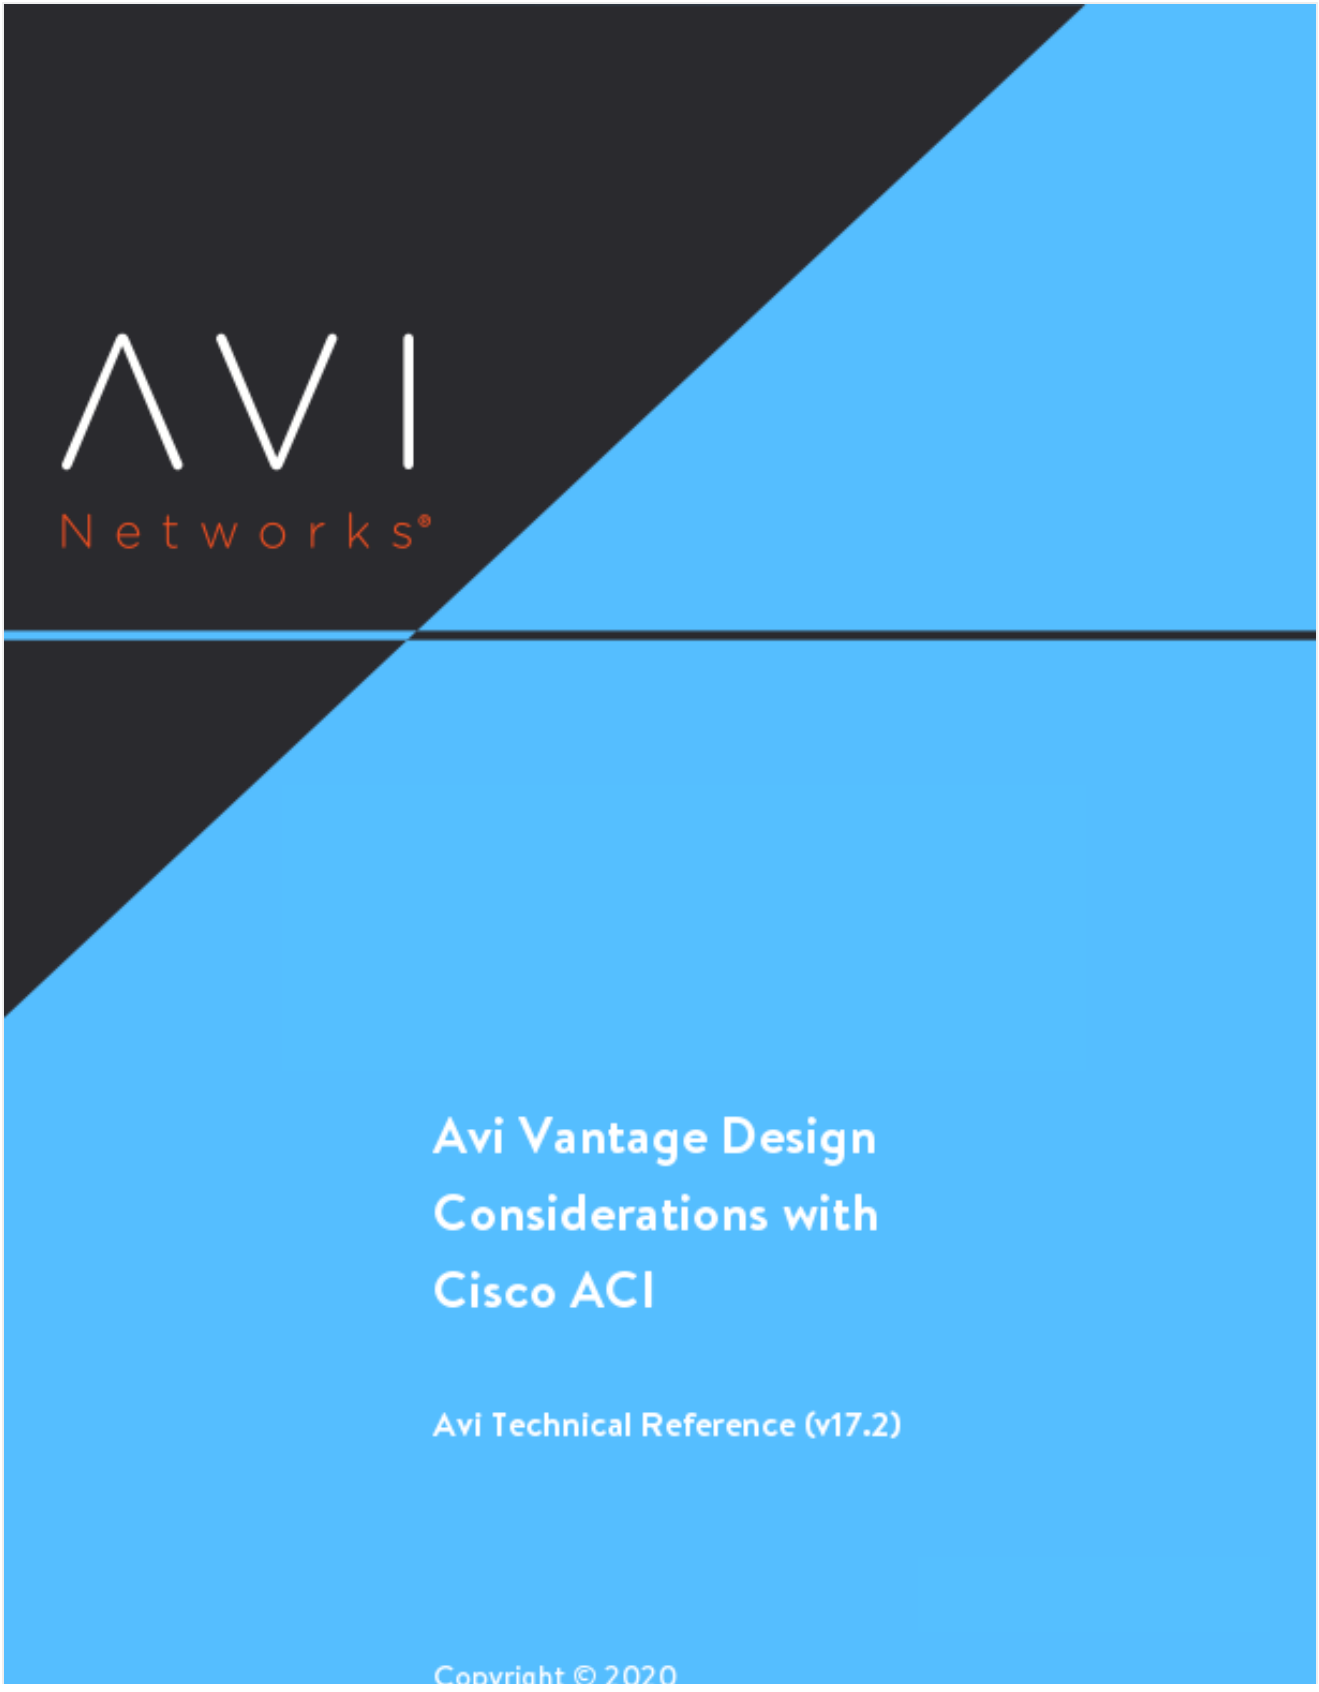

Avi Vantage Design Considerations with Cisco ACI

Avi Technical Reference (v17.2)

Avi Vantage Design Considerations with Cisco ACI

view online

Overview

Cisco Application Centric Infrastructure (ACI) is a software defined networking solution offered by Cisco for data centers and clouds which helps in increasing operational efficiency, delivering network automation, and improving security for any combination of on-premises data centers, private, and public clouds.

The Avi Vantage Platform provides enterprise-grade distributed ADC and iWAF (Intelligent Web Application Firewall) solutions for on-premises and public-cloud infrastructure. Avi Vantage also provides inbuilt analytics that enhances the end-user application experience as well as ease of operations for network administrators.

For complete information on Avi Vantage architecture, please refer to Avi Vantage Architectural Overview.

This document discusses options for deploying Avi Vantage within Cisco ACI in several host infrastructures such as, VMware, Cisco CSP, etc., along with the deployment best practices. This document does not discuss the steps for deployment. For complete deployment information, refer to the Cisco ACI with Avi Vantage Deployment Guide.

Intended Audience

This document is intended for virtualization and network architects seeking to deploy Cisco ACI along with Avi Vantage solution.

Note: A solid understanding and hands-on experience with Cisco ACI and Avi Vantage are the prerequisites to understand this design guide.

Avi Vantage deployment within ACI Integration

The recommended deployment model for Avi Vantage within Cisco ACI is referenced as Network Policy Mode. In this mode, Cisco ACI provides the network connectivity and contracts required by the EPGs that are used by the Avi Service engines to allow the traffic through.

Hosting Infrastructure

Avi Vantage can be hosted on VMware, Cisco CSP 2100 , bare-metal , public clouds, and several such platforms. The following are a few host infrastructures that support ACI fabric:

Avi Vantage on VMware with Write Access

VMware deployments where Avi Controller is configured with vCenter cloud connector. The Avi Controller has write access permissions to vCenter and handles the complete automation involved in creating Service Engines and placing them in the right network. The Controller also scales the Service Engines based on the configured threshold.

Refer to VMware write Access for more details.

Avi Vantage on VMware with Read/No Access

VMware deployments where Avi Controller has only read access or no access permission to the vCenter. In such deployments, Service Engines are manually deployed and the Avi Controller does not provide much automation.

Refer to VMware Read/No Access for more details.

Avi Vantage on Cisco CSP 2100

Cisco CSP 2100 is a NFV platform based on Intel x86 and the KVM hypervisor. Both the Avi Controller and Avi Service Engines can be deployed on Cisco CSP 2100.

Refer to Avi on Cisco CSP 2100 for more details.

Avi Vantage on OpenStack

Avi Vantage integrates with OpenStack infrastructure components to provide centralized automation, monitoring, and management of application discovery and delivery.

Refer to Avi on OpenStack for more details.

Integration Design Options

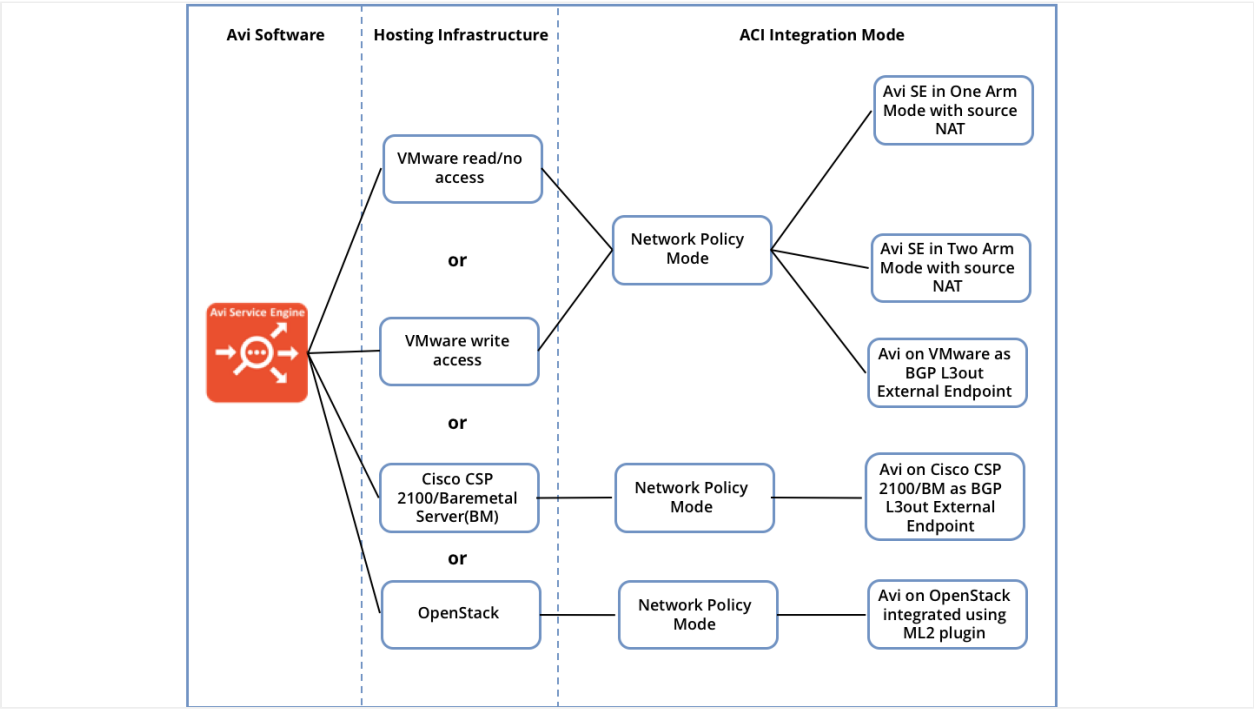

For Avi Vantage deployments within Cisco ACI, the design options are segregated among different hosting infrastructures. The following section summarizes the available hosting infrastructure and the associated integration options.

ACI Integration Mode along with Hosting Infrastructure

```
<th>Hosting Infrastructure</th>
<th>Integration Options</th>
<th>Brief Summary </th>
```

```
<td>VMware write access </td>
<td>Network policy mode</td>
<td><ul>
```

```
    <li>ACI provides network connectivity and contracts for access control</li>
    <li>Avi Vantage provides automated configuration and provisioning for its L4-L7 services</li>
    </ul>
    </td>
```

```
<td>VMware read access and VMware no access</td>
<td>Network policy mode</td>
<td><ul>
<li>Workload on any hypervisor(s), bare-metal server(s), and ACI manages reachability</li>
<li>Avi Vantage deployed in no-orchestrator mode (without automated VMware provisioning)
```

```
<li>Avi Service Engines peer with ACI fabric as a BGP Layer 3-out</li>
</ul></td>
```

```
<td>Cisco CSP 2100 and bare-metal servers</td>
<td>Network policy mode</td>
<td><ul>
    <li>Avi Vantage deployed on CSP 2100 or bare-metal server(s)</li>
    <li>Avi Service Engines peer with ACI fabric as a BGP Layer 3-out</li>
    </ul>
    </td>
```

Refer to the following documentation links for a detailed description on ACI integration options for each hosting infrastructure:

VMware Write Access

- Cisco ACI Network Policy Mode on Write Access VMware Cloud
- Cisco ACI Network Policy Mode on VMware Write Access Cloud as BGP L3 Out

VMware Read and No access * Cisco ACI Network Policy Mode on Read/No Access VMware Cloud

Design Considerations and Limitations

Below are a few basic recommendations and best practices applicable for all design options.

Each design option also has specific recommendations mentioned under respective links. Please refer to the links in the section above for more details.

Avi Controller Considerations

The Avi Controller is a single point of management and control for the Avi Vantage system, and is typically deployed as a redundant three-node cluster.

To allow control plane communication between the Avi Controller cluster and Service Engines, open the firewall ports mentioned in the table below.


```
<th>Traffic source</th>
<th>Traffic destination</th>
<th>Ports to allow</th>
```

```
<td>Avi Controller</td>
<td>Avi Controller</td>
<td>TCP 22 (SSH)<br>
```

```
<td>Avi Service Engine</td>
<td>Avi Controller</td>
<td>TCP 22 (SSH)<br>
```

TCP 443 (HTTPS) TCP 8443 (HTTPS) TCP 5098 (SSH) (if the Controller is a docker container, SSH is on port 5098)
TCP 8443 (HTTPS) UDP 123 (NTP) TCP 5098 (SSH) (if the Controller is a docker container, SSH is on port 5098)

Note: For VMware vCenter Controller-to-ESXi hosts allow port 443.

CPU and Memory Allocations

The CPU and memory sizing recommendations for the Avi Vantage Controller cluster are based on Service Engine and virtual service scale. Reference the following link for the most recent recommendations.

Avi Vantage Controller Sizing Recommendations

Avi Service Engine Considerations

Avi Service Engines handle all data plane operations within Avi Vantage by receiving and executing instructions from the Controller. The SEs perform load balancing and all client and server-facing network interactions.

For network policy mode, Service Engines can be hosted on the same infrastructure as that of the Controller, or a different infrastructure. The only requirement is to ensure connectivity between the Controller and Service Engines.

Avi Vantage Scale Out Considerations

The Avi Vantage scale out option is used to scale out the virtual services to multiple Service Engines or migrate to new Service Engines for better resource utilization. This scale out can be triggered either automatically based on different parameters like CPU, PPS or manually by using the scale out option under virtual service.

With ACI, the bridge domain with virtual service network will have endpoint learning enabled by default, by which the ACI will map the IP address to the MAC address. So, if multiple Service Engines host the same virtual service, then the ACI fabric will see the same virtual service IP address with multiple MAC addresses on different leaf ports, leading to auto flapping in the specific network.

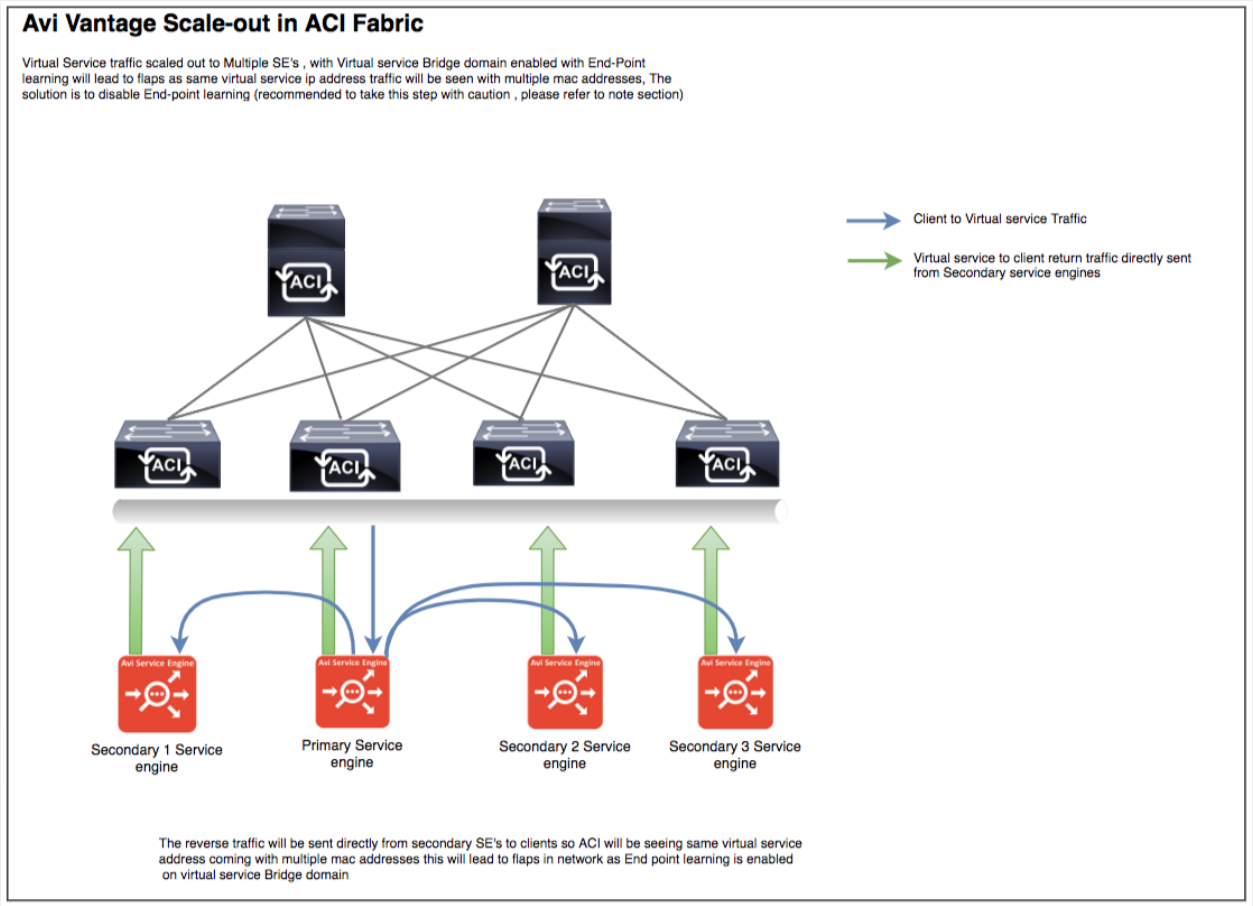

Starting with Avi Vantage release 17.2.10, there is a workaround available for this. The Service Engines can function with SE tunnel mode enabled, where the return traffic from the secondary SE will be sent to the primary SE. This avoids all IP address to MAC address conflicts, as the traffic flow will be from a single Service Engine.

The SE Tunnel Mode option can be disabled by referring to the following documentation: Autoscale Service Engines

Enabling SE Tunnel Mode is recommended only for proof of concepts in the network, where Avi Vantage is been tested in an existing ACI fabric network which has a shared bridge domain for virtual service networks. We recommend disabling this setting for production deployments and follow the section below to disable endpoint leaning on virtual service bridge domain.

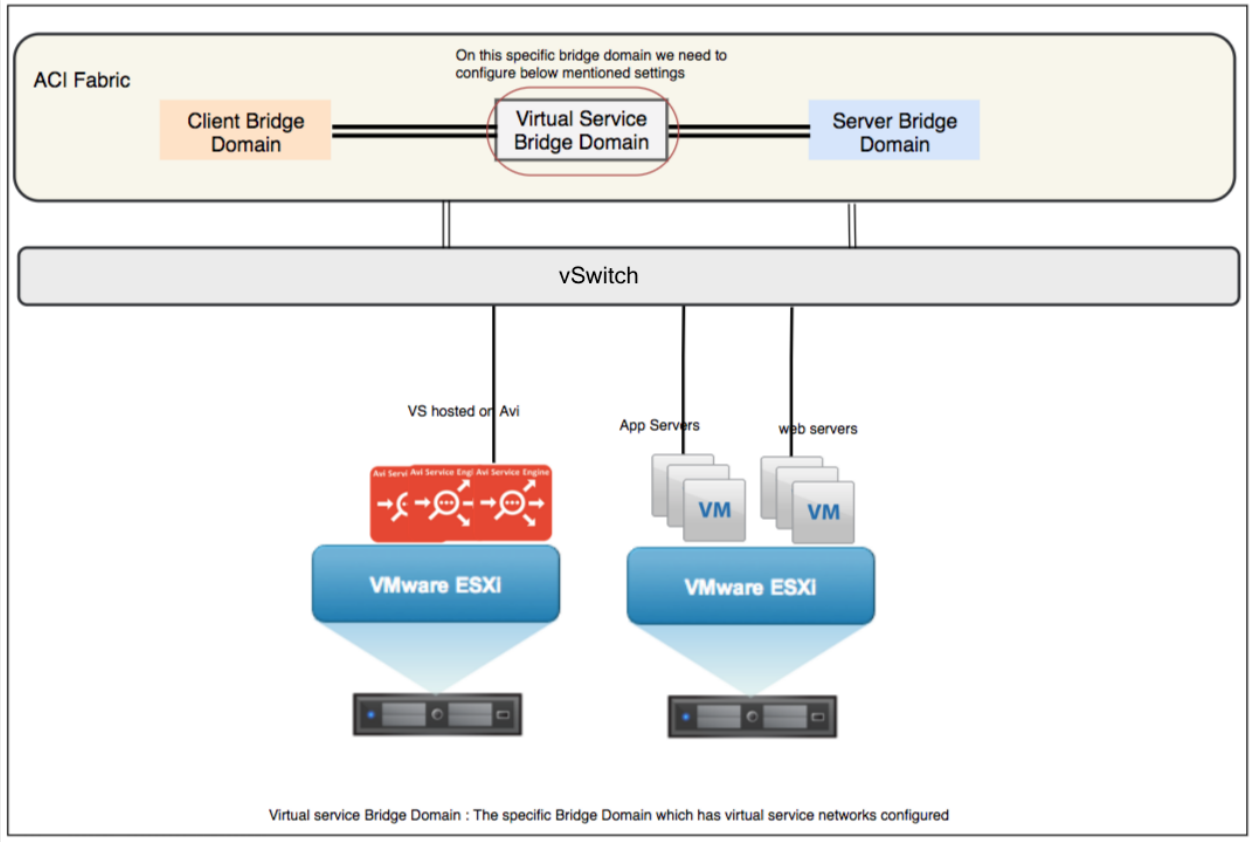

Configure the virtual service network bridge domain with the settings below, as it will allow the scaled out traffic for virtual service on multiple Service Engines to transit through the ACI fabric seamlessly.

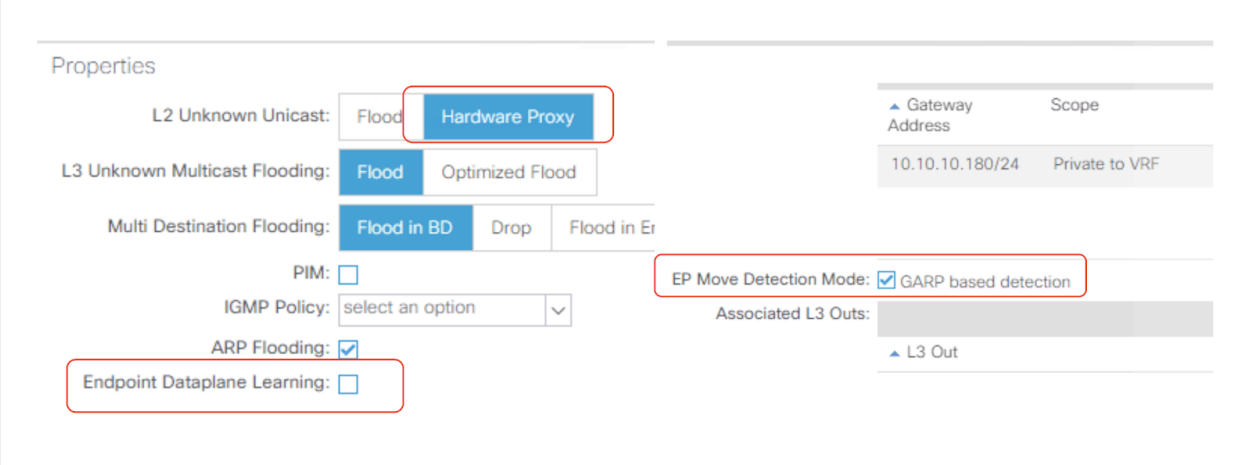

Note: Configure this setting while creating the virtual service network bridge domain. Modifying the settings for an existing bridge domain will retain some stale entries, leading to unwanted packet drops. For more information on Avi scale out, refer to Virtual Service Scaling.

Source NAT Considerations

By default, Avi Service Engines perform source NAT. You can disable this option on the Avi Controller, if needed. We recommend using source NAT if it is not required to preserve the client IP address. Disabling source NAT will disable

connection multiplexing. Refer to Connection Multiplexing for more information on connection multiplexing and its impact on other features.

Disabling source NAT on Avi Service Engines would have an impact on ACI as well, as the source IPs seen by the client network will be the same as those seen by Avi Service Engines. Enable Limit IP Learning to *subnet option* under the virtual service network bridge domain.

Also for deployments where source NAT is disabled, ensure that the default gateway on servers is pointed to Avi Service engines, so that the return traffic takes the path of the source traffic.

EPG Workload Considerations

For load balancing, it is recommended that the workload for each application be in its respective EPG. For instance, three tier applications with web, app, and database servers are recommended to have specific application servers in specific EPGs like web servers, app servers, etc. This ensures high level security with application communication and external routed clients-to-application communication.

Limitations

The following are a few limitations applicable to all design options: * Direct server return is not supported. Use NAT, if the clients and servers belong to the same network or same EPG. * Only GoTo deployment mode is supported. GoThrough, bridge, or transparent mode is not supported.