Macro-Segmentation Service™ (MSS™) Design & Deployment Guide – L3 Firewall

Version 1.0

May 2020

Table of contents

Table of contents

Table of contents

Introduction

Modern day applications deployed in the data centers have become multi-tiered and distributed. This has resulted in an increase in the amount of east-west traffic seen in the data centers. This includes traffic from physical-to-physical (P-to-P), virtual-to-virtual (V-to-V), and between physical and virtual (P-to-V) workload.

Legacy security architecture consists of perimeter firewalls, that primarily enforce security for north-south traffic. As the applications within the datacenter become increasingly distributed, it results in a rise in the number of workloads (both physical and virtual) deployed in the datacenter which in turn present many new possible entry points for security breaches. Once an attacker gets in, he can easily move laterally within the datacenter and go unnoticed, as the perimeter firewall is only inspecting north-south traffic.

Micro-segmentation solves part of the problem by providing security for east-west flows between virtual workloads (V-to-V). Arista Macro-Segmentation Service (MSS) addresses the remaining gap in security deployment models by securing the east-west traffic between physical-to-physical (P-to-P) and physical-to-virtual (P-to-V) workloads.

Arista MSS enhances the paradigm of software defined service insertion driving operational efficiencies in IT and security operations meeting the key requirements for modern cyber security, compliance, threat detection, etc.

This design and deployment guide outlines Arista Networks integration with best of breed firewalls for today's modern data center network security models. Arista Macro-Segmentation Service components include a leaf-spine switch fabric, Arista CloudVision and a firewall attached to service leaf switches. The goal of such a design is to allow for consistency in application deployment, scale, manageability and easier scalability of the network and service layers.

Solution Overview

Arista Macro-Segmentation Service (MSS) provides a software-driven dynamic and scalable network service to logically insert security devices into the path of traffic with complete flexibility on placement of security devices and workloads. It is specifically aimed at physical-to-physical (P-to-P) and physical-to-virtual (P-to-V) workloads.

What makes MSS unique is that it places the control of policy enforcement directly in the hands of security administrators. This is accomplished using standards based forwarding with no proprietary frame formats and without placing limitations on where the service devices must exist within the network.

Arista CloudVision provides a single point of integration to the network and ties in automation and orchestration capabilities with a network wide view by aggregating the entire network state. Arista MSS is a service in CloudVision that provides the point of integration between individual firewalls or security managers and the Arista network fabric. When enabled, MSS can steer interesting traffic to the firewall for further inspection or program the leaf switches to either bypass the firewall or drop the traffic based on the policy configuration on the firewall, thus reducing the load on the firewall and the network fabric.

MSS Deployment Mode

Many data centers have firewalls deployed in layer-3 mode, acting as first hop for the hosts, serving applications. The Layer 3 Firewall is connected to the network and configured to enforce policy between different security zones or endpoints.

Instead of using routing policy to attract traffic to the firewall, the Macro-Segmentation Service redirects traffic to the firewall, dynamically inserting it into the path for traffic between relevant endpoints. **L3 Firewall is not configured as a gateway for the subnet.**

This document will focus on the design and deployment of MSS with the Routed (L3) firewall deployment mode.

Use Cases

The use cases below discuss widely encountered security challenges in the modern datacenter.

- Securing East-West Traffic
- Rapid Threat Assessment and Containment
- Automated Network Segmentation
- Dynamic policy migration
- Securing Management Interfaces

Securing East-West Traffic

Prior to Micro-Segmentation or Arista's Macro-Segmentation, security was primarily enforced at the perimeter for mostly the north-south traffic flows. There was limited or no security for the east-west flows, especially for the physical devices. Security policies also depended on the network topology thus requiring tight co-ordination between the network and the security teams.

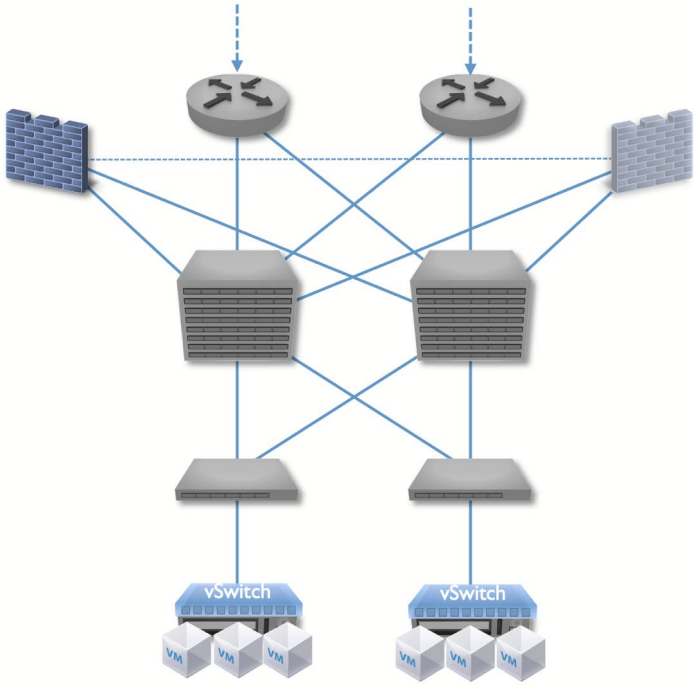

Legacy Approach

Figure 1. Traditional approach to securing traffic in data center

Traditional approaches also present management and scalability challenges as they involve manually configuring multiple TOR switches with Access Control Lists (ACLs) to steer traffic to the firewalls. This not only increases the number of touch points for configuration, increasing the probability of configuration errors, but also represents scalability challenges. There are also performance implications where sending all traffic to the firewall for inspection can unnecessarily overwhelm both the network and firewall.

One of the advantages of using Arista MSS is to remove the restriction of firewall placement, thus solving the architectural restriction of the legacy approach. As shown in figure 2, firewalls can now be attached to a leaf/service switch in the network fabric and still protect hosts irrespective of their physical location.

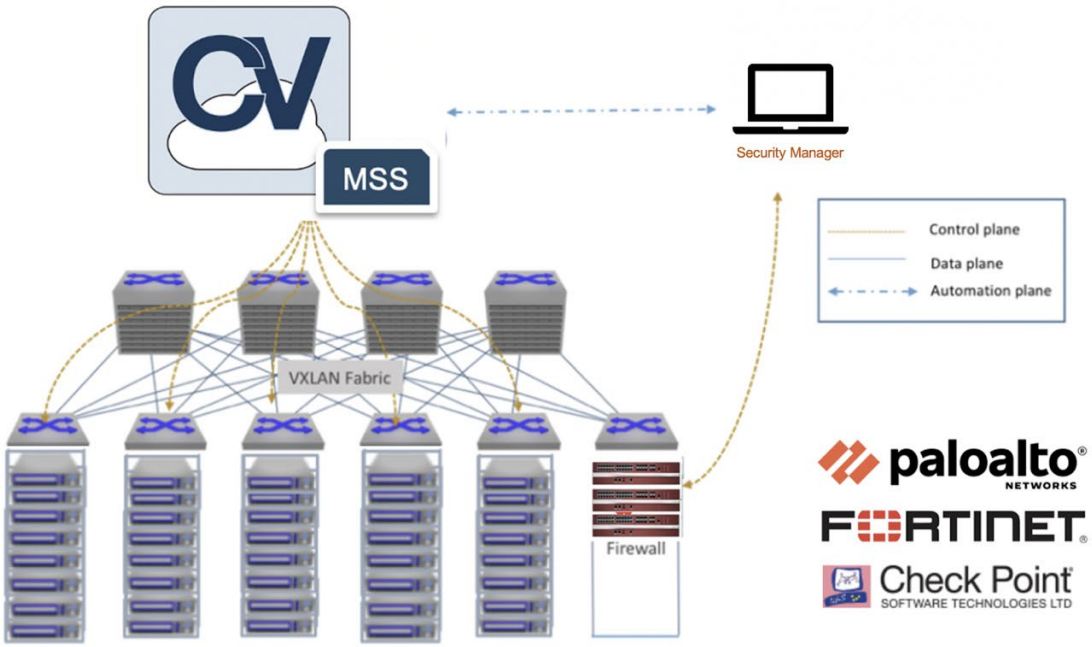

Figure 2: Modern approach to securing east-west traffic with Arista MSS

MSS reduces the number of configuration touch points by making the firewall/security manager as the single point of policy control. The policy is defined by the security administrator on the firewall/security manager. MSS reads the policy via API integration, identifies the TORs where the traffic needs to be intercepted, and seamlessly programs redirect rules to steer traffic to the firewall or offload rules to either bypass the firewall or drop the packets at the TOR, without any involvement of network administrators.

Note: Depending on the firewall vendor, MSS will be able to read the policies from the security manager or the firewall or from both.

Rapid Threat Assessment and Containment

Enterprises are faced with implementing an agile security architecture to protect critical assets from ever evolving sophisticated threats. When such threats in the form of exploits, ransomware, compromised systems etc. are detected, the SecOps is challenged to react fast to protect critical assets and infrastructure. Security vulnerability fixes or updates to mitigate malware may not be available to mitigate the threat in a timely manner.

The security administrator can leverage the power of Arista MSS to separate infected hosts from the rest of the network, block specific known communication patterns directly related to malware or redirect traffic from these hosts to the firewall by simply changing the policy on the firewall to quarantine the compromised host. The security administrator or network administrator does not need to know the physical location of the compromised host within the campus or datacenter network. CloudVision has this information in it's database and can implement the quarantine policies where they are needed.

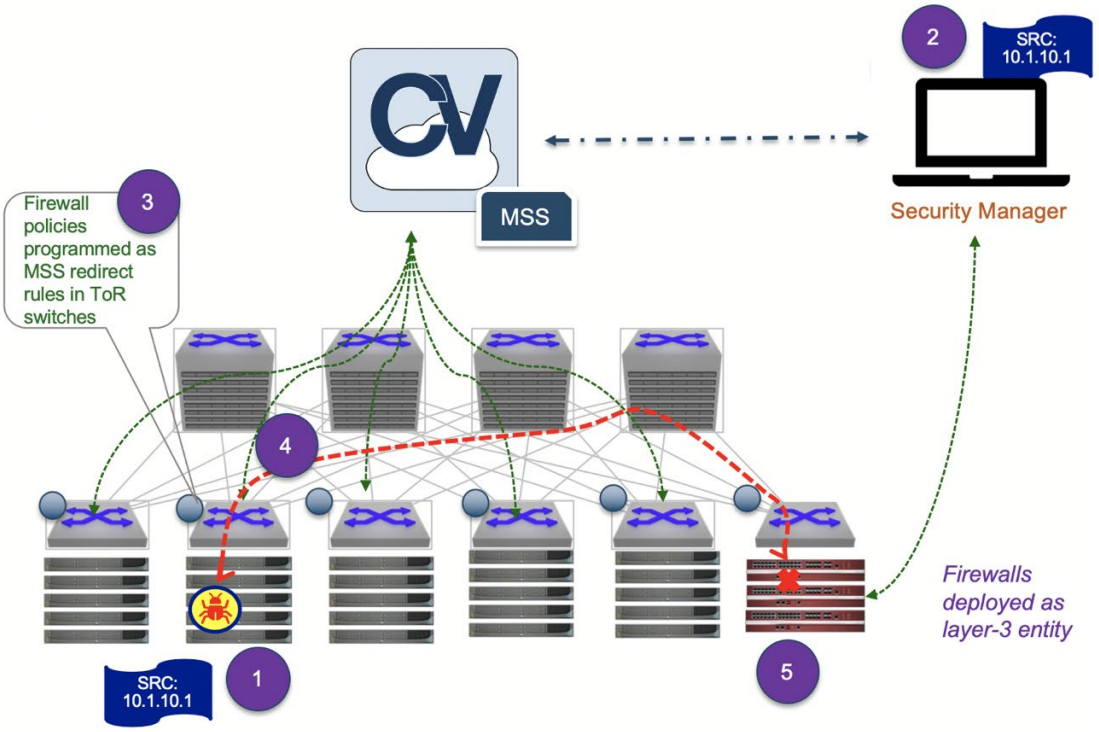

Figure 3. Rapid threat assessment and containment using MSS

As seen in the figure 3 above, the workflow for rapid containment, with MSS, is as follows

1. Suspect/ Malicious traffic detected
2. SecOps updates quarantine rules in the firewall
3. Firewall rules received by CloudVision MSS, programs intercept rules on the TOR switches
4. TOR switches intercept traffic and redirect to firewall
5. Firewall drops malware traffic

Automated Network Segmentation

When designing modern Datacenter or Campus networks, segmentation/tenancy is often an important security requirement to split the network into different security zones and apply security services between and within them. The toolkit to accomplish these tasks, from the network side, are virtual routing and forwarding instances (VRFs), VLANs and often vendor specific proprietary mechanisms. Implementing these tools in the network drastically increases its complexity and complicates operations and maintenance of the IT ecosystem.

A new automated approach would be to use firewall as the only configuration entity for segmentation and allow MSS to provide necessary network isolation for a multi-tenant environment. A new tenant can be placed in an isolated firewall zone with corresponding set of policies. The most common objective would be to restrict inter-tenant traffic. When MSS is used, no manual VRF configuration is necessary to achieve this objective. In addition, MSS can provide more granular segmentation on top of this bare minimum. One added benefit of this approach is that a tenant's workload can be placed anywhere in the fabric.

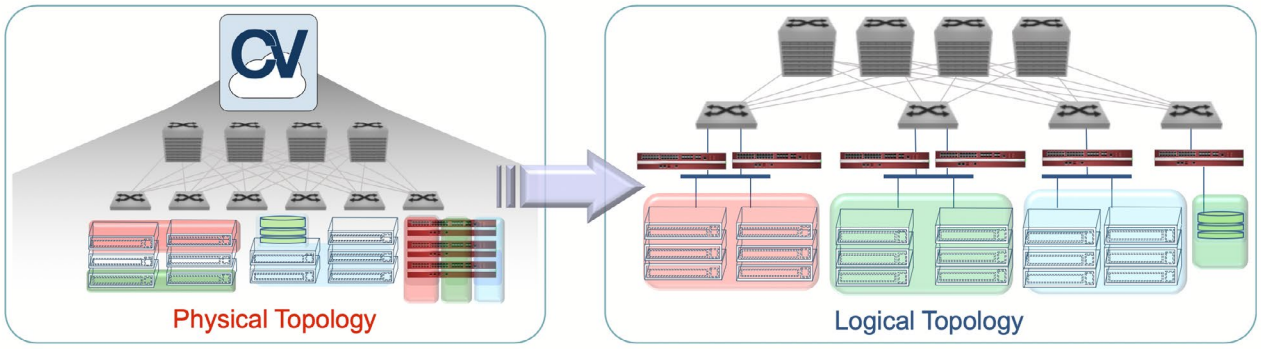

Figure 4. Automated network segmentation using MSS

Dynamic Policy Migration

One of the key issues facing modern data centers is that the security policies are often dependent on network topology. With virtualization, workload mobility allows for independence of placement of workloads, based on compute power requirements. In case of a workload mobility event, network policies will need to be reconfigured on the TOR connected to the host, where the endpoint moved to, to redirect traffic to the firewall. This makes the network and the security administrators co-dependent and prone to delays.

With MSS, security admins own security policies. No need for network admins to get involved. As the endpoint moves to another rack or leaf switch, the same security rules will be enforced. In the example in figure 5, Endpoint A is sending traffic to Endpoint B which is getting redirected to the firewall due to the redirect rule programmed on its TOR switch.

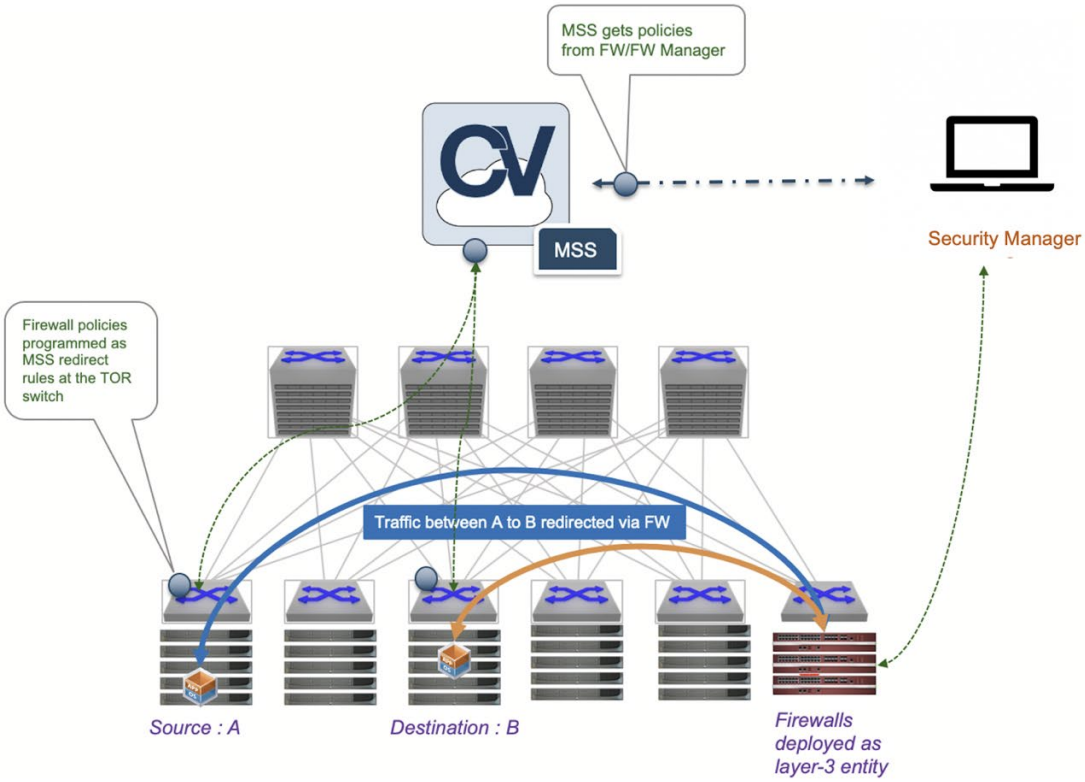

Figure 5.Traffic flow between endpoint A and B before migration

Now Endpoint A moves to a compute node connected to a different TOR switch. Cloudvision MSS, in real-time, updates the security rules for offload and redirect policies to the new ToR, connected to the host.

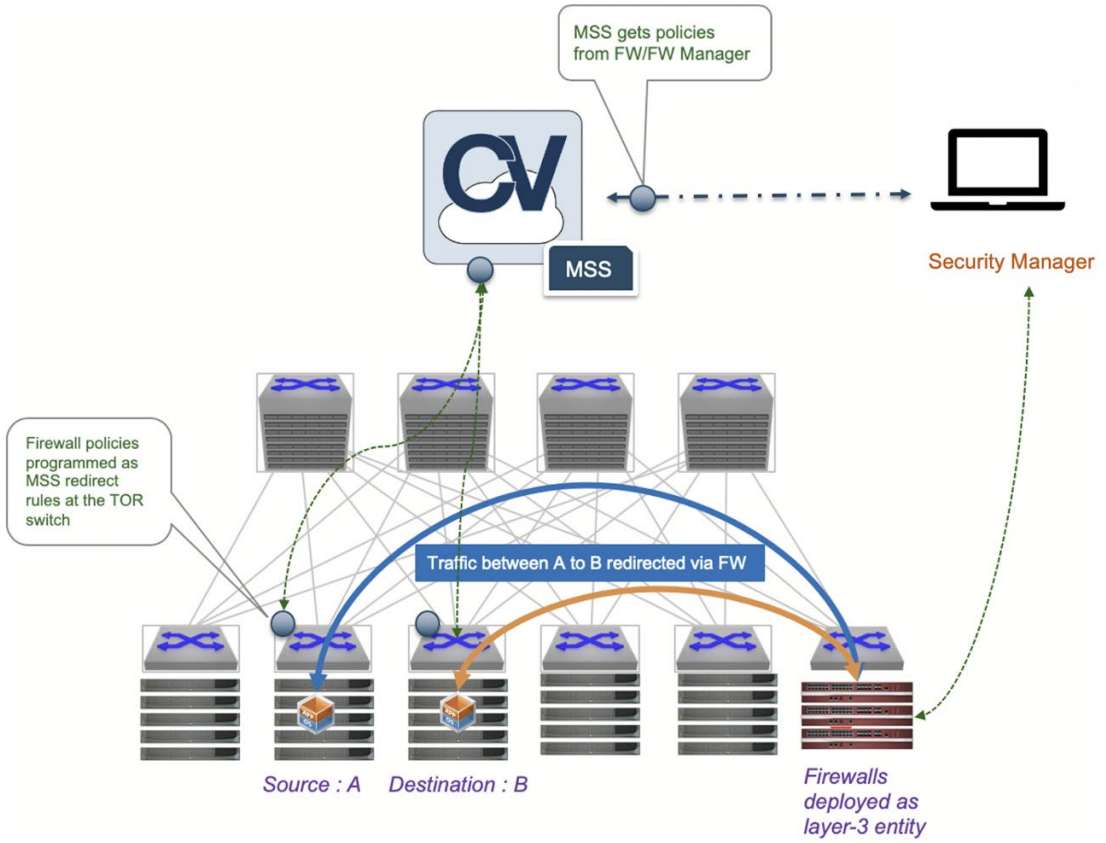

Figure 6.Traffic flow between endpoint A and B after migration

Monitoring and Securing Management Traffic

The modern data center also has an out-of-band network that caters to managing the application, storage, virtualization, network, analytics and other layers. With virtualization, the hypervisor management also needs to be secured. Should an attacker gain access to a hypervisor management interface, they could either hop to another device on the network or compromise the local virtual machines or access critical assets in databases, that is often shared across host/applications.

Arista's MSS can be used to protect management interfaces. Only explicitly allowed hosts, defined by the security administrator in firewall policy, would be allowed access, such as a jump host or administrator end user computing instances etc.

Key Benefits

Arista's Macro-Segmentation Service (MSS) for L3 Firewalls offers the following key benefits:

- Insert security between any physical and virtual workloads in data center controlled by the Security Admin
- Automatic and seamlessly orchestrated service insertion - eliminating manual steering of traffic, per workload or tenant
- Reducing the load on firewalls by blocking or allowing traffic using offloads rules directly at the TOR switch
- Security policies follows the host and application throughout the network without any manual intervention
- No proprietary frame formats, tagging, or encapsulation for implementing service insertion
- Reduce number of configuration touchpoints by providing a single point of control – e.g. the security policy manager for physical firewalls
- Faster reaction to security threats within the datacenter or campus network by separating infected hosts or denying rogue traffic pattern throughout the whole network

Security Workflow with MSS

The section below walks through MSS from a high level.

Step 1: CloudVision as single point of control

Arista's Macro-Segmentation service is enabled in CloudVision. Arista switches are configured to stream their state in real-time to CloudVision.

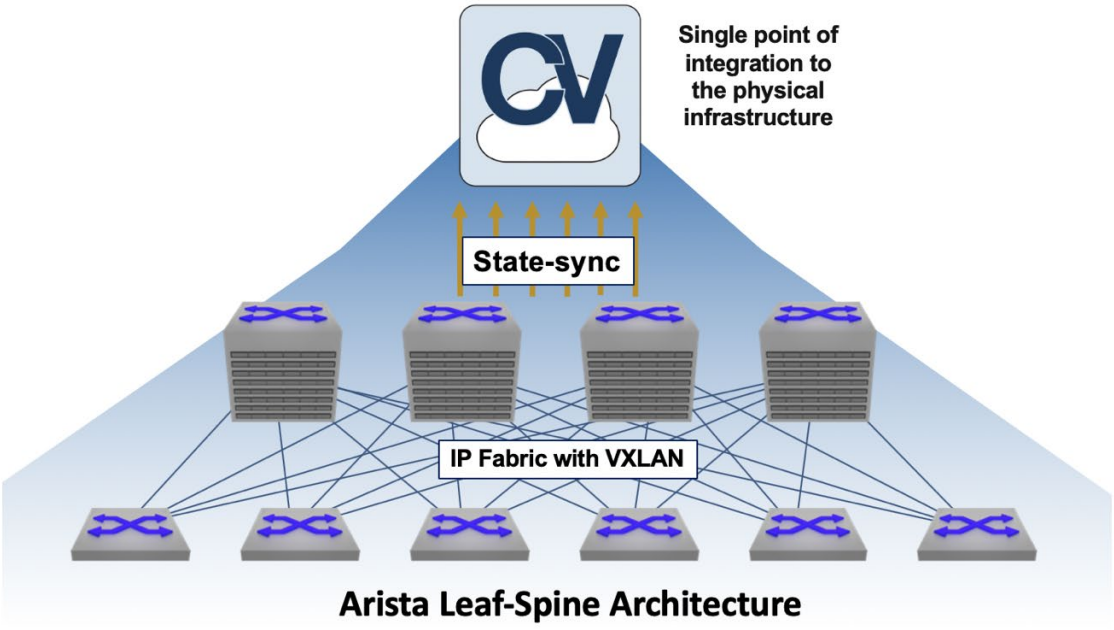

Figure 7: Arista CloudVision providing a single point of control

This allows CloudVision to build a database of hosts, network switches and service appliances such as firewalls attached to the network, identifying physical ports, IP addresses, etc.

CloudVision is also configured to communicate and sync policies from the firewall/security manager. While CloudVision has the state, it's not in the data path and just instruments the path.

Step 2: Firewall rules are implemented by the security team

Security policies are created by the security team with user-defined tags. These tags are user defined strings that are used as an identifier or comments in the FW/Security manager console and associated with the policies. These tags can be either for redirecting traffic to the firewall service for L4-L7 inspection or for offloading the rule to the Arista leaf switch. When using offload tags, the action taken on the leaf switch is controlled by the firewall action defined in the security policy. When the security administrator defines a DROP action in the rule the traffic is dropped at the Leaf Switch where the host/subnet is located. Using a PERMIT action on the firewall policy, the traffic flows matching the rule are forwarded directly by the leaf switch without redirecting to the firewall.

CloudVision will send a request to the firewall or security manager to provide the details of the security policies and will determine where traffic needs to be intercepted. CVX retrieves/polls security policies from the firewall at a configurable interval in order to ensure the two entities are in sync in near real-time.

Step 3: CloudVision applies an intercept to steer interesting traffic

Once a firewall policy has been created by the Security Administrator with the configured tags that affects a host or subnet or service that CloudVision is aware of, CloudVision then pushes redirection or offload rules to the physical switches.

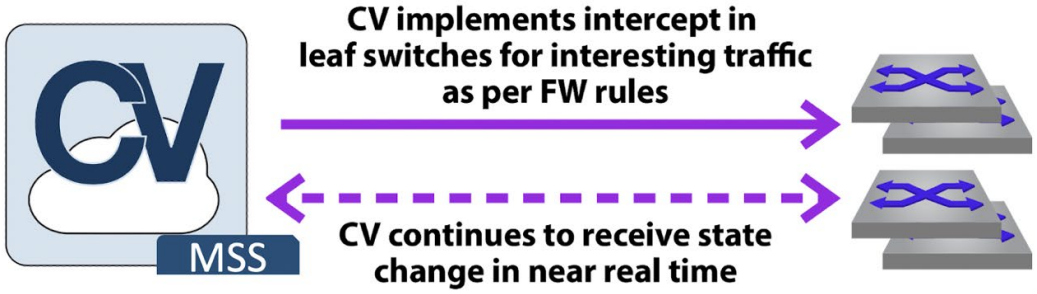

Figure 8. CloudVision to physical switch communication

Step 4: Data plane traffic steering with Macro Segmentation Service

Redirect and offload rules are configured as DirectFlow[1] rules, which are optimized to save TCAM space as compared to traditional ACLs thus achieving higher scale.

Redirect Rules

Intra-subnet and Inter-subnet traffic patterns, that need stateful L4-L7 inspection or Next Generation Firewall (NGFW) features, are tagged in the policies with a user-defined redirect tag, in the firewall. All traffic from the hosts identified in the firewall policy will be redirected bidirectionally to the firewall service for inspection.

[1] *For more information on DirectFlow, please visit this link: https://www.arista.com/en/um-eos/eos-directflow*

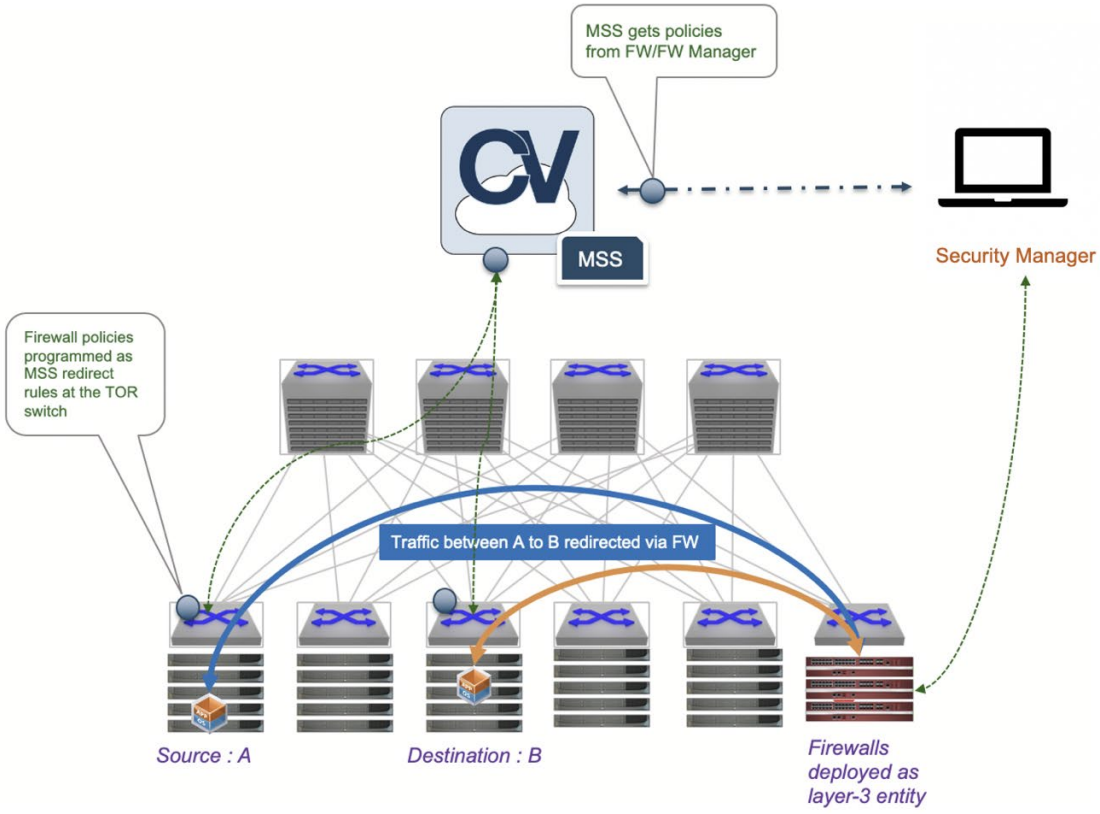

Figure 9. MSS redirect rules steering traffic to firewall

If the intent is to redirect a specific flow of traffic, and not all traffic from hosts identified in the policy, add a second tag 'verbatim' to the policy on the FW. This expresses the intent that only traffic matching the exact policy should be redirected to the FW.

Offload Rules

Firewall rules marked with a user-defined offload tag are either dropped directly on the switch where the host is connected or forwarded without firewall inspection.

For example, storage backup traffic may need no inspection, so it is forwarded directly, thus saving additional burden on the network fabric as well as the firewall. Firewall admins need to set action as allow and tag the policy with the offload tag.

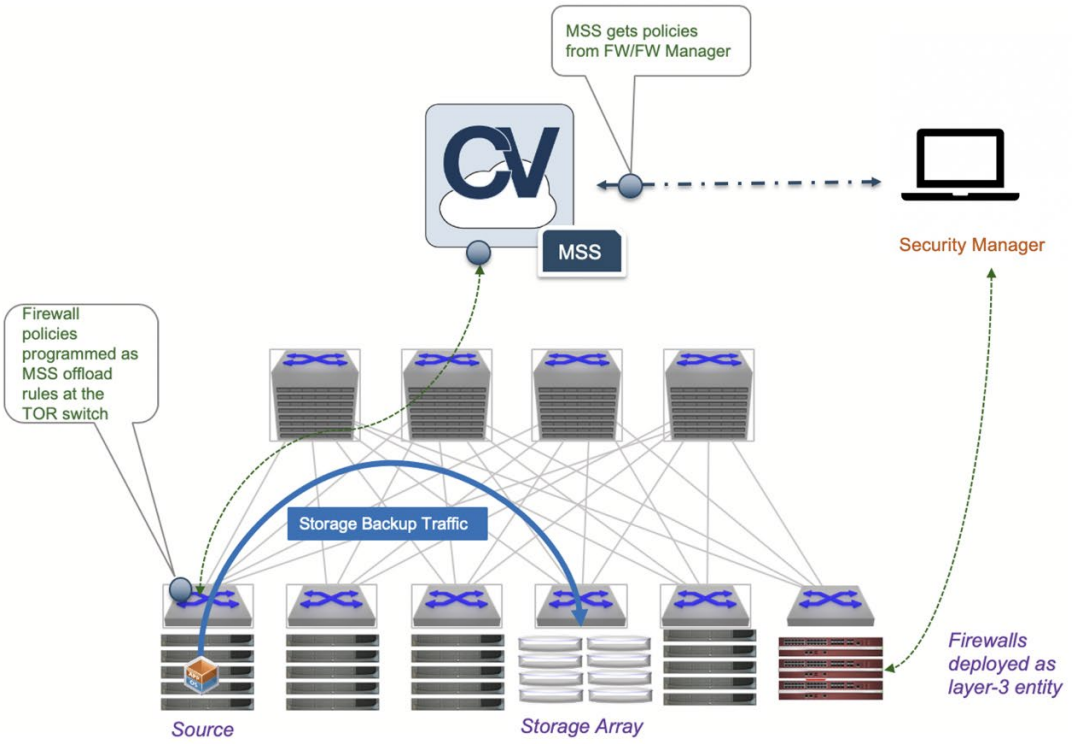

Figure 10. MSS offload rules enforced at the TOR to bypass traffic to firewall

Similarly, connections to un-secure ports or connections from outside to backend ports should be dropped close to the source. Therefore, the security administrator configures the policy with a drop action and tags the policy with an offload tag. This not only saves firewall bandwidth but also reduces network fabric bandwidth.

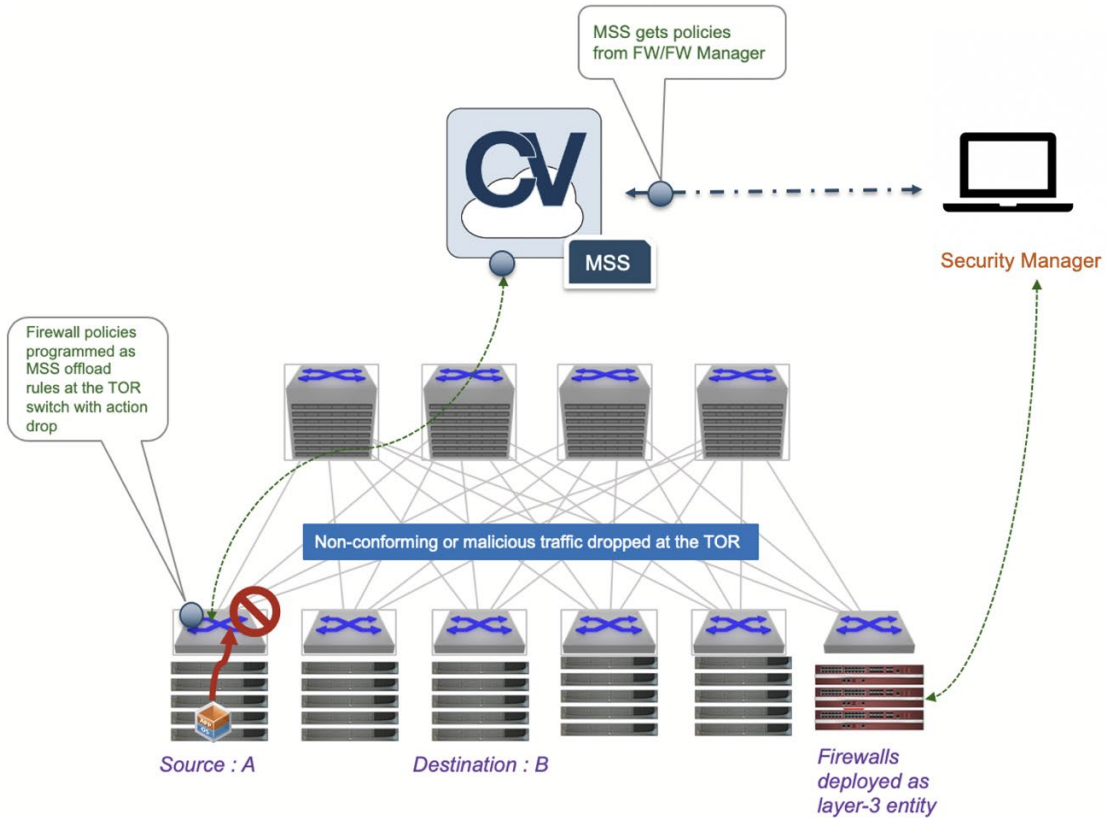

Figure 11. MSS offload rules enforced at the TOR to drop traffic

With the offload tag, the enforcement on the switches is stateless and bidirectional. In other words, if a policy that permits SSH traffic from Host 1(H1) to Host 2(H2) is marked with an offload tag on the firewall, then SSH traffic from H2 to H1 will also be permitted.

For security reasons an implicit redirect rule is always added to hosts matching a policy marked with offload tags on the switches. In case of a resource overflow or issues with programming a rule on switches or the absence of information on how to process traffic from the protected host that doesn't match this specific policy, as a last resort, the implicit redirect rule enforces the security policy by steering traffic originated or destined to the set of hosts, mentioned in the offload policy, to the firewall service for further actions.

If we add the 'verbatim' modifier tag to the policy, then the implicit redirect rule is no longer added to the hosts in policy.

Policies that are marked with Offload tag should be marked as higher priority as compared to the policies marked with the redirect tags (This is not applicable when using the 'verbatim' modifier).

For more examples on how the redirect and offload tags (with or without the verbatim modifier) affect the traffic flow, please refer to Appendix 5.

Reference Design

This reference design uses the following terminology:

- **Intercept Switch/VTEP**: Top of Rack switch and VXLAN tunnel endpoint connected to host from which traffic is intercepted. In this design, intercept-1 and intercept-2
- **Service Switch/VTEP**: Top of Rack switch and VXLAN tunnel endpoint connected to firewall. In this design, service-1 and service-2
- **Service Interface**: The interface at the Top of Rack switch that receives/sends packets to the firewall.
- **VXLAN**: Virtual eXtensible LAN - a standards based method of encapsulating Layer 2 traffic across a Layer 3 fabric
- **CVX**: Arista CloudVision eXchange (CVX) is a part of CloudVision and is a virtualized instance of the same Extensible Operating System (EOS) that runs on physical switches. It functions as a point of integration between firewall/security manager and the Arista network in order to steer interesting traffic to the firewall.

Requirements For Arista MSS Integration

- **Enable VXLAN:** Modern Data Center and Campus Networks where MSS L3 should be used to enforce policy need to support a standard based VXLAN overlay, which is leveraged by MSS to steer traffic to the firewall service. The control-plane can be either controller based (CVX) or EVPN based. There are no special requirements to the physical network topology, but it is recommended to follow state-of-the art L3 leaf and spine or campus PoD designs.
- **Routing at the TOR**: The transit network which is used to attach the firewall service needs to be adjacent to all MSS enabled TOR that hold L3 SVI's for end-hosts. That means that when using a controller-based control-plane (CVX), a direct routing model needs to be implemented. As of this writing, when using an EVPN based control-plane, asymmetric routing must be used. This restriction will be removed in future.
- **DirectFlow**: Directflow needs to be enabled on all the TORs.

Physical Topology

Figure 12 depicts the physical topology used in this reference design. The physical fabric is a Layer 3 ECMP eBGP based fabric with Layer 2 VXLAN overlays. In this scenario Tenant C and Tenant D workloads are hosted in the same physical infrastructure (Server 1 and Server 2), however they are logically isolated from communicating with each other directly. Server 1 and Server 2 are connected to Intercept VTEP 1 and Intercept VTEP 2 in MLAG configuration for HA. An active/standby pair of firewalls is attached to a Service VTEP 1 and Service VTEP 2 using a pair of physical interfaces or subinterface per zone.

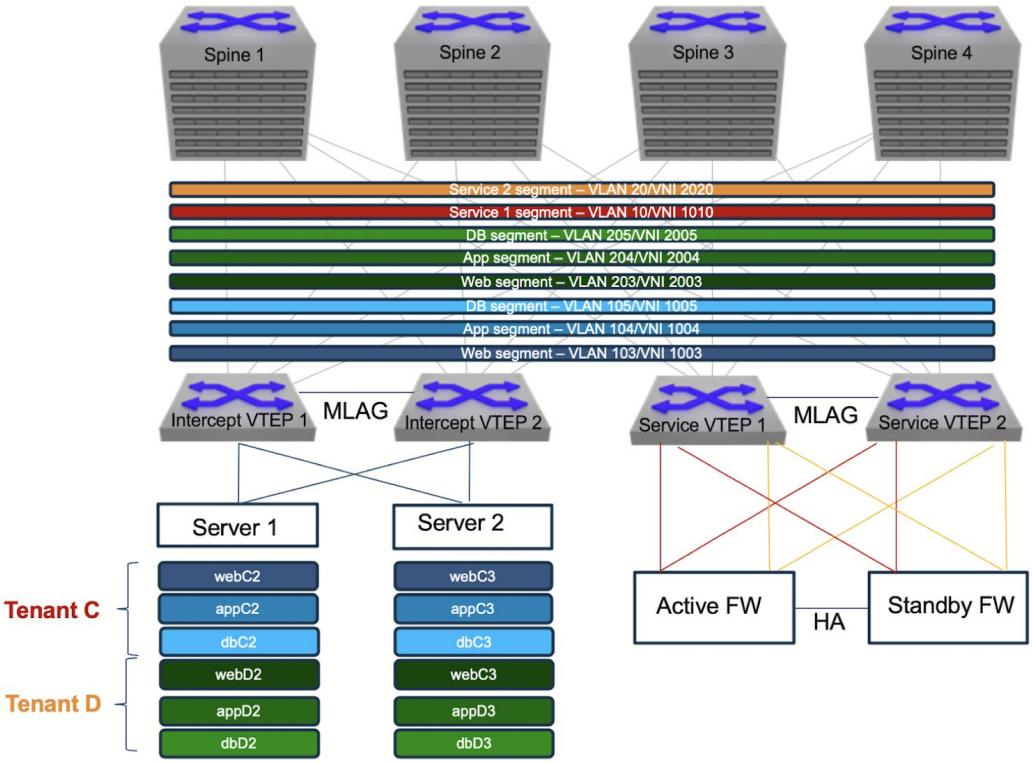

Figure 12: Physical topology for reference design

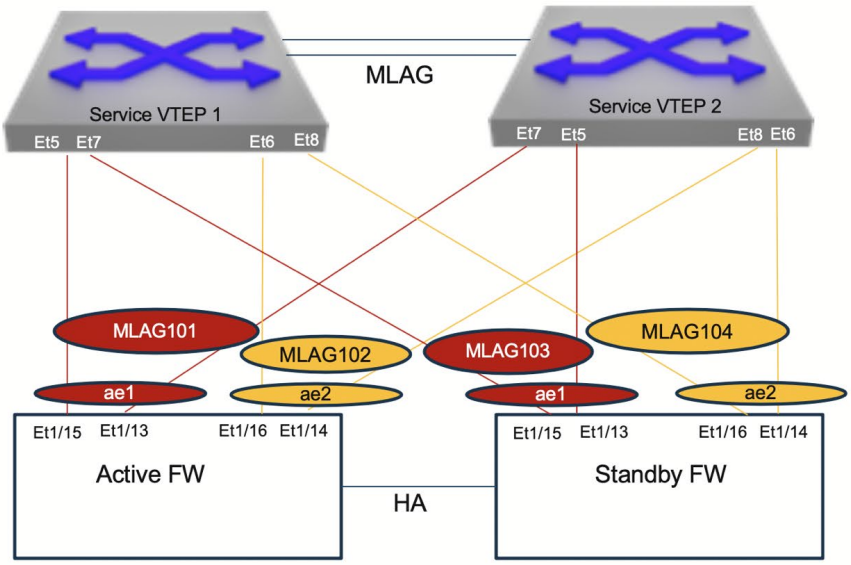

Figure 13: Connection between firewall and service VTEPs

Each tenant has a three tier application, consisting of multiple Web servers (Apache web service), App servers (Tomcat web based application service) and Database servers (MySQL database). Tenant C and D maps to two discrete zones in the Firewall Zone configuration.

The Inter-zone and intra-zone segmentation relies completely on MSS and Firewall policy configurations that are tagged appropriately to enforce security policies. There are no VRFs or Access Control Lists that are configured in the Arista switches to segment traffic between tenant's or within each tenant's applications.

Logical Topology

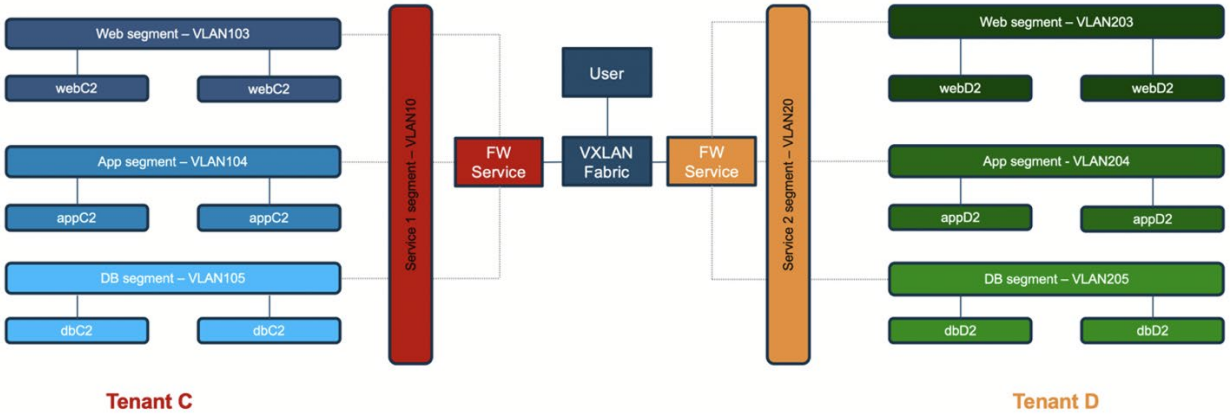

Figure 14: Logical topology for reference design

Segmentation Goal

This zone configuration isolates Tenant C and D from each other with the default policy (inter-zone default policy). Within each zone, traffic is permitted by default (intra-zone default policy).

In this particular configuration, traffic within the same zone is segmented based on the application profiles.

For Tenant C:

- Web servers can talk to each other directly only on SSL port 443, however unsecure HTTP port 80 traffic is denied even though the web servers are connected directly on the same VLAN and IP subnet.
- Web Servers can communicate with the app servers on port 8080 HTTP as well.
- Web Servers are not allowed to make any database queries over TCP port 3306 directly with the database servers.
- Only the application servers can make database queries over TCP port 3306 to the database servers.
- Enforcement point for the above policies should be on the Firewall

For Tenant D

- Similar requirements as mentioned above, except that the enforcement point for Tenant D workloads is on the intercept VTEPs.

These rules are defined in the Firewall and are enforced by the firewall or the Arista switches based on the Firewall policy tags.

For Tenant C, firewall policy tags defined as "MSS-redirect" are enforced by the Active Firewall.

For Tenant D, Firewall tags defined as "MSS-offload" are enforced by the Arista switches locally without being redirected to the Active firewall.

Configuration

The steps below outline how to configure Arista MSS.

Step 1: Deploy Arista CloudVision

The first step is to deploy Arista CloudVision and configure the Arista TOR switches to connect to it. A CVX cluster of 3 instances with hostnames of cvx01, cvx02, and cvx03 have been configured for this design.

Please see the CVX configuration guide for more information.

NOTE: It is a best practice to deploy CVX in a Highly Available (HA) cluster of three (3) instances. and have CVX clients have inband connectivity to CVX server.

Step 2: Enable the VXLAN Control Service on CVX

Once the three (3) Arista CVX instances have been deployed and the ToR switches have been configured to be managed by them, the VXLAN Control Service (VCS) must be enabled on every CVX instance.

The VXLAN control service allows hardware VXLAN Tunnel EndPoints (VTEPs) to share state with each other in order to establish VXLAN tunnels without the need for a multicast control plane.

All CVX instances

```
cvx(config-cvx)#service vxlan
cvx(config-cvx-vxlan)#no shutdown
```

Step 3: Configure Access switch and Service switch ports

This step involves configuration of the switch ports connected to the hosts, whose traffic needs to be steered to firewalls and 'service switch' which is connected to the firewalls.

Access switch configuration

For MLAG attached hosts

Access Switch (Intercept VTEP 1)

```
interface et21
description server1
channel-group 147 mode active
!
interface et22
description server2
channel-group 148 mode active
!
interface Port-Channel 147
description server1
switchport mode trunk
switchport trunk allowed vlan 103-105,203-205
mlag 147
!
interface Port-Channel 148
description server2
switchport mode trunk
switchport trunk allowed vlan 103-105,203-205
mlag 148
```

Access Switch (Intercept VTEP-2)

```
interface et21
description server1
channel-group 147 mode active
!
interface et22
description server2
channel-group 148 mode active
!
interface Port-Channel 147
description server1
switchport mode trunk
switchport trunk allowed vlan 103-105,203-205
mlag 147
!
interface Port-Channel 148
description server2
switchport mode trunk
switchport trunk allowed vlan 103-105,203-205
mlag 148
```

For single homed hosts

Access Switch (Intercept VTEP 1)

```
interface et21
description host1
switchport mode trunk
switchport trunk allowed vlan 103-105,203-205
!
interface et22
description host2
switchport mode access
```

Service Switch port configuration

A service switch is defined as the switch connecting to the firewalls.

Service Switch (Service VTEP-1)

```
interface Port-Channel1001
   description TenantA-zone1
   switchport access vlan 10
   mlag 101
!
interface Port-Channel1002
   switchport access vlan 20
   mlag 102
!
interface Port-Channel1003
   switchport access vlan 10
   mlag 103
!
interface Port-Channel1004
   switchport access vlan 20
   mlag 104
!
interface Ethernet5
   channel-group 1001 mode on
!
interface Ethernet6
   channel-group 1002 mode on
!
interface Ethernet7
   channel-group 1003 mode on
!
interface Ethernet8
   channel-group 1004 mode on
```

Service Switch (Service VTEP-1)

```
interface Port-Channel1001
   switchport access vlan 10
   mlag 101

!
interface Port-Channel1002
   switchport access vlan 20
   mlag 102
!
interface Port-Channel1003
   switchport access vlan 10
   mlag 103
!
interface Port-Channel1004
   switchport access vlan 20
   mlag 104
!
interface Ethernet5
   channel-group 1003 mode on
!
interface Ethernet6
   channel-group 1004 mode on
!
interface Ethernet7
   channel-group 1001 mode on
!
interface Ethernet8
    channel-group 1002 mode on
```

Note: The example above shows an access port. Trunk port can be used if the Firewall side is configured with L3 vlan sub interfaces.

Configuring VLAN to VNI Mapping

In VXLAN, VLAN to VNI mapping needs to be configured for all tenant and service vlans. All switches need to be configured identically:

```
interface Vxlan1
   vxlan source-interface Loopback1
   vxlan controller-client
   vxlan udp-port 4789
   vxlan vlan 10 vni 1010
   vxlan vlan 20 vni 1020
   vxlan vlan 103 vni 1003
   vxlan vlan 104 vni 1004
   vxlan vlan 105 vni 1005
   vxlan vlan 203 vni 2003
   vxlan vlan 204 vni 2004
   vxlan vlan 205 vni 2005
```

For more information about how to configure VXLAN, please see the configuration guide

Step 4: Enable DirectFlow on access switches

Arista MSS uses DirectFlow to steer interesting traffic from the intercepted host to the firewall. Directflow must be enabled on every intercept switch as well as the service switches.

The following config needs to be applied to all the intercept and service switches

```
!!!For all platforms except 7280R/R2  !!!

Switch# configure
Switch(config)# directflow
Switch(config-directflow)# no shutdown

!!! For 7280R/R2, enable directflow using below steps !!!

Step 1 : Configure the TCAM profile

!
hardware tcam
   profile direct-flow-mssl3-vxlan
      feature acl port ip
         sequence 50
         key size limit 160
         key field dscp dst-ip ip-frag ip-protocol l4-dst-port l4-ops l4-src-port src-ip
tcp-control ttl
         action count drop
         packet ipv4 forwarding bridged
         packet ipv4 forwarding routed
         packet ipv4 forwarding routed multicast
         packet ipv4 mpls ipv4 forwarding mpls decap
         packet ipv4 mpls ipv6 forwarding mpls decap
         packet ipv4 non-vxlan forwarding routed decap
         packet ipv4 vxlan eth ipv4 forwarding routed decap
         packet ipv4 vxlan eth ipv6 forwarding routed decap
         packet ipv4 vxlan forwarding bridged decap
      feature acl port ip egress mpls-tunnelled-match
         sequence 100
      feature acl port ipv6
         sequence 30
         key field dst-ipv6 ipv6-next-header ipv6-traffic-class l4-dst-port l4-ops-3b l4-src-
port src-ipv6-high src-ipv6-low tcp-control
         action count drop
         packet ipv6 forwarding bridged
         packet ipv6 forwarding routed
         packet ipv6 forwarding routed multicast
         packet ipv6 ipv6 forwarding routed decap
feature acl port mac
         sequence 60
         key size limit 160
         key field dst-mac ether-type src-mac
         action count drop
         packet ipv4 forwarding bridged
```

```
packet ipv4 forwarding bridged
         packet ipv4 forwarding routed
         packet ipv4 forwarding routed multicast
         packet ipv4 mpls ipv4 forwarding mpls decap
         packet ipv4 mpls ipv6 forwarding mpls decap
         packet ipv4 non-vxlan forwarding routed decap
         packet ipv4 vxlan eth ipv4 forwarding routed decap
         packet ipv4 vxlan forwarding bridged decap
         packet ipv6 forwarding bridged
         packet ipv6 forwarding routed
         packet ipv6 forwarding routed decap
         packet ipv6 forwarding routed multicast
         packet ipv6 ipv6 forwarding routed decap
         packet mpls forwarding bridged decap
         packet mpls ipv4 forwarding mpls
         packet mpls ipv6 forwarding mpls
         packet mpls non-ip forwarding mpls
         packet non-ip forwarding bridged
      feature acl subintf ip
         sequence 45
         key size limit 160
         key field dscp dst-ip ip-frag ip-protocol l4-dst-port l4-ops-18b l4-src-port src-ip
tcp-control ttl
         action count drop
         packet ipv4 forwarding routed
      feature acl subintf ipv6
         sequence 20
         key field dst-ipv6 ipv6-next-header l4-dst-port l4-src-port src-ipv6-high src-ipv6-
low tcp-control
         action count drop
         packet ipv6 forwarding routed
      feature acl vlan ip
         sequence 40
         key size limit 160
         key field dscp dst-ip ip-frag ip-protocol l4-dst-port l4-ops-18b l4-src-port src-ip
tcp-control ttl
         action count drop
         packet ipv4 forwarding routed
         packet ipv4 mpls ipv4 forwarding mpls decap
         packet ipv4 mpls ipv6 forwarding mpls decap
         packet ipv4 non-vxlan forwarding routed decap
         packet ipv4 vxlan eth ipv4 forwarding routed decap
         packet ipv4 vxlan eth ipv6 forwarding routed decap
      feature acl vlan ipv6
         sequence 15
key field dst-ipv6 ipv6-next-header l4-dst-port l4-src-port src-ipv6-high src-ipv6-low tcp-
control
         action count drop
         packet ipv6 forwarding routed
         packet ipv6 ipv6 forwarding routed decap
```

```
 feature acl vlan ipv6 egress
         sequence 25
         key field dst-ipv6 ipv6-next-header ipv6-traffic-class l4-dst-port l4-src-port src-
ipv6-high src-ipv6-low tcp-control
         action count drop
         packet ipv6 forwarding routed
      feature flow
         key size limit 160
         key field dst-ip ether-type in-port ip-protocol l4-dst-port l4-src-port src-ip
         action drop redirect set-fwd-header
         packet ipv4 forwarding bridged
         packet ipv4 forwarding routed
      feature forwarding-destination mpls
         sequence 105
      feature mpls
         sequence 5
         key size limit 160
         action drop redirect set-ecn
         packet ipv4 mpls ipv4 forwarding mpls decap
         packet ipv4 mpls ipv6 forwarding mpls decap
         packet mpls ipv4 forwarding mpls
         packet mpls ipv6 forwarding mpls
         packet mpls non-ip forwarding mpls
      feature mpls pop ingress
         sequence 95
      feature pbr mpls
         sequence 70
         key size limit 160
         key field mpls-inner-ip-tos
         action count drop redirect
         packet mpls ipv4 forwarding mpls
         packet mpls ipv6 forwarding mpls
         packet mpls non-ip forwarding mpls
      feature tunnel vxlan
         sequence 55
         key size limit 160
         key field in-port vxlan-inner-etype vxlan-inner-ip-options vxlan-inner-ip-ttl
         packet ipv4 vxlan eth ipv4 forwarding routed decap
         packet ipv4 vxlan eth ipv6 forwarding routed decap
         packet ipv4 vxlan forwarding bridged decap
      feature tunnel vxlan routing
         sequence 10
         packet ipv4 forwarding routed
         packet ipv4 non-vxlan forwarding routed decap
         packet ipv4 vxlan eth ipv4 forwarding routed decap
         packet ipv4 vxlan eth ipv6 forwarding routed decap


Step 2 : Enable Directflow
!
Switch#configure
Switch(config)#directflow
Switch(config-directflow)no shutdown
```

Additional configuration when using MLAG at the intercept VTEP

Access Switch (Intercept VTEP-1)

```
directflow
   no shutdown
  !
   flow bypassPeerlink
      priority 65535
      match input interface <peer-link>
      match ethertype ip
```

Access Switch (Intercept VTEP-2)

```
directflow
   no shutdown
   !
   flow bypassPeerlink
      priority 65535
      match input interface <peer-link>
      match ethertype ip
```

Additional configuration when using MLAG at the Service VTEP

Service Switch (Service VTEP-1)

```
directflow
   no shutdown
   !
   flow bypassPeerlink
      priority 65535
      match input interface <peer-link>
      match ethertype ip
!
   flow bypassServiceIntf
      priority 65535
      match input interface <service-interface>
      match ethertype ip
```

Service Switch (Service VTEP-2)

```
directflow
   no shutdown
   !
   flow bypassPeerlink
      priority 65535
      match input interface <peer-link>
      match ethertype ip
!
   flow bypassServiceIntf
      priority 65535
      match input interface <service-interface>
      match ethertype ip
```

Note: In the above configuration <peer-link> and <service-interface> can be ethernet interfaces or port-channel

Additional configuration on ALL TORs with Recirc-channel

Platforms that use recirc-channel for Vxlan routing will need an additional bypass direct flow rule. https://eos.arista.com/eos-4-15-2f/recirculation-channel/

Service Switch (Service VTEP-1)

```
directflow
  no shutdown
   !
   flow bypassRecirc-Channel
      priority 65535
      match input interface Re627
      match ethertype ip
```

Step 5: Enable routing on TOR switches

CVX uses Address Resolution Protocol (ARP) to determine where intercept hosts are physically located in the network.

It is required that VXLAN routing be configured on every ToR and service switch that will be intercepting traffic to ensure that CVX is aware of every host ARP entry. The configuration below shows the routing configuration for each tier of the application, but not the entire VXLAN configuration.

For more information on how to configure VXLAN and VXLAN routing, please see the VXLAN section in the EOS Configuration guide.

Note: For a DCS-7050X, DCS-7060X and DCS-7260X recirculation must be enabled.

For further information please refer to: https://eos.arista.com/eos-4-15-2f/vxlan-routing/

For R-Series platforms the following configuration must be configured:
https://eos.arista.com/eos-4-18-0f/arad-jericho-vxlan-routing/

Access switch (Intercept VTEP-1)

```
!
interface Vlan10
   ip address virtual 100.64.10.1/24
!
interface Vlan20
   ip address virtual 100.64.20.1/24
!
interface Vlan100
   ip address virtual 100.64.100.1/24
!
interface Vlan103
   ip address virtual 100.64.103.1/24
!
interface Vlan104
   ip address virtual 100.64.104.1/24
!
interface Vlan105
   ip address virtual 100.64.105.1/24
!
interface Vlan203
   ip address virtual 100.64.203.1/24
!
interface Vlan204
   ip address virtual 100.64.204.1/24
!
interface Vlan205
   ip address virtual 100.64.205.1/24
!
ip virtual-router mac-address 00:00:00:0a:0b:0c
```

Access switch (Intercept VTEP-2)

```
!
interface Vlan10
   ip address virtual 100.64.10.1/24
!
interface Vlan20
   ip address virtual 100.64.20.1/24
!
interface Vlan100
  ip address virtual 100.64.100.1/24
!
interface Vlan103
   ip address virtual 100.64.103.1/24
!
interface Vlan104
   ip address virtual 100.64.104.1/24
!
interface Vlan105
   ip address virtual 100.64.105.1/24
!
interface Vlan203
   ip address virtual 100.64.203.1/24
!
interface Vlan204
   ip address virtual 100.64.204.1/24
!
interface Vlan205
   ip address virtual 100.64.205.1/24
!
ip virtual-router mac-address 00:00:00:0a:0b:0c
```

Service switch (Service VTEP-1)

```
!
interface Vlan10
   ip address virtual 100.64.10.1/24
!
interface Vlan20
   ip address virtual 100.64.20.1/24
!
interface Vlan100
  ip address virtual 100.64.100.1/24
!
interface Vlan103
   ip address virtual 100.64.103.1/24
!
interface Vlan104
   ip address virtual 100.64.104.1/24
!
interface Vlan105
   ip address virtual 100.64.105.1/24
!
interface Vlan203
   ip address virtual 100.64.203.1/24
```

```
!
interface Vlan204
   ip address virtual 100.64.204.1/24
!
interface Vlan205
   ip address virtual 100.64.205.1/24
!
ip virtual-router mac-address 00:00:00:0a:0b:0c
```

Service switch (Service VTEP-2)

```
!
interface Vlan10
   ip address virtual 100.64.10.1/24
!
interface Vlan20
   ip address virtual 100.64.20.1/24
!
interface Vlan100
  ip address virtual 100.64.100.1/24
!
interface Vlan103
   ip address virtual 100.64.103.1/24
!
interface Vlan104
   ip address virtual 100.64.104.1/24
!
interface Vlan105
   ip address virtual 100.64.105.1/24
!
interface Vlan203
   ip address virtual 100.64.203.1/24
!
interface Vlan204
   ip address virtual 100.64.204.1/24
!
interface Vlan205
   ip address virtual 100.64.205.1/24
!
ip virtual-router mac-address 00:00:00:0a:0b:0c
```

Step 6: Enabling the MSS Service and Firewall configuration

The next step enables the Arista MSS service on CVX. The reference design includes three (3) CVX instances in a cluster, and the configuration must be the same for every instance.

Depending on the firewall, please refer to the appropriate Appendix for the MSS configuration.

Appendix 1A: Enabling MSS for Palo Alto Networks firewall
Appendix 2A: Enabling MSS for Fortinet firewall
Appendix 3A: Enabling MSS for Checkpoint firewall

Step 7: Firewall Configuration

For specific firewall configurations please use the appropriate Appendix.

Appendix 1B: Configuring Palo Alto Networks firewall
Appendix 2B: Configuring Fortinet firewall
Appendix 3B: Configuring Checkpoint firewall

Other Deployment Considerations

Deploying MSS with EVPN

In some deployments, the VXLAN fabric might already be using EVPN as the control plane. In such deployments, MSS is deployed by using collaboration of VCS and EVPN control planes. To support MSS in this scenario:

- There is no change to CVX configuration and VCS would still run on CVX
- Arista TORs would be configured as VCS client
- Arista TORs would also act as BGP EVPN peer
- VCS learns L2 information from EVPN and exports this info to MSS for creating intercepts on switches

Following line of configuration is added to each VTEP to make sure EVPN control plane takes priority for control plane learning

```
!
interface Vxlan1
   vxlan controller-client import vlan none
   vxlan controller-client
```

Guidelines

This section outlines some design guidelines for deploying Arista Macro-Segmentation Service.

- The firewall needs to have routes back to the original subnets in which the end hosts reside. Only static routes are supported (no dynamic protocols)
- For MSS redirect verbatim tag
  - Must have IP address specified in source or destination field if the corresponding zone is an external zone (zone towards default route)
  - Must have either zone or IP specified in both source and destination field. 'Any', 'All', or similar constructs are not supported for source or destination fields. (applicable without verbatim tag as well)
- For MSS offload verbatim tag
  - The offload tag, when applied to a policy, identifies a flow based on five-tuple and the action (permit | deny) that is then enforced on the TOR switches.
  - For policies tagged with the offload tag, at least one field (Source/Destination zone, Source/ Destination IP, protocol or L4 port) must be defined (and not = ANY).
  - Must have IP address specified in source or destination field if the corresponding zone is an external zone (zone towards default route)
- Any tagged policy cannot have source or destination IPv6 address and cannot have IPv4 address with prefix length less than /16.
- MSS is supported only with Active/Standby firewall configuration or single firewall configuration
- MSS will read policies from only one Virtual System/ Virtual Domain per device set.

Troubleshooting

This section contains a series of troubleshooting steps when traffic from a particular host are not being correctly intercepted by MSS. Please refer to the appropriate appendix for troubleshooting

Appendix 1C: Troubleshooting MSS and Palo Alto Networks firewall
Appendix 2C: Troubleshooting MSS and Fortinet firewall
Appendix 3C: Troubleshooting MSS and Checkpoint firewall

Appendix 1A: Enabling MSS for Palo Alto Networks Firewall

Requirements

- Arista CloudVision running EOS release 4.23.2F or later
- Palo Alto Networks firewall(s) running PAN-OS version 8.0.8 or greater
- Optionally Palo Alto Networks Panorama running PAN-OS version 8.0.8 or greater

MSS with Palo Alto is supported on the following Arista platforms

Platform	Minimum EOS Version required
Arista 7050X/X2/X3 and 7060X/X2 Series	4.23.2F
Arista 7280R/7280R2 Series	4.23.2F

Configure the MSS service on CVX for Palo Alto Firewalls

In this step both the active and standby Palo Alto firewalls are configured. If Panorama is used, only Panorama needs to be configured.

Command	Description
`service mss`	Enables the MSS service on CVX
`dynamic device-set <name>`	Creates a set of devices, typically a pair of firewalls or security manager
`tag redirect <list of tags>`	Specifies the tag(s) that MSS looks for when reading security policy from the firewall or security manager for redirecting traffic. Any tag in the list can be used to tag the policy
`tag redirect <list of tags>`	Specifies the tag(s) that MSS looks for when reading security policy from the firewall or security manager for offloading rules. Any tag in the list can be used to tag the policy
`tag modifier verbatim <list of tags>`	The verbatim modifier changes the behaviour of MSS to install policies exactly as they are defined on the Firewall. Without the verbatim modifier when defining a policy for redirection between two subnets A and B, all traffic within the subnet A as well as within B is redirected. Any tag in the list can be used to tag the policy
`type palo-alto panorama`	Sets the firewall type
`state active`	Allows you set the device set as active or disabled
`device <firewall Name>`	Defines a device. Note this is the hostname or IP address that MSS will use to communicate with it
`username admin password 0 admin`	Sets the username and password to access the device. Once entered the password is encrypted

All CVX instances

```
!
   service mss
      no shutdown

      !
      dynamic device-set PAN
         device 10.92.59.101
            username admin password 7 CF+X7x7GbctS7QTS+u8kaQ==
            group Arista-MSS-Stack
         state active
         type palo-alto panorama
         policy tag redirect MSS-redirect
         policy tag offload MSS-offload
         policy tag modifier verbatim MSS-verbatim

   !
   service vxlan
      no shutdown
```

Appendix 1B: Configuring Palo Alto Networks firewall

Interface Configuration

Interfaces have been configured in aggregation groups ae1 and ae2. ae1 is part of the zone 'Tenant C' and ae2 is part of the zone 'Tenant D'

```
demo@PA-MGMT37-A(active)> show interface all

total configured hardware interfaces: 11

name                    id    speed/duplex/state        mac address
--------------------------------------------------------------------------------
ethernet1/13            28    1000/full/up              00:1b:17:00:25:1c
ethernet1/14            29    1000/full/up              00:1b:17:00:25:1d
ethernet1/15            30    1000/full/up              00:1b:17:00:25:1e
ethernet1/16            31    1000/full/up              00:1b:17:00:25:1f
ae1                     48    [n/a]/[n/a]/up            00:1b:17:00:25:30
ae2                     49    [n/a]/[n/a]/up            00:1b:17:00:25:31
dedicated-ha1           5     1000/full/up              24:0b:0a:00:aa:3b
dedicated-ha2           6     1000/full/up              00:86:9c:3c:60:06
vlan                    1     [n/a]/[n/a]/up            00:1b:17:00:25:01
loopback                3     [n/a]/[n/a]/up            00:1b:17:00:25:03
tunnel                  4     [n/a]/[n/a]/up            00:1b:17:00:25:04

aggregation groups: 2
ae1 members:
  ethernet1/13 ethernet1/15
ae2 members:
  ethernet1/14 ethernet1/16


total configured logical interfaces: 7

name                id     vsys zone           forwarding           tag    address
------------------ ------ ---- -------------- -------------------- ------ ------------------
ae1                 48     1    Tenant-C       vr:default            0      100.64.10.254/24
ae2                 49     1    Tenant-D       vr:default            0      100.64.20.254/24
dedicated-ha1       5      1                   ha                    0      10.92.159.27/24
dedicated-ha2       6      1                   ha                    0      10.92.58.29/24
vlan                1      1                   N/A                   0      N/A
loopback            3      1                   N/A                   0      N/A
tunnel              4      1                   N/A                   0      N/A
```

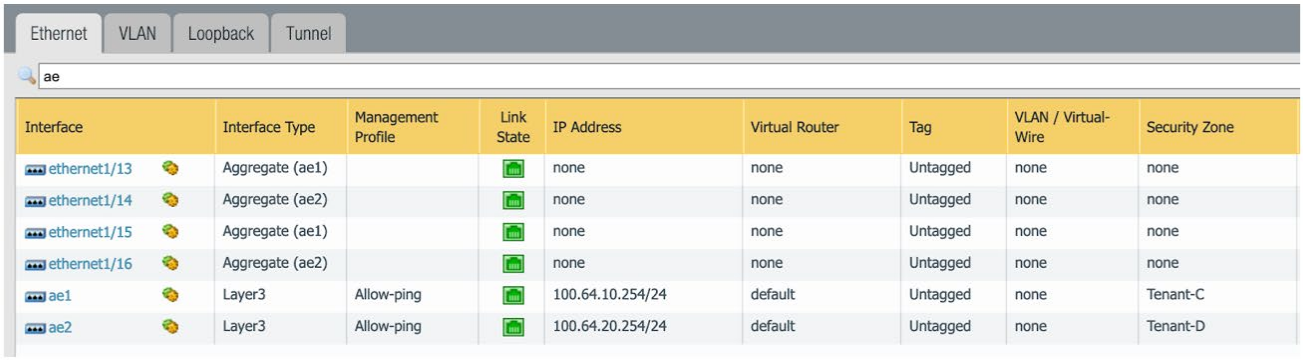

Interface	Interface Type	Management Profile	Link State	IP Address	Virtual Router	Tag	VLAN / Virtual-Wire	Security Zone
ethernet1/13	Aggregate (ae1)			none	none	Untagged	none	none
ethernet1/14	Aggregate (ae2)			none	none	Untagged	none	none
ethernet1/15	Aggregate (ae1)			none	none	Untagged	none	none
ethernet1/16	Aggregate (ae2)			none	none	Untagged	none	none
ae1	Layer3	Allow-ping		100.64.10.254/24	default	Untagged	none	Tenant-C
ae2	Layer3	Allow-ping		100.64.20.254/24	default	Untagged	none	Tenant-D

Figure 15: Interface configuration for Palo Alto Firewall

Route Configuration

The firewall needs to have routes back to the original subnets in which the end hosts reside. Static routes have been created for each subnet. Default GW for the Tenant C subnets is the VLAN 10 interface on the TOR and for Tenant D subnets its the VLAN 20 interface on the TOR

```
demo@PA-MGMT37-A(active)> show routing route
flags: A:active, ?:loose, C:connect, H:host, S:static, ~:internal, R:rip, O:ospf, B:bgp,
       Oi:ospf intra-area, Oo:ospf inter-area, O1:ospf ext-type-1, O2:ospf ext-type-2,
E:ecmp, M:multicast
  VIRTUAL ROUTER: default (id 1)
  ==========
destination                         nexthop                              metric flags
age   interface
100.64.10.0/24                      100.64.10.254                         0      A C
ae1
100.64.10.254/32                    0.0.0.0                               0      A H
100.64.20.0/24                      100.64.20.254                         0      A C
ae2
100.64.20.254/32                    0.0.0.0                               0      A H
100.64.103.0/24                     100.64.10.1                           10     A S
ae1
100.64.104.0/24                     100.64.10.1                           10     A S
ae1
100.64.105.0/24                     100.64.10.1                           10     A S
ae1
100.64.203.0/24                     100.64.20.1                           10     A S
ae2
100.64.204.0/24                     100.64.20.1                           10     A S
ae2
100.64.205.0/24                     100.64.20.1                           10     A S
ae2
total routes shown: 10
```

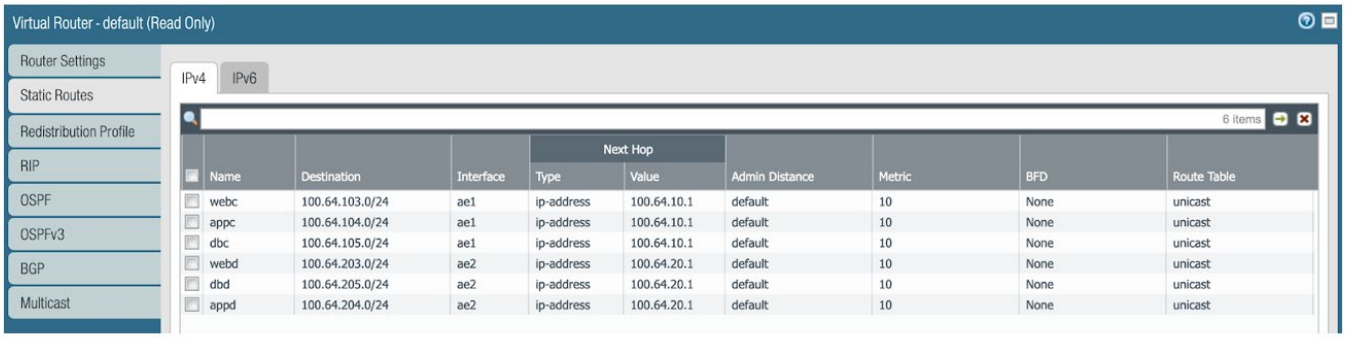

Figure 16: Route configuration for Palo Alto Firewall

Policy Configuration

For this reference design, following policies are created in addition to the default policies for interzone traffic and intrazone traffic

For Tenant C:

- Web servers can talk to each other directly only on SSL port 443, however unsecure HTTP port 80 traffic is denied even though the web servers are connected directly on the same VLAN and IP subnet.
- Web Servers can communicate with the app servers on port 8080 HTTP as well.
- Web Servers are not allowed to make any database queries over TCP port 3306 directly with the database servers.
- Only the application servers can make database queries over TCP port 3306 to the database servers.

For Tenant D

- Similar requirements as mentioned above, except that the enforcement point for Tenant D workloads is on the intercept VTEPs.

These rules are defined in the Firewall and are enforced by the firewall or the Arista switches based on the Firewall policy tags.

					Source				Destination		
	Name	Tags	Action	Type	Zone	Address	User	HIP Profile	Zone	Address	Service
1	webc2webc-http	MSS-redirect	Deny	intrazone	Tenant-C	webc	any	any	(intrazone)	webc	apache-http-80
2	webc2webc-ssl	MSS-redirect	Allow	intrazone	Tenant-C	webc	any	any	(intrazone)	webc	apache-ssl-443
3	webc2appc-http	MSS-redirect	Allow	intrazone	Tenant-C	webc	any	any	(intrazone)	appc	tomcat-http-8080
4	webc2dbc-mysql	MSS-redirect	Deny	intrazone	Tenant-C	webc	any	any	(intrazone)	dbc	mysql-3306
5	appc2dbc-mysql	MSS-redirect	Allow	intrazone	Tenant-C	appc	any	any	(intrazone)	dbc	mysql-3306
6	webd2webd-http	MSS-offload	Deny	intrazone	Tenant-D	webd	any	any	(intrazone)	webd	apache-http-80
7	webd2webd-ssl	MSS-offload	Allow	intrazone	Tenant-D	webd	any	any	(intrazone)	webd	apache-ssl-443
8	webd2appd-http	MSS-offload	Allow	intrazone	Tenant-D	webd	any	any	(intrazone)	appd	tomcat-http-8080
9	webd2dbd-mysql	MSS-offload	Deny	intrazone	Tenant-D	webd	any	any	(intrazone)	dbd	mysql-3306
10	appd2dbd-mysql	MSS-offload	Allow	intrazone	Tenant-D	appd	any	any	(intrazone)	dbd	mysql-3306
11	intrazone-default	none	Allow	intrazone	any	any	any	any	(intrazone)	any	any
12	interzone-default	none	Deny	interzone	any	any	any	any	any	any	any

Figure 17: Policy Configuration for Palo Alto Networks Firewall

For Tenant C, firewall policy tags defined as "MSS-redirect" are enforced by the Active Firewall.

For Tenant D, Firewall tags defined as "MSS-offload" are enforced by the Arista switches locally without being redirected to the Active firewall.

Appendix 1C: Troubleshooting MSS for Palo Alto Networks firewall

Check if CVX service is enabled

```
mss-l3-cvx#show cvx
CVX Server
  Status: Enabled
  UUID: c84d66ae-2d7f-11ea-8264-7f39ec151b3c
  Mode: Standalone
  Heartbeat interval: 20.0
  Heartbeat timeout: 60.0
  Client connection state preserving: Disabled
mss-l3-cvx#
```

Check if all TOR switches are registered for VCS on CVX

```
mss-l3-cvx#show cvx connections brief
          CVX Connection Table
Switch            Hostname        Status
----------------- --------------- ------
44:4c:a8:23:ba:5f intercept-vtep1 active
44:4c:a8:23:27:d1 intercept-vtep2 active
44:4c:a8:23:24:ed service-vtep    active

mss-l3-cvx#show cvx service vxlan
Vxlan
  Status: Enabled
  Supported versions: 2

  Switch            Status  Negotiated Version
  ----------------- ------- ------------------
  44:4c:a8:23:ba:5f Enabled 2
  44:4c:a8:23:24:ed Enabled 2
  44:4c:a8:23:27:d1 Enabled 2

mss-l3-cvx#show service vxlan switch all capability | grep Hos
  Hostname     : intercept-vtep1
  Hostname     : service-vtep
  Hostname     : intercept-vtep2
```

Check if all TOR switches are registered with MSS

```
mss-l3-cvx#show cvx service mss
Mss
  Status: Enabled
  Supported versions: 1

  Switch            Status  Negotiated Version
  ----------------- ------- ------------------
  44:4c:a8:23:ba:5f Enabled 1
  44:4c:a8:23:24:ed Enabled 1
  44:4c:a8:23:27:d1 Enabled 1
```

Check ARP learning via CVX

```
mss-l3-cvx#show service vxlan arp received
                        Received ARP Table
------------------------------------------------------------------------
Switch               VNI      IP Address      MAC Address        Changes
-------------------- -------- --------------- ------------------ ---------
44-4c-a8-23-24-ed    10070    10.10.70.254    5849.3b1f.4412     0
44-4c-a8-23-24-ed    10080    10.10.80.254    5849.3b1f.4413     0
44-4c-a8-23-24-ed    10090    10.10.90.254    5849.3b1f.4413     0

Total IP addresses for this criterion: 3
```

Ensure the MSS service is enabled

The MSS service should be enabled on every CVX instance. To verify, run the following commands on CVX:

```
mss-l3-cvxC#show service mss status
State: Enabled
<snip>
```

The MSS service should be enabled on every CVX instance. To verify, run the following commands on CVX:

```
mss-l3-cvx#show service mss dynamic
Total policies processed: 5064
Policy Source      Device Set Service Device State
------------------ ---------- -------------- ----------
palo-alto-panorama PANORAMA   10.92.59.101   active

mss-l3-cvx#show service mss dynamic status
Service Device Policy Monitoring Status:
Device: 001801053738
  IP address: 10.90.164.169
  HA peer IP: 10.90.164.170
  Policy source type: PaloAltoPanorama
  Accessed via Aggregation Manager: 10.92.59.101
  Device set name: PANORAMA
  Device set state: Active
  Total policies processed: 0
  Last seen at time: 2020 Apr 08, 18:46:42
Device: 001801053832
  IP address: 10.90.164.170
  HA peer IP: 10.90.164.169
  Policy source type: PaloAltoPanorama
  Accessed via Aggregation Manager: 10.92.59.101
  Device set name: PANORAMA
  Device set state: Active
  Total policies processed: 3556070
  Last seen at time: 2020 Apr 08, 18:46:33
Device: 10.92.59.101
  IP address: 10.92.59.101
  Policy source type: PaloAltoPanorama
  Aggregation Manager: True
  Device group member(s):
    001801053738
    001801053832
  Device set name: PANORAMA
  Device set state: Active
  Last seen at time: 2020 Apr 08, 18:46:43
```

The command shows the connection to the firewall or security manager. The output of state active might be misleading since it just reflects the configuration state of the device and does not reflect any API status.

To be sure that API connection is fine we can use Total policies processed. These counters must rise every time the firewall gets queried for policies (every 15s in default) and reflects the overall number of policies analyzed by MSS whether a MSS relevant tag is set or not.

Also the Last seen at time output shows if it ever had a working connection to the firewall.

Policy is not fetched from the firewall correctly

cvx# show service mss dynamic device-set <device_set_name> device <device_name> policies

If output for above command no policies are seen by Arista MSS, check if CVX is able to communicate with the firewall or security manager using the following command:

Below command shows policies that are valid andare installed into the network.

If policy is missing, check the policy tag configured on the firewall policy.

```
CVX-01#show service mss policy
                Macro-Segmentation L2 Policy Table
------------------------------------------------------------------------------
   Source       Device       Policy       Configured State    Operational State
------------    ------------ ------------ ------------------  -----------------
                Macro-Segmentation L3 Policy Table
------------------------------------------------------------------------------
  Source              Device               Policy           Offload   Redirect    Unconverged
                                                                        status
status         IPs
---------------------- --------------------  ----------------- --------- ---------- -------
   PaloAltoPanorama     001801053832_HAPair   appc2dbc-mysql    N/A       Active      0 of 4
   PaloAltoPanorama     001801053832_HAPair   appd2dbd-mysql    Active    Active      0 of 4
   PaloAltoPanorama     001801053832_HAPair   webc2appc-http    N/A        Active     0 of 4
   PaloAltoPanorama     001801053832_HAPair    webc2dbc-mysql   N/A        Active     0 of 4
   PaloAltoPanorama     001801053832_HAPair    webc2webc-http   N/A        Active     0 of 2
   PaloAltoPanorama     001801053832_HAPair    webc2webc-ssl    N/A        Active     0 of 2
   PaloAltoPanorama     001801053832_HAPair    webd2appd-http   Active     Active     0 of 4
   PaloAltoPanorama     001801053832_HAPair    webd2dbd-mysql   Active     Active     0 of 4
   PaloAltoPanorama     001801053832_HAPair    webd2webd-http   Active     Active     0 of 2
   PaloAltoPanorama     001801053832_HAPair    webd2webd-ssl    Active     Active     0 of 2
```

IP-MAC binding not learned by CVX

Check the status of the policy to ensure that CVX has the necessary information to redirect traffic:

```
CVX-01#show service mss policy name webd2appd-http
             Macro-Segmentation L2 Policy Table
------------------------------------------------------------------------------
   Source  Device         Policy       Configured State    Operational State
------------ ------------ ------------ -------------------- -----------------
             Macro-Segmentation L3 Policy Table
------------------------------------------------------------------------------
      Source                    Device                    Policy          Offload   Redirect
Unconverged                                                                         status
status     IPs
---------------------- ------------------------ -------------------- ------------- -------
   PaloAltoPanorama 001801053832_HAPair  webd2appd-http          Active   Active  0 of 4
```

If the policy status is not "active," check the ARP table information received by CVX.

```
cvx# show service vxlan arp received
Received ARP Table
------------------------------------------------------------------------
Switch               VNI      IP Address      MAC Address        Changes
-------------------- -------- --------------- ------------------ ---------
00-00-91-02-00-00    1000     10.10.10.102    00:00:01:02:00:00  0
00-1c-73-00-e2-16    2000     10.10.20.103    00:00:01:03:00:00  0
```

If the IP Address of the host is not seen in the CVX ARP table, ICMP ping the host which is not on the same subnet (or VNI) as the intercept host and verify ARP table information again. If ARP information for the host is learned by the CVX after the ping, check the status of the policy and ensure it's "active".

If the situation still persists, run the following commands on the intercept VTEP. If the host MAC address is learned on the VXLAN interface (Vxlan1), this indicates that there is a Layer 2 (L2) loop in the network. Resolve the loop and verify the policy status again.

```
intercept-switch# show arp
Address         Age (min)  Hardware Addr   Interface
10.10.100.1           N/A  0000.0101.0000  Vlan100, Vxlan1

intercept-switch# show mac address-table
Mac Address Table
------------------------------------------------------------------

Vlan    Mac Address       Type        Ports      Moves   Last Move
----    -----------       ----        -----      -----   ---------
100    0000.0101.0000    DYNAMIC     Vx1        186     0:23:02 ago
```

Check if Direct Flow is enabled on all TOR switches

```
intercept-vtep1.09:35:02#sh directflow
DirectFlow configuration: Enabled
Total matched: 347 packets
Total programmed flows: 16 flows
intercept-vtep1.09:35:04#
```

Appendix 2A: Enabling MSS for Fortinet firewall

Requirements

- Arista CloudVision running EOS release 4.23.2F or later
- FortiManager running 5.6.3 or later
- Fortigate firewall(s) running FortiOS version 5.6.3 or later

MSS with Fortinet is supported on the following Arista platforms

Platform	Minimum EOS Version required
Arista 7050X/X2/X3 and 7060X/X2 Series	4.23.2F
Arista 7280R/7280R2 Series	4.23.2F

Configure the MSS service on CVX for Fortinet Fortigate Firewalls

In this step both the active and standby Fortinet firewalls and the Fortimanager are configured.

In this reference design, the Fortimanager has a DNS name of "fm-node-1". The username and password are both "admin".

Command	Description
`service mss`	Enables the MSS service on CVX
`dynamic device-set fnet`	Creates a set of devices, typically a pair of firewalls with the name 'fnet'.
`device <IP or Hostname>`	Defines a device. Note this is the hostname or IP address that MSS will use to communicate with Fortimanager
`username admin password 0 <PW>`	Sets the username and password to access the device. Once entered the password is encrypted
`group Demo`	Defines the device group from the Fortimanager that will be used
`type fortinet fortimanager`	Sets the type to Fortinet Fortimanager
`tag redirect <list of tags>`	Specifies the tag(s) that MSS looks for when reading security policy from the firewall or security manager for redirecting traffic. Any tag in the list can be used to tag the policy
`tag redirect <list of tags>`	Specifies the tag(s) that MSS looks for when reading security policy from the firewall or security manager for offloading rules. Any tag in the list can be used to tag the policy
`tag modifier verbatim <list of tags>`	The verbatim modifier changes the behaviour of MSS to install policies exactly as they are defined on the Firewall. Without the verbatim modifier when defining a policy for redirection between two subnets A and B, all traffic within the subnet A as well as within B is redirected. Any tag in the list can be used to tag the policy
`admin domain Demo`	Set the admin domain on the Fortimanager to 'Demo'
`state active`	Allows you set the device set as active or disabled
`virtual domain default`	Specifies the name of the VDOM that will be used with MSS
`device member <cluster-name>`	Adds the Fortinet Firewall or Firewall Cluster name that will be used in the above defined device-set

All CVX instance

```
!
   service mss
      no shutdown
      !
      dynamic device-set fnet
         device 10.90.164.220
            username admin password 7 CF+X7x7GbctS7QTS+u8kaQ==
            group Demo
         !
         device member DEMO-FORT-A
         !
         device member DEMO-FORT-B
         !
         device member MSS-Demo
         state active
         type fortinet fortimanager
         policy tag redirect MSS-redirect
         policy tag offload MSS-offload
         policy tag modifier verbatim MSS-verbatim
         admin domain Demo
         virtual domain default
   !
   service vxlan
      no shutdown
```

Appendix 2B: Configuring Fortinet firewall

Interface Configurations

Interfaces have been configured in aggregation groups ae1 and ae2. ae1 is part of the zone 'Tenant C' and ae2 is part of the zone 'Tenant D'

DEMO-FORT-A : default | System : Interface | Router

View | Collapse All | More | Column Settings

Name	Type	Mapped Policy Interface	IP/Netmask	Virtual Domain	Status	Administrative Status
Aggregate (2)						
ae1	Aggregate	ae1	0.0.0.0/0.0.0.0	default		Up
ae2	Aggregate	ae2	0.0.0.0/0.0.0.0	default		Up
Tunnel (1)						
ssl.default (SSL VPN interface)	Tunnel	ssl.default	0.0.0.0/0.0.0.0	default		Up
VLAN (2)						
VLAN10	VLAN	VLAN10	100.64.10.254/255.255.255.0	default	Up	Up
VLAN20	VLAN	VLAN20	100.64.20.254/255.255.255.0	default	Up	Up

Figure 18: Interface Configuration for Fortinet Firewall

Route Configuration

The firewall needs to have routes back to the original subnets in which the end hosts reside. Static routes have been created for each subnet. Default GW for the Tenant C subnets is the VLAN 10 interface on the TOR and for Tenant D subnets its the VLAN 20 interface on the TOR.

DEMO-FORT-A : default | System | Router : Static Route

View | Column Settings

ID	Destination	Gateway	Interface
Static Route (2)			
1	100.64.0.0/255.192.0.0	100.64.20.1	VLAN20
2	100.64.0.0/255.192.0.0	100.64.10.1	VLAN10

Figure 19: Route Configuration for Fortinet Firewall

Policy Configuration

For this reference design, following policies are created in addition to the default policies for interzone traffic and intrazone traffic

For Tenant C:

- Web servers can talk to each other directly only on SSL port 443, however unsecure HTTP port 80 traffic is denied even though the web servers are connected directly on the same VLAN and IP subnet.
- Web Servers can communicate with the app servers on port 8080 HTTP as well.
- Web Servers are not allowed to make any database queries over TCP port 3306 directly with the database servers.
- Only the application s ervers can make database queries over TCP port 3306 to the database servers.

For Tenant D

- Similar requirements as mentioned above, except that the enforcement point for Tenant C workloads is on the intercept VTEPs.

These rules are defined in the Firewall and are enforced by the firewall or the Arista switches based on the Firewall policy tags. Tags must be added in the Comments section of the policy using "tags(<comma separated list of tags>)".

Create New | Edit | Delete | Section | Column Settings

#	Name	From	To	Source	Destination	Schedule	Service	Comments	Action
TENANT-C (1-5 / Total:5)									
1	webe2webe-http	any	any	webe	webe	always	HTTP	tags(MSS-redirect)	Deny
2	webe2webe-ssl	any	any	webe	webe	always	HTTPS	tags(MSS-redirect)	Accept
3	webe2appe-http	any	any	webe	appe	always	TOMCAT	tags(MSS-redirect)	Accept
4	webe2dbe-mysql	any	any	webe	dbe	always	MYSQL	tags(MSS-redirect)	Deny
5	appe2dbe-mysql	any	any	appe	dbe	always	MYSQL	tags(MSS-redirect)	Accept
TENANT-D (6-5 / Total:5)									
6	webf2webf-http	any	any	webf	webf	always	HTTP	tags(MSS-offload)	Deny
7	webf2webf-ssl	any	any	webf	webf	always	HTTPS	tags(MSS-offload)	Accept
8	webf2appf-http	any	any	webf	appf	always	TOMCAT	tags(MSS-offload)	Accept
9	webf2dbf-mysql	any	any	webf	dbf	always	MYSQL	tags(MSS-offload)	Deny
10	appf2dbf	any	any	appf	dbf	always	MYSQL	tags(MSS-offload)	Accept
Implicit (11-11 / Total:1)									
11	Implicit Deny	any	any	all	all	always	ALL		Deny

Figure 20: Policy Configuration for Fortinet Firewall

For Tenant C, firewall policy tags defined as "MSS-redirect" are enforced by the Active Firewall.

For Tenant D, Firewall tags defined as "MSS-offload" are enforced by the Arista switches locally without being redirected to the Active firewall.

Appendix 2C: Troubleshooting MSS for Fortinet firewall

Check if CVX service is enabled

```
mss-l3-cvx#show cvx
CVX Server
  Status: Enabled
  UUID: c84d66ae-2d7f-11ea-8264-7f39ec151b3c
  Mode: Standalone
  Heartbeat interval: 20.0
  Heartbeat timeout: 60.0
  Client connection state preserving: Disabled
mss-l3-cvx#
```

Check if all TOR switches are registered for VCS on CVX

```
mss-l3-cvx#show cvx connections brief
          CVX Connection Table
Switch            Hostname        Status
----------------- --------------- ------
44:4c:a8:23:ba:5f intercept-vtep1 active
44:4c:a8:23:27:d1 intercept-vtep2 active
44:4c:a8:23:24:ed service-vtep    active
mss-l3-cvx#show cvx service vxlan
Vxlan
  Status: Enabled
  Supported versions: 2

  Switch            Status  Negotiated Version
  ----------------- ------- ------------------
  44:4c:a8:23:ba:5f Enabled 2
  44:4c:a8:23:24:ed Enabled 2
  44:4c:a8:23:27:d1 Enabled 2
mss-l3-cvx#show service vxlan switch all capability | grep Hos
  Hostname     : intercept-vtep1
  Hostname     : service-vtep
  Hostname     : intercept-vtep2
```

Check if all TOR switches are registered with MSS

```
mss-l3-cvx#show cvx service mss
Mss
  Status: Enabled
  Supported versions: 1

  Switch            Status  Negotiated Version
  ----------------- ------- ------------------
  44:4c:a8:23:ba:5f Enabled 1
  44:4c:a8:23:24:ed Enabled 1
  44:4c:a8:23:27:d1 Enabled 1
```

Check ARP learning via CVX

```
mss-l3-cvx#show service vxlan arp received
                        Received ARP Table
------------------------------------------------------------------------
Switch              VNI      IP Address      MAC Address        Changes
------------------- -------- --------------- ------------------ ---------
44-4c-a8-23-24-ed   10070    10.10.70.254    5849.3b1f.4412     0
44-4c-a8-23-24-ed   10080    10.10.80.254    5849.3b1f.4413     0
44-4c-a8-23-24-ed   10090    10.10.90.254    5849.3b1f.4413     0

Total IP addresses for this criterion: 3
```

Ensure the MSS service is enabled

The MSS service should be enabled on every CVX instance. To verify, run the following commands on CVX:

```
CVX-01#show service mss status
State: Enabled
<snip>
```

The state should be "Enabled", and reflect the configured VNI range.

Check if API connection is working

```
CVX-04#show service mss dynamic
Total policies processed: 0
Policy Source         Device Set Service Device State
--------------------- ---------- -------------- ----------
fortinet-fortimanager fnet       10.90.164.220  active
```

The state should be "active". ge.

Policy is not fetched from the firewall correctly

The following command will list all the policies retrieved from the firewall by Arista MSS:

```
cvx# show service mss dynamic device-set <device_set_name> device <device_name> policies
```

If no policies are seen by Arista MSS, check if CVX is able to communicate with the firewall or security manager using the following command:

```
cvx# show service mss dynamic status
Service Device Policy Monitoring Status:

Device: fm-node-1
  Policy source type: FortinetFortiManager
  Aggregation Manager: True
  Device group member(s):
    FGT80
  Device set name: fnet
  Device set state: Active
  Last seen at time: 2019 Feb 14, 16:09:44
```

```
Device: fw-ha-node-1
  Policy source type: FortinetFortiManager
  Accessed via Aggregation Manager: 192.168.1.99
  Device set name: fnet
  Device set state: Active
  Total policies processed: 4
  Last seen at time: 2019 Feb 14, 16:09:41
```

NOTE: The number of policies processed is an aggregate number and should increment every time MSS polls for new policies.

If a specific policy is missing, check the policy tag configured on the firewall policy.

IP-MAC binding not learned by CVX

Check the status of the policy to ensure that CVX has the necessary information to redirect traffic:

```
CVX-01#show service mss policy name policy1
               Macro-Segmentation L2 Policy Table
------------------------------------------------------------------------------
   Source      Device         Policy         Configured State       Operational State
------------ ------------ ------------ --------------------- -----------------

               Macro-Segmentation L3 Policy Table
------------------------------------------------------------------------------
       Source               Device               Policy               Offload       Redirect      Unconverged
                                                                          status        status
IPs
---------------------- ------------------------ -------------------- ------------- -------------- ---------
--
   fortinet-fortimanager      10.90.164.220             policy1       Active         Active      0 of 4
```

If the policy status is not "active," check the ARP table information received by CVX.

```
cvx# show service vxlan arp received
Received ARP Table
------------------------------------------------------------------------
Switch              VNI      IP Address      MAC Address        Changes
------------------- -------- --------------- ------------------ ---------
00-00-91-02-00-00   1000     10.10.10.102    00:00:01:02:00:00  0
00-1c-73-00-e2-16   2000     10.10.20.103    00:00:01:03:00:00  0
```

If the IP Address of the host is not seen in the CVX ARP table, ICMP ping a host which is not on the same subnet (or VNI) as the intercept host and verify ARP table information again. If ARP information for the host is learned by the CVX after the ping, check the status of the policy and ensure it's "active".

If the situation still persists, run the following commands on the intercept VTEP. If the host MAC address is learned on the VXLAN interface (Vxlan1), this indicates that there is a Layer 2 (L2) loop in the network. Resolve the loop and verify the policy status again.

```
intercept-switch# show arp
Address         Age (min)  Hardware Addr   Interface
10.10.100.1           N/A  0000.0101.0000  Vlan100, Vxlan1

intercept-switch# show mac address-table
Mac Address Table
------------------------------------------------------------------

Vlan    Mac Address       Type        Ports      Moves   Last Move
----    -----------       ----        -----      -----   ---------
100    0000.0101.0000    DYNAMIC     Vx1        186     0:23:02 ago
```

Check if Direct Flow is enabled on all TOR switches

```
intercept-vtep1.09:35:02#sh directflow
DirectFlow configuration: Enabled
Total matched: 347 packets
Total programmed flows: 16 flows
intercept-vtep1.09:35:04#
```

Appendix 3A: Enabling MSS for Checkpoint firewall

Requirements

- Arista CloudVision running EOS release 4.23.2F or later
- Management Server Versions
- Version R80.30 with API version 1.5 (and above). In addition to this Management Server version, Check Point has provided a "hot fix" that provides a "Proxy API" ability which allows the user to access the Gateway APIs through a special URL on the Management Server. This hot-fix is required for MSS to work.
- Gateway Versions: Version R80.30 with API version 1.2 (and above)

MSS with Checkpoint is supported on the following Arista platforms

Platform	Minimum EOS Version required
Arista 7050X/7050X2/7060X/X2 Series	4.23.2F
Arista 7280R/7280R2 Series	4.23.2F

Configure the MSS service on CVX for Checkpoint Firewalls

In this step Checkpoint Management server is configured.

Command	Description
`service mss`	Enables the MSS service on CVX
`dynamic device-set checkpoint`	Creates a set of devices
`device <IP or Hostname>`	Defines a device. Note this is the hostname or IP address that MSS will use to communicate with Checkpoint Management server
`username admin password 0 <PW>`	Sets the username and password to access the device. Once entered the password is encrypted
`group Demo`	Defines the device group from the Checkpoint management server that will be used
`type check-point management-server`	Sets the type to checkpoint management server
`tag redirect <list of tags>`	Specifies the tag(s) that MSS looks for when reading security policy from the firewall or security manager for redirecting traffic. Any tag in the list can be used to tag the policy
`tag redirect <list of tags>`	Specifies the tag(s) that MSS looks for when reading security policy from the firewall or security manager for offloading rules. Any tag in the list can be used to tag the policy
`tag modifier verbatim <list of tags>`	The verbatim modifier changes the behaviour of MSS to install policies exactly as they are defined on the Firewall. Without the verbatim modifier when defining a policy for redirection between two subnets A and B, all traffic within the subnet A as well as within B is redirected. Any tag in the list can be used to tag the policy
`admin domain Demo`	Set the admin domain to 'Demo'
`state active`	Allows you set the device set as active or disabled

Check if Direct Flow is enabled on all TOR switches

```
!
cvx
   no shutdown
   !
   service mss
      no shutdown
      policy enforcement rules verbatim
      !
      dynamic device-set checkpoint
         device 172.28.134.208
            username admin password 7 7Y37xXCLlyi6OVpsSVIPqw==
            protocol https 4434
            group poc-mss
         state active
         type check-point management-server
         policy tag redirect MSS_redirect
         policy tag offload MSS_offload
         policy tag modifier verbatim MSS_verbatim
   !
   service vxlan
      no shutdown
```

Appendix 3B: Configuring Checkpoint firewall

Interface Configurations

Interfaces configured in aggregation groups bond1.100 and bond1.200. bond1.100 is part of the zone 'Tenant C' and bond1.200 is part of the zone 'Tenant D'

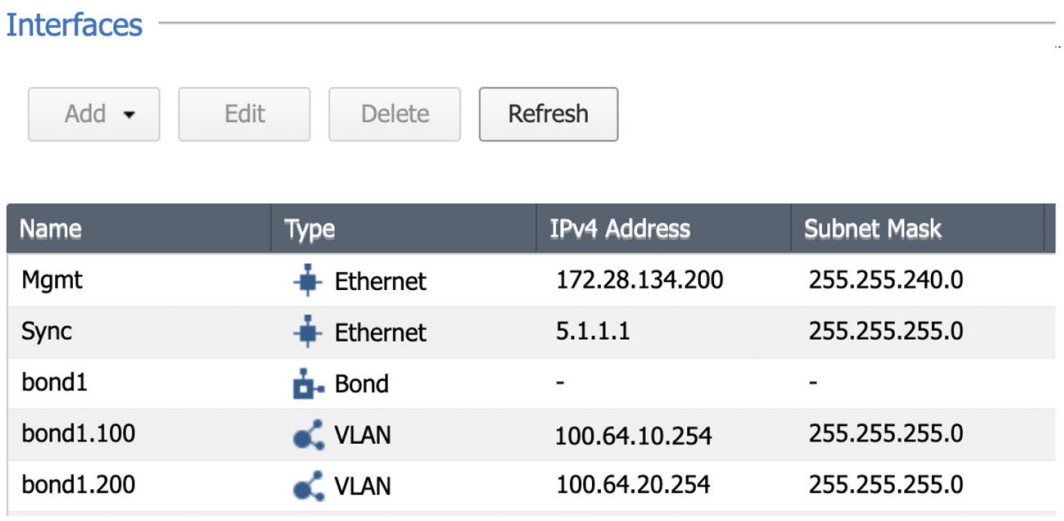

Figure 21: Interface Configuration for Checkpoint Firewall

Route Configuration

The firewall needs to have routes back to the original subnets in which the end hosts reside. Static routes have been created for each subnet. Default GW for the Tenant C subnets is the VLAN 10 interface on the TOR and for Tenant D subnets its the VLAN 20 interface on the TOR.

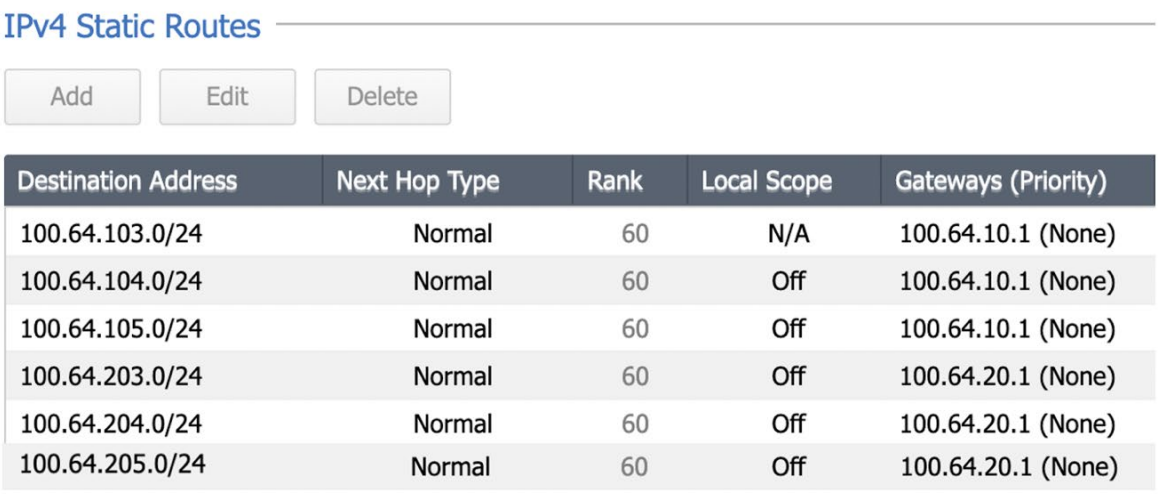

Figure 22: Route Configuration for Checkpoint Firewall

Policy Configuration

For Tenant C:

- Web servers can talk to each other directly only on SSL port 443, however unsecure HTTP port 80 traffic is denied even though the web servers are connected directly on the same VLAN and IP subnet.
- Web Servers can communicate with the app servers on port 8080 HTTP as well.
- Web Servers are not allowed to make any database queries over TCP port 3306 directly with the database servers.
- Only the application servers can make database queries over TCP port 3306 to the database servers.

For Tenant D:

- Similar requirements as mentioned above, except that the enforcement point for Tenant E workloads is on the intercept VTEPs.

These rules are defined in the Firewall and are enforced by the firewall or the Arista switches based on the Firewall policy tags.

No.	Name	Source	Destination	Services & Applications	Action	Install On	Comments
1	webc2webc-http	webc	webc	http	Drop	fwchk101	tags(MSS_redirect)
2	webc2webc-ssl	webc	webc	ssl_v3	Accept	fwchk101	tags(MSS_redirect)
3	webc2appc-http	webc	appc	HTTP_proxy	Accept	fwchk101	tags(MSS_redirect)
4	webc2dbc-mysql	webc	dbc	MySQL	Drop	fwchk101	tags(MSS_redirect)
5	appc2dbc-mysql	appc	dbc	MySQL	Accept	fwchk101	tags(MSS_redirect)
6	webd2webd-http	webd	webd	http	Drop	fwchk101	tags(MSS_offload)
7	webd2webd-ssl	webd	webd	ssl_v3	Accept	fwchk101	tags(MSS_offload)
8	webd2appd-http	webd	appd	HTTP_proxy	Accept	fwchk101	tags(MSS_offload)
9	webd2dbd-mysql	webd	dbd	MySQL	Drop	fwchk101	tags(MSS_offload)
10	appd2dbd-mysql	appd	dbd	MySQL	Accept	fwchk101	tags(MSS_offload)
11	Intrazone-TenantC	Tenant C	Tenant C	Any	Accept	fwchk101	
12	Intrazone-TenantD	Tenant D	Tenant D	Any	Accept	fwchk101	
13	Interzone-TenantC2TenantD	Tenant C	Tenant D	Any	Drop	fwchk101	
14	Interzone-TenantD2TenantC	Tenant D	Tenant C	Any	Drop	fwchk101	
15		Any	Any	Any	Drop	fwchk101	

Figure 23: Policy Configuration for Checkpoint Firewall

For Tenant C, firewall policy tags defined as "MSS-redirect" are enforced by the Active Firewall.

For Tenant D, Firewall tags defined as "MSS-offload" are enforced by the Arista switches locally without being redirected to the Active firewall.

Appendix 3C: Troubleshooting Checkpoint firewall

Check if CVX service is enabled

```
mss-l3-cvx#show cvx
CVX Server
  Status: Enabled
  UUID: c84d66ae-2d7f-11ea-8264-7f39ec151b3c
  Mode: Standalone
  Heartbeat interval: 20.0
  Heartbeat timeout: 60.0
  Client connection state preserving: Disabled
mss-l3-cvx#
```

Check if all TOR switches are registered for VCS on CVX

```
mss-l3-cvx#show cvx connections brief
          CVX Connection Table
Switch            Hostname        Status
----------------- --------------- ------
44:4c:a8:23:ba:5f intercept-vtep1 active
44:4c:a8:23:27:d1 intercept-vtep2 active
44:4c:a8:23:24:ed service-vtep    active


mss-l3-cvx#show cvx service vxlan
Vxlan
  Status: Enabled
  Supported versions: 2

  Switch            Status  Negotiated Version
  ----------------- ------- ------------------
  44:4c:a8:23:ba:5f Enabled 2
  44:4c:a8:23:24:ed Enabled 2
  44:4c:a8:23:27:d1 Enabled 2

mss-l3-cvx#show service vxlan switch all capability | grep Hos
  Hostname     : intercept-vtep1
  Hostname     : service-vtep
  Hostname     : intercept-vtep2
```

Check if all TOR switches are registered with MSS

```
mss-l3-cvx#show cvx service mss
Mss
  Status: Enabled
  Supported versions: 1

  Switch            Status  Negotiated Version
  ----------------- ------- ------------------
  44:4c:a8:23:ba:5f Enabled 1
  44:4c:a8:23:24:ed Enabled 1
  44:4c:a8:23:27:d1 Enabled 1
```

Check ARP learning via CVX

```
mss-l3-cvx#show service vxlan arp received
                        Received ARP Table
------------------------------------------------------------------------
Switch              VNI      IP Address      MAC Address        Changes
------------------- -------- --------------- ------------------ ---------
44-4c-a8-23-24-ed   10070    10.10.70.254    5849.3b1f.4412     0
44-4c-a8-23-24-ed   10080    10.10.80.254    5849.3b1f.4413     0
44-4c-a8-23-24-ed   10090    10.10.90.254    5849.3b1f.4413     0

Total IP addresses for this criterion: 3
```

Ensure the MSS service is enabled

The MSS service should be enabled on every CVX instance. To verify, run the following commands on CVX:

```
CVX-01#show service mss status
State: Enabled
<snip>
```

The state should be "Enabled", and reflect the configured VNI range.

Check if API connection is working

```
CVX-04#show service mss dynamic
Total policies processed: 0
Policy Source           Device Set Service Device State
----------------------- ---------- -------------- ----------
check-point-mgmt-server checkpoint 172.28.134.208 active
```

The state should be "active".

Policy is not fetched from the firewall correctly

The following command will list all the policies retrieved from the firewall by Arista MSS:

```
show service mss dynamic device-set <device-set name> device <firewall serial#> policies
```

The following command will list the Service Leaf port details:

```
show service mss dynamic device-set <device-set name> device <firewall serial#> neighbors
```

If no policies are seen by Arista MSS, check if CVX is able to communicate with the firewall or security manager using the following command:

```
cvx#show service mss dynamic status
Service Device Policy Monitoring Status:

Device: 172.28.134.200
  IP address: 172.28.134.200
  HA peer IP: 172.28.134.205
  Policy source type: CheckPointMgmtServer
  Accessed via Aggregation Manager: 172.28.134.208
  Device set name: checkpoint
  Device set state: Active
  Total policies processed: 112280
  Last seen at time: 2020 Apr 06, 07:32:51

Device: 172.28.134.205
  IP address: 172.28.134.205
  HA peer IP: 172.28.134.200
  Policy source type: CheckPointMgmtServer
  Accessed via Aggregation Manager: 172.28.134.208
  Device set name: checkpoint
  Device set state: Active
  Total policies processed: 5
  Last seen at time: 2020 Apr 06, 07:33:27

Device: 172.28.134.208
  IP address: 172.28.134.208
  Policy source type: CheckPointMgmtServer
  Aggregation Manager: True
  Device group member(s):
    172.28.134.205
    172.28.134.200
  Device set name: checkpoint
  Device set state: Active
  Last seen at time: 2020 Apr 06, 07:33:19
```

NOTE: The number of policies processed is an aggregate number and should increment every time MSS polls for new policies.

If only corresponding policy is missing, check the policy tag configured on the firewall policy.

IP-MAC binding not learned by CVX

Check the status of the policy to ensure that CVX has the necessary information to redirect traffic:

```
CVX-01#show service mss policy name policy1
               Macro-Segmentation L2 Policy Table
------------------------------------------------------------------------------
   Source       Device          Policy          Configured State Operational State
------------ ------------ ------------ ---------------------- -----------------

               Macro-Segmentation L3 Policy Table
------------------------------------------------------------------------------
        Source                  Device     Policy                Offload       Redirect Unconverged
                                                                                 status         status         IPs
---------------------- ------------------------ -------------------- ------------- -------------- -----------
   check-point-mgmt     172.28.134.208          policy1               Active       Active       0 of 4
```

If the policy status is not "active," check the ARP table information received by CVX.

```
cvx# show service vxlan arp received
Received ARP Table
------------------------------------------------------------------------------
Switch               VNI       IP Address       MAC Address         Changes
------------------- -------- --------------- ------------------ ----------
00-00-91-02-00-00    1000      10.10.10.102     00:00:01:02:00:00   0
00-1c-73-00-e2-16    2000      10.10.20.103     00:00:01:03:00:00   0
```

If the IP Address of the host is not seen in the CVX ARP table, ICMP ping a host which is not on the same subnet (or VNI) as the intercept host and verify ARP table information again. If ARP information for the host is learned by the CVX after the ping, check the status of the policy and ensure it's "active".

If the situation still persists, run the following commands on the intercept VTEP. If the host MAC address is learned on the VXLAN interface (Vxlan1), this indicates that there is a Layer 2 (L2) loop in the network. Resolve the loop and verify the policy status again.

```
intercept-switch# show arp
Address         Age (min)  Hardware Addr   Interface
10.10.100.1           N/A  0000.0101.0000  Vlan100, Vxlan1

intercept-switch# show mac address-table
Mac Address Table
------------------------------------------------------------------

Vlan    Mac Address       Type        Ports      Moves   Last Move
----    -----------       ----        -----      -----   ---------
100    0000.0101.0000    DYNAMIC     Vx1        186     0:23:02 ago
```

Check if Direct Flow is enabled on all TOR switches

```
intercept-vtep1.09:35:02#sh directflow

DirectFlow configuration: Enabled

Total matched: 347 packets

Total programmed flows: 16 flows

intercept-vtep1.09:35:04#
```

Appendix 4: Supported Deployment Models

As MSS processes traffic (to perform redirect to firewall, bypass firewall or drop actions) on ingress; the ingress (to the MSS deployment) TOR device needs to support MSS.

The following table summarises all supported configurations in different deployment models:

Deployment Model	Policy Enforcement Scheme	Tags on Firewall Policy
Both Compute and service TORs: DCS-7050X, DCS-7050X2, DCS-7050X3, DCS-7060X, DCS-7060X2	group, verbatim	• redirect • offload • redirect, verbatim • offload, verbatim
Both Compute and service TORs: DCS-7020R, DCS-7280R, DCS-7280R2, DCS-7500R, DCS-7500R2	verbatim	• redirect, verbatim • offload, verbatim
Both Compute and service TORs: DCS-7050X, DCS-7050X2, DCS-7050X3, DCS-7060X, DCS-7060X2, DCS-7020R, DCS-7280R, DCS-7280R2, DCS-7500R, DCS-7500R2	verbatim	• redirect, verbatim • offload, verbatim
DCS-7050X, DCS-7050X2, DCS-7050X3, DCS-7060X, DCS-7060X2 as compute TOR and other series as service TOR (with no intercepted hosts connected)	group, verbatim	• redirect • offload • redirect, verbatim • offload, verbatim

The below mentioned platforms are supported only in verbatim qualifier

- DCS-7020
- DCS-7500R
- DCS-7500R2
- DCS-7280R
- DCS-7280R2

To enable verbatim support, following should be configured:

```
cvx
   service mss
     policy enforcement rules group verbatim
```

Appendix 5: Understanding MSS tags

he following four firewall policies are tagged with four supported tag combinations:

ID	Policy Enforcement Scheme	Tags on Firewall Policy
1.	Source: 10.10.20.2 Destination: 10.10.20.3 Application: SSH Action: PERMIT	MSS_redirect, MSS_verbatim
2.	Source: 10.10.20.2 Destination: 10.10.20.3 Application: ANY Action: DENY	MSS_offload, MSS_verbatim
3.	Source: 10.10.20.1 Destination: 10.10.20.1 Application: ANY Action: PERMIT	MSS_offload
4	Source: 10.10.20.1 Destination: 10.10.20.3 Application: SSH Action: DENY	MSS_redirect

The following describes the traffic forwarding behavior after MSS programs the hardware for all of these policies. For the sake of simplicity, we do not describe the treatment that any traffic undergoes after hitting the firewall.

- **Policy1**: Only SSH traffic from 10.10.20.2 and 10.10.20.3 is redirected to the firewall. Note that tagging only this single rule allows MSS to redirect both forward and reverse directions traffic of this flow.
- **Policy2**: Non-SSH traffic from 10.10.20.2 and 10.10.20.3 is dropped at the TOR, and thus bypasses the firewall. Note that MSS considers the relative priority between the policies when configuring the switches. Thus, DirectFlow rules installed for Policy2 carry lower priority than the ones for Policy1.
- **Policy3**: Any traffic between 10.10.10.1 and 10.10.20.1 is allowed at the TOR and thus bypasses the firewall. In addition, 10.10.10.1 and 10.10.20.1 are treated as intercepted hosts for implicit redirect feature. Any remaining traffic involving any of these two IP addresses is redirected to the firewall. Examples of such traffic are: regular ping traffic between 10.10.10.1 and 10.10.20.3, SSH traffic between 10.10.20.1 and 10.10.20.2, etc.
- **Policy4**: Both 10.10.10.1 and 10.10.10.2 are treated as intercepted hosts. Thus, all traffic originating at or destined to these two IPs are forwarded to the firewall. For example, traffic from 10.10.10.1 to 10.10.20.3 is also redirected.

Summary and Conclusion

In summary, Arista is leveraging its experience with delivering "always-on" data center networks for the largest, most critical networks in the world, and bringing this innovation to the Enterprise with the UCN for Campus architecture.

The key elements of the Arista solution delivered within this design are:

- **Available Architecture:** architectural leadership, born in the cloud
- **Agile Workloads:** standards-based, highly scalable and resilient platform with Any Workload
- **Analytics:** real-time streaming network telemetry, unparalleled in the industry, providing network operators insight not previously seen in the campus
- **Automation:** fundamentals of Arista EOS and CloudVision, as well as the deep programmability of the Arista platforms, introduces meaningful automation to the operation of the campus network
- **AnyCloud API:** Being open and programmable with full API support across campus to data center to public cloud - all with a single binary image and common CloudVision platform

Santa Clara—Corporate Headquarters
5453 Great America Parkway,
Santa Clara, CA 95054

Phone: +1-408-547-5500
Fax: +1-408-538-8920
Email: info@arista.com

Ireland—International Headquarters
3130 Atlantic Avenue
Westpark Business Campus
Shannon, Co. Clare
Ireland

Vancouver—R&D Office
9200 Glenlyon Pkwy, Unit 300
Burnaby, British Columbia
Canada V5J 5J8

San Francisco—R&D and Sales Office 1390
Market Street, Suite 800
San Francisco, CA 94102

India—R&D Office
Global Tech Park, Tower A & B, 11th Floor
Marathahalli Outer Ring Road
Devarabeesanahalli Village, Varthur Hobli
Bangalore, India 560103

Singapore—APAC Administrative Office
9 Temasek Boulevard
#29-01, Suntec Tower Two
Singapore 038989

Nashua—R&D Office
10 Tara Boulevard
Nashua, NH 03062

Copyright © 2018 Arista Networks, Inc. All rights reserved. CloudVision, and EOS are registered trademarks and Arista Networks is a trademark of Arista Networks, Inc. All other company names are trademarks of their respective holders. Information in this document is subject to change without notice. Certain features may not yet be available. Arista Networks, Inc. assumes no responsibility for any errors that may appear in this document. June 13, 2019 07-0012-01