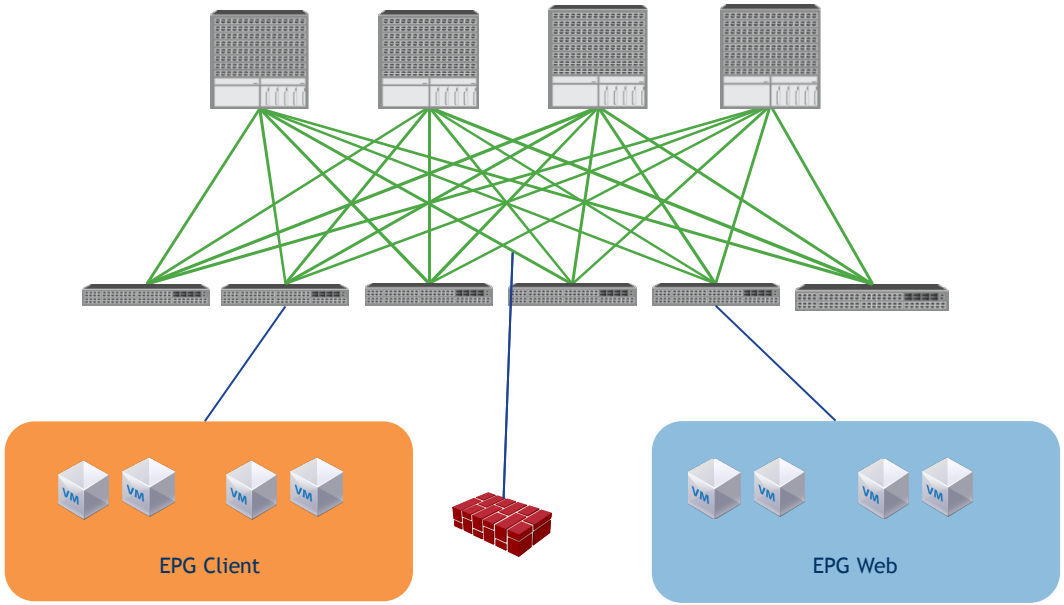

Policy Based Redirect (PBR)
Overview

- APIC 2.0 or later
- Use case 1: Inspect specific traffic by FW

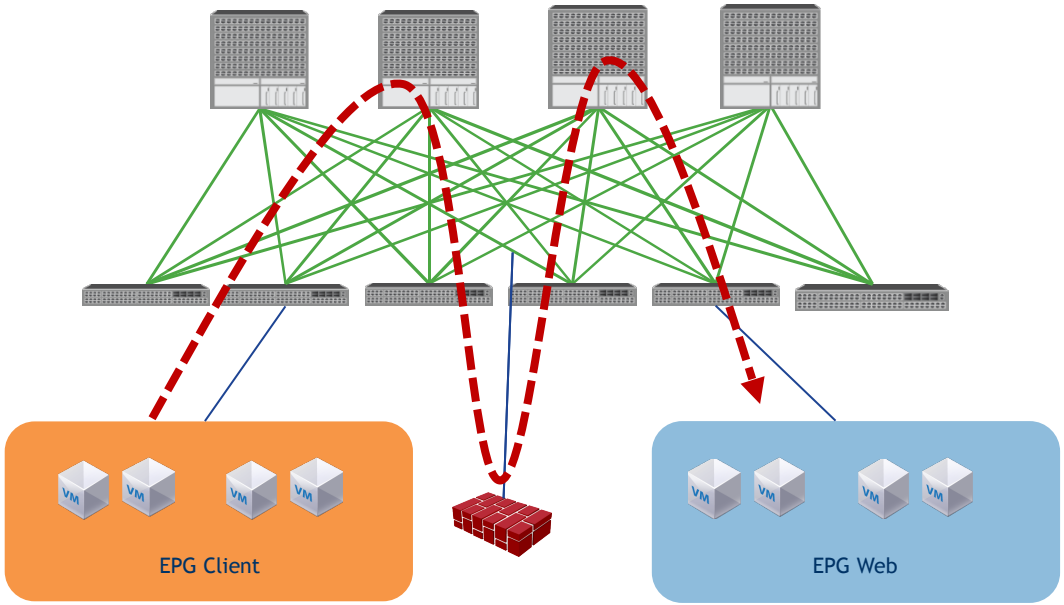

Policy Based Redirect (PBR)
Overview

- APIC 2.0 or later
- Use case 1: Inspect specific traffic by FW

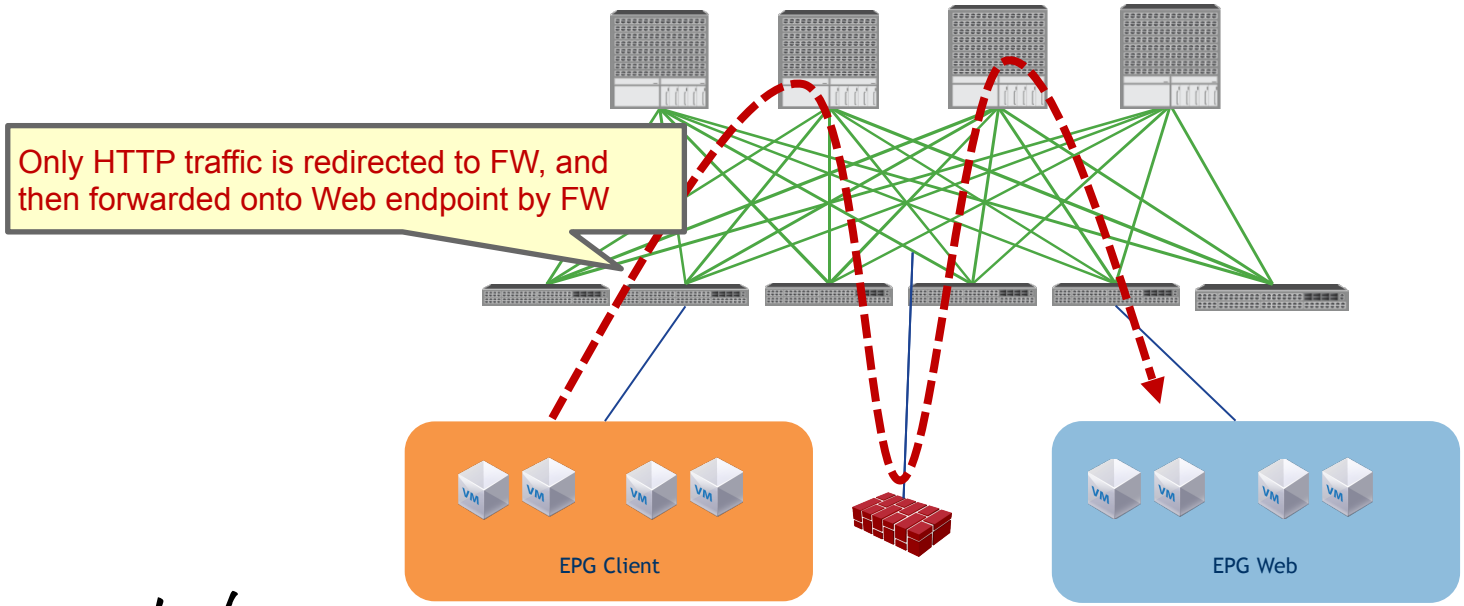

Policy Based Redirect (PBR)
Overview

- APIC 2.0 or later
- Use case 1: Inspect specific traffic by FW

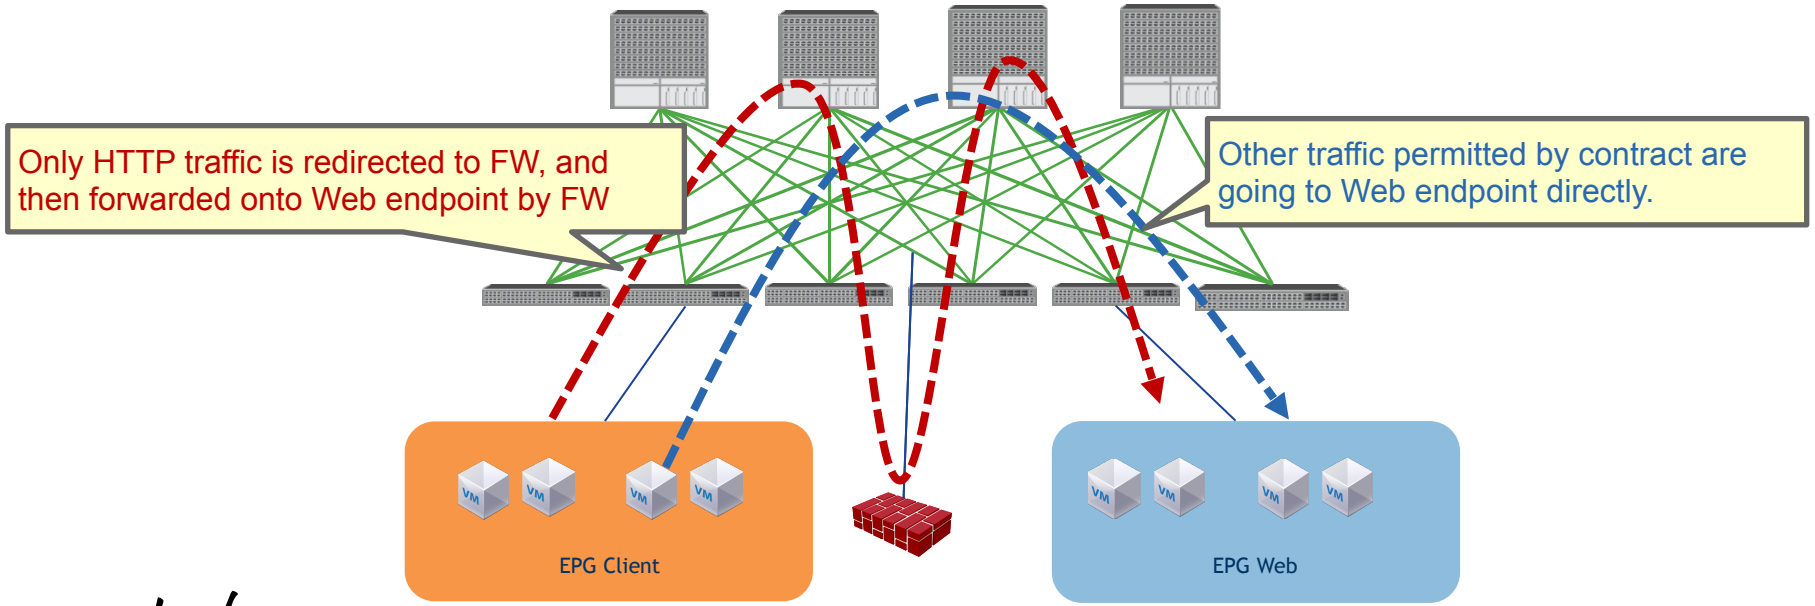

Policy Based Redirect (PBR)
Overview

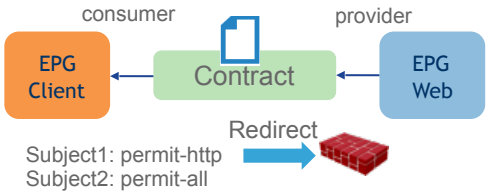

- APIC 2.0 or later
- Use case 1: Inspect specific traffic by FW

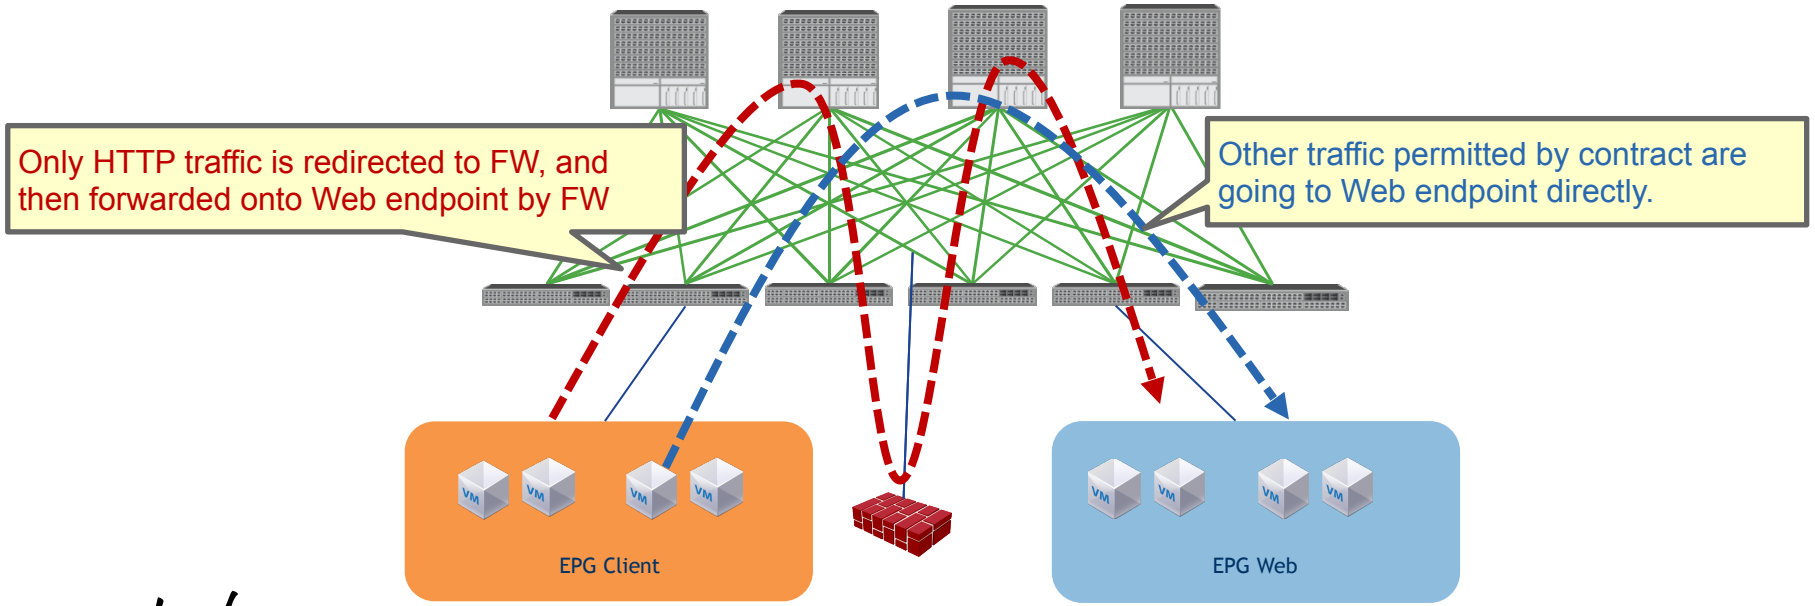

Policy Based Redirect (PBR)
Overview

- Use case 2: Use different Firewall based on source

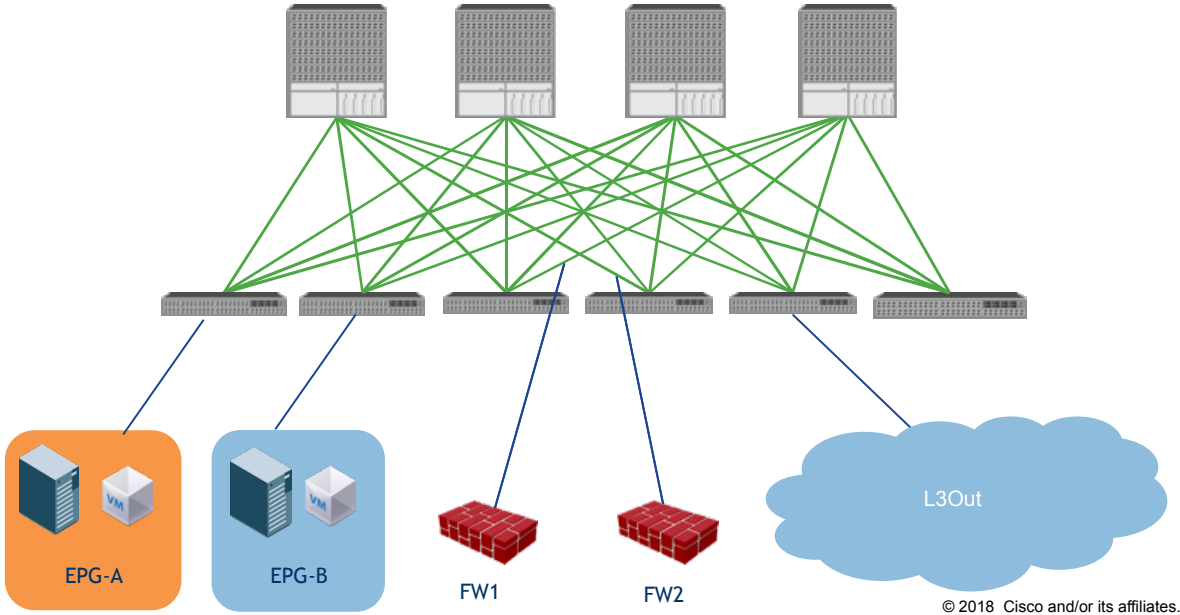

Policy Based Redirect (PBR)
Overview

- Use case 2: Use different Firewall based on source

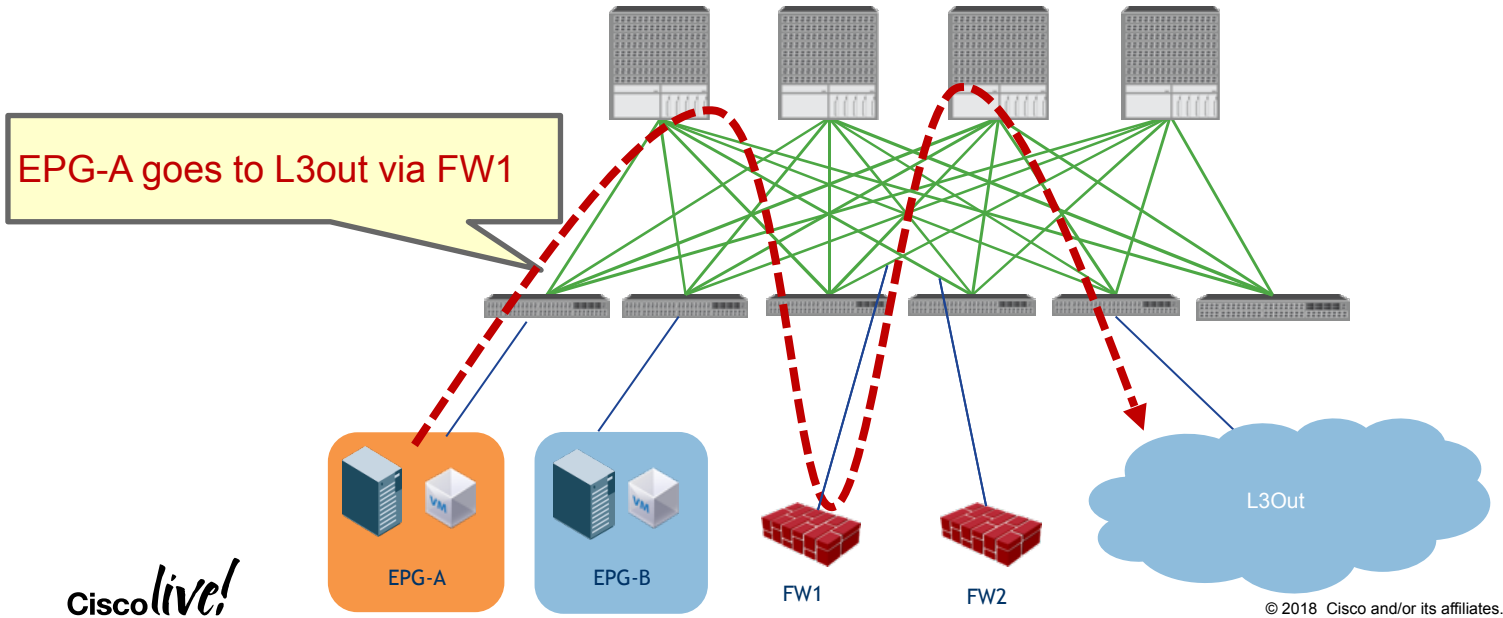

Policy Based Redirect (PBR)
Overview

- Use case 2: Use different Firewall based on source

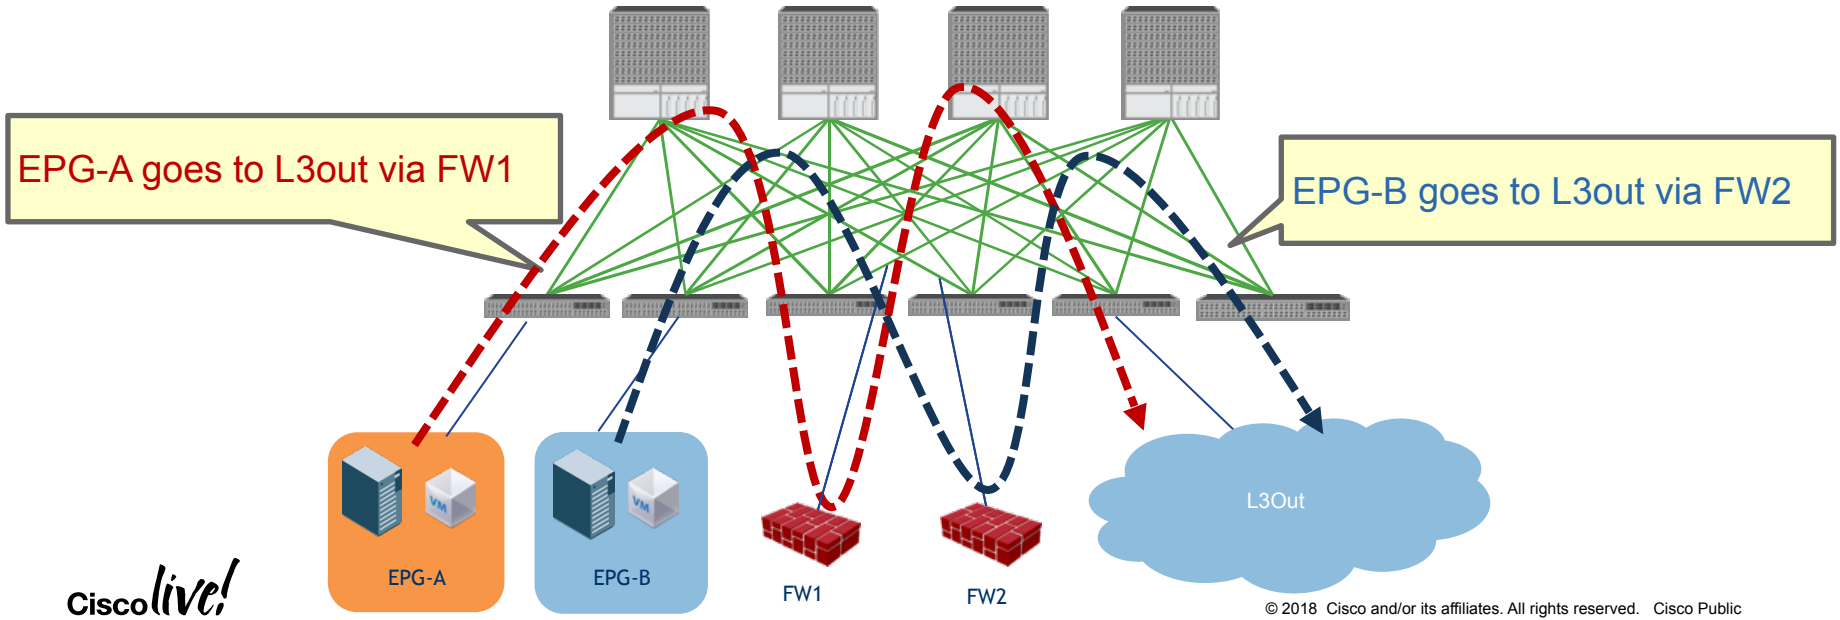

Policy Based Redirect (PBR)
Overview

- Use case 2: Use different Firewall based on source

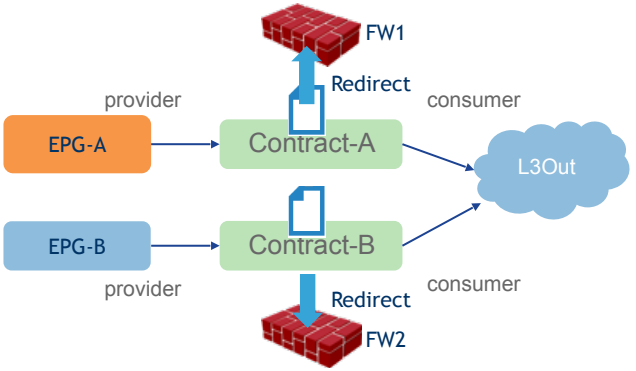

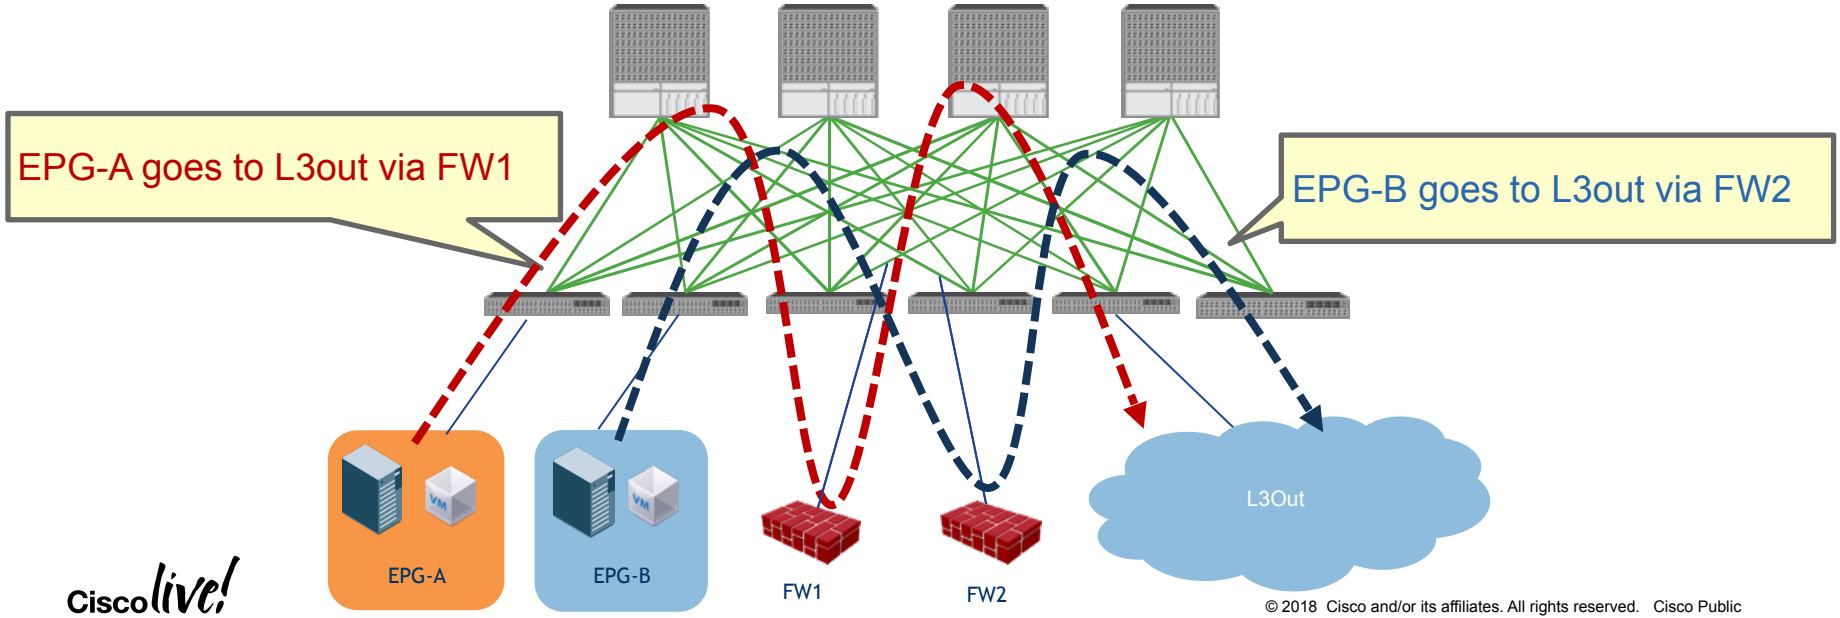

Copy Service
Overview

- APIC 2.0
- Service Graph is mandatory and EX/FX hardware is required
- Copy specific traffic

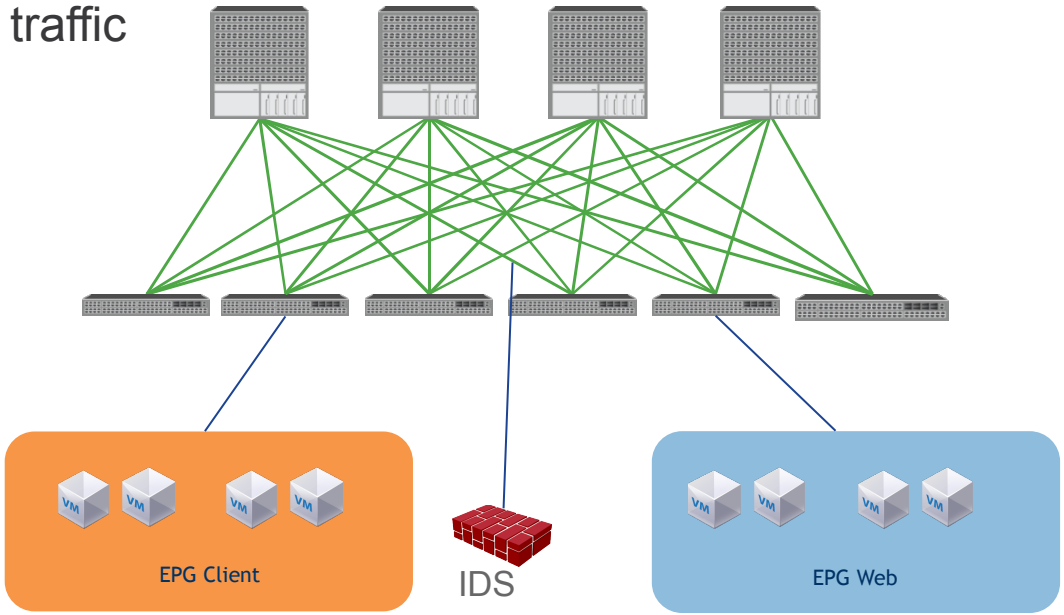

Copy Service
Overview

- APIC 2.0
- Service Graph is mandatory and EX/FX hardware is required
- Copy specific traffic

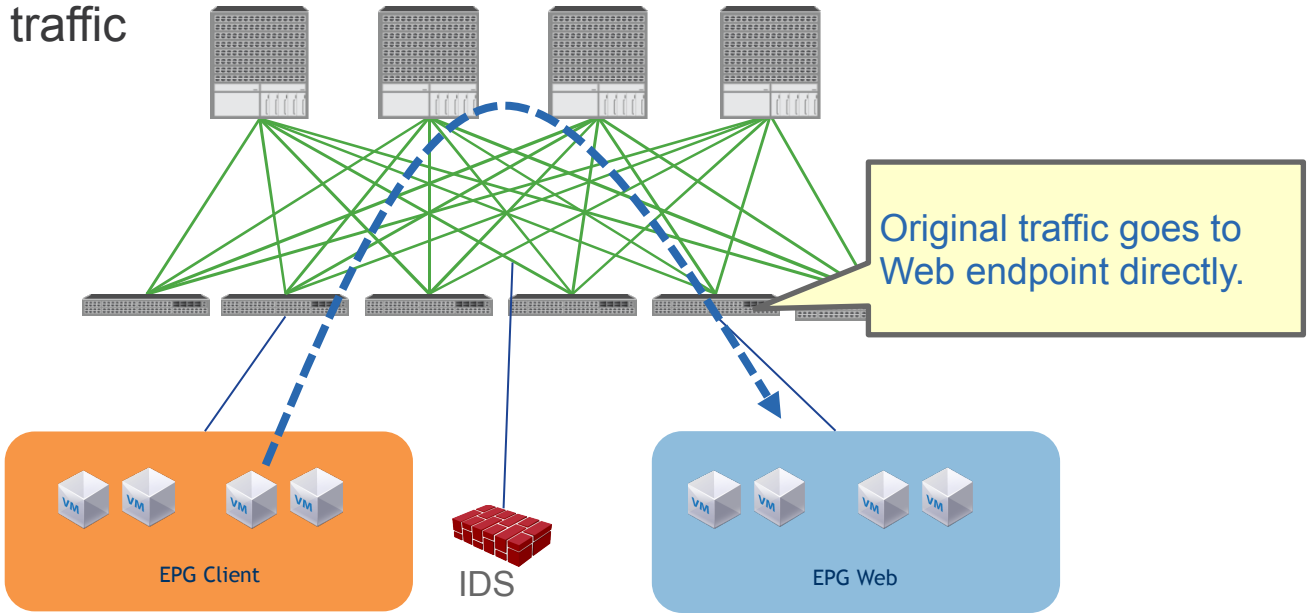

Copy Service
Overview

- APIC 2.0
- Service Graph is mandatory and EX/FX hardware is required
- Copy specific traffic

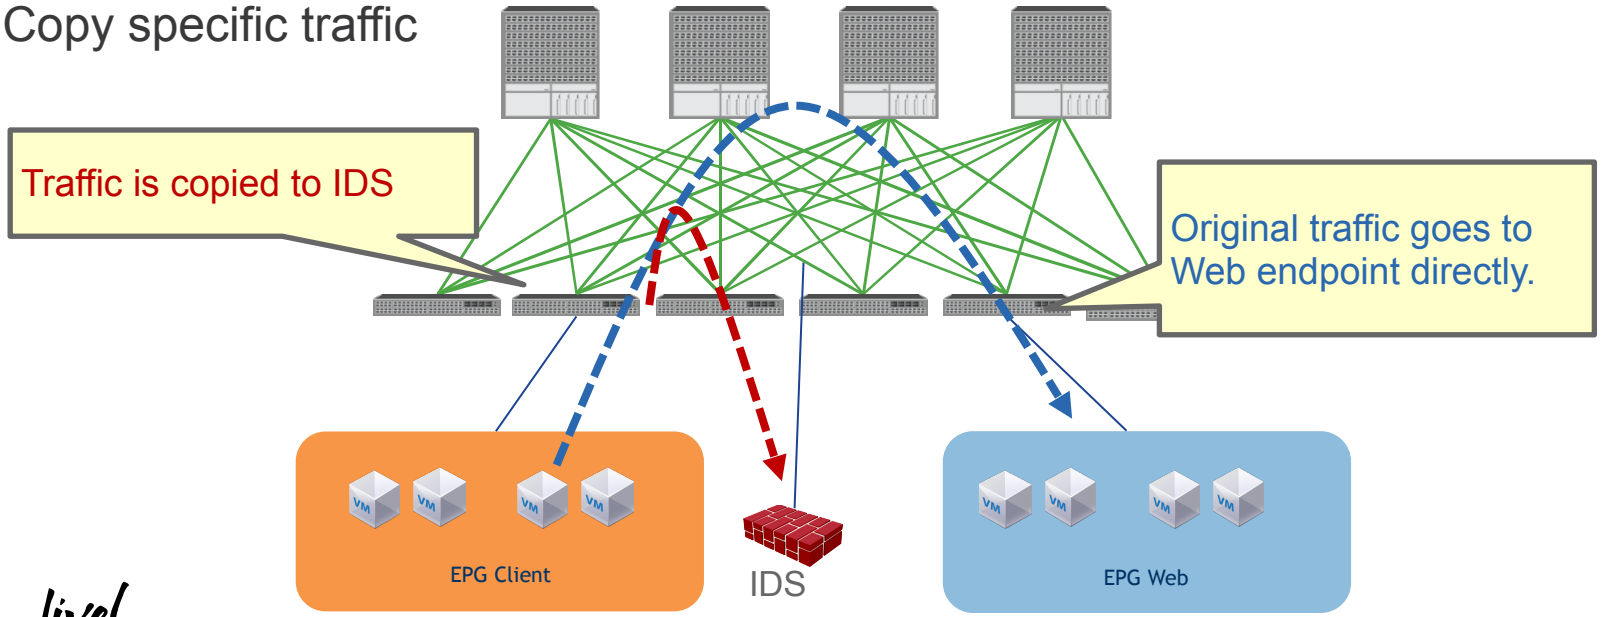

Copy Service
Overview

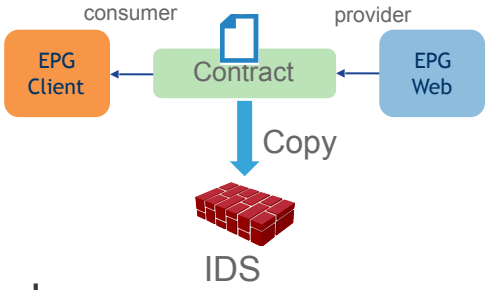

- APIC 2.0
- Service Graph is mandatory and EX/FX hardware is required
- Copy specific traffic

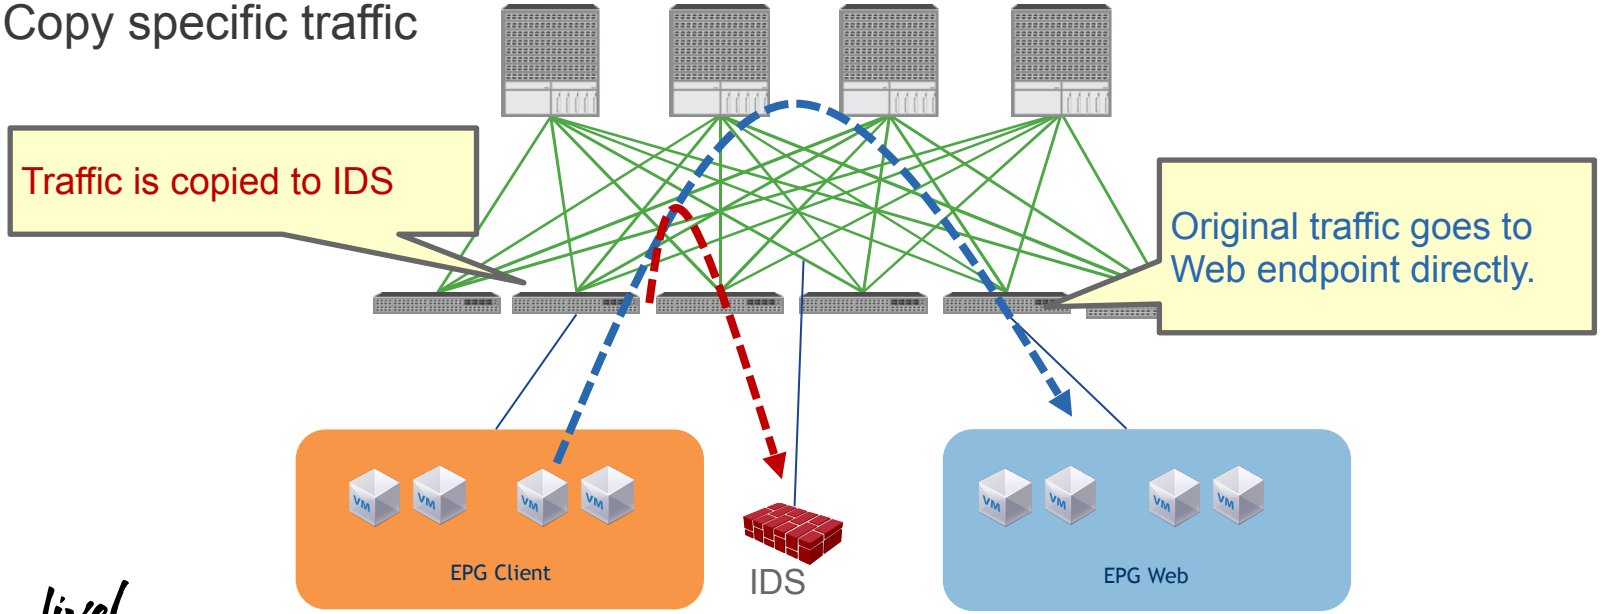

Design Use Cases and Considerations

- Firewall
- ADC
- IPS/IDS

- PBR (Policy Based Redirect)
- Multi-pod + Service

Firewall

Firewall Mode Options

	L2 (Transparent)	L3 (Routed)
	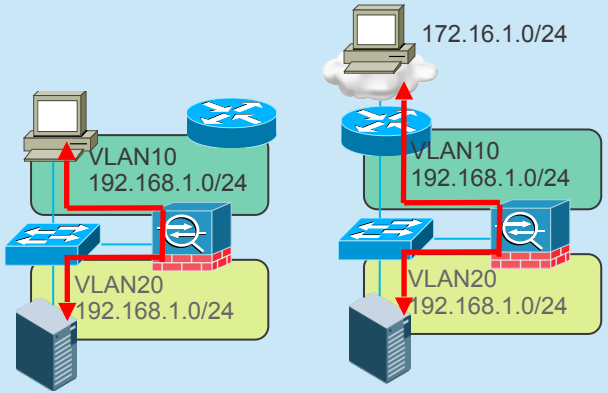	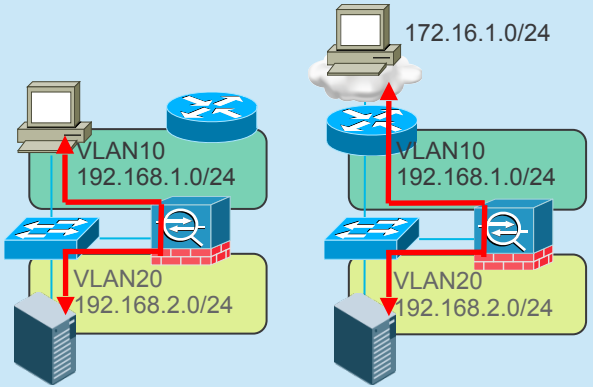
Default GW of servers	Router	Firewall alias IP of server-side (in case of ASA, active IP)
When to use	Want to use same subnet for client-side and server-side	Common case when we can have separate VLAN/subnet for client-side and server-side.

ACI L2 FW Design

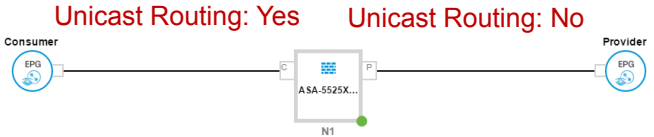

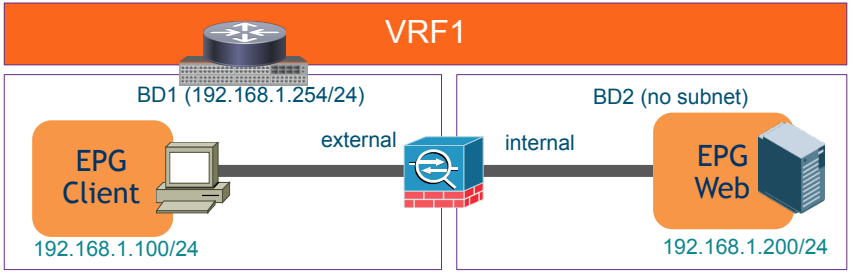

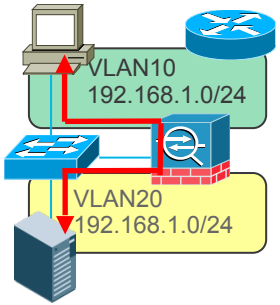

*Go-through mode will overwrite BD settings to Flooding/ARP enabled automatically.

ACI L2 FW Design

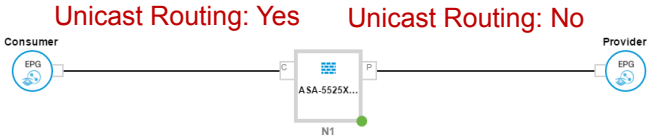

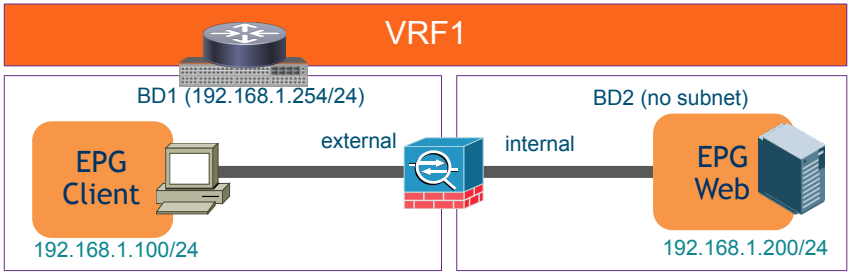

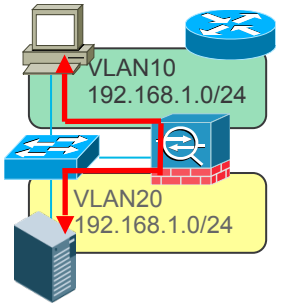

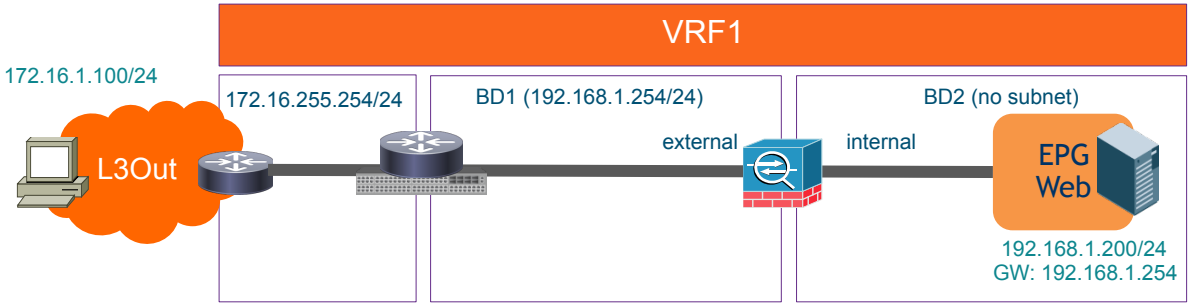

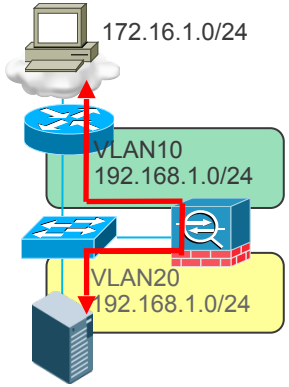

*Go-through mode will overwrite BD settings to Flooding/ARP enabled automatically.

ACI L3 FW Design

FW as Gateway for All (ACI as L2)

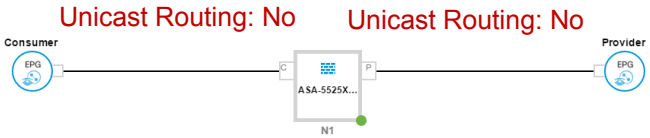

- If you want to insert services in the path between client and servers you normally need to make firewalls be the default gateway.
  - This means that you need to design your bridge domains in order to fit the firewall and load balancing placement requirements
  - All traffic goes to Firewall because Firewall is the default gateway.

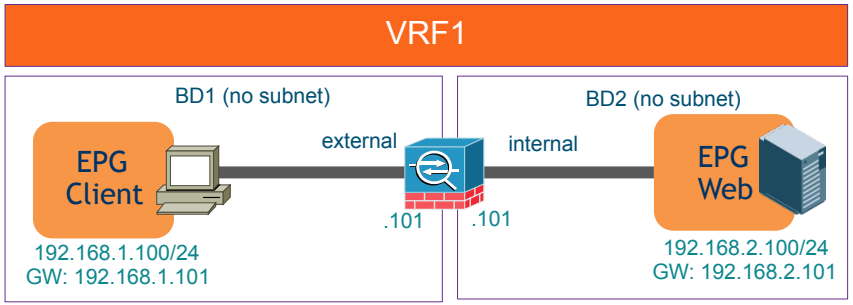

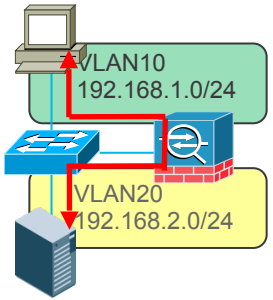

ACI L3 FW Design

ACI as Gateway (ACI as L3) with L3out

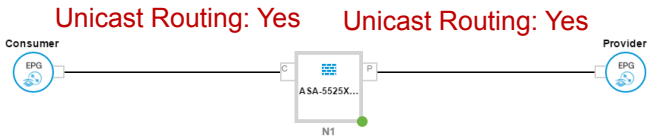

- If you want routers/switches or ACI to be the default gateway you need to use designs such as "VRF-sandwich":
  - This means that you need to redesign your network with more VRFs than you need just to make the traffic flow via the L4-L7 device
  - All traffic goes to Firewall based on routing design.

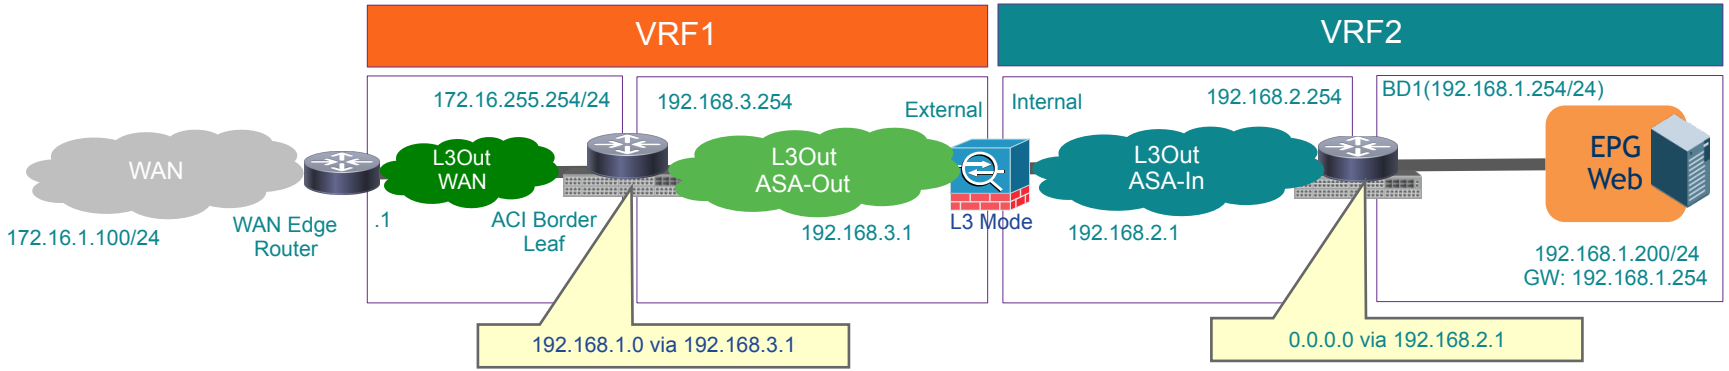

ACI L3 FW Design

ACI as Gateway (ACI as L3) with PBR

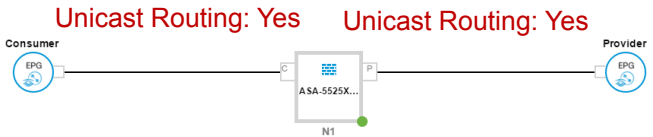

- All BDs are in same VRF
- Support selective redirect

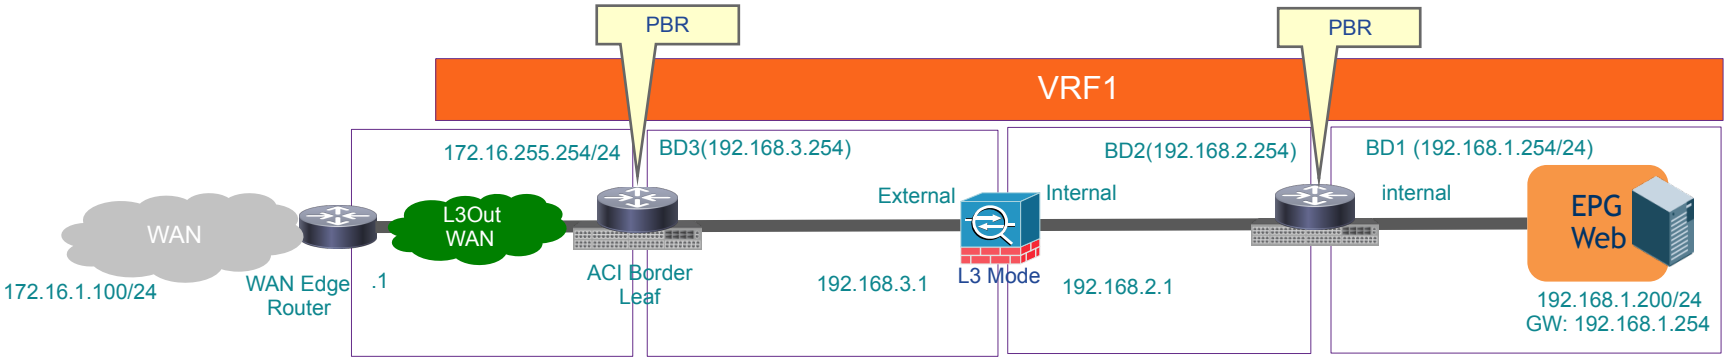

PBR

Design Considerations

For Your Reference

- Supported for both managed node and unmanaged node
- Works only with GoTo (L3) devices. Transparent device (e.g IDS, IPS) is not supported today.
- Both Active/Standby service nodes should have same VMAC
- For non-EX/FX leafs, service node can not be under either consumer or provider EPG Leaf
- Prior to ACI version 3.1:
  - Service node has to be in BD. The BD must be different from provider/consumer EPG BD.
  - “Endpoint Dataplane Learning” should be disabled on BDs for service node interfaces
- Prior to ACI version 3.2:
  - Can be enabled for only one node in multi-node service graphs

ADC

ADC Design Options

	L2 (Bridged)	L3 (Routed) Two-Arm	One-Arm	DSR
	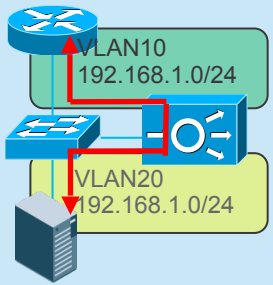	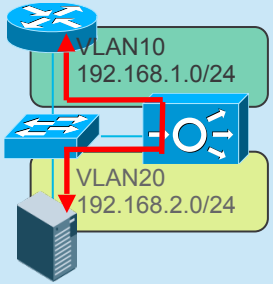	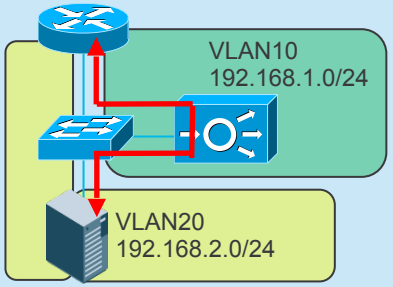	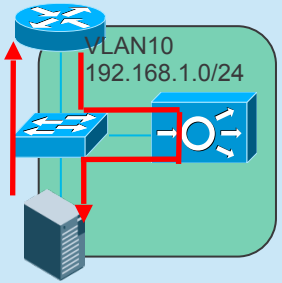
Default GW of server	Router	ADC alias IP (server side)	Router	Router
NAT	ADC does DNAT (VIP <-> real server IP)	ADC does DNAT (VIP <-> real server IP)	ADC does DNAT and SNAT (VIP <-> real server IP, client IP <-> ADC)	ADC doesn't do NAT
When we use	Want to use same subnet for client-side and server-side	Common case when we can have separate VLAN/subnet for client-side and server-side.	Don't want to insert new VLAN/ subnet between router and server.	Return traffic is huge and want to reduce load of ADC (ex streaming server)
comment	Current ADC device package doesn't support it.			APIC version 1.2

ACI ADC Design
Two-arm Mode

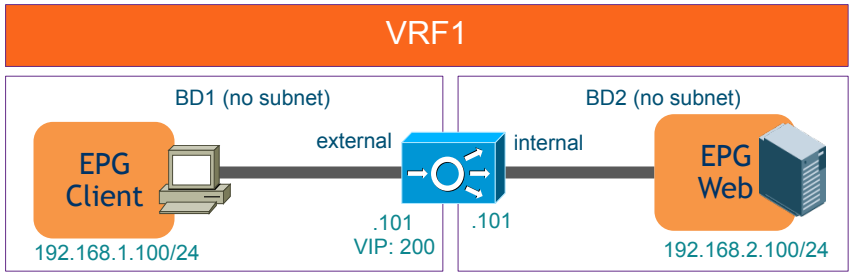

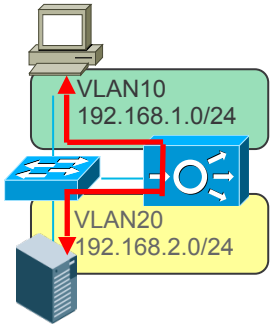

ACI ADC Design
Two-arm Mode

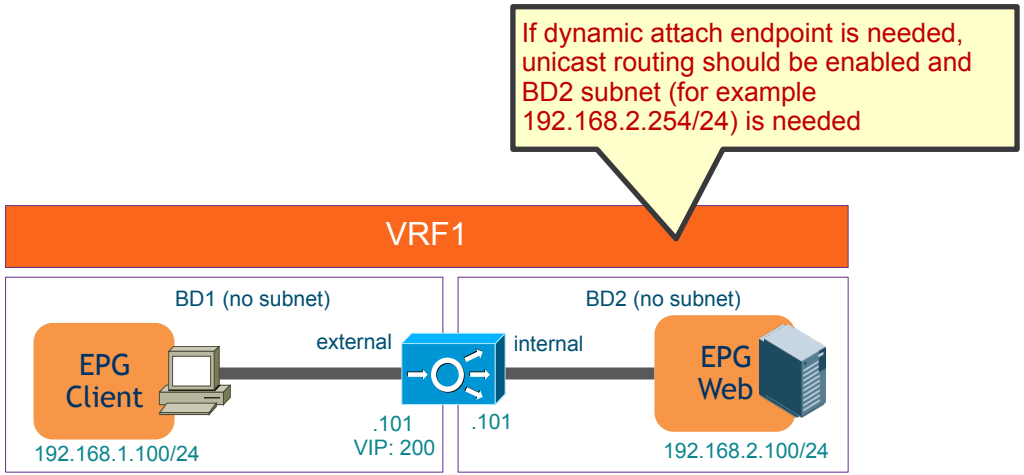

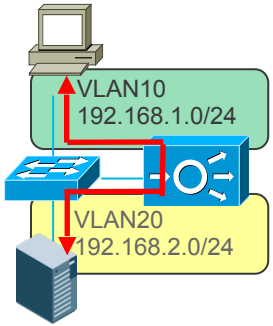

ACI ADC Design
Two-arm Mode

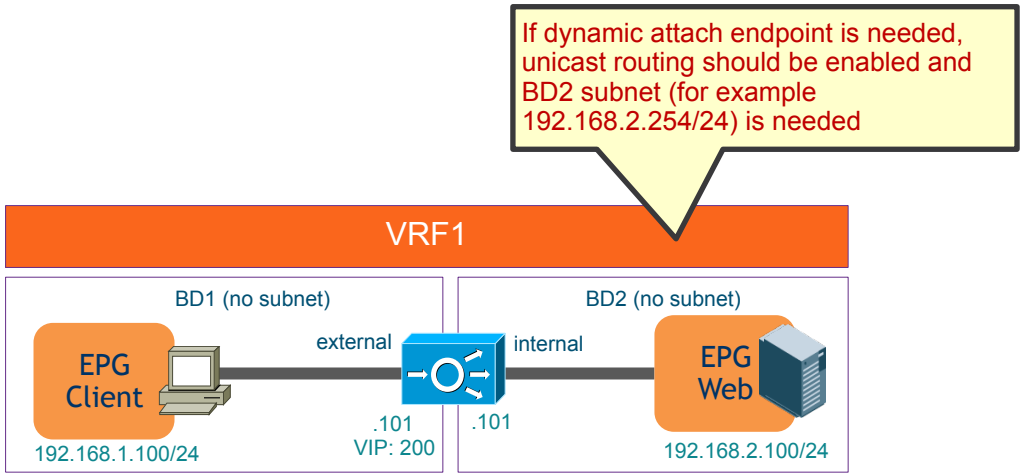

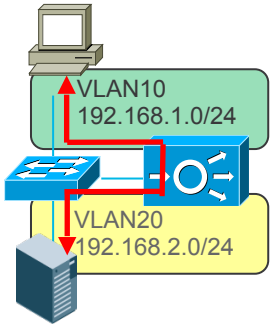

No SNAT on ADC

ACI ADC Design
Two-arm Mode

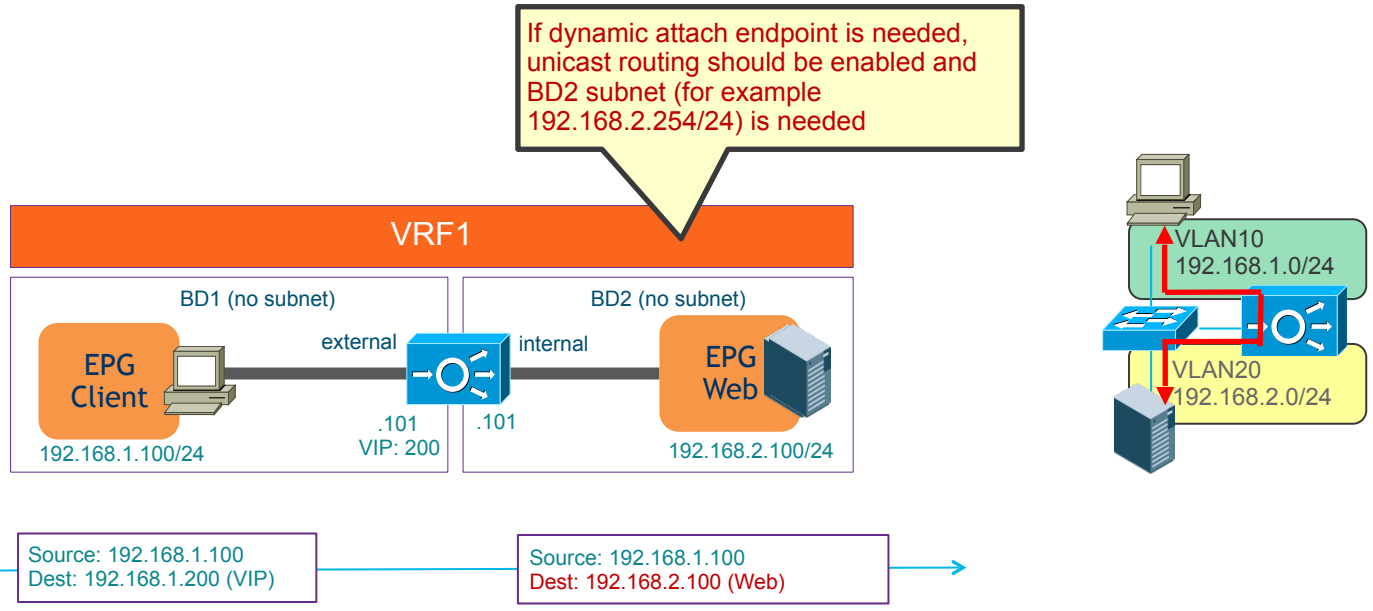

ACI ADC Design
Two-arm Mode

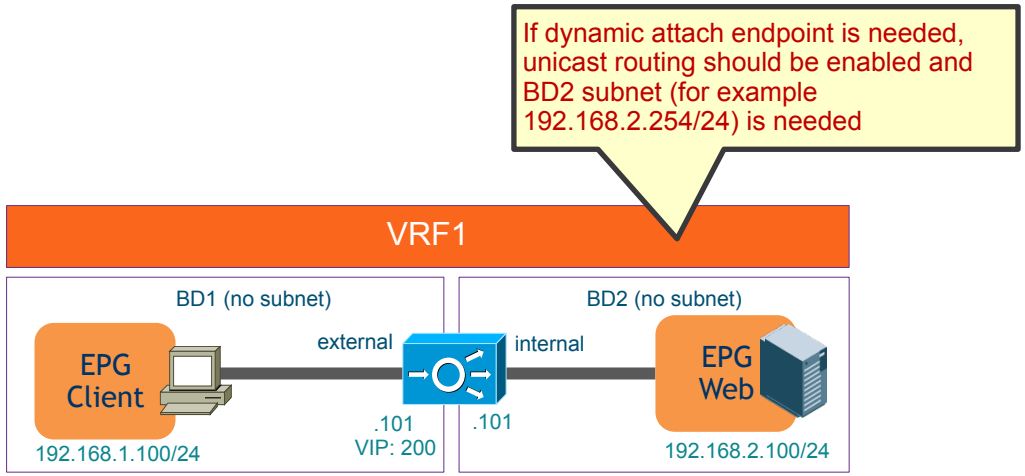

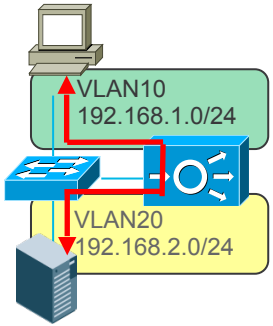

No SNAT on ADC

SNAT on ADC

ACI ADC Design
Two-arm Mode

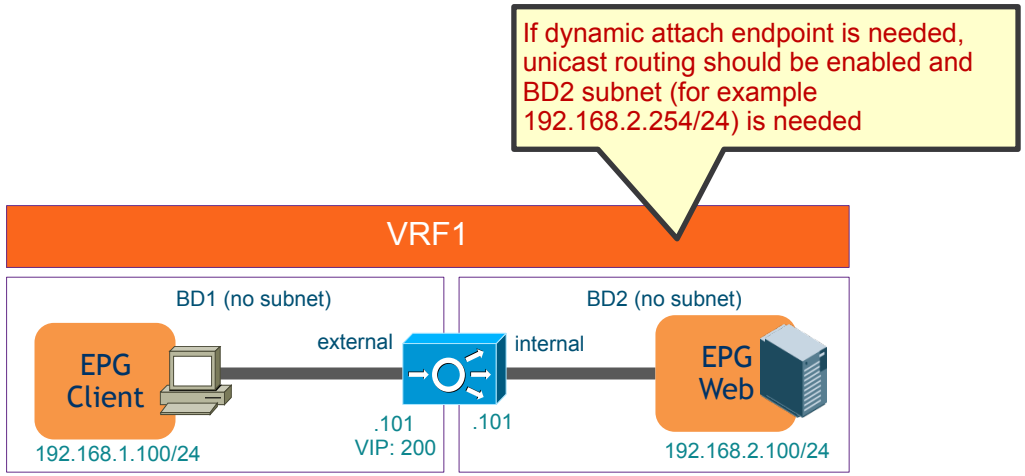

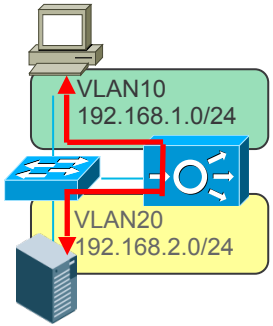

No SNAT on ADC

Source: 192.168.1.100
Dest: 192.168.1.200 (VIP)

Source: 192.168.1.100
Dest: 192.168.2.100 (Web)

SNAT on ADC

Source: 192.168.1.100
Dest: 192.168.1.200 (VIP)

Source: 192.168.2.101 (ADC)
Dest: 192.168.2.100 (Web)

ACI ADC Design
One-arm with SNAT

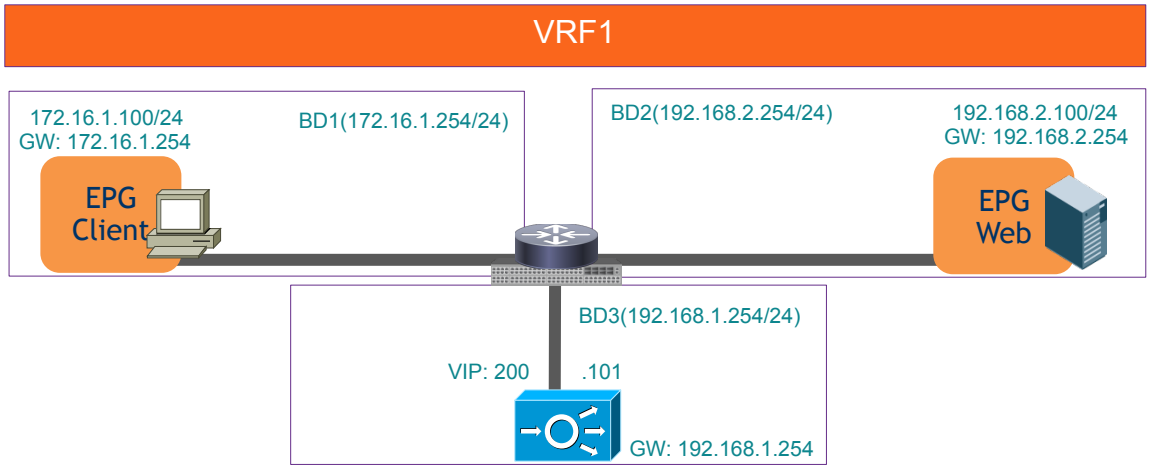

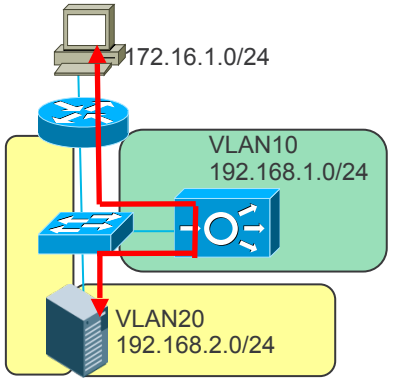

ACI ADC Design
One-arm with SNAT

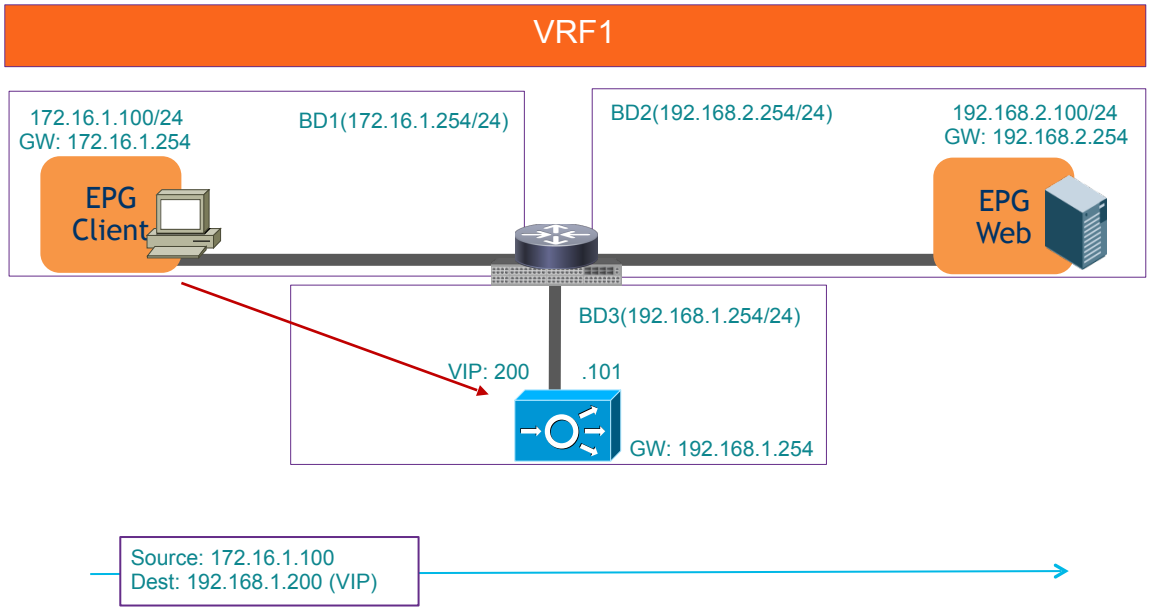

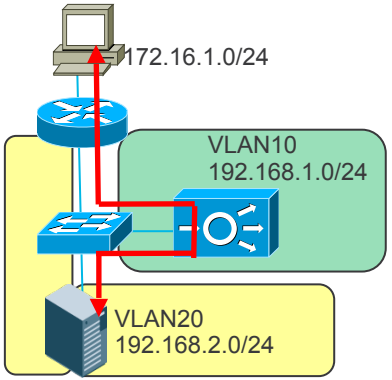

ACI ADC Design
One-arm with SNAT

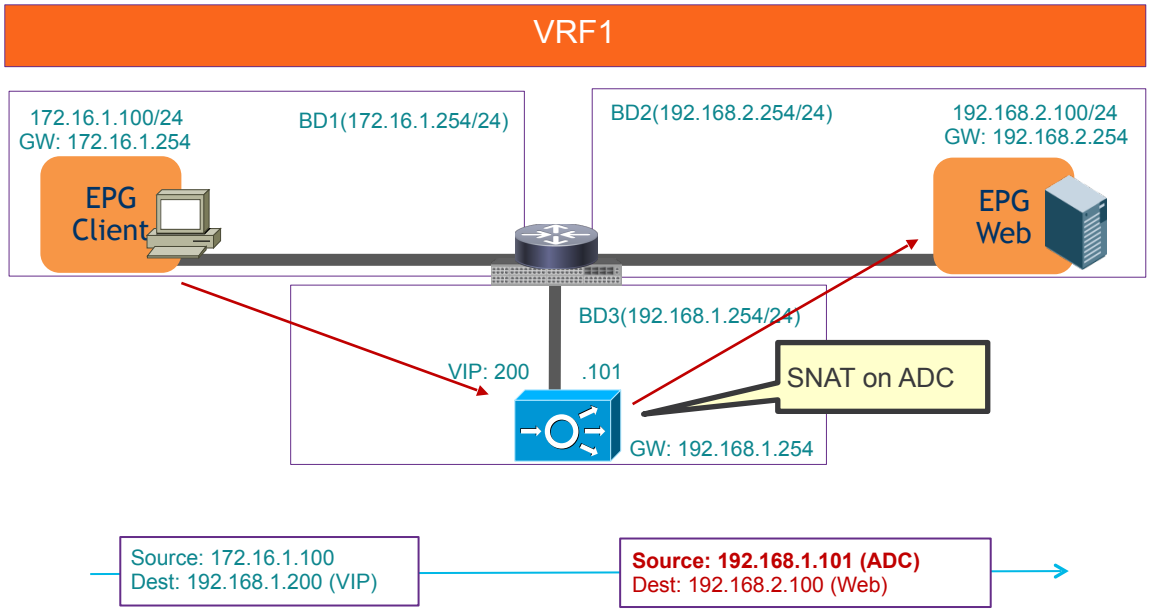

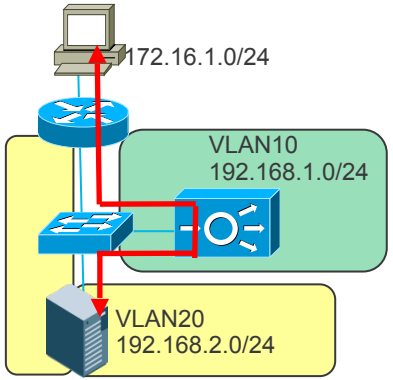

ACI ADC Design
One-arm with PBR

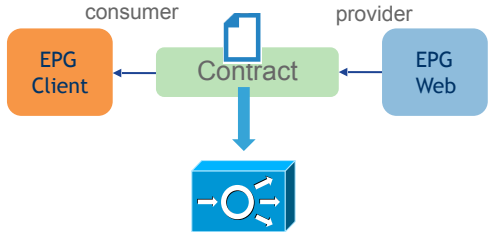

SNAT is not required and server sees real client IP

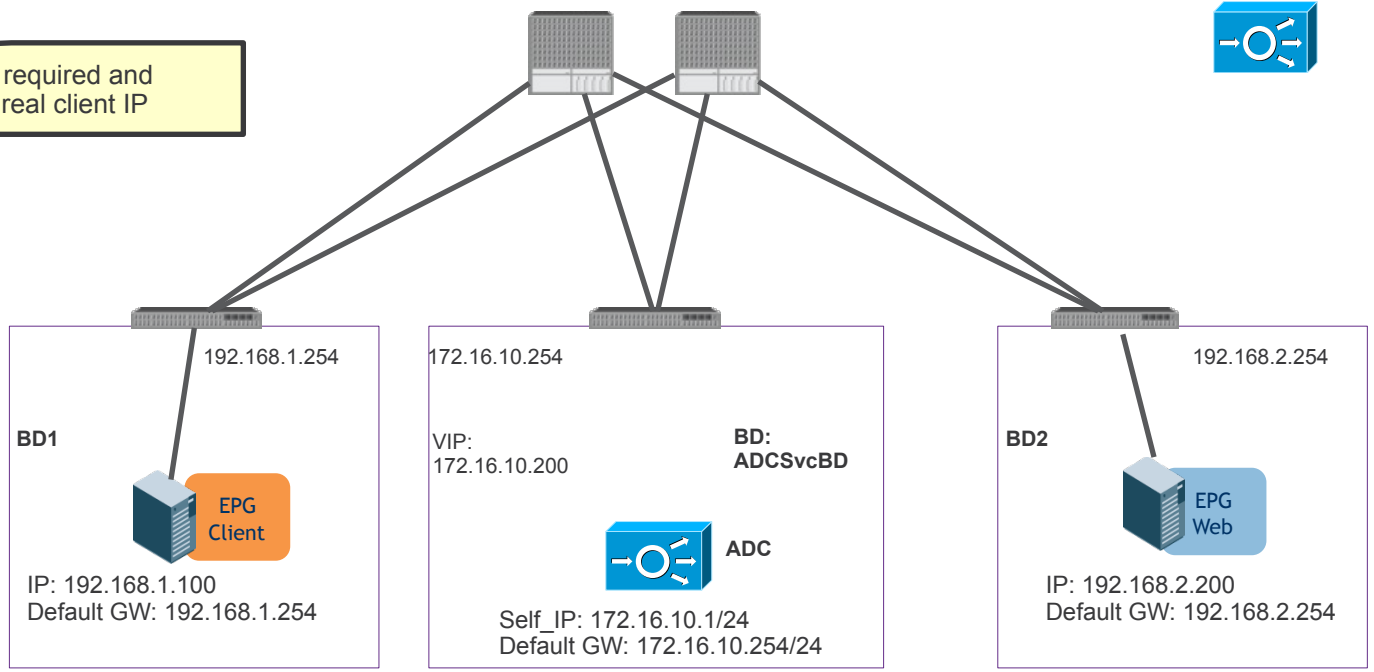

All BDs are in same VRF

ACI ADC Design
One-arm with PBR

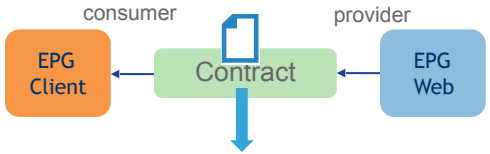

SNAT is not required and server sees real client IP

Incoming Traffic
Source: 192.168.1.100
Dest: 172.16.10.200 (VIP)

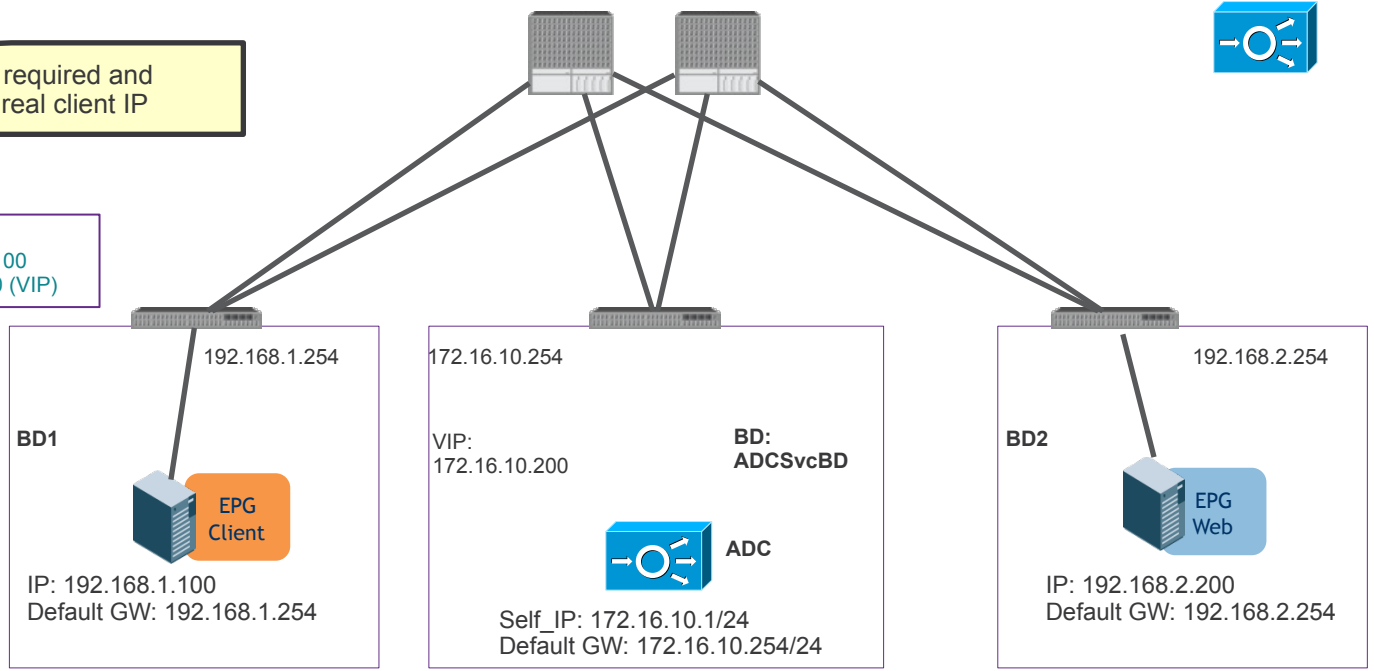

All BDs are in same VRF

ACI ADC Design
One-arm with PBR

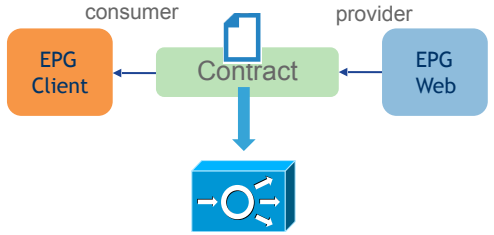

SNAT is not required and server sees real client IP

Incoming Traffic
Source: 192.168.1.100
Dest: 172.16.10.200 (VIP)

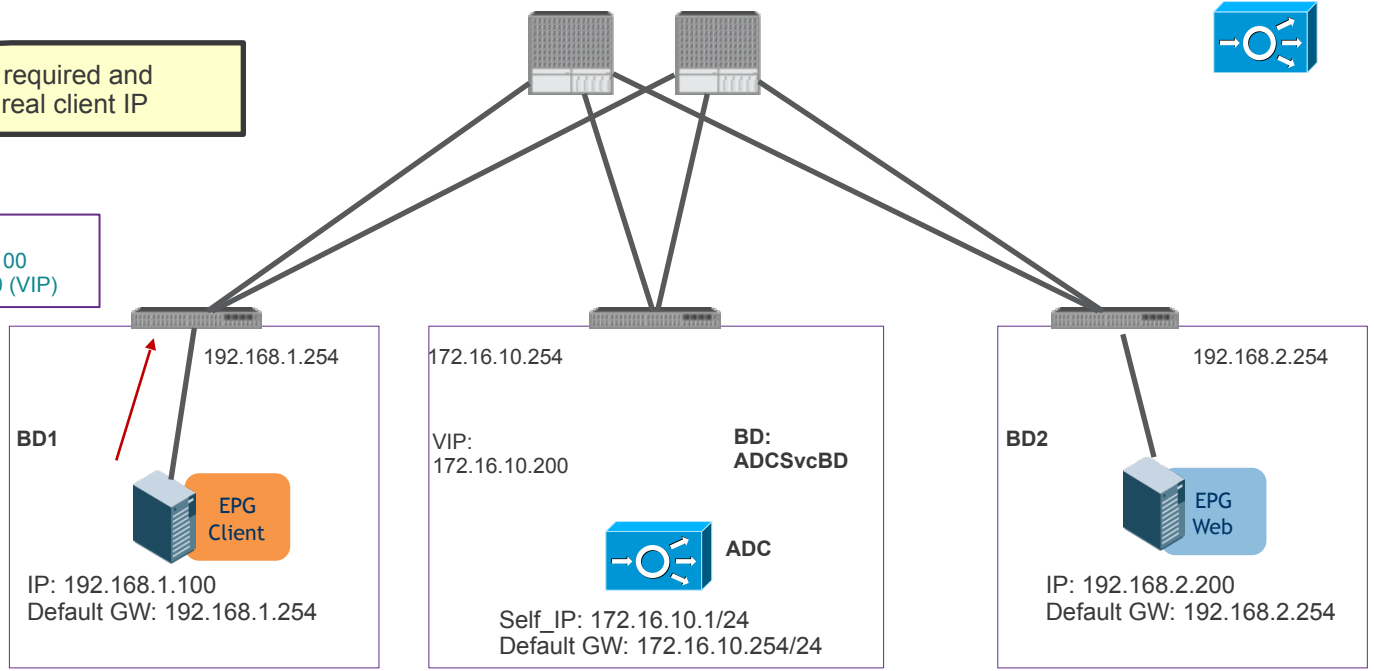

All BDs are in same VRF

ACI ADC Design
One-arm with PBR

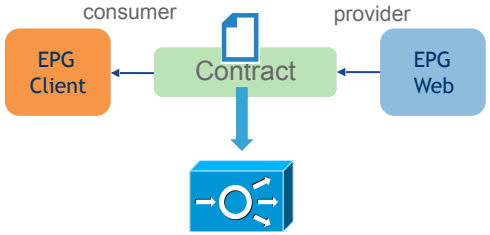

SNAT is not required and server sees real client IP

Incoming Traffic
Source: 192.168.1.100
Dest: 172.16.10.200 (VIP)

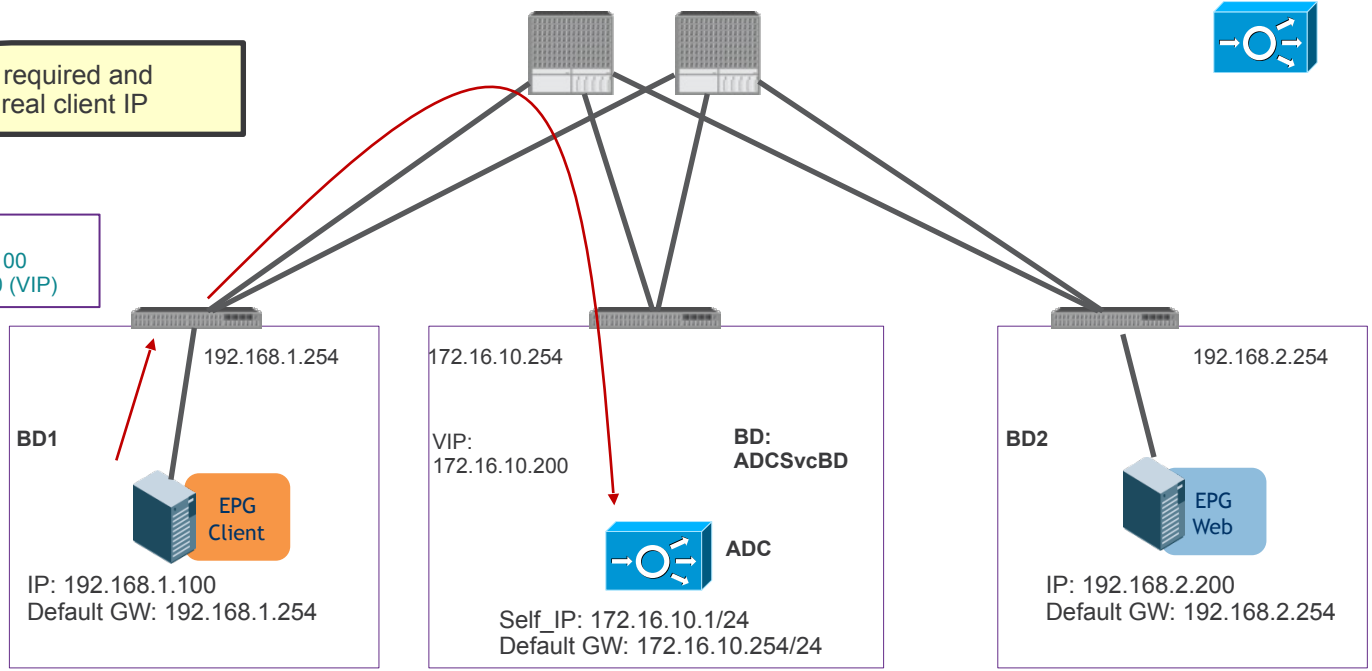

All BDs are in same VRF

ACI ADC Design
One-arm with PBR

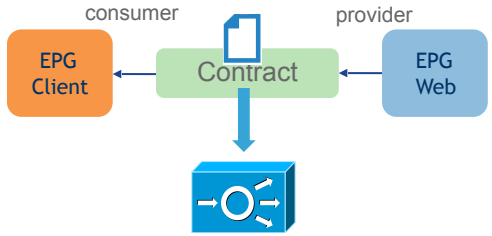

SNAT is not required and server sees real client IP

Incoming Traffic
Source: 192.168.1.100
Dest: 172.16.10.200 (VIP)

Return Traffic
Source: 192.168.2.200
Dest: 192.168.1.100

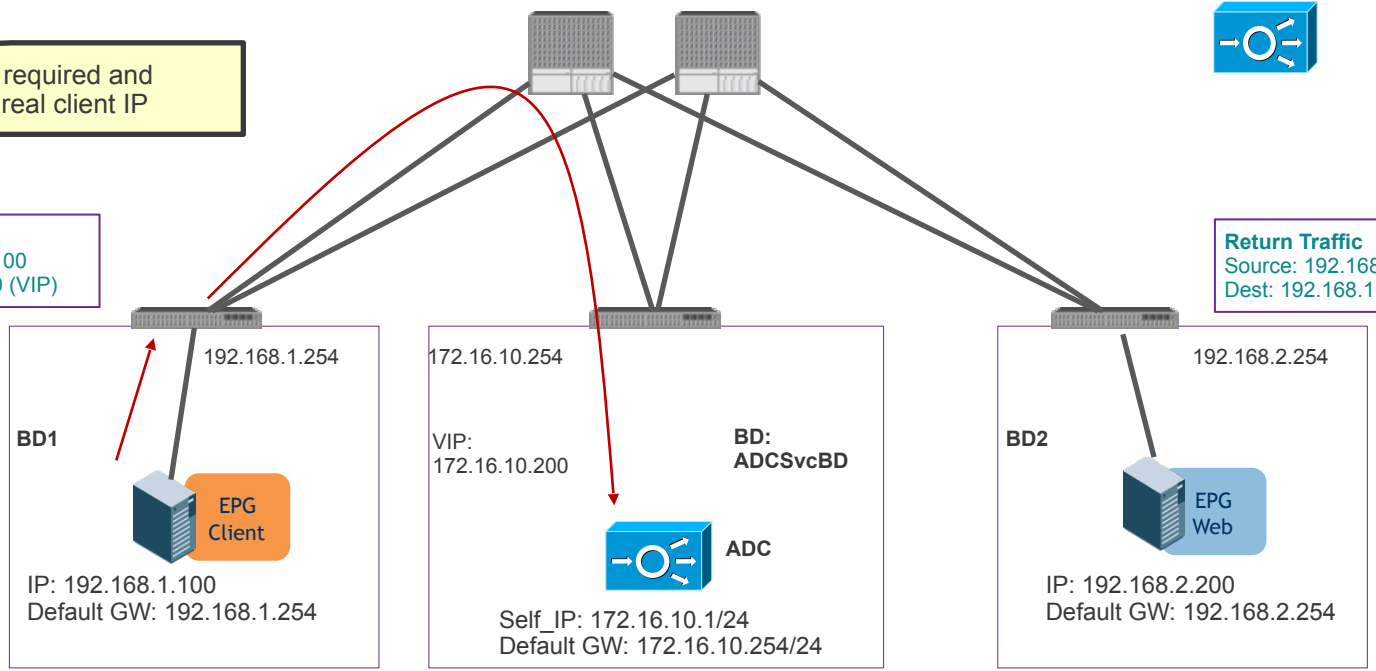

VRF1

All BDs are in same VRF

ACI ADC Design
One-arm with PBR

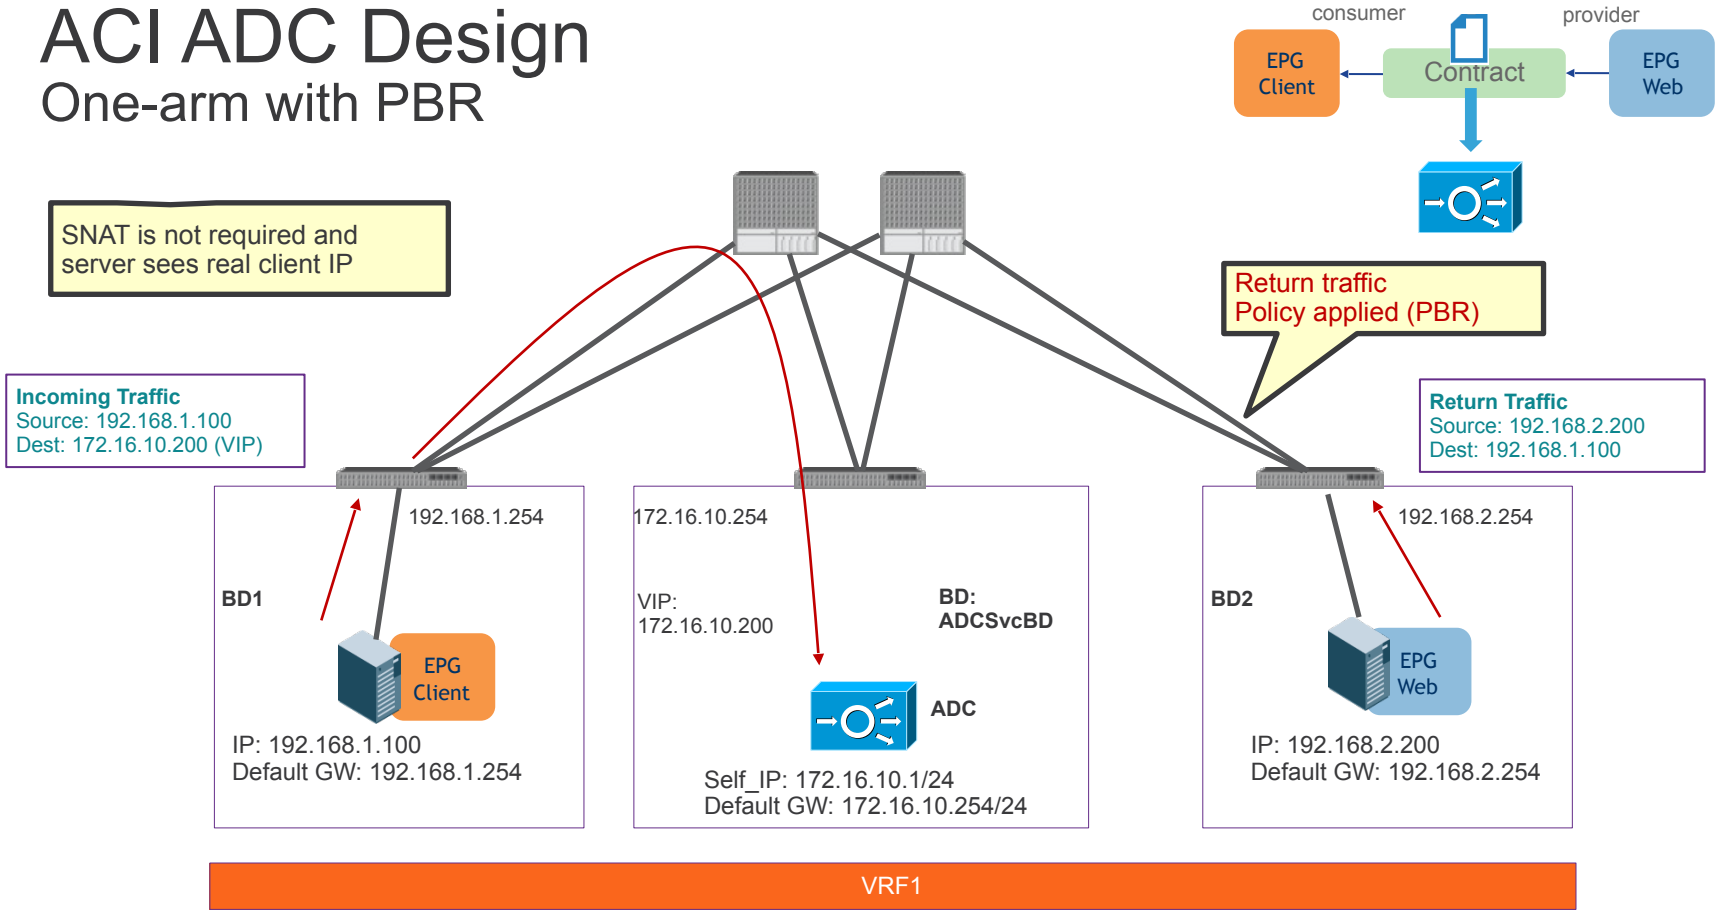

All BDs are in same VRF

ACI ADC Design
One-arm with PBR

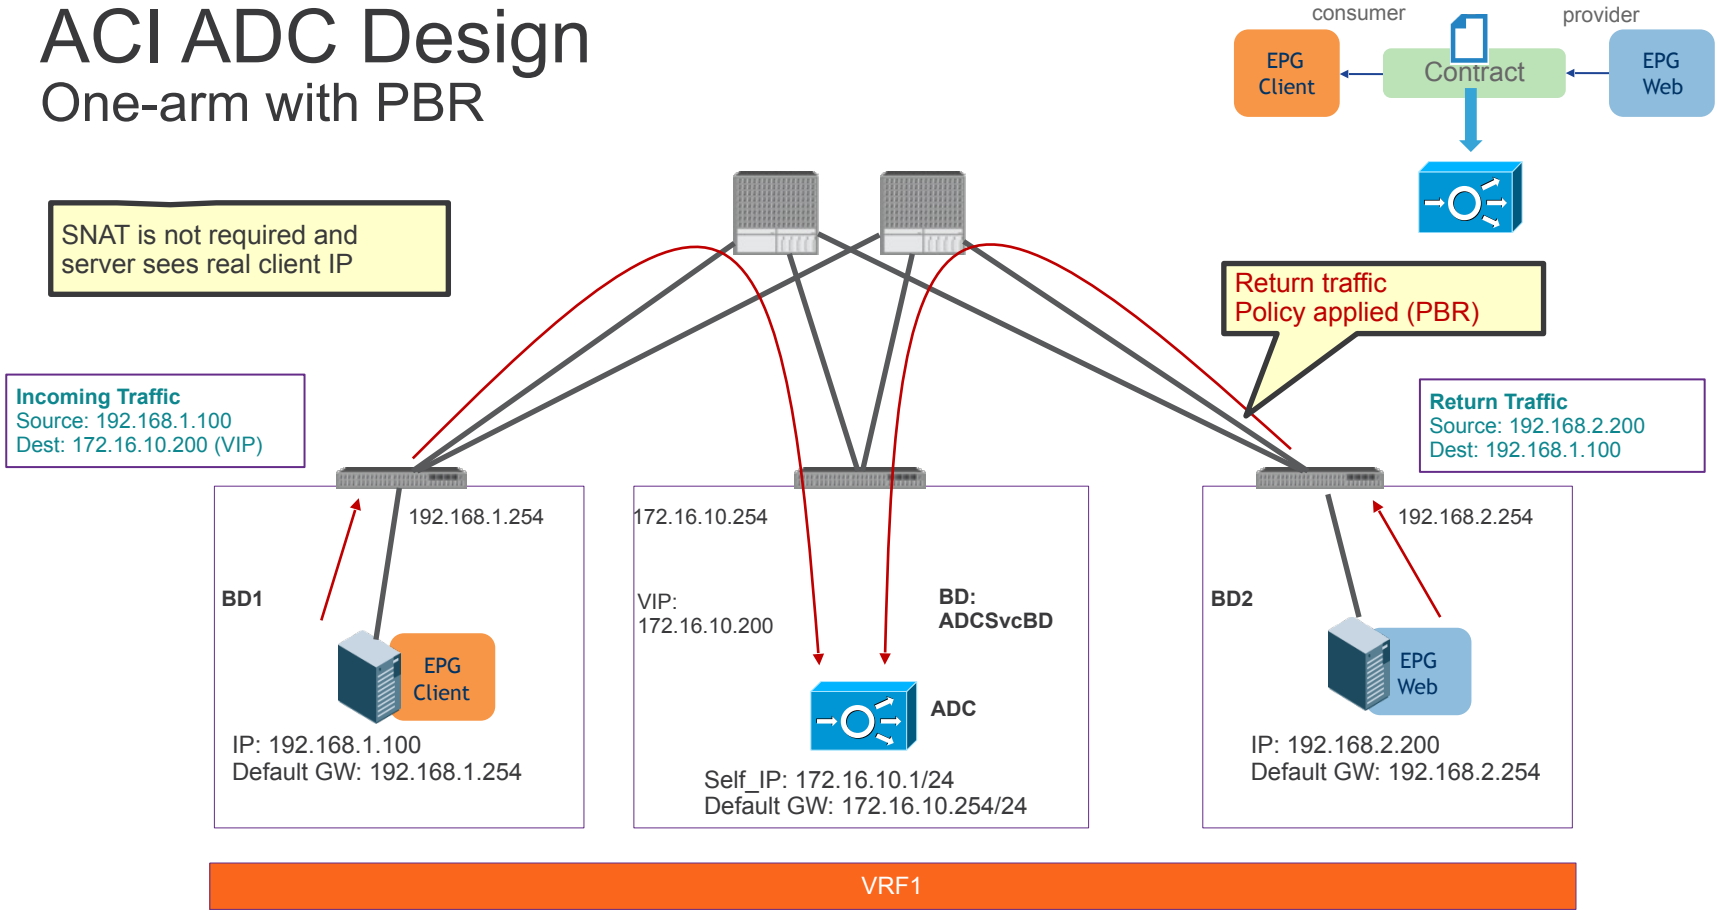

All BDs are in same VRF

DSR (Direct Server Return)

- It is commonly deployed when Return traffic is huge and want to reduce load of ADC
- Servers must suppress ARP response to VIP ARP requests
- ACI Dataplane end-point learning

  causes VIP move within fabric

  and should be disabled

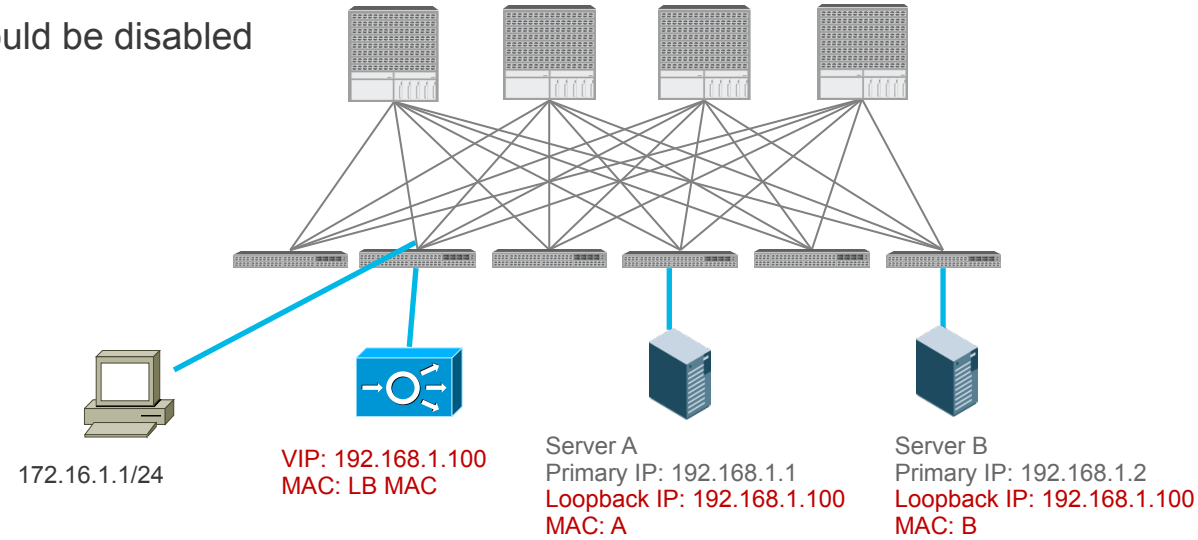

DSR (Direct Server Return)

- It is commonly deployed when Return traffic is huge and want to reduce load of ADC
- Servers must suppress ARP response to VIP ARP requests
- ACI Dataplane end-point learning

  causes VIP move within fabric

  and should be disabled

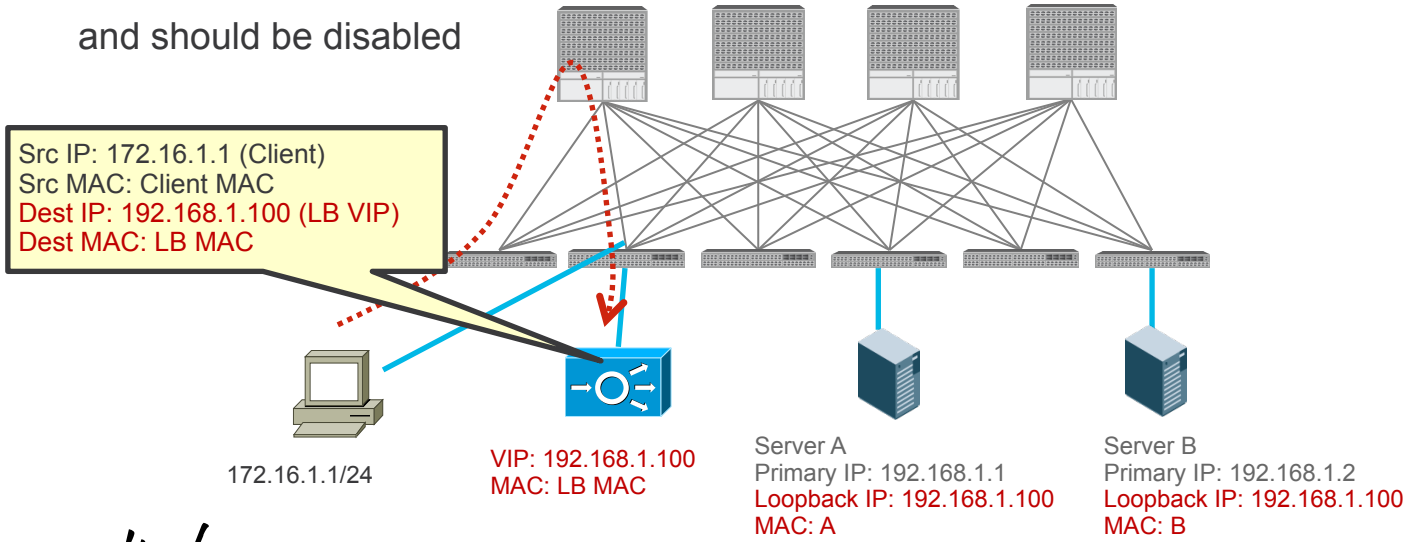

DSR (Direct Server Return)

- It is commonly deployed when Return traffic is huge and want to reduce load of ADC
- Servers must suppress ARP response to VIP ARP requests
- ACI Dataplane end-point learning

  causes VIP move within fabric

  and should be disabled

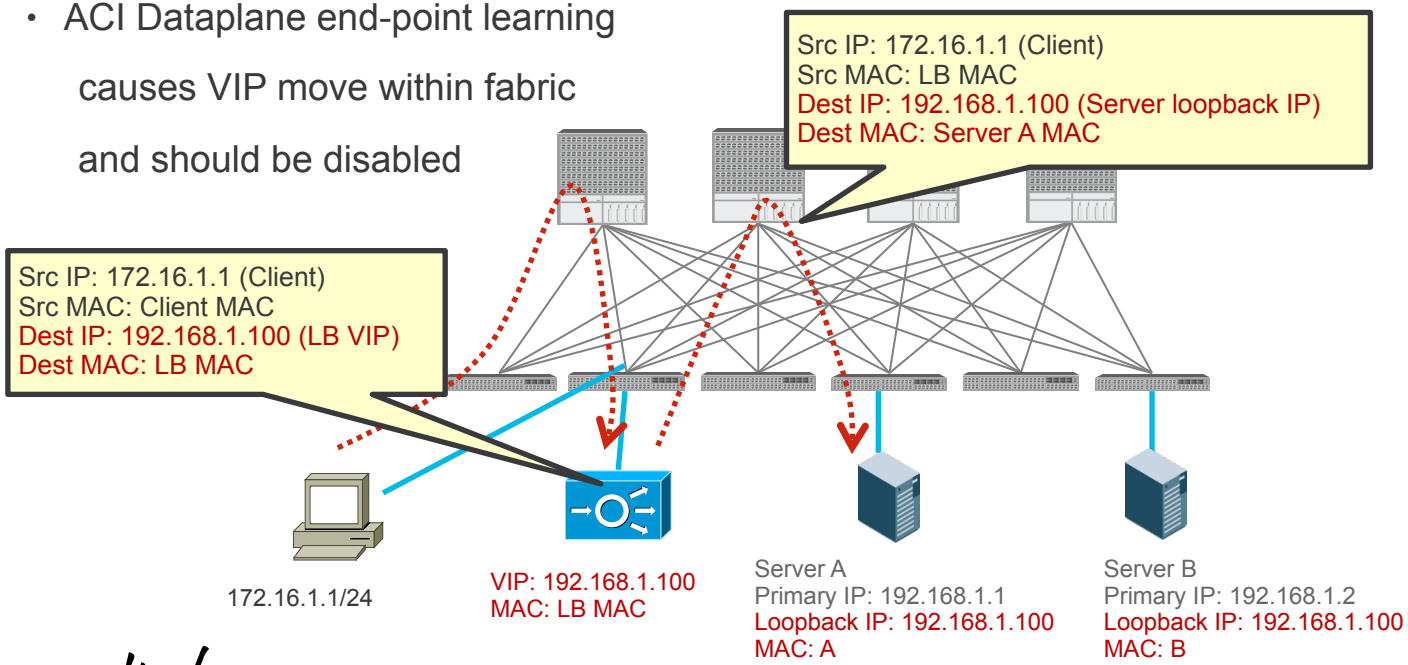

DSR (Direct Server Return)

- It is commonly deployed when Return traffic is huge and want to reduce load of ADC
- Servers must suppress ARP response to VIP ARP requests
- ACI Dataplane end-point learning

  causes VIP move within fabric

  and should be disabled

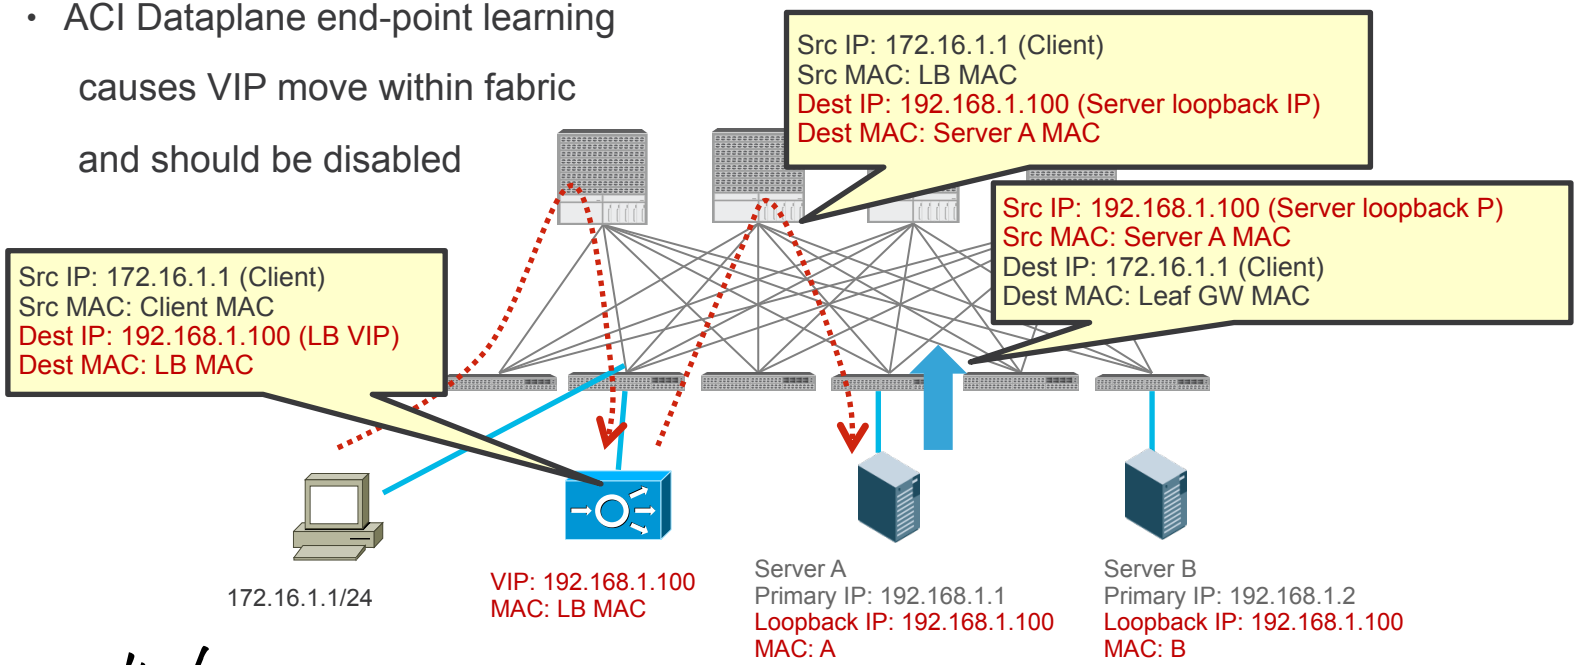

DSR (Direct Server Return)

- It is commonly deployed when Return traffic is huge and want to reduce load of ADC
- Servers must suppress ARP response to VIP ARP requests
- ACI Dataplane end-point learning

  causes VIP move within fabric

  and should be disabled

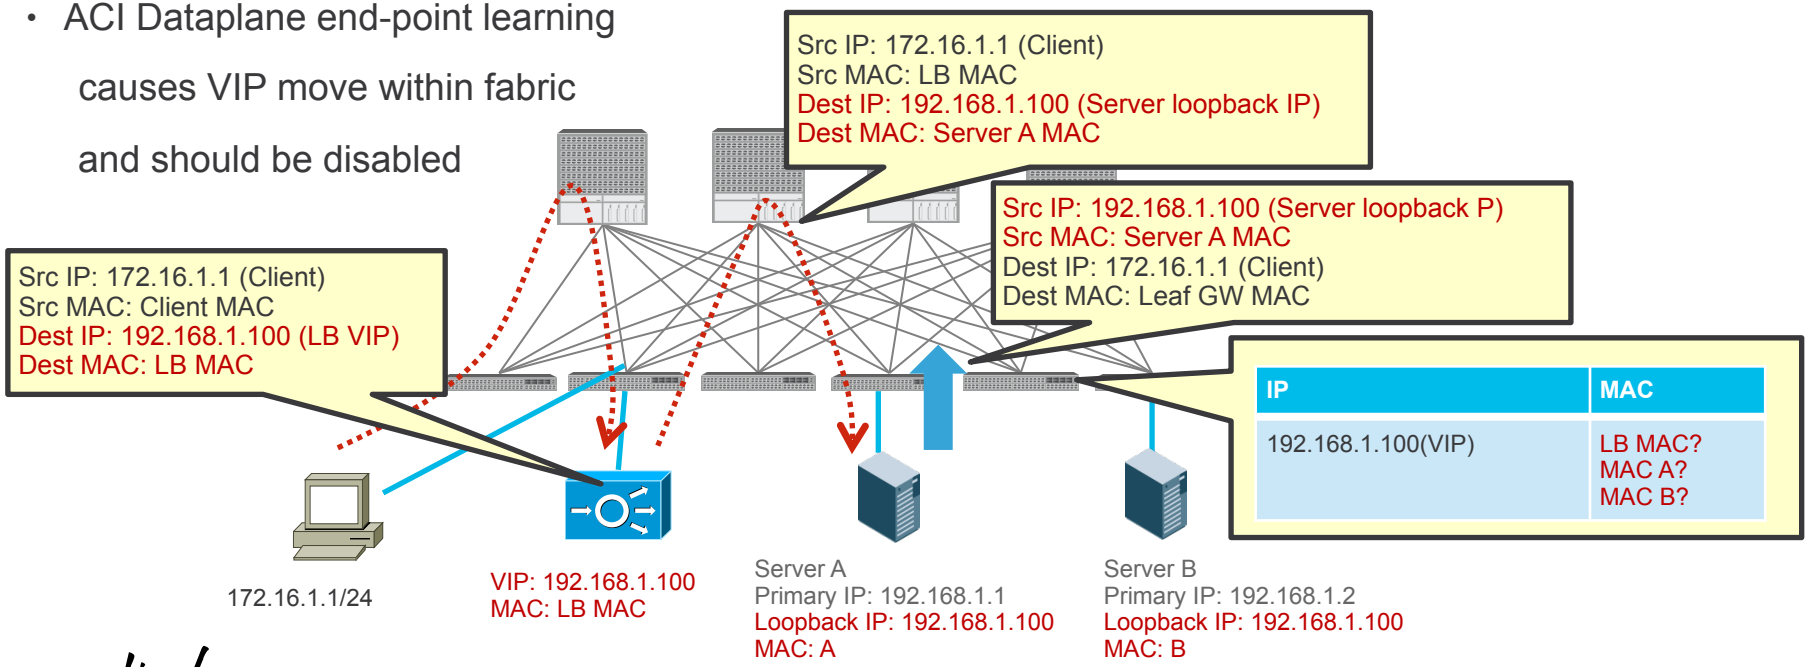

DSR Design Consideration

- Stop Data-path learning of VIP shared by LB and servers.
- Configure the VIP as static under server EPG > L4-L7 Virtual IPs
- Allow ACI Fabric to learn VIP to MAC address binding only from ARP, GARP, or ND
- Server and load-balancer EPGs must be in the same bridge domain.
- GARP learning needs to be enabled
- This feature applies to the L2-DSR scenario only. Not to L3-DSR.

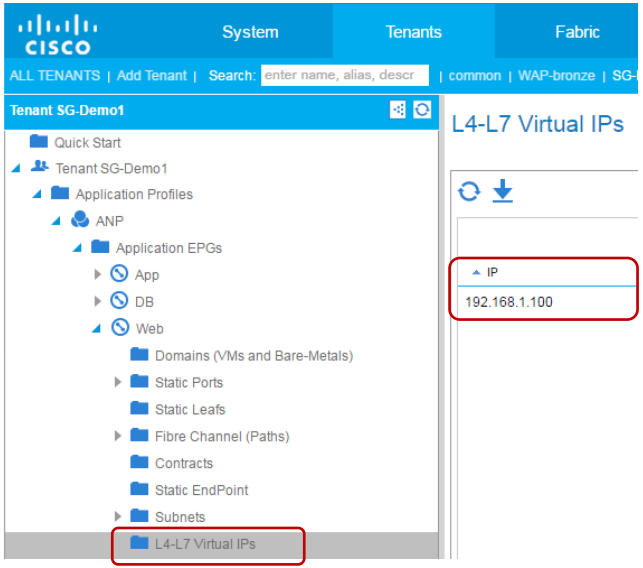

IPS/IDS Designs

IPS Mode Options

	L1 (Fail-wire)	L2 (Transparent)	L3 (Routed)
	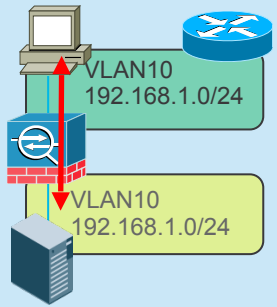	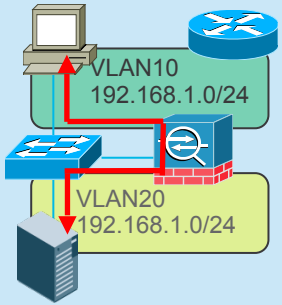	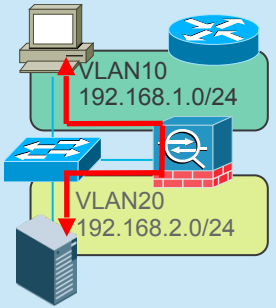
Default GW of servers	Router	Router	IPS alias IP of server-side (in case of FirePOWER, active IP)
When to use	Want to insert IPS inline	Want to use same subnet for client-side and server-side	Common case when we can have separate VLAN/subnet for client-side and server-side.
comment	Firepower appliances support L1 bypass network modules		

L2 and L3 mode are pretty much same with Firewall design. Next slides will focus on L1 mode.

IPS Design Consideration

- For L1 mode, service device doesn't change VLAN ID. Traffic from both legs of service node uses same VLAN encap.
- To put inline device accordingly, use of different BD is needed.

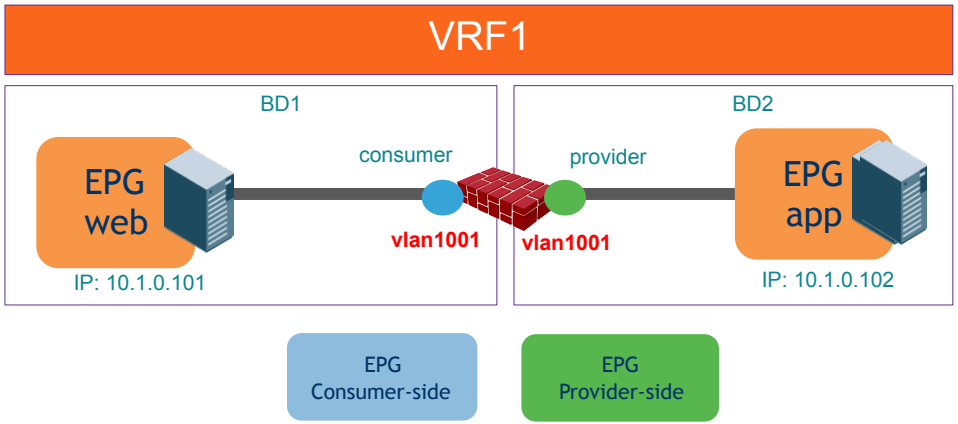

PBR

PBR
Frequently Asked Design Questions

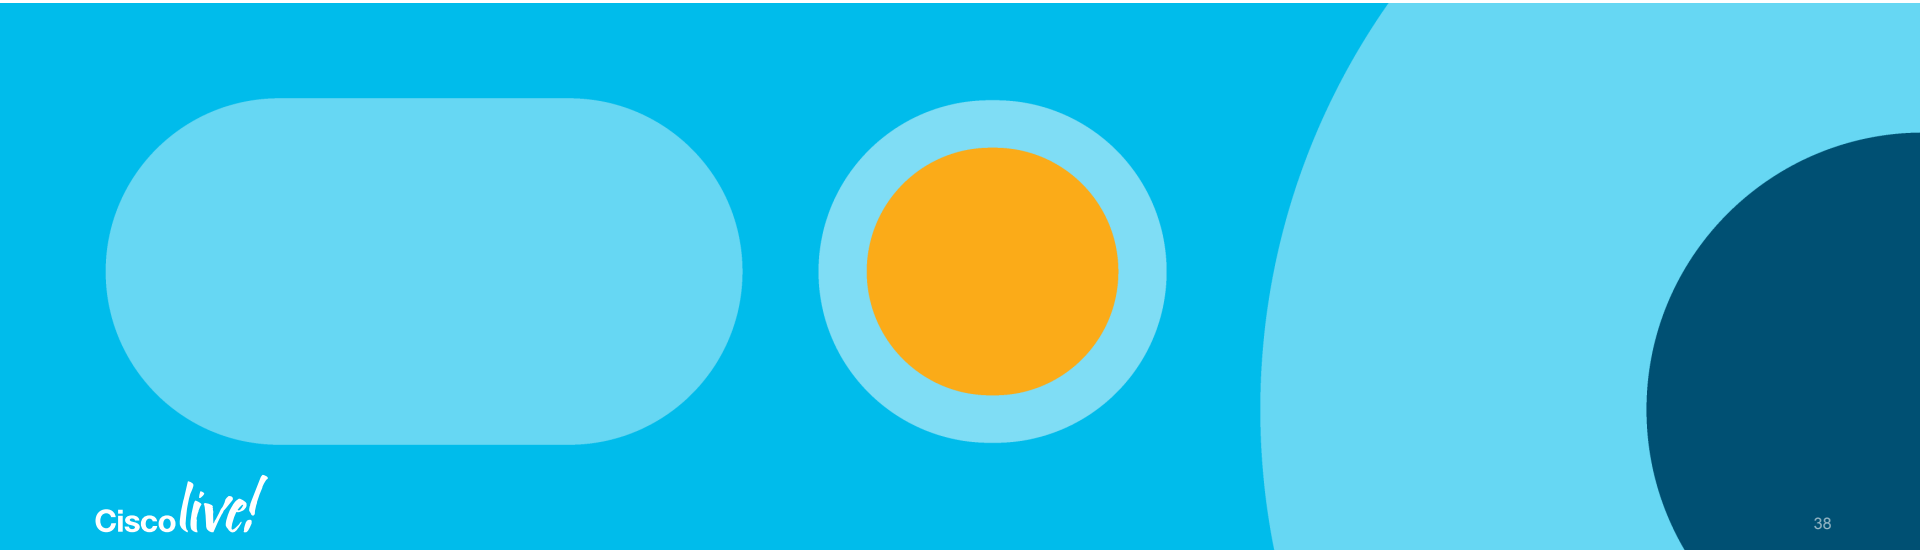

Can we use PBR for L3out EPG?

Can we use PBR for L3out EPG?

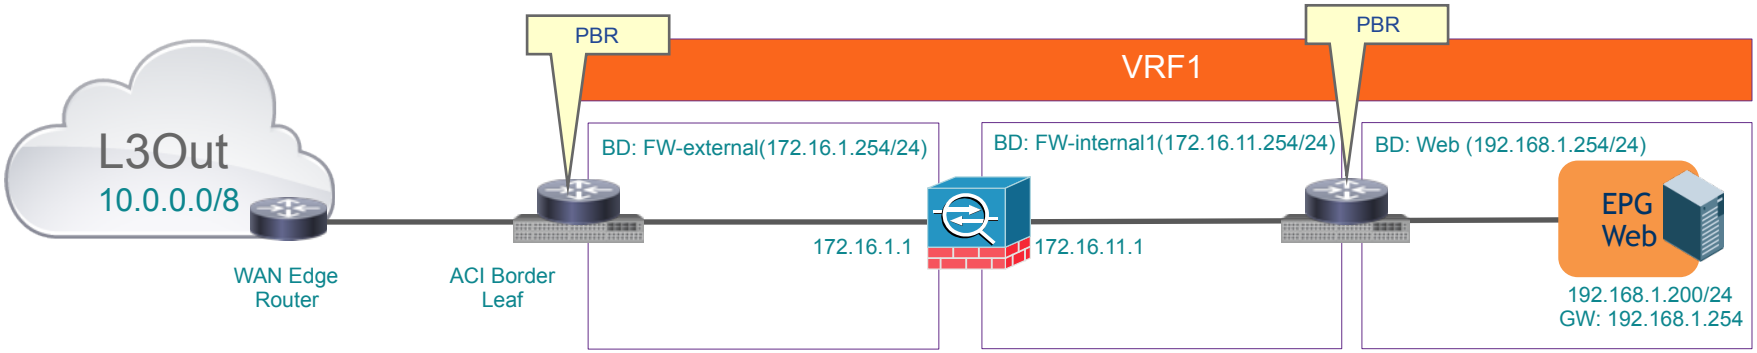

Can we use PBR for L3out EPG?

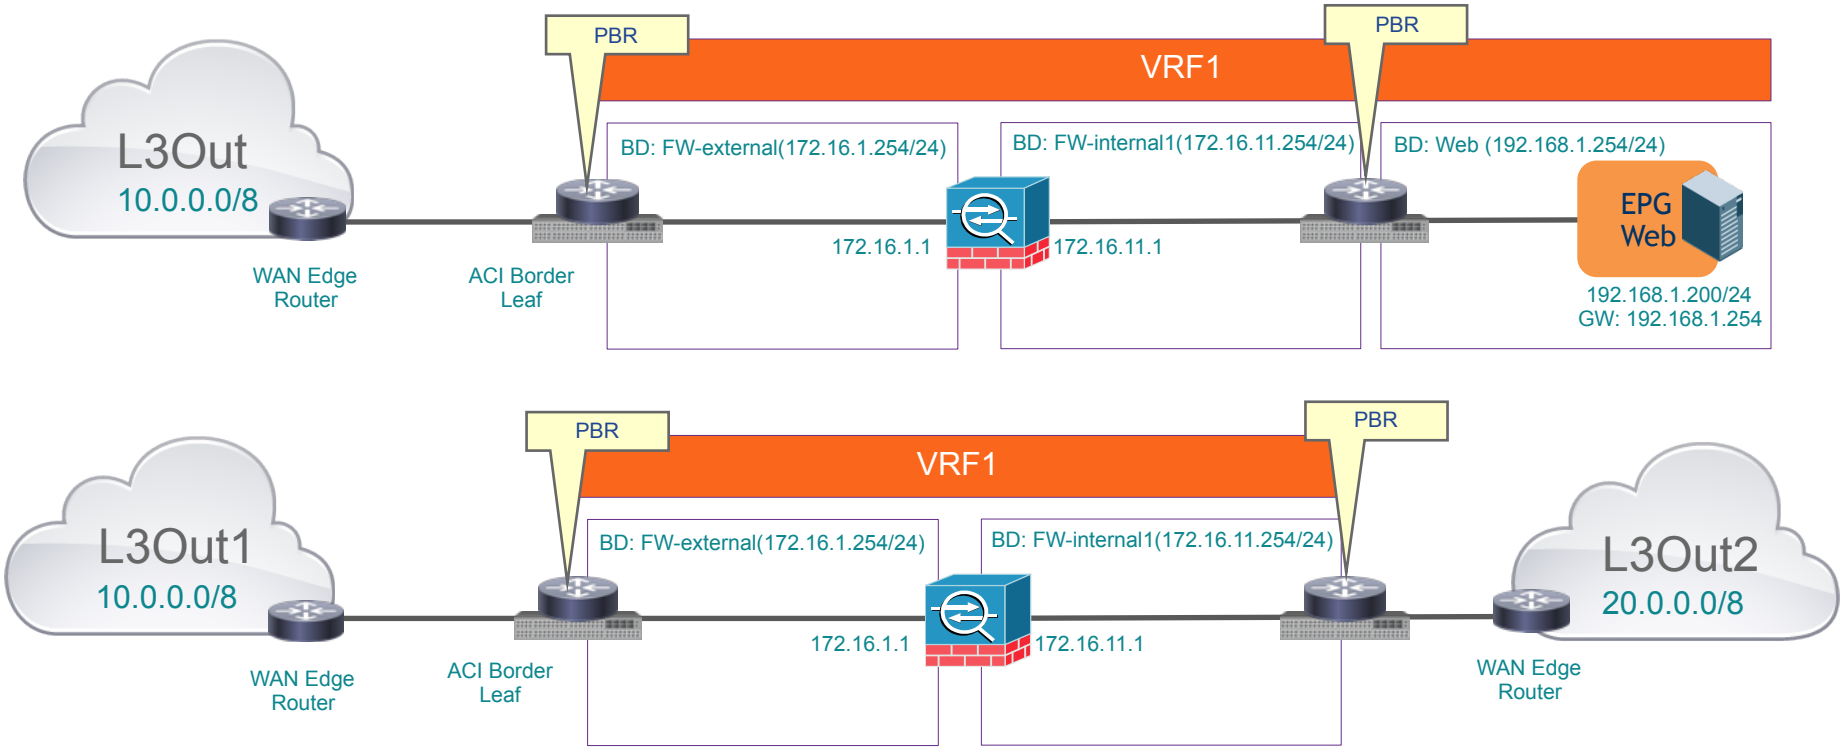

Can we use PBR for EPGs in same subnet?

Can we use PBR for EPGs in same subnet?

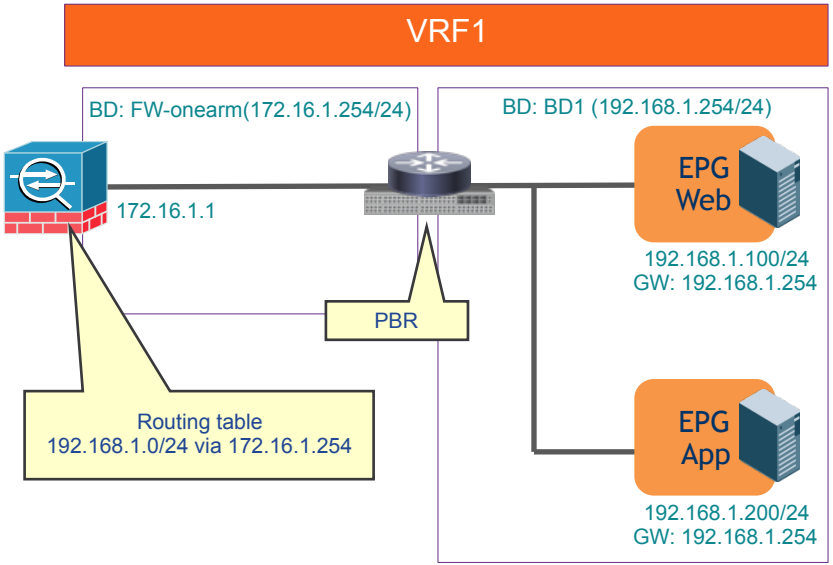

Can we use PBR for EPGs in same subnet?

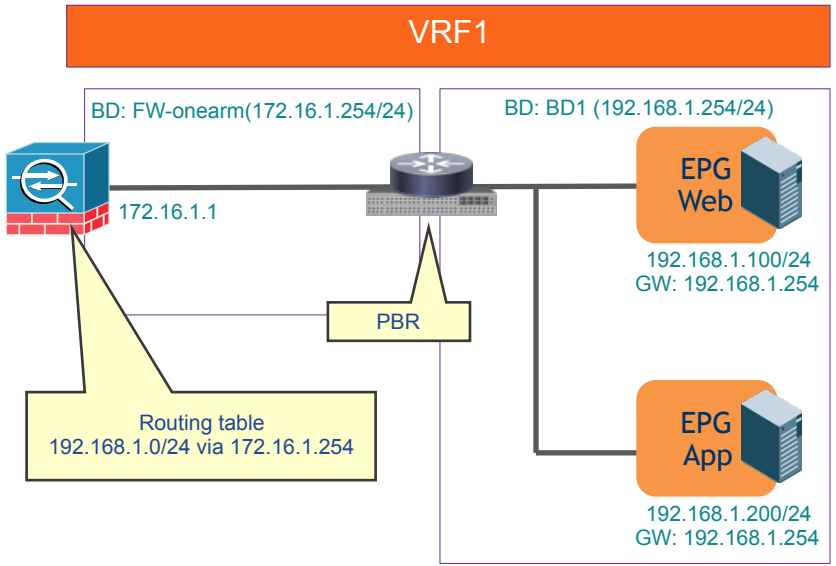

- Consumer and provider EPGs can be in same BD subnet.

Can we use PBR for EPGs in same subnet?

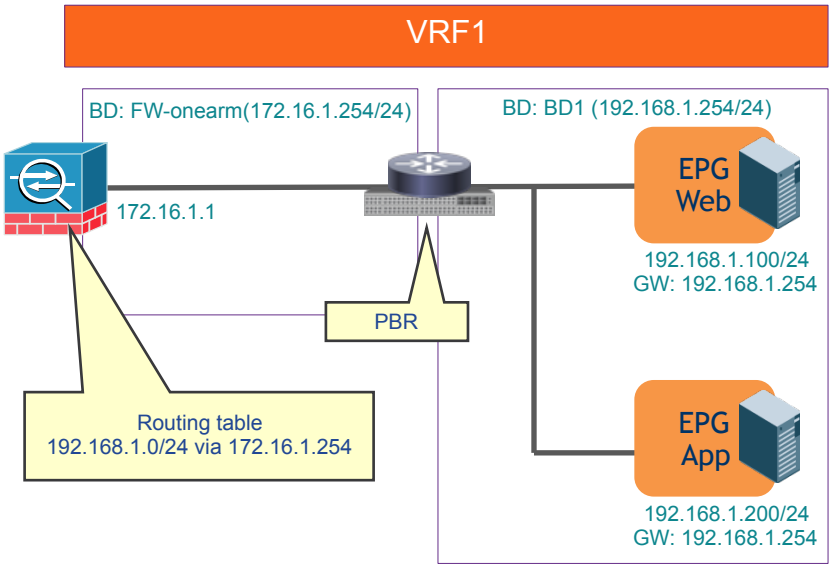

- Consumer and provider EPGs can be in same BD subnet.
- One-arm mode deployment

Can we use PBR for EPGs in same subnet?

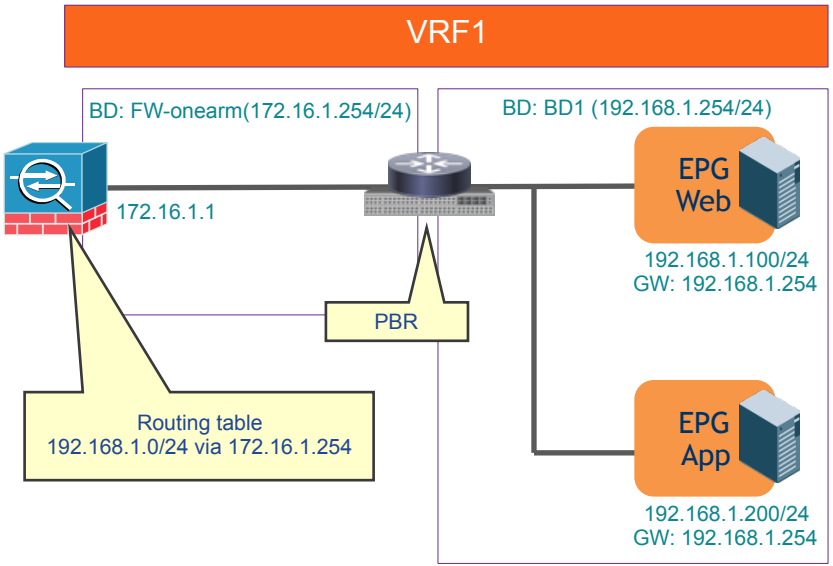

- Consumer and provider EPGs can be in same BD subnet.
- One-arm mode deployment

Can we use PBR for EPGs in same subnet?

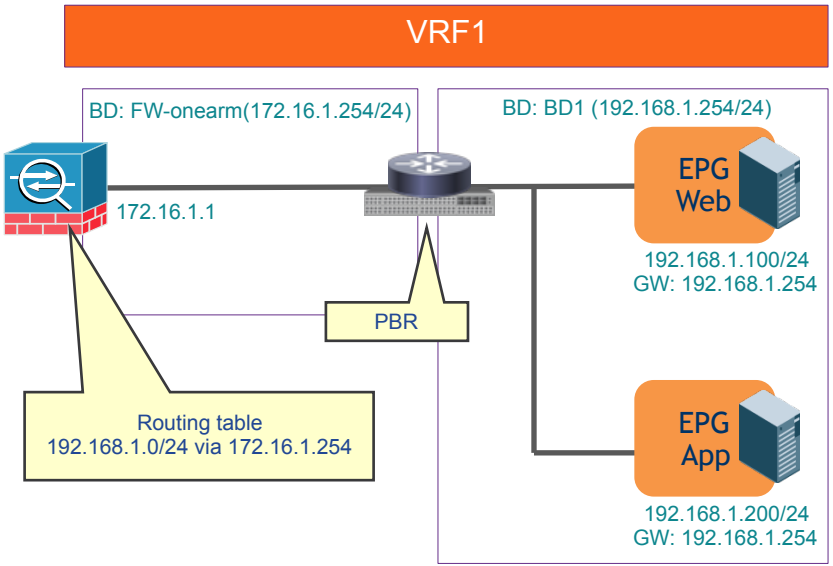

- Consumer and provider EPGs can be in same BD subnet.
- One-arm mode deployment
- Note
  - ARP must be excluded from PBR
  - Service Graph for intra-EPG contract is not supported today.

Can we reuse same PBR node for multiple times?

Can we reuse same PBR node for multiple times?

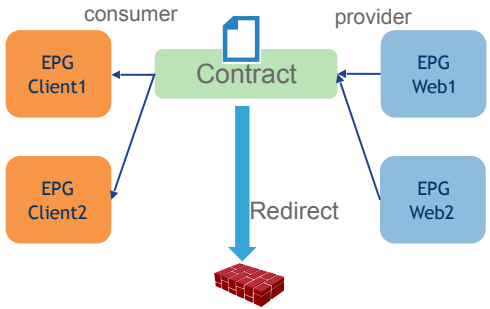

Can we reuse same PBR node for multiple times?

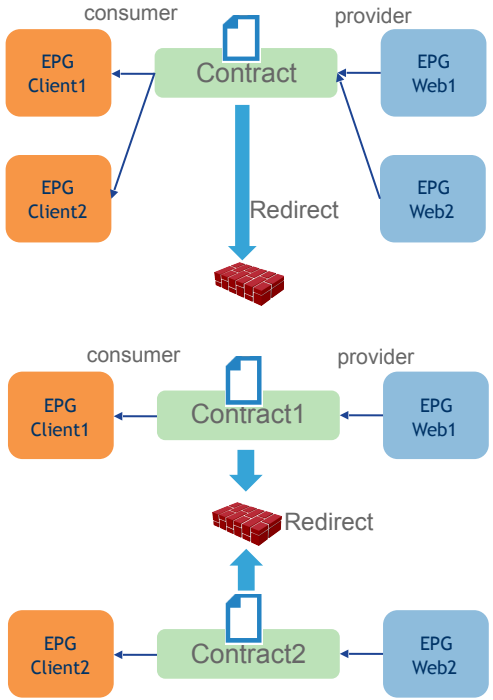

Can we reuse same PBR node for multiple times?

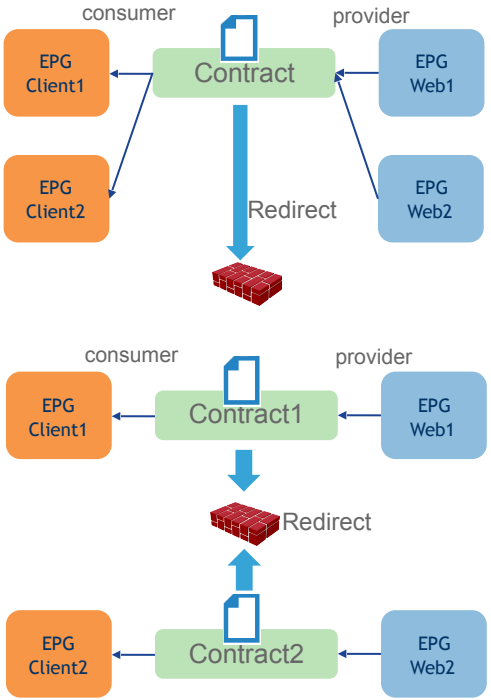

- Multiple consumer/provider EPGs

Can we reuse same PBR node for multiple times?

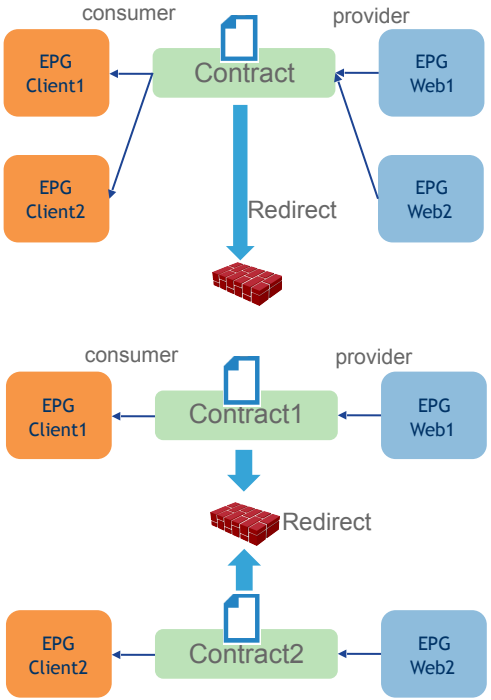

- Multiple consumer/provider EPGs
- Multiple contracts using same PBR destination.

Can we reuse same PBR node for multiple times?

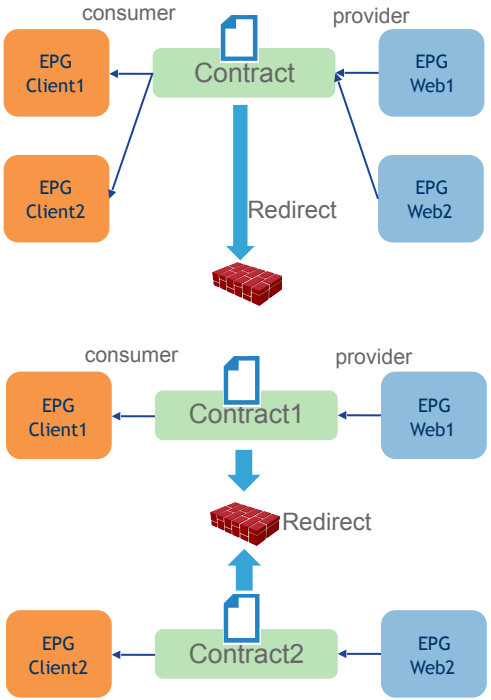

- Multiple consumer/provider EPGs
- Multiple contracts using same PBR destination.

Can we reuse same PBR node for multiple times?

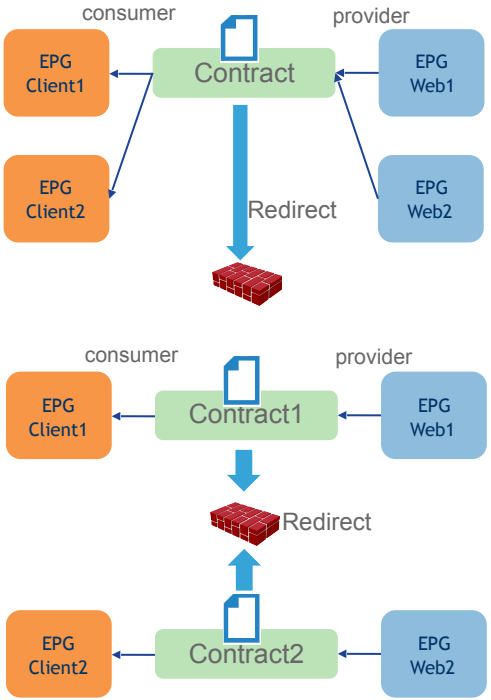

- Multiple consumer/provider EPGs
- Multiple contracts using same PBR destination.

- Note
  - Depending on routing design, one-arm mode deployment may be required.

Can PBR node be in consumer/provider subnet?

Can PBR node be in consumer/provider subnet?

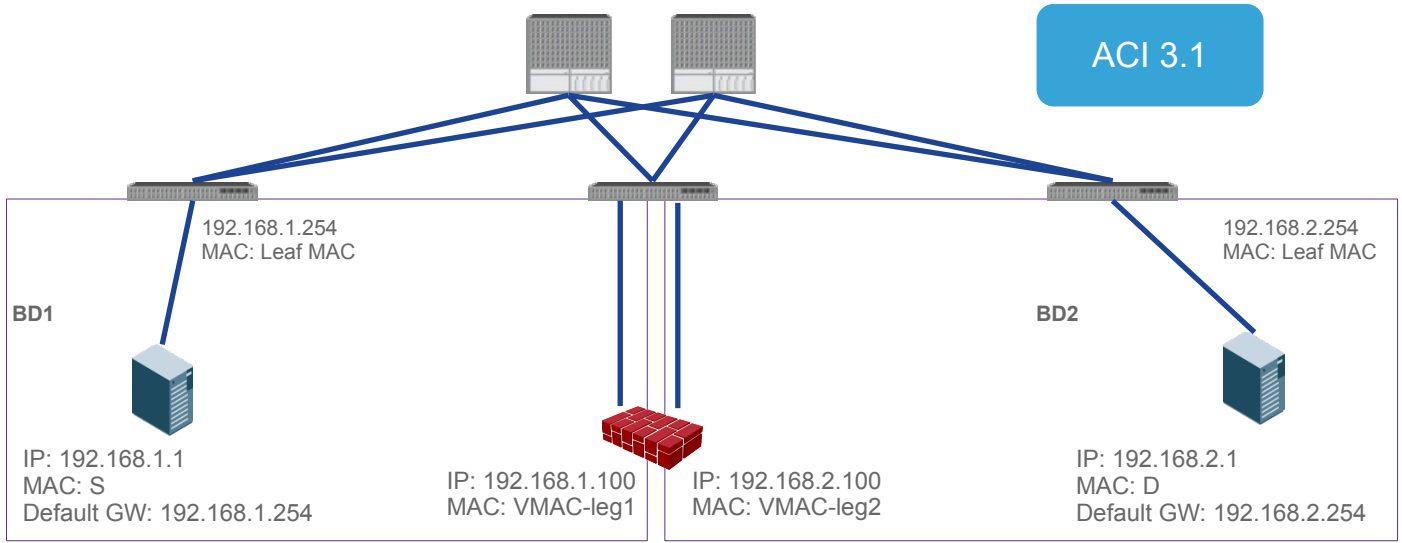

Can PBR node be in consumer/provider subnet?

- Prior to APIC version 3.1, PBR node must be different than the consumer/provider BDs.

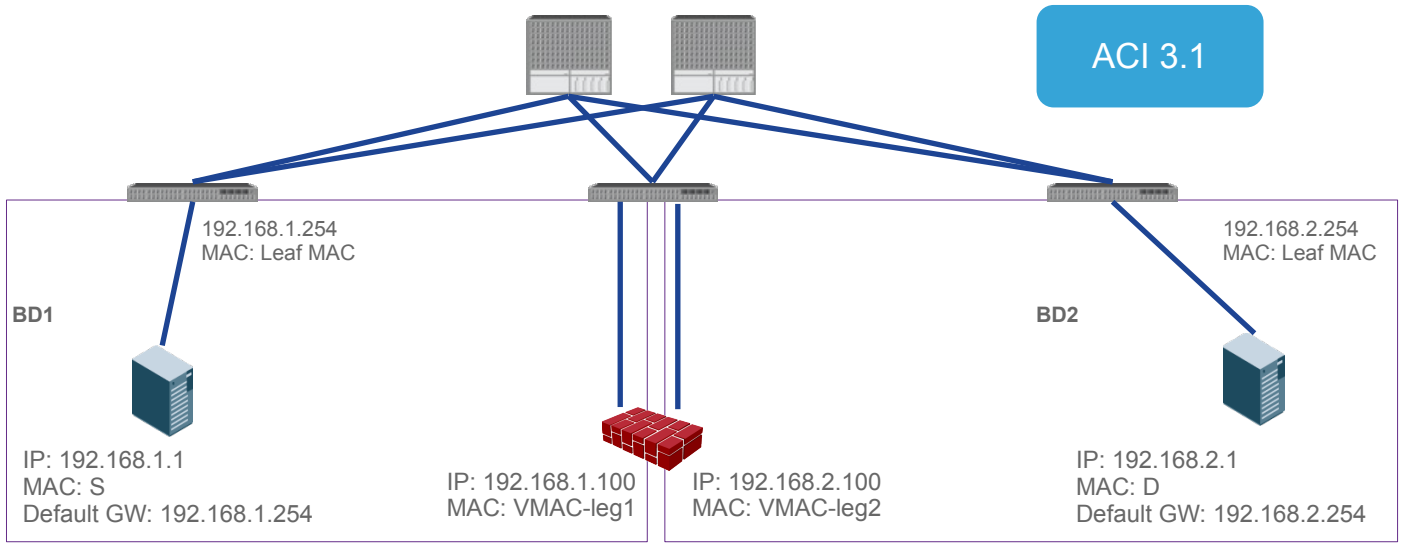

Can PBR node be in consumer/provider subnet?

- Prior to APIC version 3.1, PBR node must be different than the consumer/provider BDs.
- Starting from APIC version 3.1, this requirement no longer mandatory. (Need EX/FX Leaf)

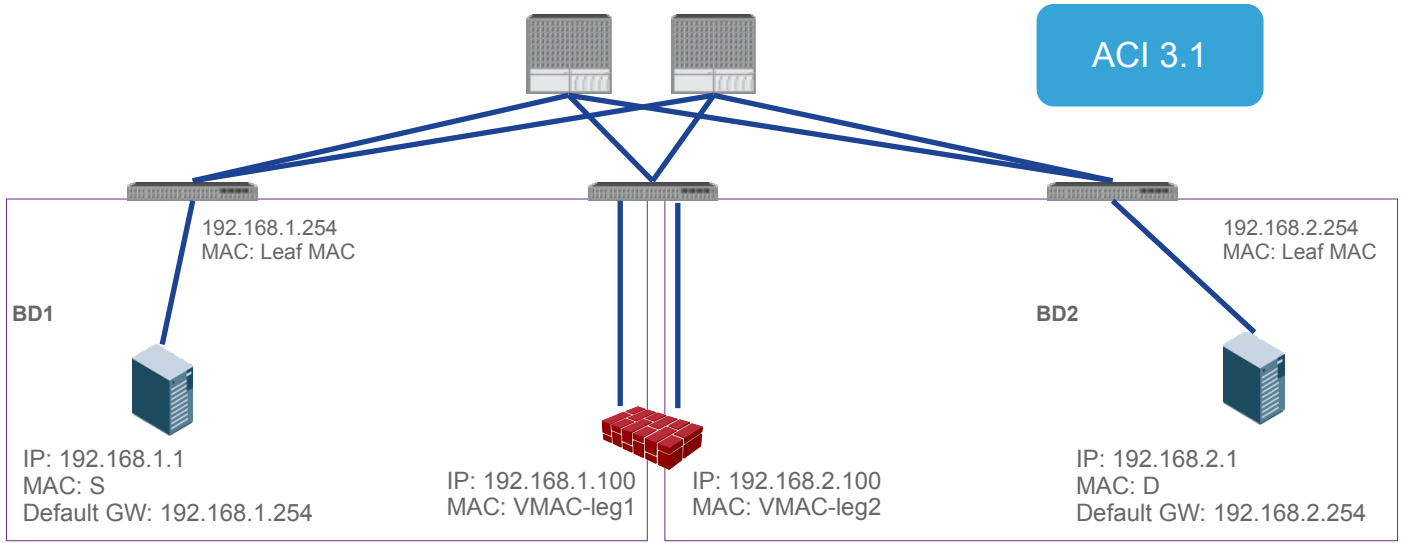

Can we use PBR for inter VRF contract?

Can we use PBR for inter VRF contract?

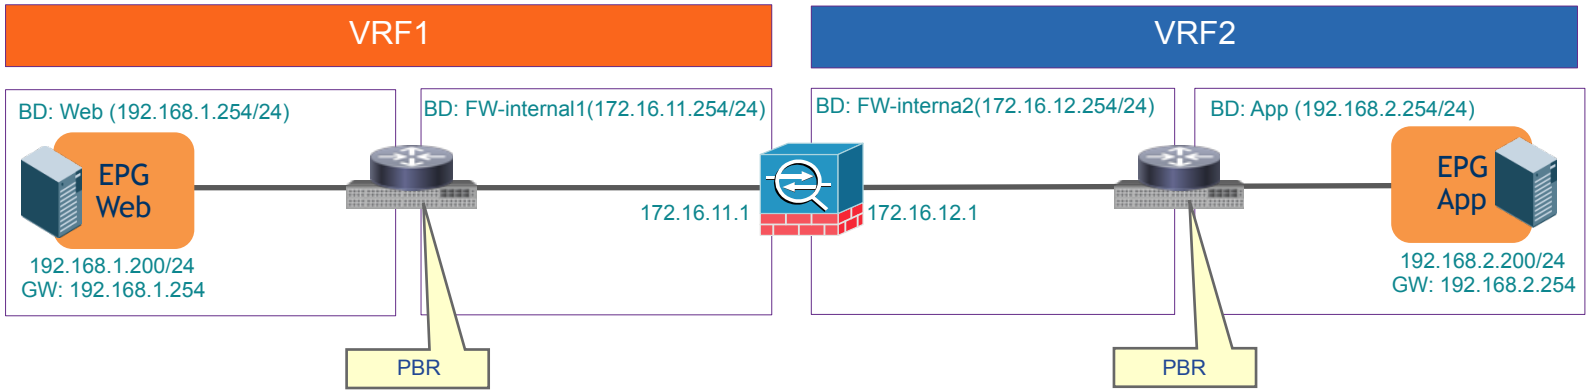

Can we use PBR for inter VRF contract?

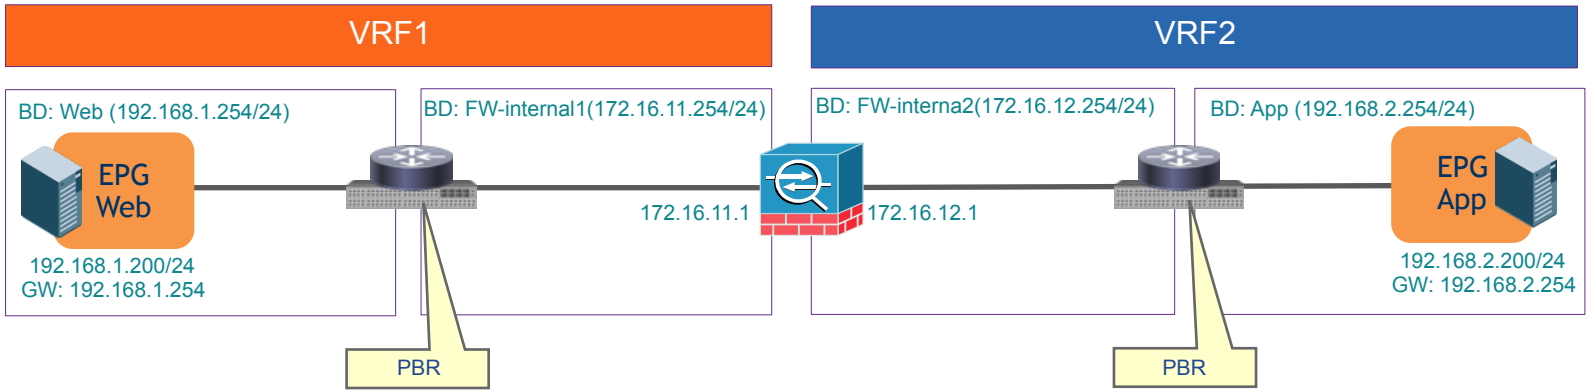

- The PBR node can be between VRF instances or within one of the VRF instances. The PBR node must be in either the consumer or provider VRF instance. For example, you cannot put the PBR node in VRF3, which is neither a consumer nor a provider VRF instance.

How PBR works for FW + LB Service Graph?

- Incoming traffic

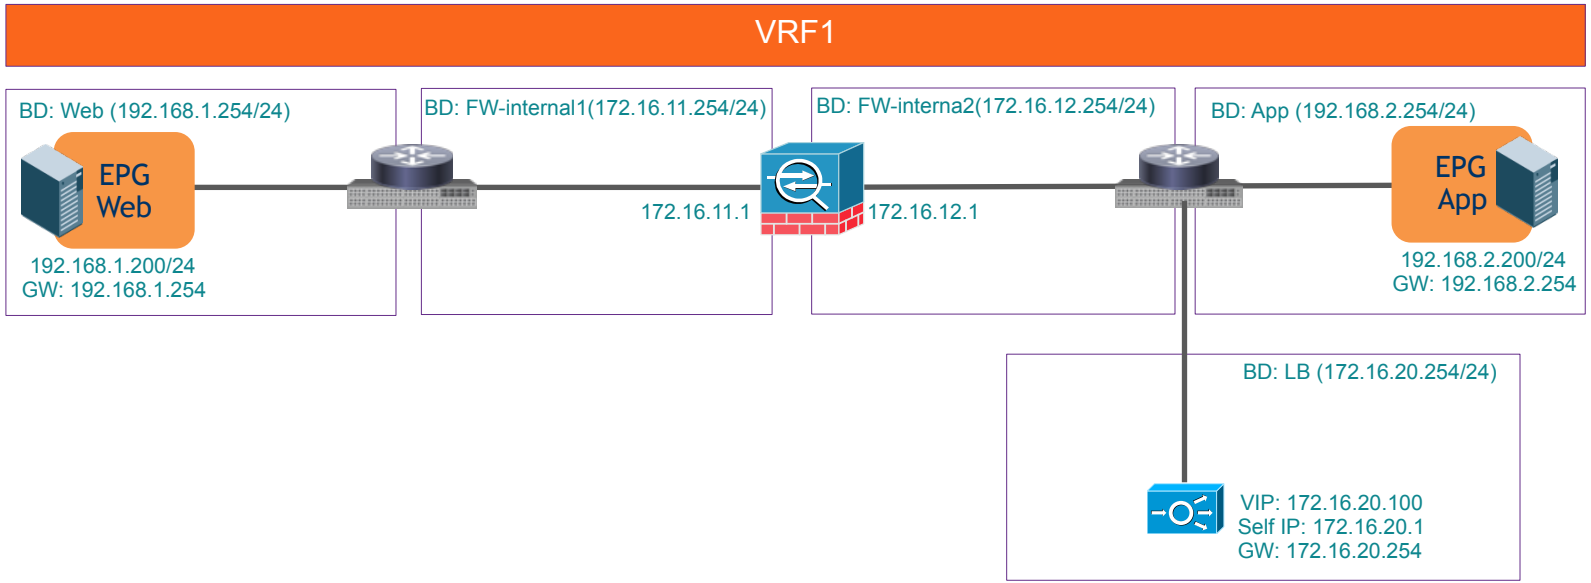

How PBR works for FW + LB Service Graph?

- Incoming traffic

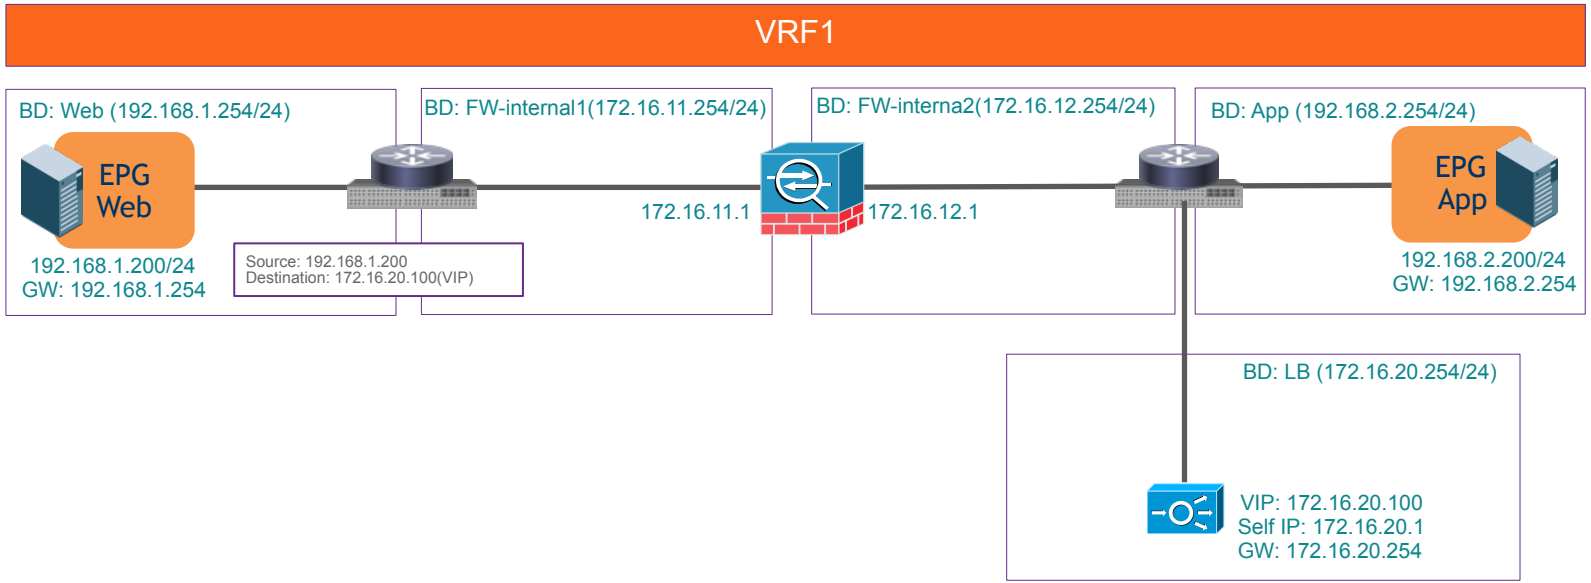

How PBR works for FW + LB Service Graph?

- Incoming traffic

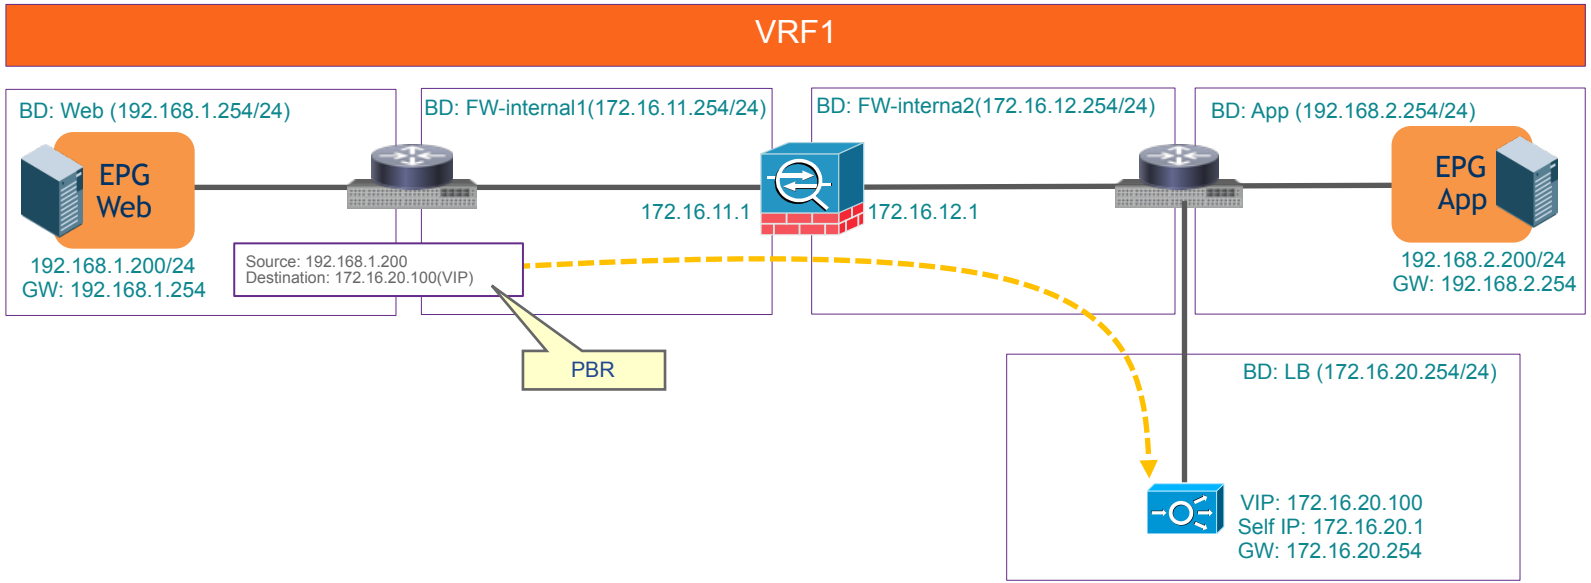

How PBR works for FW + LB Service Graph?

- Incoming traffic

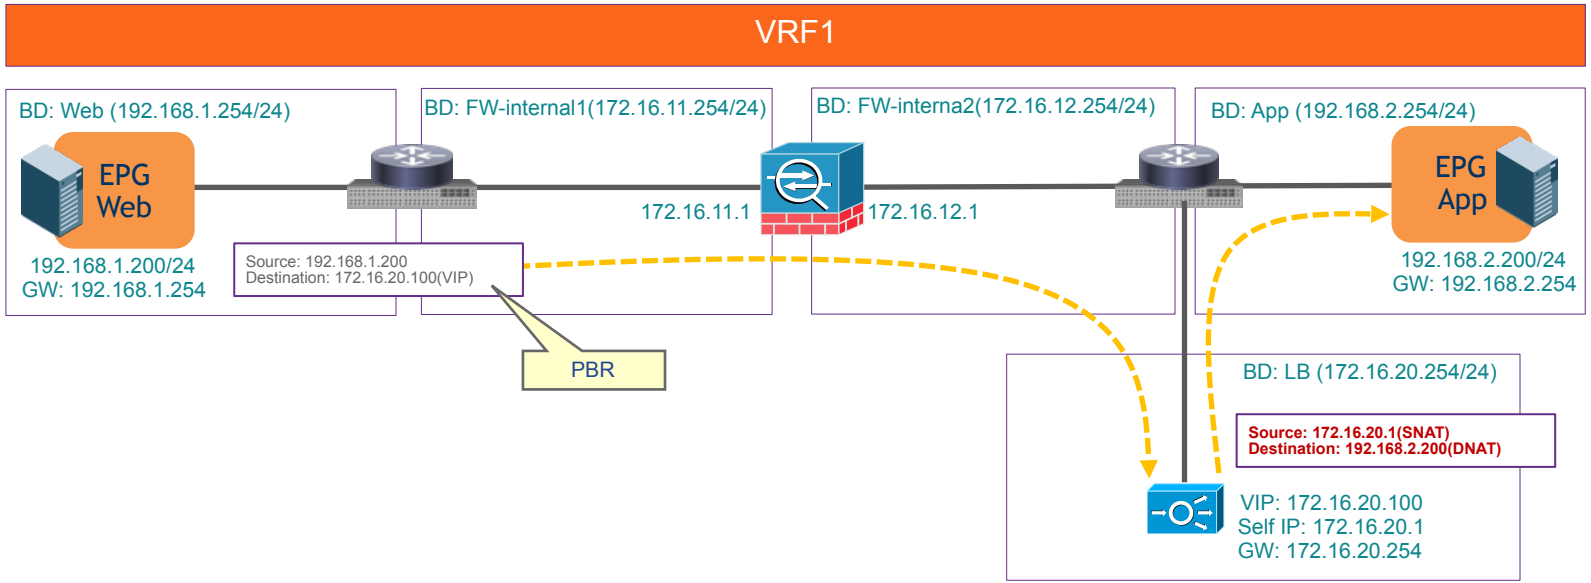

How PBR works for FW + LB Service Graph?

- Return traffic

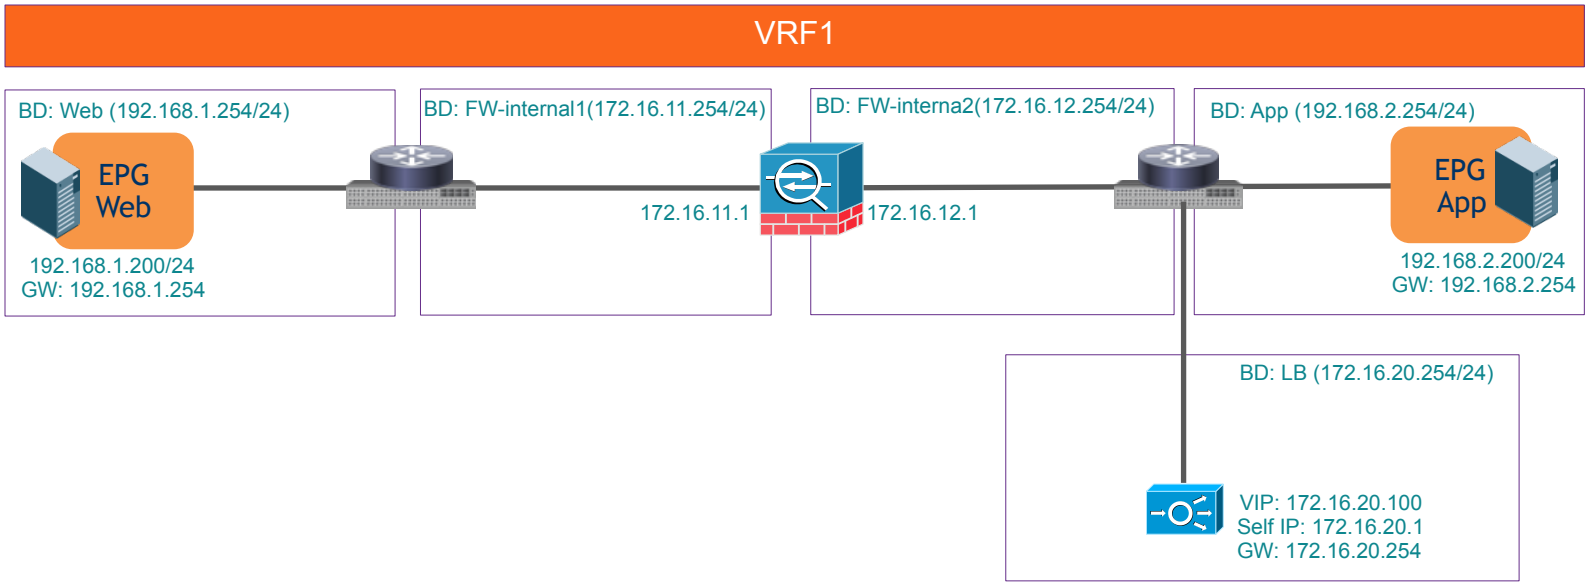

How PBR works for FW + LB Service Graph?

- Return traffic

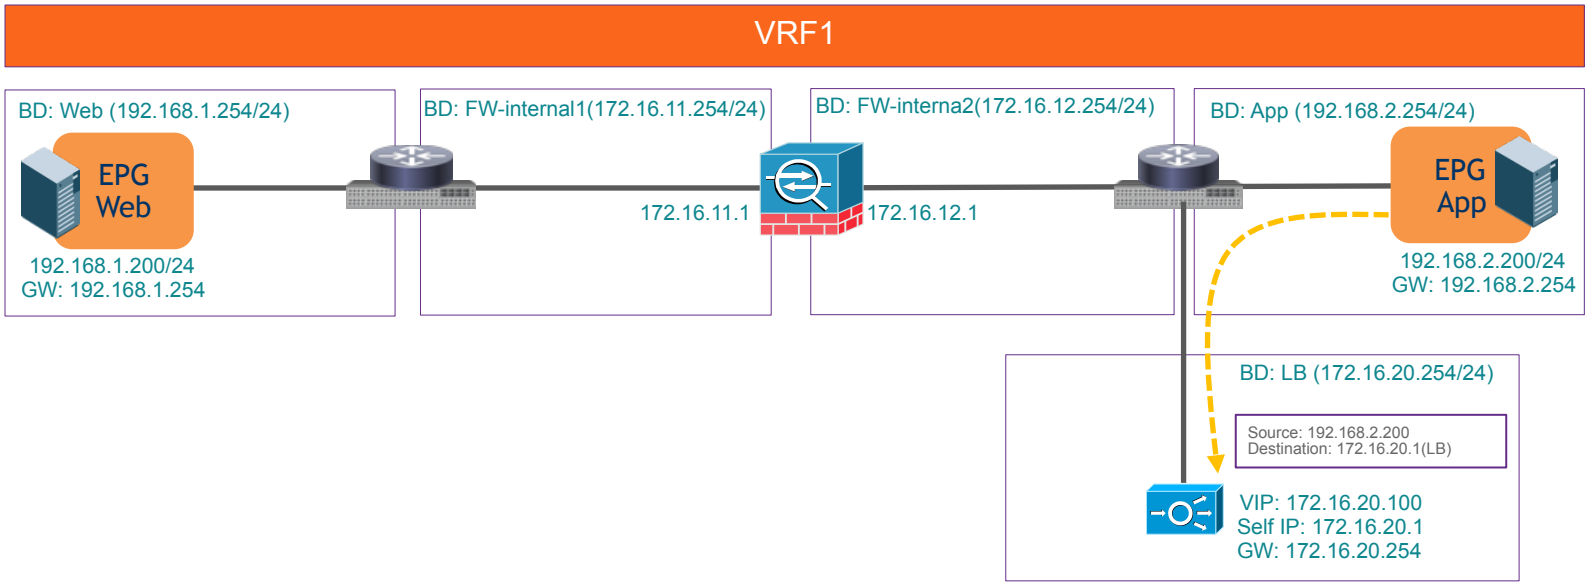

How PBR works for FW + LB Service Graph?

- Return traffic

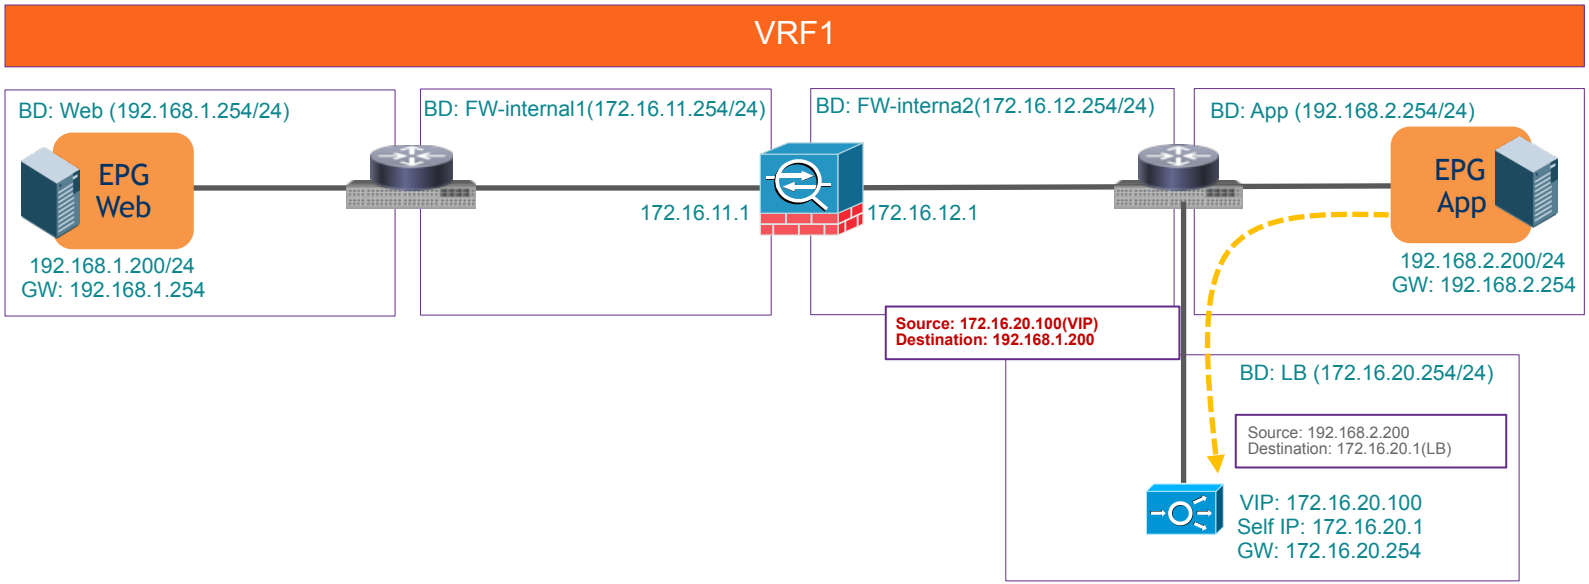

How PBR works for FW + LB Service Graph?

- Return traffic

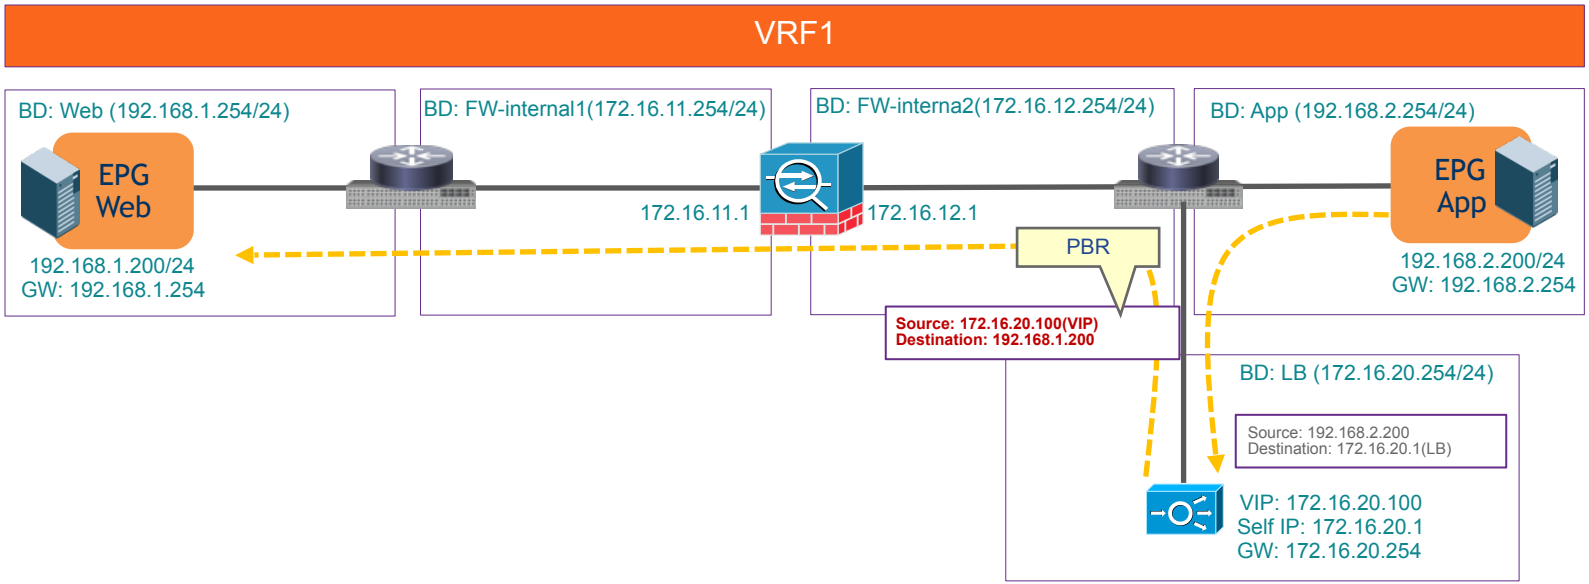

Symmetric PBR
Scale-out service

Symmetric PBR

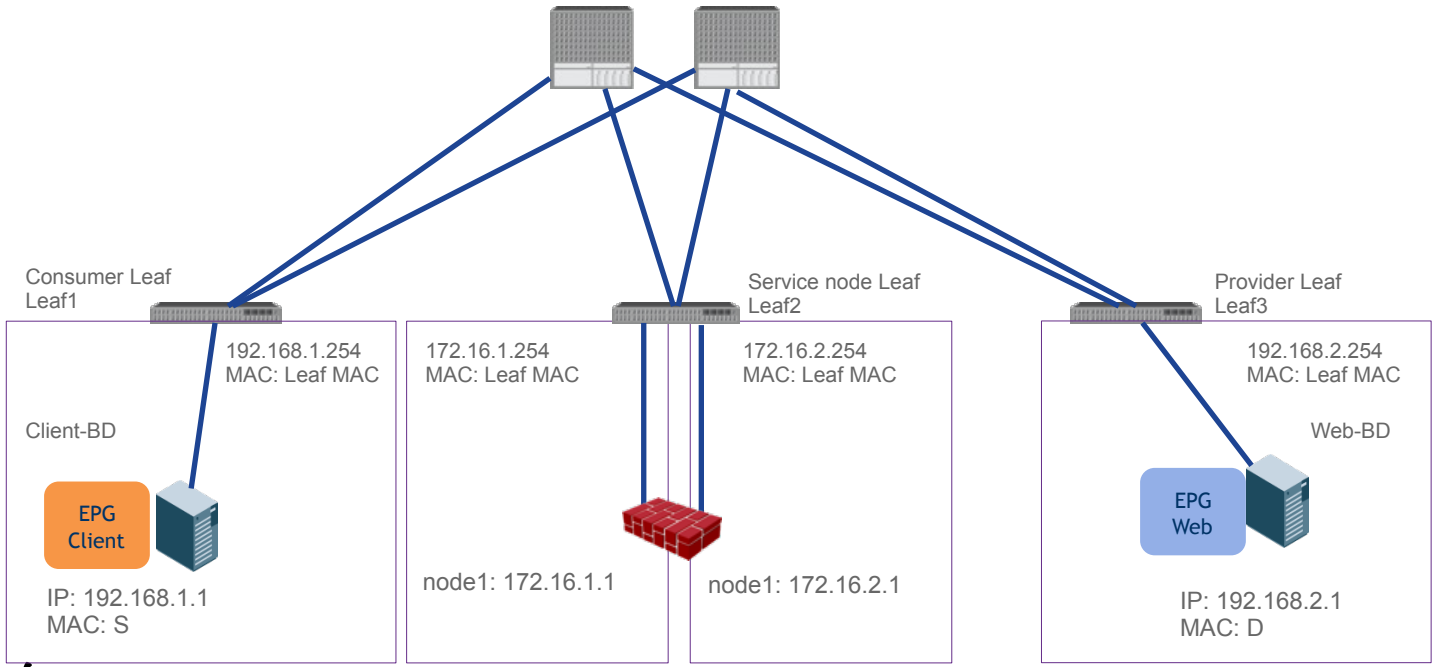

Symmetric PBR

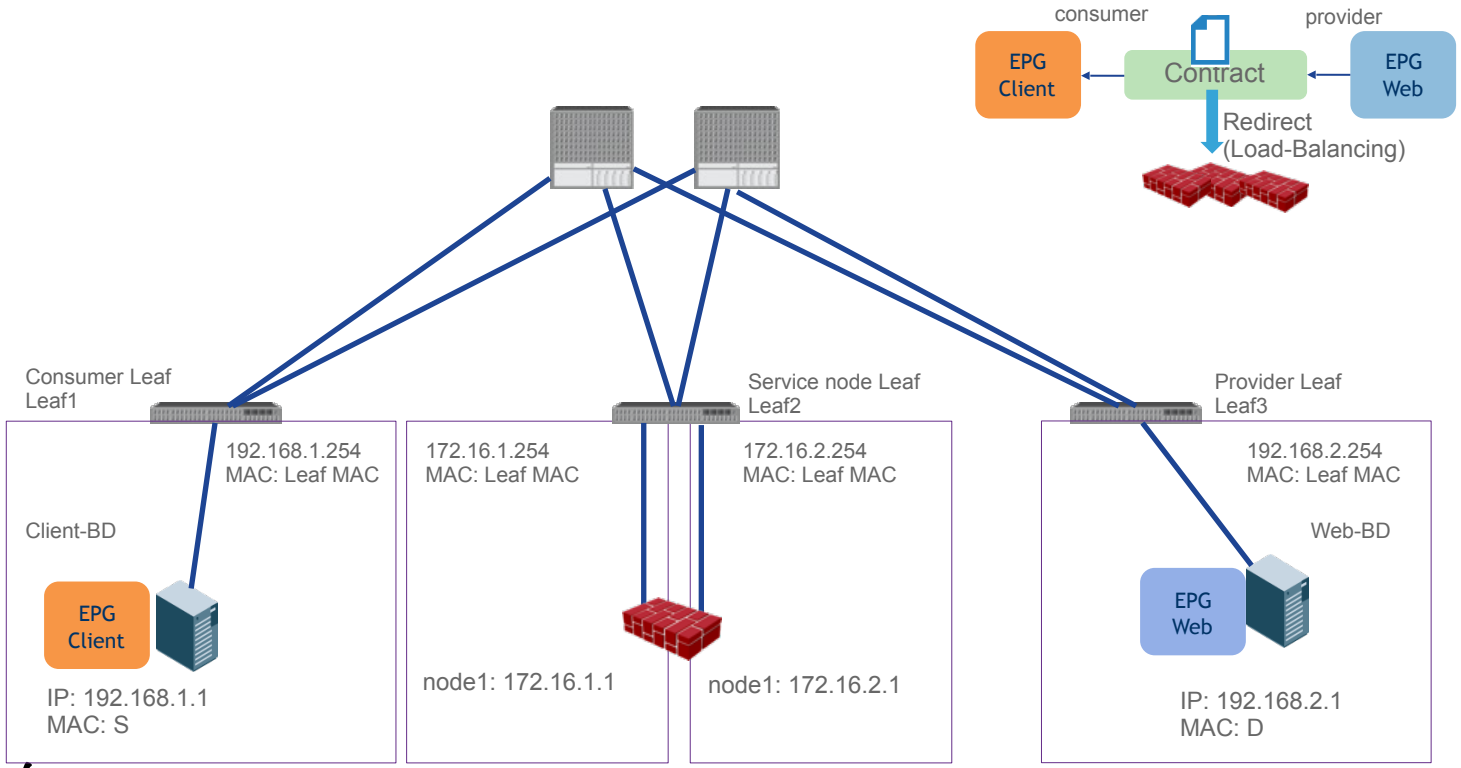

Symmetric PBR

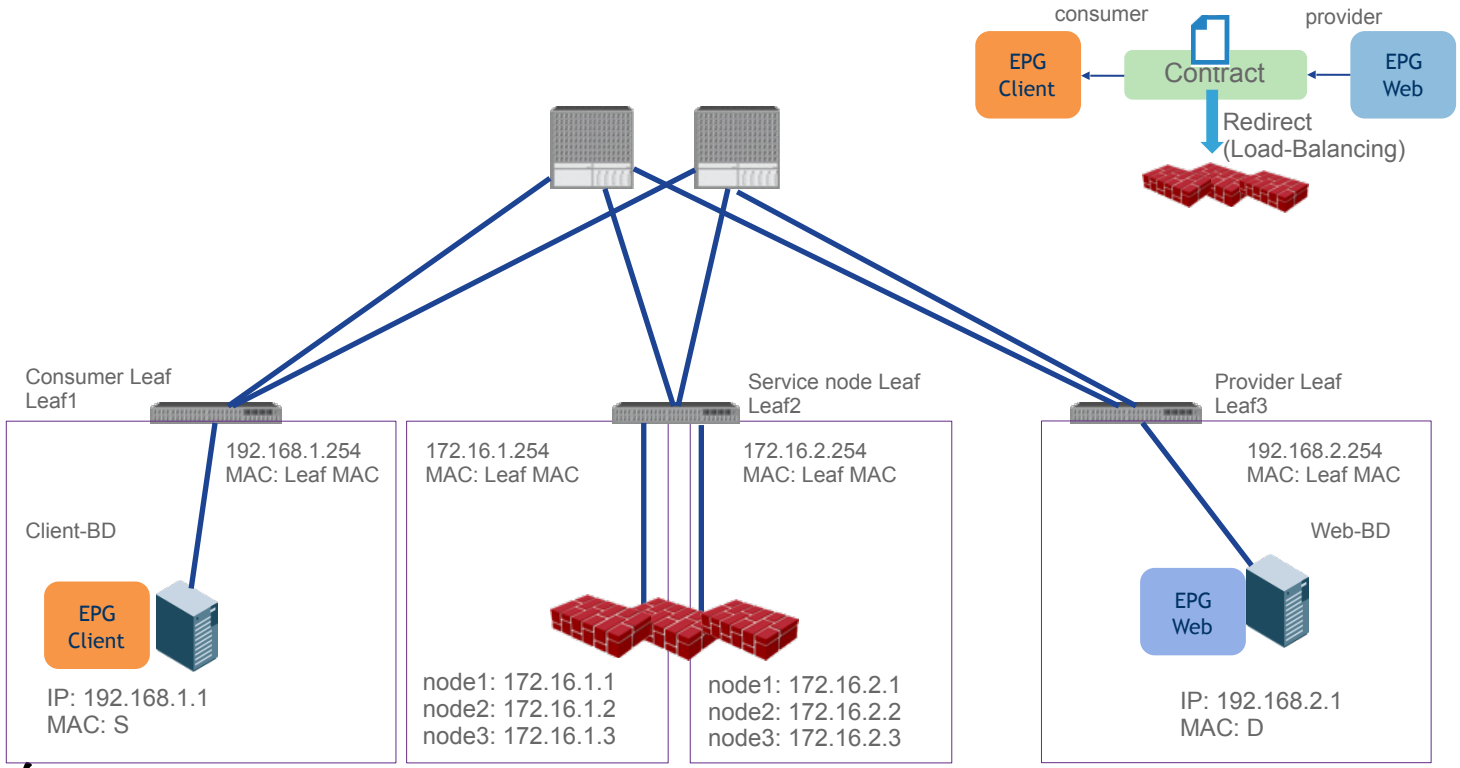

Symmetric PBR

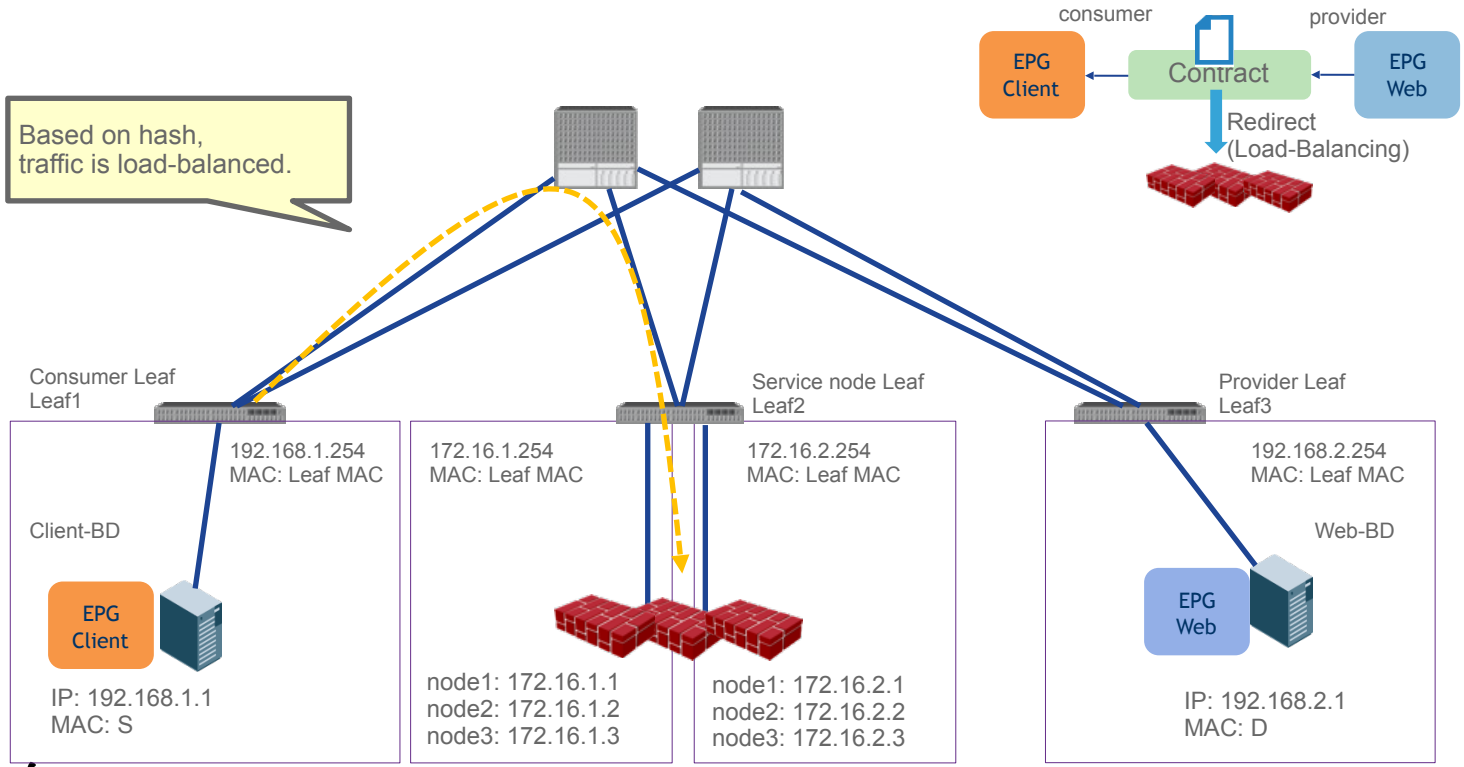

PBR with source IP/dest IP only hash

ACI 2.2(3) 3.1

Use case

- Hashing is based on Source IP, Destination IP and Protocol Type by default.
- Source IP only, Destination IP only hash option is available.
  - If you wants to make same user (same source IP address) go to the same service node.
    - Use of source IP based hash for incoming traffic
    - Use of destination IP based hash for return traffic

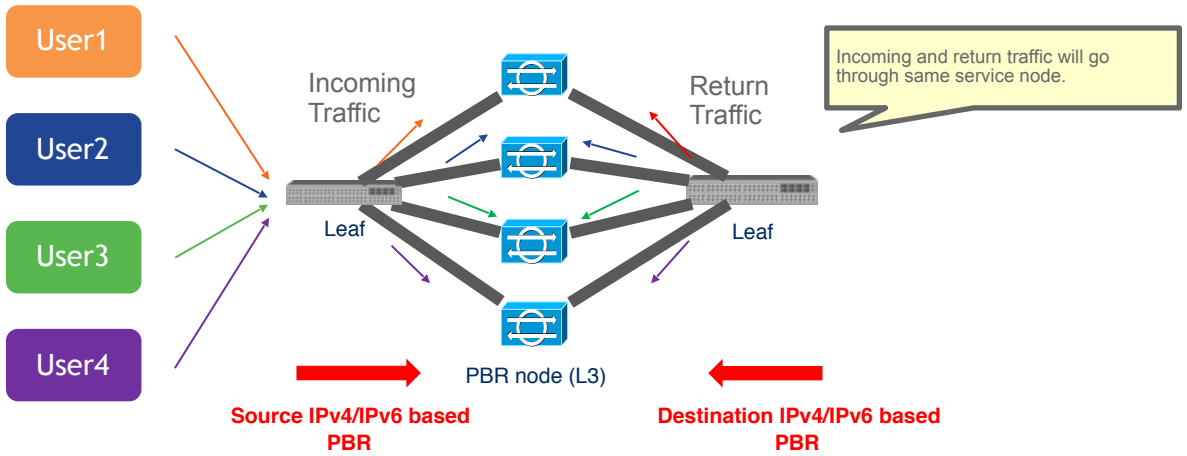

PBR with Tracking

Tracking behavior

ACI 2.2(3) 3.1

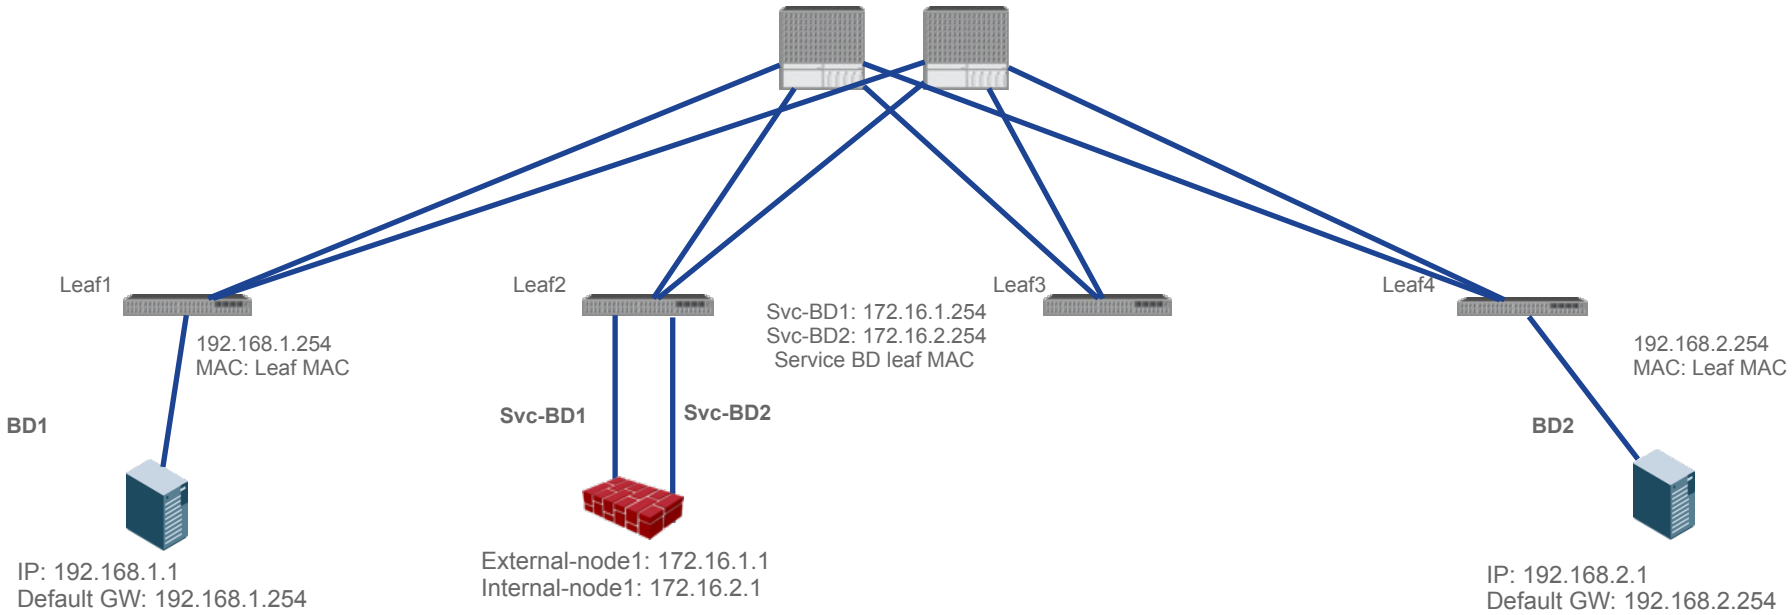

PBR with Tracking

Tracking behavior

ACI 2.2(3) 3.1

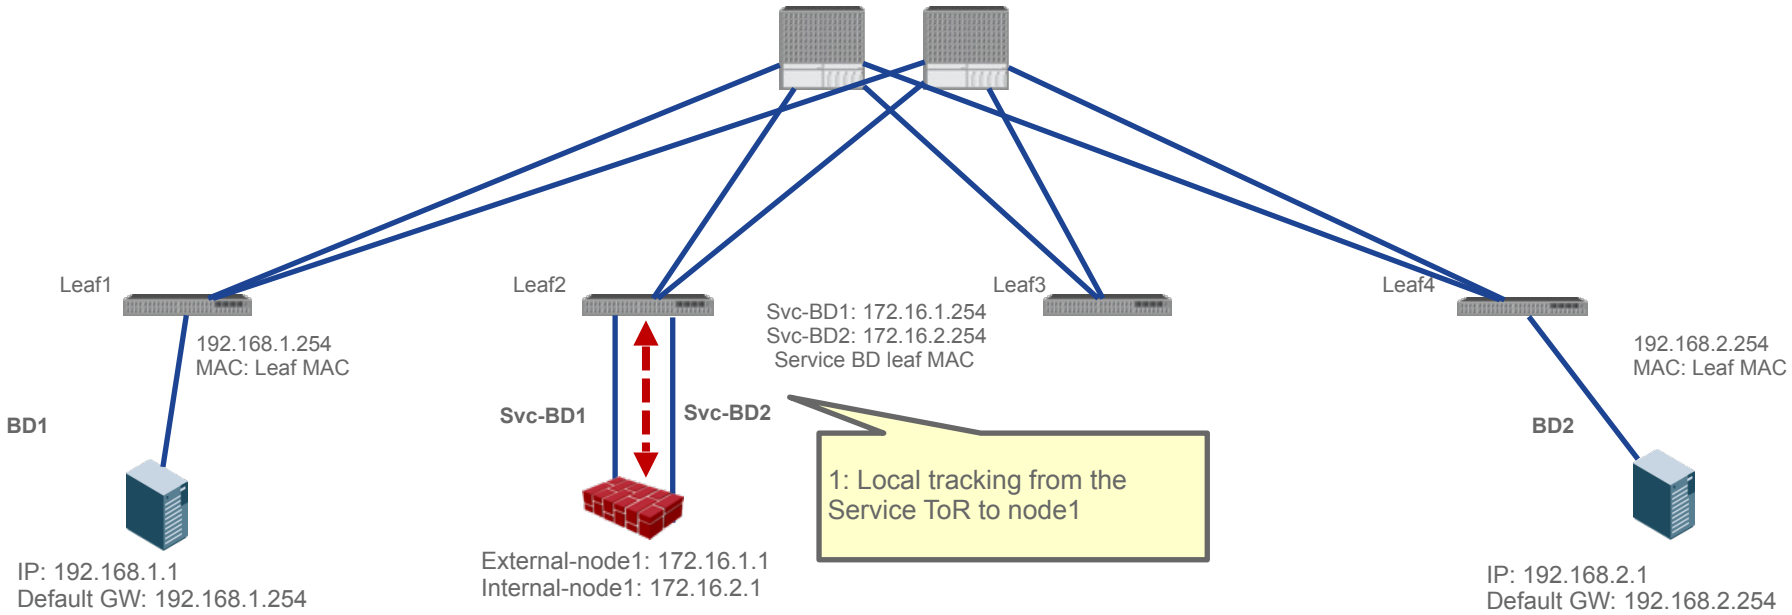

PBR with Tracking

Tracking behavior

ACI 2.2(3) 3.1

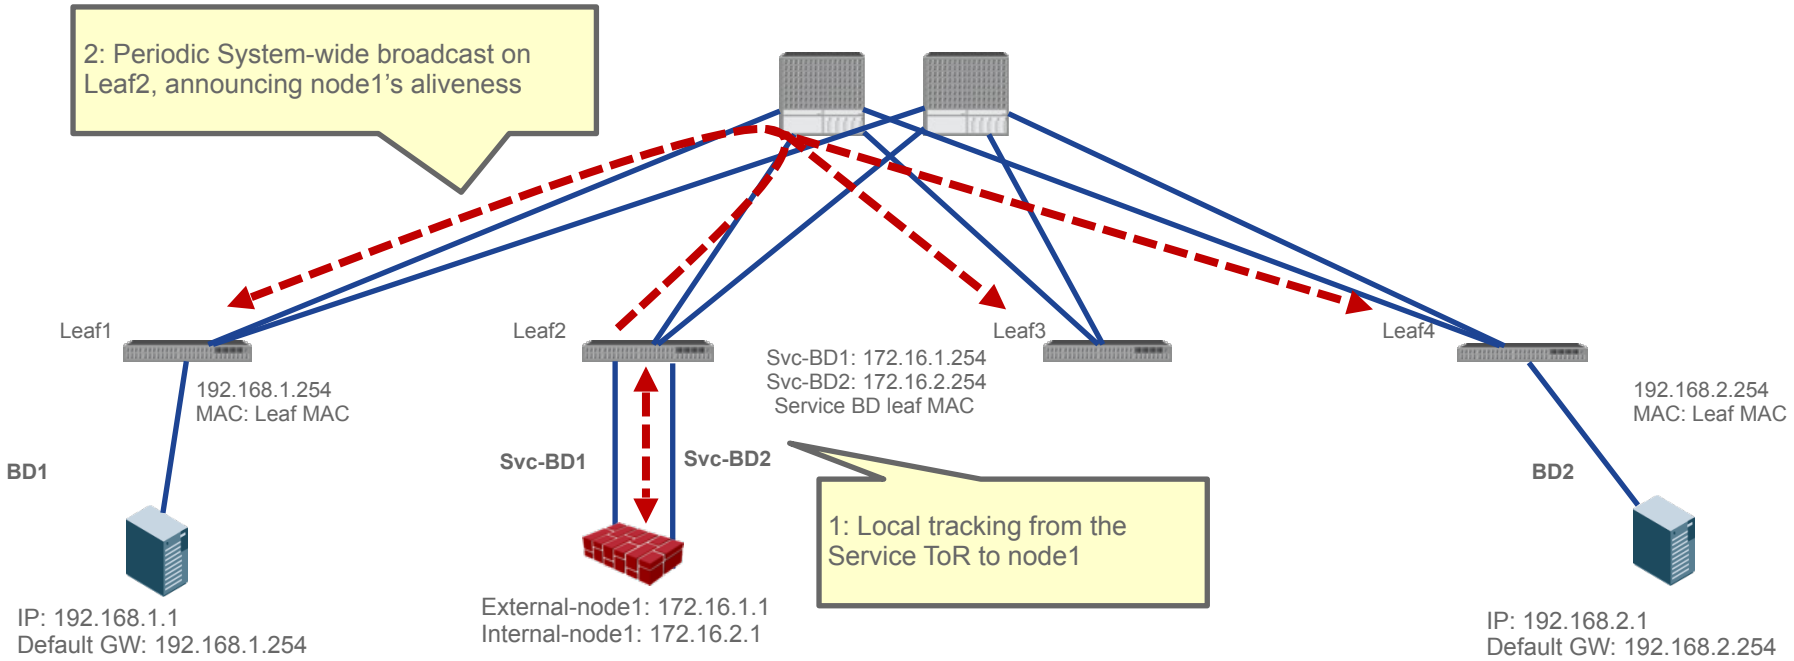

PBR with tracking

Health group

ACI 2.2(3) 3.1

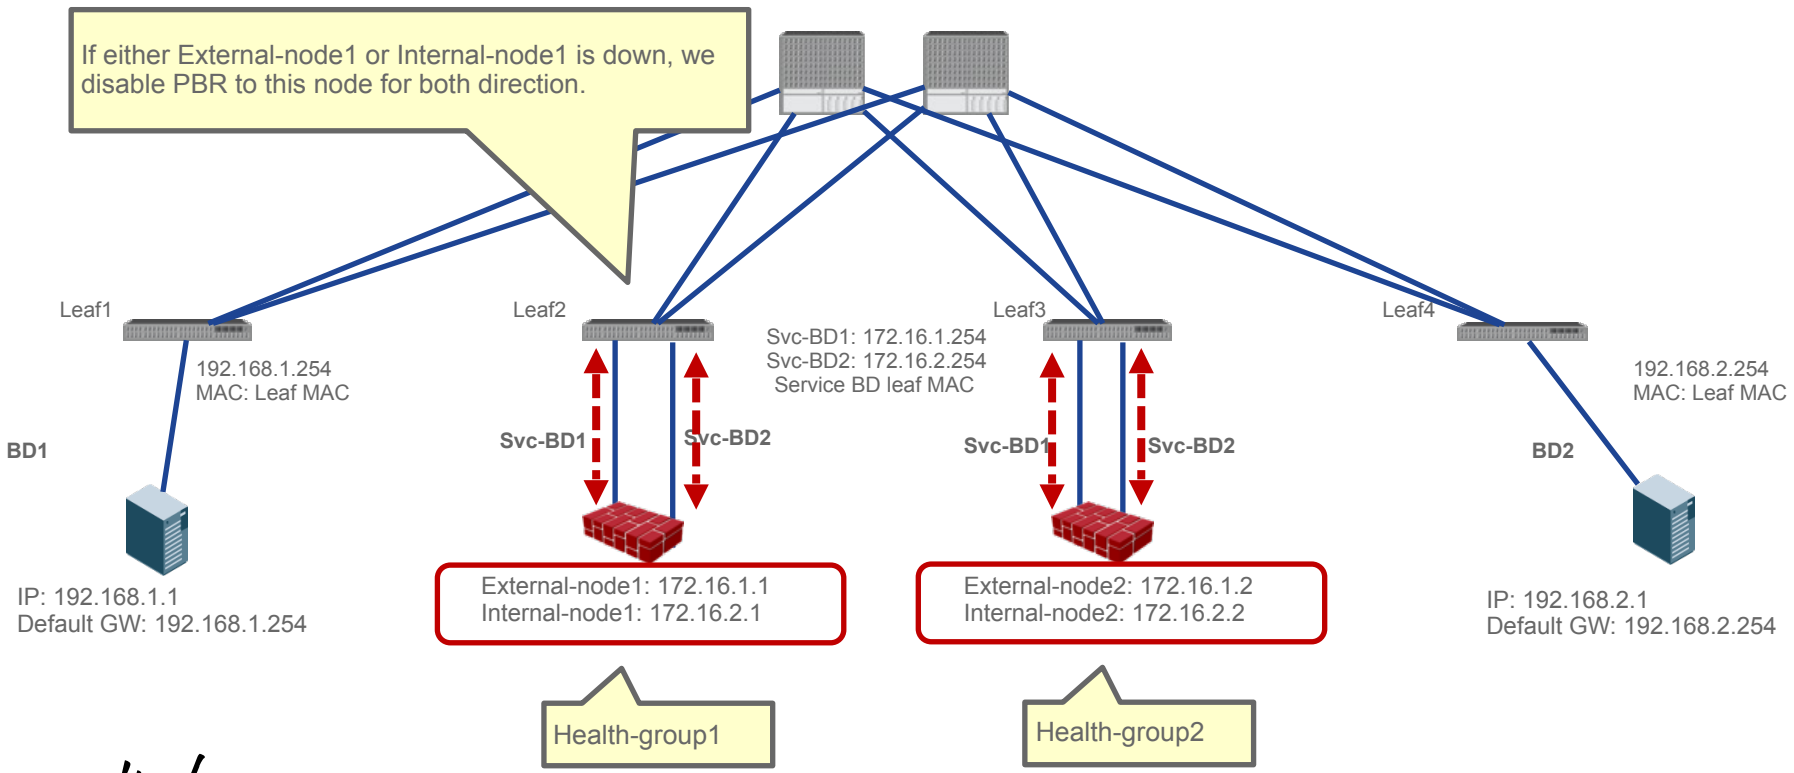

PBR with tracking

Threshold

ACI 2.2(3)
3.1

PBR is enabled
Available node =100%

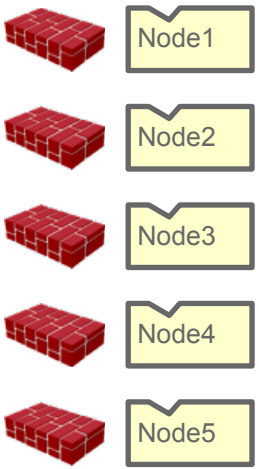

PBR with tracking

Threshold

ACI 2.2(3)
3.1

PBR is enabled
Available node =100%

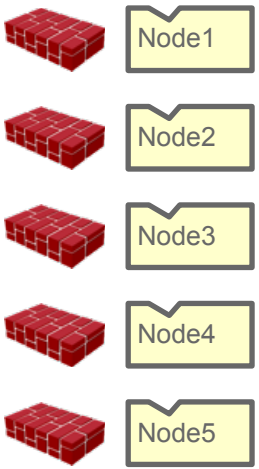

PBR is disabled
Available node =20%
(<= 20%)

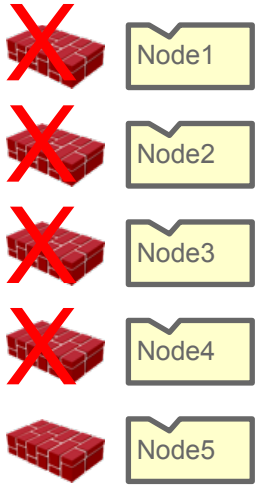

PBR with tracking

Threshold

ACI 2.2(3) 3.1

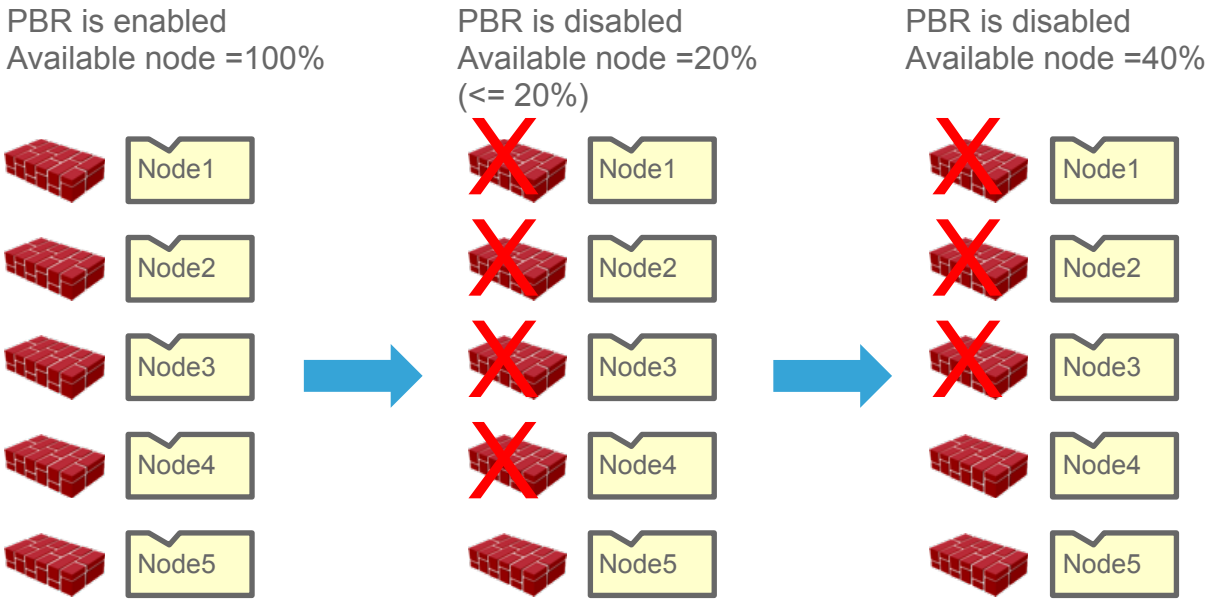

PBR with tracking

Threshold

ACI 2.2(3) 3.1

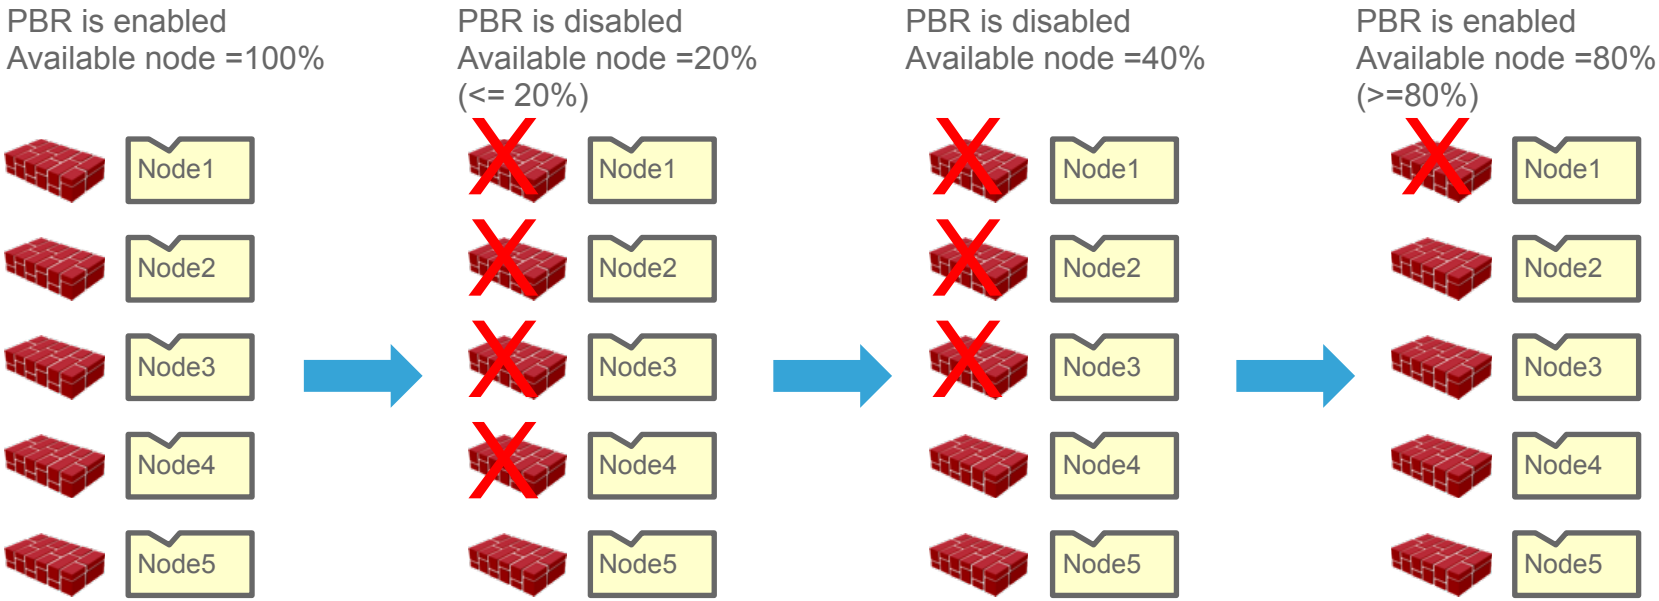

PBR with tracking

Down Action: Permit

ACI 2.2(3) 3.1

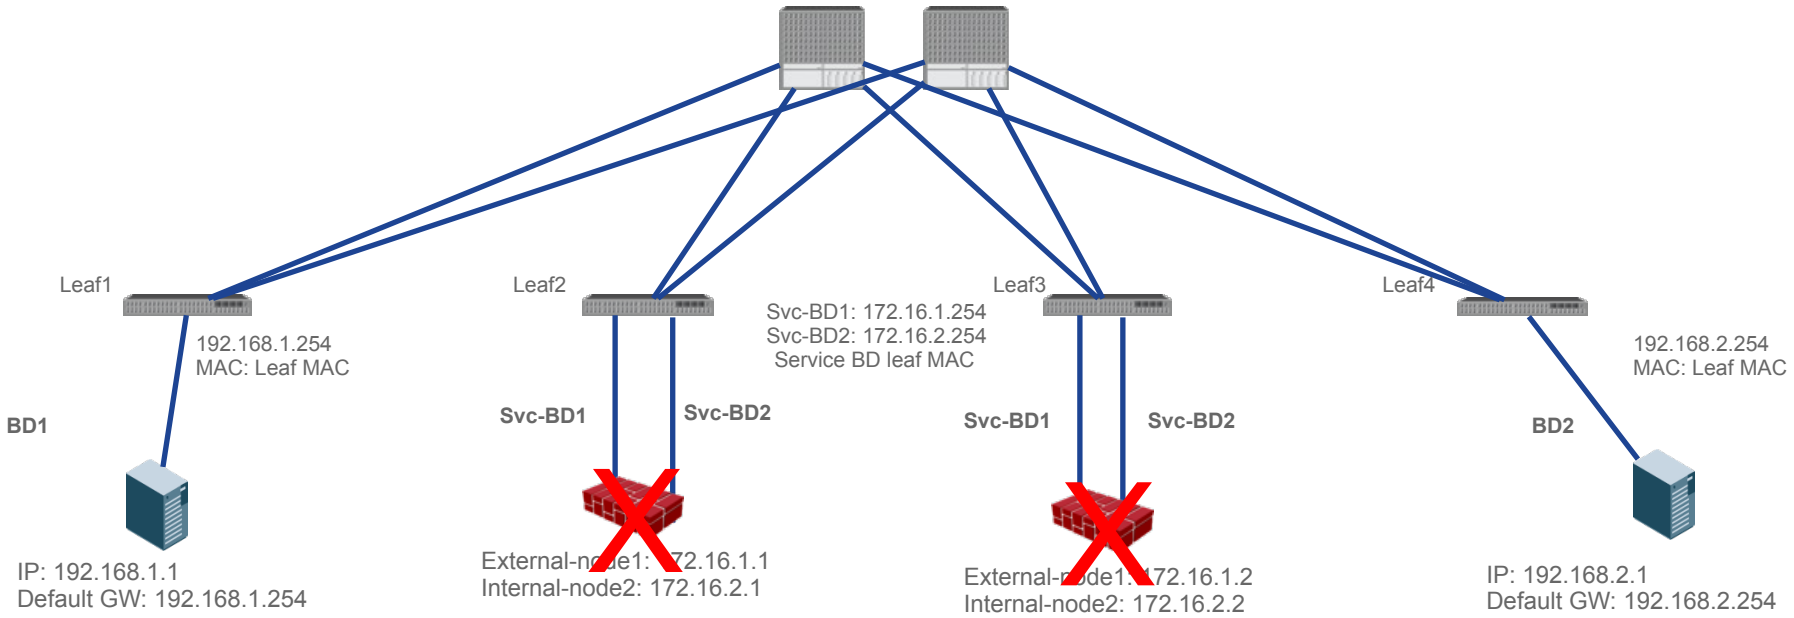

PBR with tracking

Down Action: Permit

ACI 2.2(3) 3.1

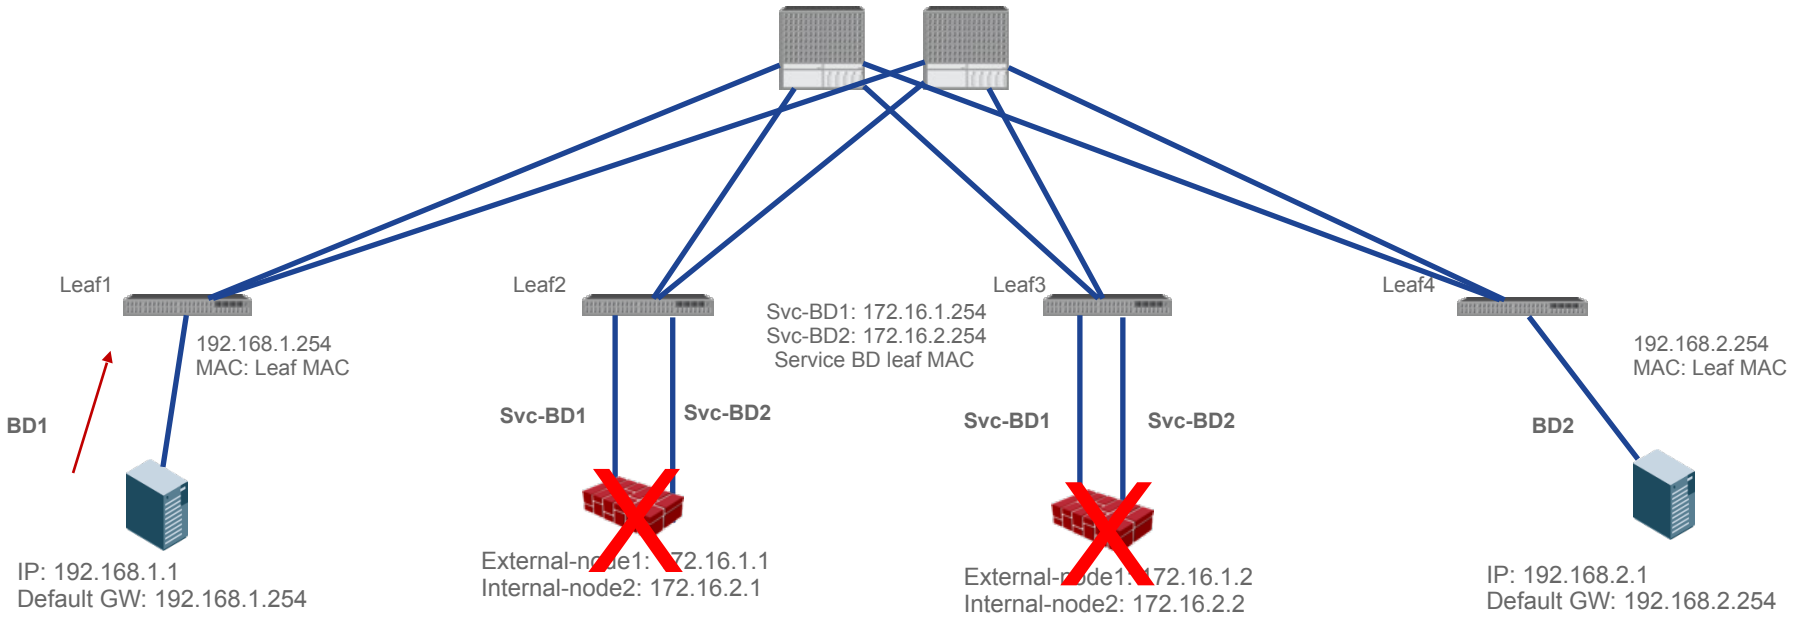

PBR with tracking

Down Action: Permit

ACI 2.2(3) 3.1

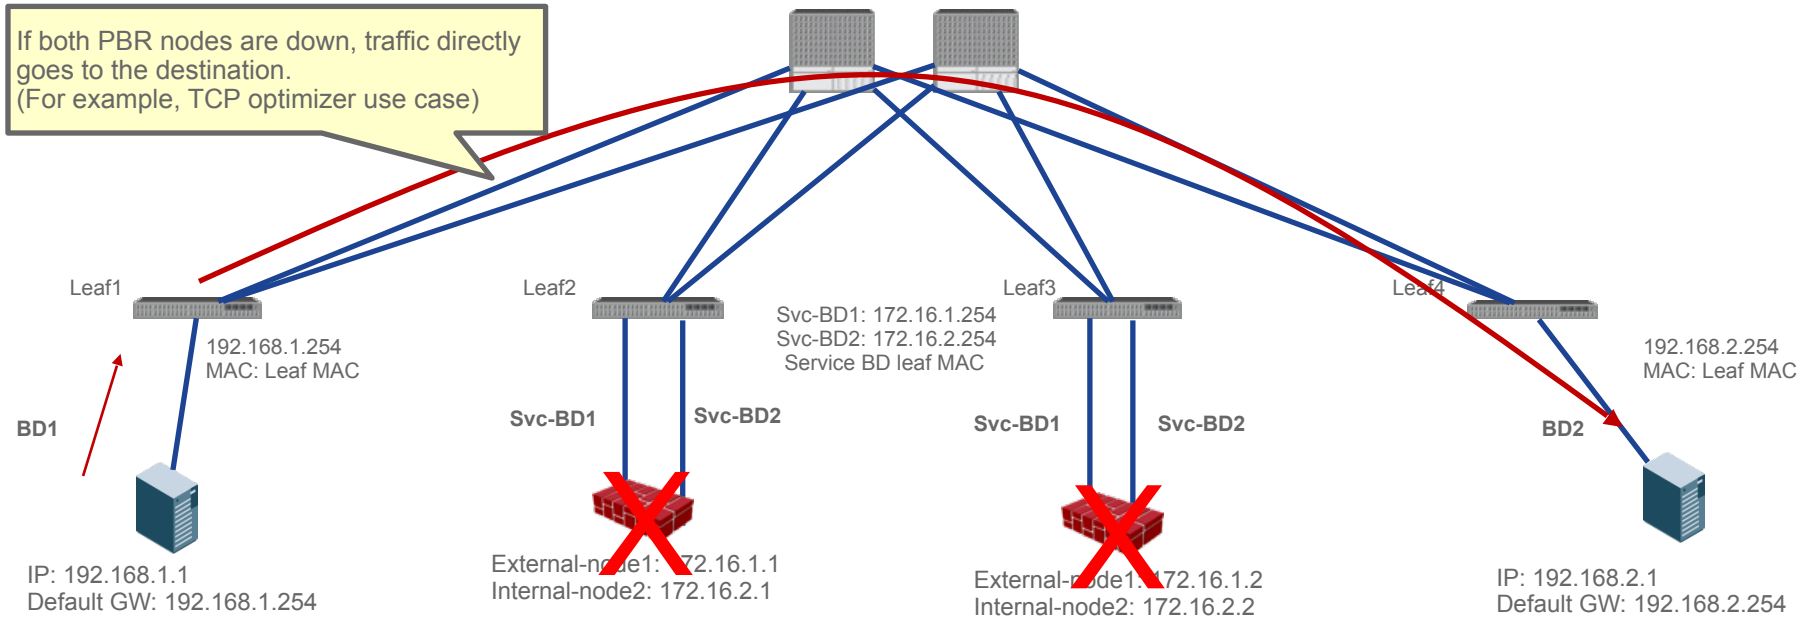

PBR with tracking

Down Action: Deny

ACI 2.2(3) 3.1

If all PBR node is down, traffic is dropped.
(For example, FW/IPS or security device use case)

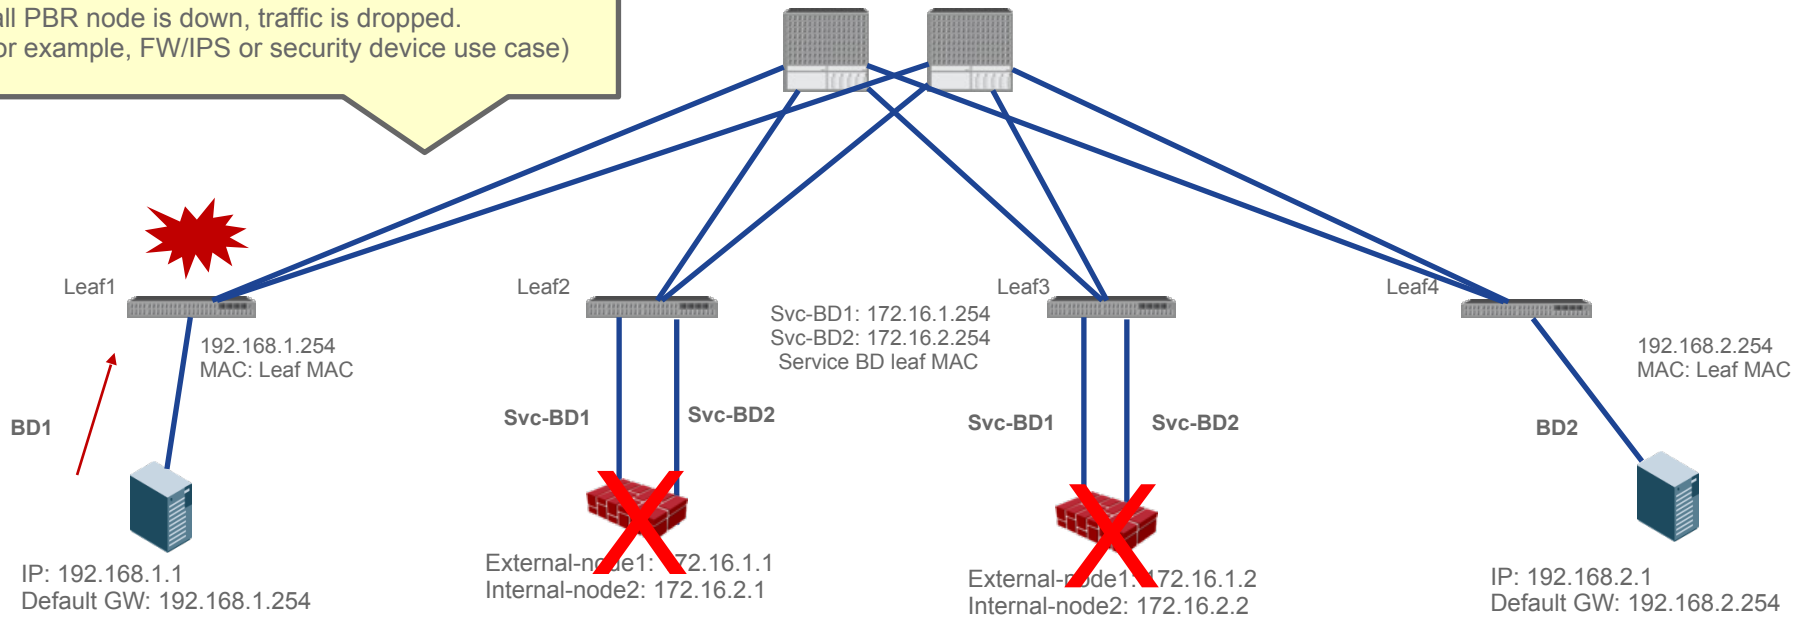

Multi-pod Design Considerations

ACI Multi-Pod
Network Services Integration Models

For Your Reference

ACI Multi-Pod
Network Services Integration Models

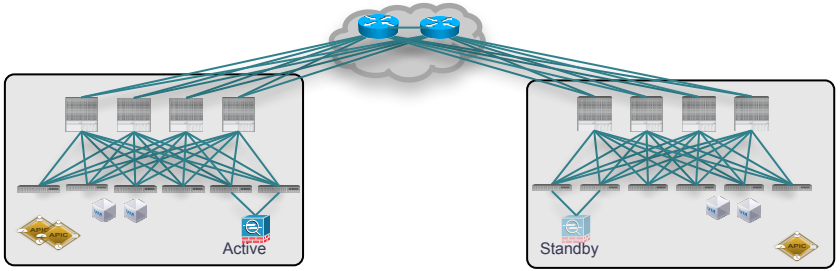

- Active and Standby pair deployed across Pods
- No issues with asymmetric flows but causes traffic hair-pinning across the IPN

ACI Multi-Pod
Network Services Integration Models

For Your Reference

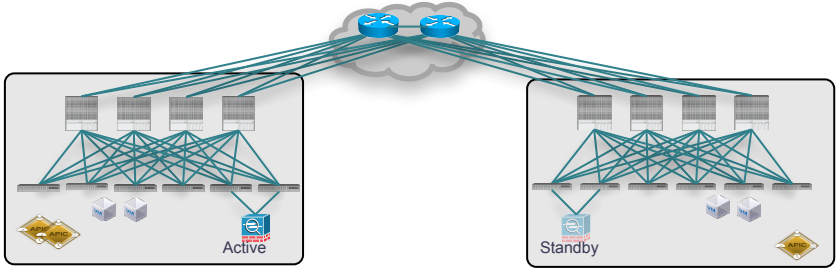

- Active and Standby pair deployed across Pods
- No issues with asymmetric flows but causes traffic hair-pinning across the IPN

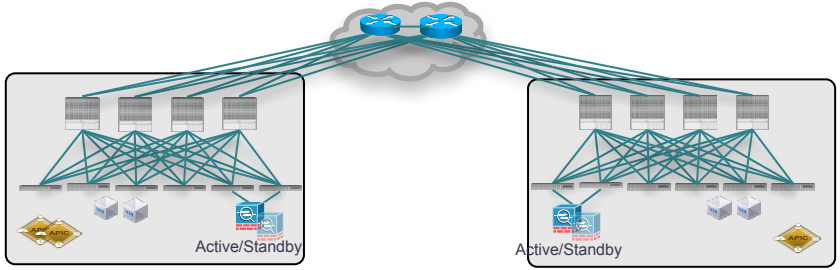

- Independent Active/Standby pair deployed in each Pod
- Only for perimeter FW use case assuming proper solution is adopted to keep symmetric ingress/ egress traffic flows

ACI Multi-Pod
Network Services Integration Models

For Your Reference

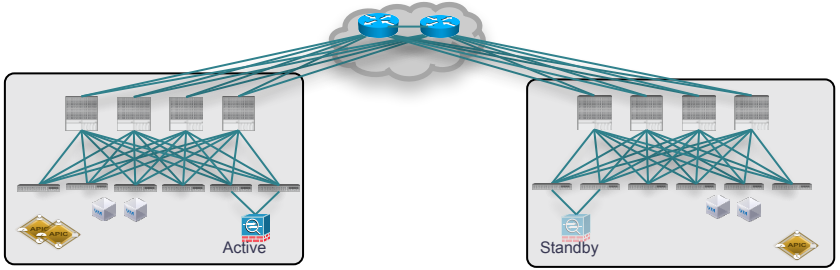

- Active and Standby pair deployed across Pods
- No issues with asymmetric flows but causes traffic hair-pinning across the IPN

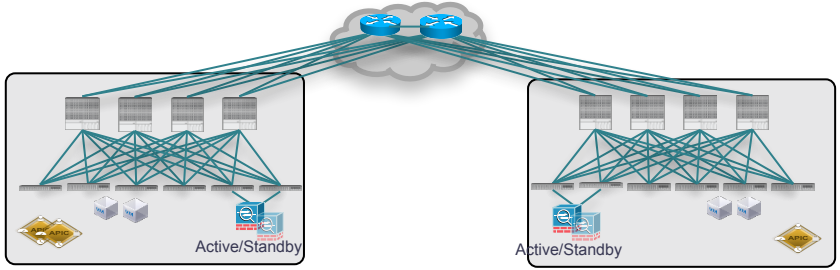

- Independent Active/Standby pair deployed in each Pod
- Only for perimeter FW use case assuming proper solution is adopted to keep symmetric ingress/ egress traffic flows

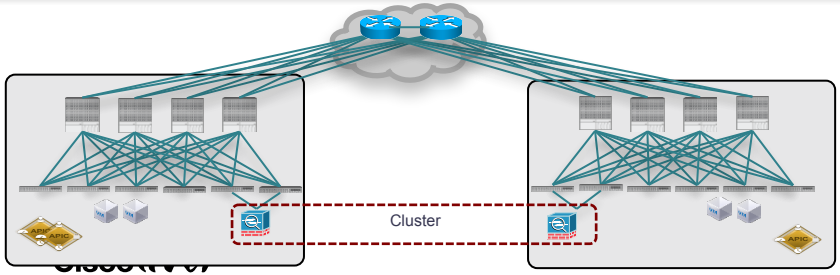

- FW cluster deployed across Pods
- Not currently supported (scoped for 1HCY18)

Active/Standby FW Pair across Pods

Multi-Pod and Network Services

Active/Standby Pair across Pods

- Easiest deployment models, no risk of creating asymmetric traffic paths across stateful services
- Support service nodes in both Layer 2 and Layer 3 modes
- Leverage Multi-Pod capabilities of seamlessly providing L2/L3 connectivity across Pods
- Applies to both east-west and north-south communication
- Supports both traditional Border Leaf L3Outs and GOLF L3Outs (for north-south communication)
- Future migration path to an Active/Active cluster FW solution (Q2CY18)

Option 1
FW in Layer 2 Mode

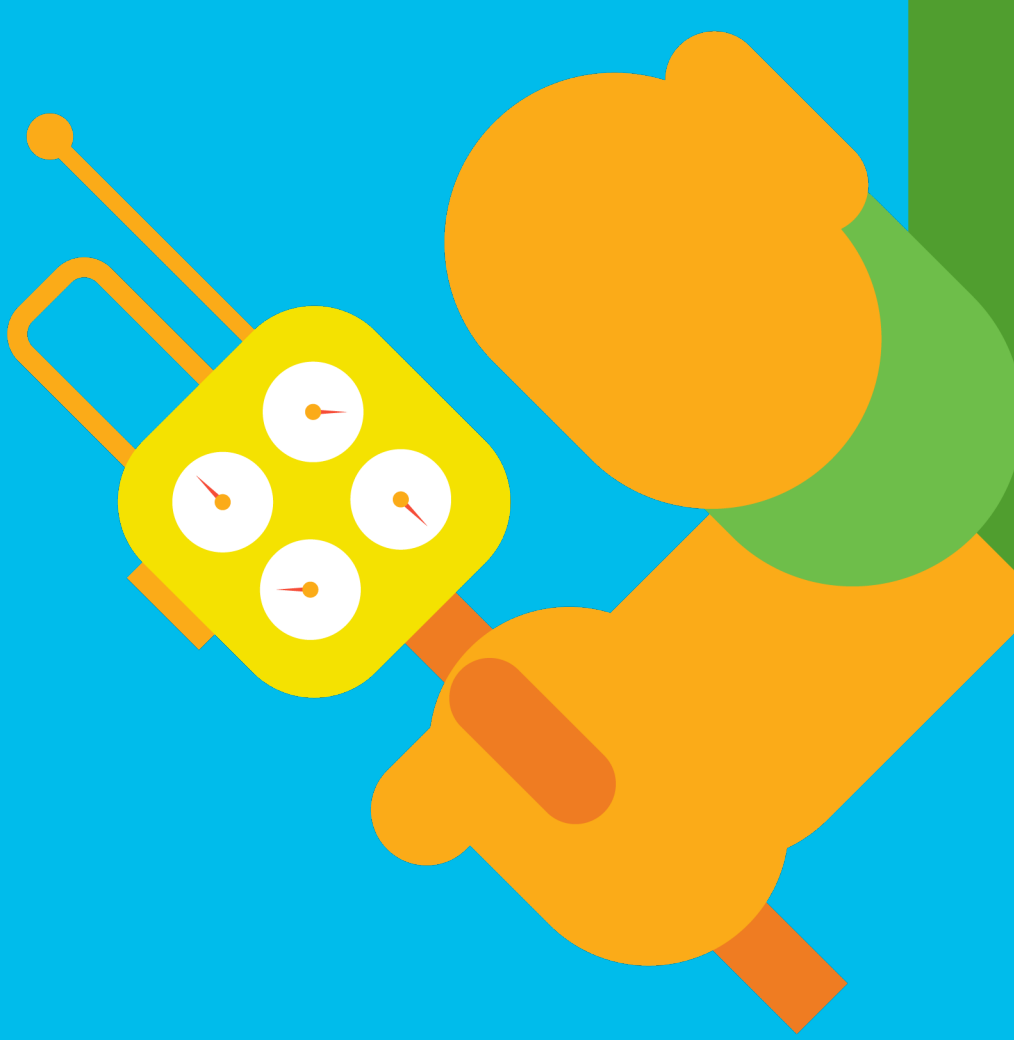

Active/Standby Pair across Pods

Option 1: FW in L2 Mode

North-South

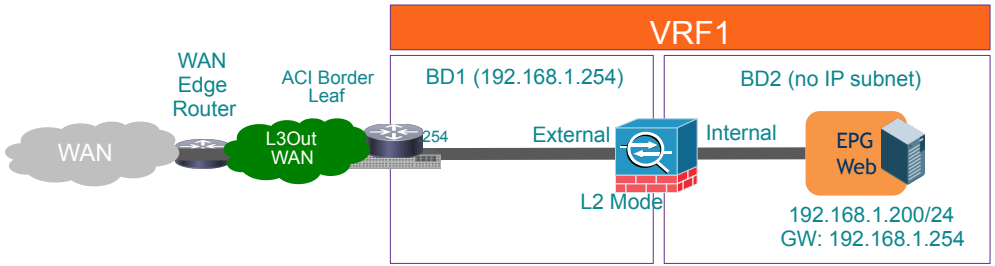

East-West

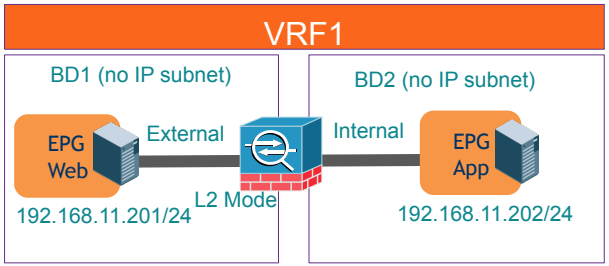

- All traffic leaving the EPG is forced through the FW via L2 lookup (i.e. 'BD stitching')
- Support for all deployment models (service-graph managed and un-managed, no service-graph)

Active/Standby Pair across Pods

Option 1: FW in L2 Mode

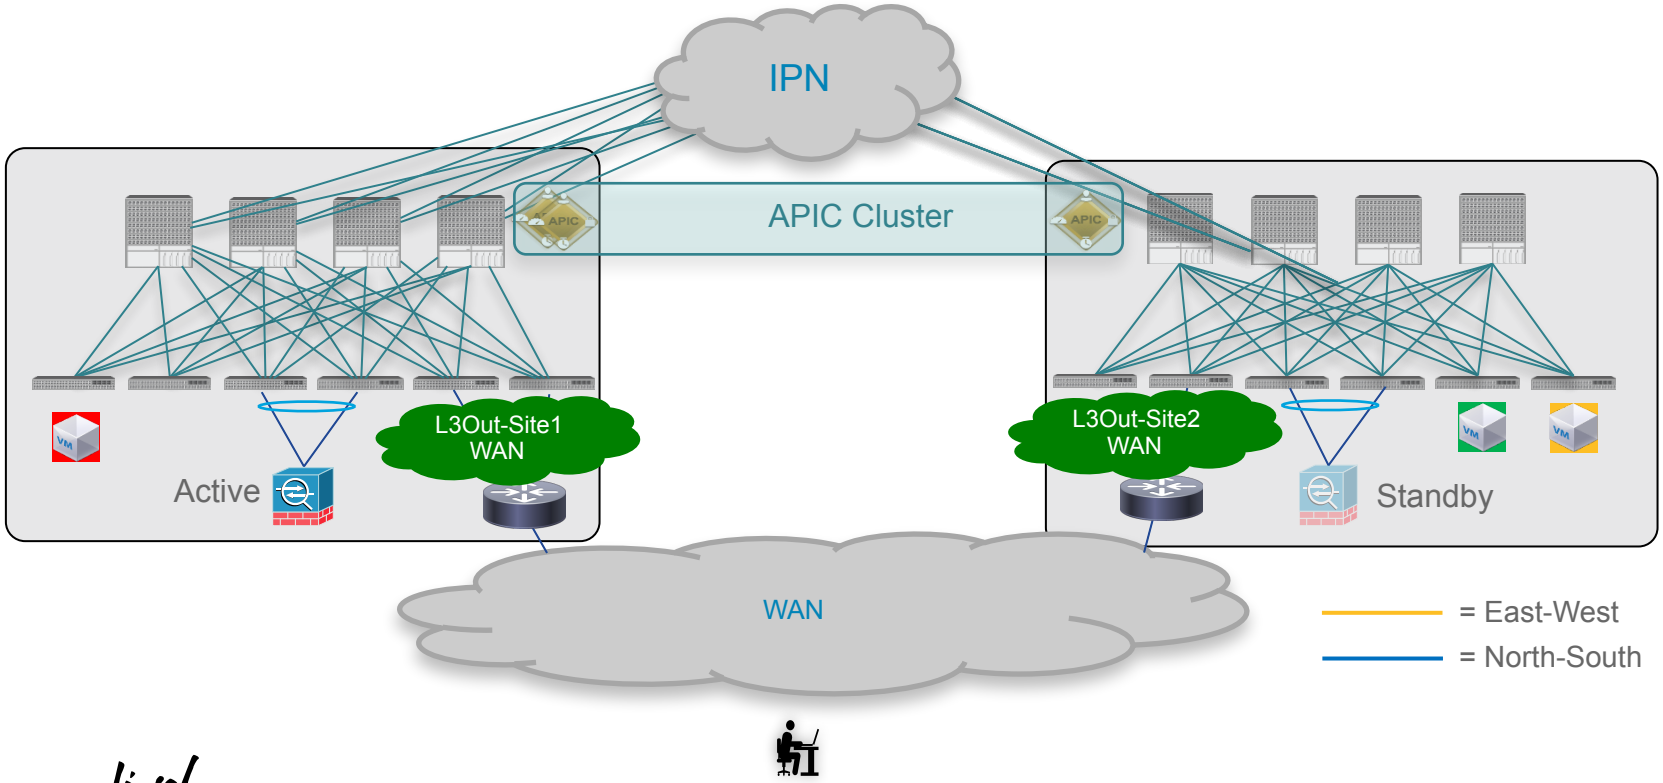

Active/Standby Pair across Pods

Option 1: FW in L2 Mode

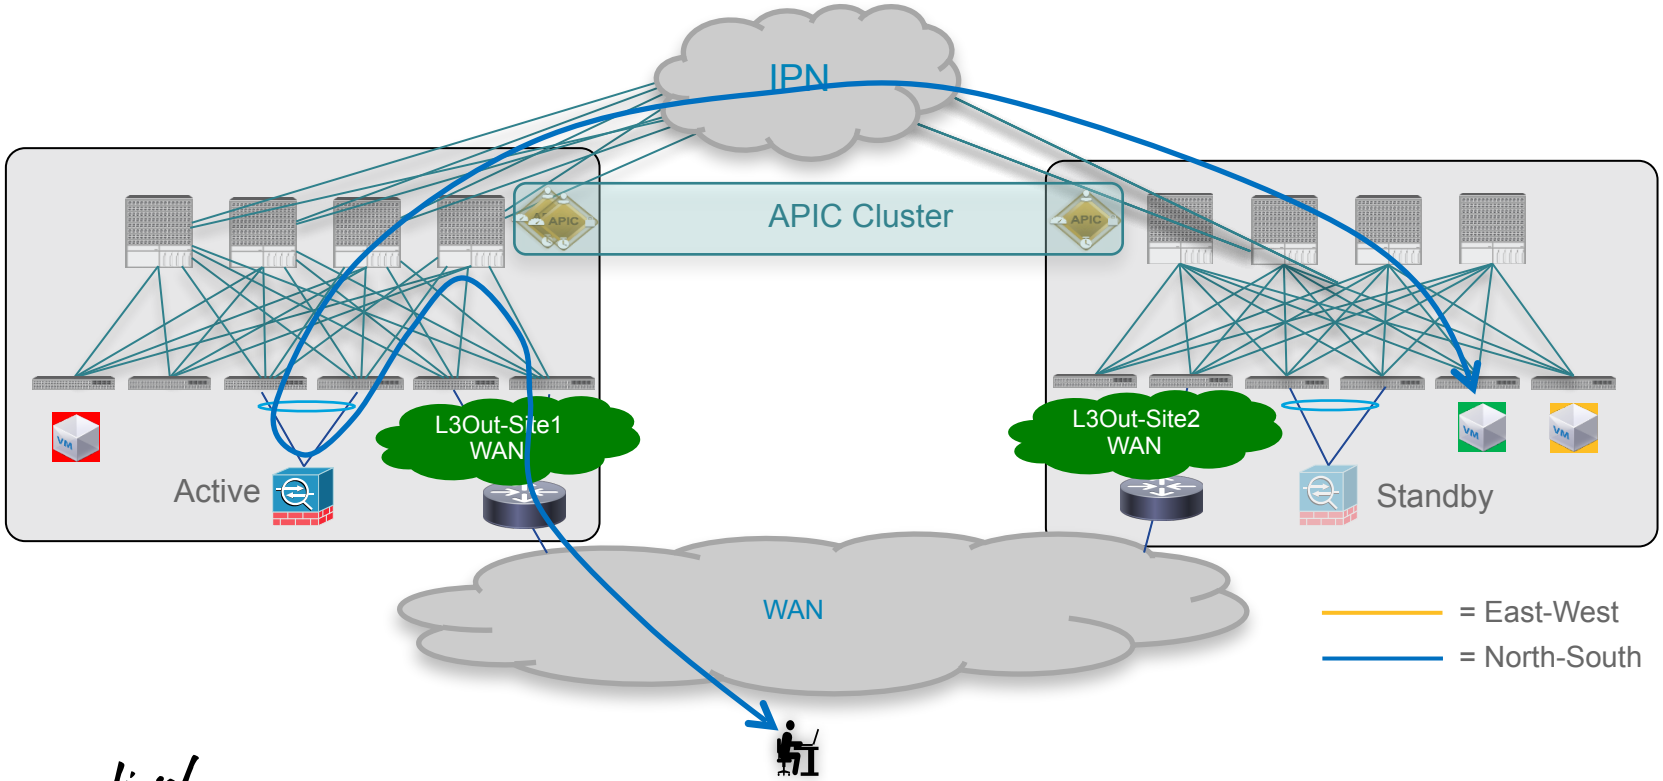

Active/Standby Pair across Pods

Option 1: FW in L2 Mode

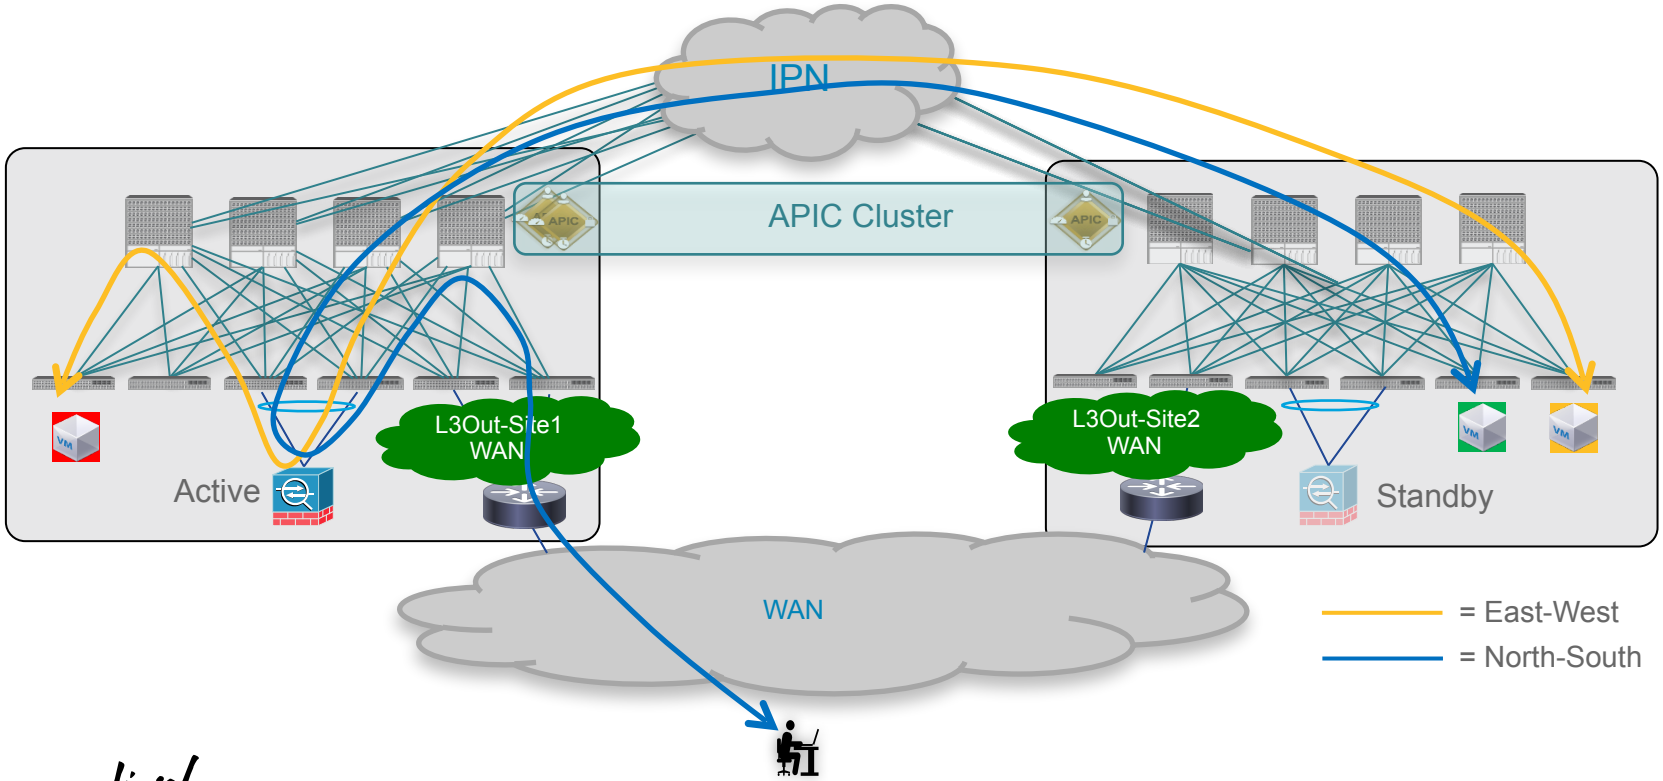

Option 2
FW in Layer 3 Mode and PBR

Active/Standby Pair across Pods

Option 2: FW in L3 Mode and PBR

North-South

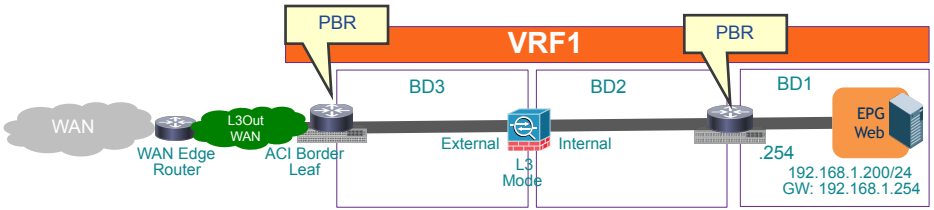

East-West

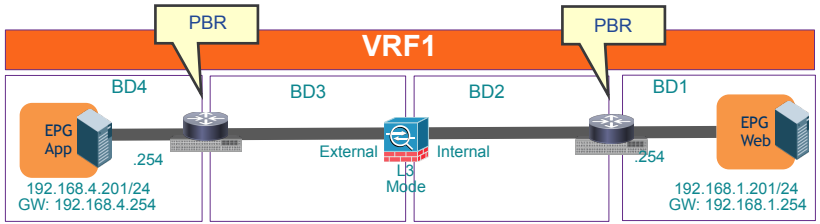

- The default gateway for the endpoints is always deployed on the ACI leaf nodes (Anycast gateway)
- Traffic is forced through the Active FW via PBR policy
- Support for all deployment models (service-graph managed and un-managed, no service-graph)

Active/Standby Pair across Pods

Option 2: FW in L3 Mode and PBR

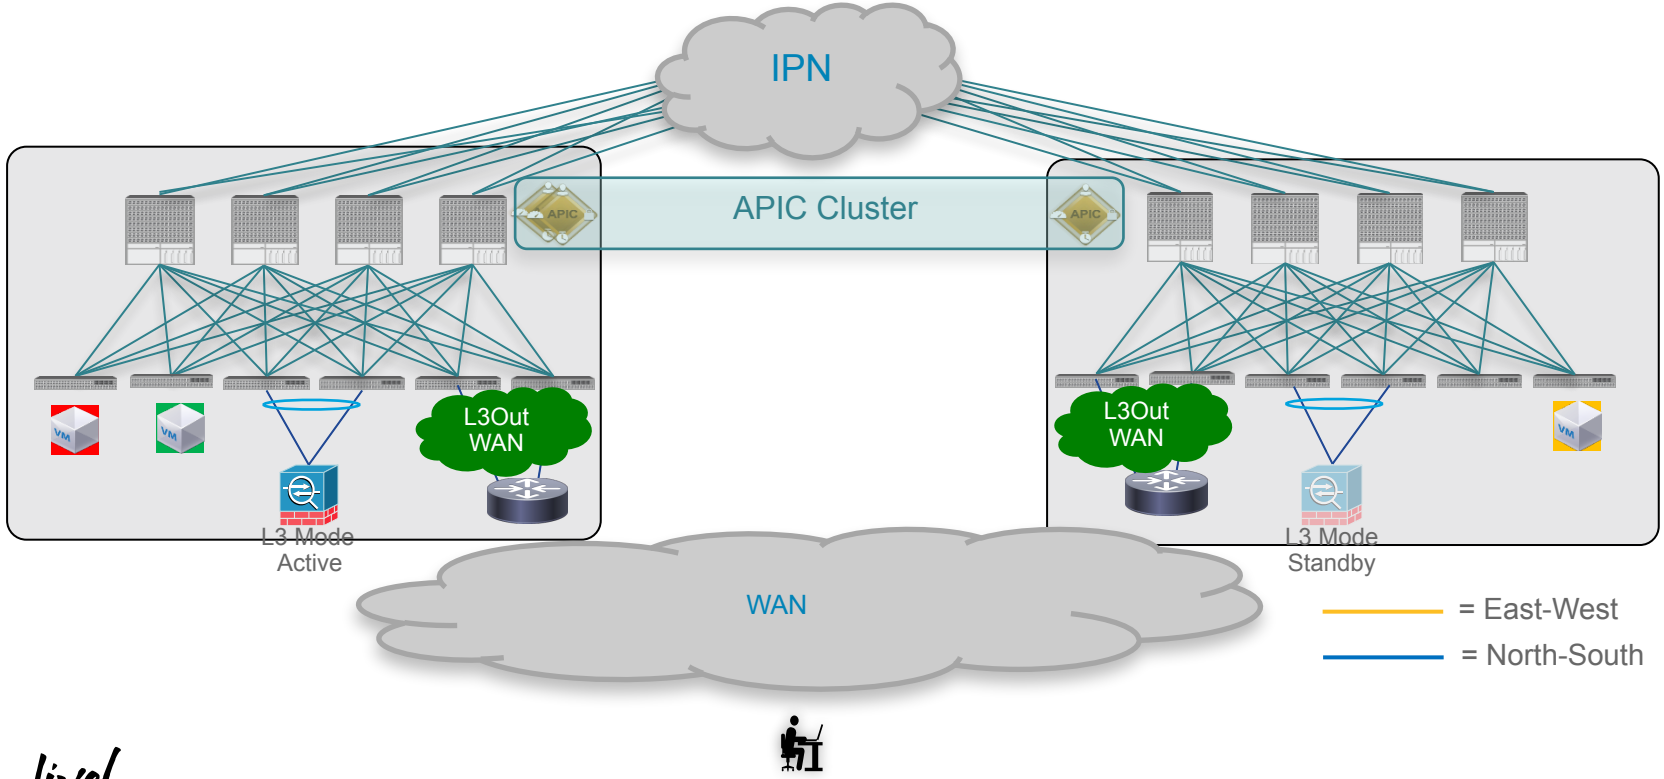

Active/Standby Pair across Pods

Option 2: FW in L3 Mode and PBR

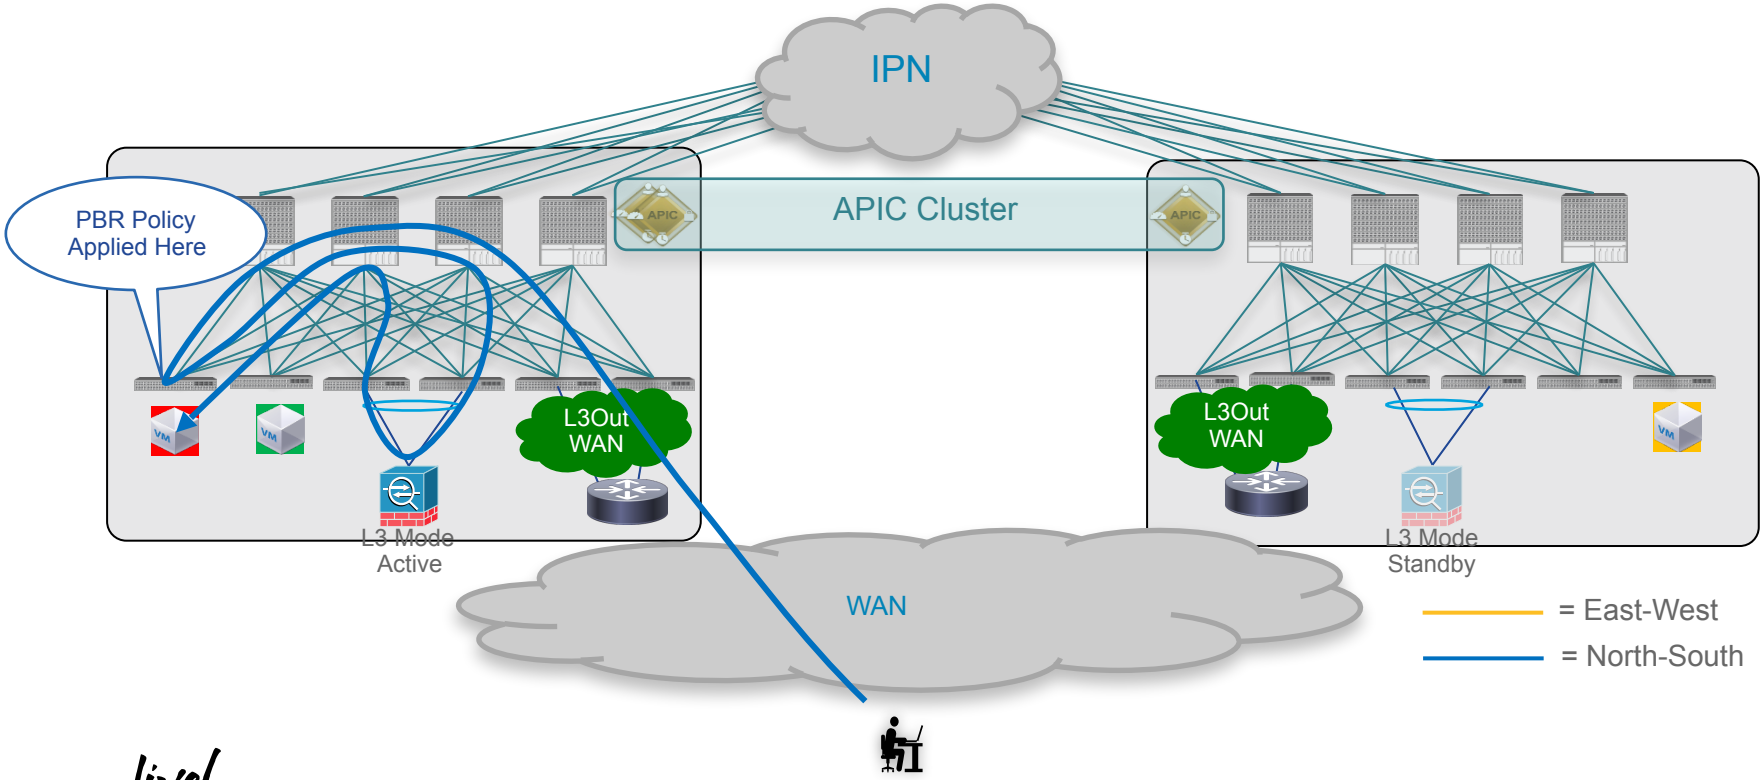

Active/Standby Pair across Pods

Option 2: FW in L3 Mode and PBR

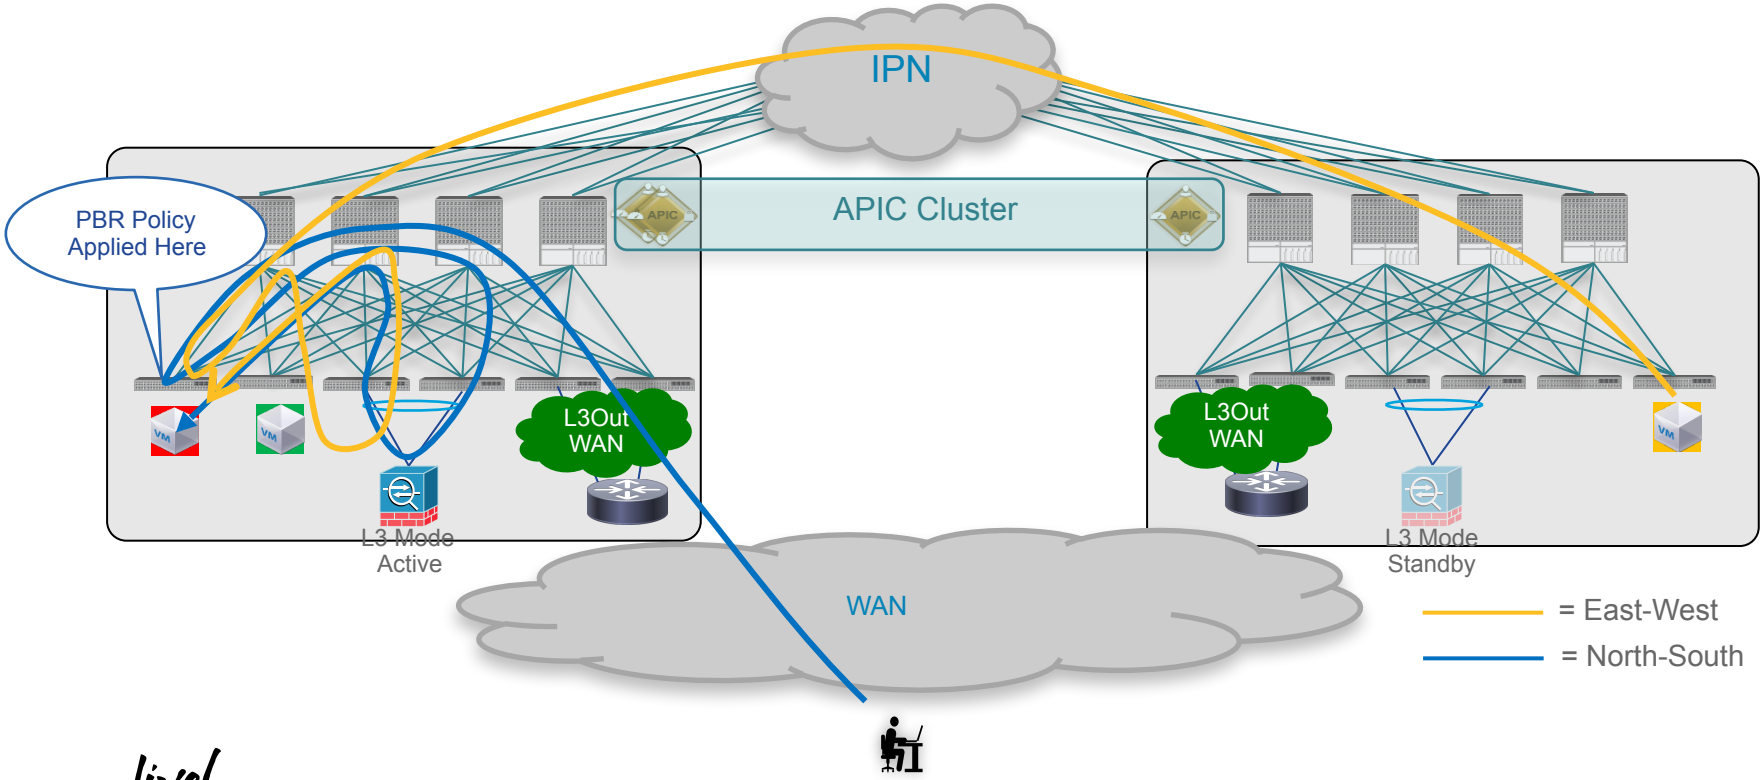

Option 3
FW in Layer 3 Mode and L3Outs

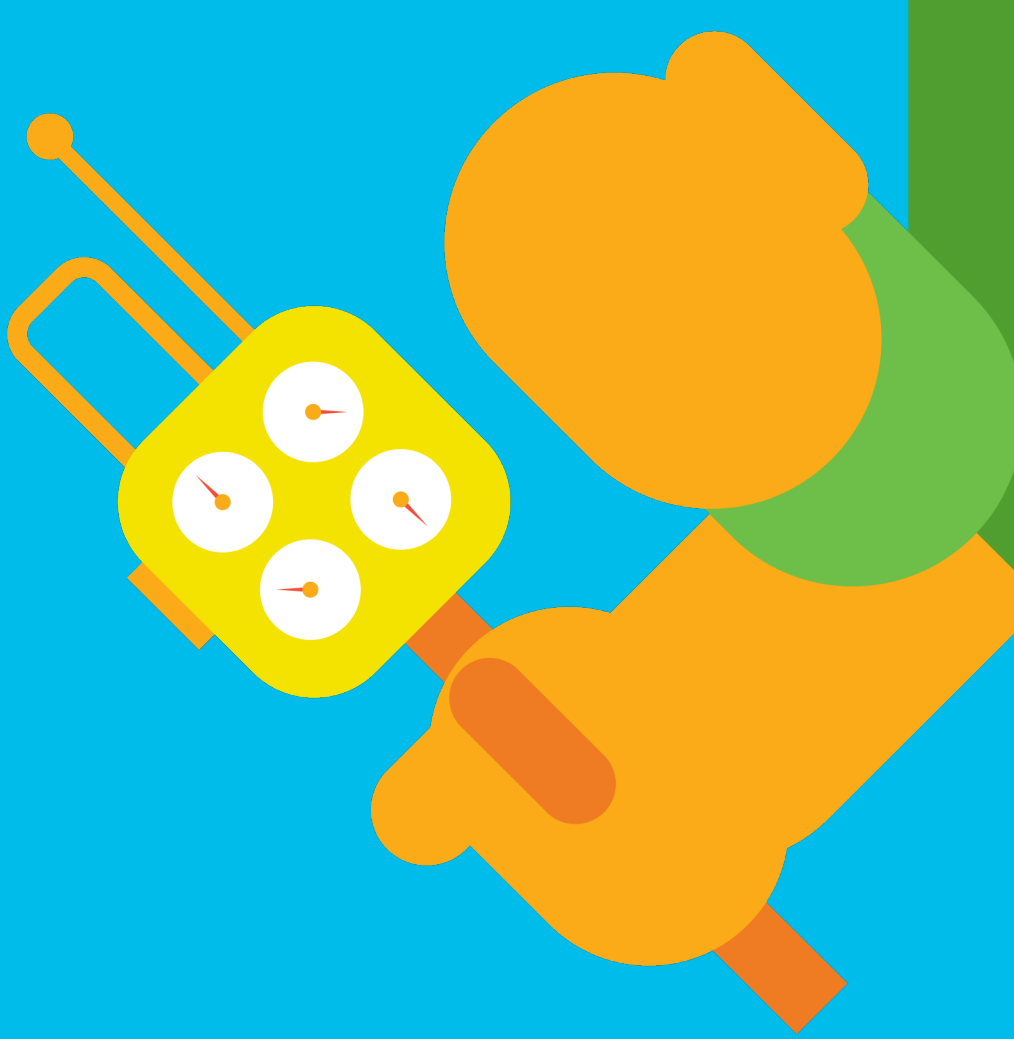

Active/Standby Pair across Pods

Option 3: FW in L3 Mode and L3Outs

North-South

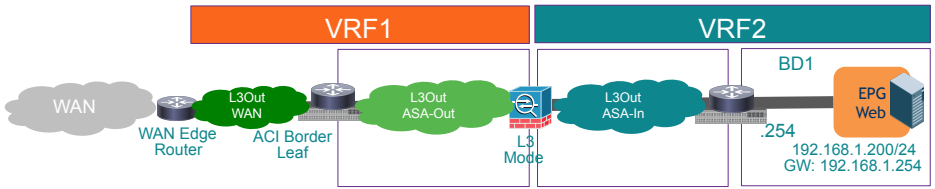

East-West

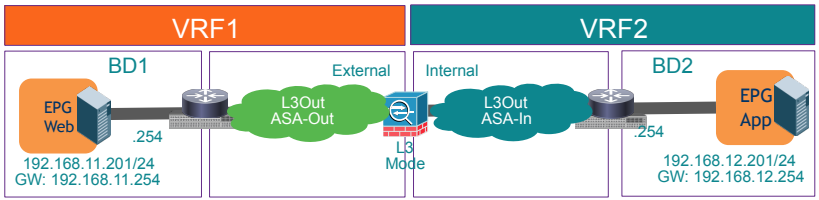

- The default gateway for the endpoints is always deployed on the ACI leaf nodes (Anycast gateway)
- Traffic is forced through the Active FW via L3Out connections (VRF-sandwich) leveraging static or dynamic routing
- For the N-S scenario, alternative option is to connect the external FW interface directly to the WAN Edge router
- Support for all deployment models (service-graph managed and un-managed, no service-graph)

Cisco live!

Active/Standby Pair across Pods

Option 3: FW in L3 Mode and L3Outs

Single L3Out Defined across Pods

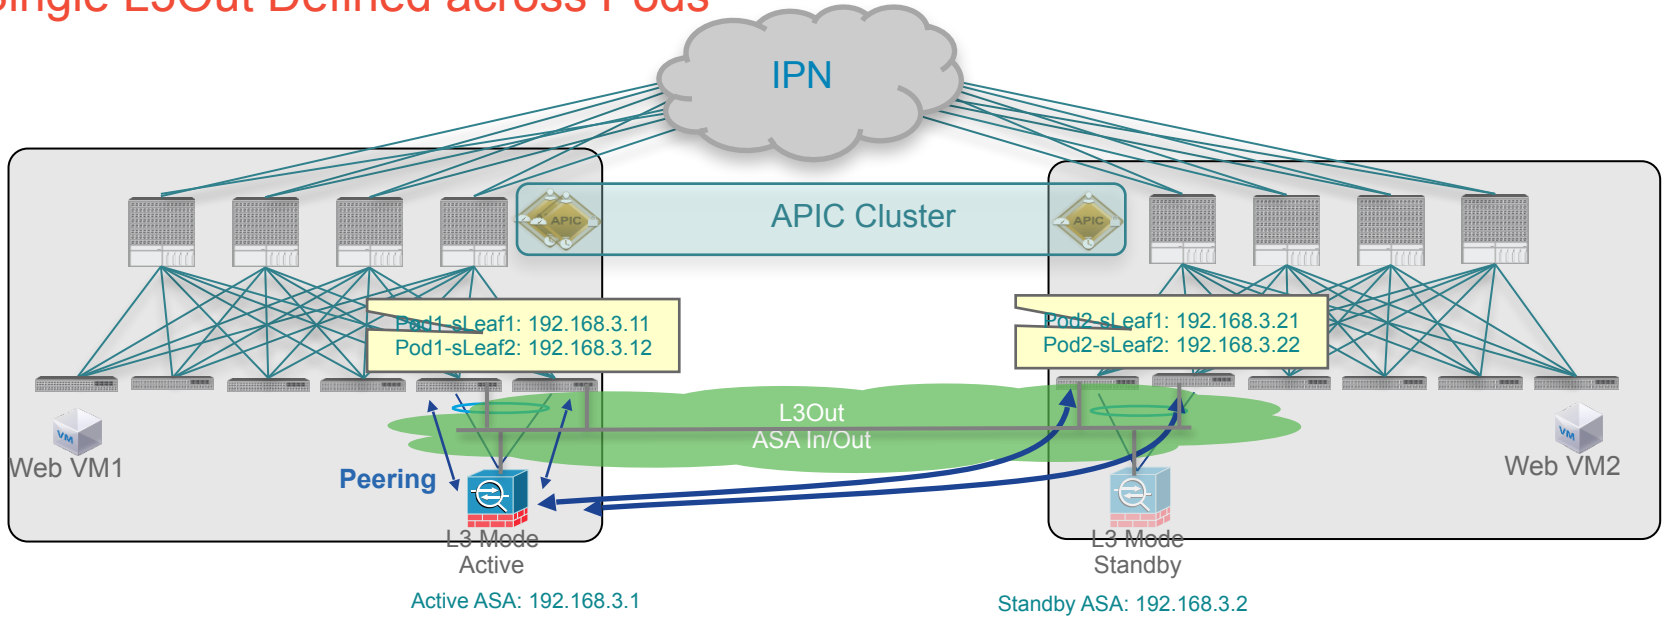

Note: supported from ACI SW releases 2.1(3), 2.2(3), 2.3(1) and 3.0(1) and deploying EX/FX HW for ACI service leaf nodes

Active/Standby Pair across Pods

Option 3: FW in L3 Mode and L3Outs

Single L3Out Defined across Pods

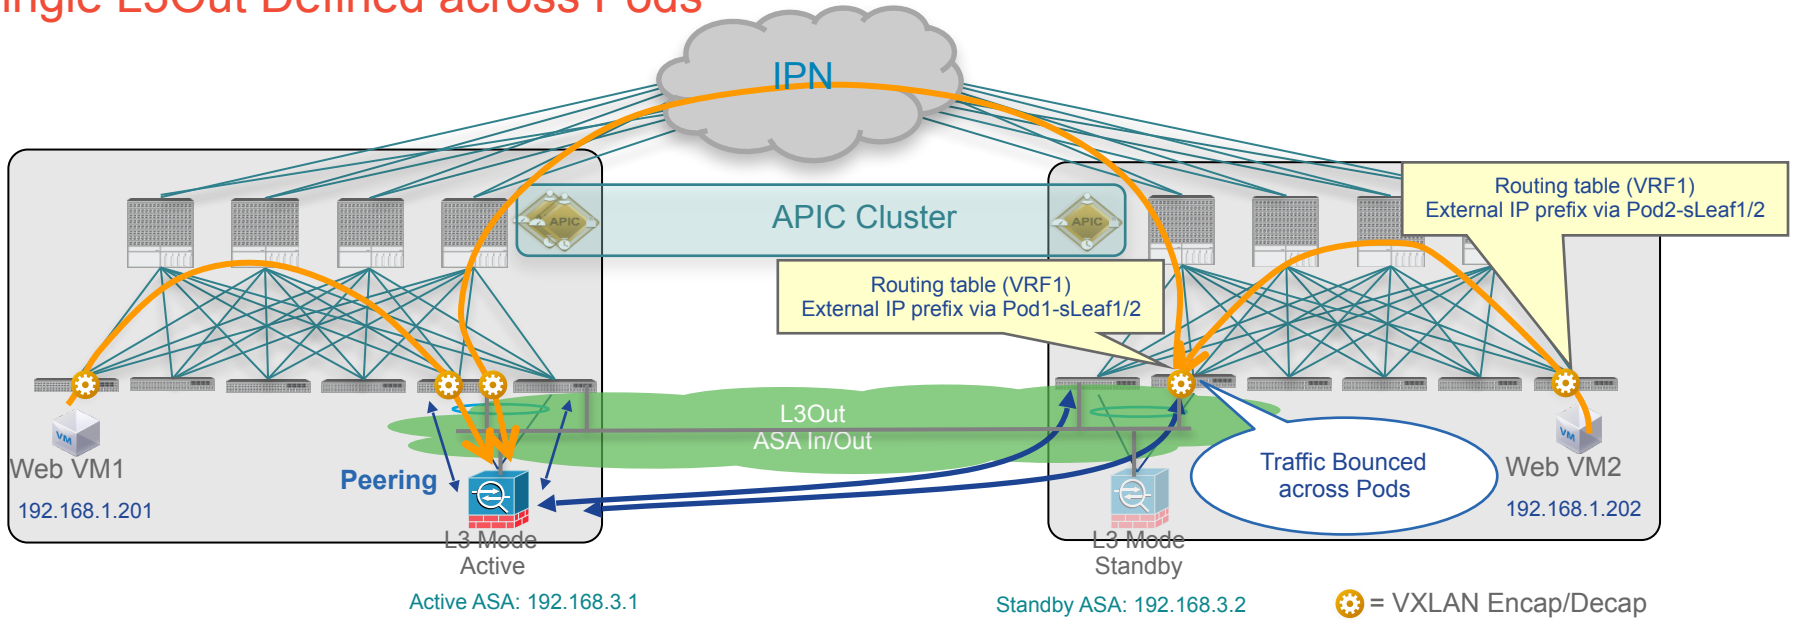

Note: supported from ACI SW releases 2.1(3), 2.2(3), 2.3(1) and 3.0(1) and deploying EX/FX HW for ACI service leaf nodes

Active/Standby Independent FW Pairs in each Pod

Multi-Pod and Network Services

Active/Standby Independent Pairs in each Pod

- Only supported with network services in Layer 3 mode
- Must pay specific attention to avoid the creation of asymmetric traffic path through independent FWs (i.e. not sharing connection state)
- Recommendation is to deploy Symmetric PBR to enforce both legs of the traffic flows through the same FW
- Symmetric PBR can be used for both North-South and East-West flows
- Solution improvements in 3.1 release with the introduction of location/Pod aware PBR policy (North-South flows only)

FW in Layer 3 Mode and Symmetric PBR

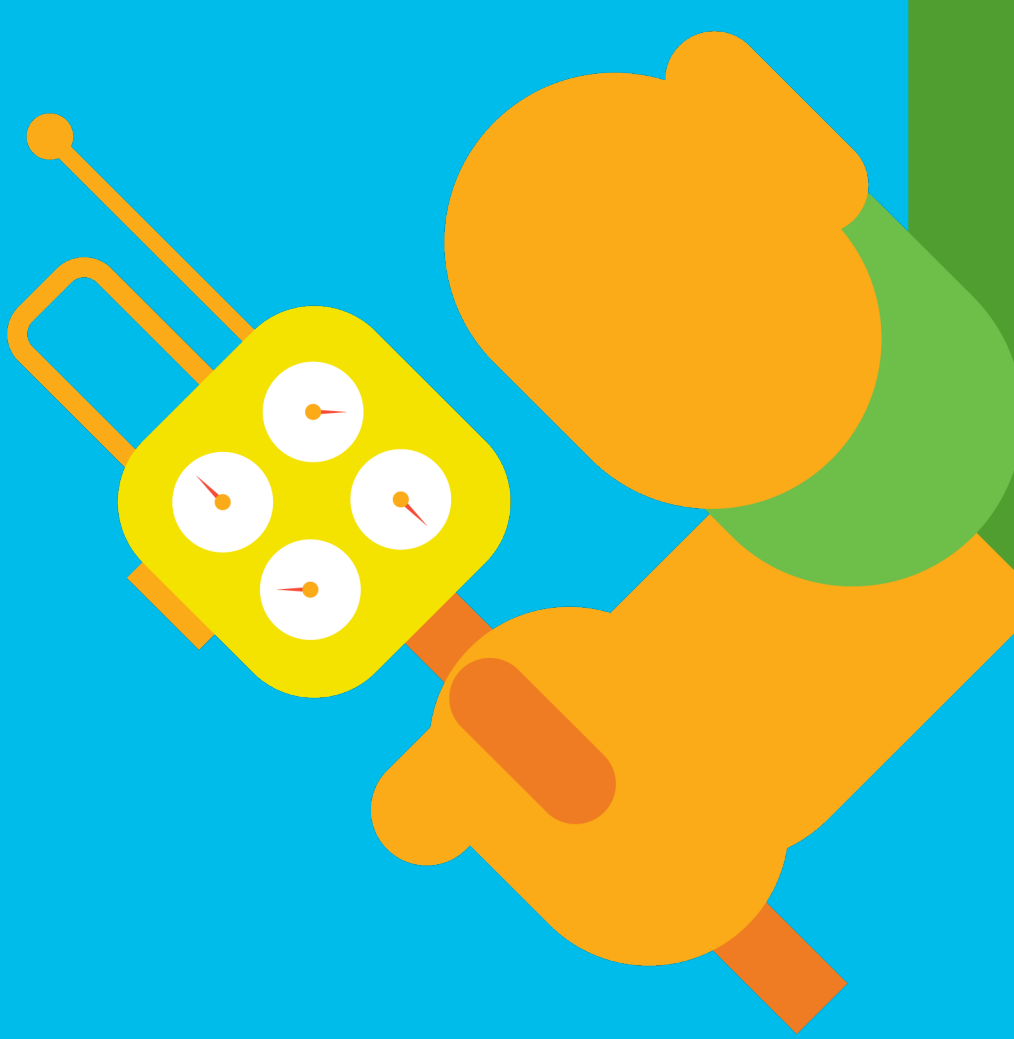

Active/Standby Independent Pairs in each Pod

FW in Layer 3 Mode and Symmetric PBR

North-South

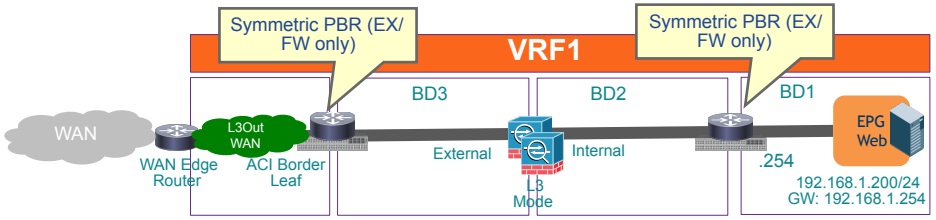

East-West

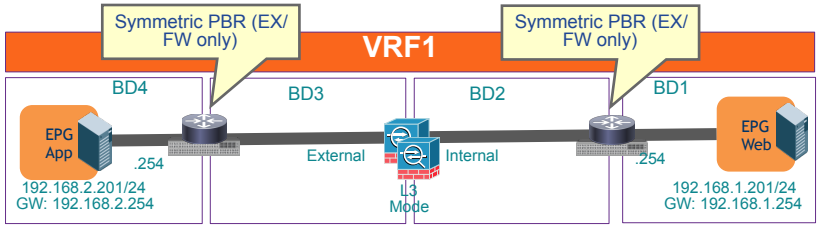

- The default gateway for the endpoints is always deployed on the ACI leaf nodes (Anycast gateway)
- Traffic is forced through the same Active FW via Symmetric PBR policy (supported only with EX/ FX ACI leaf HW)
  - Two IP/MAC nodes are specified in the same PBR policy
  - Traffic is redirected to one of the two nodes based on hashing (Source_IP, Dest_IP, Protocol Type)
- Supported only with service-graph in un-managed mode

Cisco live!

Active/Standby Independent Pairs in each Pod

FW in Layer 3 Mode and Symmetric PBR

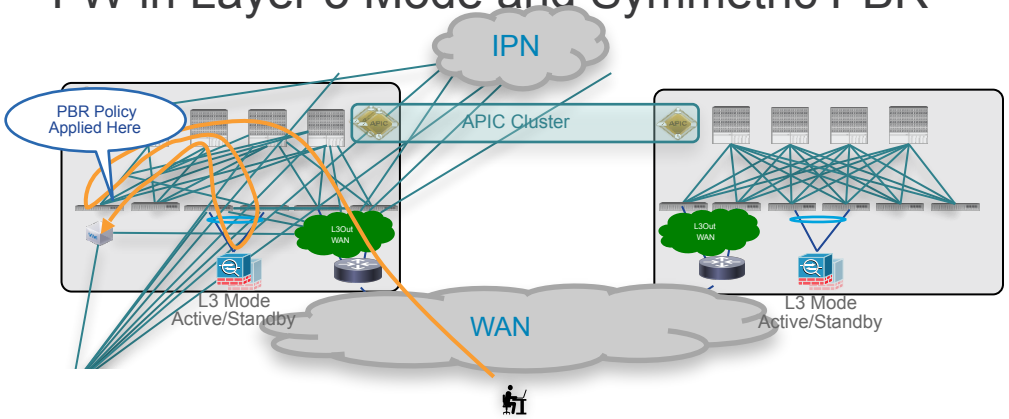

Optimized behavior where the FW located in the same Pod with the destination endpoint is selected by the PBR policy

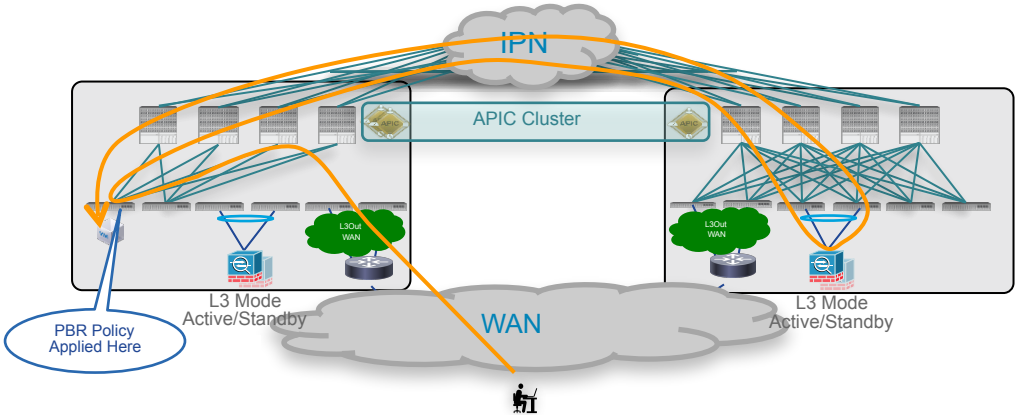

Sub-optimal traffic flow if the FW located in the remote Pod is selected by the PBR policy

Active/Standby Independent Pairs in each Pod

Location Based PBR (ACI 3.1 Release)

- The suboptimal behavior shown in the previous slide could be avoided introducing ‘location based PBR’
  - The leaf node where the PBR policy is applied would always preferably select a local PBR node
  - A tracking functionality is also available to verify the liveliness of local PBR nodes
  - Remote PBR nodes start getting used if/when all the local PBR nodes have failed
- Coupled with a solution of ingress path optimization (i.e. host routes advertisement on GOLF L3Outs) helps minimizing traffic hair-pinning across the IPN

ACI 3.1

Active/Standby Independent Pairs in each Pod

Location Based PBR (ACI 3.1 Release)

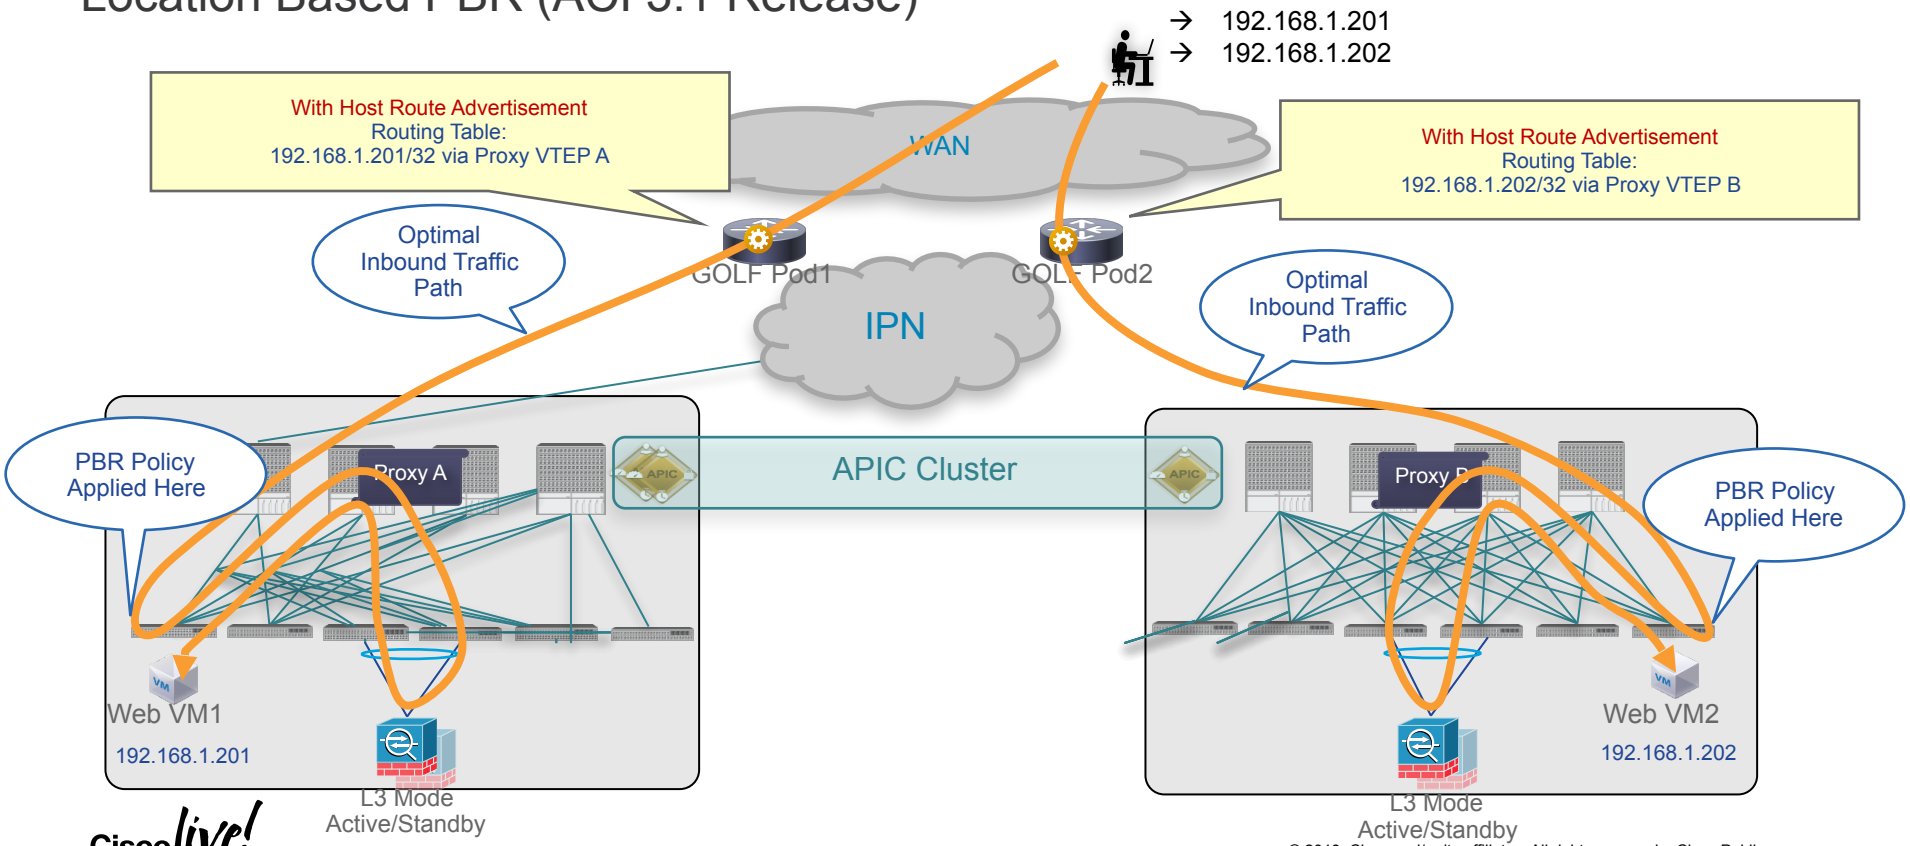

Upcoming New Features

L4-L7 Service Related Enhancement in ACI 3.2

- Multi-node PBR
- vzAny with PBR
- Resilient Hash PBR

- Anycast IP/MAC support in Multi-pod(Stretch)
- Multi-site with Service Graph (PBR)

Multi-node PBR

ACI 3.2

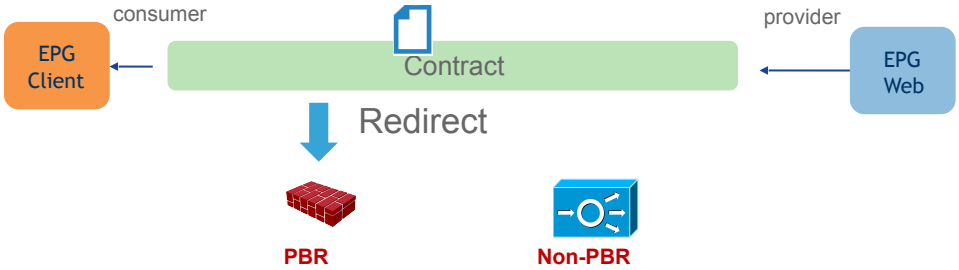

Multi-node PBR

ACI 3.2

- Prior to ACI 3.2: Concatenating PBR nodes is not supported.
  - For example, both 1st and 2nd node can't be PBR nodes. Either one of them can be.

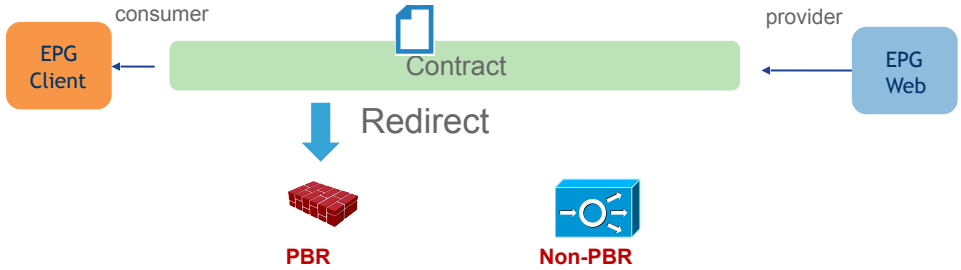

- ACI 3.2: Support more than 1 node PBR in a Service Graph.

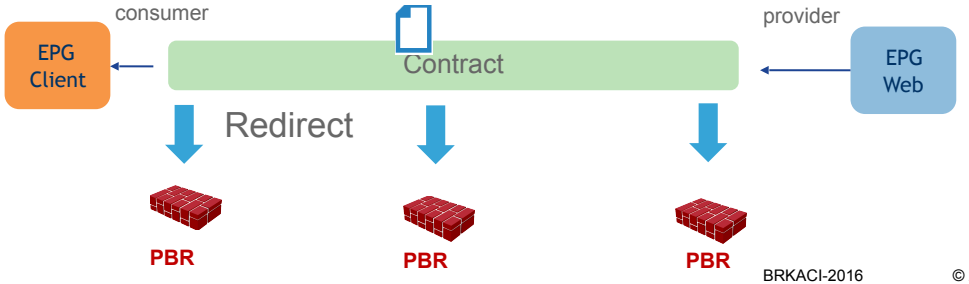

PBR with vzAny

Today

- vzAny is useful if we have a security requirement that is applied to all EPGs in same VRF and also it helps to reduce policy TCAM consumption.
- Today, PBR with vzAny (provider) is not supported.
- vzAny (consumer) can be used for shared service use case.

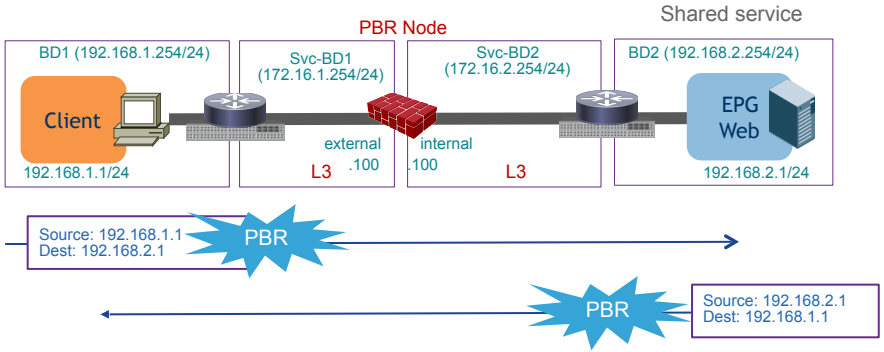

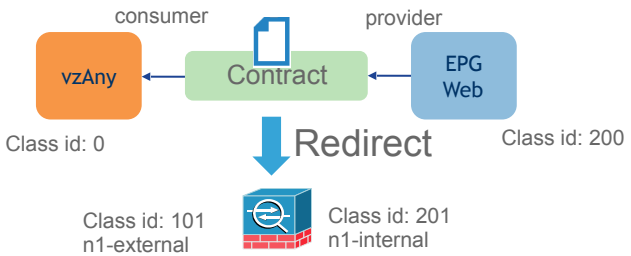

PBR with vzAny

ACI 3.2

- In ACI 3.2, PBW with vzAny (provider) is also supported.
- Use case: Insert Firewall everywhere.

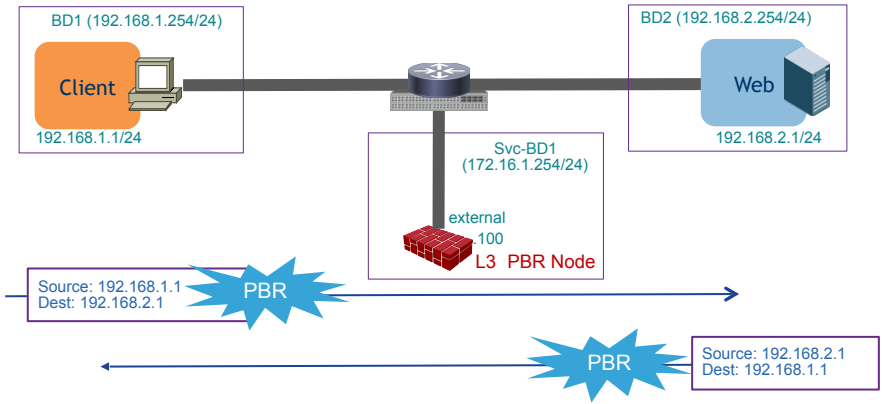

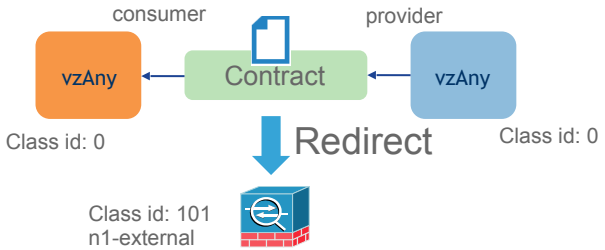

Resilient Hash PBR

Today

- Symmetric PBR is supported today, but if one of the PBR nodes is down, traffic will be re-hashed. So existing connection having been going through available PBR nodes could be affected.

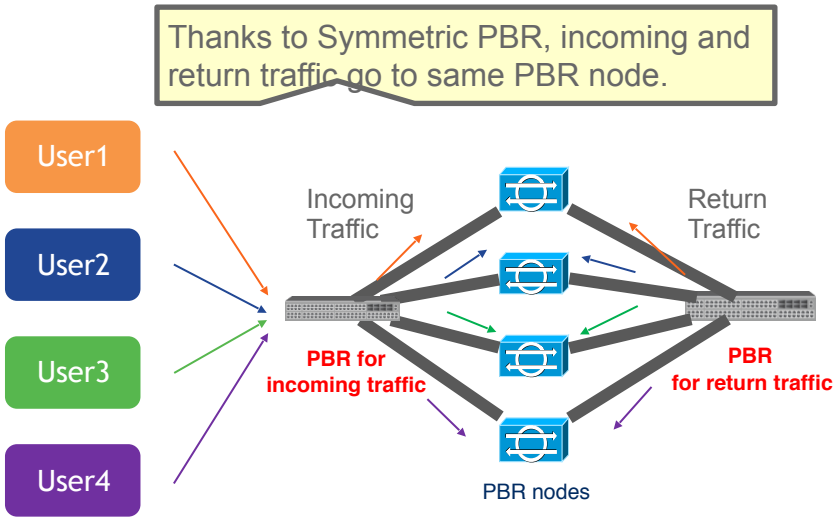

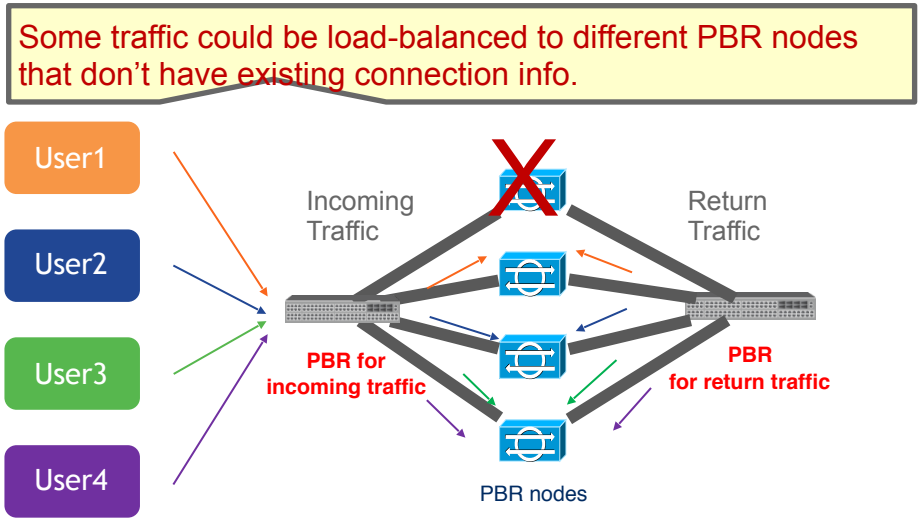

Resilient Hash PBR

3.2

- With Resilient Hash PBR, only the traffics that went though failed node will be re-hashed.

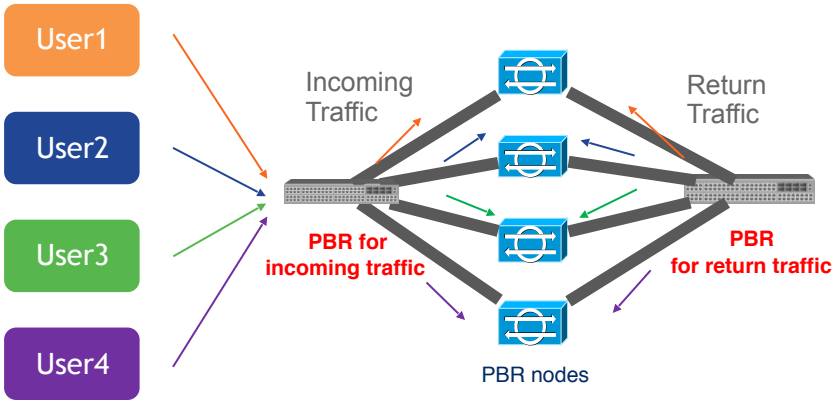

Some traffic could be load-balanced to different PBR nodes that don't have existing connection info.

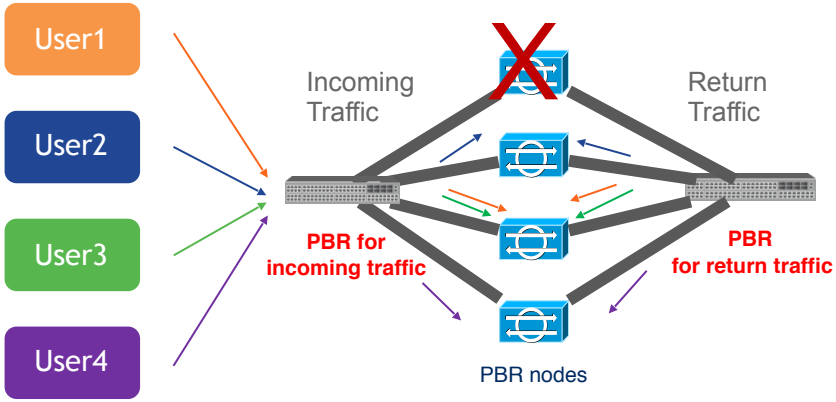

Anycast IP/MAC in Multi-pod

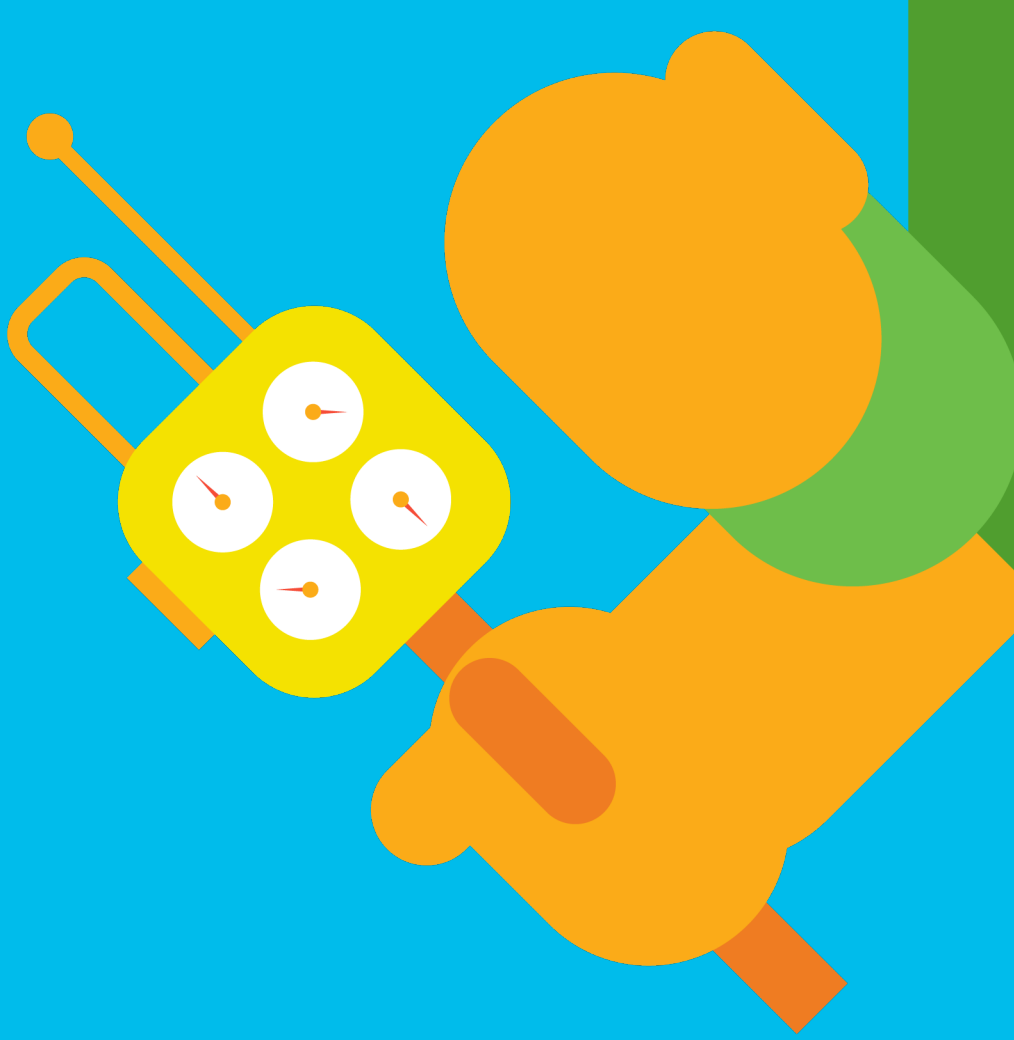

Active/Active Firewall Cluster across pods
without Anycast IP/MAC

NOT SUPPORTED
(Prior to 3.2)

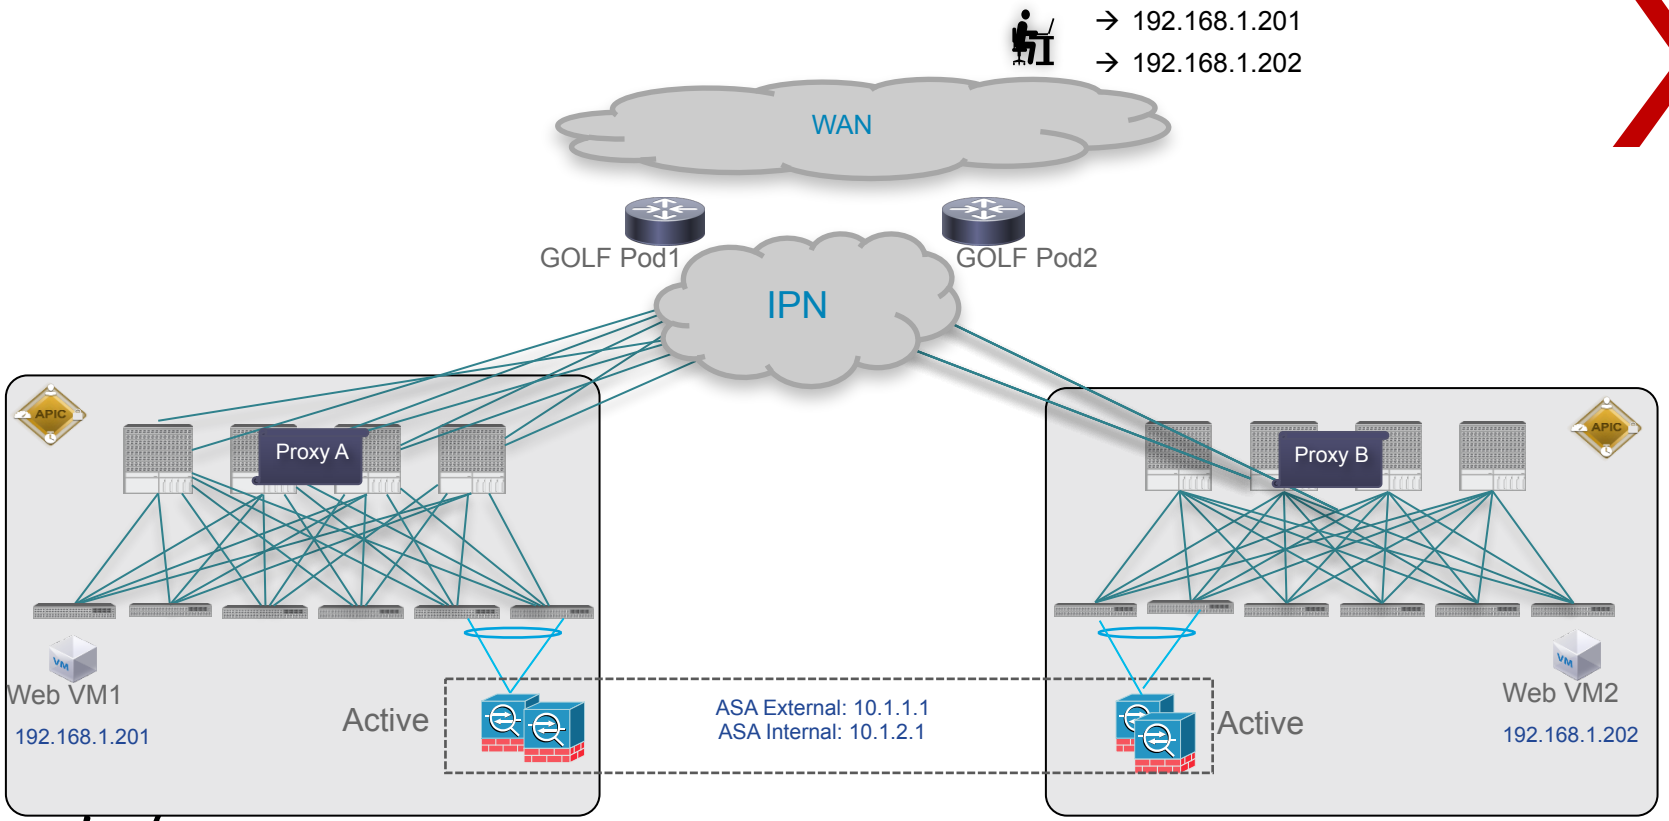

Active/Active Firewall Cluster across pods
without Anycast IP/MAC

NOT SUPPORTED
(Prior to 3.2)

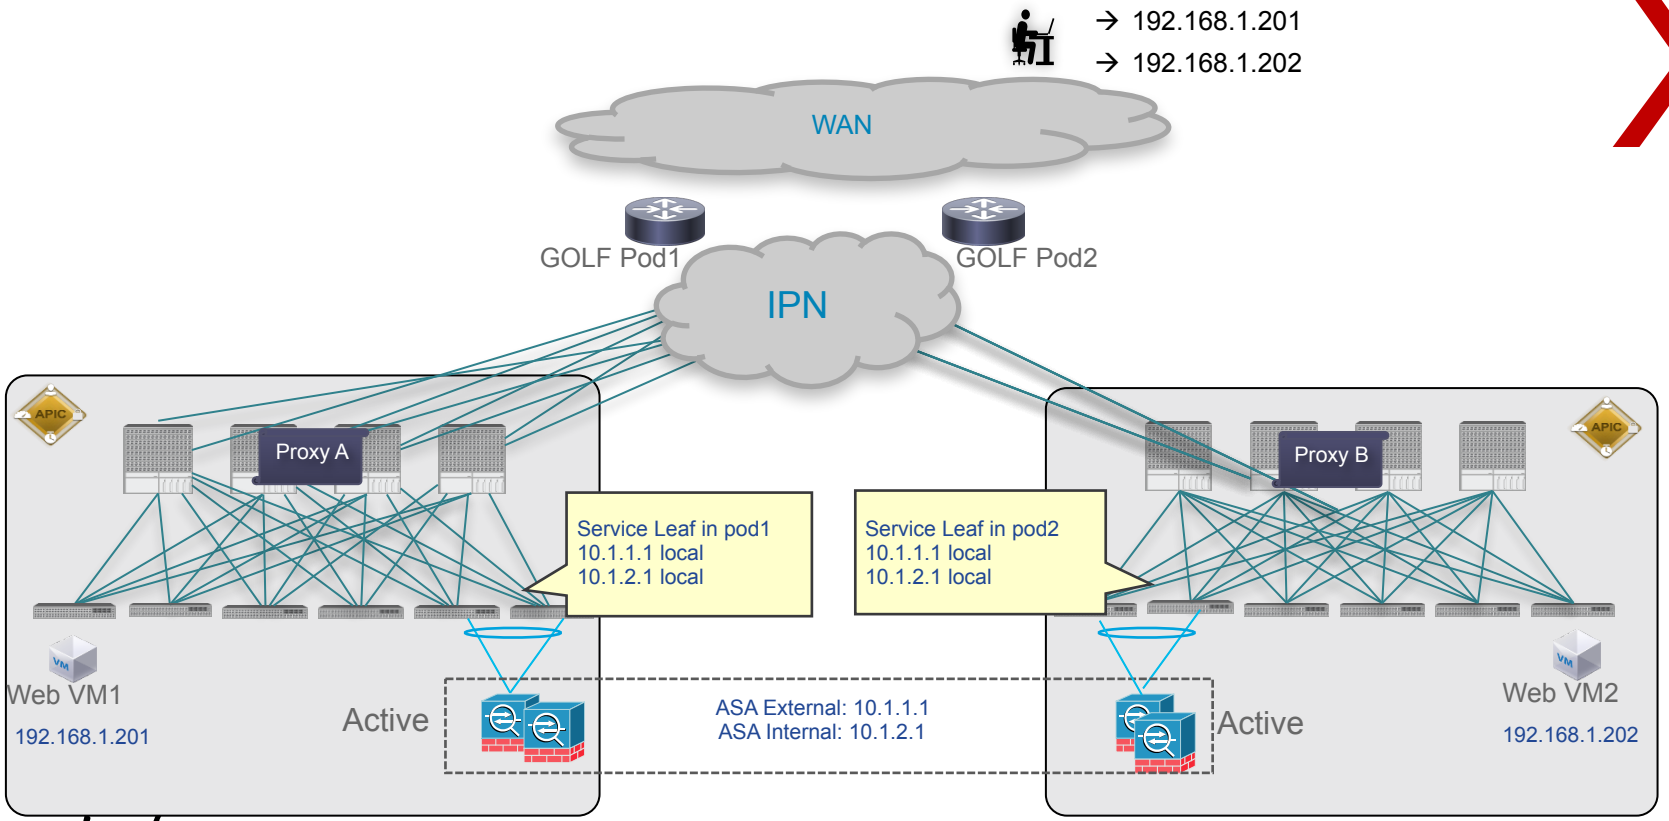

Active/Active Firewall Cluster across pods
without Anycast IP/MAC

NOT SUPPORTED
(Prior to 3.2)

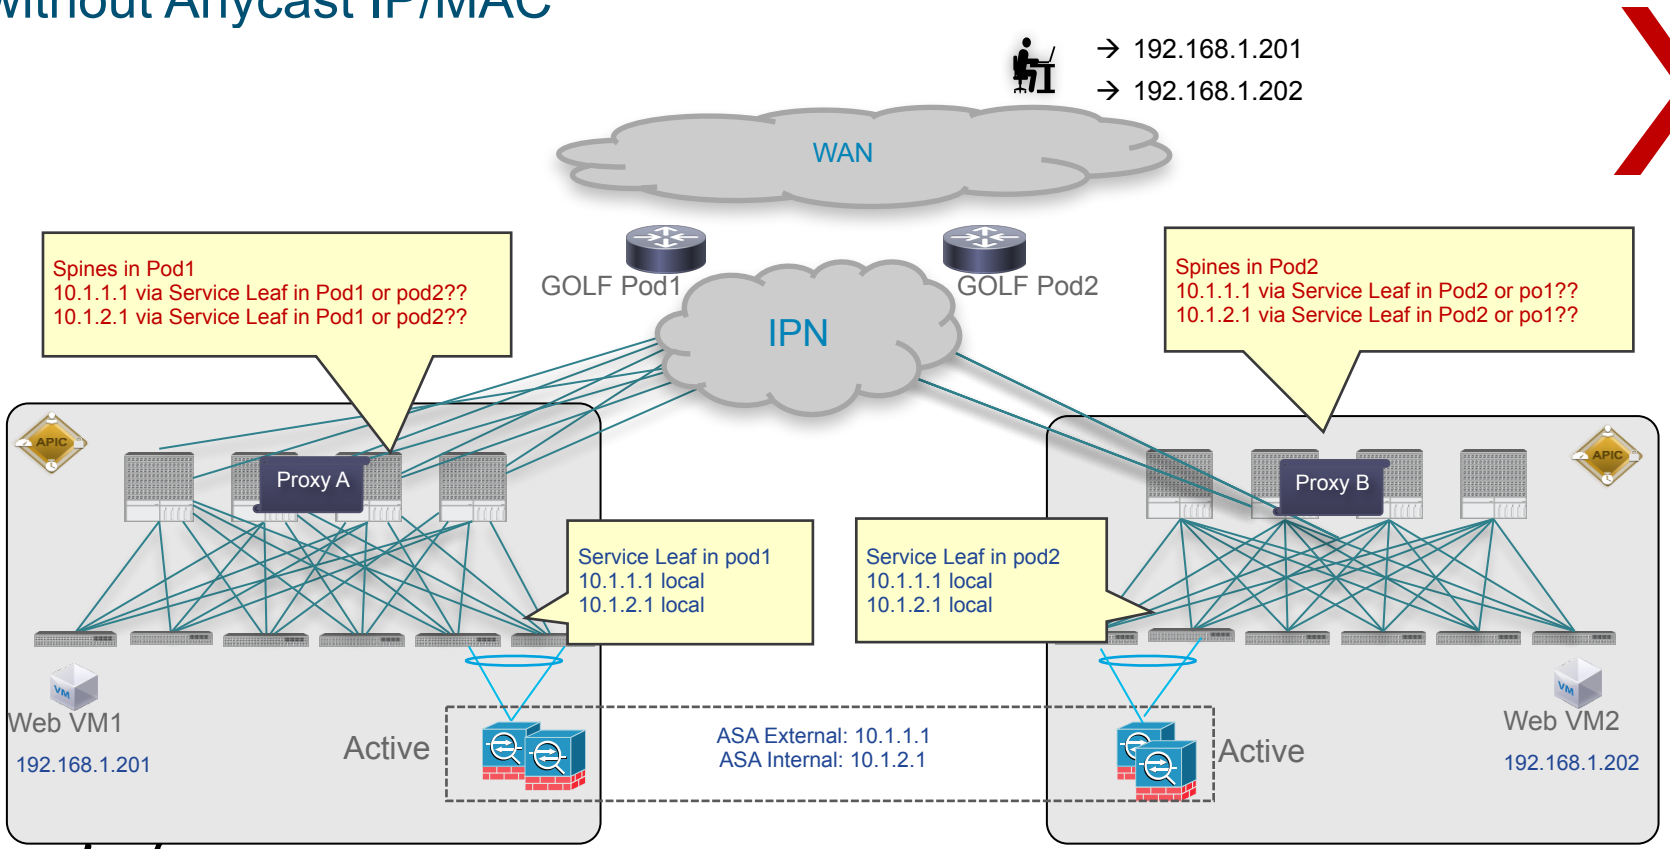

Active/Active Firewall Cluster across pods
with Anycast IP/MAC

3.2

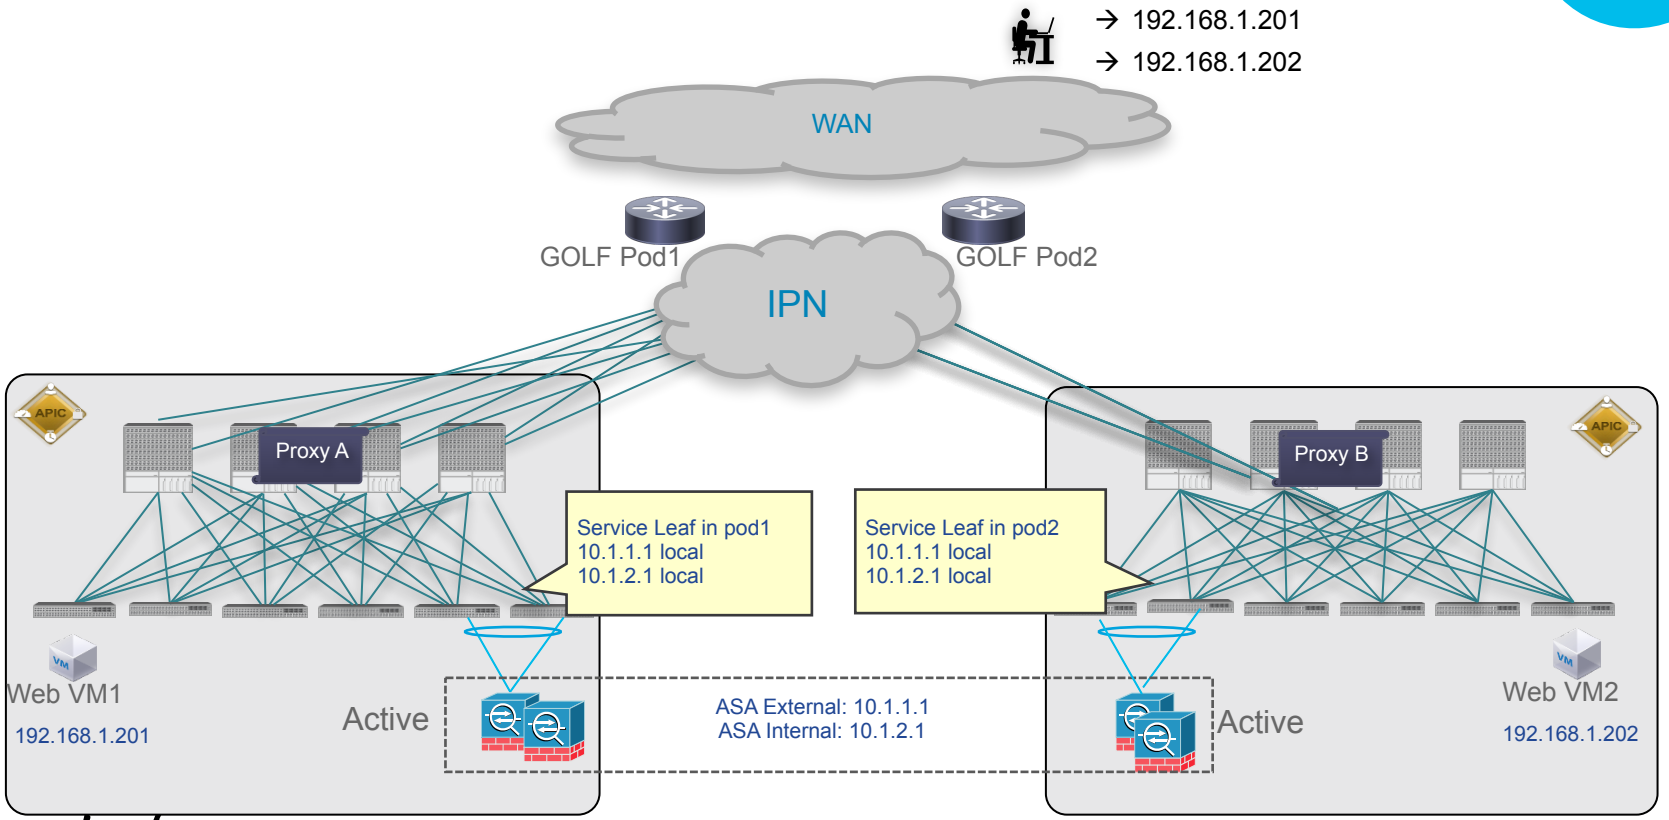

Active/Active Firewall Cluster across pods
with Anycast IP/MAC

3.2

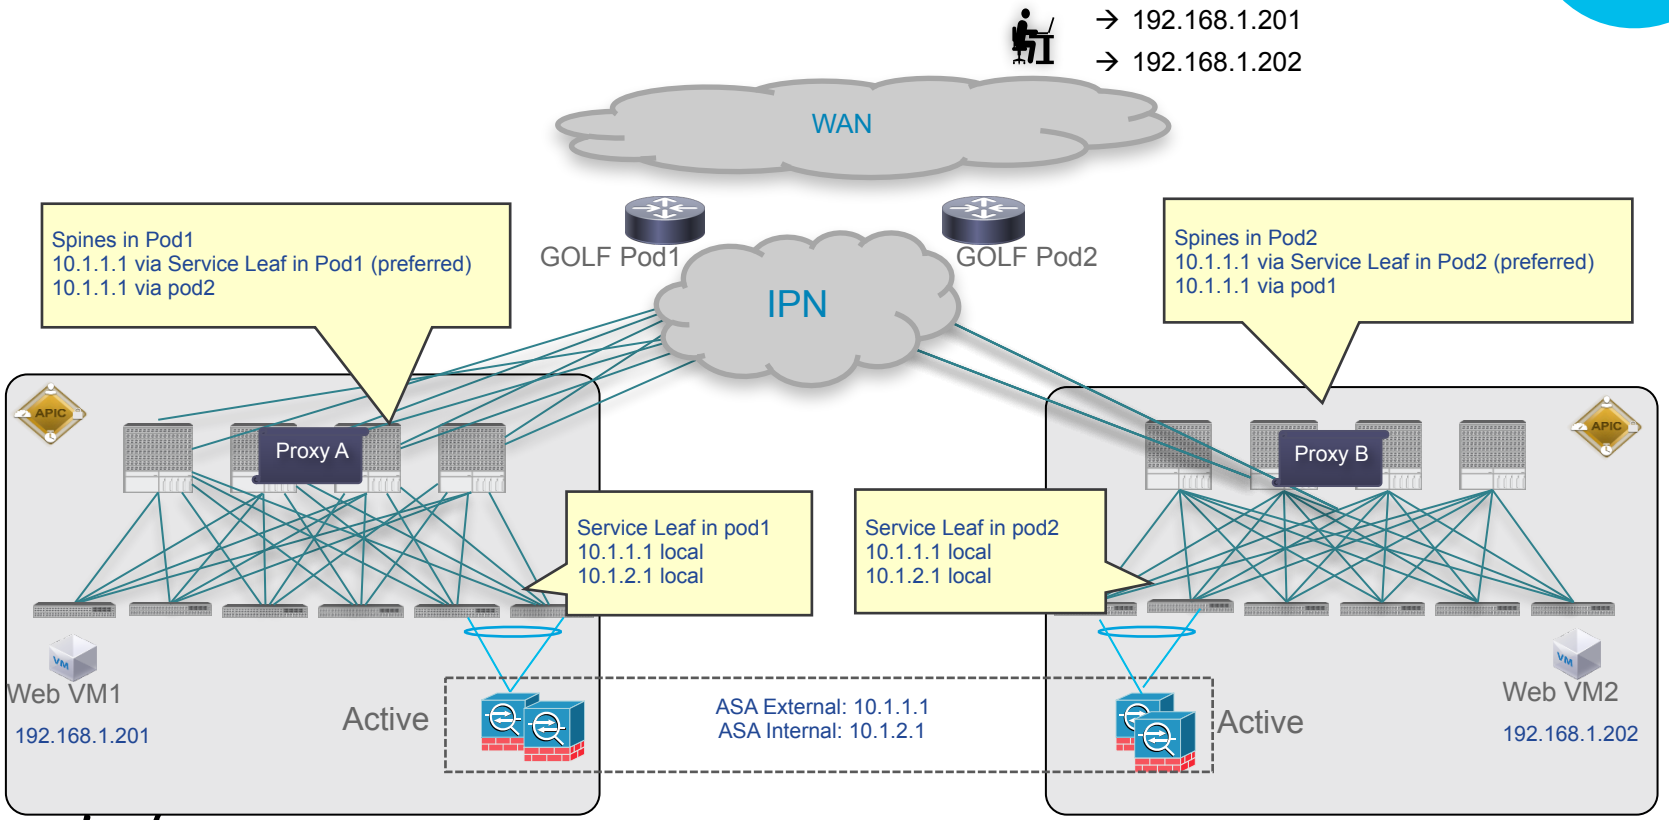

Useful Links

- Service Graph Design with Cisco Application Centric Infrastructure White Paper
  - https://www.cisco.com/c/en/us/solutions/collateral/data-center-virtualization/application-centric-infrastructure/white-paper-c11-734298.html
- Cisco Application Centric Infrastructure Policy-Based Redirect Service Graph Design White Paper
  - https://www.cisco.com/c/en/us/solutions/data-center-virtualization/application-centric-infrastructure/white-paper-c11-739971.html
- ACI Fabric Endpoint Learning White Paper
  - https://www.cisco.com/c/en/us/solutions/collateral/data-center-virtualization/application-centric-infrastructure/white-paper-c11-739989.html
- ACI Multi-pod white paper
  - https://www.cisco.com/c/en/us/solutions/collateral/data-center-virtualization/application-centric-infrastructure/white-paper-c11-737855.html

Q & A

Complete Your Online Session Evaluation

- Give us your feedback and receive a **Cisco Live 2018 Cap** by completing the overall event evaluation and 5 session evaluations.
- All evaluations can be completed via the Cisco Live Mobile App.

Don’t forget: Cisco Live sessions will be available for viewing on demand after the event at www.CiscoLive.com/Global.

Ciscolive!

Thank you


cisco
You're
Cisco live!