CHECK POINT™

YOU DESERVE THE BEST SECURITY

Private Cloud Security for Cisco ACI Infrastructure

Rel 2.0

TABLE OF CONTENTS

Introduction

As companies are embarking on their application and data modernization programs and considering cloud and infrastructure requirements, they will most likely opt for a hybrid cloud strategy, with application and data workloads spread across both public and private clouds.

In hybrid deployments, hyper-scalers combine the cloud benefits of innovation, speed, consumption, and scale of the public cloud with the benefits of private clouds, such as regulatory compliance, performance, data gravity, and recouping existing investments. Furthermore, hybrid deployments provide the same level of operation and management in both public and private cloud environments, e.g., unified management, flexibility, agility.

Cisco Application Centric Infrastructure (ACI) is a mature SDN (Software Defined Network) technology that offers enterprises of all sizes "cloud-like" performance, availability, resilience, monitoring, and automation. Enterprises that want to build their own on-premise private clouds will find Cisco ACI provides most if not all the features they need to do so. The cloud-like features of Cisco ACI enable customers to leverage a fundamentally more secure approach to data and network security by moving to a security model independent of routing and network topology.

Check Point CloudGuard for Cisco ACI delivers industry-leading security management and enforcement tailored to protecting customer information assets. Security service insertion in modern, application-centric private and hybrid cloud networks is sophisticated, yet simple, way to design, deploy, scale and operate in a complex environment.

.

Cisco ACI Overview

Cisco ACI provides an open security policy framework that expresses policies using the language of applications rather than networks. So, instead of using classical networking constructs like VLANs, IPs, and MAC addresses, policies are defined in a language that is natural for application owners. The security policy and segmentation are decoupled from the underlying topology of the network through a group-policy approach.

Terminology and Definitions

Application Centric Infrastructure - ACI

A software-defined data center solution that applies an application-centric policy model to enable rapid application deployment. ACI data center infrastructure should be deployed in a spine-leaf topology and run on Cisco Nexus 9000 series switches.

Cisco APIC

Cisco APIC is the main architectural component of Cisco ACI. It automates and manages the Cisco ACI fabric, enforces policies, and monitors health. Cisco APIC establishes, stores, and enforces Cisco ACI application policies based on the application's network requirements. Cisco APIC also provides policy authority and resolution mechanisms.

It is important to distinguish between two views when looking at Cisco ACI and Check Point Security Gateways integration: Logical and Infrastructure. Infrastructure - relates to all the physical components: switches, routers, etc. Logical is used to set up communication between workloads within the switch fabric.

ACI Policy Model

ACI policy models provide a convenient means of specifying application requirements, which APIC then translates into a network infrastructure. A number of constructs are included in this object-oriented model, including tenants, contexts, bridge domains, endpoint groups, and service graphs.

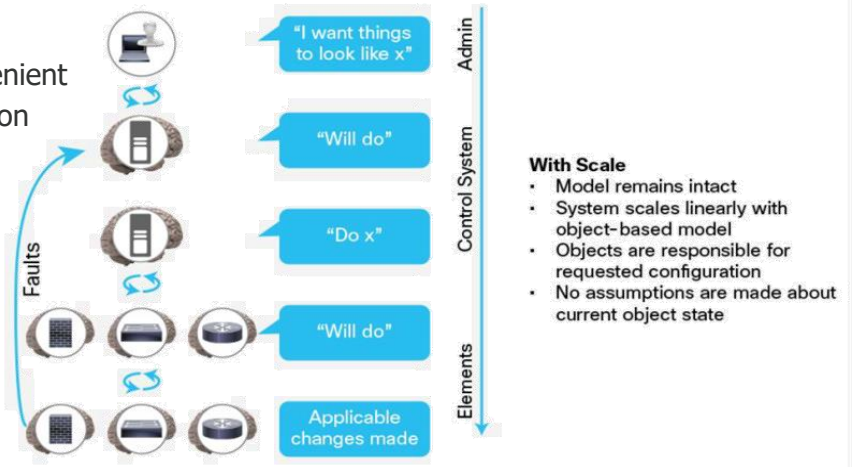

Figure 1: Cisco ACI Policy Model[1], Source: Cisco Systems

Policy models are based on promise theory, allowing declarative, scalable control of intelligent objects. Promise theory relies on the underlying objects handling configuration state changes initiated by the control system. This reduces the complexity of the controller and allows for greater scalability.

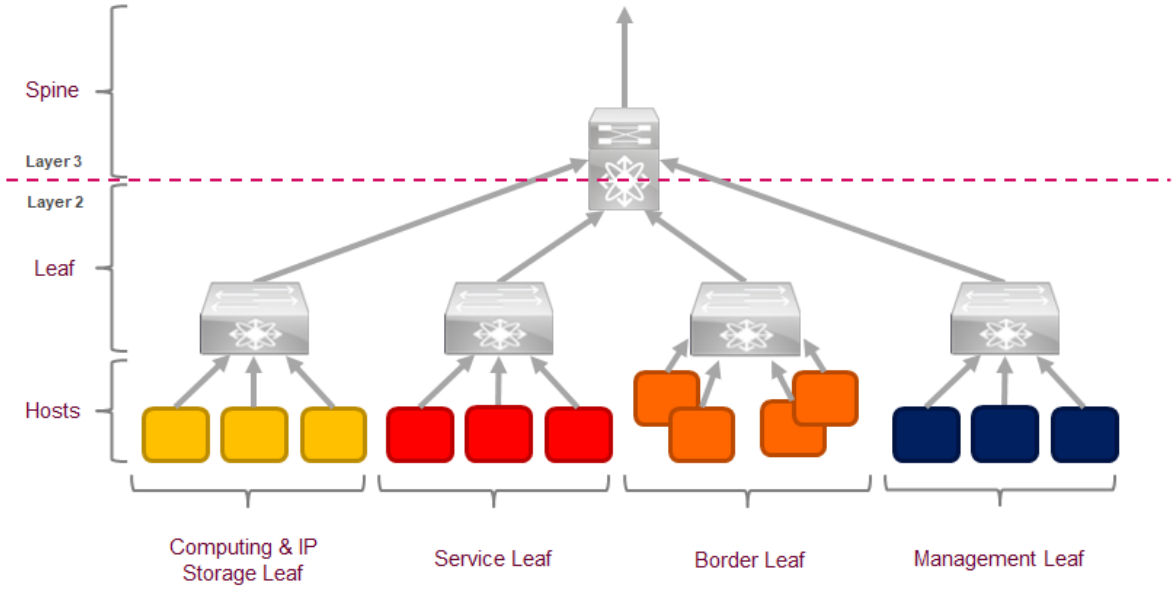

Figure 2: Spine and Leaf Topology[1]

Spine

Spines are special switches that form the backbone of ACI networks. All leaf switches must be connected to spines, and spines handle leaf-to-leaf communication. Spine switches typically contain a large number of high bandwidth (40 /100 GbE) aggregation ports. The ACI fabric relies on these ports for bandwidth throughput.

Leafs

ACI's spine-leaf topology uses leaf switches to connect all endpoint devices, such as servers, routers or firewalls, to the ACI fabric. Leaf switches can be defined based on the roles attached to endpoints by enabling all of them to connect at the same layer:

- **Computing Leaf Switches -** Used to connect to computer systems.
- **Service Leaf Switches -** Used to connect to Layer 4-7 service devices, such as application load balancers and firewalls.
- **IP Storage Leaf Switches -** Used to connect to IP storage systems.
- **Border Leaf Switches -** Used for external connectivity. External routers are supported for routing and policy enforcement for traffic between internal and external endpoints.

[1] Source: Cisco Data Center Spine-and-Leaf Architecture: Design overview White paper: https://www.cisco.com/c/en/us/products/collateral/switches/nexus-7000-series-switches/white-paper-c11-737022.html

- **Transit Leaf Switches -** Used to connect to the spines on other data centers. It exists only in stretched fabric topologies.
- **Management Leaf Switches** - Used to connect to the OOB Network for the Operations and Management Services for the infrastructure.

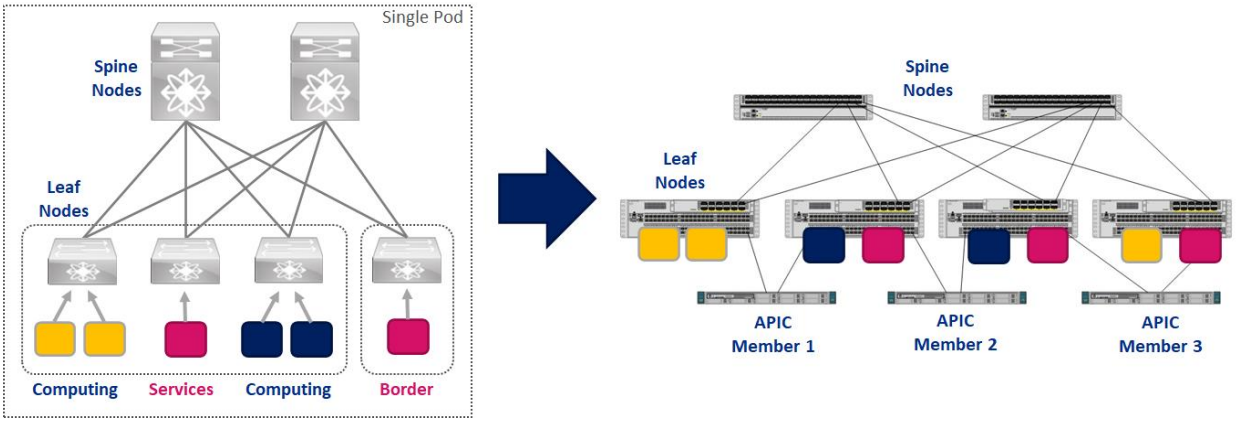

Figure 3: Logical mapping of components to the infrastructure in the Switch Fabric.

Deployment Modes

In traditional data centers design, Cisco Systems typically used to refer to three tiers: core, aggregation, and access, while modern and advanced data centers design is typically based on a two-tier spine-leaf architecture. The new approach offers a more optimized design to accommodate east-to-west traffic flows, which are predominant in the new application based on the following design patterns.

Single Pod

APIC pods are sets of interconnected leaf and spine switches (ACI Fabrics) that are under the control of an APIC cluster.

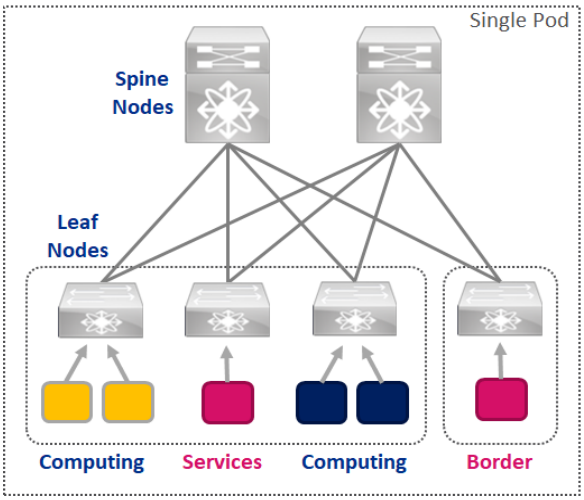

Figure 4: Single Pod Architecture

Multi-Pod

The Multi-pod[2] is an architectural design that has multiple ACI fabrics under the control of single management or administration.

InterPod Network (IPN)

The pods within the topology are all connected via IP-routed Inter-Pod Network[3], a transport type connection that enables IP Routing and Multicast in order to allow interconnection between pods, and connectivity within pods to the IPN occurs on spine nodes. IPN is not managed by APIC and needs to be configured independently.

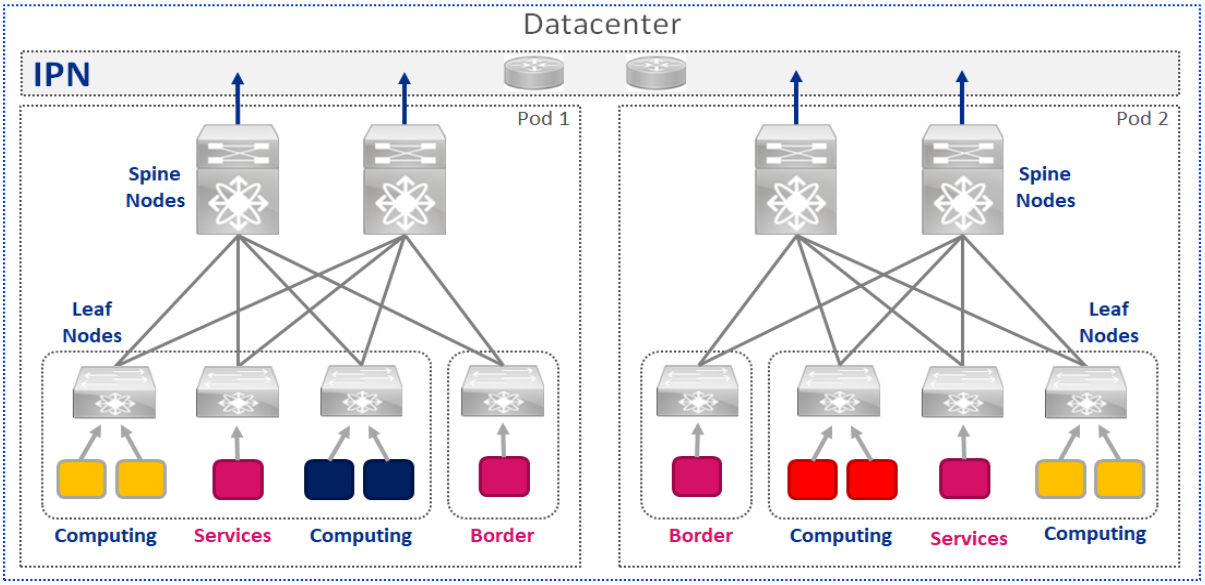

Figure 5: Multi-Pod & InterPod Network Architecture

[2] Source: Cisco ACI Multi-Pod White Paper - URL: https://www.cisco.com/c/en/us/solutions/collateral/data-center-virtualization/application-centric-infrastructure/white-paper-c11-737855.html

[3] Source: Cisco ACI Multi-Pod White Paper, Inter-Pod Connectivity Deployment Considerations - URL: https://www.cisco.com/c/en/us/solutions/collateral/data-center-virtualization/application-centric-infrastructure/white-paper-c11-737855.html#InterPodConnectivityDeploymentConsiderations

Multi-Site

The Multi-site[4] architecture is the interconnection of APIC cluster domains with their associated pods. Multi-Site designs may also be called Multi-Fabric designs since they interconnect separate availability zones (ACI fabrics), deployed either as single pods or multiple pods (Multi-Pod design).

InterSite Network (ISN)

All communications between endpoints (EPG's) can be accomplished using site-to-site VPNs (Virtual Extensible Local Area Network) over a generic IP network that connects different sites with the InterSite Network[5]. Using VXLAN encapsulation for the InterSite IP network greatly simplifies the setup of the configuration. Other than routing capabilities and increased maximum transmission unit (MTU) size (given the overhead created by the VXLAN encapsulation), this IP network does not have any other specific functional requirements.

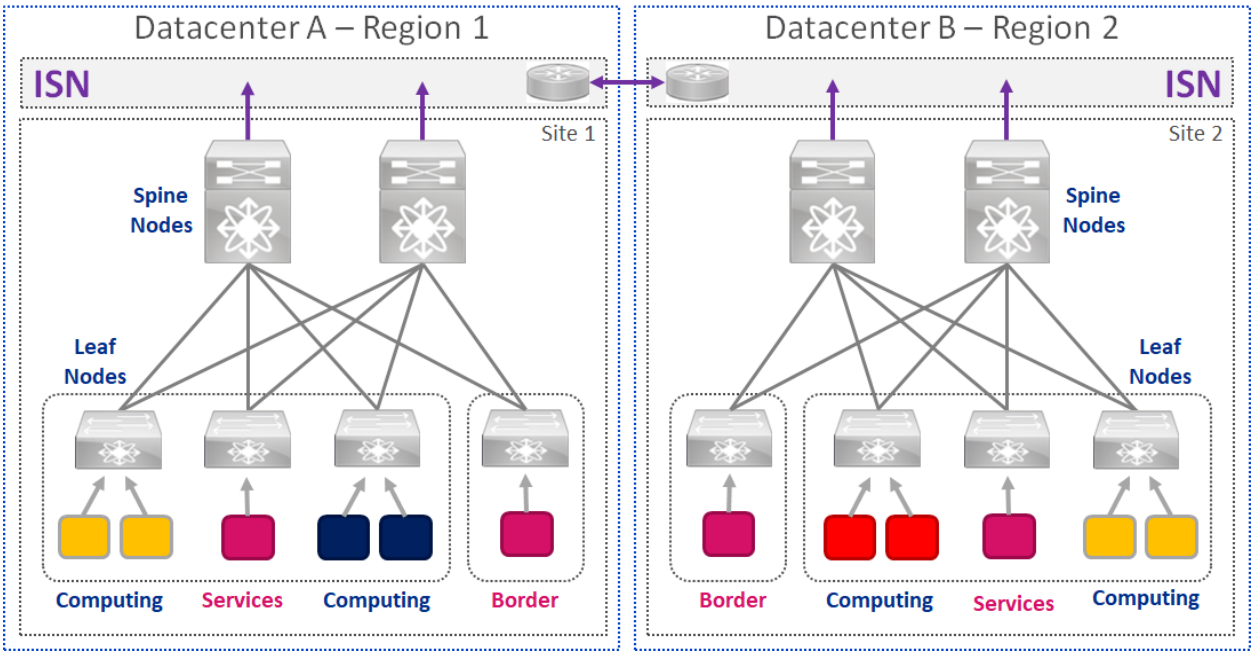

Figure 6: Multi-Site and Inter-Site Network Architecture

In a Multi-Site topology, each fabric could be considered a separate availability zone. These availability zones are managed cohesively by the Multi-Site Orchestrator. The nature of the

[4] Source: Cisco ACI Multi-Site White Paper - URL: https://www.cisco.com/c/en/us/solutions/collateral/data-center-virtualization/application-centric-infrastructure/white-paper-c11-739609.html

[5] Source: Cisco ACI Multi-Site White Paper, Intersite Network (ISN) deployment considerations - URL: https://www.cisco.com/c/en/us/solutions/collateral/data-center-virtualization/application-centric-infrastructure/white-paper-c11-739609.html#IntersiteNetworkISNdeploymentconsiderations

architecture ensures that whatever happens to one site (availability zone) in terms of network-level failures and configuration mistakes will not impact other site(s) or availability zones. This guarantees business continuance at the highest level.

There is a misconception that Multi-Site somehow supersedes Multi-Pod or that the Multi-Pod architecture is no longer relevant. In reality, they are two separate technologies applicable to different use cases.

Furthermore, there is no valid reason for Multi-Pod and Multi-Site topologies not to work together. For example, there could be multiple data centers deployed all over the world, but each can have its own ACI Multi-Pod fabric and tied together through the Multi-Site Orchestrator. These two architectures are built to work harmoniously, so you are no longer faced with an either/or decision and will ultimately have a high degree of deployment flexibility.

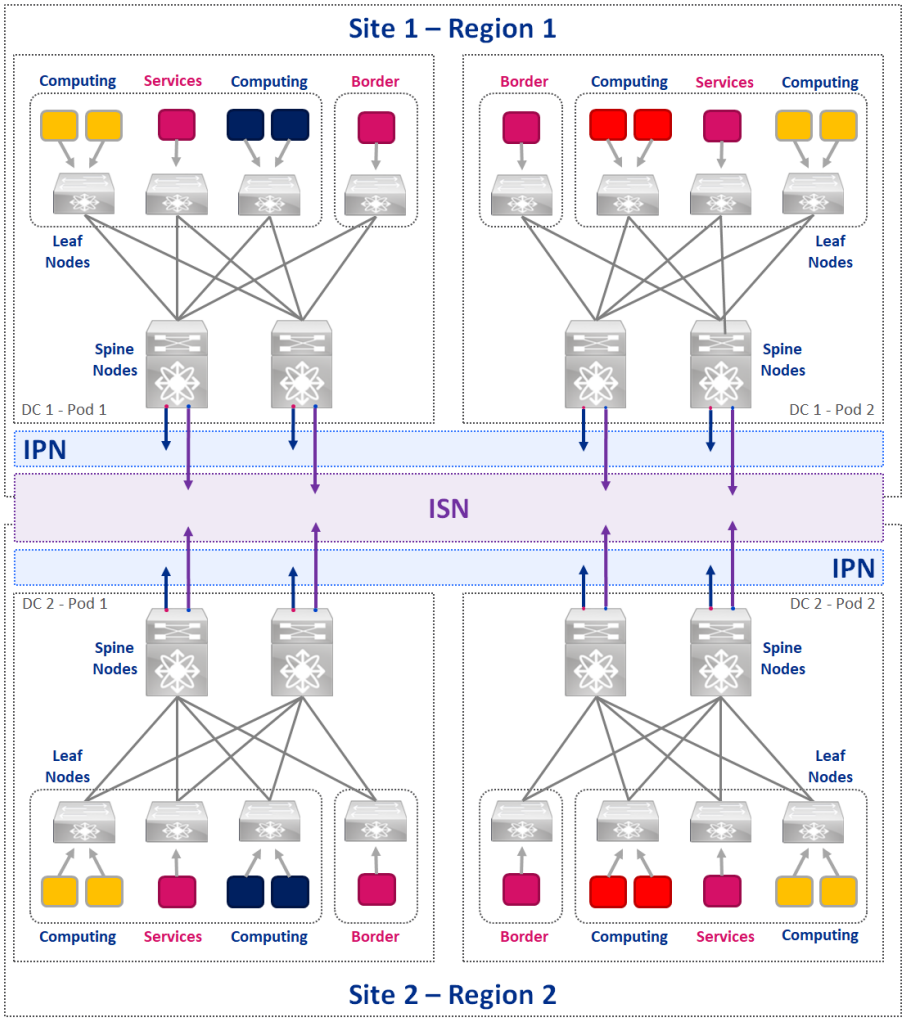

Figure 7: Multi-Pod and Multi-site architecture

Cisco ACI environment

Tenant

Tenants act as containers for other elements of the policy model (such as contexts, bridge domains, contracts, filters, and application profiles). Each tenant can be virtually isolated from the rest of the environment, or tenants can have some shared resources. Furthermore, it is a unit of isolation from a policy perspective, but it is not a private network. Depending on the environment, tenants can represent different customers, organizations, domains, or simply a convenient way of policies grouping.

Note: VRFs are also known as contexts; each VRF can be associated with multiple bridge domains.

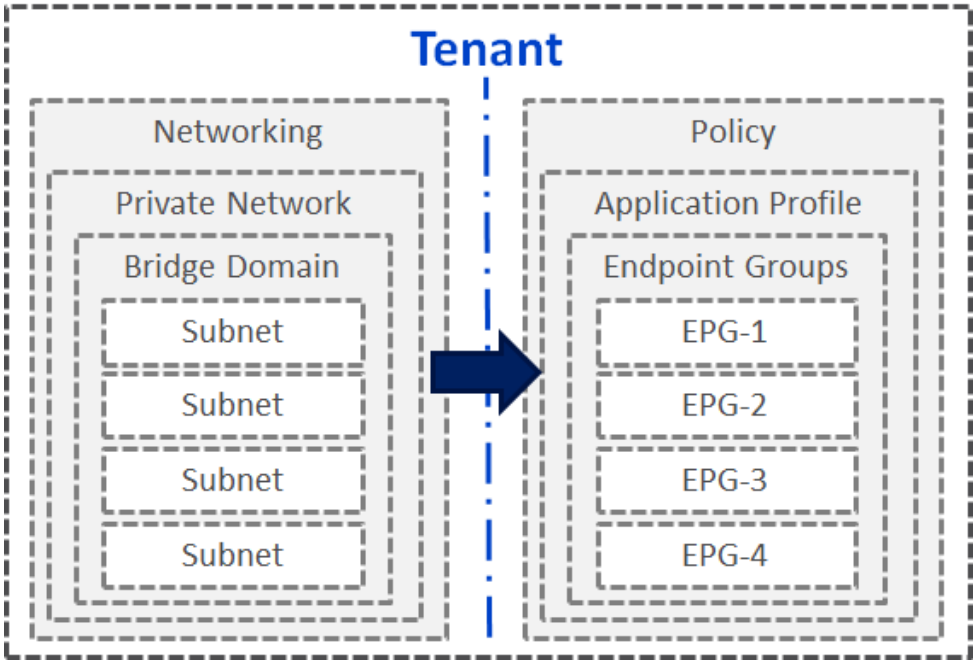

Figure 8: Mapping between Networking and Policy for Tenant configuration[6]

Tenants mainly serve as a logical separator for customers, business units, groups or similar entities. They can be used to separate traffic, visibility, or admin separation. For example, private networks that are intended for use by multiple tenants and are not created in the common tenant require explicit configuration to be shared.

[6] Source: Operating Cisco Application Centric Infrastructure, Tenants - URL: https://www.cisco.com/c/en/us/td/docs/switches/datacenter/aci/apic/sw/1-x/Operating_ACI/guide/b_Cisco_Operating_ACI/b_Cisco_Operating_ACI_chapter_0111.html

The following Tenant distribution is considered to be best practices:

- **Common:** The common tenant is usually used as a shared services tenant. Objects created inside the common tenant are available to other tenants.
- **Infrastructure:** It is the infrastructure tenant's responsibility to expand the infrastructure. In Multi-Pod and Multi-Site fabric deployment, the infra tenant is used to link the pods or sites.
    - Examples: Internal DMZ, External DMZ, Restricted Data Center
- **Management:** Most of the management configuration is performed in the management tenant. Assigning management IP addresses to switches and configuring the contracts that will limit access to the fabric management interfaces would be completed in this tenant.

VRF/Private Network

The Virtual Routing and Forwarding (VRF) object, or context, represents a tenant network (a private network in the APIC GUI). A tenant can have multiple VRFs. A VRF is a unique Layer 3 forwarding and application policy domain.

It is used to define a unique layer 3 forwarding domain within the fabric. One or more VRF can be created inside a tenant, also known as 'private networks', and can be viewed as the equivalent of a VRF in the traditional networking world. Each context defines a separate layer 3 domain, which means IP addresses within a context can overlap with addresses within other contexts.

Bridge Domains and Subnets

A Bridge Domain (BD)[7] is a construct used to define a layer 2 boundary within the fabric. BDs can be viewed as somewhat similar to regular VLANs in a traditional switching environment. Bridge domains possess several enhancements such as better handling of ARP requests and no flooding by default. Furthermore, a Bridge domain may span multiple switches and contain multiple subnets, but each subnet can only operate within a single bridge domain.

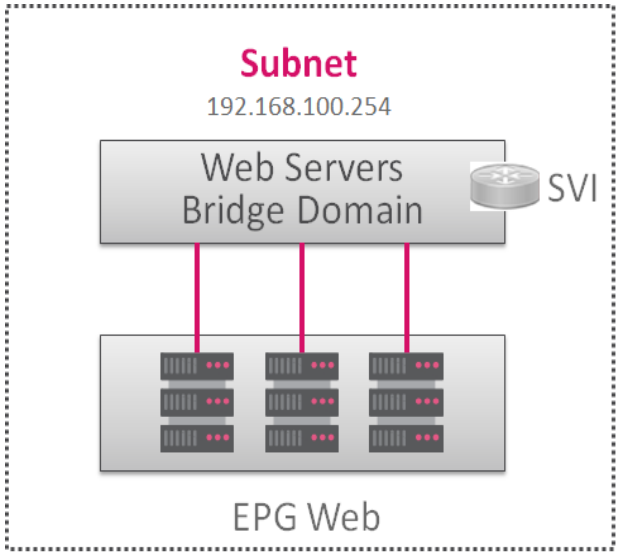

Figure 9: Example of a Bridge domain and SVI.

A BD is essentially a container for subnets. A Switched Virtual Interface (SVI) is a logical router that is configured for a VLAN as the default gateway in order to allow traffic to be routed between VLANs. A subnet is used to designate which gateway (SVI) will be used within

[7] Source: Cisco Application Centric Infrastructure Fundamentals, Bridge Domains and Subnets - URL: https://www.cisco.com/c/en/us/td/docs/switches/datacenter/aci/apic/sw/1-x/aci-fundamentals/b_ACI-Fundamentals/b_ACI-Fundamentals_chapter_010001.html#concept_8FDD3C7A35284B2E809136922D3EA02B

a bridge domain. This gateway will typically be used by hosts associated with a bridge domain as their next-hop gateway. A bridge domain's gateways are available on all leaf switches where the bridge domain is active.

Physical & VMM (Virtual Machine Manager) domains

Typically, physical domain[8] profiles are used for bare metal server attachment and management access, while a VLAN pool is associated with a domain. Endpoint groups (explained below) are then configured to use the VLANs associated with the domain.

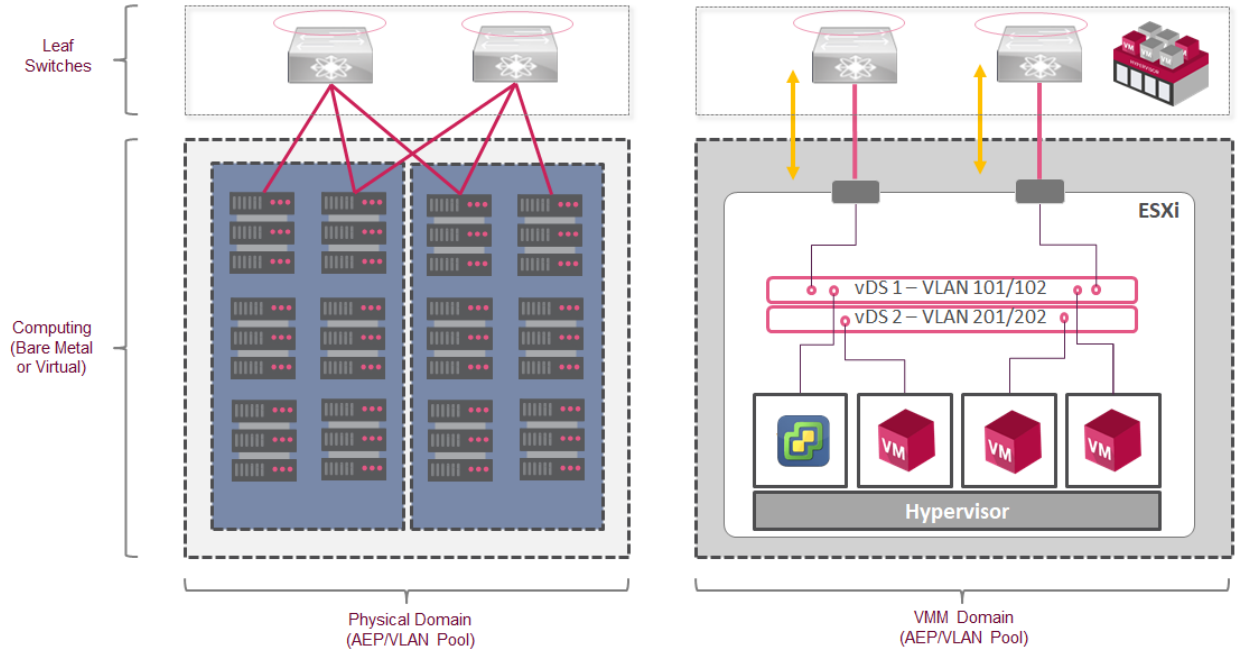

Figure 10: Example of Physical & VMM Domains

A Virtual Machine Manager (VMM)[9] domain profile specifies the policies for connecting virtual machine controllers to the ACI fabric. The VMM domain policy is created in APIC and pushed into the leaf switches. VMM domains contain VM controllers, such as VMware vCenter, and the credential(s) required for the ACI API to interact with the VM controller. A VMM domain enables VM mobility within the domain but not across different domains. A single VMM domain can contain multiple instances of VM controllers but they must be the same kind.

---

[8] Source: Cisco APIC Layer 2 Networking Configuration Guide - URL: Cisco APIC Layer 2 Networking Configuration Guide, Release 3.x and Earlier - Networking Domains [Cisco Application Policy Infrastructure Controller (APIC)] - Cisco

[9] Source: Configure VMM Domain Integration with ACI and UCS B Series, Create the VMM Domain - URL: https://www.cisco.com/c/en/us/support/docs/cloud-systems-management/application-policy-infrastructure-controller-apic/118965-config-vmm-aci-ucs-00.html#anc5

Endpoint Groups (EPGs)

In simple words, the End Point Group is a group of devices/endpoints that share common policy requirements. It provides a new model for mapping application resources to the network. Rather than using forwarding constructs such as addressing or VLANs to apply connectivity and policy, EPGs use a grouping of application endpoints.

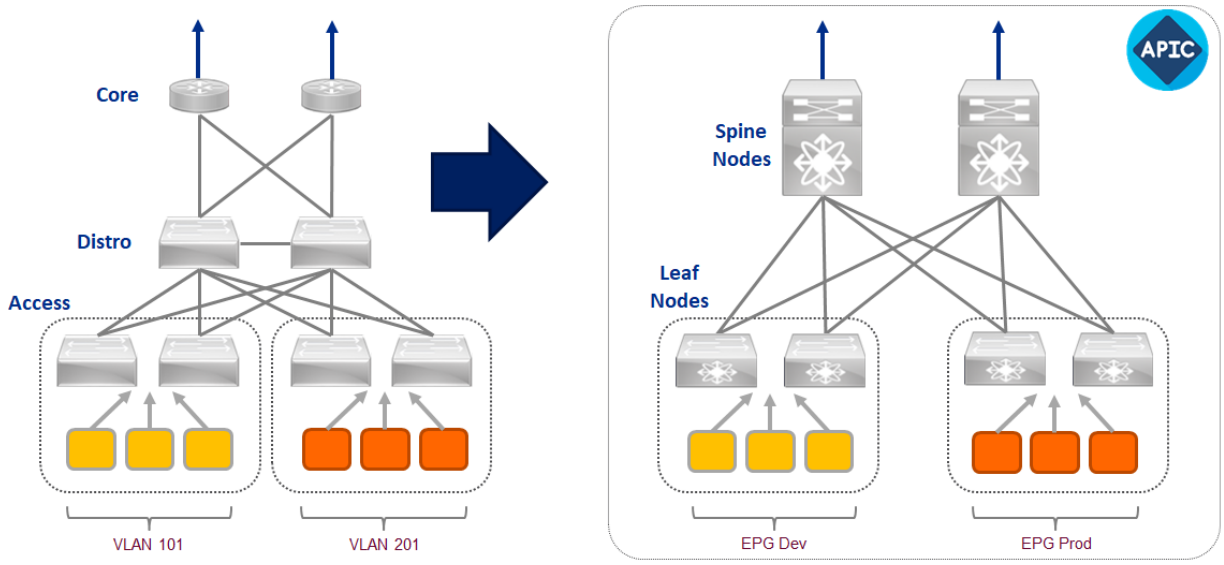

Figure 11: EPG mapping with traditional VLAN approach, Source: Cisco Systems

EPGs act as a container for collections of applications, or application components and tiers, which can be used to apply forwarding and policy logic. They allow the separation of network policy, security, and forwarding from addressing and instead apply it to logical application boundaries.

There are multiple types of EPGs:

- Application endpoint group - This is the regular EPG we all know and love
- L2 external EPG - An EPG used to contain endpoints from external L2 connectivity (used when extending a BD to an external L2 network)
- L3 external EPG - An EPG used for external L3 connectivity (external routes)
- Management EPGs for out-of-band and in-band access

EPGs are designed to abstract the instantiation of network policy and forwarding from basic network constructs (VLANs and subnets.) This allows applications to be deployed on the network in a model consistent with their development and intent. Endpoints assigned to an EPG can be defined in several ways. Endpoints can be defined by virtual port, physical port, IP address, DNS name, and in the future through identification methods such as IP address plus Layer 4 port and others.

There is no dedicated manner in which EPGs should be deployed and utilized; however, the rest of this document will cover some typical EPG uses.

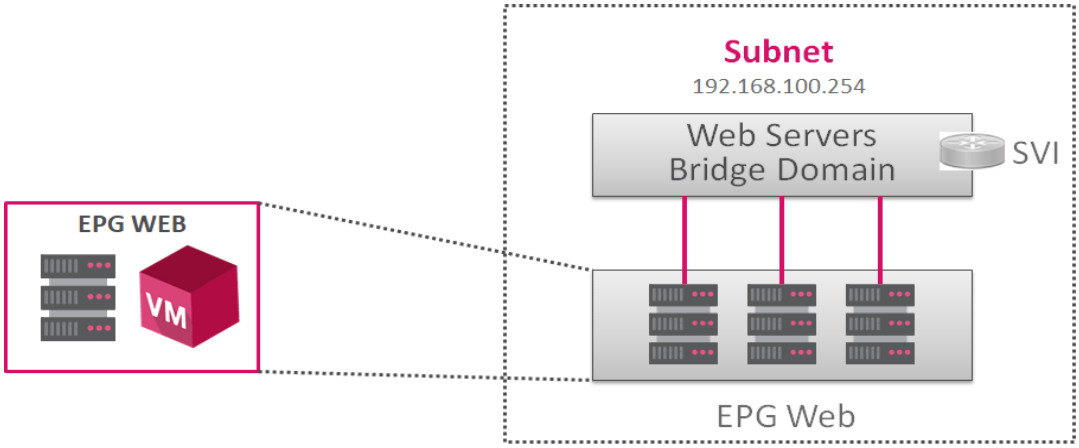

Figure 12: EPG grouping devices and Endpoints that share a common policy

A few typical examples:

Mapping traditional network constructs to the ACI fabric

- EPG as VLAN
- EPG as a subnet (model classic networking using EPGs)
- EPG as a virtual extensible LAN (VXLAN)/Network Virtualization using Generic Routing Encapsulation (NVGRE) virtual network identifier (VNID)
- EPG as a VMware port group

Utilizing the ACI fabric for stateless network abstraction

- EPG as an application component group (web, app, database, etc.)
- EPG as a development phase (development, test, production)
- EPG as a zone (internal, DMZ, shared services, etc.)

MicroEPG (uEPG)

Micro-segmentation[10] is the method of creating zones in data centers and cloud environments to isolate workloads from one another and secure them individually. By default, Endpoints inside the same EPG can communicate freely without any restrictions. A Micro EPG (uEPG) is equivalent to a regular EPG for all intents and purposes (as Service Graphs and PBRs), but the classification is based on endpoint attributes (and dynamic in nature). This enables the organization the capability to filter with those attributes and apply more dynamic policies and traffic inspection through the Service Graphs using Check Point Firewalls applying policies to any endpoints within the tenant.

[10] Cisco ACI Virtualization Guide 3.0 - URL: Cisco ACI Virtualization Guide, Release 3.0(1) - Microsegmentation with Cisco ACI [Cisco Application Policy Infrastructure Controller (APIC)] - Cisco

- For endpoints on Physical Domains (bare metal), you can use IP or MAC addresses
- For endpoints on VMware or Microsoft VMM Domains, you can use IP, MAC addresses or VM-attributes

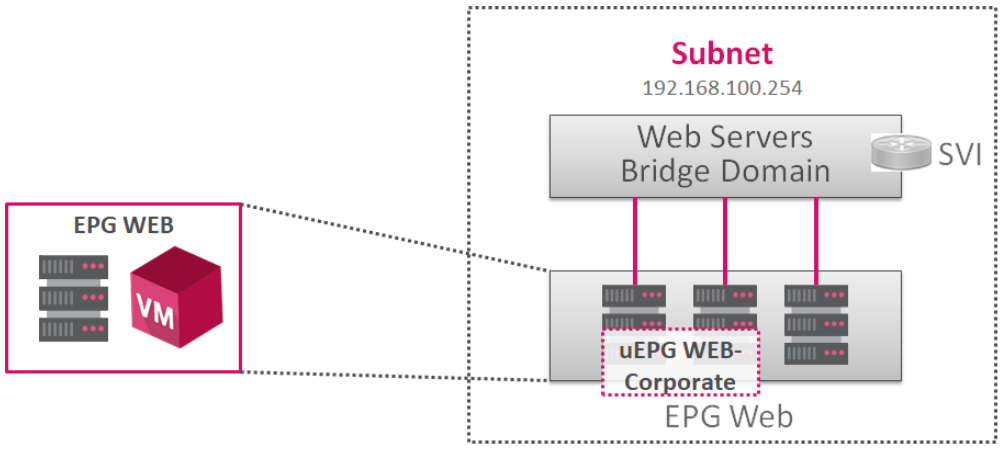

Figure 13: Example of Micro-segmentation with Cisco ACI and Check Point

Endpoint Security Group (ESG[11])

By definition, the EPGs are associated with a single bridge domain and used to define security zones within it, EPG are used to define both forwarding and security segmentation at the same time.

The direct relationship between the bridge domain and an EPG limits the possibility of an EPG spanning more than one Bridge Domain. This kind of limitation can be resolved by using a new ESG construct because it will allow the relationship between endpoints from multiple BD / EPGs but still limited to a single VRF.

The Endpoint Security Group (ESG) enables organizations to move towards with an Application Centric model approach, instead of spending a lot of time on preparing for a migration from a Network Centric to an Application Centric model.

Some typical uses for ESGs:

- ESG and ESG
- ESG and L3Out EPG
- ESG and inband-EPG
- ESG and vzAny

[11] Cisco APIC Security Configuration Guide 5.2 - URL: Cisco APIC Security Configuration Guide, Release 5.2(x)

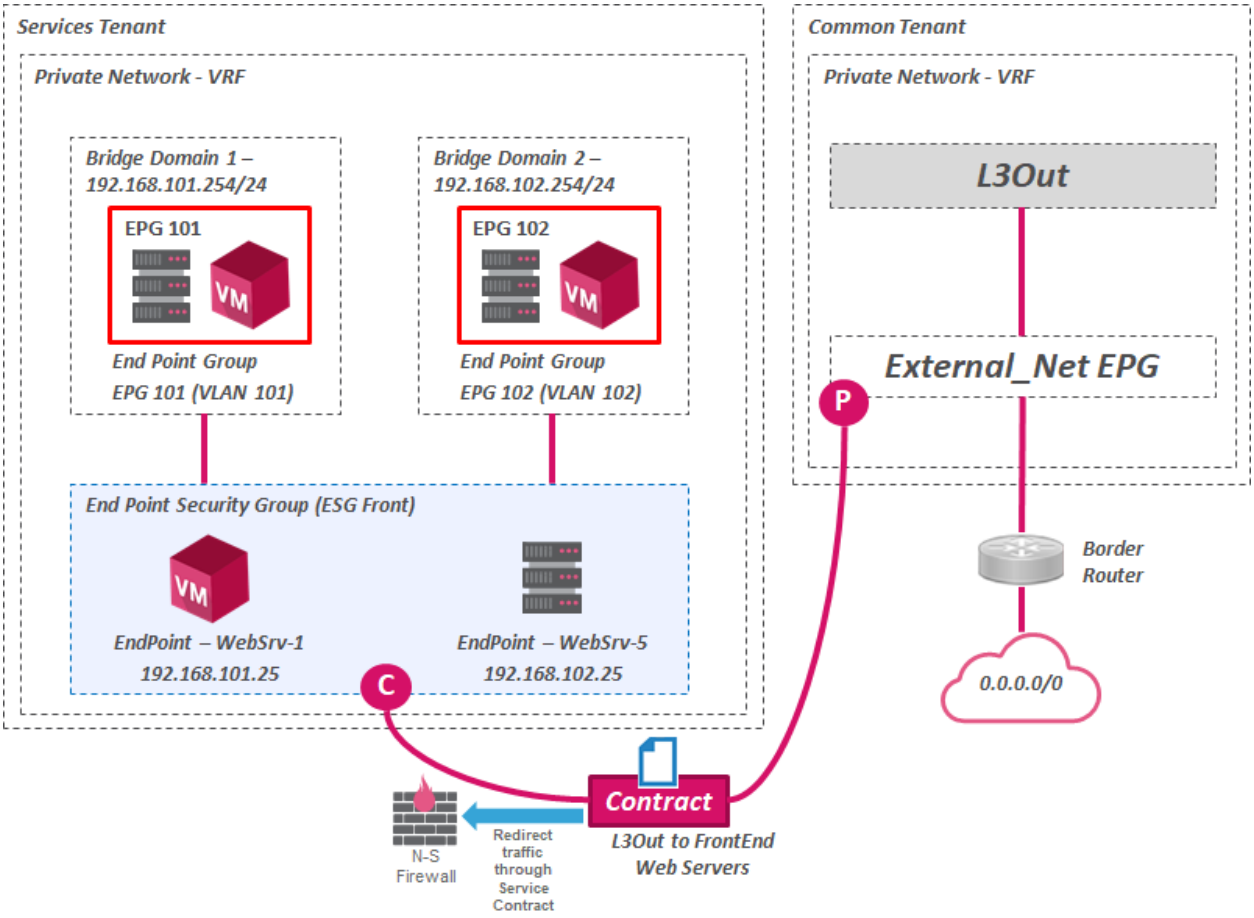

Figure 14: Example of a practical use case of ESG with Check Point Firewalls

Application Profile

An application profile[12] defines the policies, services, and relationships between endpoint groups (EPGs). Application profiles can contain one or more EPGs. Modern applications typically contain multiple components. For example, an e-commerce application could require a web server, a database server, data located in a storage area network, and access to outside resources that enable financial transactions.

The application profile includes as many (or as few) EPGs as necessary to provide necessary the functionality of an application.

[12] Source: Cisco Application Centric Infrastructure Fundamentals, Application Profile - URL: https://www.cisco.com/c/en/us/td/docs/switches/datacenter/aci/apic/sw/1-x/aci-fundamentals/b_ACI-Fundamentals/b_ACI-Fundamentals_chapter_010001.html#concept_6914B5520ECA4731962F30F93E5A77A6

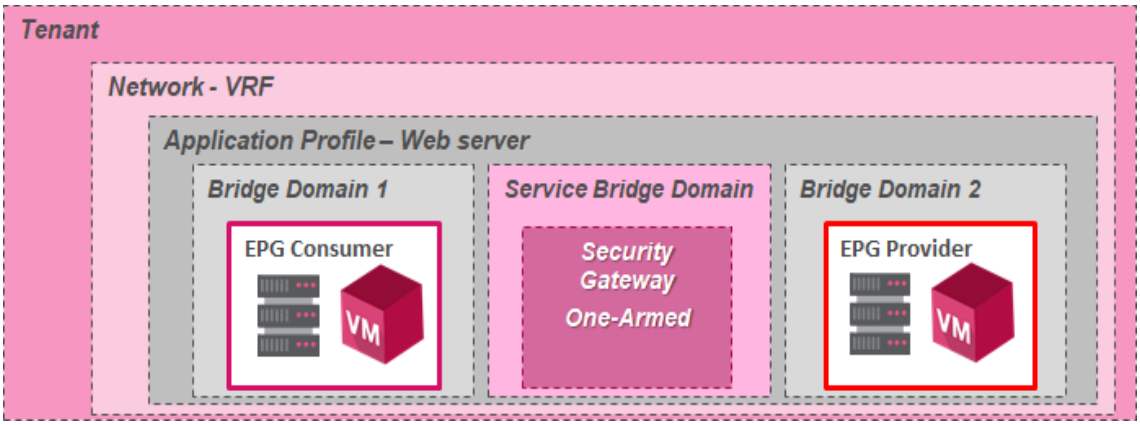

Figure 15: Application Profile and its interaction with other layers

Service Contract

A Service contract[13] within Cisco ACI defines how EPGs can communicate with each other, defining the Ingress and Egress traffic flows. This is based on an allow list - without a permit contract (by default) traffic between different EPGs is not allowed. A contract consists of subjects, each made up of filters, actions, and labels. A contract can have many subjects.

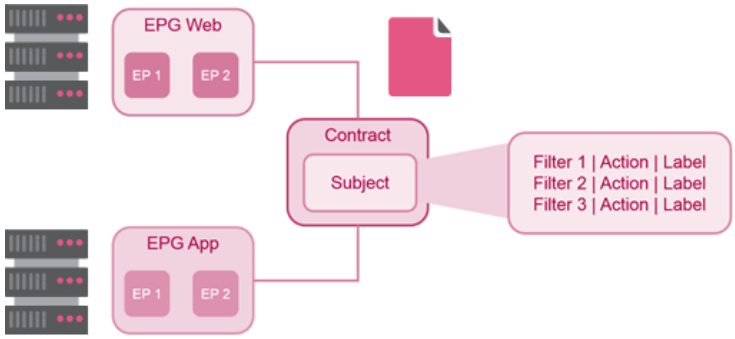

Figure 16: Service Contract example

Service Graph

By using Cisco ACI's service graph, traffic between different security zones within the fabric can be redirected to a firewall or load balancer, eliminating the need to configure the firewall or load balancer as the default gateway for servers. Furthermore, Cisco ACI can selectively send traffic to L4-L7 devices (for example Check Point firewall).

[13] Source: Cisco ACI Contract Guide, How contracts work - URL: https://www.cisco.com/c/en/us/solutions/collateral/data-center-virtualization/application-centric-infrastructure/white-paper-c11-743951.html#Howcontractswork

Firewall inspection also can be transparently inserted in a Layer 2 domain with almost no modification to existing routing and switching configurations. Moreover, Cisco ACI allows increasing the capacity of L4-L7 devices by creating a pool of devices to which Cisco ACI can distribute traffic using Symmetric PBR mechanism.

With the service graph, Cisco ACI introduces changes in the operating model. A configuration can now include not only network connectivity—VLANs, IP addresses, and so on, but also the configuration of Access Control Lists (ACLs), load-balancing rules, etc.

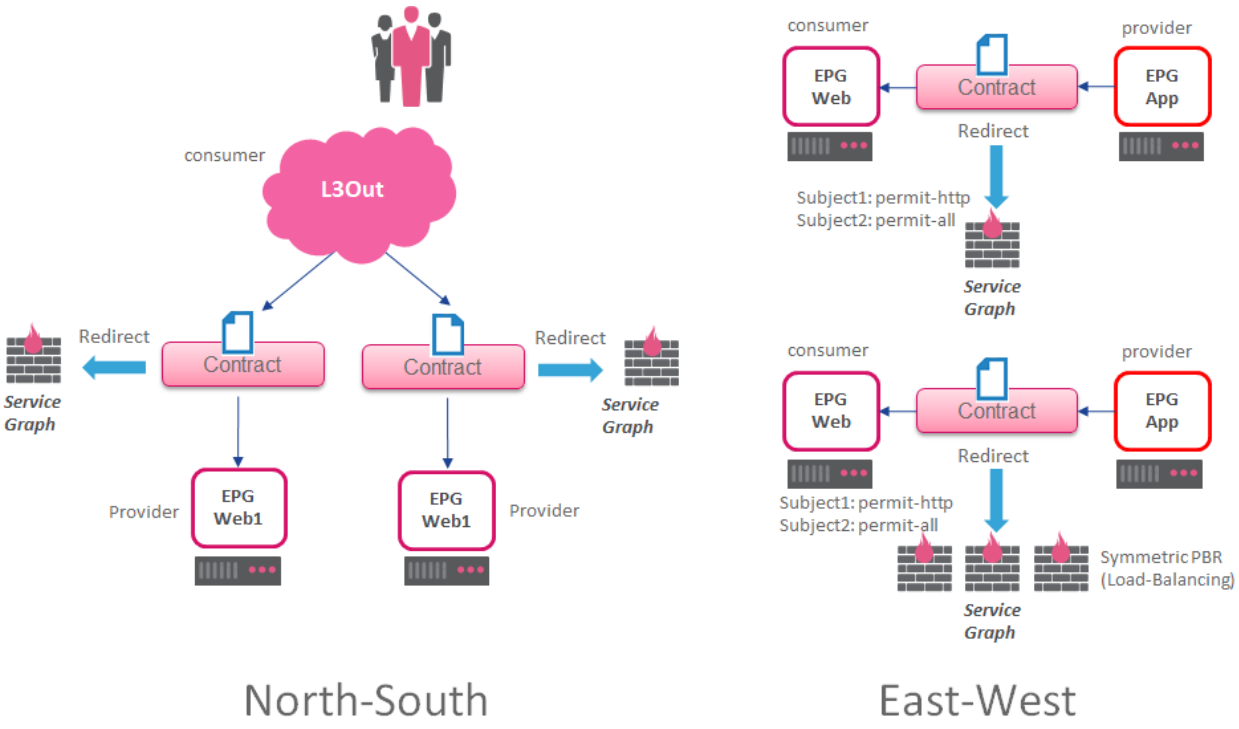

Figure 17: Examples of Service Graphs for North-South and East-West traffic flows

Check Point Integration with Cisco ACI

Check Point CloudGuard for ACI[14] is the Check Point Advanced Security solution for the Cisco ACI fabric. Check Point CloudGuard is designed to enforce advanced threat prevention within the ACI fabric and integrates seamlessly with Cisco APICs and Check Point Security Management Server. It proactively stops malware and zero-day attacks inside the Data Center environment and outside of the fabric. Unified management of virtual and physical gateways simplifies security management in the hybrid network environment.

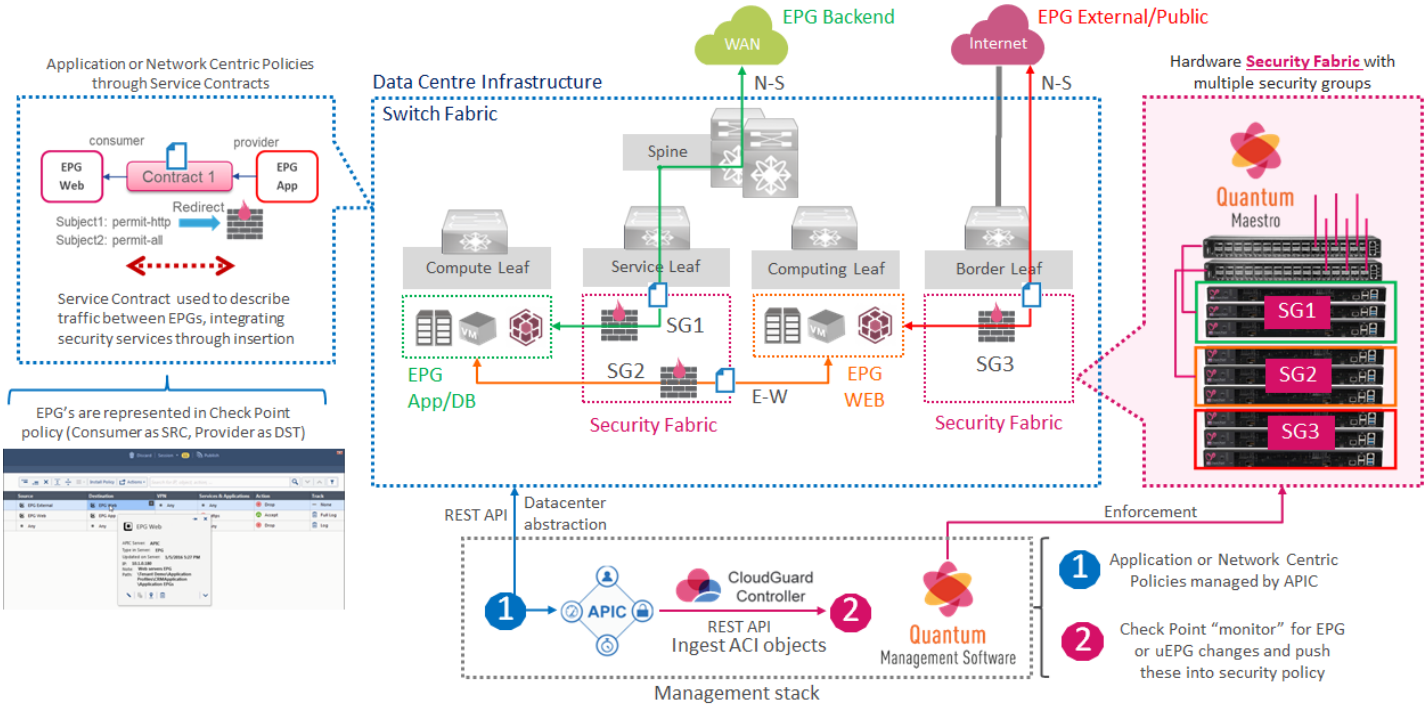

Figure 18: Check Point Integration with Cisco ACI and Maestro

Check Point CloudGuard for Cisco ACI has two main components:

- **The CloudGuard Controller**
  - With CloudGuard Controller, Check Point Security Management Server can be integrated with Cisco APIC, as well as with other leading SDN controllers and cloud managers, including VMware vCenter, in order to make dynamic security policies for ACI objects and VMs. The Controller automatically syncs any object changes directly into the dynamic security policies, without the need for a policy push. It manages CloudGuard gateways as well as physical gateways and gives complete

[14] Source: CloudGuard for ACI R80.10 Administration Guide, URL: https://sc1.checkpoint.com/documents/R80.10/WebAdminGuides/EN/CP_R80.10_vSEC_for_ACI_AdminGuide/html_frameset.htm?topic=documents/R80.10/WebAdminGuides/EN/CP_R80.10_vSEC_for_ACI_AdminGuide/171241

visibility into Data Center security. The CloudGuard Controller can be used to generate security policies for installations on any Check Point Security Gateway across the network.

- **The Check Point Security Gateway**
  - The Security Gateway can be is a hyper-scale system (Maestro), physical or virtual Check Point appliances deployed inside the ACI fabric enforcing the Check Point security policy.

CloudGuard Controller and Cisco APIC

Check Point CloudGuard for ACI requires a license attached to the Security Management Server or the Multi-Domain Server. The license is based on the total number of Cisco ACI leaf switches managed by the APICs that are integrated with the Check Point Security Management Server or Multi-Domain Server. The CloudGuard for ACI license includes ACI integration functionality. Additional licenses aren't required on the gateways for this functionality.

The license covers Management High Availability for the Security Management Server and the Multi-Domain Server. All processes not associated with ACI integration must have a separate license. For example, licenses to enable typical management and/or gateway functions or capabilities. The license is perpetual and cumulative, which means it is always possible to add more leaf licenses.

The CloudGuard Controller is a component of every Security Management Server, which integrates with the Cisco APIC at the management level. It allows consumption of various ACI metadata which can be used in the Check Point access and threat prevention security policies.

For example, constructs such as EPGs are discovered by the CloudGuard Controller from the APIC. When these EPGs are used as a part of already provisioned security policies, CloudGuard Controller monitors and updates their membership properties on the relevant security gateways in real-time (within a few seconds). In this way, organizations benefit

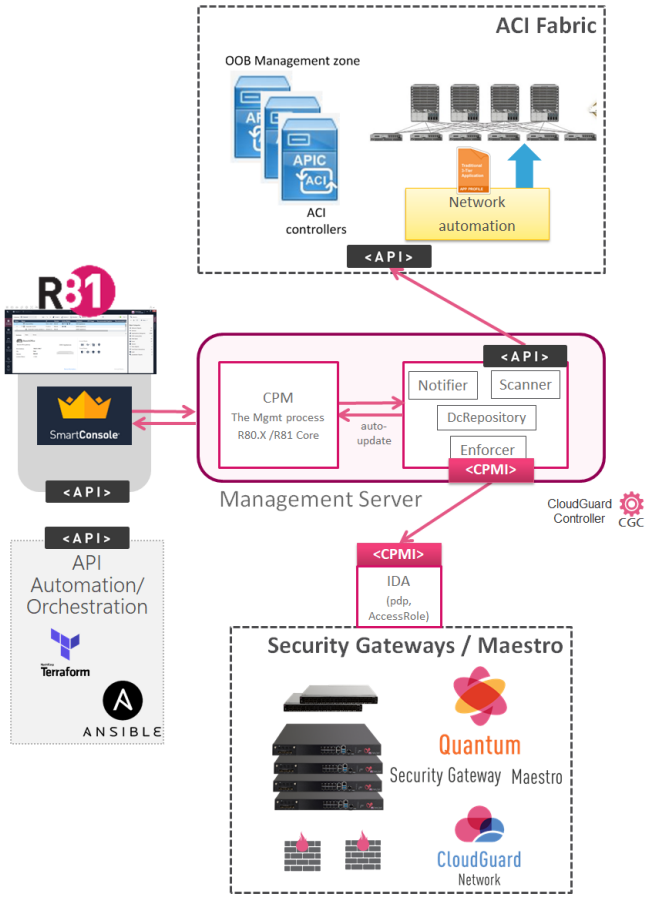

Figure 19: Check Point CloudGuard Controller integration with Cisco APIC.

from using a single abstraction view of objects and policies within the datacenter. Furthermore, Cloudguard Controller supports integration with Azure, AWS, Alibaba, VMware vCenter, VMware NSX-T, Kubernetes, and others in the same way, making unified multi-cloud policies possible.

Mapping Service Contracts with Check Point Security Policies

The mapping between an Application Profile and Application Control Policy and Threat Prevention Policy is crucial to the construction of security policies in Check Point Security Gateways. However, the mapping process is very simple.

Comparison of Cisco ACI and Check Point Software Technologies constructs used in the policy:

- EPG Consumer → Source/From
- EPG Provider → Destination/To
- Filters → Ports/Applications/Signatures
- Actions → Action (Allow, Deny, Drop)

When the Application Profile is built using a Service Graph, we can import EPG objects through the Datacenter configuration in the Check Point Management Console.

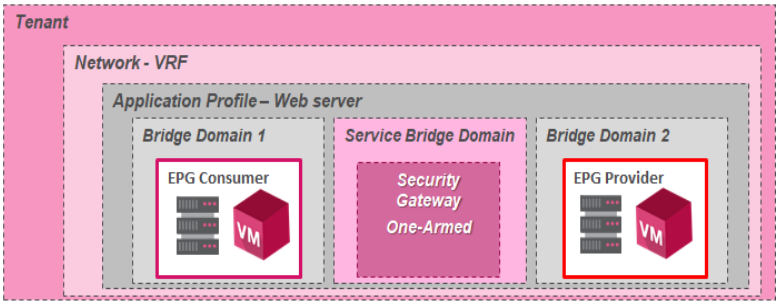

ACI-aware Security Policy

Check Point Unified Security Policy

Rule	From	To	Application	Action
3	Web_Serve1 (ACI EPG object)	App_Server1 (ACI EPG object)	Tomcat	Allow
4	App_Server1 (ACI EPG object)	Database_Servers (ACI EPG object)	MSSQL	Allow

Figure 20: Mapping Cisco APIC Service Contract with Check Point Security Policies

ACI and Check Point Gateways physically start interacting only after their logical configurations are completed.

Therefore, it is vital to consider both perspectives in order to create both Network Policies and Security Policies.

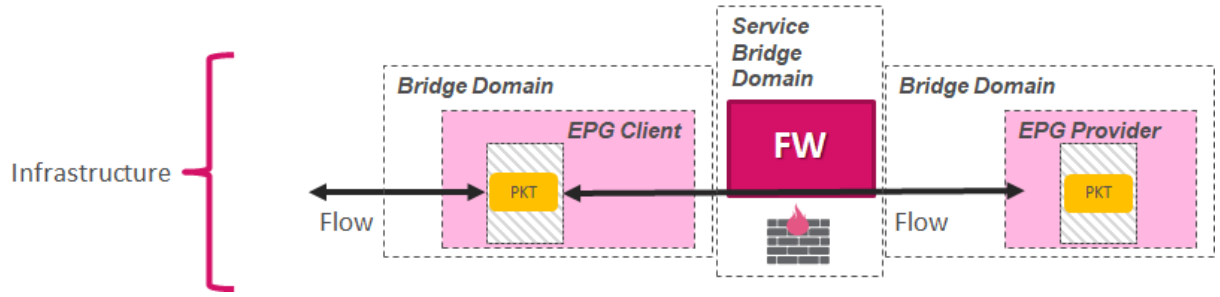

Figure 21: Infrastructure view for the traffic flow between EPG's and Service Bridge

Below is an example of a Security Policy that is mapped from the Logical Perspective above.

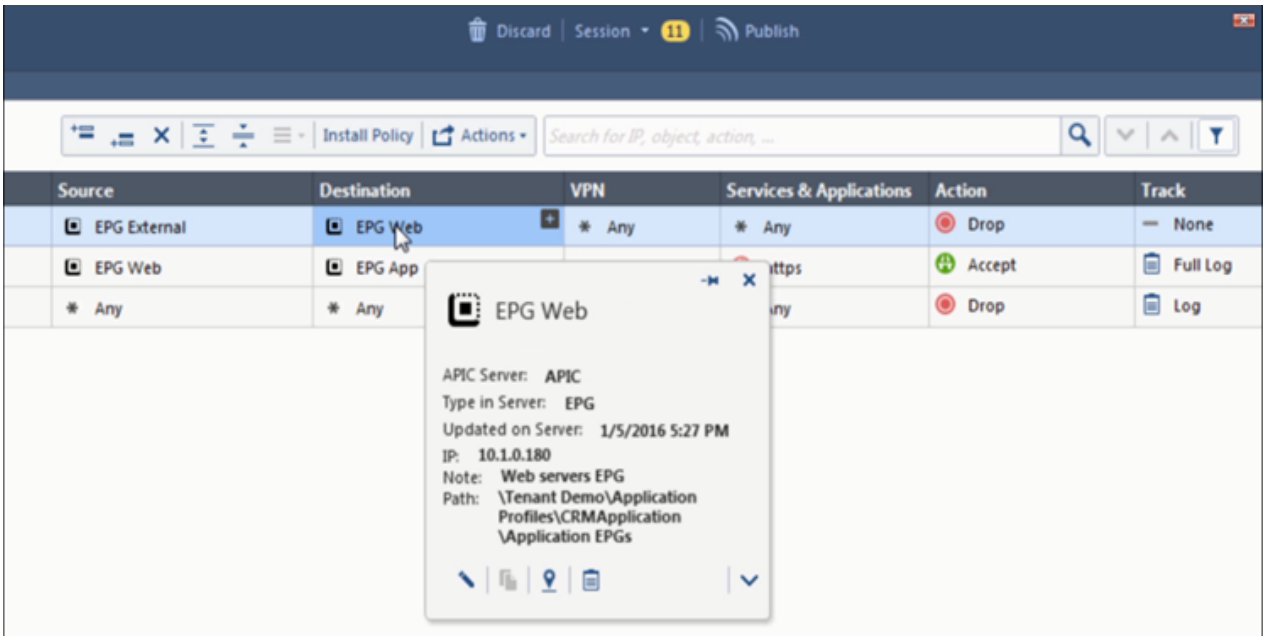

Figure 22: Check Point Security policy example - EPG objects propagated from APIC to the CloudGuard Controller

All ACI objects (EPGs, Application Profiles, Tenants, and Application EPGs) used in the unified security policy get updated automatically in nearly real-time via API provided by the CloudGuard Controller without any manual EPG object modifications in the Check Point policy.

Check Point and Cisco ACI Integration Benefits

Together, Cisco and Check Point provide a powerful solution that gives customers comprehensive traffic visibility and reporting capabilities in addition to proactive protection from even the most advanced threats within virtual network environments. The joint solution creates an intelligent application delivery network architecture that keeps security on top of application workloads while accelerating application deployment and maintaining multi-tenancy and operational workflows. The benefits provided by this power solution extend beyond the comprehensive threat prevention provided by Check Point, to name a few:

Auto-quarantine of Infected Hosts

Hosts identified by CloudGuard as infected can be automatically isolated and quarantined. In order to accomplish this, CloudGuard tags the infected hosts and shares this information with the ACI fabric. An orchestration platform can also be used to trigger automated remediation services. A threat can be contained quickly, and a remediation service can be applied to the affected hosts.

Context-Aware Security Policies

Integrating Cisco's Application Policy Infrastructure Controller (APIC) with Check Point CloudGuard allows endpoint groups (EPG) to be imported and reused within Check Point security policies. This reduces security policy creation time from minutes to seconds. Any changes or additions to endpoint groups are automatically tracked without administrator intervention in real-time context sharing.

Complete Visibility and Control

CloudGuard for Cisco ACI provides consolidated logging and reporting of threats and security events. Check Point logs are further enriched with ACI context including EPG names and security tags. Furthermore, Check Point SmartEvent offers advanced incident tracking and threat analysis across both physical and virtual datacenter networks.

Centralized and Unified Management

Security management is simplified with centralized configuration and monitoring of Check Point CloudGuard. The traffic is logged and can be viewed within the same dashboard as other gateways. To track security compliance across the data center network, security reports can be easily generated via the same dashboard. A layered approach to policy management allows administrators to segment a single policy into sub-policies for customized protections and delegation of duties per application or segment. By integrating all aspects of security management such as policy management, logging, monitoring, event analysis, and reporting into a single dashboard, security administrators receive a holistic view of security posture across their organization.

Whenever an organization converts its data centers into a hybrid model, it must be able to manage a dramatic increase in lateral network traffic between applications, both dynamically and automatically. Check Point CloudGuard Network Security for Cisco ACI provides these capabilities delivering comprehensive and dynamic security specifically architected for Cisco ACI-enabled data centers.

Maestro Hyper-Scale Virtual Security Fabric

Check Point Maestro is an easy-to-manage Hyperscale network security solution that lets businesses leverage current hardware investments and maximize appliance capacity. By using Maestro, organizations can streamline workflow orchestration and scale up their existing Check Point security gateways on demand - just as they can spin up servers and compute resources on the public cloud.

Maestro orchestration enables companies to expand from a single Check Point gateway deployment to 52 gateways in minutes, providing unprecedented levels of flexibility and enabling massive more than 1 Terabit per second firewall performance.

Check Point Maestro brings the scale, agility, and elasticity of the cloud to on-premise. Check Point HyperSync technology provides an efficient (N+1) clustering model, maximizing the capabilities of your existing security gateways. It has never been easier to create your own virtualized private cloud by stacking multiple Check Point security gateways together. Group them by security feature set, policy, or the assets they protect and, if necessary, further virtualize them with virtual systems technology. With the Maestro Hyperscale Orchestrator, businesses of all sizes can have cloud-level security on-premise.

The Maestro Web UI or RESTful APIs will help you scale compute to meet your needs while minimizing downtime and maximizing efficiency.

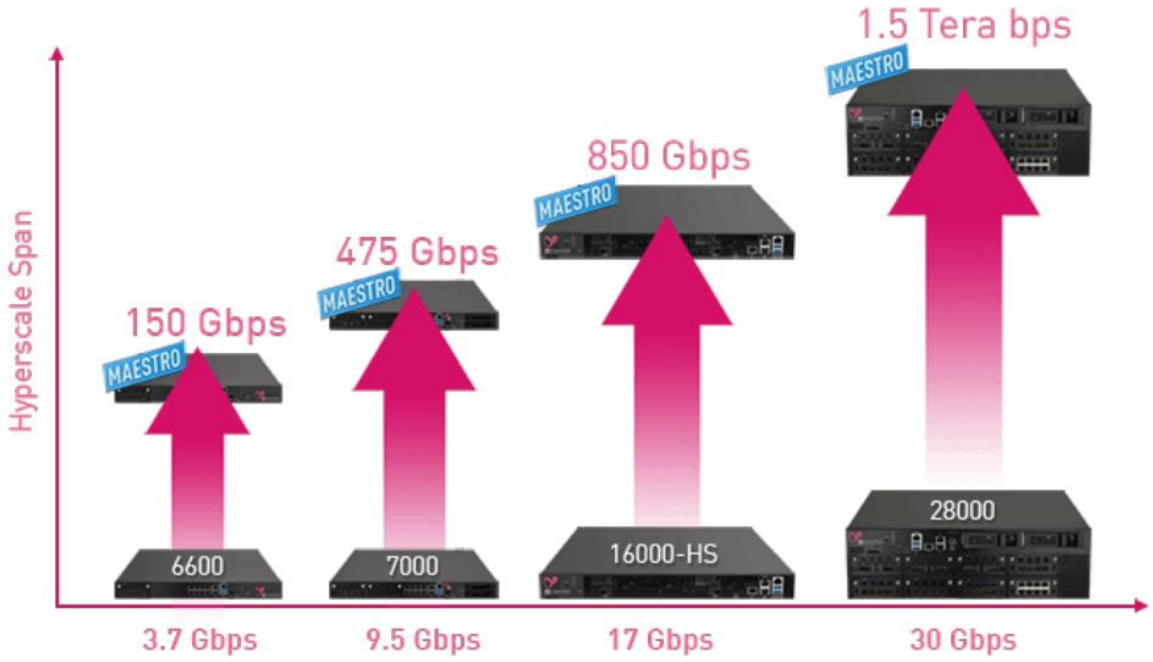

Figure 23: Hyperscale Network Security with Check Point Maestro

Virtual Systems[15] Security Gateways (VSX)

Check Point security gateways are recommended to be deployed inside of the ACI fabric to benefit from the flexibility offered by Cisco ACI and PBR. It is recommended that Check Point Gateways deployed within ACI fabric, whether they are deployed as CloudGuard Network Security Gateways or as standard Check Point physical appliances, be configured in Virtual System eXtension mode (VSX). All standard main-train Next Generation Threat Prevention (NGTP) security gateway blades offer industry-leading features for firewalls and threat prevention. They are utilized to block threats entering the DC and prevent lateral movement between applications inside the data center.

Both physical and virtual gateways are fully supported and are functionally identical, so choosing to deploy appliances or Check Point VM's would depend more on the projected load, customer preference, flexibility requirements, automation stack, and commercial considerations.

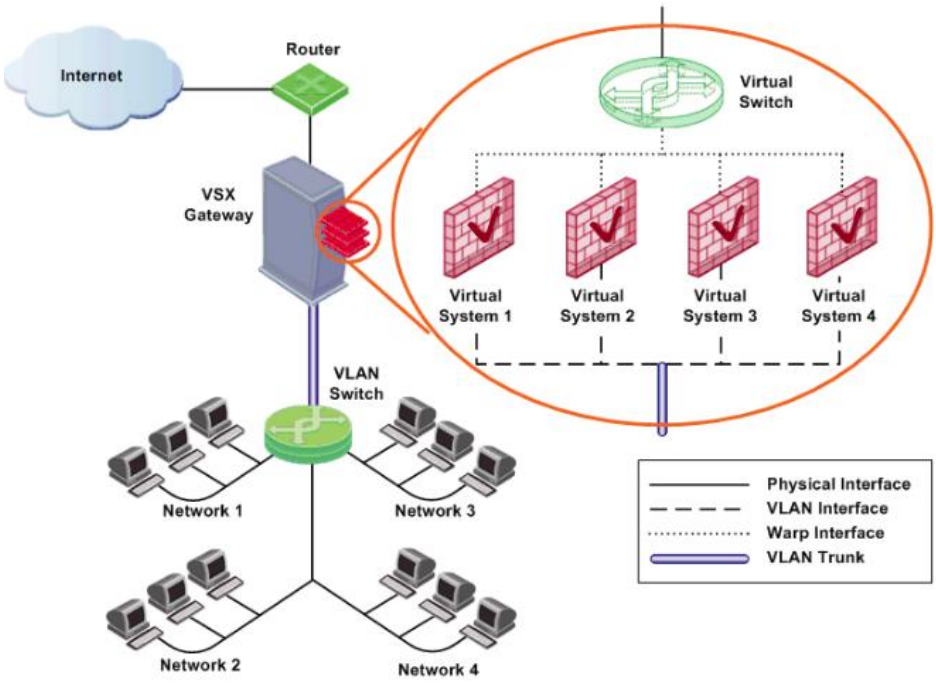

Figure 24: Virtual Systems Security Gateways architecture

Supported Check Point Security gateways

Various Check Point Security Gateway appliances can be integrated with the ACI fabric and managed by the R80 / R81 Security Management Server. They can be clustered and virtualized systems using VSX VSLS (Virtual System Load Sharing) or this can be a standalone deployment, depending on the requirements. The Check Point Maestro product family is also supported. Check Point Maestro deployment has many additional benefits when true hyper-scale performance and the ability to segregate different workloads or tenants into different security groups are required.

[15] Source: Check Point Virtual Systems - URL: https://www.checkpoint.com/downloads/products/virtual-systems-datasheet.pdf

Virtual CloudGuard Gateways

The Check Point CloudGuard Network Security gateways integrate with Cisco ACI Virtual Machine Manager (VMM) to provide L4-7 threat prevention using general-purpose virtualized compute. This is attractive for customers wishing to follow a strict "converged" design pattern over the alternatives

Single and Multi-Pod Overview

The ACI Underlay network, also known as the Physical network, is the composition of all the required protocols that make the overlay network function. The principle of Clos Network defines a circuit switching that arranges a dedicated communications path for a connection between Endpoints for the duration of the connection. This principle is the standard for all modern packet-switched networks focused on Application connectivity.

Some basic principles in the design for all ACI topology:

- Every Spine switch must be connected to every leaf switch
- Every Leaf switch must be connected to every Spine switch
- Leaf Switches must not be connected to other leaf switches and spine switches must not be connected to other spine switches
- Everything that is not a leaf switch must not be connected to a spine.

In the following sections, three separate Check Point designs are presented to illustrate Check Point's integration with Cisco ACI.

Single Pod Design

This is the simplest and most basic deployment of ACI. It consists of a number of spine switches and several leaf switches connected together typically in a single geographical location.

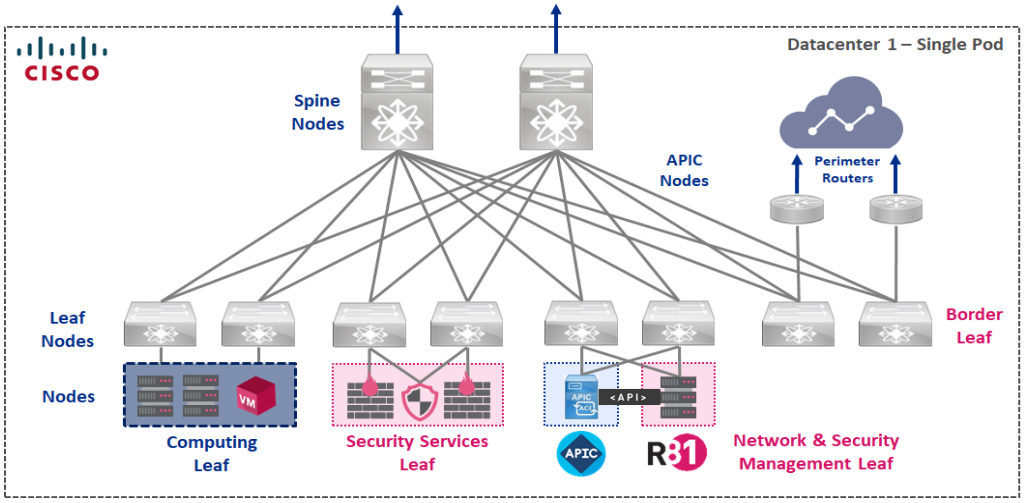

Figure 25: Single Pod Architecture - Check Point and Cisco Systems integration.

Multi-Pod Design

ACI Fabric can scale in a Multi-Pod topology but be still managed by a single APIC cluster. This kind of design deployment would require a special network between the pods - the InterPod Network (IPN). This special IPN is essentially an external network to the fabric and for this reason, it is not managed by the APIC.

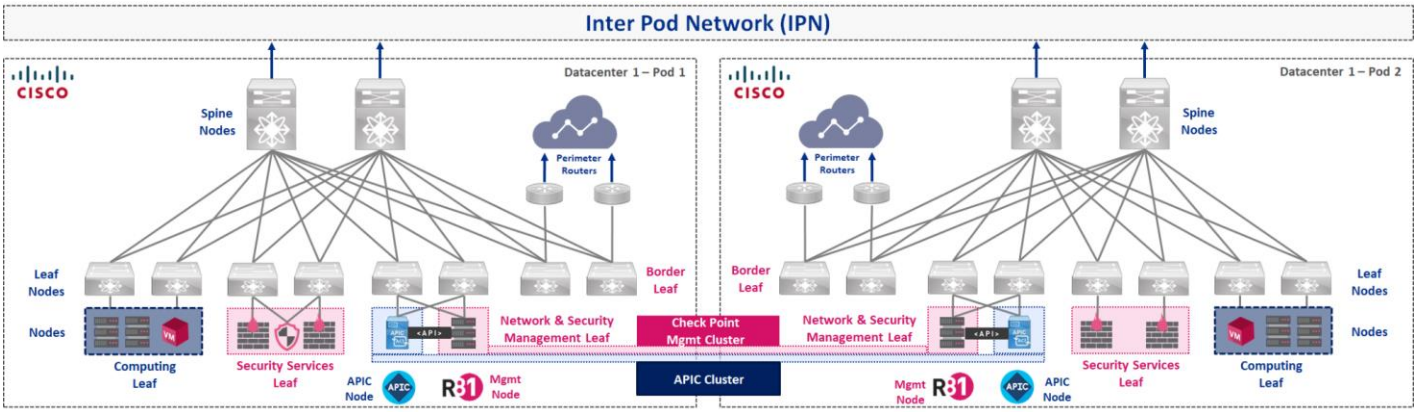

Figure 26: Multi Pod Architecture - Check Point and Cisco Systems integration.

Multi-Site Design

This design consists of multiple separate fabrics, each with its own APIC cluster and an orchestrator that runs on top. Separate fabrics in this deployment can be connected together using L3Out connectivity. A Multi-site Orchestrator (MSO) running on top of it enables administrators to configure the connectivity between sites and to have tenants stretching the sites (stretched bridge domains and EPGs). This kind of architecture can be useful for Multi-cloud environment to stretch Tenants between Private Cloud to the Public Cloud (Azure, Amazon, Google)

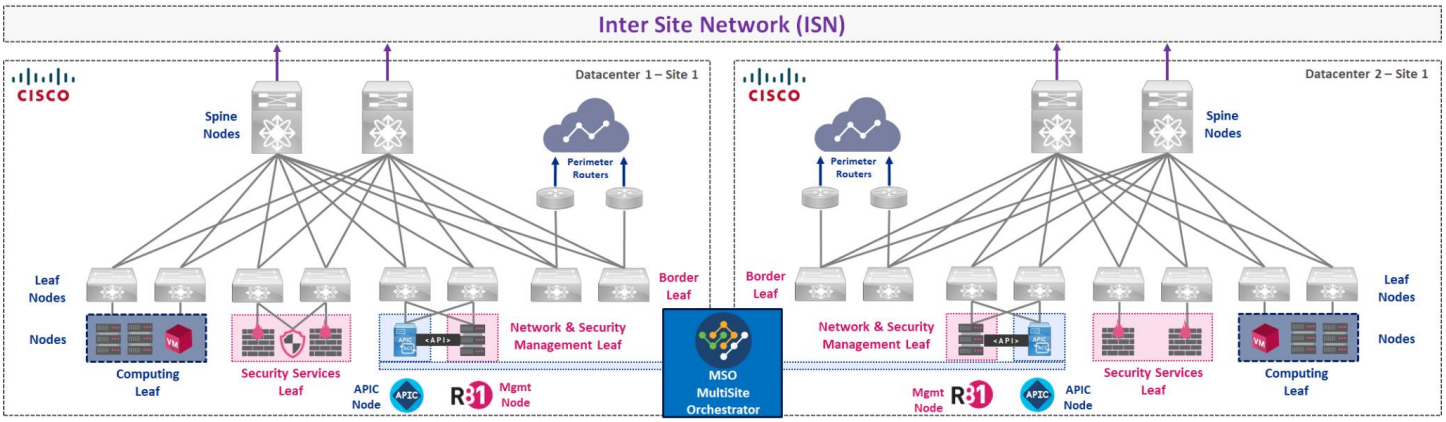

Figure 27: Multi-Site Architecture - Check Point and Cisco Systems integration.

There are other topologies that Cisco ACI supports but are not within the scope of this whitepaper:

- Remote Leaf
- vPod
- cPod
- Cloud Only (Azure/Amazon/Google)

Integrating Check Point Firewalls to the ACI Infrastructure

There are two different supported ways to integrate Check Point firewall with Cisco ACI:

- Policy-Based Redirect (PBR)
- L3 Go-To Mode

In the case of PBR, Cisco ACI technology allows integration with the Check Point Firewall through a technique called service graph deployment based on Layer 4 - Layer 7 Policy-Based Redirect where specific traffic can be redirected for firewall inspection based on application requirements instead of the network path.

L3 Go-To (Routed Mode)

In this deployment option, the CloudGuard Gateway functions as a Layer 3 gateway with a NIC in each network, and traffic is redirected to the gateway by way of a network route. This is typically used to control traffic between different bridge domains and communications outside of the fabric (e.g. the Internet). One common scenario is for Datacenter segmentation between different bridges domains. This deployment scenario is often relevant to situations where Physical segmentation is required and/or specific compliance regulations. This design should be primarily used for securing N-S traffic flow (L3Out deployments) and it is not recommended for securing E-W communication due to its potential complexity.

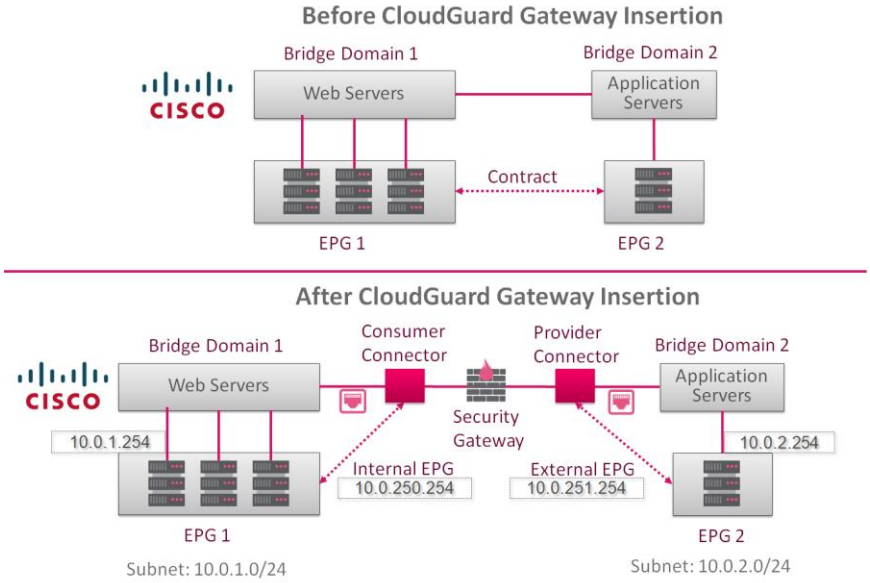

Figure 28: Check Point Service provisioning for L3 Go-To (Routed Mode)

The Check Point L3 Go-To (Routed-Mode) security devices can be deployed in the ACI fabric via service graph. A virtual system instantiated from the L4-L7 routed (GoTo) device is used as the default gateway of endpoints (EPs) contained in the Bridge Domain connected to its connector. With this deployment, firewall can provide all the Check Point capabilities, including HTTPS inspection, NAT, IPsec VPN, etc. for the endpoints connected via the associated bridge domain. However, in this scenario, the firewall network configuration as the default gateway of the application EPs is necessary.

Please note: configuring the device as a default gateway for the EPs requires disabling Unicast Routing capability on the bridge domain (as it does not provide routing services for the EPs). Disabling unicast routing on a Bridge Domain (BD) can cause a design constraint as it prevents the fabric from learning EPs IP addresses.

Policy-Based Redirect (Routed-Mode)

By using PRB, the switch fabric offers the ability to forward traffic between distinct security zones via Check Point Firewalls (Appliances, Virtual Systems, or Security Groups), by avoiding traditional networking parts such as Virtual Routing and Forwarding (VRF) sandwiching or VLAN stitching, firewalls do not need to act as default gateways for servers.

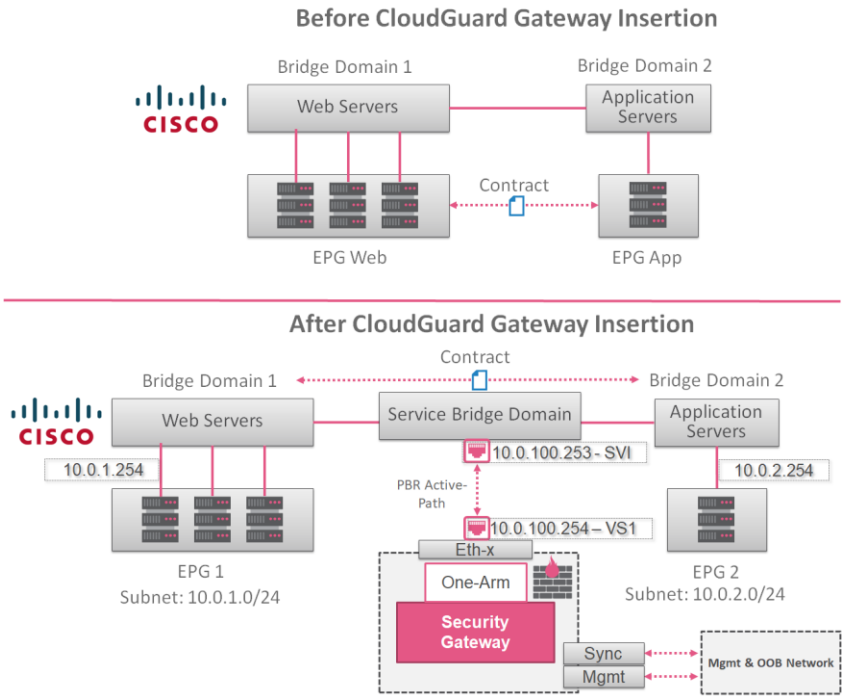

Figure 29: Check Point Service provisioning, before and after insertion

Switch Fabric can send traffic to Check Point Firewalls selectively, for instance, only on specific Layer 4 ports. The firewall inspection can also be transparently inserted into a Layer 2 domain

(Security Service Bridge Domain) without requiring any adjustments to switching or routing configuration.

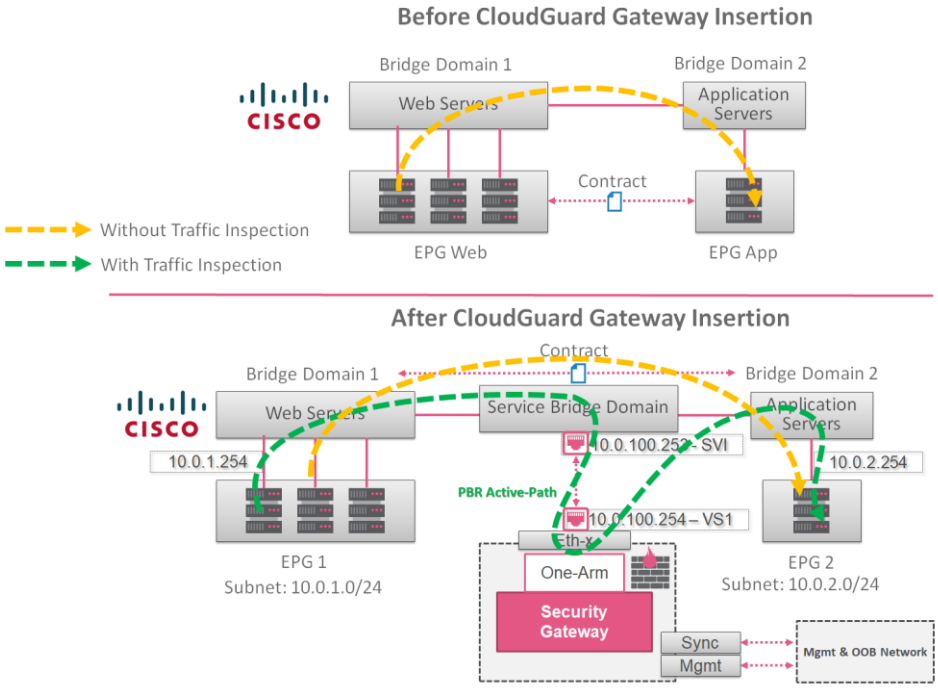

Figure 30: Check Point Service provisioning, traffic flows configured in the Service Graph

Policy-Based Redirect (PBR)[16]

The APIC policy configures and enforces the PBR to control traffic between different EPGs regardless of whether they are on the same IP address subnet or a different one.

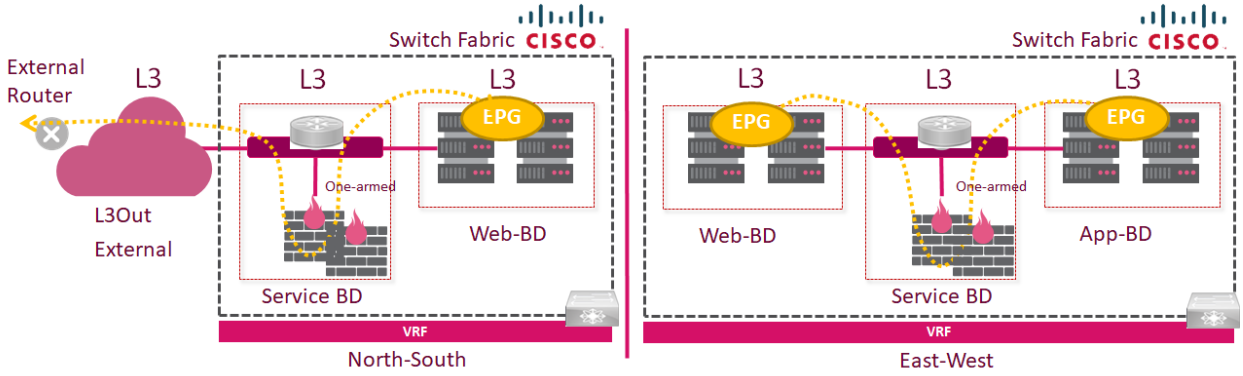

Figure 31: Policy-Based Redirect enables provisioning service appliances

[16] Source: Cisco APIC Layer 4 to Layer 7 Services Deployment Guide, About Policy-Based Redirect - URL: https://www.cisco.com/c/en/us/td/docs/switches/datacenter/aci/apic/sw/1-x/L4-L7_Services_Deployment/guide/b_L4L7_Deploy_ver211/b_L4L7_Deploy_ver211_chapter_01001.html

ACI contracts define which traffic is permitted without inspection done by the gateway. Whenever traffic is permitted by a contract, the administrator can apply PBR based on the service-graph template (SGT) and the specific service node for inspection. Once the traffic is classified at the ingress leaf and the PBR service node is identified, it is switched to the PBR service node rather than going directly to the destination leaf. After processing on the service node, the flow is passed back to ACI which then switches to the destination leaf switch.

The CloudGuard Gateways are deployed on their own bridge domain. The same Check Point Gateway can be configured to secure datacenter internal networks (east-west traffic) as well as network perimeter (north-south) traffic flows. Furthermore, Check Point gateways can be provisioned with advanced PBR configurations to provide an extra level of resilience with active-backup PBR, and additional capacity with symmetric PBR.

Symmetric Policy-Based Redirect

Symmetric PBR allows load-balance traffic and delivers to multiple destinations such as multiple firewalls for the purpose of load distribution. Hash tuple is used for the source IP address, destination IP address, and protocol number. Configuring Check Point Firewalls for flow tracking is essential because the firewall will track both directions of the traffic flow.

In the following diagram, we have an example of Symmetric PBR provisioning to provide these advanced capabilities

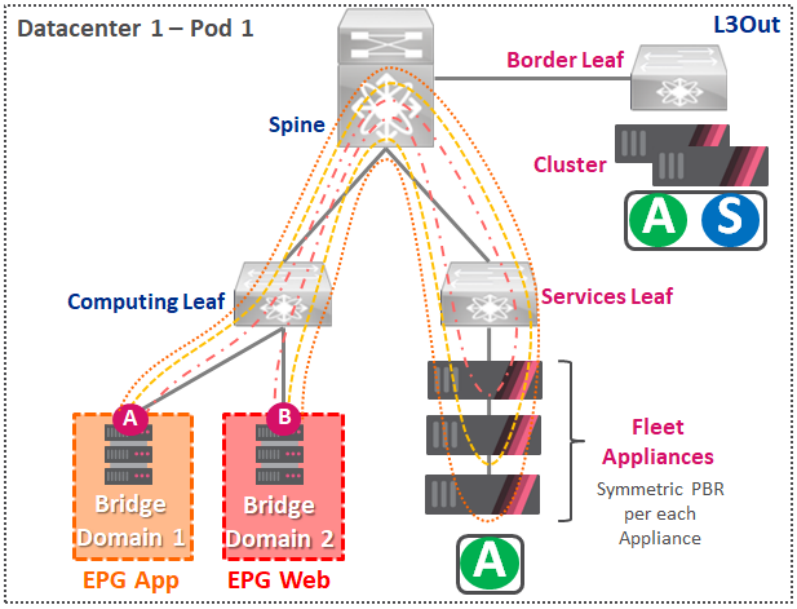

Figure 32: Symmetric PBR for Firewall Load Balancing

Single Pod Security Design

Single Pod overview

In this scenario, there is a Single Pod where multiple different tenants can be created according to the datacenter design. Within the switch fabric, computing, service, monitoring, and border leaf functionality can be distributed among different switches. Check Point Security Gateways can secure North-South and East-West traffic flows utilizing PBR configuration.

To provide high availability and an additional level of resiliency in this environment Check Point firewall Clusters can be deployed in High-Availability configuration.

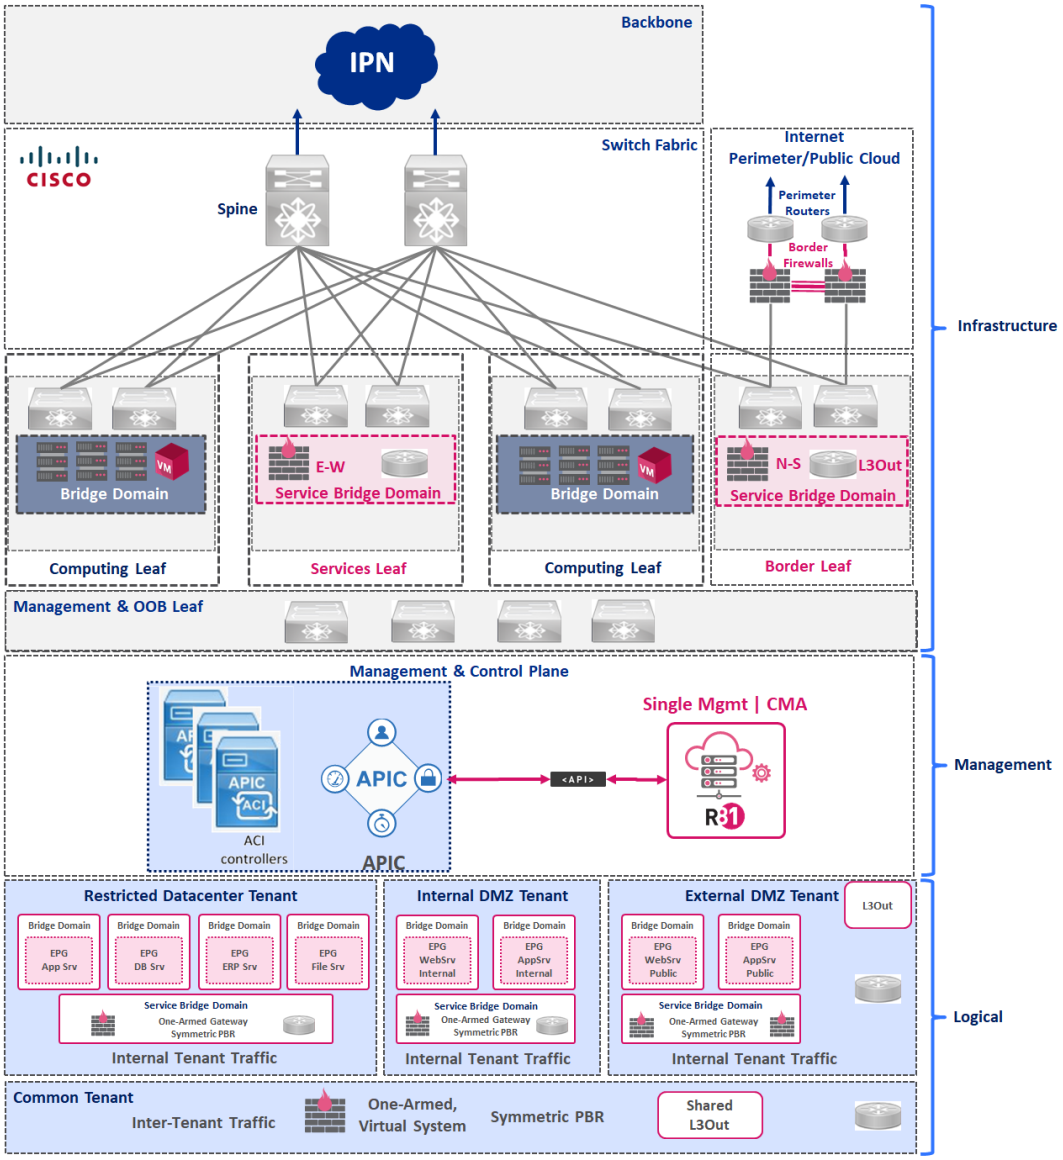

Figure 33: Single POD design

This architecture design has the Check Point Security Gateway cluster deployed in a one-arm configuration as the simplest network topology in order to connect to service leaf switches to secure east-west traffic within a tenant. A different firewall cluster would be connected to border leaf switches to provide north-south traffic flow inspection.

The most important part of the overlay configuration for Leaf and Spine switches is relevant to the logical perspective on segmentation and services separations based on the risk or criticality. This is a list of examples of various types of tenants:

- Restricted Datacenter Tenant
  - Where the Highly-Restricted and Restricted systems should be placed.
  - Databases and other highly transactional systems should be placed
  - Micro segmentation can be configured using the Access-Control that could be generated in the Service Contract.
  - Traffic inspection makes sense only for specific applications and flows, not for all the traffic.
- Internal DMZ Tenant
  - Where the Internal Systems like Intranet or another Business support services should be placed.
  - Traffic inspection makes sense for applications that should be protected for the Ingress and Egress traffic. East-West traffic inspection should be placed only and only if we have interaction with the Restricted Data Center tenant.
- External DMZ Tenant
  - Where the external or public faced services should be placed and security controls should be the most restrictive for Ingress and Egress traffic
- Management Tenant (for In-Band/Out-Band)
  - Where it provides a convenient means to configure access policies for the fabric nodes being accessible and configurable through the API. Management interfaces of the Check Point appliances should be connected in this tenant.
- Common Tenant
  - Where "common services" to the other tenants in the Fabric like Active Directory, DHCP, DNS, NTP and Firewall/IPS.

Check Point Security Appliances for Single Pod

Within single POD design, deployment standard Check Point appliances can be utilized to secure E-W and N-S traffic flows. As explained earlier because of the potential complexity it is recommended to utilize separate clusters in Active-Standby mode for E-W and N-S traffic flows. Also, these physical appliances would be typically connected to different functional leaves.

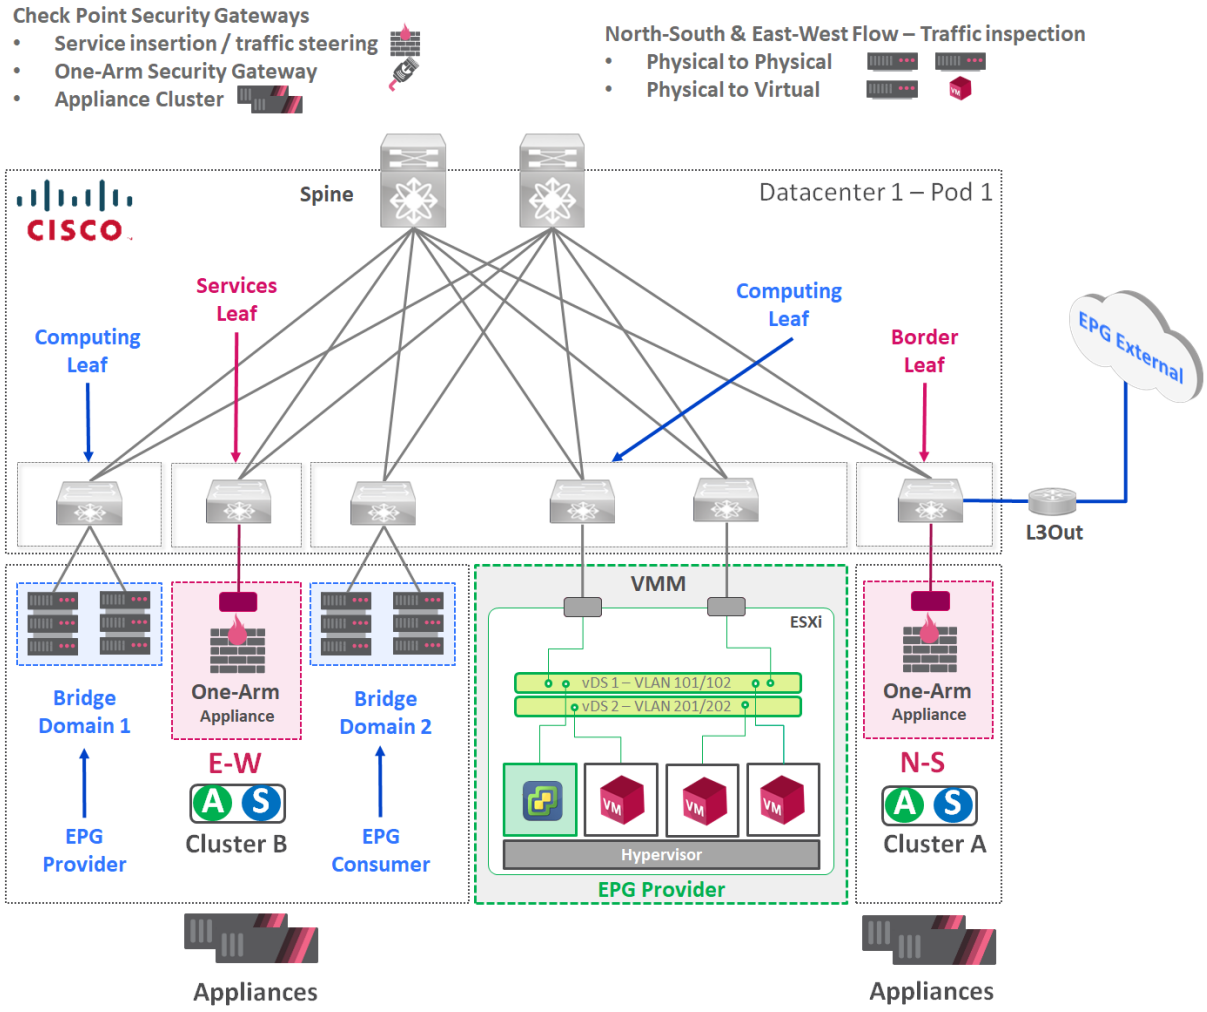

Figure 34: Single POD with physical appliances

VSX Cluster Design for Single Pod Security Deployment

VSX (Virtual System Extension) is a security and VPN solution for large-scale environments, based on the proven security of Check Point Security Gateway. VSX provides comprehensive protection for multiple networks within complex infrastructures. It allows to share access to resources at scale and interact with each other within one physical system using virtual switches or routers.

Check Point Smart-1 appliance typically provides a management layer for VSX deployment where all multiple virtual systems can be easily administered through a single interface.

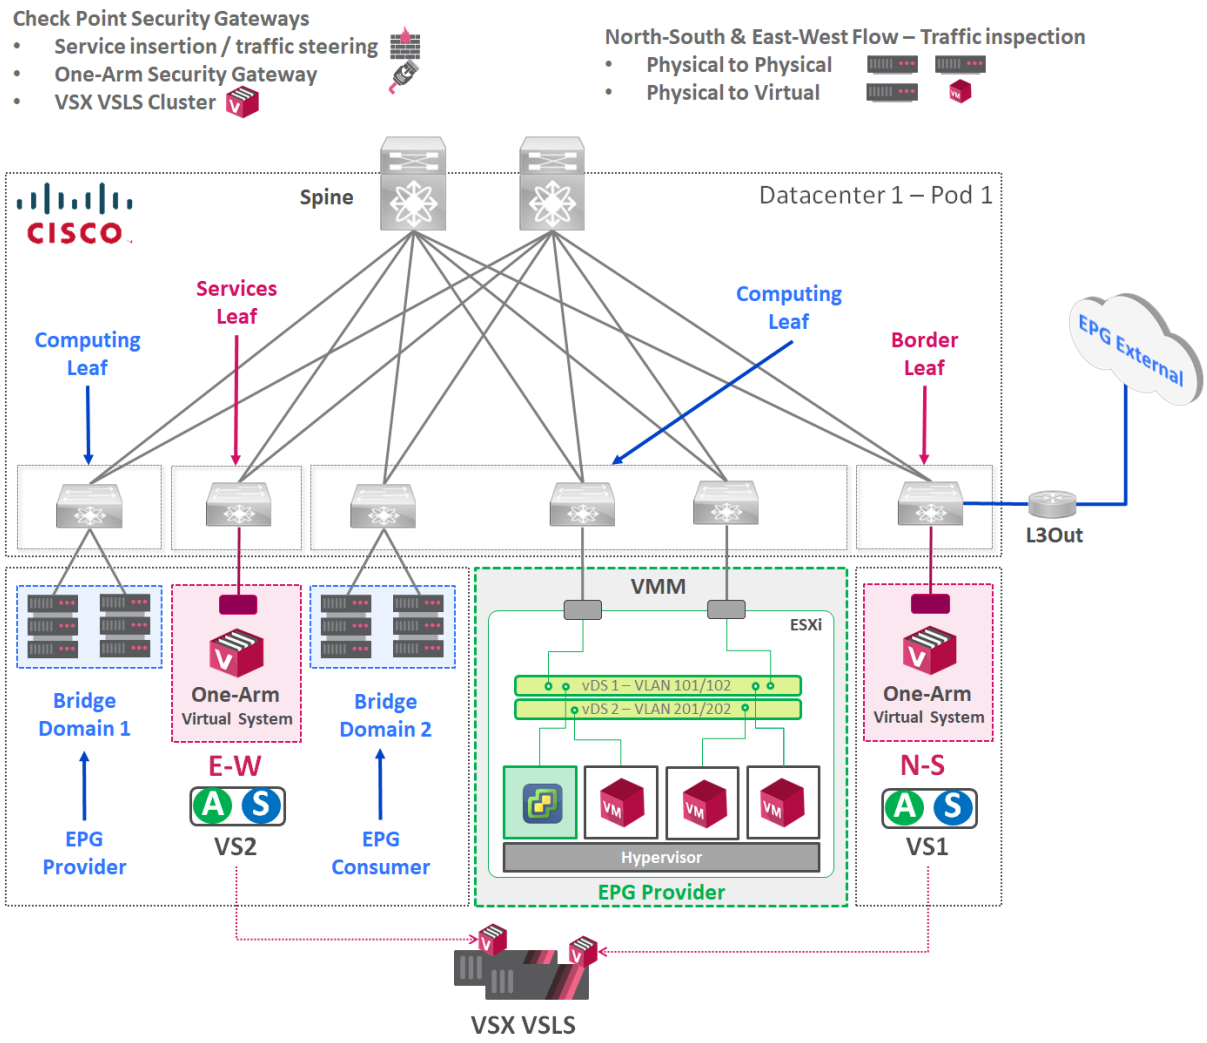

Figure 35: Single POD with VSX/VSLS Cluster

Each Virtual System would act as a full functional Security Gateway, typically protecting specified network segments. In order to reduce overall complexity and contain traffic locally, it is recommended to use internal VSX resources like virtual switches or routers.

Traffic Flows in the Single Pod Architecture

In the following diagram, we have the use case to provide protection with to the N-S and E-W

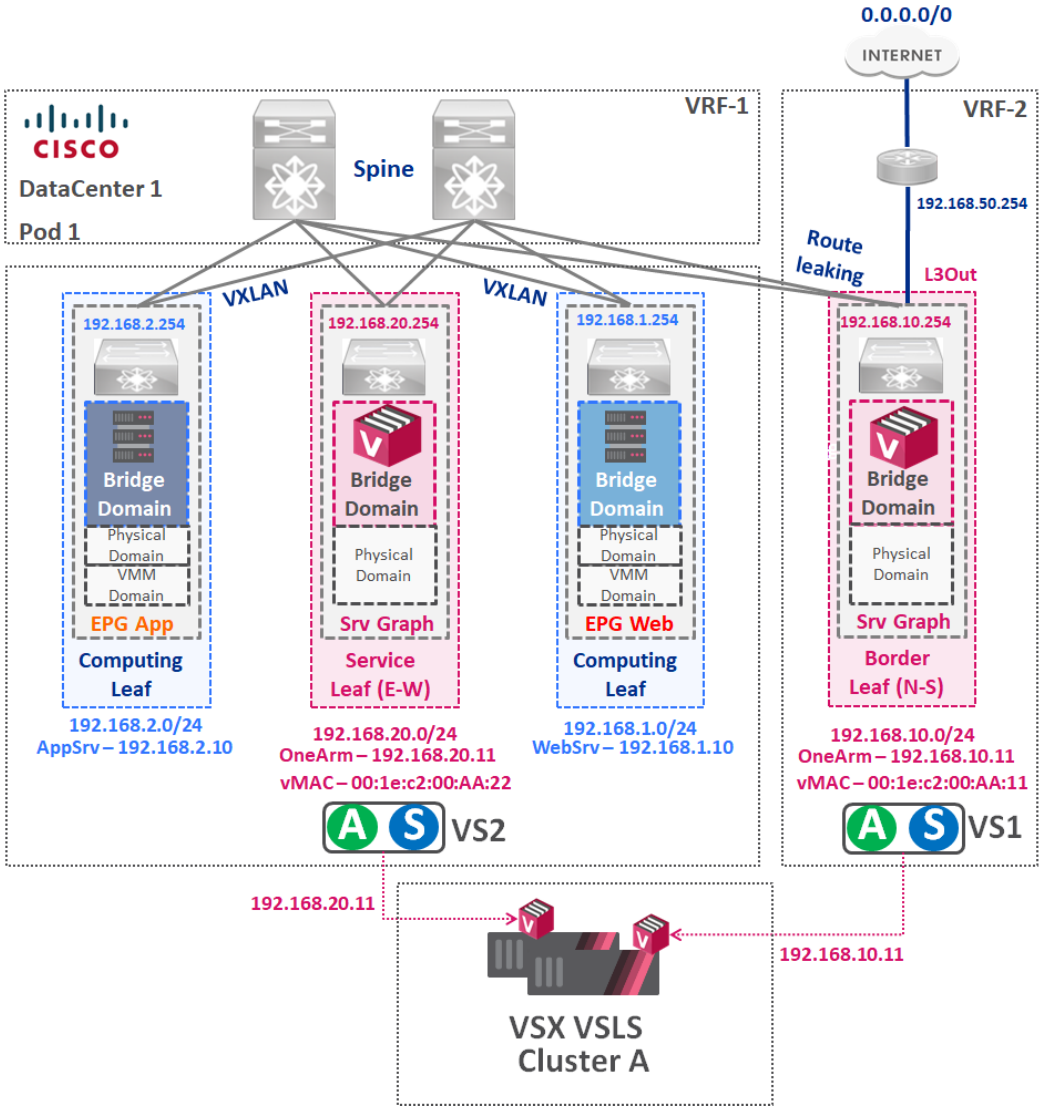

Figure 36: North-South and East-West Firewalls for Single Pod Architecture

Location-based PBR can be utilized for the traffic inspection on the Incoming traffic from the Border Leaf to check Point gateways using the One-Arm mode. This way security services can be inserted without changing routing or redesigning network topology within the datacenter. Concerns regarding network segmentation within the datacenter can be addressed by Cisco ACI's ability to create "virtual" network segments where traffic forwarding can only occur if specifically permitted by the Service Contracts. Service Graph with traffic diversion to Check Point Gateways will provide more advanced and deep traffic inspection capabilities than what can be done within native Cisco ACI.

The following diagrams will demonstrate how different traffic flows can be managed by PBR.

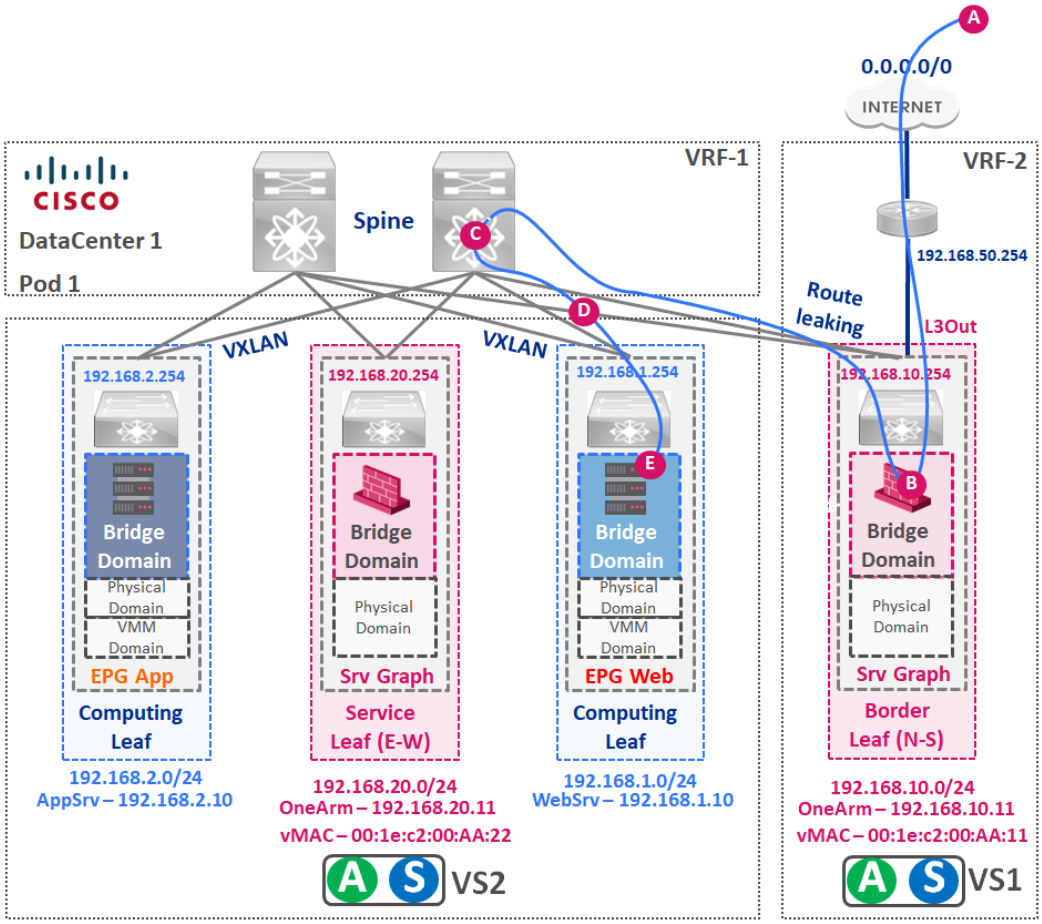

Figure 37: North-South Traffic Flow from L3Out to Web Server through Service Contract and protected by Service Graph (Firewall Service Insertion)

The traffic flow sequence:

A. Traffic is originated from the Internet and is routed to the L3Out (external connection of ACI) in the Border Leaf.
B. Service Contract defines that External EPG should access to the EPG WEB allowing the traffic forwarding, Service Graph redirects the traffic to the Check Point Gateway (Virtual System, Security Group in Maestro or regular appliance Cluster).
C. Once traffic is redirected, processed, and inspected, traffic is forwarded to the Spine.
D. Spine "Knows" that EPG Web is located in one of the Computing Leafs where the traffic is forwarded.
E. Traffic is delivered to the final destination in the EPG Web.

Following is a diagram showing how traffic flows between workloads connected to different leaves in a Switch Fabric.

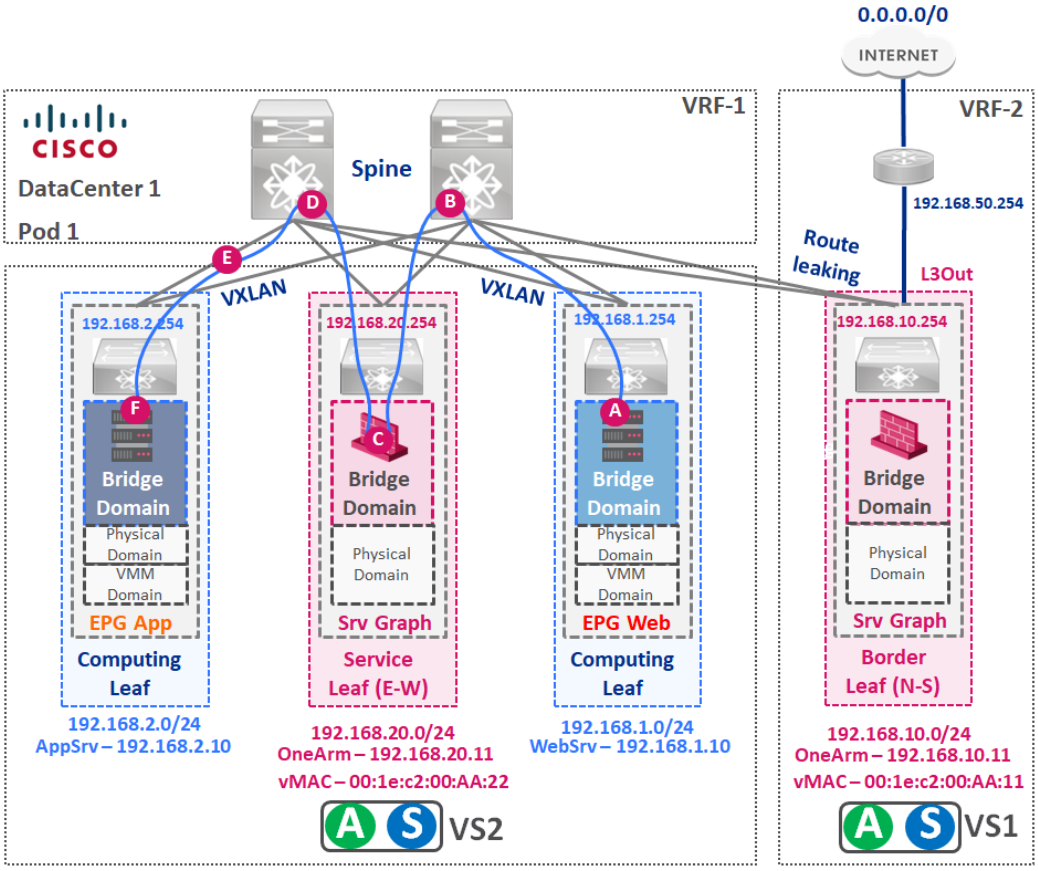

Figure 38: East-West Traffic Flow between EPG Web to EPG App through Service Contract and inspected by Firewall using Service Insertion

The traffic flow sequence:

A. Traffic is originated from EPG Web where the destination is located in the EPG App.
B. Service Contract allows the traffic between the Leafs, however, in the Service Contract we must redirect the traffic first in the Service Leaf.
C. Traffic in the Service Leaf is processed and inspected with the Check Point Security Gateway in One-Arm deployment mode.
D. Once the traffic is allowed according to the Security Policy, it is forwarded to the Spine.
E. Spine knows the computing Leaf where is located the EPG App and forwards the traffic.
F. Traffic is delivered to the final destination in the EPG App.

In the following diagram, we can see how traffic inspection for N-S and E-W can be integrated to provide a full Web Application with their Application Server providing a service. Two different security gateways are suggested in this case. Furthermore, Symmetric PBR can be utilized to scale the number of Security gateways to increase throughput of the security Gateways.

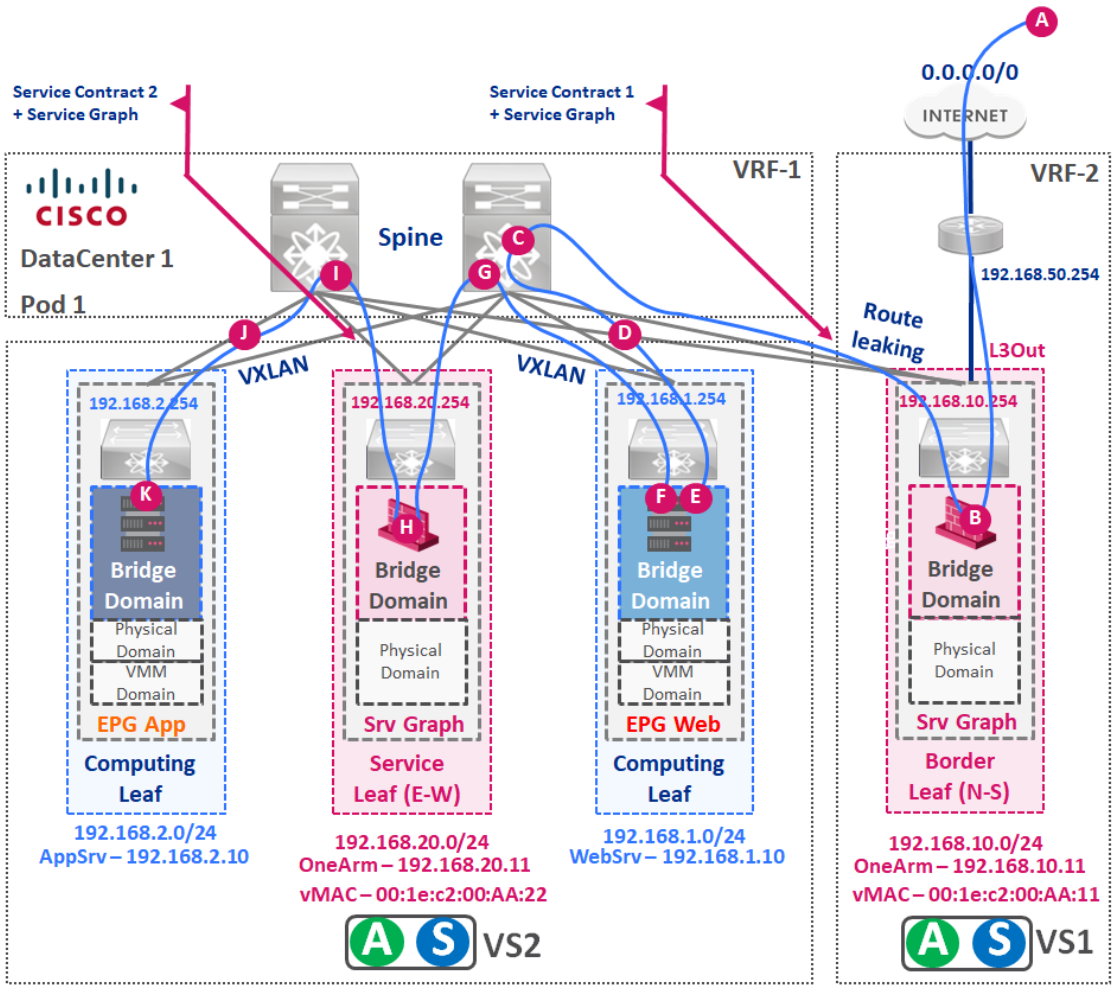

Figure 39: North-South and East-West traffic flows between L3Out, Web EPG Web and App EPG via the Service Contracts and protected using Service Graph (Firewall Insertion for different services)

Traffic flow sequence for two different services:

Service Contract Number 1:

A. Traffic is originated from the Internet and is routed to the L3Out
B. The Service Contract stipulates that External EPG should have access to the EPG web for traffic forwarding, and the Service Graph redirects that traffic to the Check Point Gateway
C. Traffic is forwarded to the Spine after Firewall inspection.
D. Spine forwards traffic to Web EPG
E. Traffic is delivered to the Web EPG.

Service Contract Number 2:

F. Traffic is originated from Web EPG with App EPG as the destination
G. This traffic is allowed according to the Service Contract, but before forwarding to the destination Leaf it needs to be redirected to the Service Leaf in order to be inspected by Check Point Security Gateway (one-arm deployment and using PBR)
H. Traffic is forwarded to the Spine after Firewall inspection

I. Spine forwards traffic to App EPG
J. Traffic is delivered to the final destination (App EPG)

Check Point Maestro for Single Pod

Check Point Maestro offers the industry a new way to maximize appliance capacity and utilize existing hardware investments. Check Point Maestro provides an easy way to manage hyperscaling - Check Point security gateways can be scaled up on-demand, in the same way as a new compute resource can be provisioned in a public cloud. Maestro orchestrator can enable capacity expansion from a single Check Point Gateway orchestration to 52 gateways in a matter of minutes.
It can deliver elastic flexibility by combining the performance of multiple gateways and splitting them into virtual Security Groups re-allocating the amount of processing power depending on the requirements. Each of these Security Groups (SG) would be managed and act as a separate fully functional gateway.

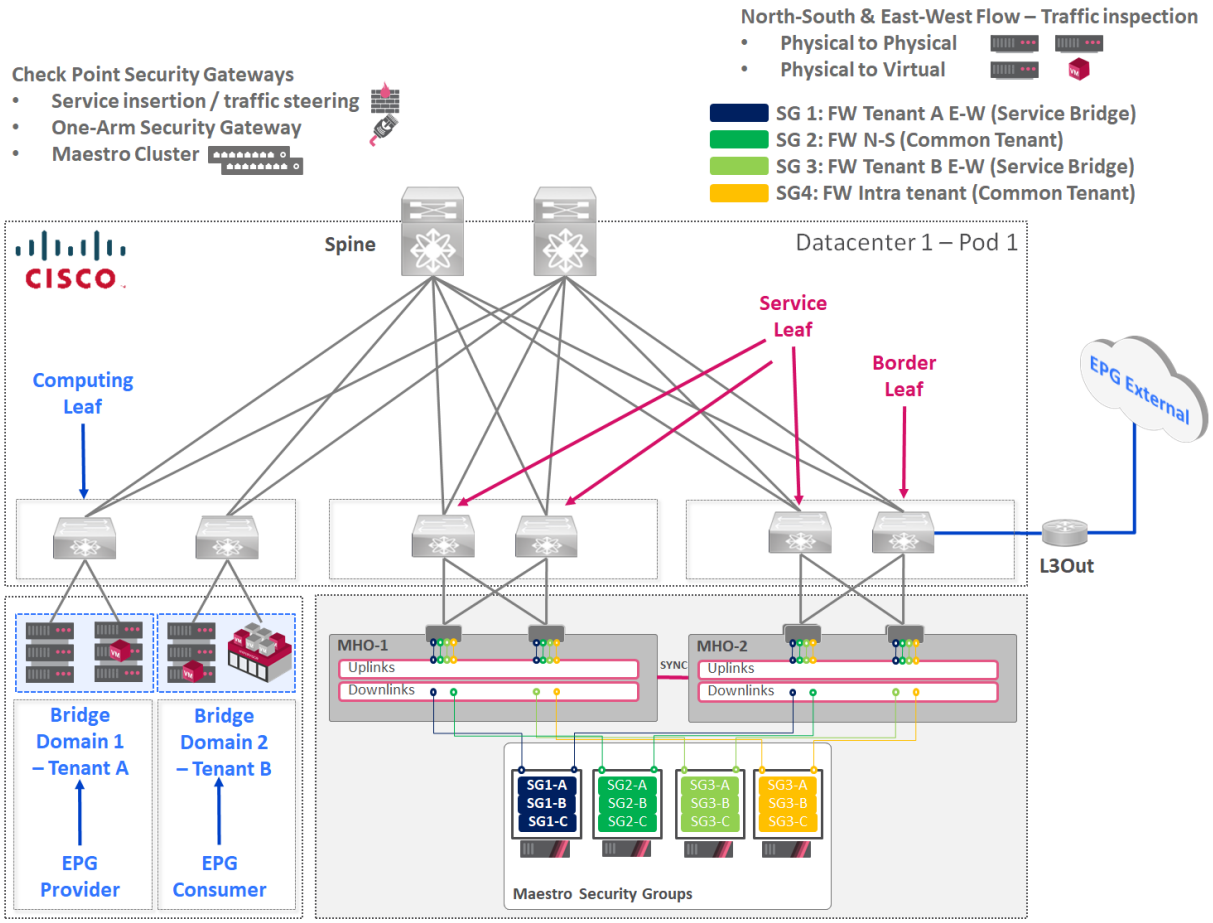

Figure 40: Single POD with Maestro Cluster

The diagram above demonstrates how various security groups can be utilized in order to process different traffic flows.

Check Point Maestro with VSX/VSLS for Single Pod

Combining the power Maestro with Virtual System Load Sharing (VSLS) would deliver even more flexibility and more optimized resource utilization per physical gateway. Each Virtual System can run is a separate firewall with it is own policy and other features which would allow to deploy a very granular level of access and inspection removing a lot of interdependencies associated with one policy for multiple "disconnected" environments and deploying them independently per tenant, zone, environment, the direction of connectivity (egress/ingress), etc.

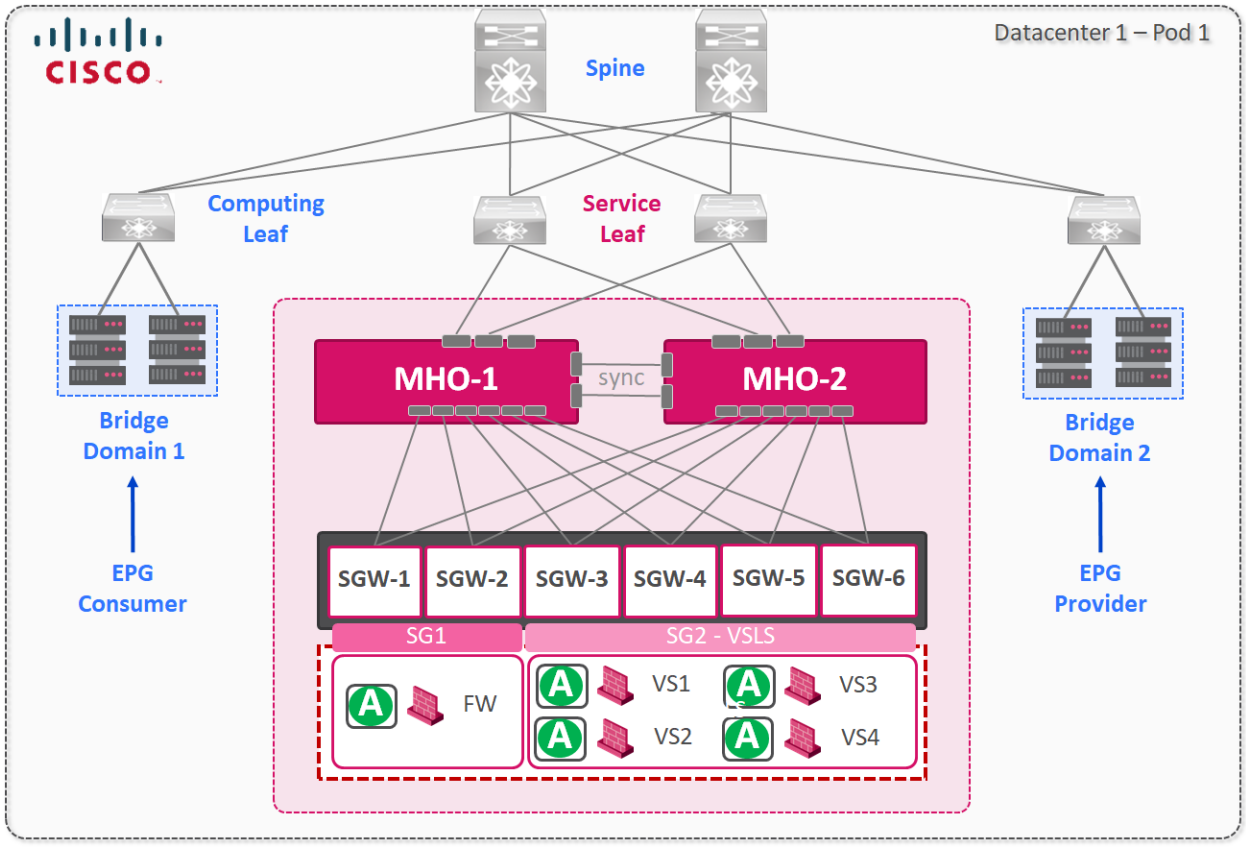

Figure 41: Check Point Maestro Topology with Multiple SGs using regular (left) and VSLS (right) systems deployment

In this architecture deployment, multiple security groups are provisioned within the Maestro fabric. Each of them is dedicated to protecting specific parts or traffic flow within the Cisco ACI fabric. One security group is connected to the service lead and responsible for the intra-tenant security (East-West traffic). Another SG can be responsible for a similar East-West traffic security but in another tenant. A different SG could be responsible for the North-South traffic associated with a traffic flow via border leaf. One more security group can be responsible for security in the common tenant.

Security Appliances Fleet with Symmetric PBR Load Balancing design

Using Symmetric Policy-Based Routing with a fleet of appliances bundled into one virtual system would deliver Telco-Grade Technology combining up to 32 gateways into N+1 clustering for capacity management with more than 1Tbps throughput.

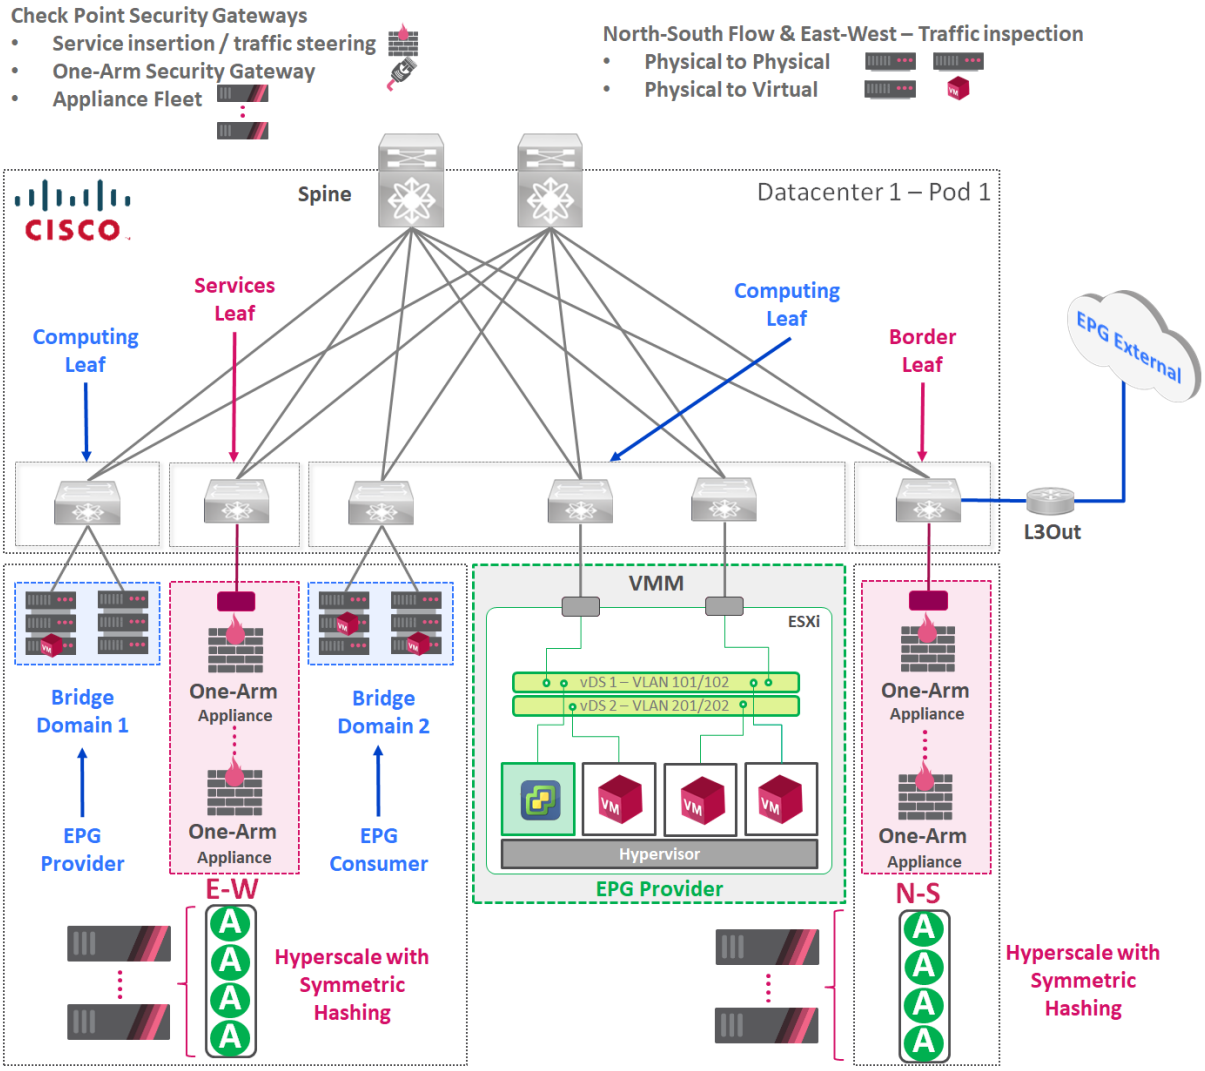

Figure 42: Hyperscale capabilities with a Fleet of Appliances and Symmetric PBR

While symmetric PBR can be useful to accommodate large deployments, it has a limitation, in the event of a firewall failure, existing connections will not survive the failover and they will need to be reestablished due to the fact that traffic load sharing is managed outside of the firewall's control plane.

Multi-Pod Security Design

ACI Multi-Pod represents the natural evolution of the original ACI Stretched Fabric design and allows to interconnect and centrally manage separate ACI networks. It is typically deployed as a single APIC cluster to manage all interconnected ACI networks.

The same APIC cluster can manage multiple Pods and to increase resiliency the various controller nodes (that make up the cluster) can be deployed across different Pods. Check Point and ACI integration allow one firewall management server to be fully integrated with multiple POD or multiple ACI fabrics as in both cases management platforms communicating with each other via API have full visibility of the environment.

For example, in the following multi-POD scenario there are two Cisco ACI PODs interconnected via the IPN (InterPod Network). As with a single pod design, multiple tenants can be accommodated in this environment, however, there are significant variations in how the HA model can be implemented.

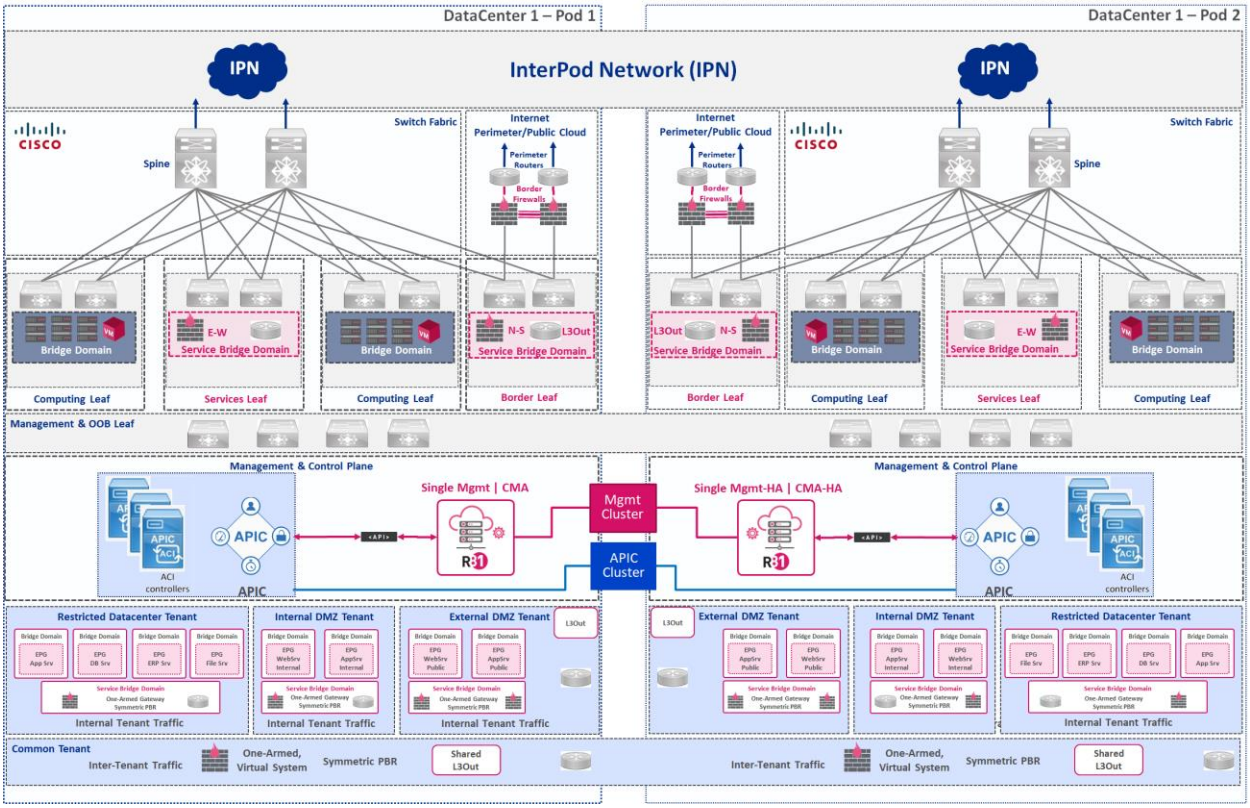

Figure 43: Multi-Pod Architecture

In order to prevent any potential complications associated with returning traffic or asymmetric routing, firewall deployment in a multi-Pod environment requires special considerations concerning firewall HA state configuration and connections synchronizations.

The following design patterns can be applied to Multi-Pod deployment:

- Active-Standby FW service node pair stretched across pods
- Active-Active FW service node cluster stretched across separate pods
- Independent Active-Standby FW service node pair in each pod

In the case of *Active-Standby FW* cluster deployment with a stretched firewall cluster across two PODs traffic always flows via the Active cluster member and the whole design and traffic flow is relatively straightforward.

However, there could be a potential issue with the additional overuse of the IPN network if the point of traffic initiation and the destination are located within a POD where the firewall module is not Active. In this case, IPN would be used for request and reply traffic forwarding between two PODs. This can result in additional latency and IPN bandwidth saturation.

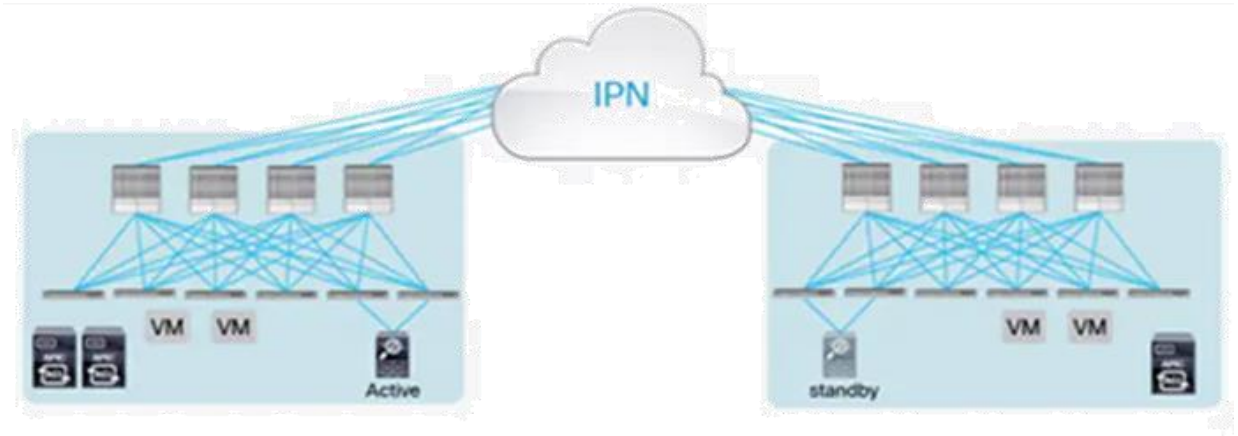

Figure 44: Network services deployment[17]: Active-Standby Firewalls pair stretched across pods - Source: Cisco Systems

Active-Active stretched cluster deployment provides an active firewall module in each POD. Traffic traversing firewall would stay within local POD as long as both source and destination belong to the POD. Firewall cluster members would process traffic independently synchronizing connections across the cluster.

Please note that traffic is not going to be load-balanced between the cluster members. Also, in some scenarios hairpin traffic forwarding would still be occurring. Additional IPN capacity planning would be required to avoid any capacity issues in a fail-safe scenario.

[17] Source: Cisco ACI Multi-Pod and Service Node Integration White Paper, ACI service integration - URL: https://www.cisco.com/c/en/us/solutions/collateral/data-center-virtualization/application-centric-infrastructure/white-paper-c11-739571.html#CiscoACIserviceintegration

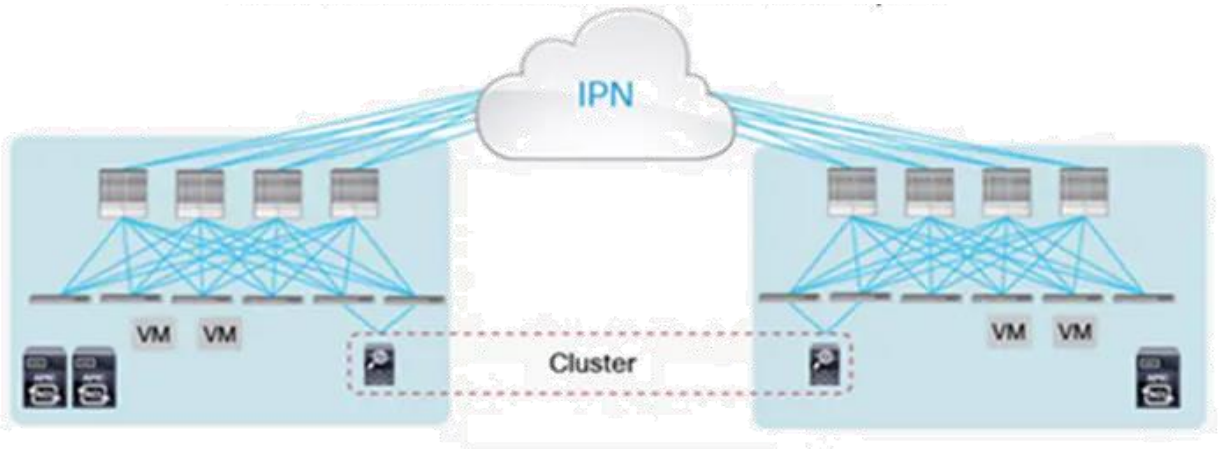

Figure 45: Network services deployment[18]: Active-Active Firewalls pair stretched across pods - Source: Cisco Systems

The last scenario is when there is an *Active-Standby cluster in each POD*. This design pattern provides an extra level of resiliency, but it would require some additional routing and NATing considerations to avoid problems caused by asymmetrical routing and hairpin traffic flows.

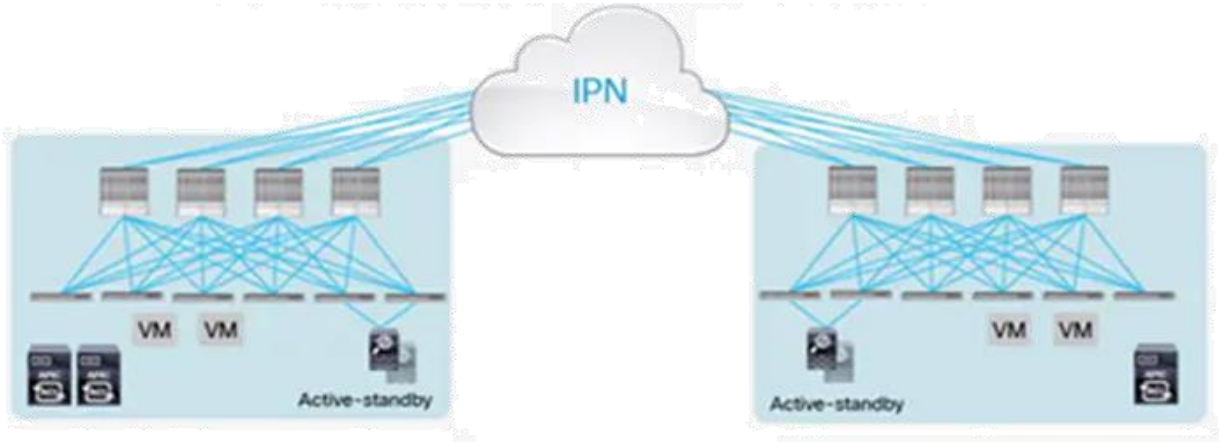

Figure 46: Network services deployment[19]: Active-Standby Firewalls per each pod - Source: Cisco Systems

As highlighted above there is a variety of ways of how firewalls in High Availability can be deployed in the multi-pod environment. The final decision should be made based on the specific requirements regarding the level of resiliency within the Pod as well as considerations regarding additional latency and bandwidth utilization of the IPN link.

---

[18] Source: Cisco ACI Multi-Pod and Service Node Integration White Paper, ACI service integration - URL: https://www.cisco.com/c/en/us/solutions/collateral/data-center-virtualization/application-centric-infrastructure/white-paper-c11-739571.html#CiscoACIserviceintegration

[19] Source: Cisco ACI Multi-Pod and Service Node Integration White Paper, ACI service integration - URL: https://www.cisco.com/c/en/us/solutions/collateral/data-center-virtualization/application-centric-infrastructure/white-paper-c11-739571.html#CiscoACIserviceintegration

Multi-Pod Security Design with dedicated Bridge Domains

This is a scenario where there is a Multi-Pod design with two active pods connected via IPN, however, networks in each POD remain dedicated to the local Bridge Domains and VRFs are deployed in a non-stretched topology. Normal PBR node can be used for firewall service insertion. Check Point recommends deploying Firewalls in the Active-Standby mode per pod in this scenario to ensure proper connectivity for L3Out and East-West traffic flows handled by two different firewall clusters running in high availability.

Depending on the application high availabilities design requirements the service graph configuration for E-W traffic inspection can be done by a firewall closest to the consumer or the provider side of the flow (Cluster A-2 or Cluster B-2).

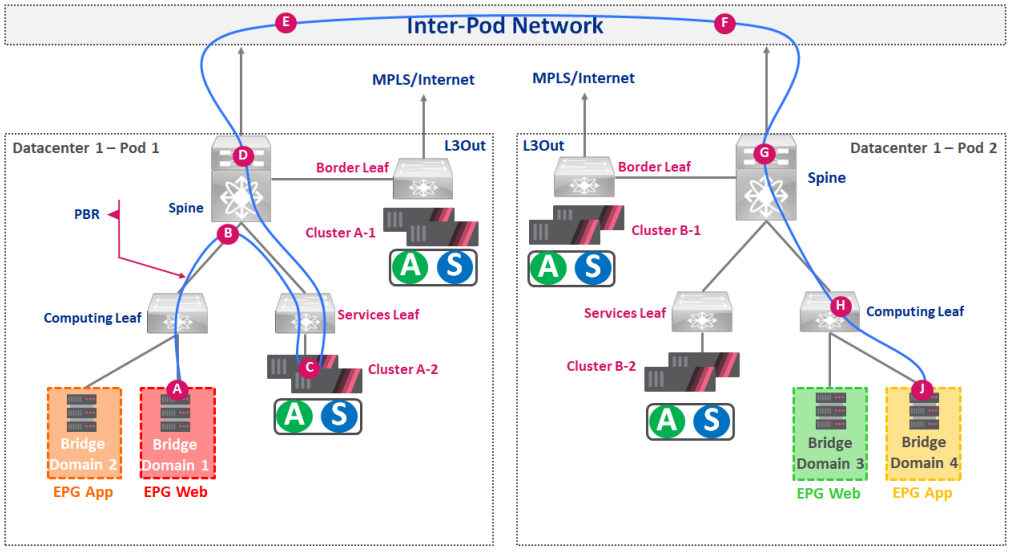

Figure 47: Multi-Pod Architecture – Inter-pod flow option A

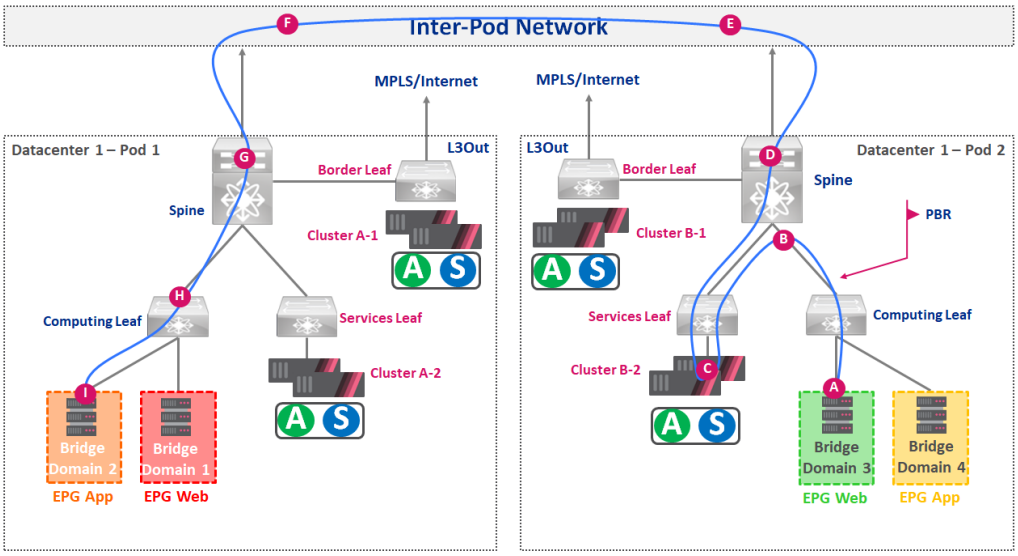

Figure 48: Multi-Pod Architecture – Inter-Pod flow option B

The recommendation would be to use firewall cluster closest to the source and local to the pod, however this is not mandatory.

Which of the firewalls is going to be utilized in the service graph is predetermined in the PBR policy configuration in the Device Selection Policies under L4-L7 Policy-Based Redirect option.
The returning traffic will be automatically forwarded by ACI using PBR to the same firewall cluster.

Maestro design with Active/Standby MHOs per pod

Using Maestro in Multi-Pod topology, similarly to Single Pod environment, can deliver elastic flexibility by combining the performance of multiple gateways and splitting them into virtual Security Groups re-allocating the amount of processing power depending on the requirements. Each of the Security Groups (SG) would be managed and act as a separate fully functional gateway in each Pod.

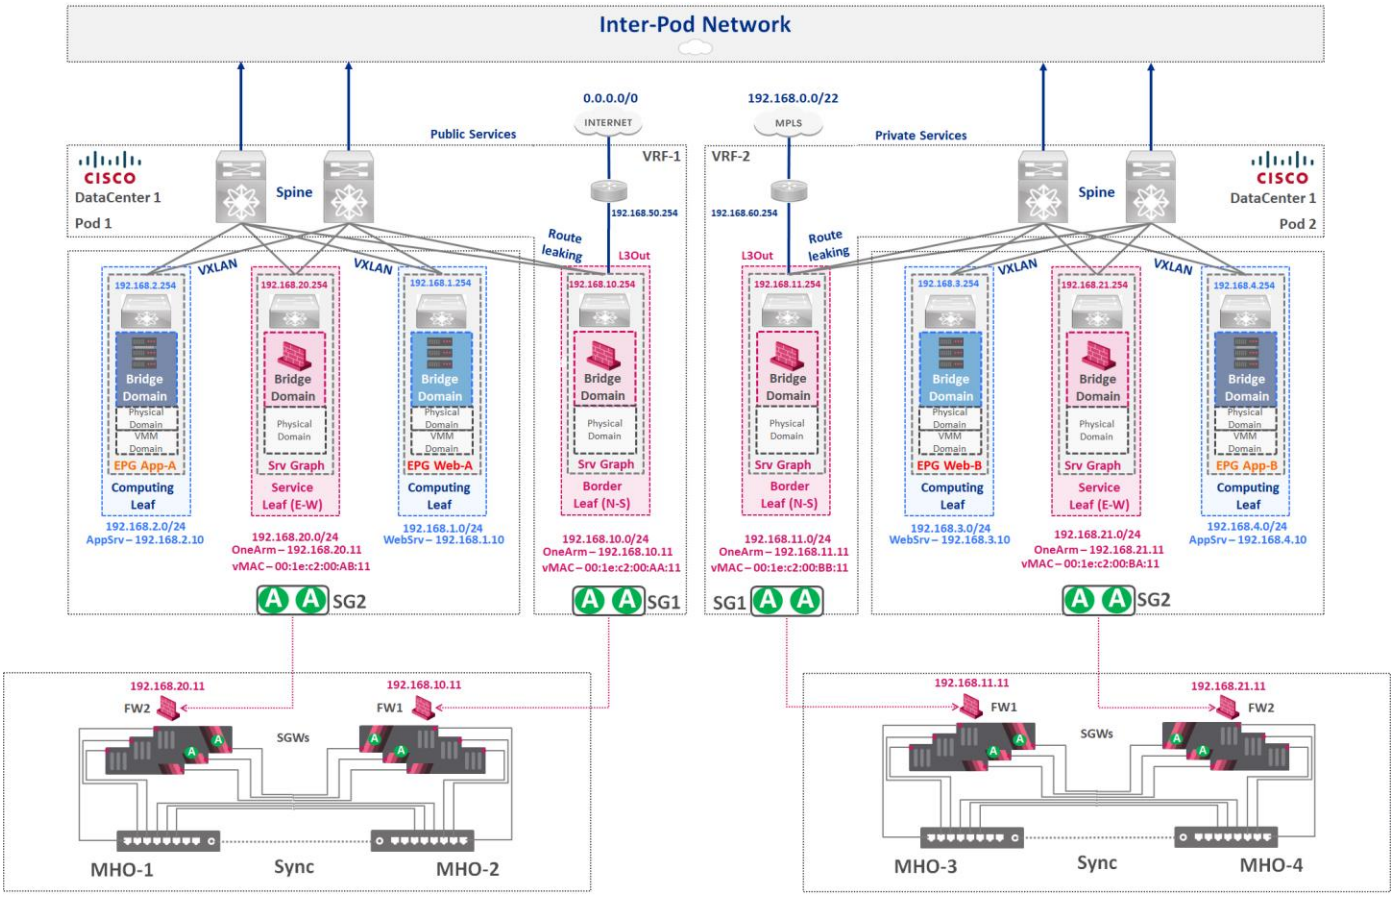

Figure 49: Multi-Pod Architecture – Network topology

VSX/VSLS design with the cluster per pod

Deploying Virtual Systems in Multi-Pod environment would allow an extra level of scalability and granularity when each Virtual System acts as a fully functional Security Gateway protecting a specific network segment. Virtual System Load Sharing would manage and control traffic distribution between Virtual Systems

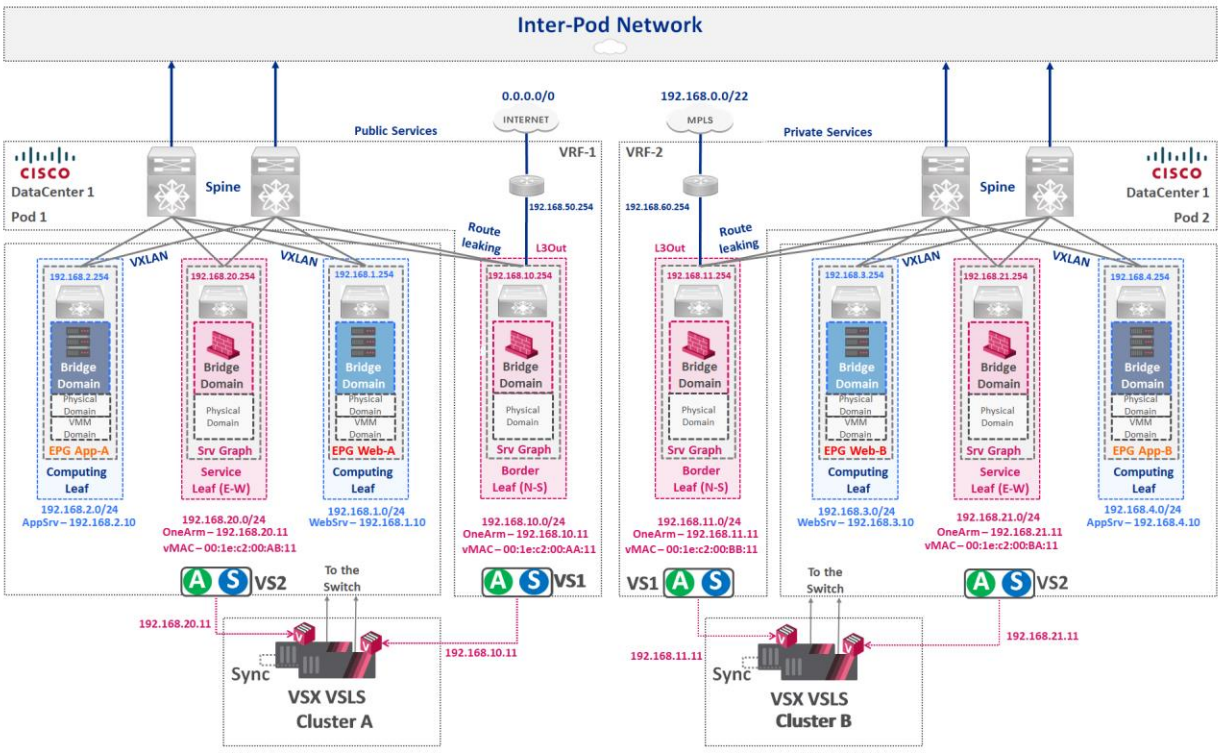

Figure 50: Multi-Pod architecture with VSX VSLS Cluster

Maestro plus VSLS Cluster design with MHO Cluster

When delivering even more flexibility and more optimized resource utilization per physical gateway is required combining the power Maestro with Virtual System Load Sharing (VSLS) should be considered. This would allow delivering an ever so granular level of access and inspection without interdependencies associated with one policy for multiple environments. Virtual System Load Sharing would manage and control traffic distribution between Virtual Systems.

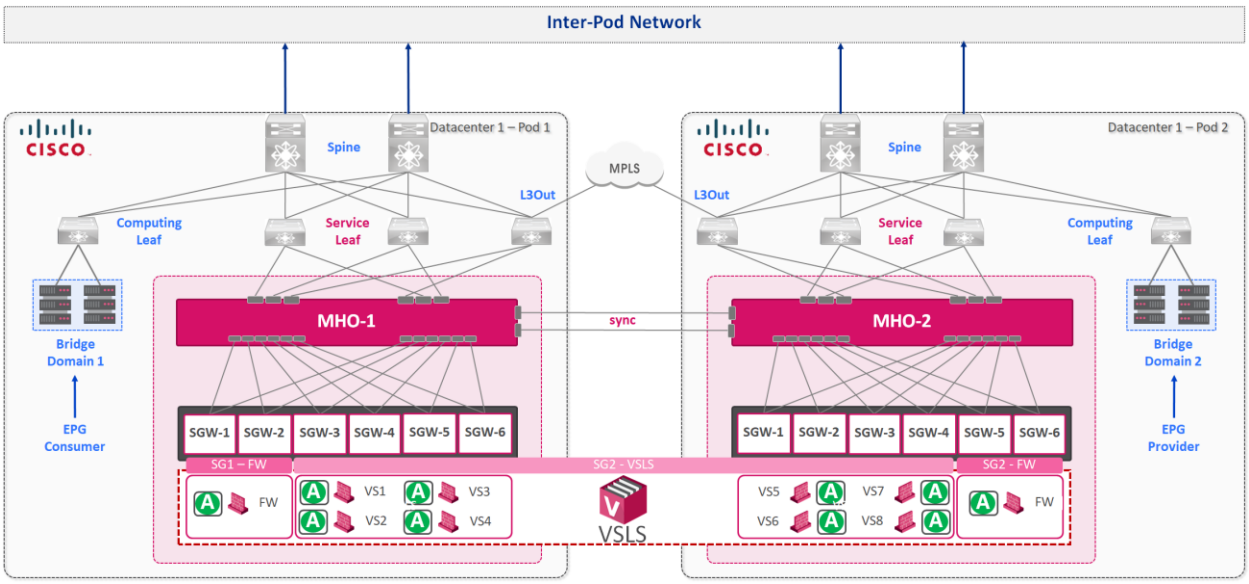

Figure 51: Multi-Pod Architecture with Maestro Cluster

This scenario also addresses the challenge of the distance between two pods deployed in different geo-locations, where it would be impossible to connect security gateways (SG) with DAC cables to each Maestro Orchestrator (MHO). Among the Security Groups, only those in the same pod were connected to the MHO representing that pod. The MHOs in different Pods are in sync typically via dual fiber optic links to providing connections synchronization within Maestro Cluster.

VSLS would be provisioned on top of security groups with MHO being Active only in one of the Pods. Traffic is directed by service graph a certain Virtual System (VS) using SPBR to keep proper proximity function and avoid a hairpin type of issues. A secondary (backup) firewall module for a particular VS would be utilized in case of a failover.

Multi-Pod Security Architecture with stretched Bridge Domains

The Multi-Pod scenario with stretched Bridge Domains is based on the Active-Active pods' concept where resources and applications networks are being stretched across multiple pods and remain active on both sites. Such design brings a lot of value in terms of resource flexibility and a well-balanced environment leveraging all Cisco ACI and computing infrastructure resources, however, at the same time, it brings some security design challenges – like Hairpin and connection drop by non-synced FW in asymmetric traffic flow.

Multi-Pod with stretched networks challenges

In the below example SPBR is configured with Service Graph policy set to have two firewalls (FW-A and FW-B) to be service nodes for connections between two bridge domains (BD1 and BD2) stretched across two PODs (Pod 1 and Pod 2).

For the East-West traffic flow the best case scenario would be when the closest firewall (Cluster A) is chosen to process the traffic like shown below:

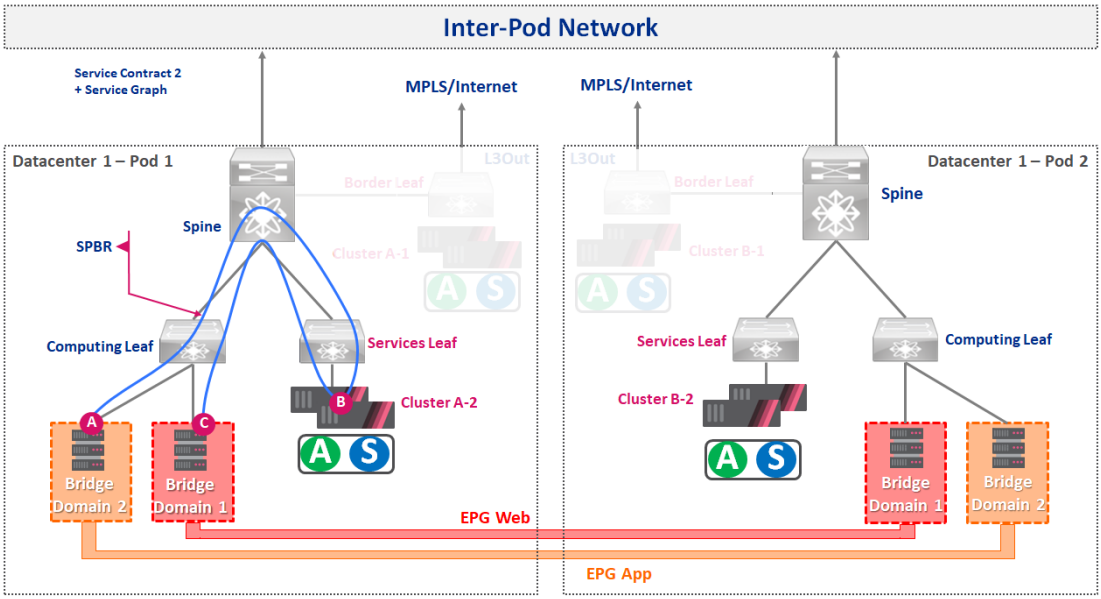

Figure 52: E-W Traffic inside the Pod with SPBR in the best case scenario

However, due to the way SPBR selection mechanism works, the worst-case scenario would be when the nearest firewall not chosen to handle connectivity, thus creating an alternative traffic path via IPN with a hairpin scenario (via Cluster B in a different Pod) is.

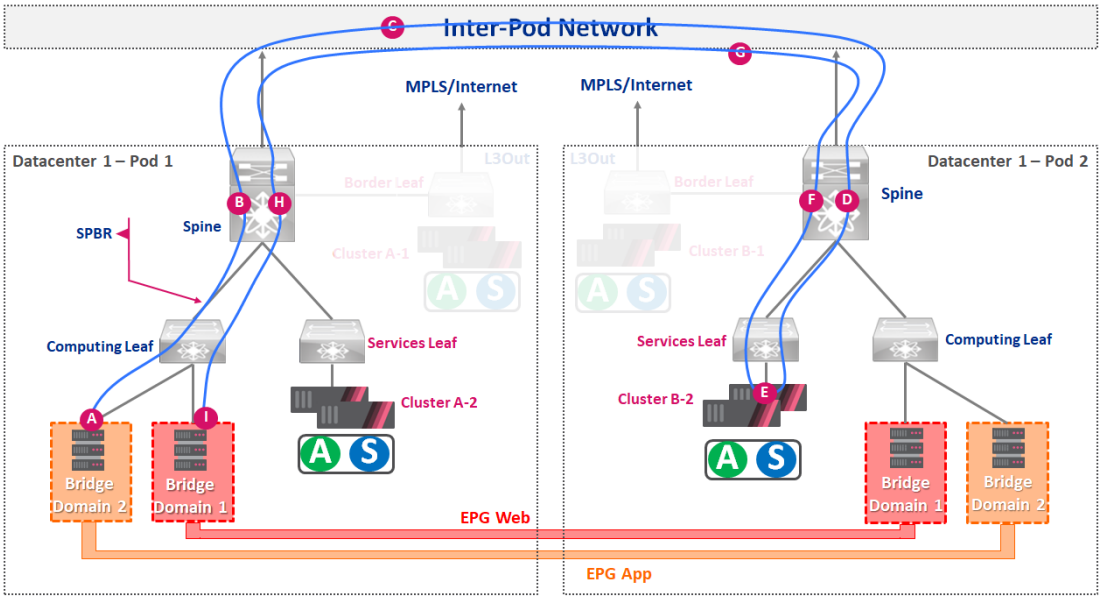

Figure 53: E-W Traffic between Pods with SPBR in the worst case scenario (hairpin)

Location-based PBR (LPBR) deployment scenarios and challenges

Location-based PBR (LPBR) – can be used to improve the proximity selection. LPBR is designed to select the closet FW based on the traffic source. One of the most common use cases is leveraging LPBR for L3Out (North-South) connectivity from the POD to the Internet or MPLS wider network. In this case, when connections are initiated from BD1 on POD-1 the FW-A will be used on POD-1, and sessions from endpoints located in BD2 in POD-2 will be forwarded through the closest FW-B located in the same pod.

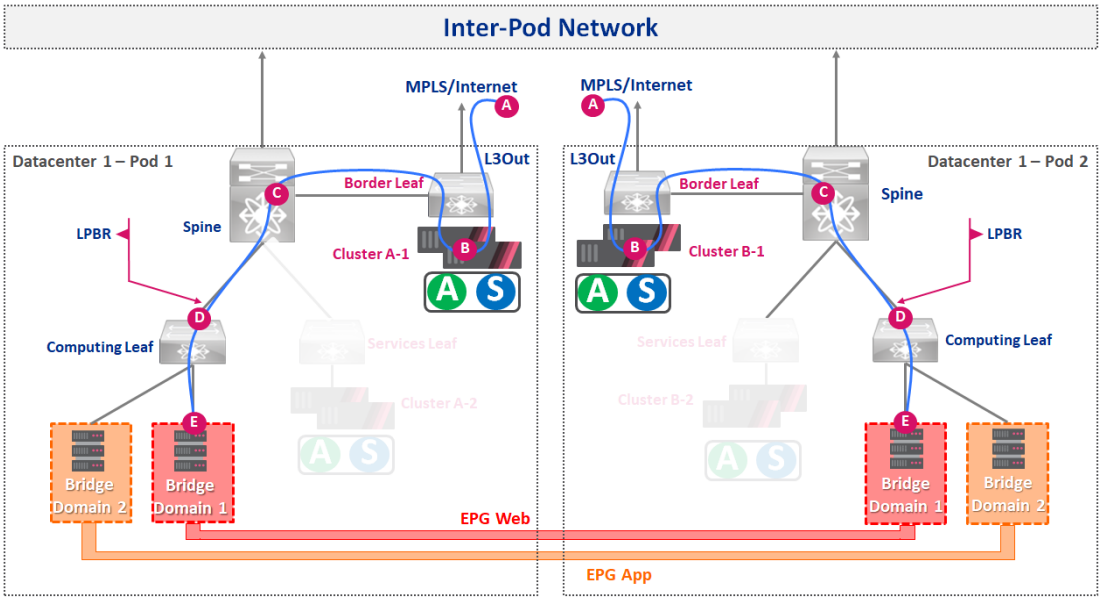

Figure 54: N-S Traffic per each Pod with location-based PBR

However, LPBR is not supported for the East-West traffic between the pods, because the return traffic from the destination pod will be sent back to the nearest (local to the Pod) firewall, which will drop the connection due to a lack of synchronization state.

Here is an example of such scenario:

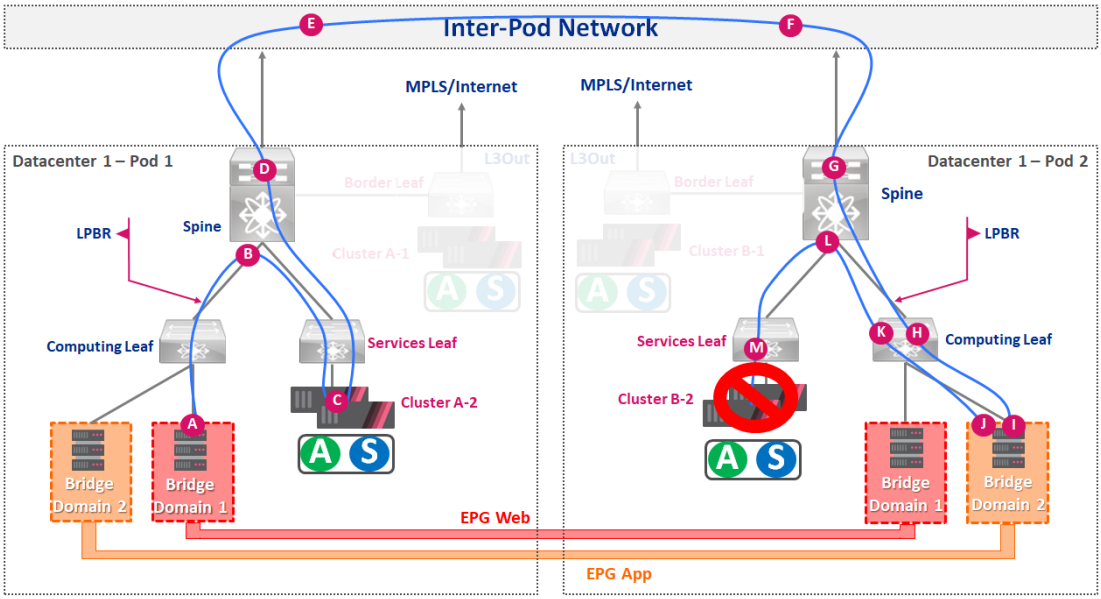

Figure 55: E-W Return connection is dropped by stateful inspection FW mechanism

Firewall deployment options for Multi-Pod stretched networks

How to address the problem described above? To support stretched bridge domains between Pods, Check Point recommends deploying Active-Active Firewalls with synced connection tables. This way, they will be able to process connections even in scenarios with asymmetric routing.

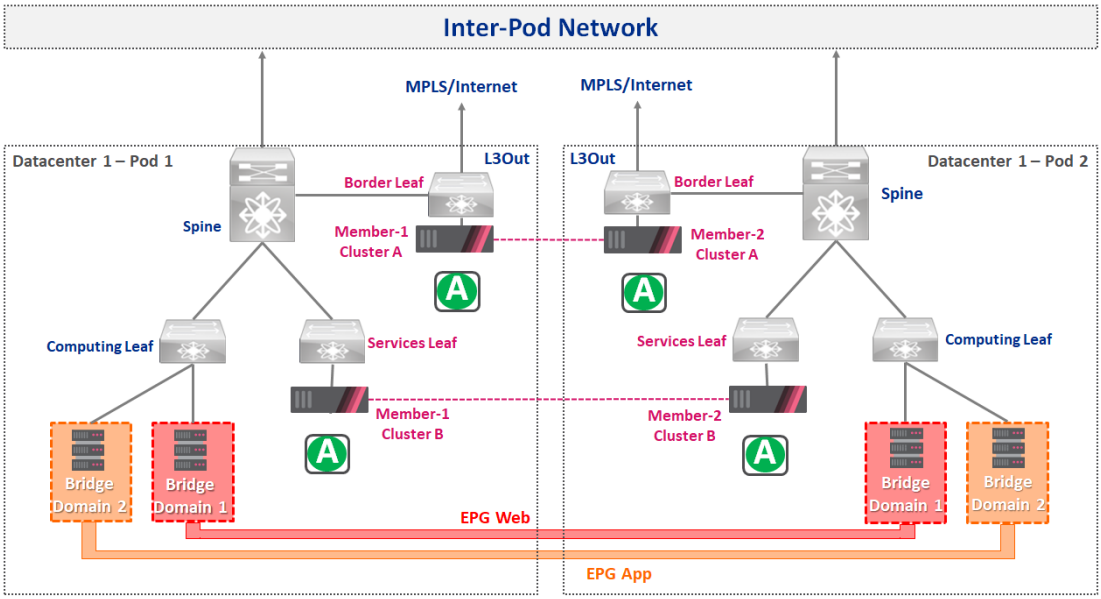

Figure 56: Firewall deployment options for Multi-Pod stretched networks

In the example below, traffic flowing from BD2 to BD1 passes through both firewalls located at POD1 and POD2

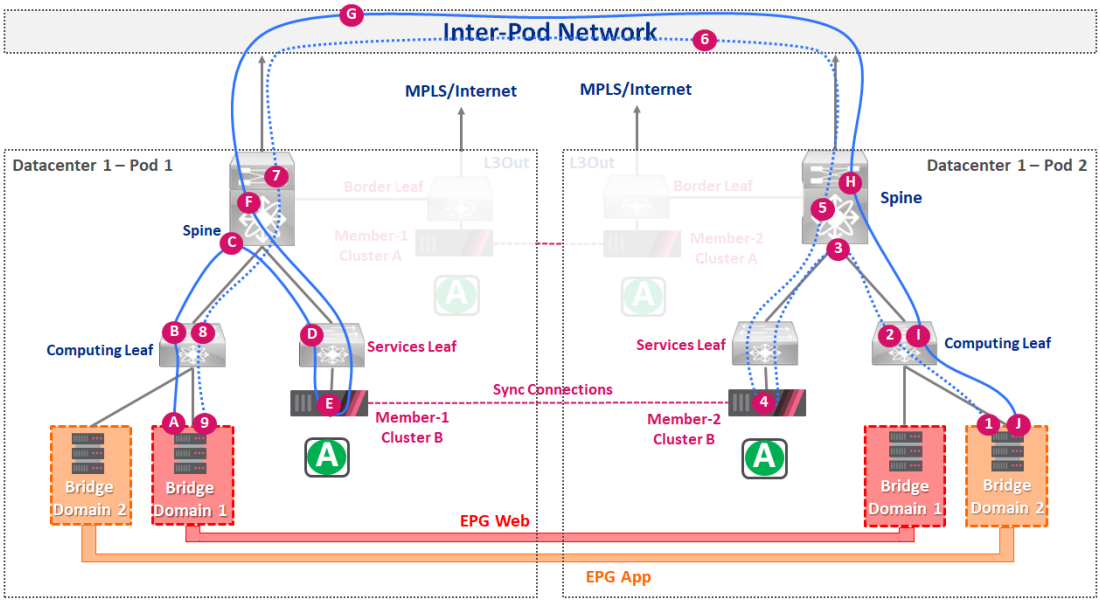

Figure 57: Firewall deployment option for Multi-Pod stretched networks E-W communications

There are two deployment options for firewalls to be synced and in Active-Active state:

1) Active-Active Firewall with different IP / MAC addresses using LPBR function to select the node
2) Active-Active Firewall with the same IP/MAC addresses using Cisco Anycast mechanism to select the node.

Option #1: Active-Active Firewalls with Location-based PBR

In a location-based PBR design, one of the Active-Active firewalls is deployed in each pod, with a sync link over an L2 stretched network between them. The firewalls would have different IP and MAC addresses. Firewall selection for traffic processing will be based on the traffic source – the nearest to the source firewall will be selected.

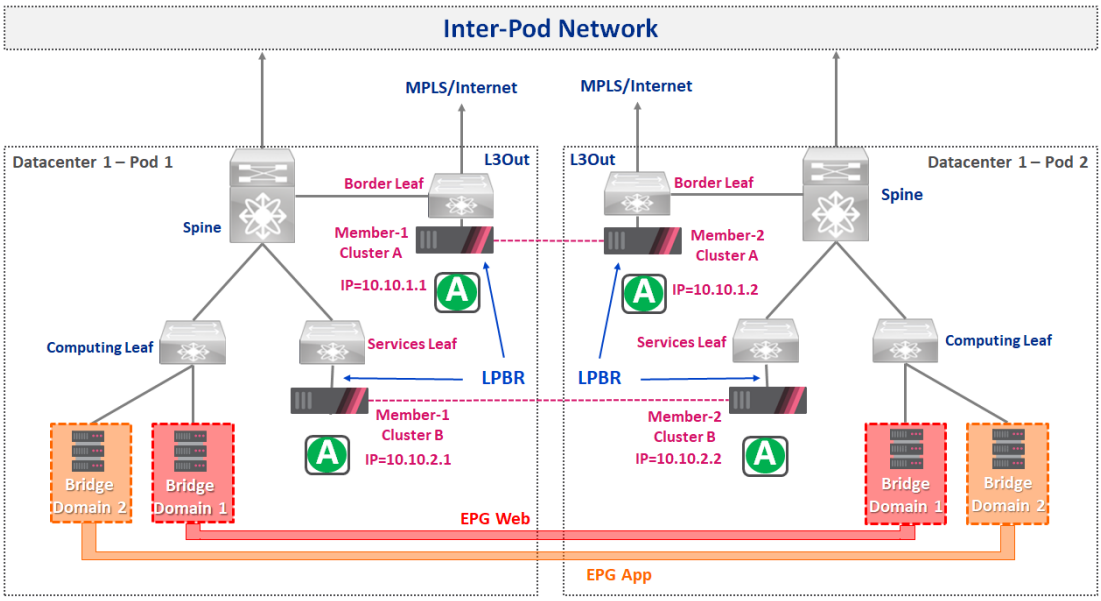

Figure 58: Active-Active Firewalls with Location-based PBR

Diagram demonstrating traffic flow in Active-Active Firewalls deployment with Location-based PBR:

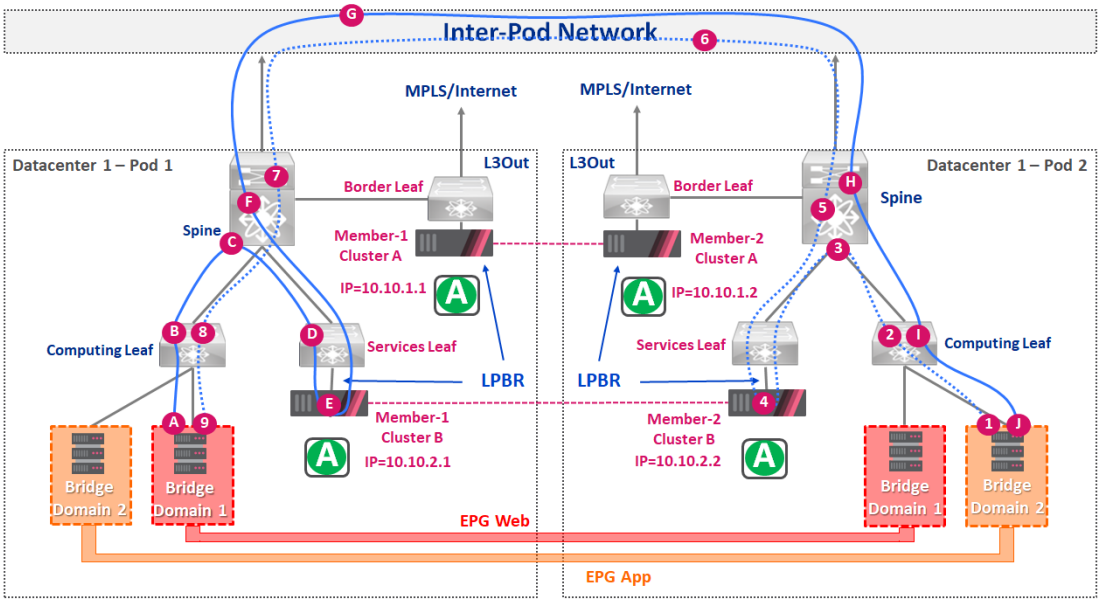

Figure 59: Active-Active Firewalls with Location-based PBR – Traffic flows

Essentially both firewalls will complement each other from the traffic processing perspective.

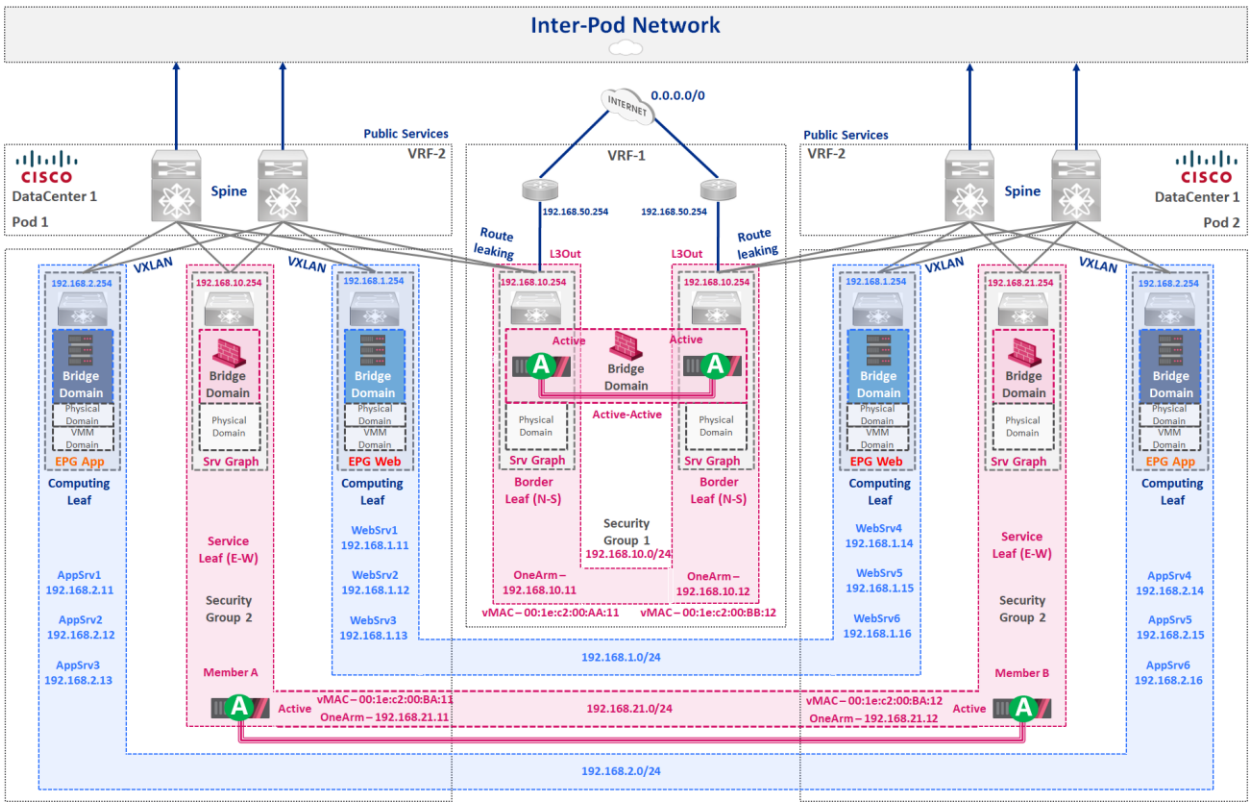

Figure 60: Active-Active Firewalls for Multi-Pod traffic with Local PBR - Topology

In this scenario, we have an Active-Active Firewall deployment with R80.40[20] or R81[21]; both cluster members have different IP addresses (also could be part of different subnets) and Sync interfaces between them. The traffic sync link should be configured outside of the Switch Fabric as a best practice. This scenario is applicable for Appliances only; Maestro Security Groups or VSX VSLS are not applicable.

> *Please note: LPBR should have IP SLA tracking configured to allow failover between the different Security gateways.*

Cisco ACI IP SLA[22] tracking allows to collect information about network performance in real time tracking an IP address using ICMP and TCP probes. Tracking configurations can influence the

[20] Active-Active Cluster XL R80.40 Admin Guide – URL: https://sc1.checkpoint.com/documents/R80.40/WebAdminGuides/EN/CP_R80.40_ClusterXL_AdminGuide/Topics-CXLG/Active-Active-Mode.htm

[21] Active-Active Cluster XL R81 Admin Guide – URL: https://sc1.checkpoint.com/documents/R81/WebAdminGuides/EN/CP_R81_ClusterXL_AdminGuide/Topics-CXLG/Active-Active-Mode.htm

[22] IP SLA Tracking, Cisco APIC Layer 3 Networking Configuration Guide, Release 5.0(x)- URL: Cisco APIC Layer 3 Networking Configuration Guide, Release 5.0(x) - IP SLAs [Support] - Cisco

PBRs, allowing to be removed when tracking results come in negative and returning the route to the table when the results become positive again.
ACI IP SLAs are available for Policy-based redirect (PBR) tracking:

- Automatically remove or add a next -hop
- Track the next-hop IP address using ICMP and TCP probes (in this case for the Check Point FWD ports like TCP port 256)
- Redirect traffic to the PBR node based on the reachability of the next-hop

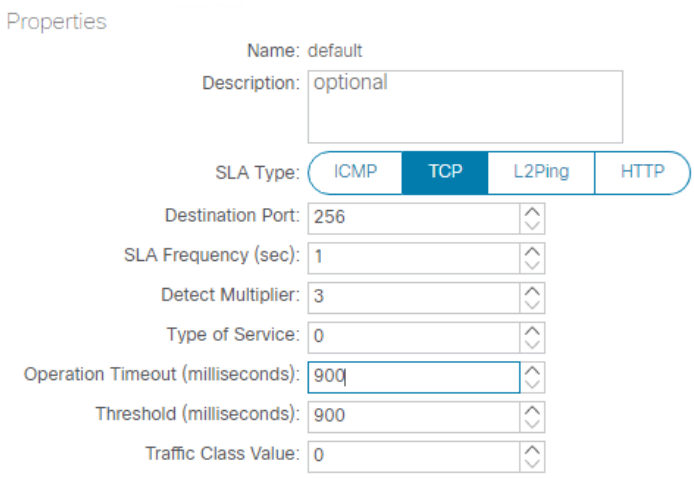

Figure 61: IP SLA Tracking – Configuration example

IP SLA Monitoring Policy
IP Service Level Agreements (SLAs) use active traffic monitoring to generate traffic in a continuous, reliable, and predictable manner, and analyze it to measure the network performance. Measurement statistics that are provided by the IP SLA monitoring policy operations can be used for troubleshooting, problem analysis, and designing network topologies.

With Cisco ACI, the IP SLA monitoring policy for Check Point Security Gateways is associated with:

- Service Redirect Policies: All the destinations under a service redirect policy are monitored based on the configurations and parameters that are set in the monitoring policy.

In the visual flow representation, we can overview how the Internet traffic can ingress to the routers in this shared L3Out, then according to the workload distribution (Pod-1 or Pod2).

North-South (external to fabric) traffic flow overview

Service Contract for N-S traffic flow:

A. Traffic is originated from the Internet and is routed to the L3Out (external connection of ACI) shared Border Leaf between the Pod 1 and Pod 2

B. Service Contract defines that External EPG should access the EPG WEB located in Pod 1 or Pod 2, allowing the traffic forwarding, important to note that the Service Graph redirects the traffic to the relevant Check Point Gateway using the LPBR located in Pod 1 or Pod 2.
C. Once traffic is redirected, processed, and inspected, traffic is forwarded to the Spine. According to the workload distribution, traffic will be forwarded to Pod 1 or Pod 2.
D. Spines "Knows" that EPG Web is located in one of the Computing Leafs in Pod 1 or Pod 2, where the traffic is forwarded.
E. Traffic is delivered to the final destination in the EPG Web.

EPG's has the Name, IP Address, and MAC Address information per each workload connected in the switch fabric. This information is critical because the traffic flow can be forwarded/redirected into the relevant transit components (Spines, Leafs, and Security Gateways). We have another use case applicable, similar principles, but a different way to do the configurations.

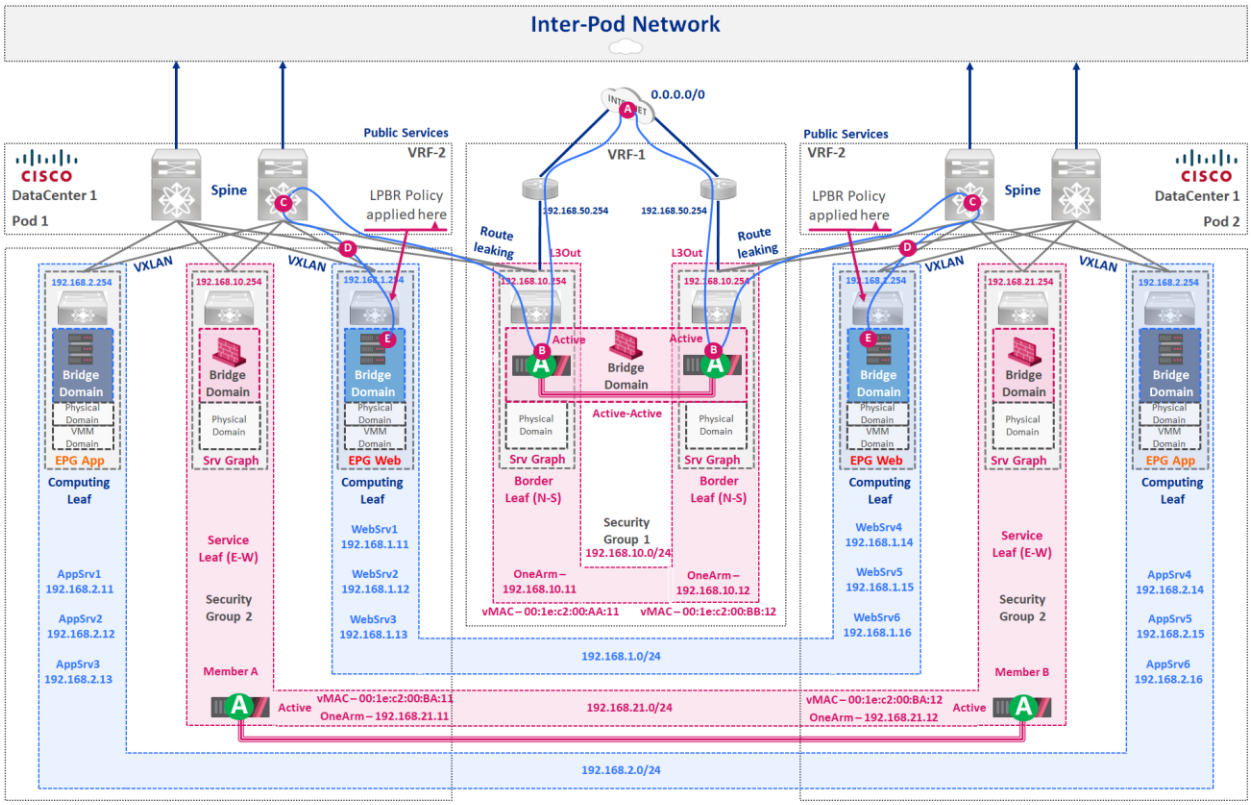

Figure 62: Visual representation of N-S traffic flows between stretched Bridge domains with Local PBR

East-West (internal to fabric) traffic flow overview

For the East-West scenario, this principle works similarly; we have the stretched bridge domains between the Web EPG Web and Apps EPG. This design can support an Active-Active data center scenario as the connections will be forwarded to Pod 1 or Pod 2 according to the workload distribution.

Service Contract for E-W traffic flow:

A. Traffic is originated from EPG Web located in the Pod-1 or Pod-2, according to the workload distribution in the Datacenter, the destination is the EPG App located in the Pod-1 or Pod-2; this is a typical Active-Active Datacenter example.
B. Service Contract allows the traffic between the Leafs, however in the Service Graph traffic must be redirected first in the Service Leaf located in the Pod 1 or Pod-2 (according to the proximity), traffic is synchronized between the cluster members.
C. Traffic in the Service Leaf is processed and inspected with the Check Point Security Gateway in One-Arm mode using LPBR; also, connections are synchronized with the other member in the Cluster.
D. Once the traffic is allowed according to the Security Policy, traffic is forwarded to the Spine located in Pod-1 or Pod-2. Traffic is forwarded to relevant Computing Leafs located in Pod-1 or Pod-2.
E. Spine (located in Pod-1 or Pod-2) knows the computing Leaf in the EPG App located in Pod-1 or Pod-2 and forwards the traffic.
F. Traffic is delivered to the final destination (G, H, I) in the EPG App located in Pod-1 or Pod-2.

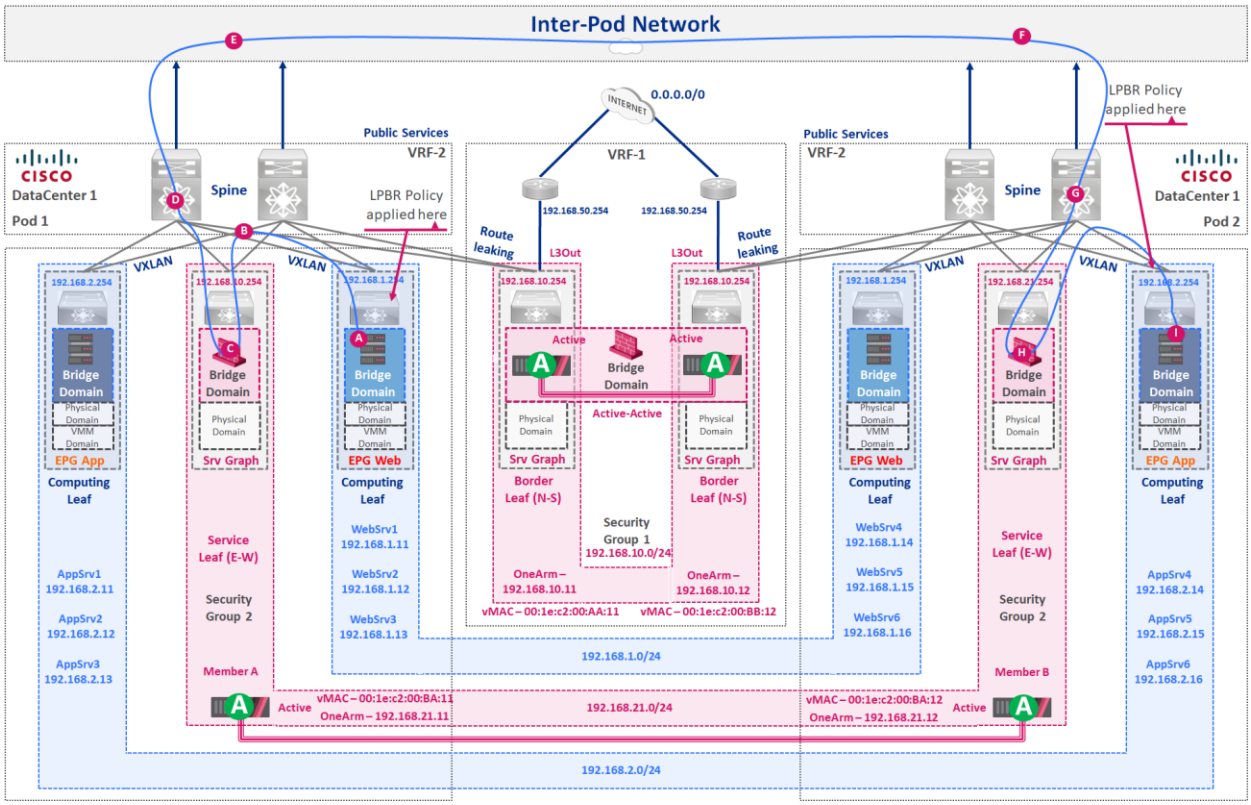

Figure 63: Active-Active Firewalls for Multi-Pod traffic inspection – East-West Traffic for stretched bridge domains

Maestro deployment with Location-based PBR scenario (*Not General Availability)

> *Please note:* *1) This solution is not part of the GA release and could be delivered through the RFE process on demand.*
>
> *2) In order to maintain a failover, make sure to configure the IP SLA mechanism as described in the previous section.*

In this scenario, Maestro MHO (orchestrators) are being deployed per pod with the synced interface. Each MHO has a Security Gateway attached with a required number of security gateways. The Security Group is configured in Active mode on both MHO / PODs.

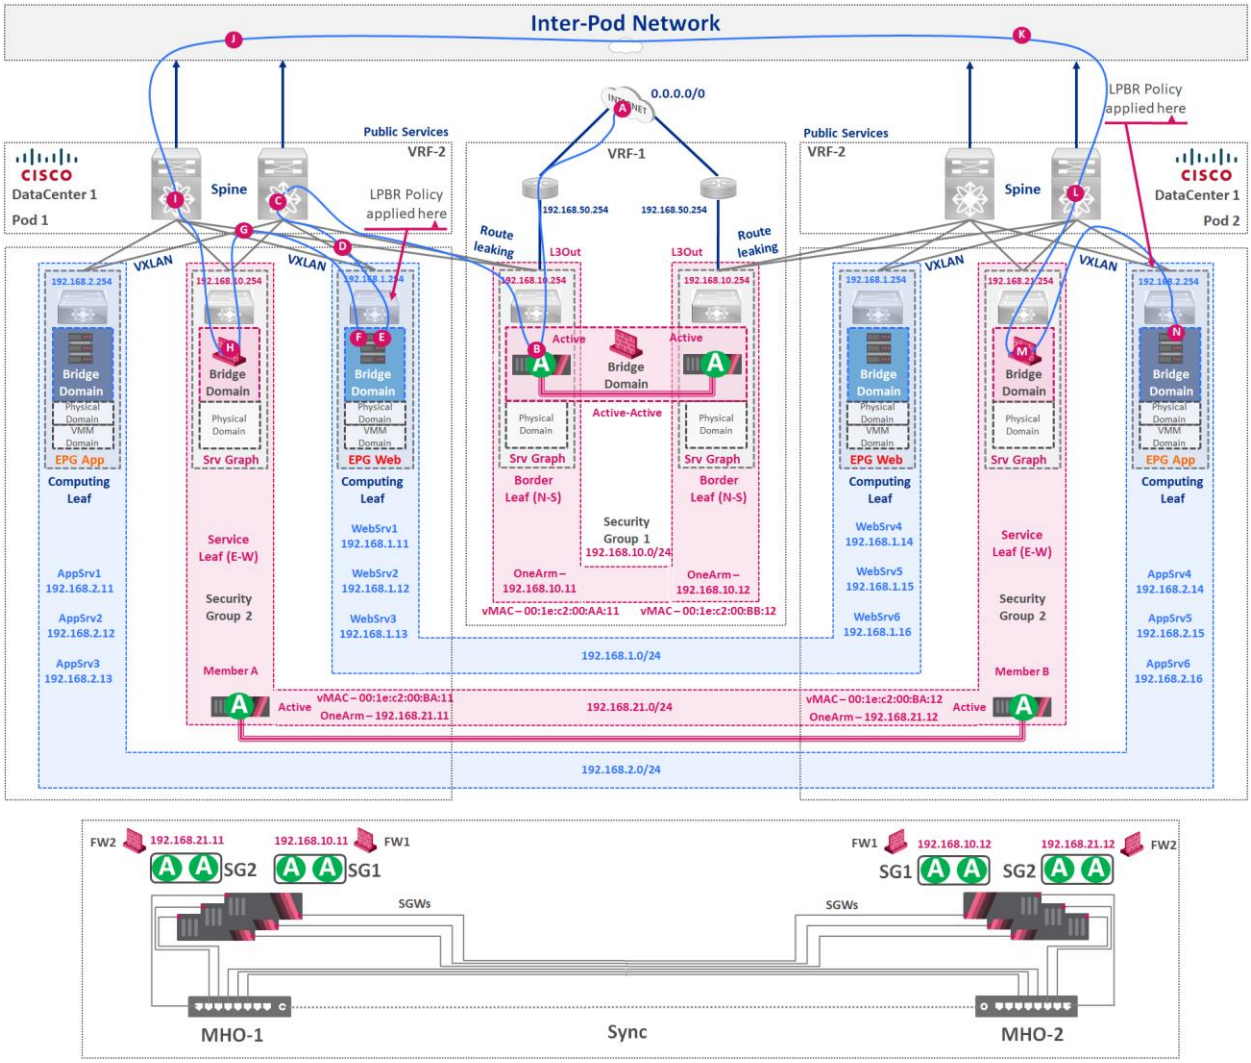

Figure 64: Maestro deployment in stretched domains for Multi-Pod scenario with LPBR

Option #2: Active-Active Firewall with Anycast support

In this design firewalls deployed per pod as an A/A cluster synced over L2 link. The firewalls should have the same IP and MAC addresses. Cisco ACI Anycast mechanism will select the closest FW to the source of the connection. This way returning traffic is not going to be dropped because connections are out of sync.

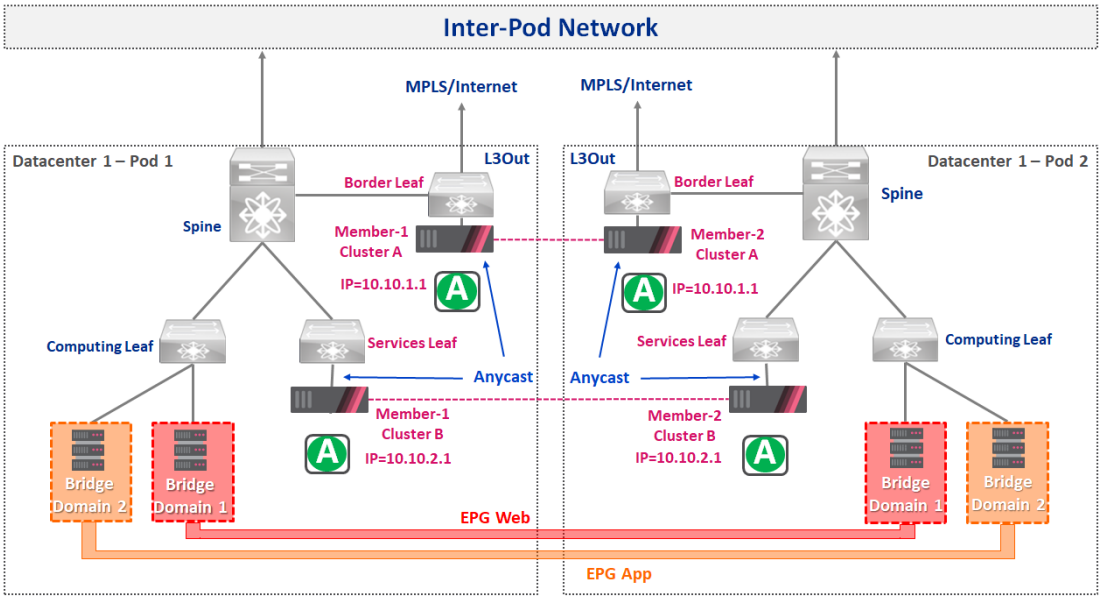

Figure 65: Active-Active Firewall with Anycast support

Diagrams demonstrating traffic flows via the second Pod traversing through different firewall modules for the initial and the returning flow:

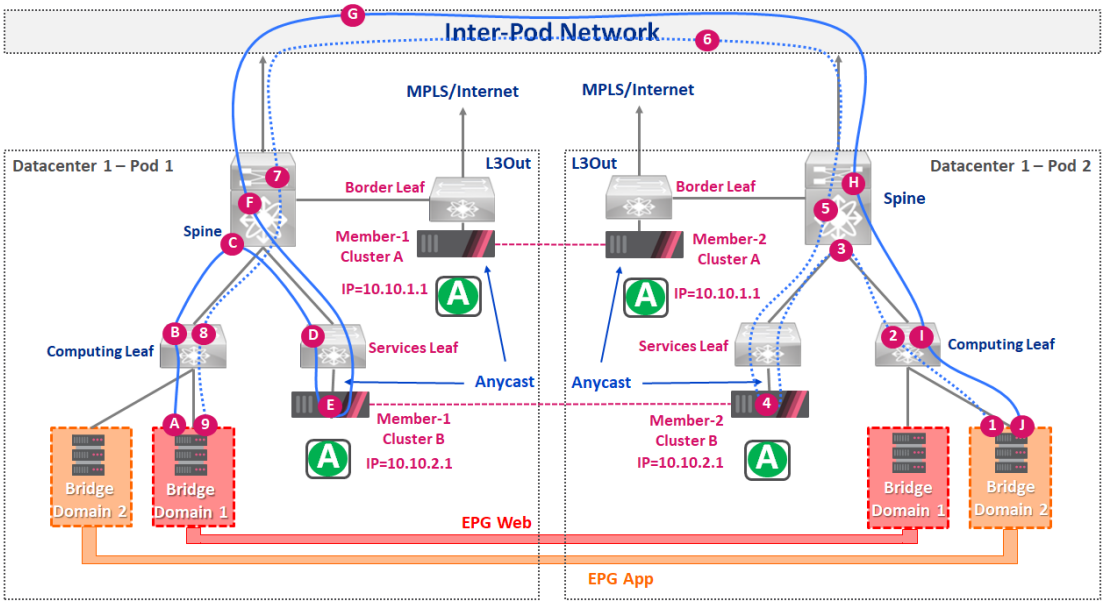

Figure 66: Active-Active Firewall with Anycast support for E-W traffic flow

Maestro deployment with Anycast scenario (*Not General Availability)

Please note: *1) This solution is not part of the GA release and could be delivered through the RFE process on demand.*

2) Cisco Anycast, at the moment, does not support IP SLA probing mechanism, therefore the failover will occur only when the FW physical port goes down.

In this scenario, Maestro MHO (orchestrators) are being deployed per pod with a synchronization link across the pods. Each MHO has a number of Security Gateways attached with a necessary number of Security Groups provisioned. Similar to the previous design pattern SC (firewalls) in each pod would be Active-Active and have the same IP and MAC addresses as on the diagram below:

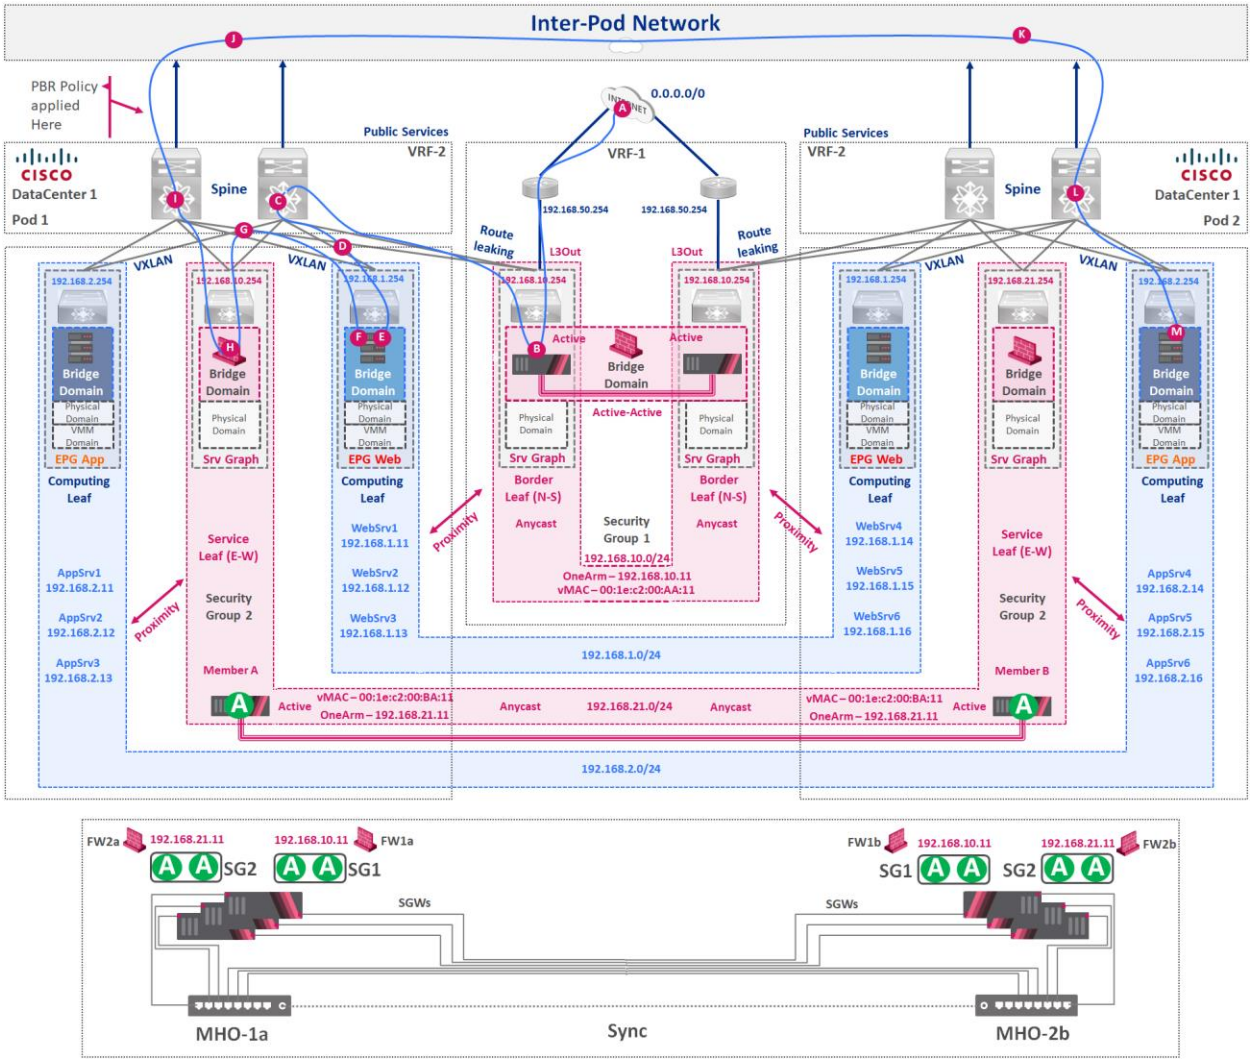

Figure 67: Maestro Active-Active for Multi-Pod traffic inspection with Anycast configuration

Anycast mechanism would select the closest FW to the source of the connection and the returning traffic is not going to be dropped because connection tables are synchronized.

Multi-Site Security Design

Multi-Site Cisco ACI deployment delivers the next level of scalability of Data centers deployments. This level can be considered to deliver multiple availability zones operating pretty much independently with very little or no interdependency in the configuration or used constructs.

Here is a quick comparison table to demonstrate most of the differences between multi-Pod and Multi-Site in order to choose the best deployment option:

	Multi-Pod	Multi-Site
Availability Zones	Single	Multiple
Redundancy	Redundant nodes, interfaces, and devices within a fabric	Full site deployment with end-to-end policy definition and enforcement
Configuration Change	APIC cluster pushes configuration changes into the entire Multi-Pod fabric while preserving tenant isolation	Multi-Site can selectively push configuration changes to specified sites enabling staging/validating while preserving tenant isolation
Interconnects	Typically uses lower latency IP network between pods	Multi-Site can deploy policies in fabrics across continents
L4-L7 Services	Services stitching across pods	Site local L4-L7 services stitching

The architecture allows you to interconnect separate Cisco ACI APIC cluster domains (fabrics), each representing a different region, all part of the same Cisco ACI Multi-Site domain. Doing so helps ensure multitenant Layer 2 and Layer 3 network connectivity across sites, and it also extends the policy domain end-to-end across the entire system. This design is achieved by using the following functional components:

- **Multi-Site Orchestrator** (MSO) - is the intensity policy manager.
  It provides single-pane management, monitoring the health-score state for all the interconnected sites. It also allows you to define, in a centralized place, all the intersite policies that can then be pushed to the different APIC domains.
- **Intersite control plane** - is used to exchange endpoint reachability information across sites. Essentially, MAC and IP address information is exchanged between the endpoints communicating across the sites in this method.
- **Intersite data plane** - all (Layer 2 or Layer 3) communications between endpoints connected to different sites is achieved by establishing site-to-site Virtual Extensible LAN (VXLAN) tunnels across a generic IP network that interconnects the various sites.

In summary, this architecture allows organizations to interconnect separate Cisco ACI APIC cluster domains (fabrics), each representing a different region or Data Centers, all as part of the same Cisco ACI Multi-Site domain. Doing so helps to ensure multitenant Layer 2 and Layer 3 network

connectivity across sites, and it also extends the policy domain end-to-end across the entire system.

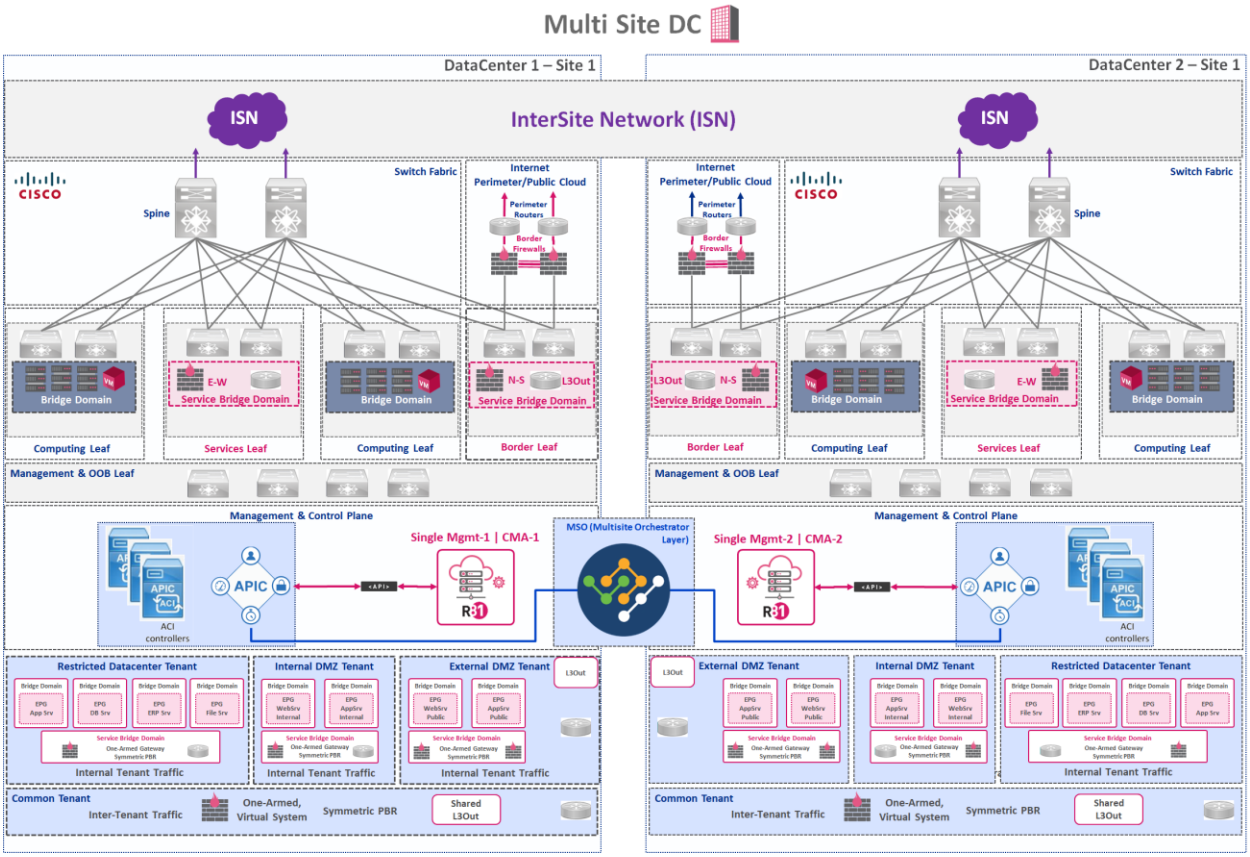

Figure 68: Multi-Site Architecture integrating the Infrastructure, Management, and Logical views

Check Point integration with Cisco ACI is based on API integration and provisioning between the Check Point management server and Cisco API Controller (APIC). So, in Multi-Site deployment, each APIC would need to be integrated in order to cover the whole ACI fabric network. Firewall gateway integration in each side would be primarily based on consideration related to the Pod topologies covered in the sections above.

Summary

Cisco ACI provides effective micro-segmentation for next generation datacenters through the integration of physical and virtual environments under a common policy model for networks, servers, storage and security. Cisco ACI's application-aware policy model and native security capabilities are leveraged by Check Point CloudGuard to dynamically insert, deploy and orchestrate advanced security protections within software-defined datacenters.

Together, Cisco and Check Point provide a powerful solution that gives customers complete traffic visibility and reporting in addition to proactive protection from even the most advanced threats within virtual network environments. The joint solution forms the foundation of an intelligent application delivery network architecture where security seamlessly follows application workloads and accelerates application deployment while maintaining reliability, multi-tenancy and operational workflows.

Worldwide Headquarters
5 Ha'Solelim Street, Tel Aviv 67897, Israel | Tel: 972-3-753-4555 | Fax: 972-3-624-1100 | Email: info@checkpoint.com

U.S. Headquarters
959 Skyway Road, Suite 300, San Carlos, CA 94070 | Tel: 800-429-4391; 650-628-2000 | Fax: 650-654-4233

www.checkpoint.com