ACI L3-out design and best-practise

Massimiliano Sbaraglia

External End-Point EPG with subnet IP defined in ACI

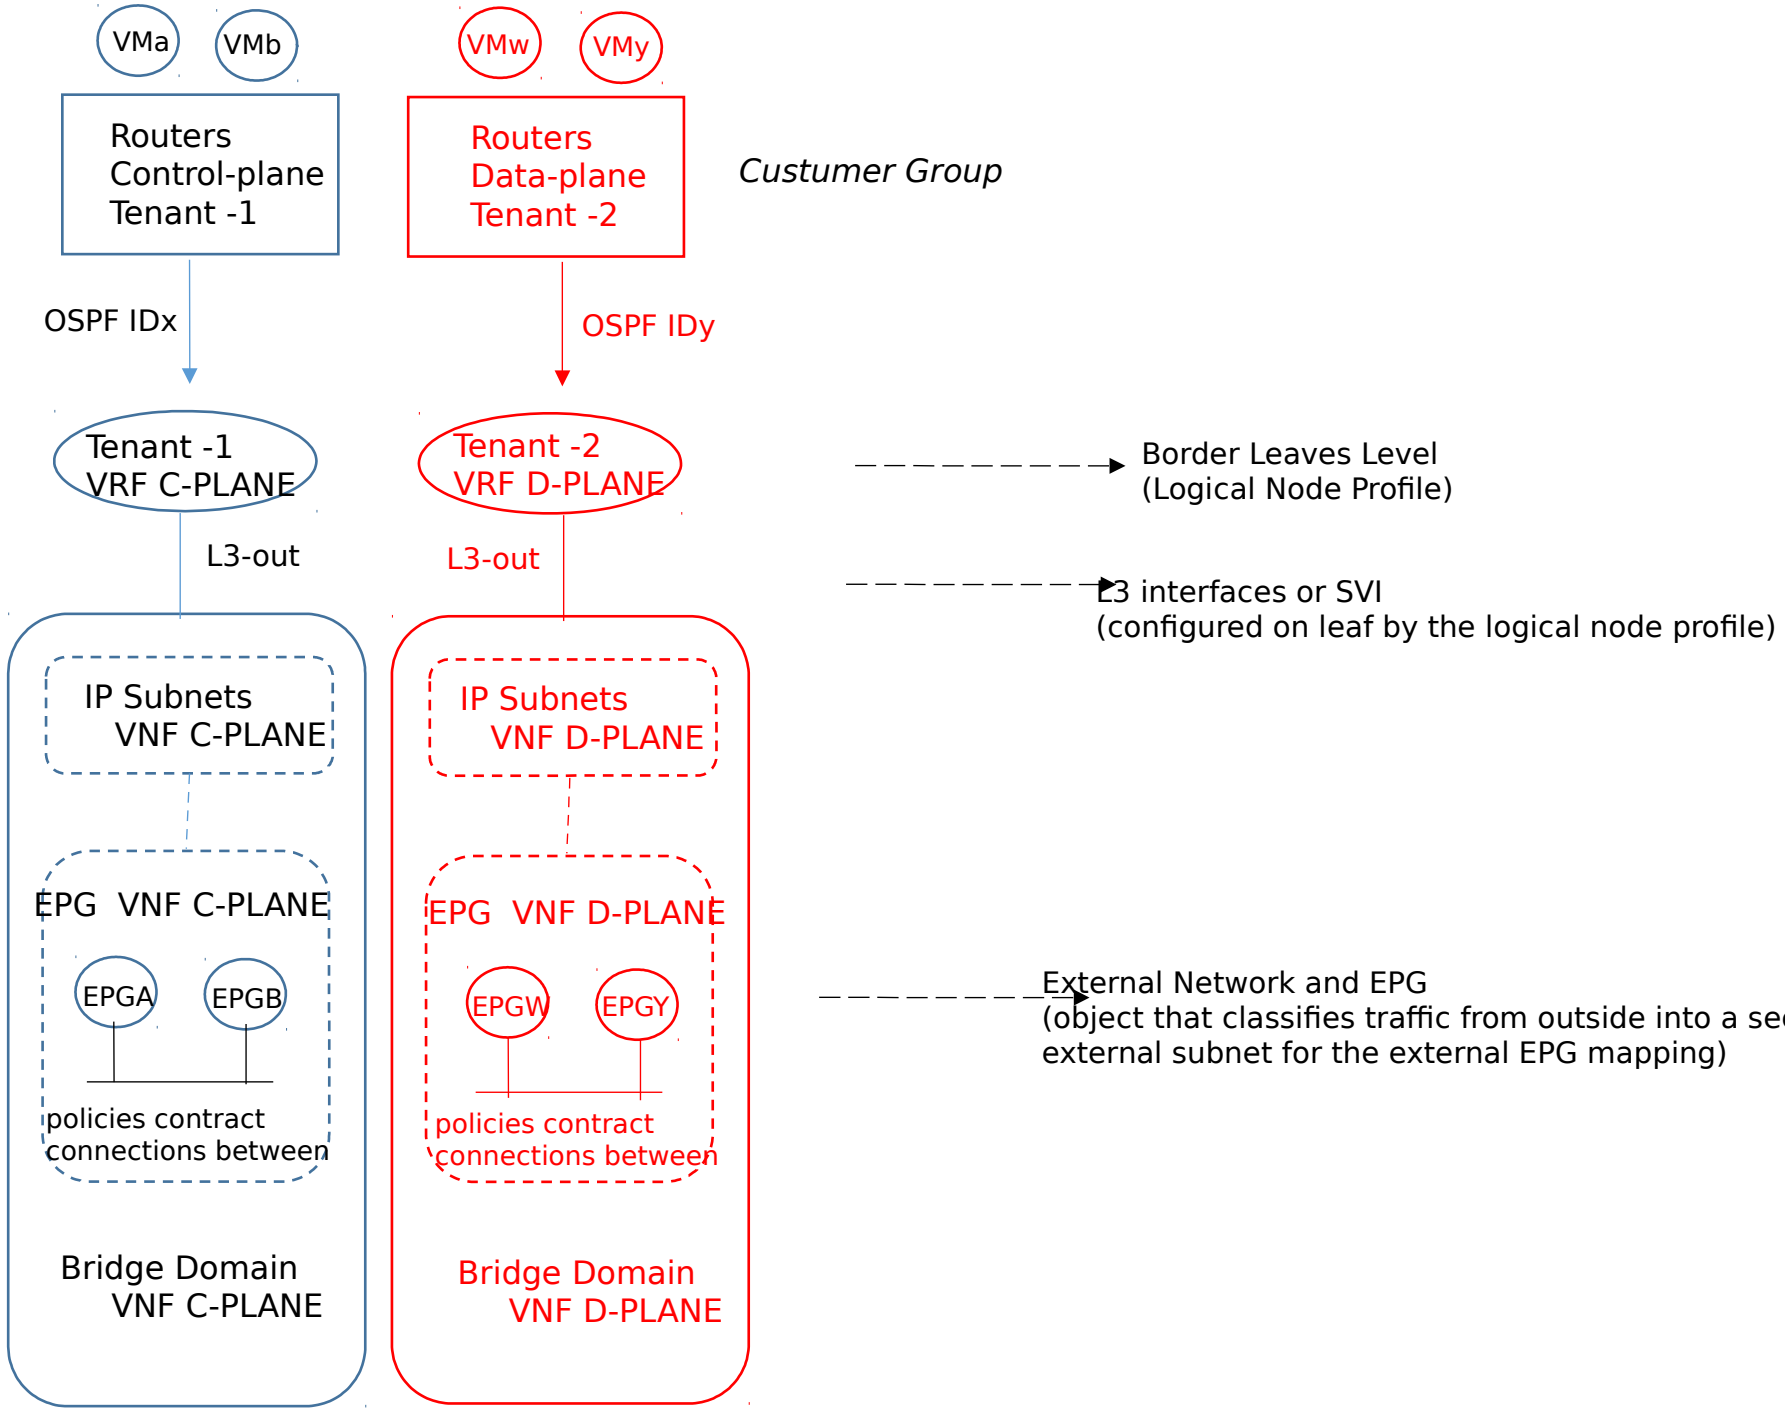

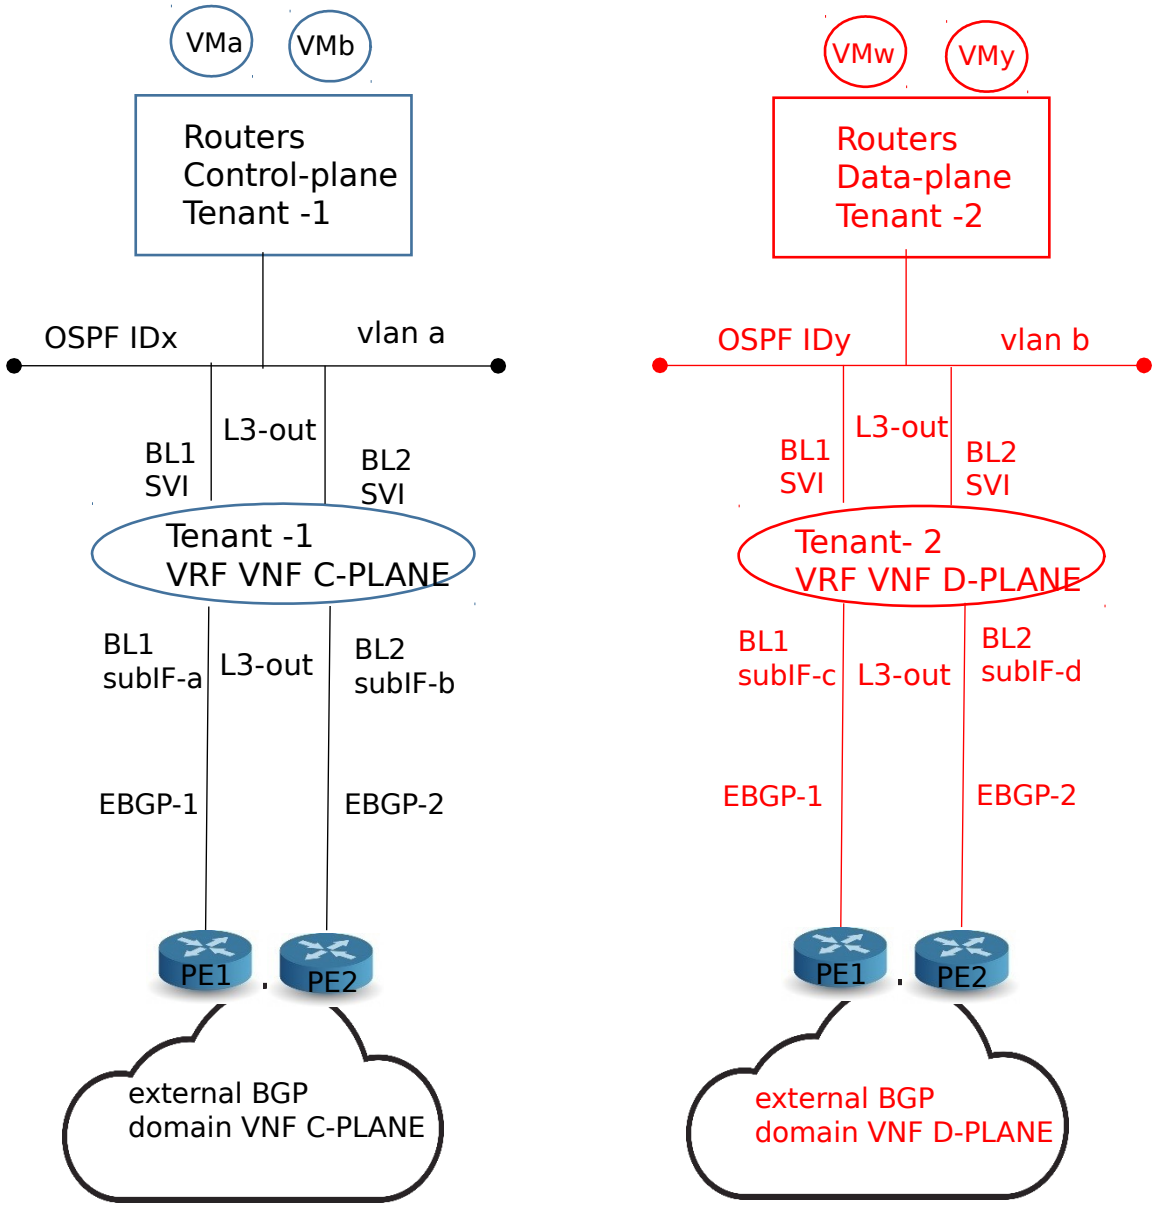

1) Configure a VRF for each customer Tenant

2) Configure L3 out policy associated with a VRF

2a) define the logical node profile ␣ Border Leaves

2b) logical interface profile ␣ SVI interface on the BL defined by the logical node profile

2c) external network and EPG: object that classifier traffic from the outside into the fabric (security zone)

3) L3-out must be referred by the Bridge Domain whose subnet need to be advertised to the outside

4) L3 out policies provide IP connectivity between a VRF and an external IP netw each L3-out is associated with one VRF instances only.

5) For subnet defined in the BD to be announced to the outside router, follow:

5a) the sunbet need to be defined as advertised externally

5b) the BD must have a relationship with the L3 out connection

5c) a contract must exist between layer 3 external EPG and the EPG associated with BD; if this contract is not in place, the advertisment of the subnet cannot occur.

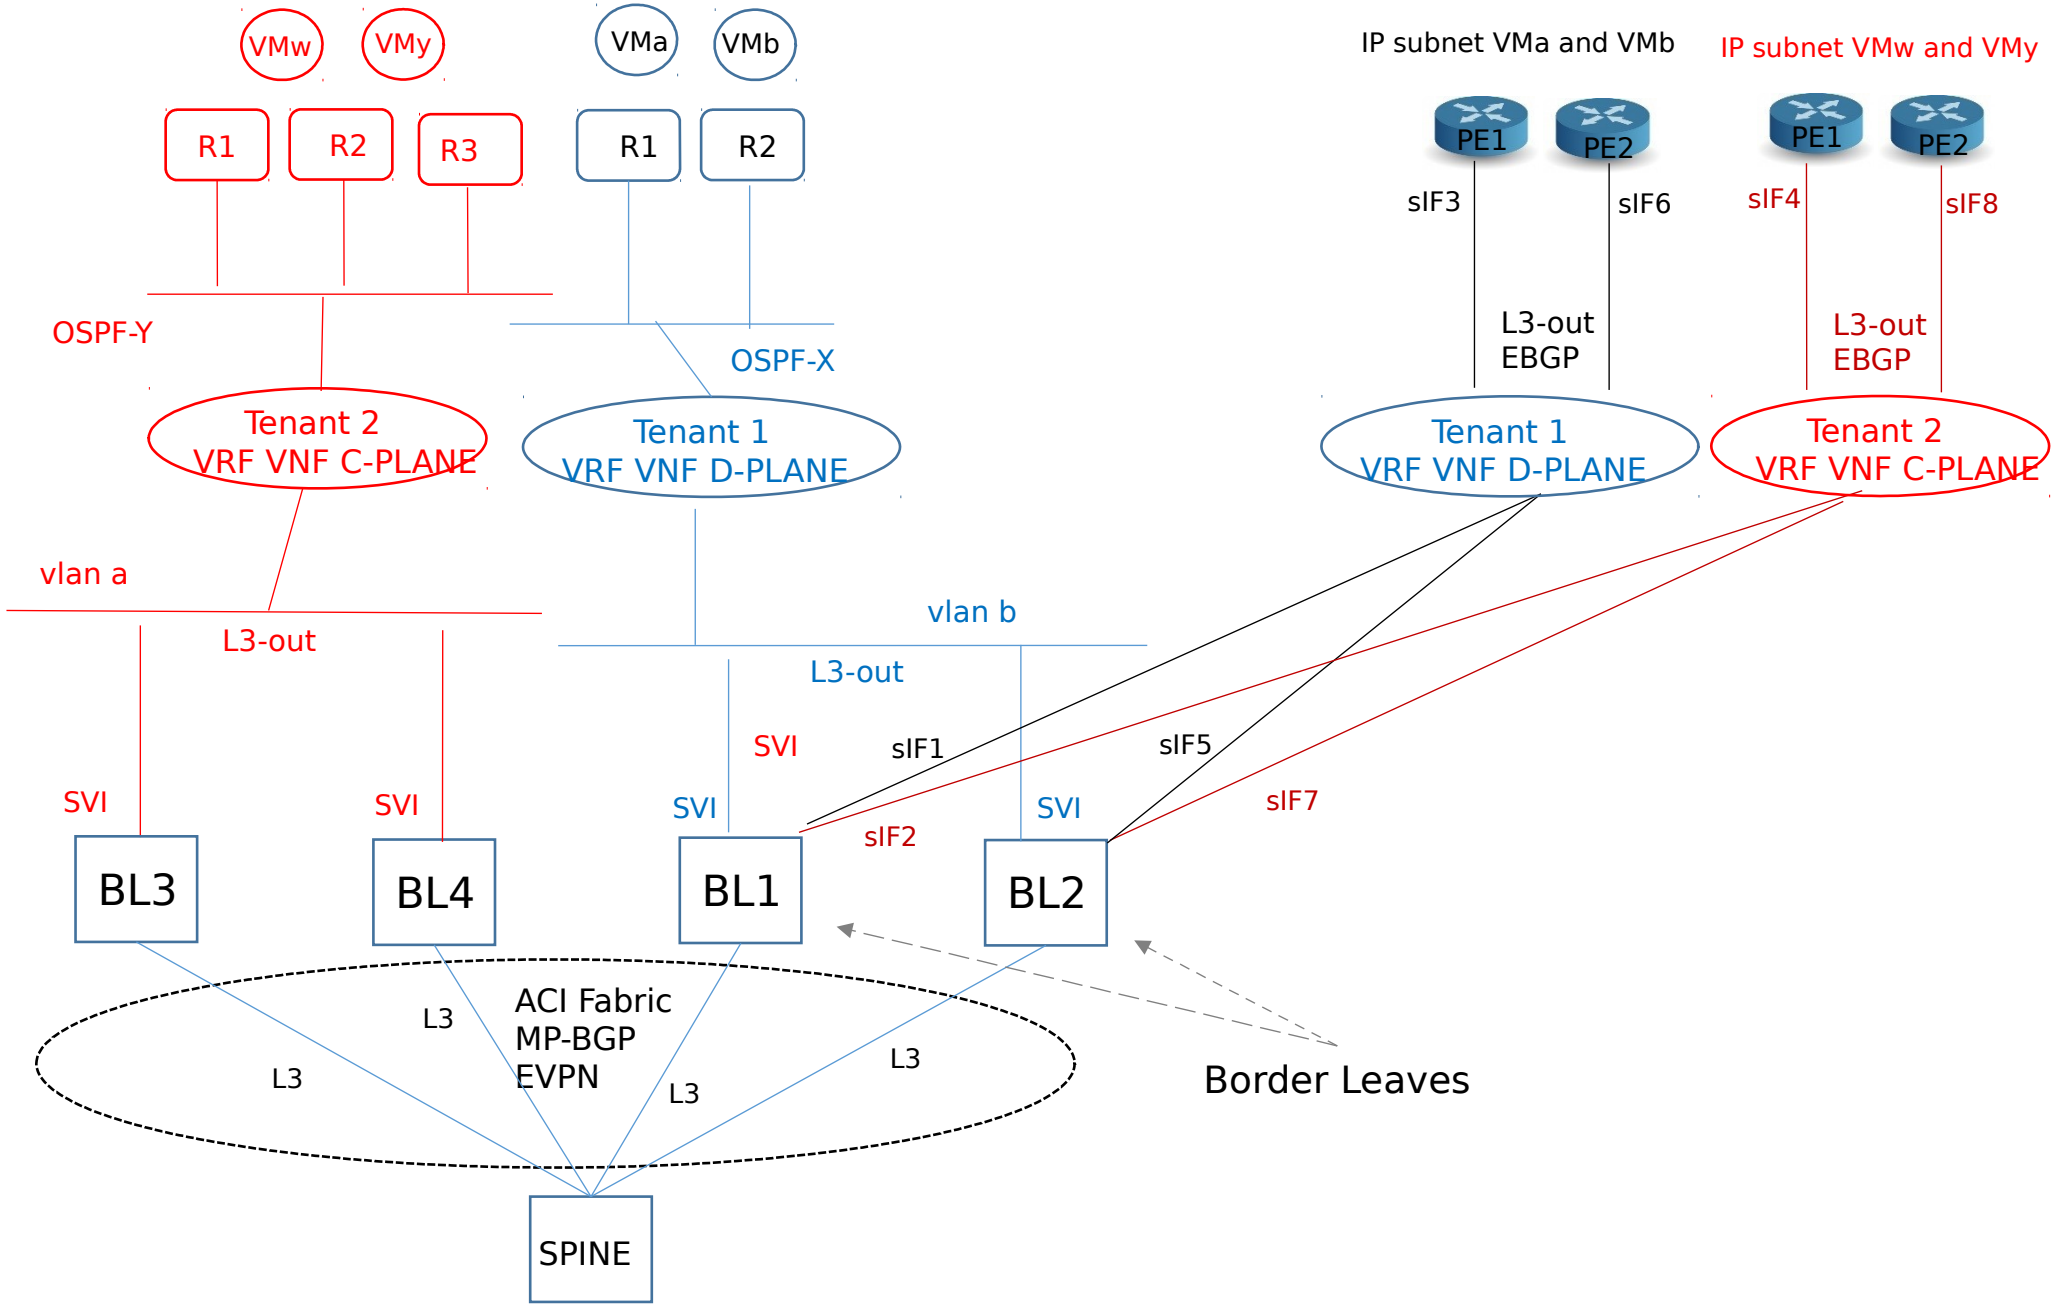
VMw
VMy
VMa
VMb
R1
R2
R3
R1
R2
OSPF-Y
OSPF-X
Tenant 2
VRF VNF C-PLANE
Tenant 1
VRF VNF D-PLANE
vlan a
L3-out
vlan b
L3-out
SVI
SVI
SVI
SVI
SVI
BL3
BL4
BL1
BL2
sIF1
sIF2
sIF5
sIF7
IP subnet VMa and VMb
IP subnet VMw and VMy
PE1
PE2
PE1
PE2
sIF3
sIF6
sIF4
sIF8
L3-out
EBGP
L3-out
EBGP
Tenant 1
VRF VNF D-PLANE
Tenant 2
VRF VNF C-PLANE
ACI Fabric
MP-BGP
EVPN
L3
L3
L3
L3
SPINE
Border Leaves

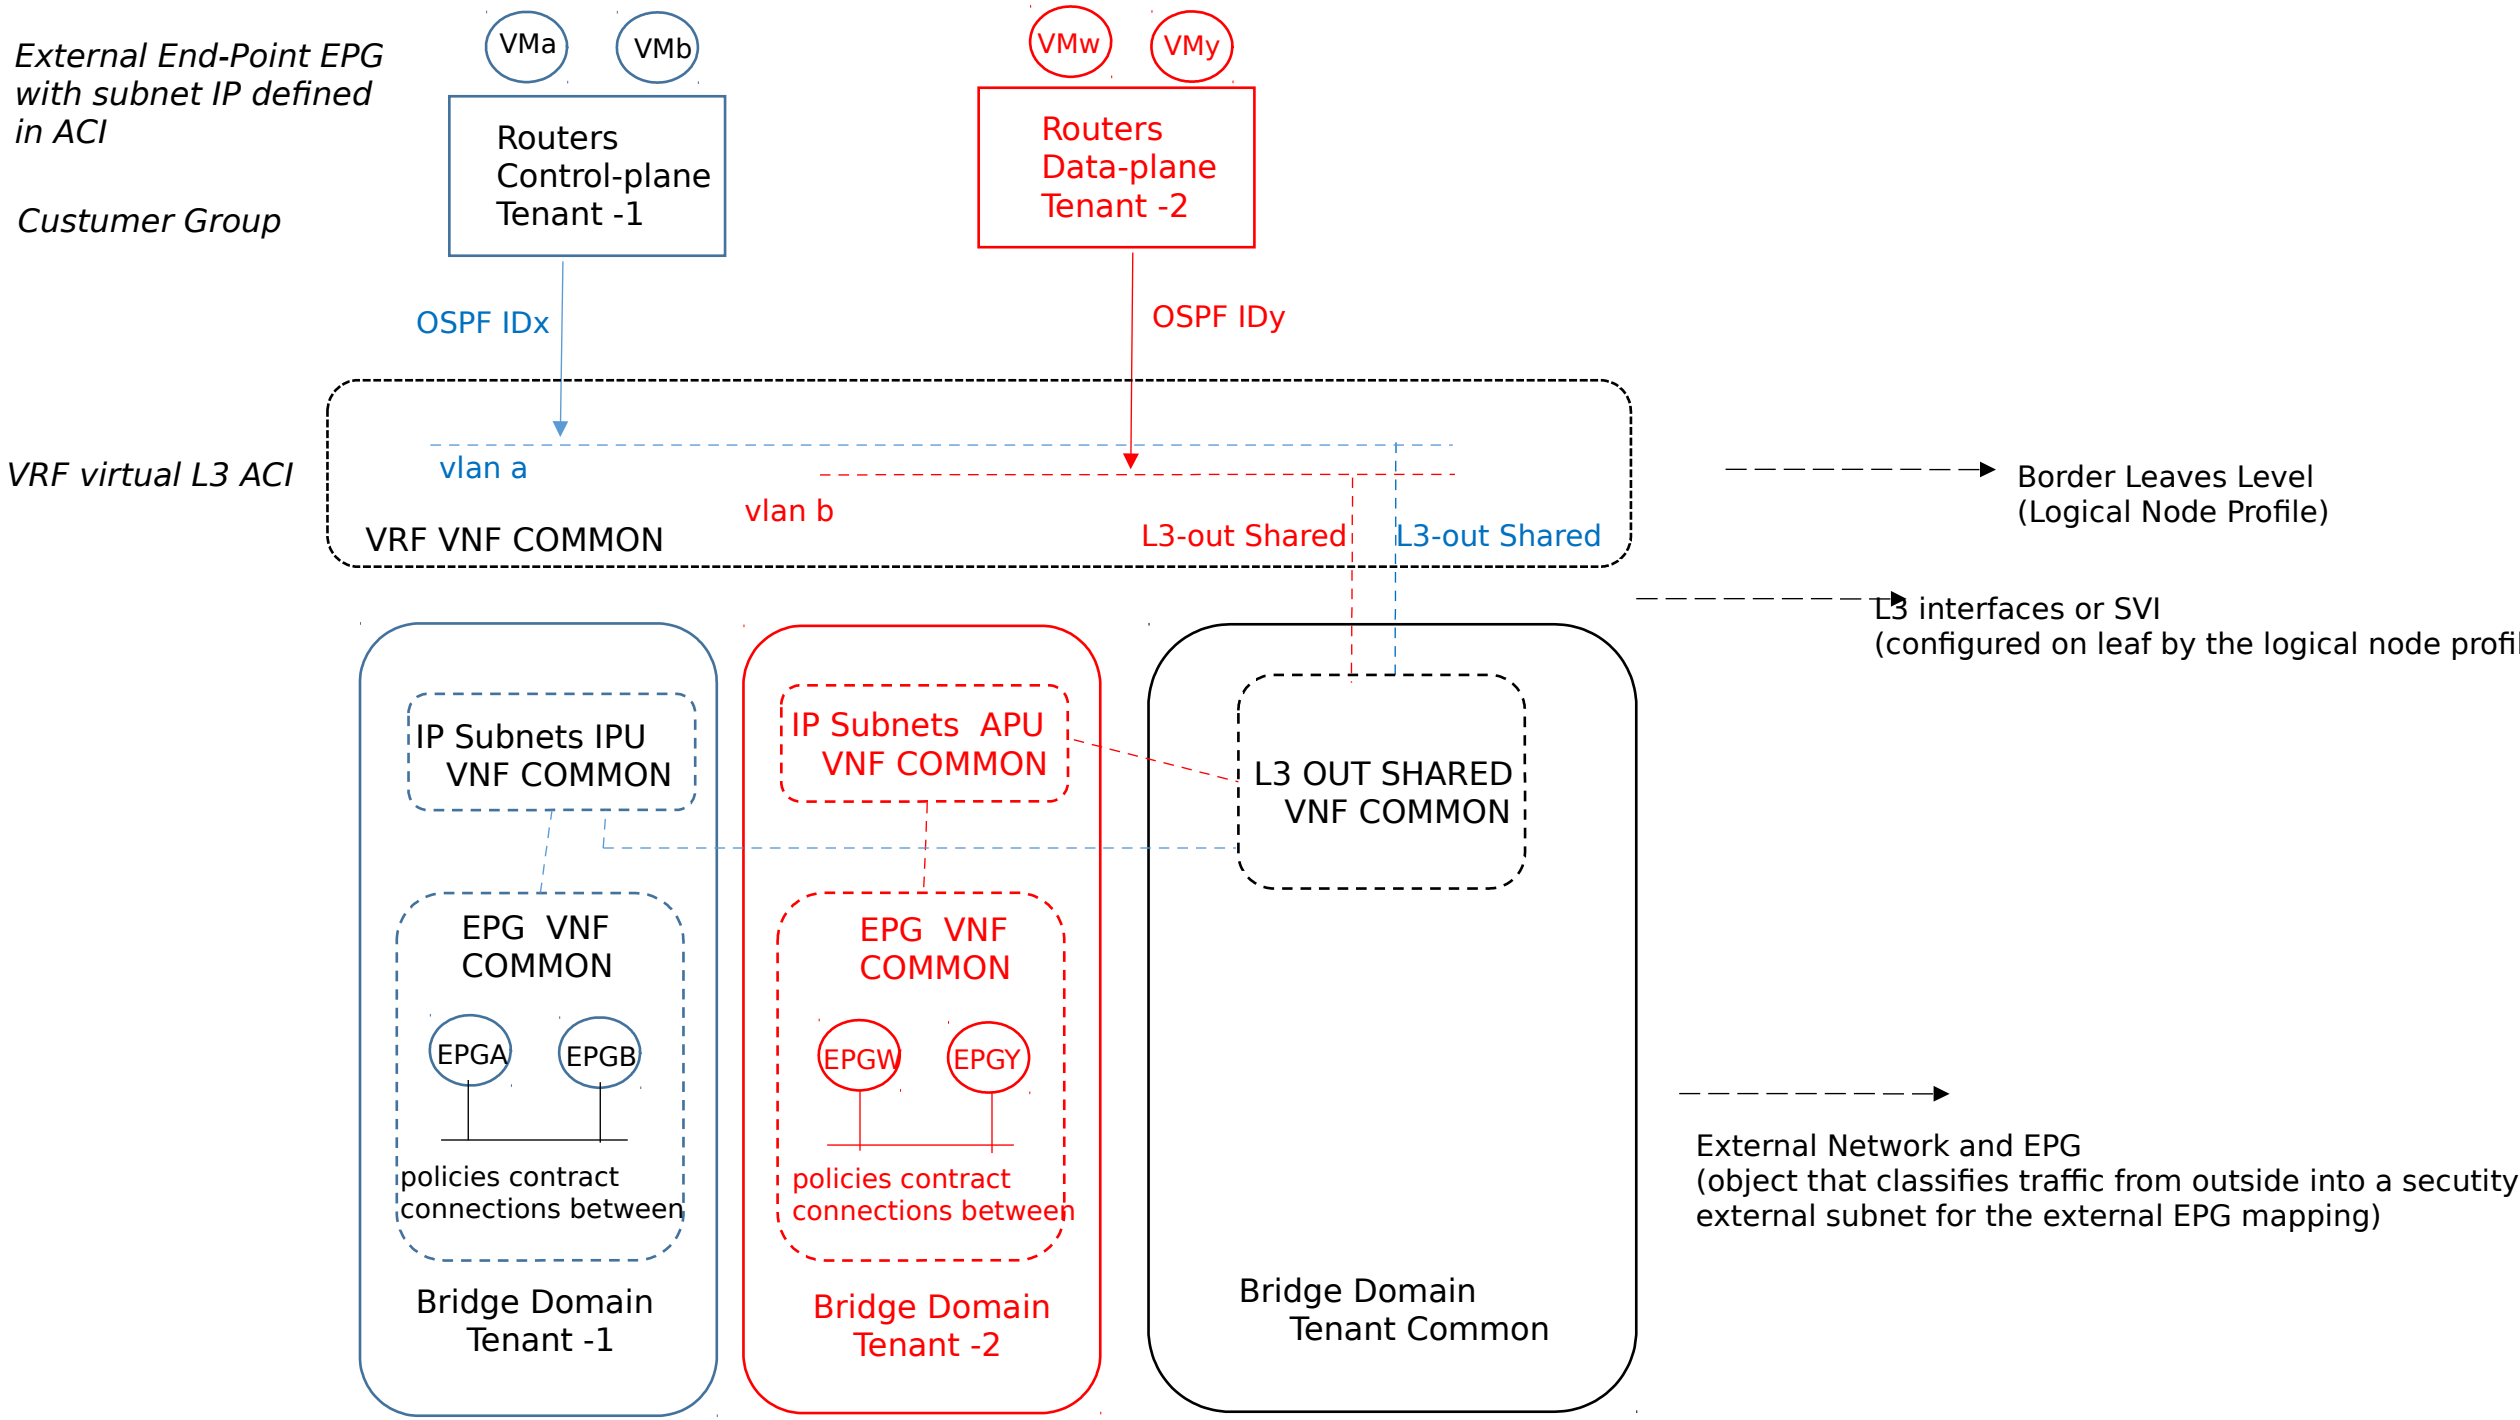
External End-Point EPG with subnet IP defined in ACI
Custumer Group
VRF virtual L3 ACI
VMa
VMb
Routers Control-plane Tenant -1
VMw
VMy
Routers Data-plane Tenant -2
OSPF IDx
OSPF IDy
vlan a
vlan b
VRF VNF COMMON
L3-out Shared
L3-out Shared
Border Leaves Level (Logical Node Profile)
L3 interfaces or SVI (configured on leaf by the logical node profi
IP Subnets IPU VNF COMMON
IP Subnets APU VNF COMMON
L3 OUT SHARED VNF COMMON
EPG VNF COMMON
EPGA
EPGB
policies contract connections between
EPG VNF COMMON
EPGW
EPGY
policies contract connections between
Bridge Domain Tenant -1
Bridge Domain Tenant -2
Bridge Domain Tenant Common
External Network and EPG (object that classifies traffic from outside into a secutity external subnet for the external EPG mapping)

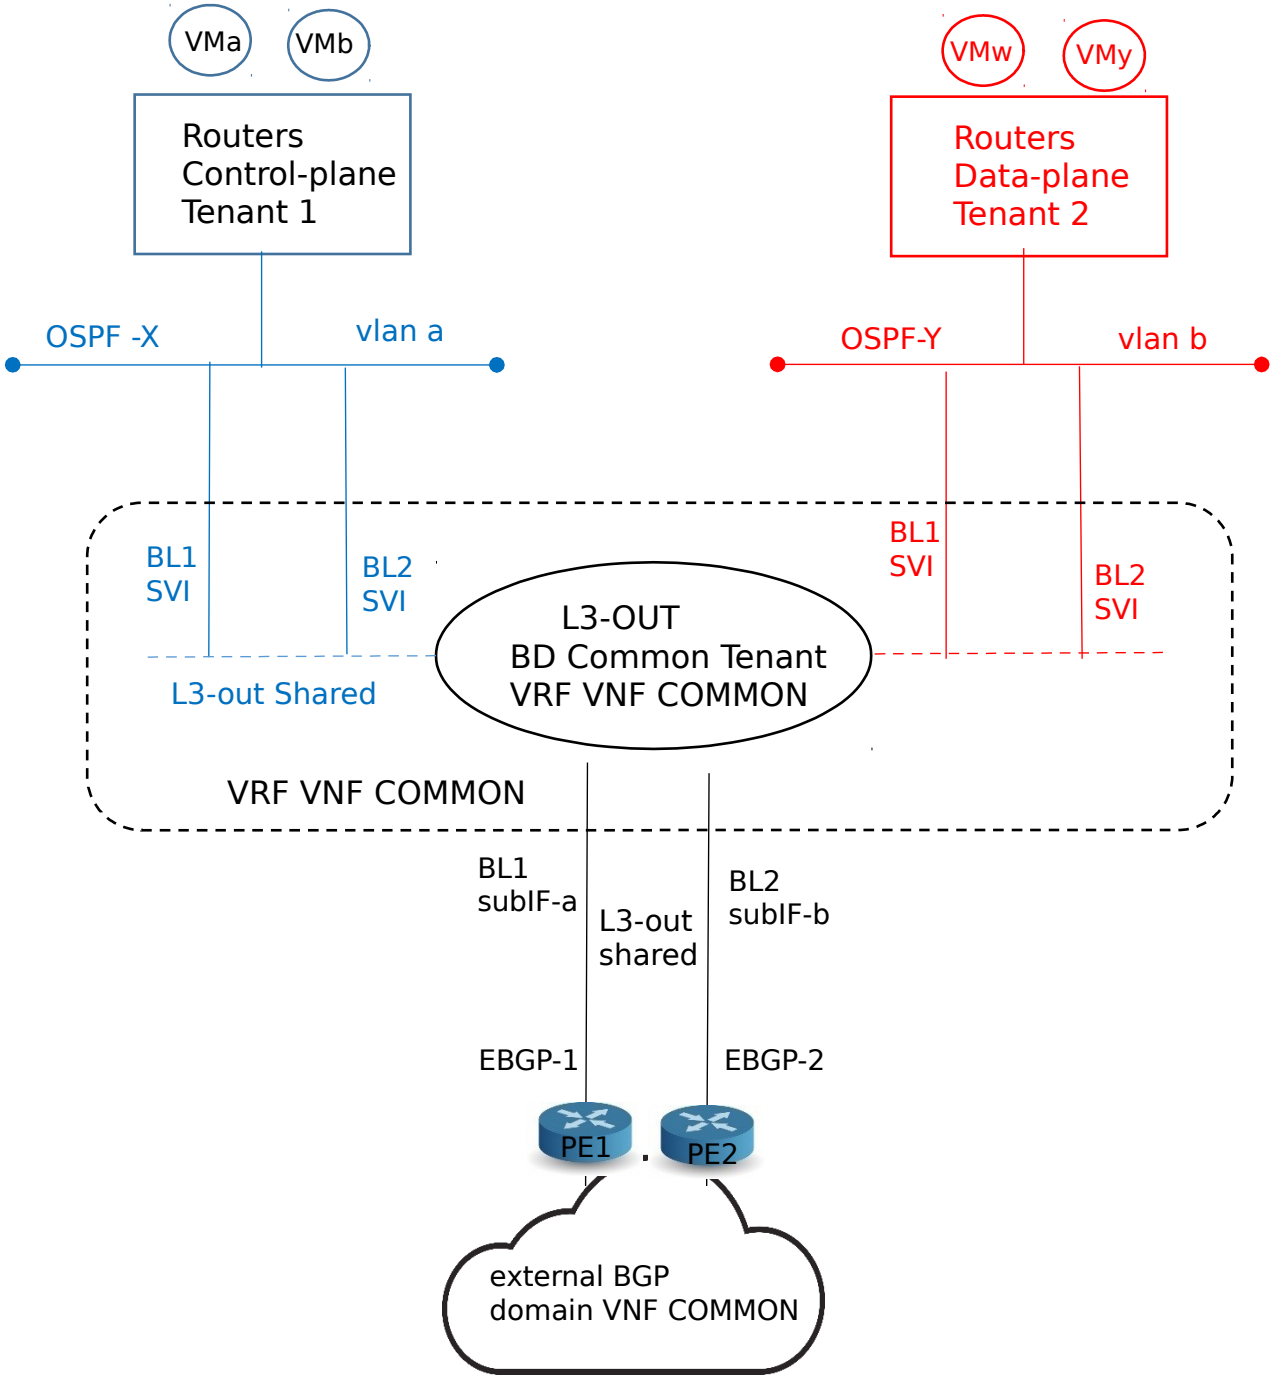

1) Configure a VRF under the common Tenant

2) Configure L3 out connection under the common Tenant and associate it with the VRF instances

3) Configure a Bridge Domain and subnet under each customer Tenant

4) Associate the Bridge Domain with a VRF in the common Tenant and the L3-out connection

5) Under each Tenant configure EPG and associate the EPG with a BD in the Tenant itself

6) Configure contracts and application profiles under each Tenant

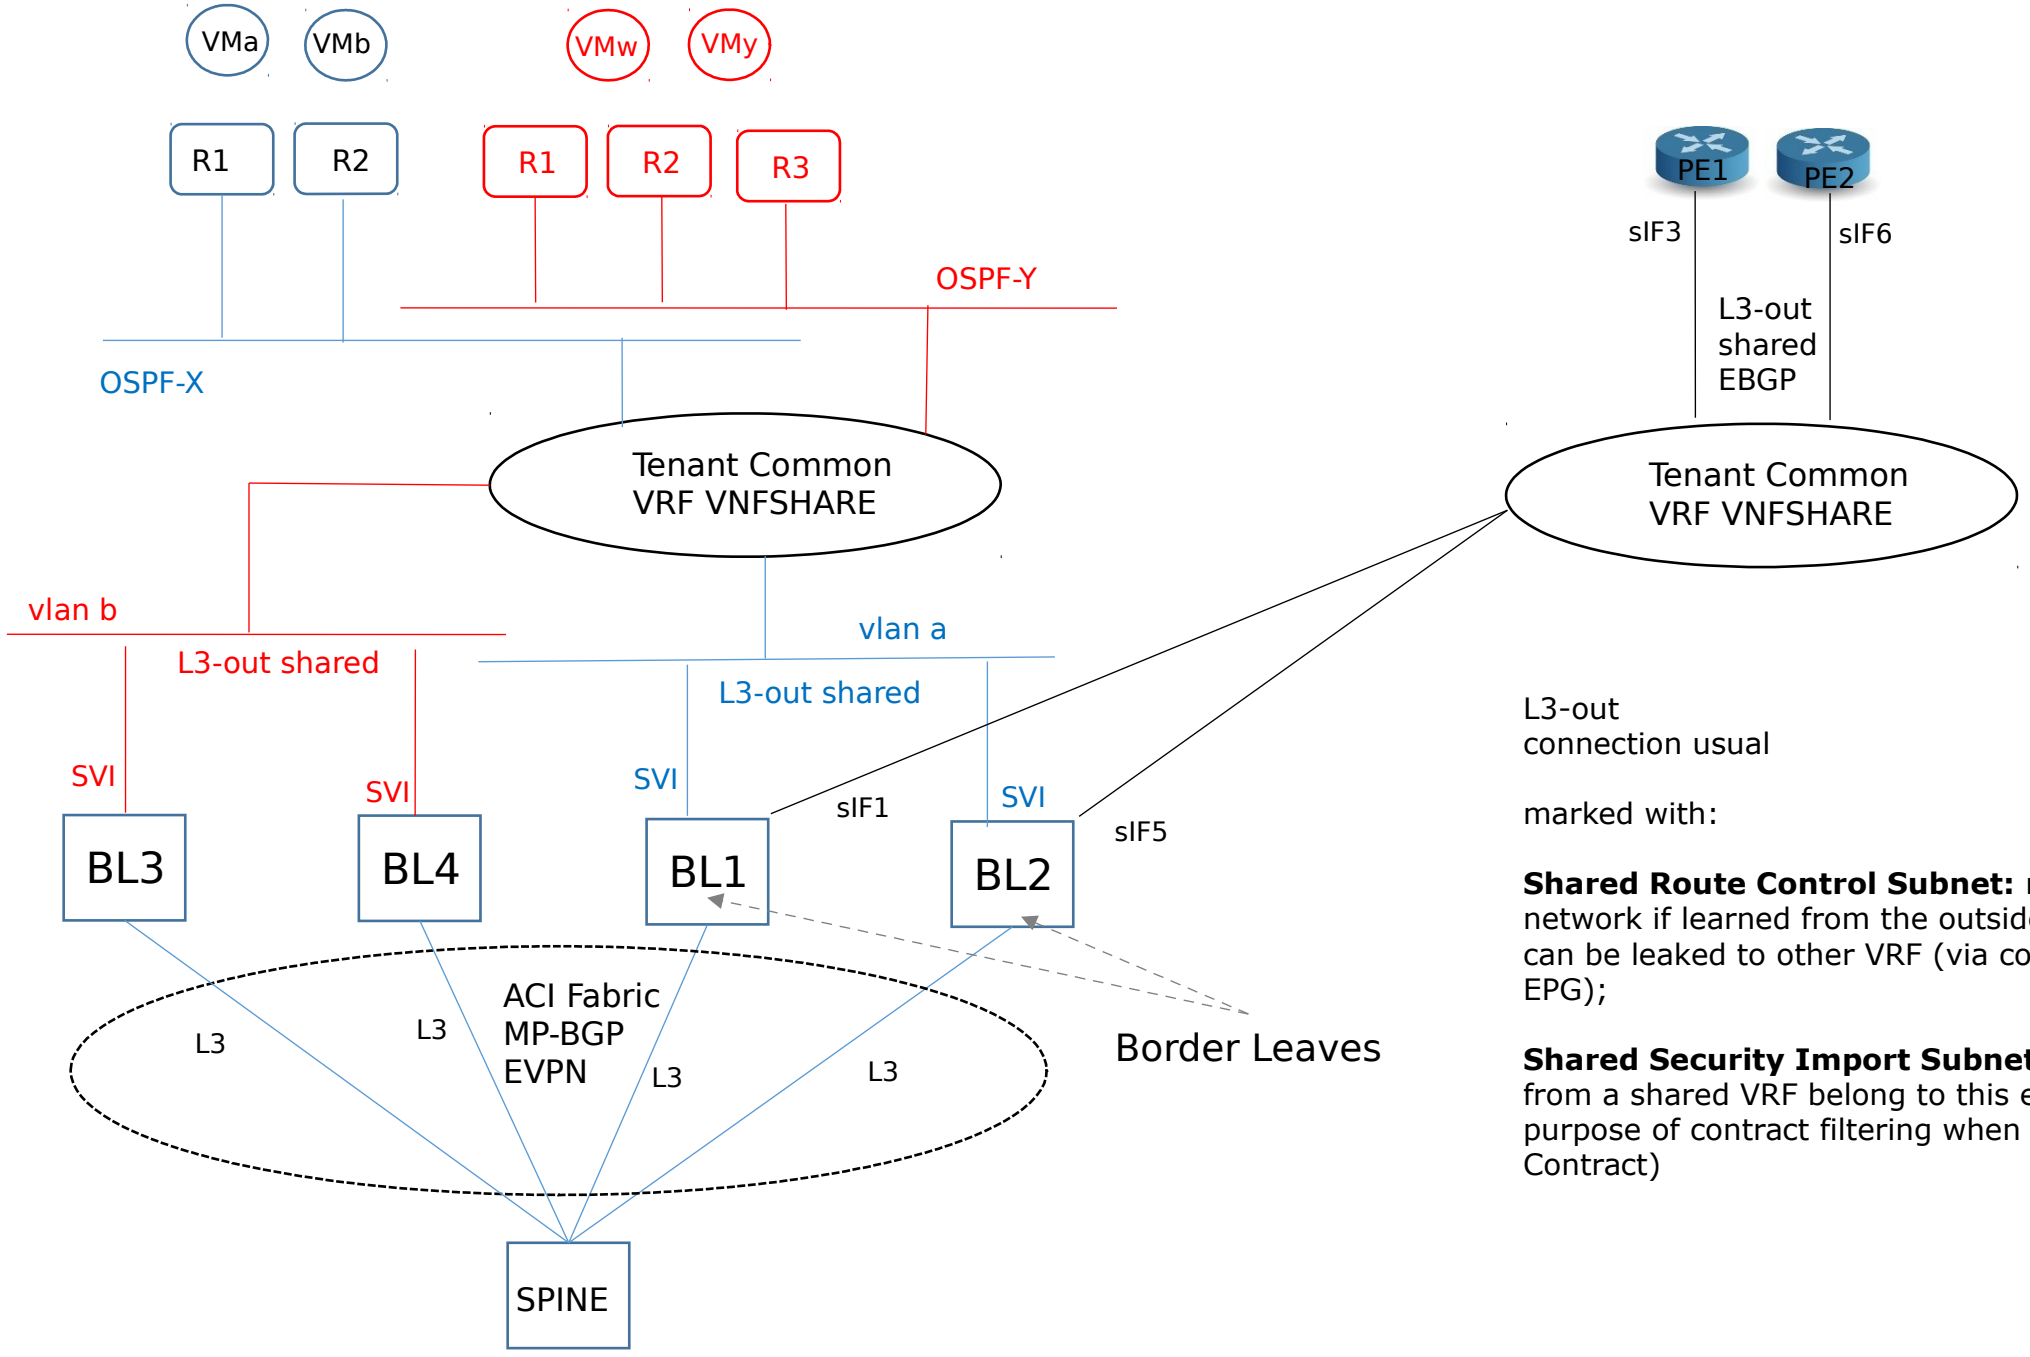

L3-out
connection usual

marked with:

Shared Route Control Subnet: mean that the network if learned from the outside through the VRF can be leaked to other VRF (via contract with the external EPG);

Shared Security Import Subnet: mean which subnet learned from a shared VRF belong to this external EPG for the purpose of contract filtering when establishing a cross-VRF Contract)

OSPF areas on different Border Leaf are different OSPF areas

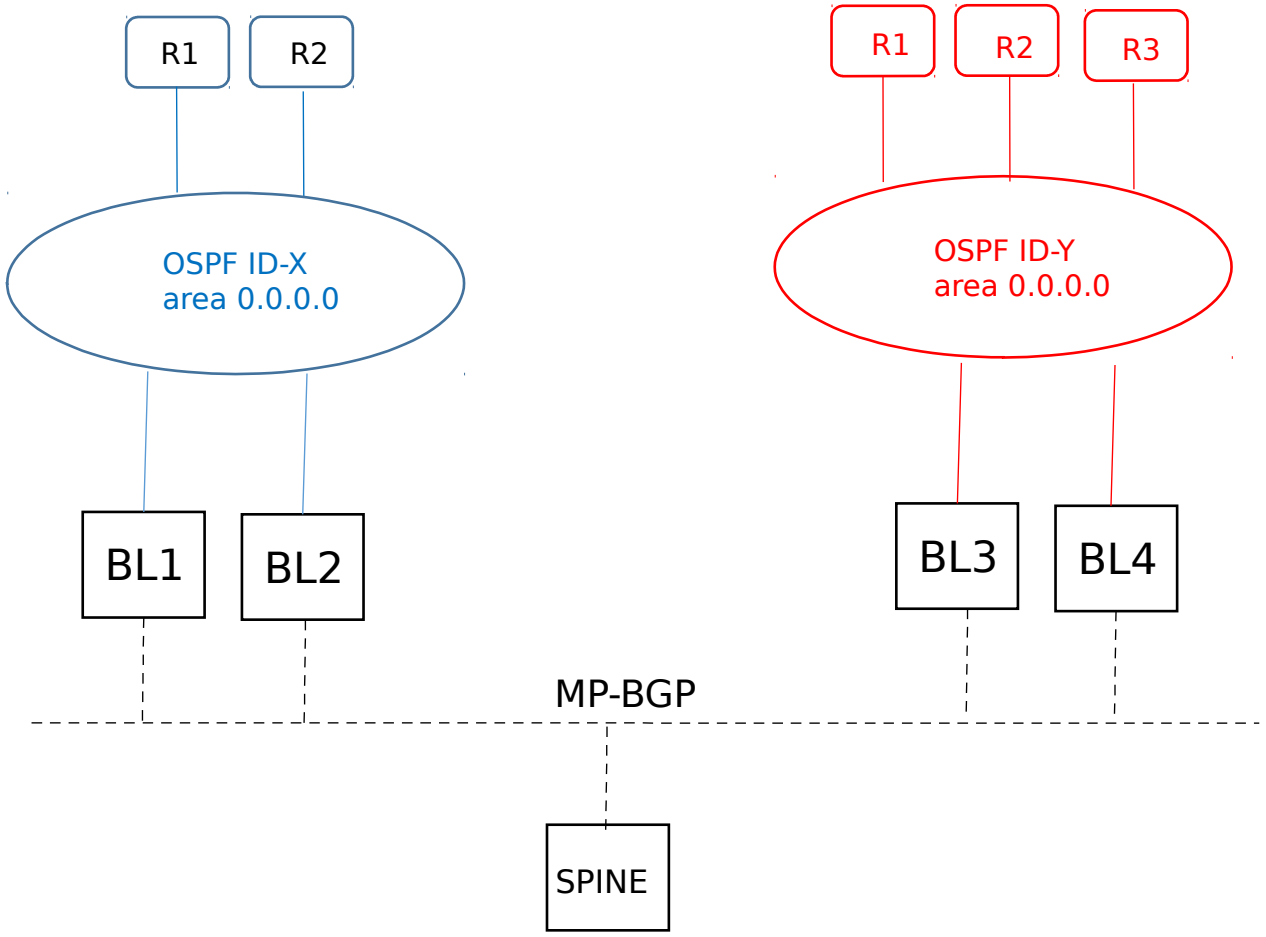

ACI border leaf running OSPF are always AS boundary (ASBR)

all external routes learned in OSPF are redistribuite into MP-BGP

MP-BGP routes are redistribuite into OSPF as external type2

OSPF areas on different border leaf (pairs BL) are different OSPF areas, even if ID match

IPv4 and IPv6 support

OSPF areas on the same Border Leaf need different area type to be advertised

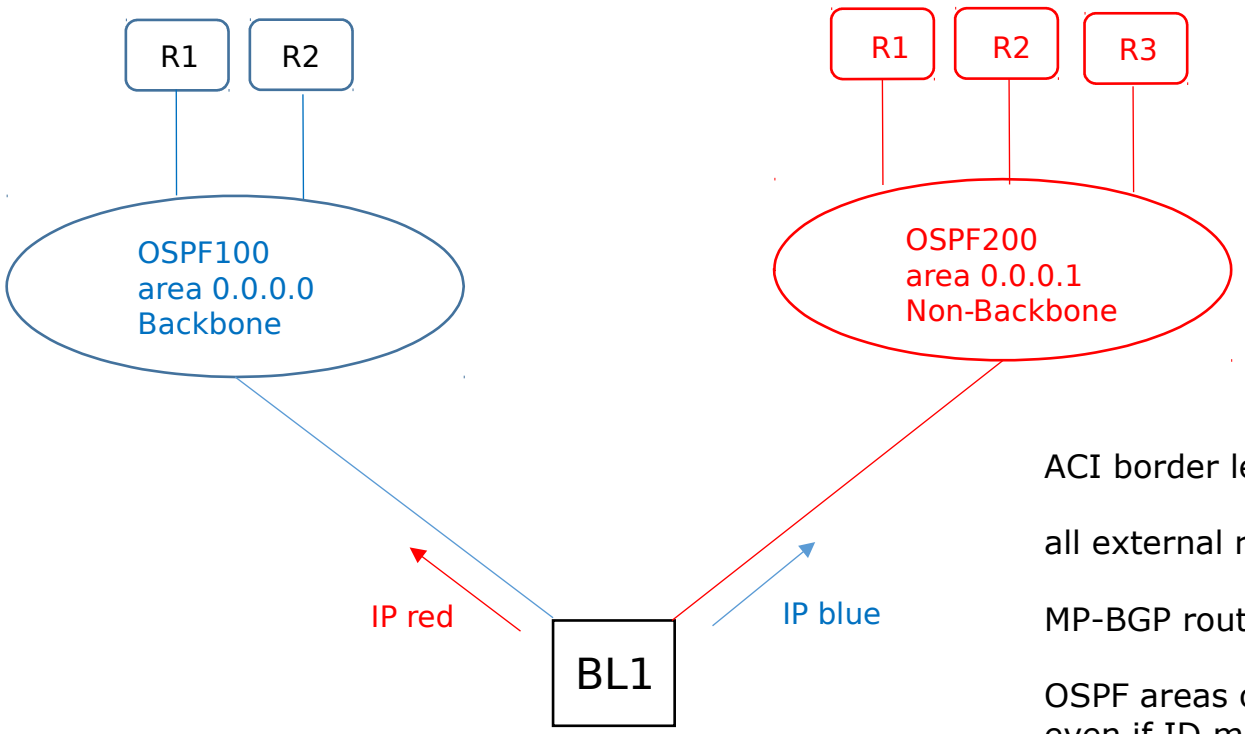

ACI border leaf running OSPF are always AS boundary (ASBR)

all external routes learned in OSPF are redistribuite into MP-BGP

MP-BGP routes are redistribuite into OSPF as external type2

OSPF areas on different border leaf (pairs BL) are different OSPF areas, even if ID match

IPv4 and IPv6 support

ACI Border Leaf follow OSPF rules which as:

multiple areas but NO backbone (both areas)area the routes are not advertised between areas;

No backbone are and backbone area are advertised between them

OSPF summarization rules

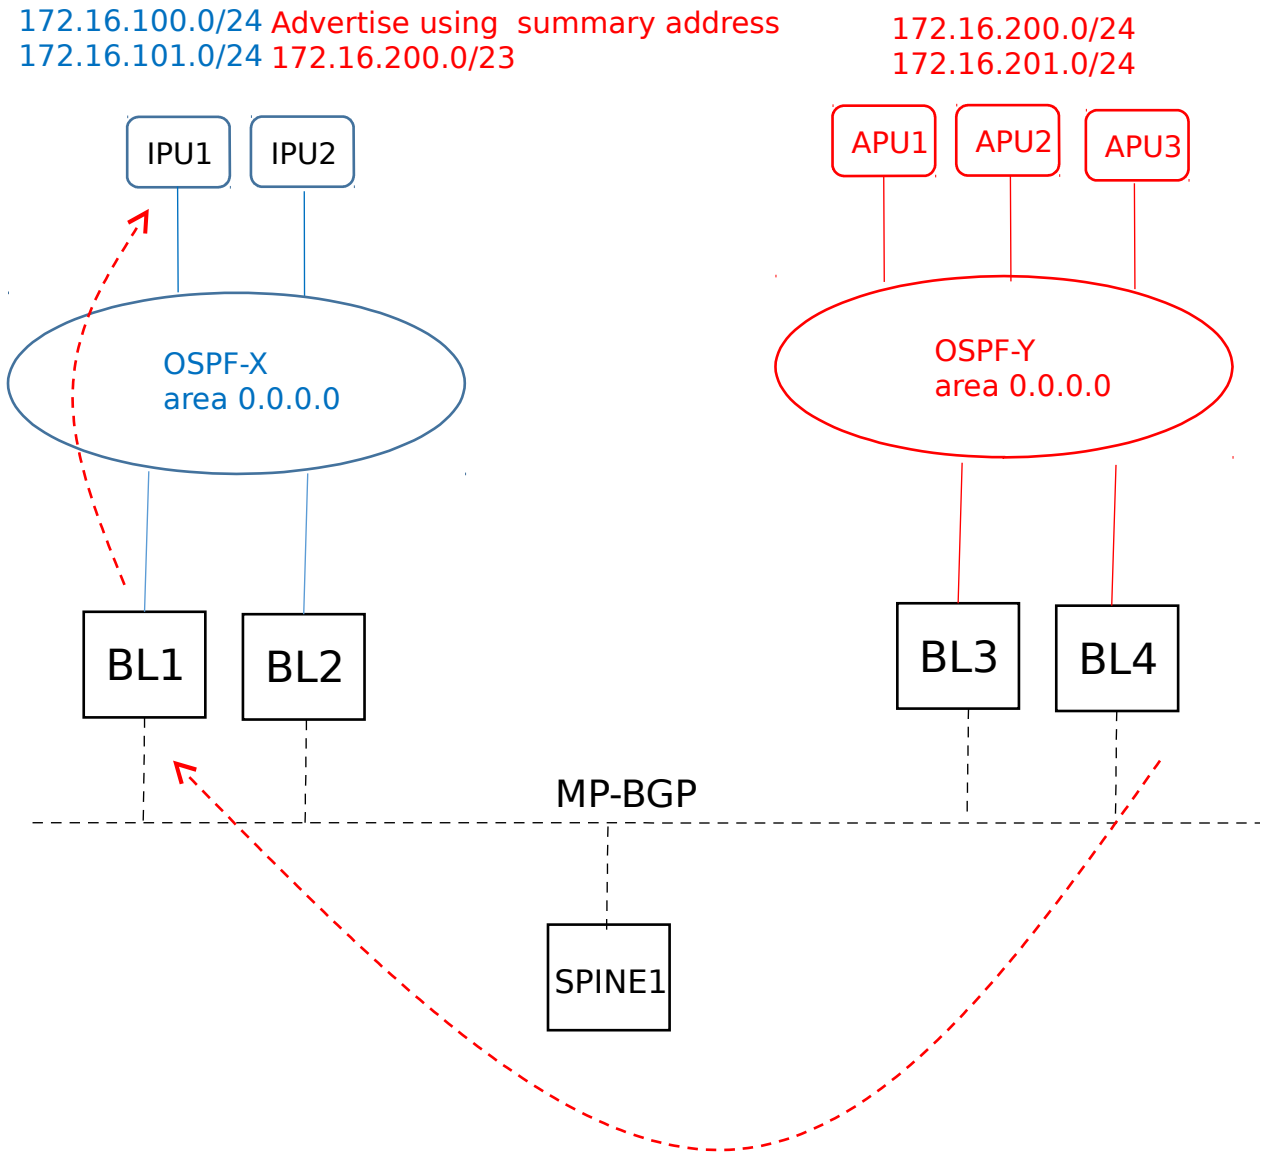

Two otions are available with ACI:

External route summarization (equivalent to the summary address config)

Inter-area summarization (equivalent to the area range config)

When Tenant routes are injected into OSPF, ACI Leaf where L3-out connection resides is acting as an ASBR; in this case the summary address config (thats is external route sumarization) should be used.

For scenario where there are two L3-out connection and each using a different area and attached to the same border leaf switch, the area range config will be used to summarize.

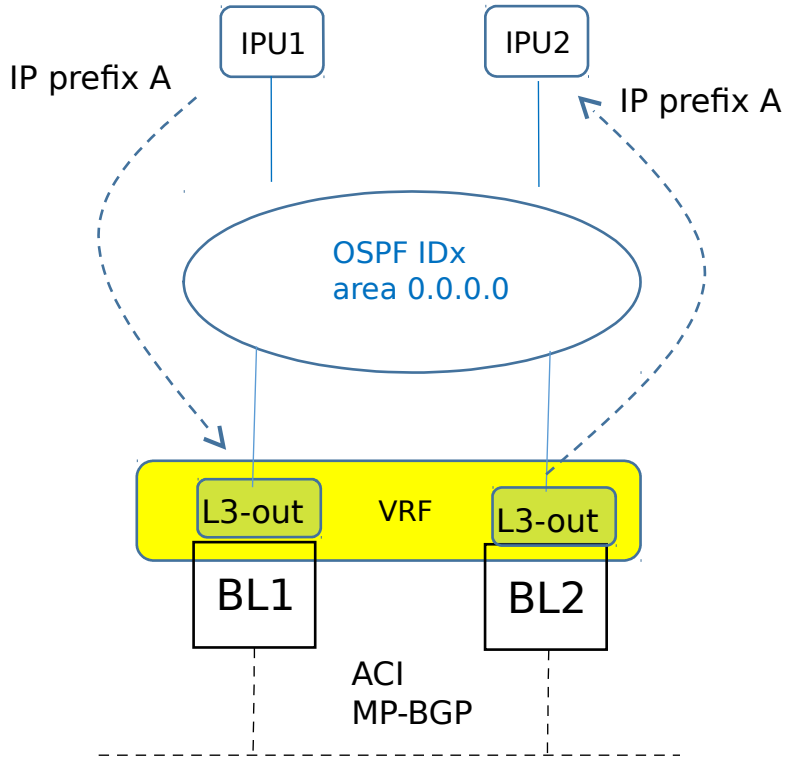
IPU1
IPU2
IP prefix A
IP prefix A
OSPF IDx
area 0.0.0.0
L3-out
VRF
L3-out
BL1
BL2
ACI
MP-BGP