Deploying ACI as Application Centric

Contents

Introduction

This document describes the approach to achieve micro-segmentation and security within/between the applications leveraging the Cisco ACI SDN solution.

Constraints using Traditional Network

- In traditional networks, segmentation within a VLAN/Subnet is impossible.
- The application gateways are on core switches. If two applications want to communicate, then complex Access Control Lists (ACLs) are required on the core switch.
- The Spanning-Tree loop between the switches breaks the data center flow and results in a traffic drop.
- The same IP Subnet contains multiple applications, which does not provide security between them. Managing these communications is not possible on traditional networks.
- Consider an example that is also depicted using the diagram. You have three applications EP_MB, EP_IB, and EP_UPI which are part of the same VLAN and IP subnet. With any L2 traffic, the traffic is always flooded to all the applications even though communication between them is not required. The restrictions between the two applications are not possible in this scenario.

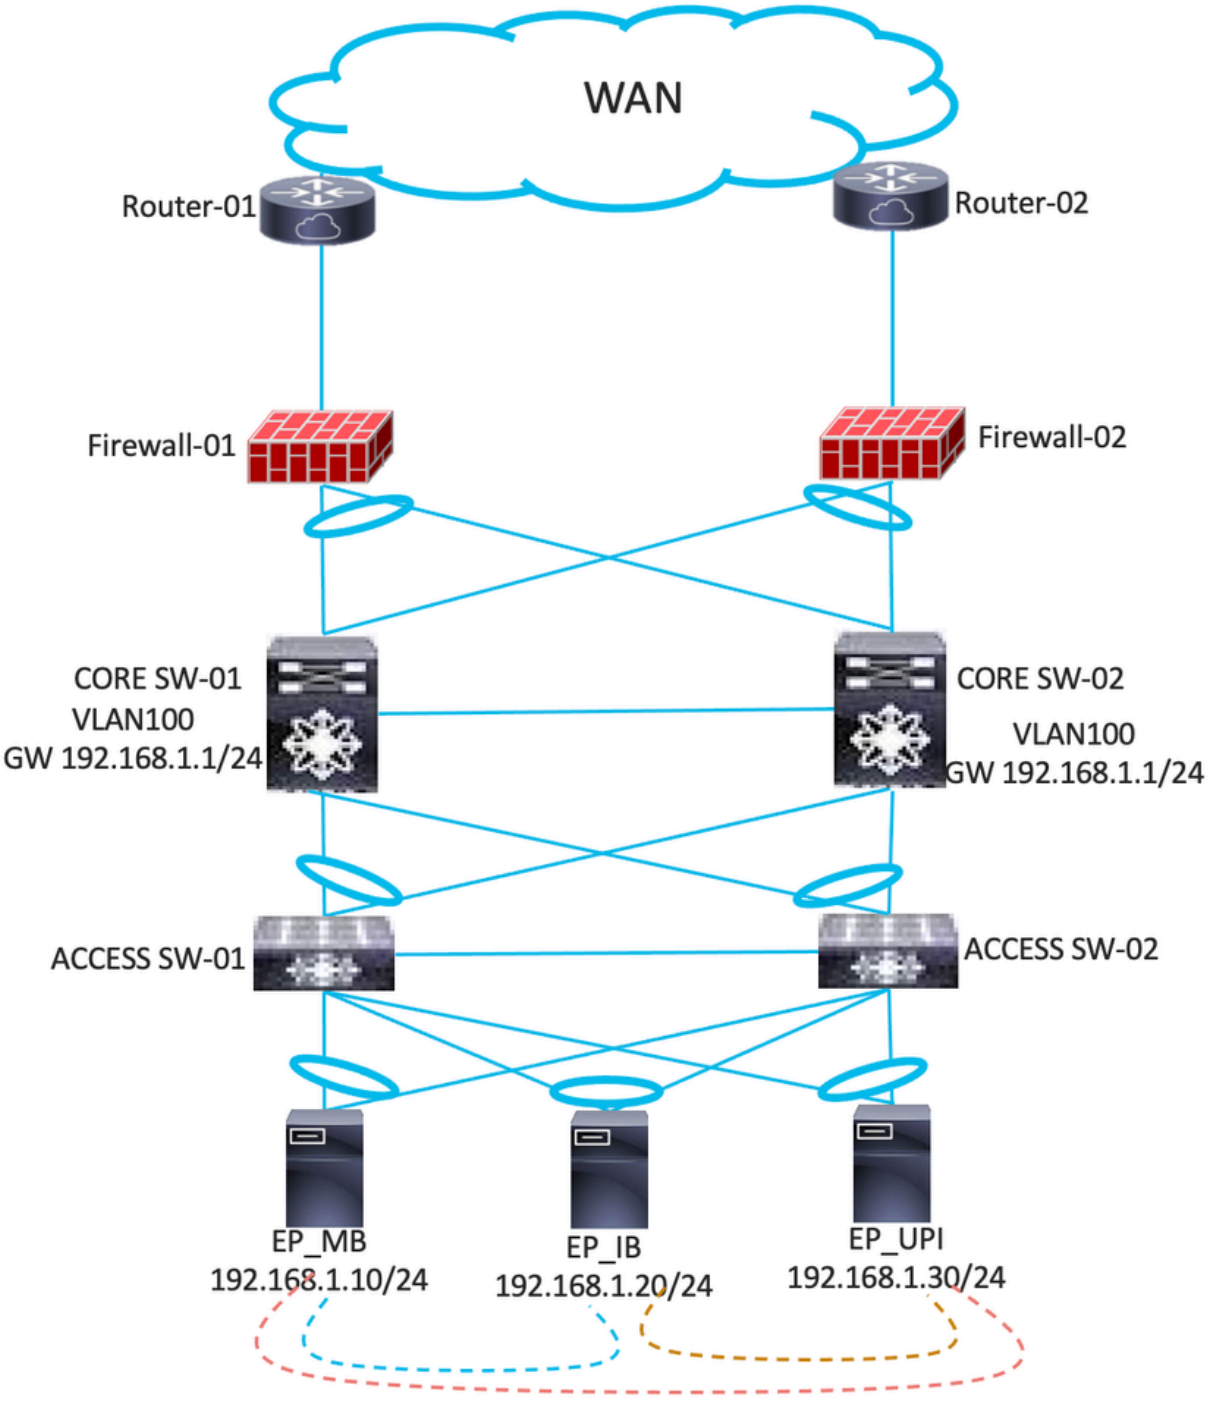

Prerequisites

Requirements

Cisco recommends that you have knowledge of these topics:

- Cisco Secure Workload (CSW)/Tetration (Secure Workload) must be deployed in the environment in order to gather the traffic flow data between the applications.
- Agents must be deployed on the servers in order to gather the data. Hence, this is only possible in the case of brownfield deployment.
- The agents must be deployed on the servers for at least 3-4 weeks for the data collection.
- If any Application Dependency Mapping (ADM) tools are unavailable, then the relevant data must be

provided.
- Server Gateway must be configured using the Application Centric Infrastructure (ACI) Fabric.

Components Used

This document is not restricted to specific software and hardware versions.

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

Solution Overview

In order to achieve micro-segmentation, you must first migrate the network to a Cisco SDN solution from traditional infrastructure and re-design the network from an application-centric view. This section describes the two phases of design in order to achieve the segmentation as desired based on the application flow which is captured via the ADM tool. At first, the Cisco ACI solution is deployed in the Network Centric Mode (as-is to existing design) and then moved towards application-centric mode.

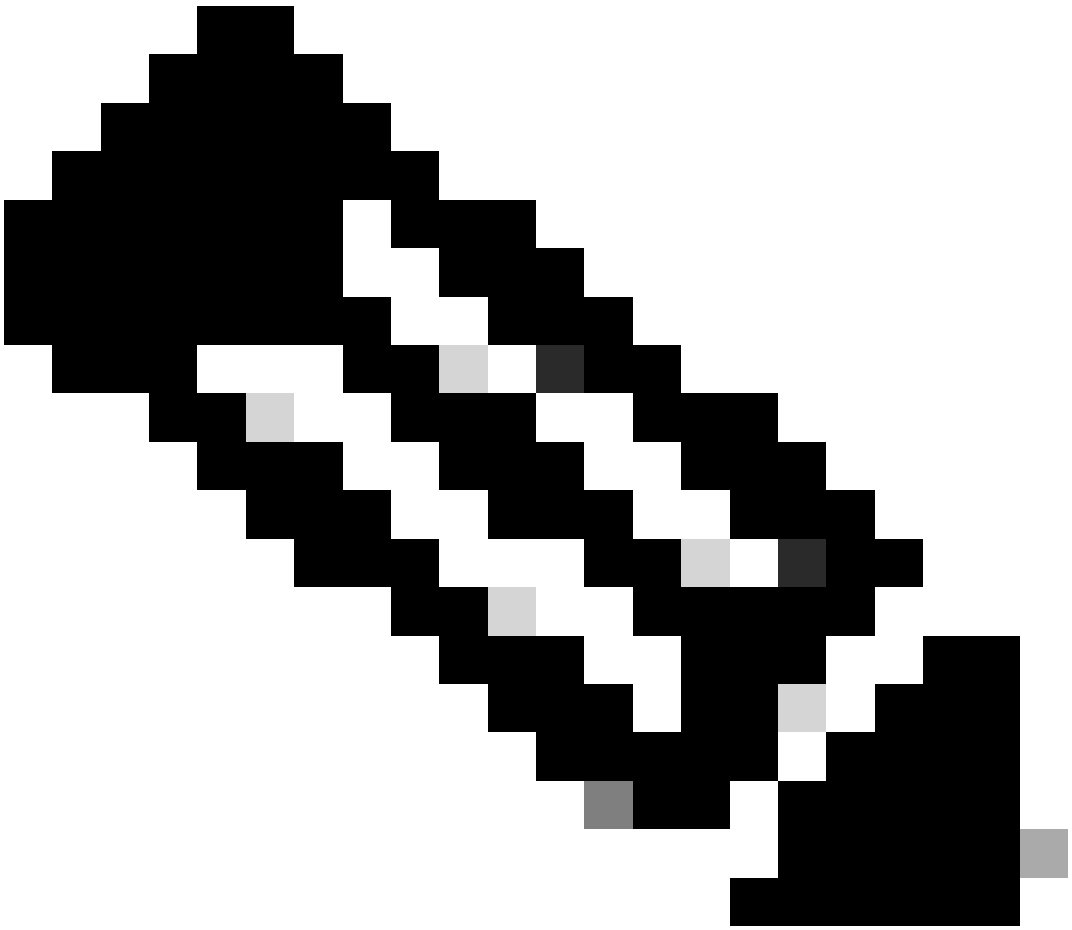

Note: You can also combine this deployment mode together in order to directly migrate services

---

from the traditional network to the application-centric mode.

Network Centric Design

In the example shown in the diagram, EPG_VL-100 contains three applications, EP_MB, EP_IB, and EP_UPI, and shares the same IP subnet and uses VLAN 100.

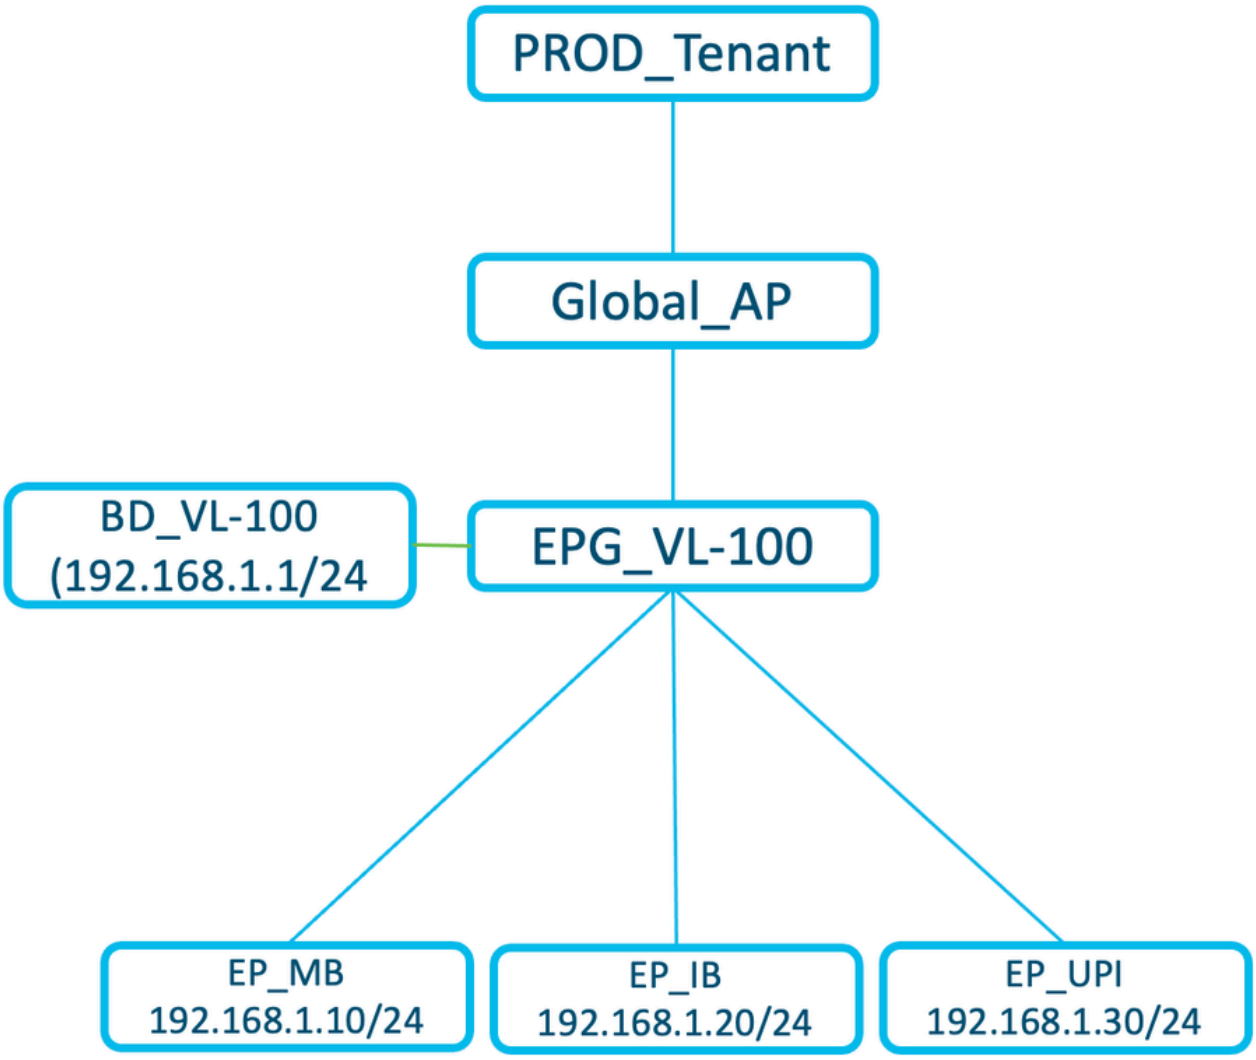

- As-Is migration from traditional network to ACI.
- One Endpoint Group (EPG) can contain multiple applications.
- No application segmentation within the same EPG in this deployment type.
- 1 BD = 1 EPG = 1 VLAN

Application Centric Design

The example shown in the diagram is a separate EPG for three applications EP_MB, EP_IB, and EP_UPI sharing the same IP subnet and using different VLANs mapped to each EPG.

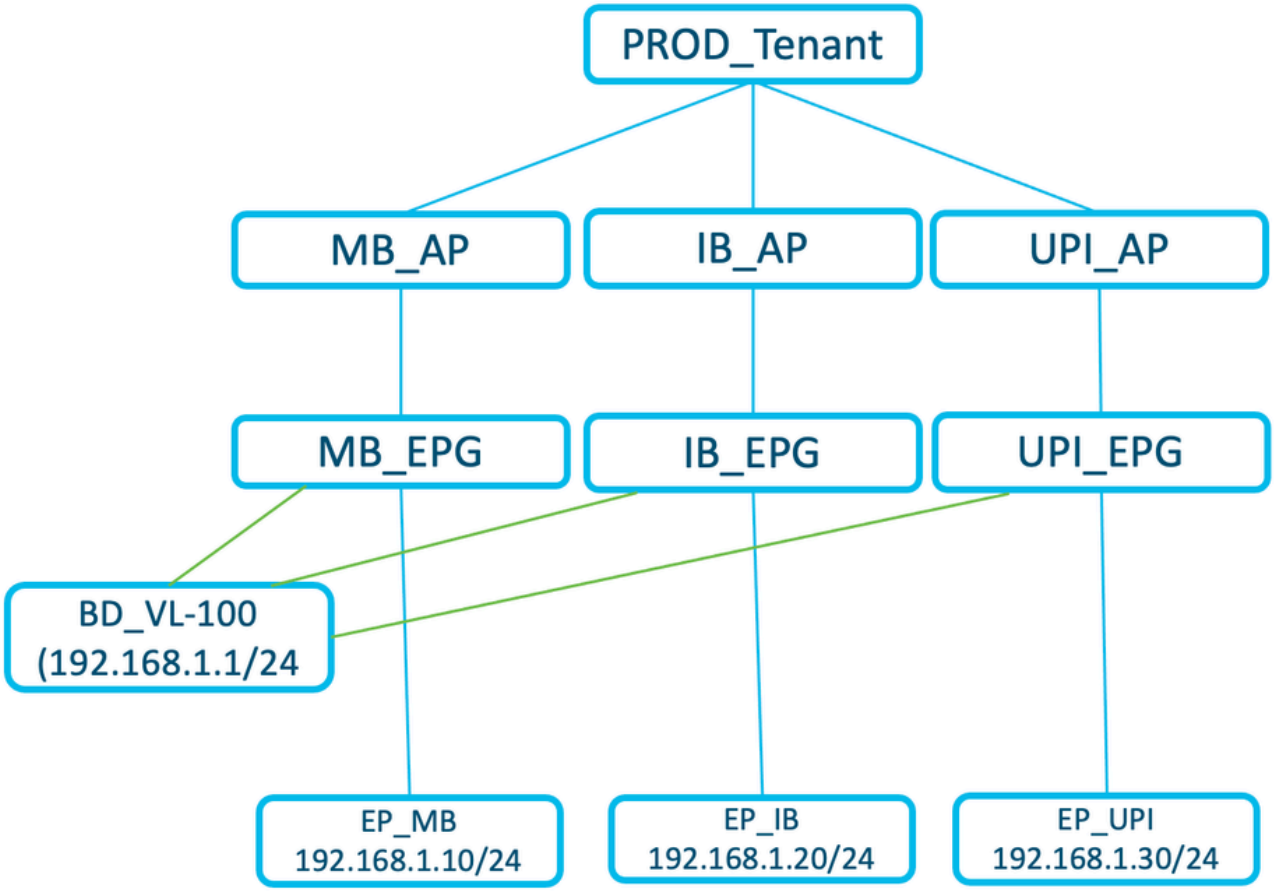

- In the Application-Centric deployment type, different EPGs are configured as per the application.
- The applications continue using the same IP subnet and its gateway.
- The segmented application EPGs to use a new VLAN.
- 1 BD to be configured with IP subnet and mapped to multiple application EPGs.
- 1 BD = N EPG = N VLAN
- Now two EPGs (applications) can communicate with each other via Contract.

Migration Approaches

Before deploying ACI as Application-centric, ACI can be deployed as Network-centric and further, the applications can be segmented.

Network Centric Migration Approach: Phase-1

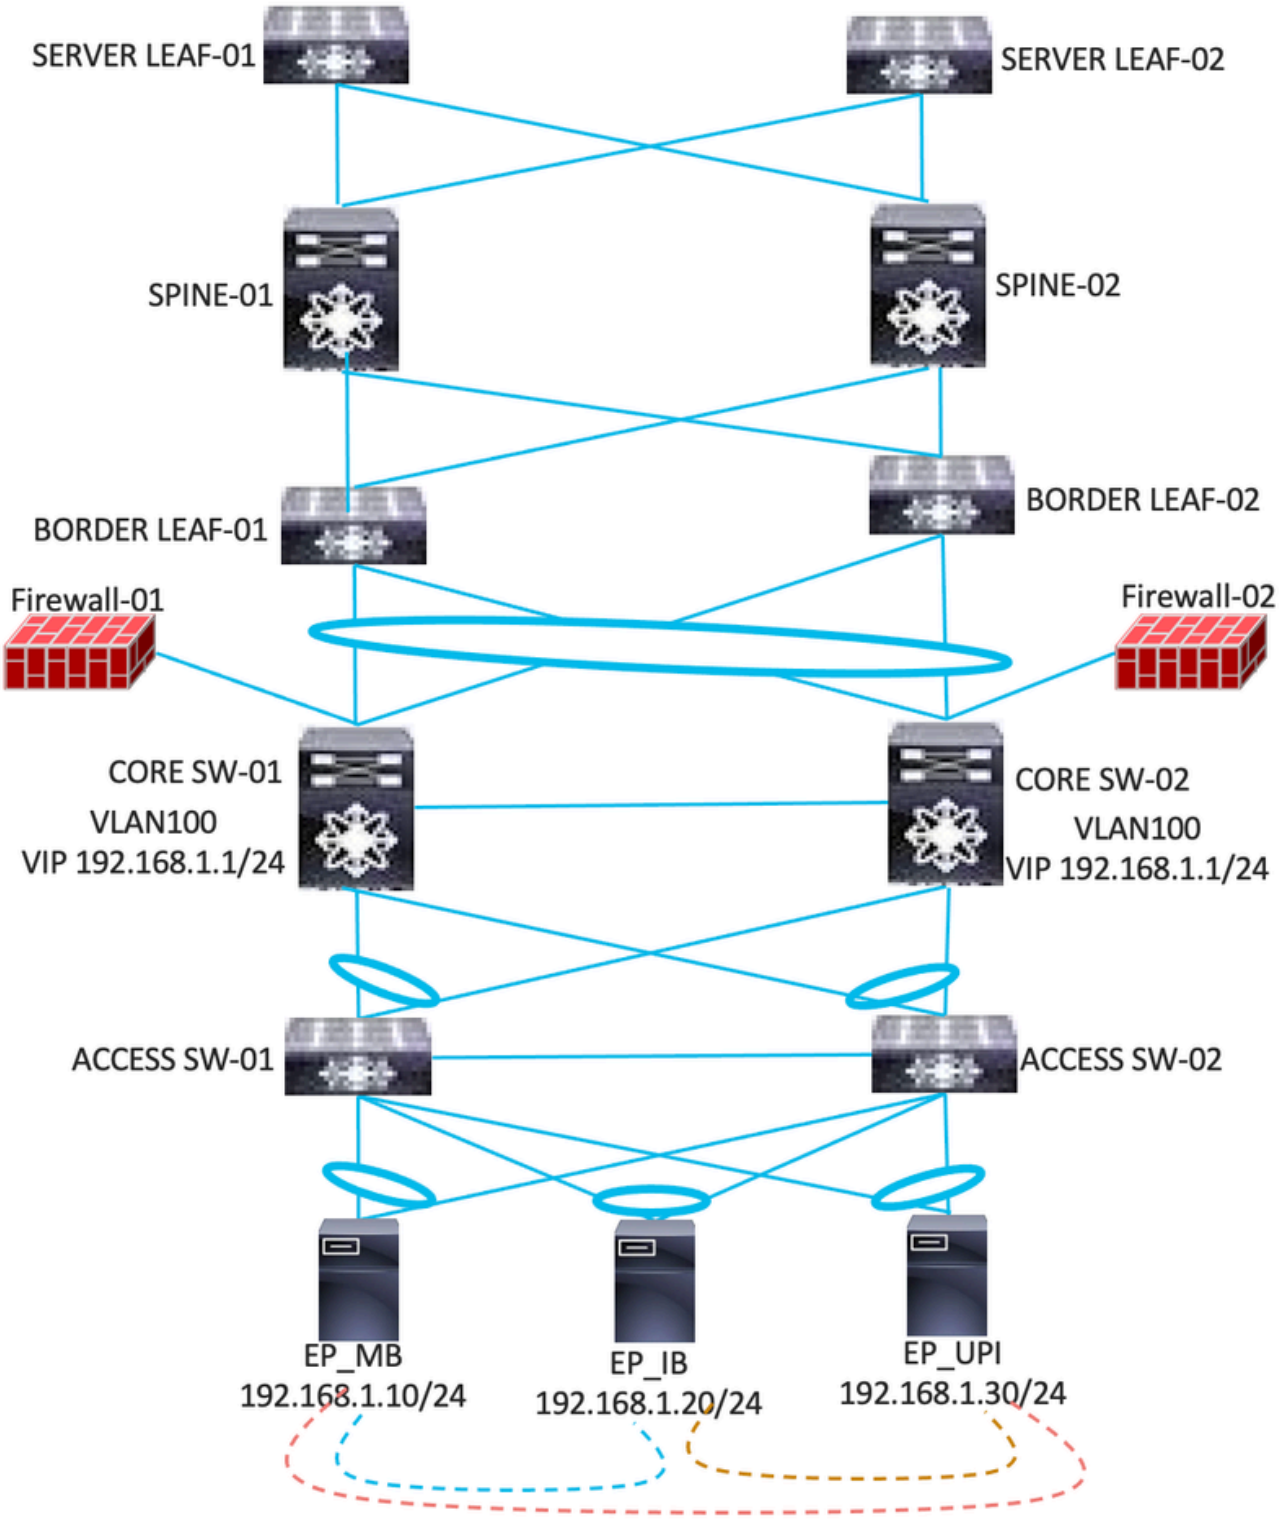

- Layer-2 interim link must be established between Border Leaf and Core Switches.
- Configure Layer-2 Bridge Domain and Endpoint Group on ACI according to the existing VLANs configured in traditional networks.
- Configure all these VLANs on the Layer-2 interim link between Border Leaf and Core Switches.
- ACI must be learning all the endpoints that are present on core switches.
- The Gateway remains on the core switches.
- Firewall connectivity remains on Core Switches.

Network Centric Migration Approach: Phase-2

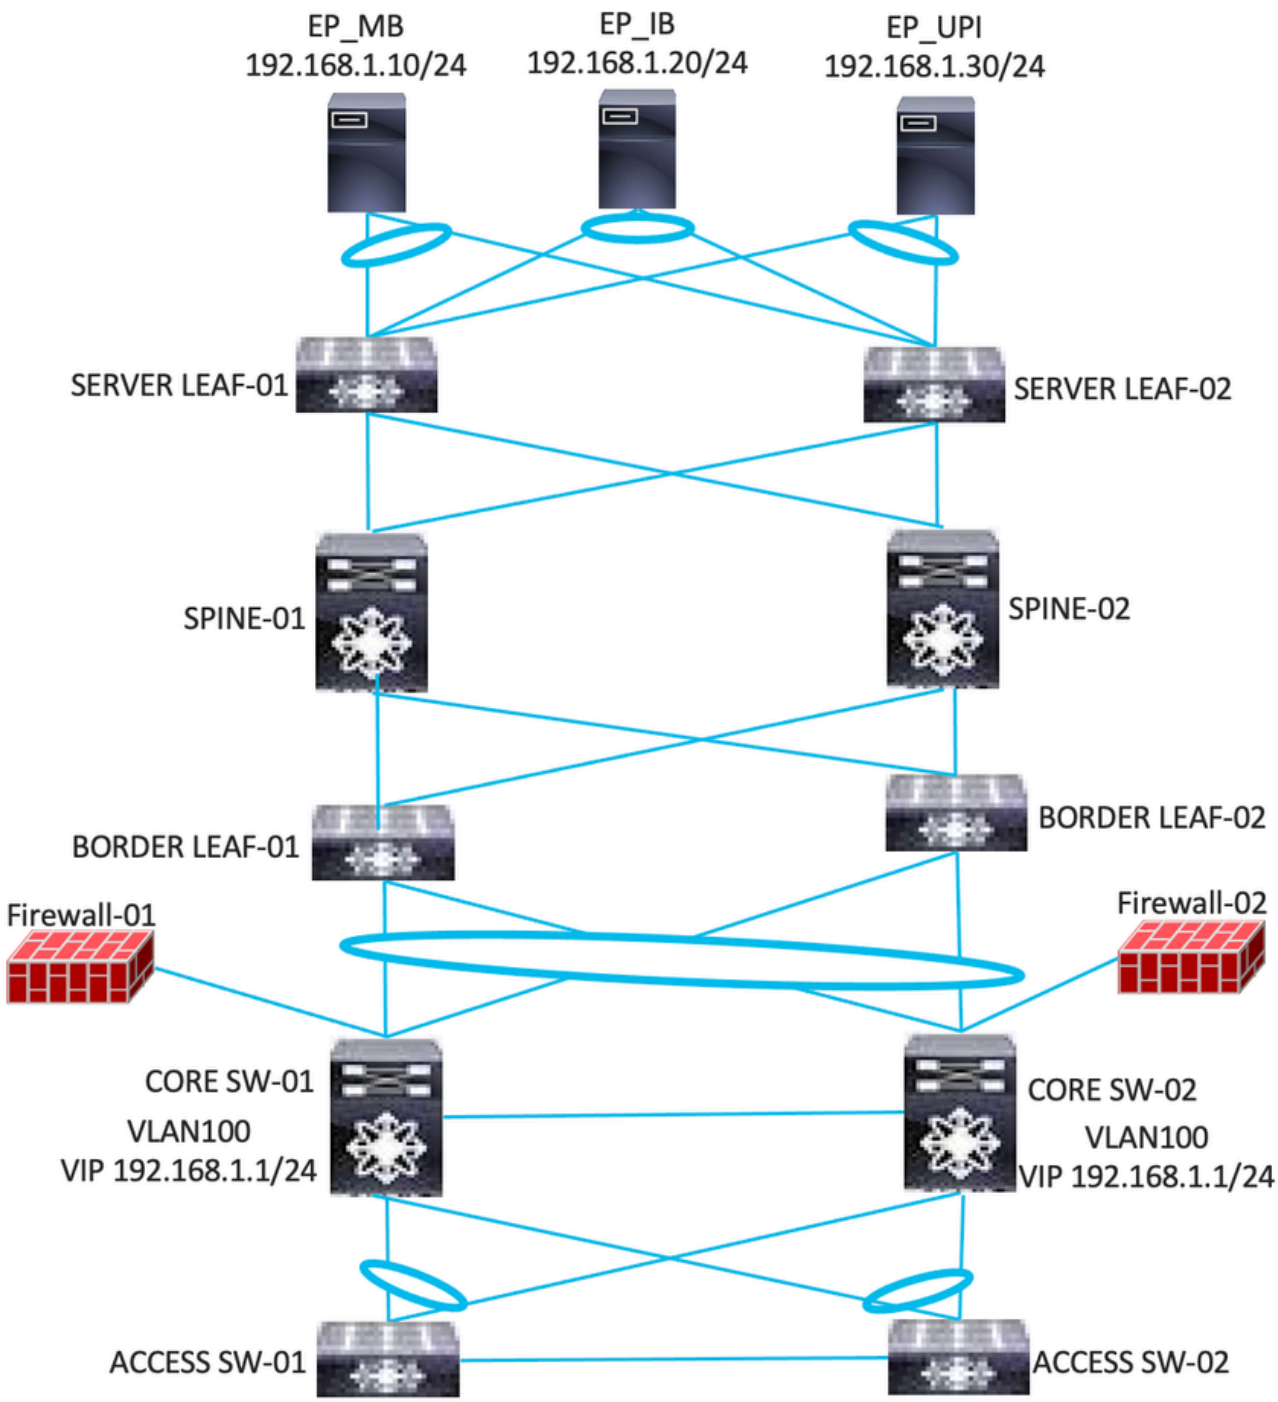

- Shift the workloads from Access Switches to Server Leaf.
- Gateway remains on Core Switches.
- Verify the Gateway is reachable from Servers.
- Verify the Server/Application is reachable.

Network Centric Migration Approach: Phase-3

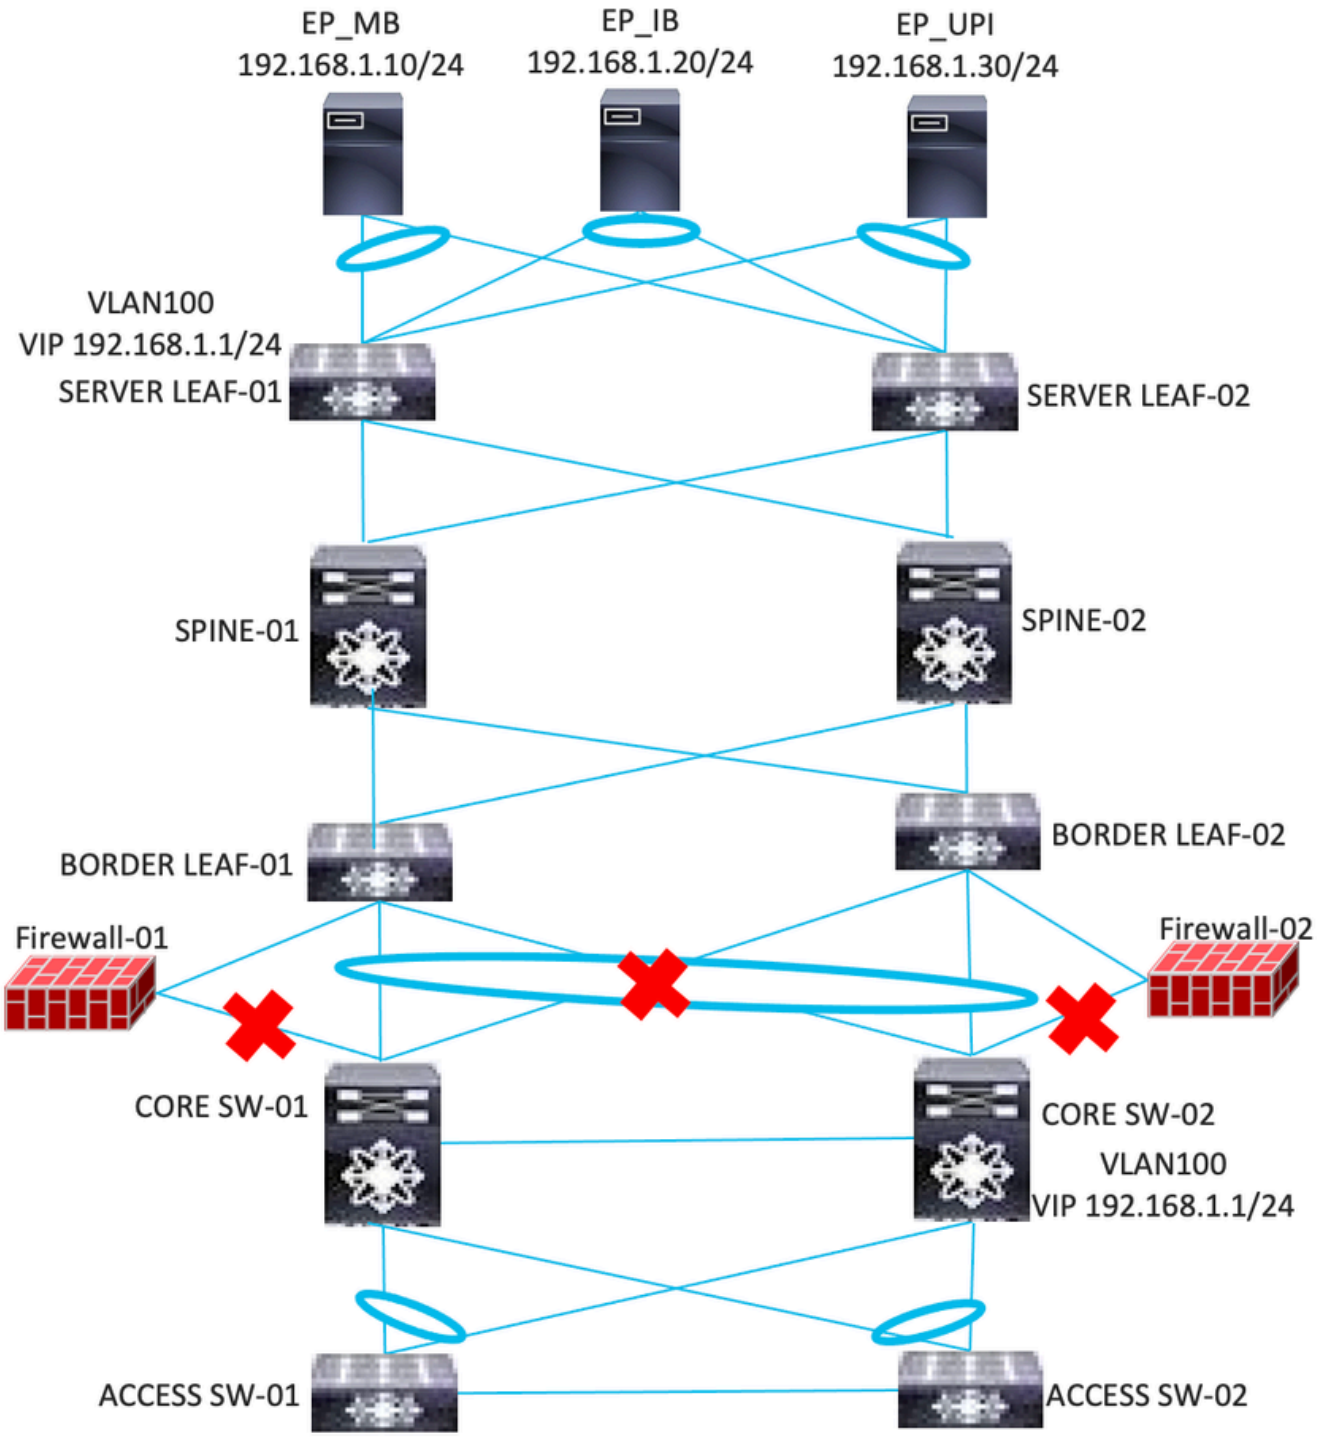

- Shut the Gateways on Core Switches and configure on ACI.
- Shift the Firewall link from Core Switches to ACI Leaf.
- Configure the L3out towards the Firewall/Router.
- Add the routes in the Firewall/Router and ACI Leaf.
- Shut down the link between Border Leaf and Core Switches.
- Verify the Server/Application is reachable.

Logical representation of ACI after Network Centric Migration approach.

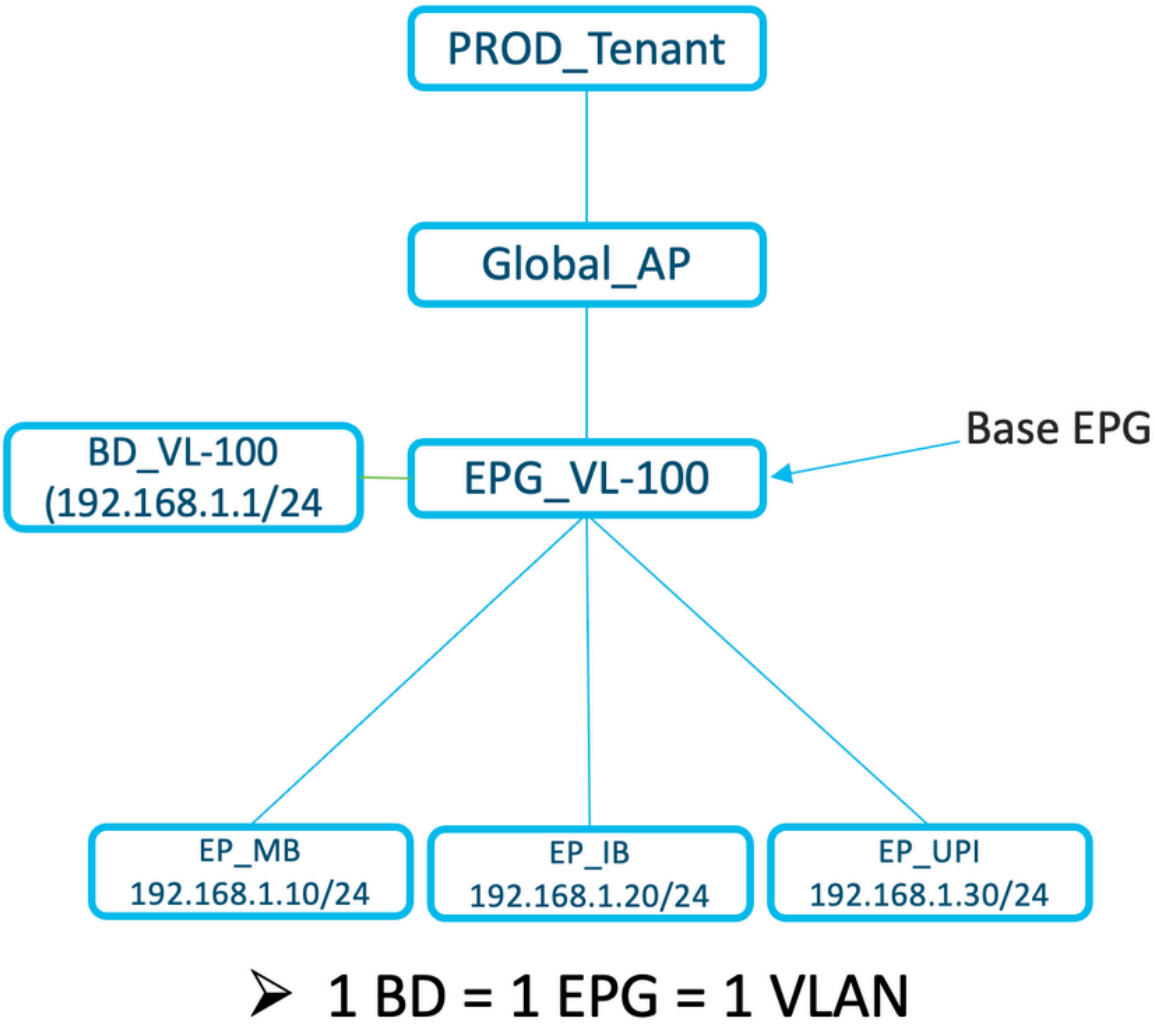

Application-Centric Migration Approach: Phase-1

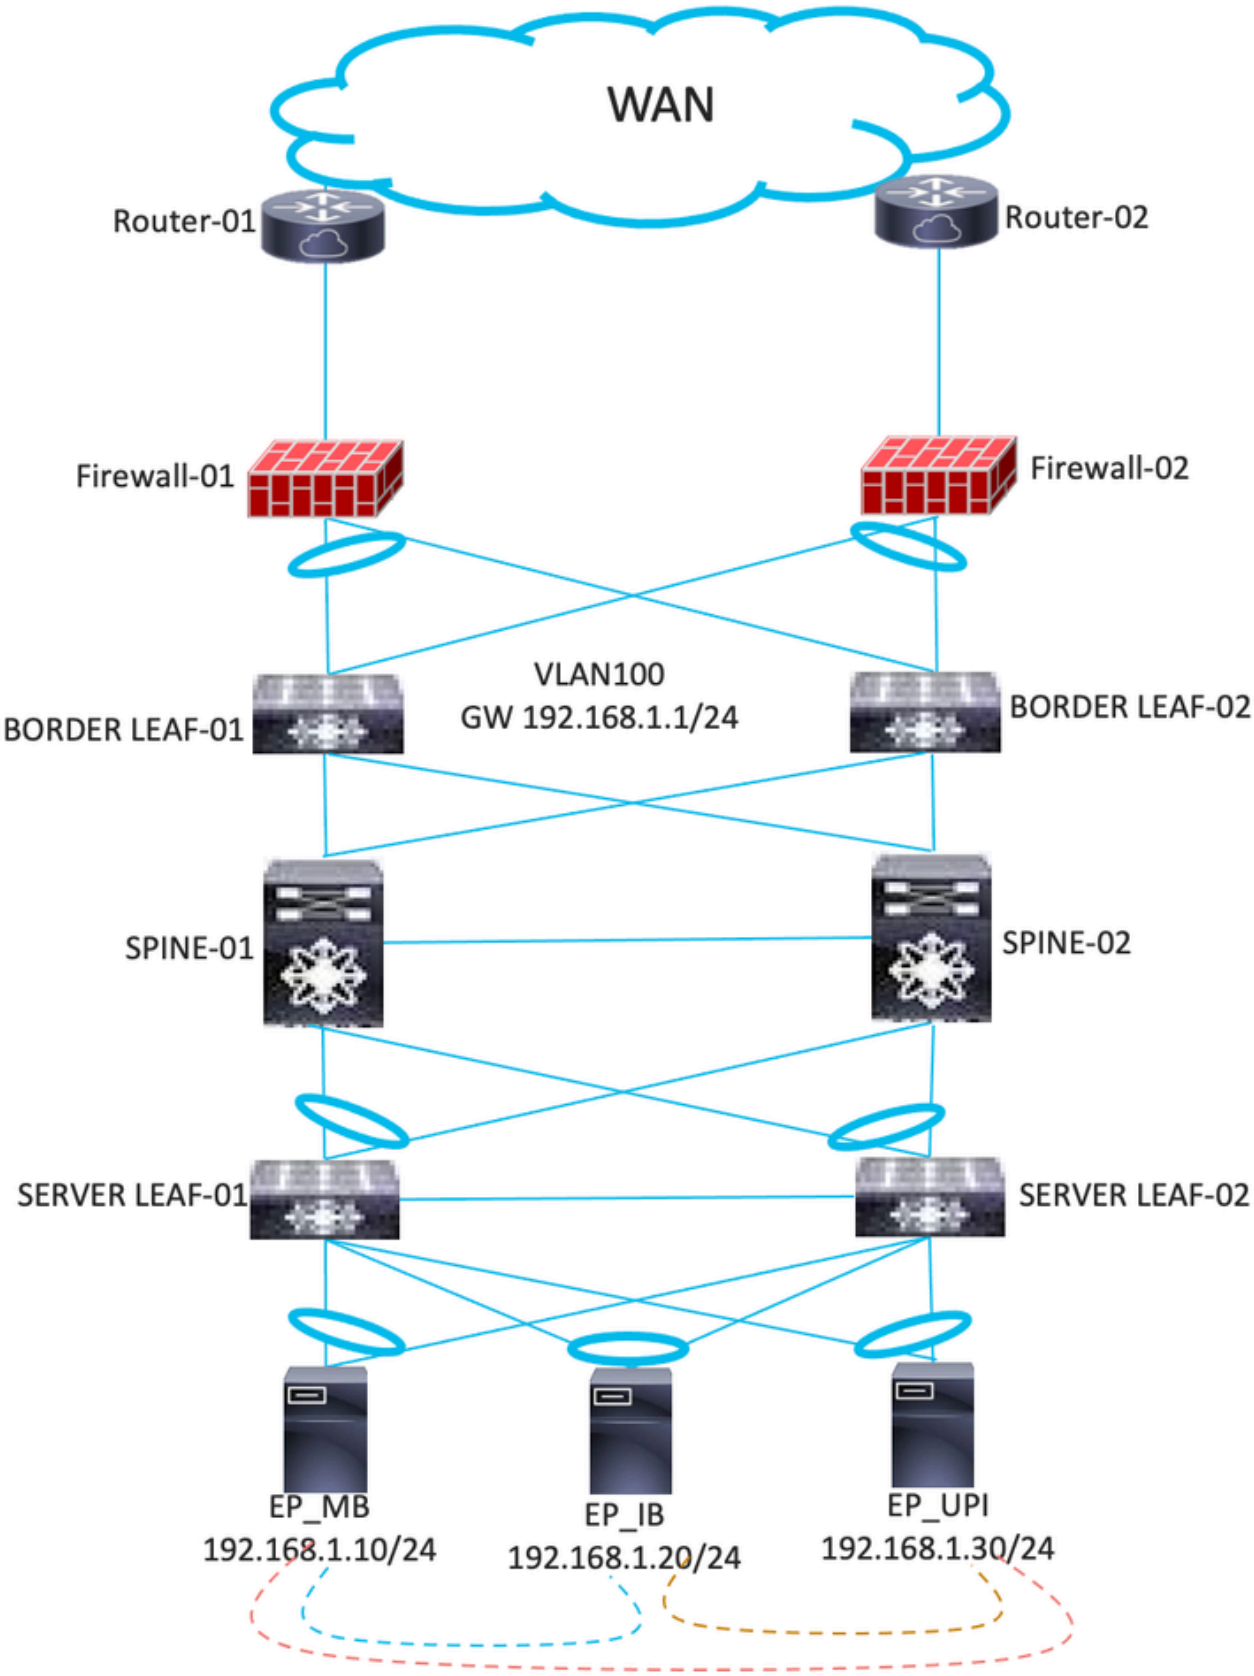

- Collection and analysis of CSW/Tetration Data.
- New EPG configuration as per CSW/Tetration Data (WEB, APP, and DB).
- For example, for the MB application, three EPGs are created such as EPG_MB_WEB, EPG_MB_APP, and EPG_MB_DB. These EPGs must be configured under one Application Profile AP_MB.

- In the case of Virtual Machine Manager (VMM) integration, vDS configuration is required for mapping the servers in the new EPG with the new VLAN.
- Map the Virtual Machine (VM) to the new vDS which is pushed through VMM integration.
- For baremetals, the server team must change the VLAN ID on the server.
- The IP addressing be the same for these deployments.
- Contract Configuration between EPGs as per CSW/Tetration data.

CSW/Tetration Data Analysis

Example of the analysis based on the CSW/Tetration data:

src_ip	consumer_scope	dst_ip	provider_scope	protocc
192.168.34.248	Default:Internal:Headquarter	192.168.20.81	PRODAPP	TCP
192.168.78.45	Default:Internal:Headquarter	192.168.20.81	PRODAPP	TCP
192.168.78.16	Default:Internal:Headquarter	192.168.20.81	PRODAPP	TCP
192.168.78.25	Default:Internal:Headquarter	192.168.20.81	PRODAPP	TCP
192.168.44.69	Default:Internal:Datacenter:DC:Application:Prod:Discovery	192.168.20.81	PRODAPP	UDP
192.168.44.69	Default:Internal:Datacenter:DC:Application:Prod:Discovery	192.168.20.81	PRODAPP	TCP
192.168.32.173	Default:Internal:Datacenter:DC:Application:Prod:DMZ	192.168.20.81	PRODAPP	TCP
192.168.44.47	Default:Internal:Datacenter:DC:Application:Prod:Monitoring	192.168.20.81	PRODAPP	TCP
192.168.44.47	Default:Internal:Datacenter:DC:Application:Prod:Monitoring	192.168.20.81	PRODAPP	UDP
192.168.44.48	Default:Internal:Datacenter:DC:Application:Prod:Monitoring	192.168.20.81	PRODAPP	UDP
192.168.44.47	Default:Internal:Datacenter:DC:Application:Prod:Monitoring	192.168.20.81	PRODAPP	TCP
192.168.44.47	Default:Internal:Datacenter:DC:Application:Prod:Monitoring	192.168.20.81	PRODAPP	TCP
192.168.44.48	Default:Internal:Datacenter:DC:Application:Prod:Monitoring	192.168.20.81	PRODAPP	TCP

192.168.44.47	Default:Internal:Datacenter:DC:Application:Prod:Monitoring	192.168.20.81	PRODAPP	TCP
192.168.44.47	Default:Internal:Datacenter:DC:Application:Prod:Monitoring	192.168.20.81	PRODAPP	TCP
192.168.44.47	Default:Internal:Datacenter:DC:Application:Prod:Monitoring	192.168.20.81	PRODAPP	TCP
192.168.44.29	Default:Internal:Datacenter:DC:Application:Prod:Monitoring	192.168.20.81	PRODAPP	TCP
192.168.44.30	Default:Internal:Datacenter:DC:Application:Prod:Monitoring	192.168.20.81	PRODAPP	TCP
192.168.44.21	Default:Internal:Datacenter:DC:Application:Prod:AAA	192.168.20.81	PRODAPP	ICMP
192.168.103.80	Default:Internal:Datacenter:DC:Application:Prod:DHCP	192.168.20.81	PRODAPP	TCP
192.168.103.71	Default:Internal:Datacenter:DC:Application:Prod:DHCP	192.168.20.81	PRODAPP	TCP
192.168.103.20	Default:Internal:Datacenter:DC:Application:Prod:DHCP	192.168.20.81	PRODAPP	TCP
192.168.103.21	Default:Internal:Datacenter:DC:Application:Prod:DHCP	192.168.20.81	PRODAPP	TCP
192.168.44.68	Default:Internal:Datacenter:DC:Application:Prod:Discovery	192.168.20.85	PRODDB	UDP
192.168.44.69	Default:Internal:Datacenter:DC:Application:Prod:Discovery	192.168.20.85	PRODDB	UDP
192.168.44.68	Default:Internal:Datacenter:DC:Application:Prod:Discovery	192.168.20.85	PRODDB	TCP
192.168.44.69	Default:Internal:Datacenter:DC:Application:Prod:Discovery	192.168.20.85	PRODDB	TCP
172.16.32.173	Default:Internal:Datacenter:DC:Application:Prod:MZ	192.168.20.85	PRODDB	TCP
192.168.44.47	Default:Internal:Datacenter:DC:Application:Prod:Monitoring	192.168.20.85	PRODDB	TCP
192.168.44.47	Default:Internal:Datacenter:DC:Application:Prod:Monitoring	192.168.20.85	PRODDB	UDP
192.168.44.48	Default:Internal:Datacenter:DC:Application:Prod:Monitoring	192.168.20.85	PRODDB	UDP
192.168.44.47	Default:Internal:Datacenter:DC:Application:Prod:Monitoring	192.168.20.85	PRODDB	UDP

192.168.44.47	Default:Internal:Datacenter:DC:Application:Prod:Monitoring	192.168.20.85	PRODDB	TCP
192.168.44.48	Default:Internal:Datacenter:DC:Application:Prod:Monitoring	192.168.20.85	PRODDB	TCP
192.168.44.47	Default:Internal:Datacenter:DC:Application:Prod:Monitoring	192.168.20.85	PRODDB	TCP
192.168.44.47	Default:Internal:Datacenter:DC:Application:Prod:Monitoring	192.168.20.85	PRODDB	TCP
192.168.44.47	Default:Internal:Datacenter:DC:Application:Prod:Monitoring	192.168.20.85	PRODDB	TCP
192.168.44.30	Default:Internal:Datacenter:DC:Application:Prod:Monitoring	192.168.20.85	PRODDB	TCP
192.168.44.29	Default:Internal:Datacenter:DC:Application:Prod:Monitoring	192.168.20.85	PRODDB	TCP
192.168.44.21	Default:Internal:Datacenter:DC:Application:Prod:Monitoring	192.168.20.85	PRODDB	ICMP

Example of EPG recommendation from the CSW/Tetration:

EPG	IP
PRODAPP	192.168.20.81
RODDB	192.168.20.85

Based on the details, the data must be analyzed for contract configuration. Example of analyzed data:

src_ip	consumer_scope	consumer_EPG	dst_IP	provider_EPG	protocol	port
192.168.44.69	Default:Internal:Datacenter: DC:Application:Prod:Disocvery	EPG_DISCOVERY	192.168.20.81	EPG-PROD-APP	UDP	137
192.168.44.69	Default:Internal:Datacenter: DC:Application:Prod:Disocvery	EPG_DISCOVERY	192.168.20.81	EPG-PROD-APP	TCP	445
192.168.44.47	Default:Internal:Datacenter: DC:Application:Prod:Monitoring	EPG_MONITORING	192.168.20.81	EPG-PROD-APP	TCP	135
192.168.44.47	Default:Internal:Datacenter:	EPG_MONITORING	192.168.20.81	EPG-PROD-	UDP	137

	DC:Application:Prod:Monitoring			APP		
192.168.44.48	Default:Internal:Datacenter: DC:Application:Prod:Monitoring	EPG_MONITORING	192.168.20.81	EPG-PROD-APP	TCP	443
192.168.44.47	Default:Internal:Datacenter: DC:Application:Prod:Monitoring	EPG_MONITORING	192.168.20.81	EPG-PROD-APP	TCP	445
192.168.44.47	Default:Internal:Datacenter: DC:Application:Prod:Monitoring	EPG_MONITORING	192.168.20.81	EPG-PROD-APP	TCP	5985
192.168.44.47	Default:Internal:Datacenter: DC:Application:Prod:Monitoring	EPG_MONITORING	192.168.20.81	EPG-PROD-APP	TCP	4915
192.168.44.47	Default:Internal:Datacenter: DC:Application:Prod:Monitoring	EPG_MONITORING	192.168.20.81	EPG-PROD-APP	TCP	4916
192.168.44.48	Default:Internal:Datacenter: DC:Application:Prod:Monitoring	EPG_MONITORING	192.168.20.81	EPG-PROD-APP	TCP	4750
192.168.44.47	Default:Internal:Datacenter: DC:Application:Prod:Monitoring	EPG_MONITORING	192.168.20.81	EPG-PROD-APP	ICMP	0
192.168.103.21	Default:Internal:Datacenter: DC:Application:Prod:DHCP	EPG_VL_157	192.168.20.81	EPG-PROD-APP	TCP	7777
192.168.44.68	Default:Internal:Datacenter: DC:Application:Prod:Disocvery	EPG_DISCOVERY	192.168.20.85	EPG-PROD-DB	UDP	137
192.168.44.68	Default:Internal:Datacenter: DC:Application:Prod:Disocvery	EPG_DISCOVERY	192.168.20.85	EPG-PROD-DB	TCP	445
192.168.44.69	Default:Internal:Datacenter: DC:Application:Prod:Monitoring	EPG_MONITORING	192.168.20.85	EPG-PROD-DB	TCP	135
192.168.44.69	Default:Internal:Datacenter: DC:Application:Prod:Monitoring	EPG_MONITORING	192.168.20.85	EPG-PROD-DB	UDP	137
192.168.44.47	Default:Internal:Datacenter: DC:Application:Prod:Monitoring	EPG_MONITORING	192.168.20.85	EPG-PROD-DB	UDP	161

192.168.44.47	Default:Internal:Datacenter: DC:Application:Prod:Monitoring	EPG_MONITORING	192.168.20.85	EPG-PROD-DB	TCP	445
192.168.44.48	Default:Internal:Datacenter: DC:Application:Prod:Monitoring	EPG_MONITORING	192.168.20.85	EPG-PROD-DB	TCP	5985
192.168.44.47	Default:Internal:Datacenter: DC:Application:Prod:Monitoring	EPG_MONITORING	192.168.20.85	EPG-PROD-DB	TCP	4915
192.168.44.47	Default:Internal:Datacenter: DC:Application:Prod:Monitoring	EPG_MONITORING	192.168.20.85	EPG-PROD-DB	TCP	6080
192.168.44.48	Default:Internal:Datacenter: DC:Application:Prod:Monitoring	EPG_MONITORING	192.168.20.85	EPG-PROD-DB	TCP	4750
192.168.44.47	Default:Internal:Datacenter: DC:Application:Prod:Monitoring	EPG_MONITORING	192.168.20.85	EPG-PROD-DB	ICMP	0
192.168.48.45	Default:Internal:Datacenter: DC:Application:Prod:Backup	EPG_VL_71	192.168.20.85	EPG-PROD-DB	TCP	5555

Based on the IP address, the consumer and provider EPGs are mentioned. Duplicate entries and North-South traffic (such as the Internet, inter-DC, inter-zones traffic, and so on) must be excluded from this data. There are some EPGs named with VLANs such as EPG_VL_157, EPG_VL_71, and so on. This means these servers are not moved to the target EPG as part of Application-centric migration. So, the contract between them is to be configured with the current mapping of EPG. Once these servers are migrated to target EPG then these existing contracts must be deleted as part of the cleanup process and the appropriate contract must be added to target EPG.

Contract

Contracts are required for communication between the EPGs. Implementation flow during the contract configuration process is captured in this section.

1. Initially VzAny contract must be applied on the Virtual Routing and Forwarding (VRF) level.

2. According to CSW/Tetration data, specific EPG contracts must be created.

3. Configure the Deny_All rule with low priority so that the VzAny contract does not allow unspecified traffic communication. For the applications, that are not yet migrated as application-centric, the communication happens via VzAny Contract.

4. After all the migration, delete the VzAny contract from the VRF.

The analysis of CSW/Tetration data and converting it into appropriate ACI objects is a very critical step. Hence, after the initial analysis, it is important to discuss our observation with the concerned and get re-confirmation on the same. Also during the implementation, careful consideration must be done in order to

ensure that all traffic is allowed as expected. For troubleshooting, you can enable logging on the contract and also track for any packet drop on a specific port using a GUI interface or CLI.

leaf# **show logging ip access-list internal packet-log deny**
[Tue Oct 1 10:34:37 2019 377572 usecs]: CName: Prod1:VRF1(VXLAN: 2654209), VlanType: Unknown, Vlan-Id: 0, SMac: 0x000c0c0c0c0c, DMac:0x000c0c0c0c0c, SIP: 192.168.21.11, DIP: 192.168.22.11, SPort: 0, DPort: 0, Src Intf: Tunnel7, Proto: 1, PktLen: 98
[Tue Oct 1 10:34:36 2019 377731 usecs]: CName: Prod1:VRF1(VXLAN: 2654209), VlanType: Unknown, Vlan-Id: 0, SMac: 0x000c0c0c0c0c, DMac:0x000c0c0c0c0c, SIP: 192.168.21.11, DIP: 192.168.22.11, SPort: 0, DPort: 0, Src Intf: Tunnel7, Proto: 1, PktLen: 98

contract_parser

An on-device Python script that produces an output that correlates the zoning rules, filters and hit statistics while performing name lookups from IDs. This script is extremely useful in that it takes a multi-step process and turns it into a single command that can be filtered to specific EPGs/VRFs or other contract-related values.

leaf# **contract_parser.py**
Key:
[prio:RuleId] [vrf:{str}] action protocol src-epg [src-l4] dst-epg [dst-l4] [flags][contract:{str}] [hit=count]

[7:4131] [vrf:common:default] permit ip tcp tn-Prod1/ap-Services/epg-NTP(16410) tn-Prod1/l3out-L3Out1/instP-extEpg(25) eq 123 [contract:uni/tn-Prod1/brc-external_to_ntp] [hit=0]
[7:4156] [vrf:common:default] permit ip tcp tn-Prod1/l3out-L3Out1/instP-extEpg(25) eq 123 tn-Prod1/ap-Services/epg-NTP(16410) [contract:uni/tn-Prod1/brc-external_to_ntp] [hit=0]
[12:4169] [vrf:common:default] deny,log any tn-Prod1/l3out-L3Out1/instP-extEpg(25) epg:any [contract:implicit] [hit=0]
[16:4167] [vrf:common:default] permit any epg:any tn-Prod1/bd-Services(32789) [contract:implicit] [hit=0]

The packet drops can also be shown in GUI using the path: **Tenant > Tenant_Name > Operational > Flows/Packets**.

Consideration

Recommendation while applying the contracts between the EPGs:

1. ACI can not be considered a Firewall in terms of policy mapping which can cause high Ternary Content Addressable Memory (TCAM) utilization.

2. Use a range of filters instead of a large number of individual filters.

3. Any contract must not use more than four ranges of filters. It can consume high Overflow Ternary Content Addressable Memory (OTCAM).

4. If any EPGs require a large number of ports, try to use a 'permit any' contract.

5. As part of the solution, if you foresee the deployment of a large number of contracts, consider modifying the Forwarding Scale Profile (FSP) accordingly.

6. Before deploying a bulk number of contracts, calculate the TCAM using the formula: **No. of Provide EPG * No. of Consumer EPG * Number of rules**.

7. The existing TCAM size can be checked on ACI UI using the path: **Operations > Capacity Dashboard**

> Leaf Capacity or

LEAF-101# vsh_lc

module-1# show platform internal hal health-stats | grep _count

mcast_count : 0

max_mcast_count : 8192

policy_count : 221

max_policy_count : 65536

policy_**otcam**_count : **322**

max_policy_**otcam**_count : **8192**

policy_label_count : 0

max_policy_label_count : 0

Some Challenges of App-Centric Deployment and Solution

1. A larger number of contracts can lead to high TCAM utilization of leaf switches.

Hence, it is important to actively track the TCAM Utilization and also prepare an estimated increase in TCAM value when a large amount of configuration deployment is done. It is good to have a maker checker process in order to ensure the configuration being pushed is appropriate. Also, it is recommended to carry out the changes with a proper scheduled maintenance window.

2. Bulk configuration (more than 50k TCAM) in a single push of contract can lead to a Policy Manager memory crash.

It is recommended to push the configuration in smaller chunks especially when the configuration is large in size. This provides a systematic and risk-free approach towards contract configuration. Also, with each configuration push, measure the increase in TCAM values.

3. The traffic flow is not captured if the applications do not communicate during the CSW/Tetration deployment time interval (3-4 weeks).

In order to avoid such a situation, the best approach is to get the CSW/Tetration data reverified from the Application owners before the change activity. Also, post-implementation, verify the logs for any failure hit count.

Value Addition

1. All the applications have been segmented and restricted according to the Central Banking guidelines.

2. Visibility of inter-application communication after migrating to application-centric deployment.

3. Micro-segmentation of the application is achieved.

4. One view of application flow. In one Application Profile, the EPGs are mapped according to the traffic flow such as Application Profile AP_Banking, in order to have three EPGs

(EPG_Banking_WEB, EPG_Banking_APP, and EPG_Banking_DB) regardless of their IP subnet.

4. One view of application flow makes troubleshooting easier.

5. Infra is more secure.

6. Structured approach for implementation and future expansion.