



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

Avi Technical Reference (v18.1)

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

[view online](#)

Overview

In this design option, Avi Vantage is deployed in VMware write access mode with the Service Engines configured as BGP L3 outs in APIC. The leaf or spine switches in ACI fabric will learn the routes and forward the virtual service traffic to Avi Service Engines.

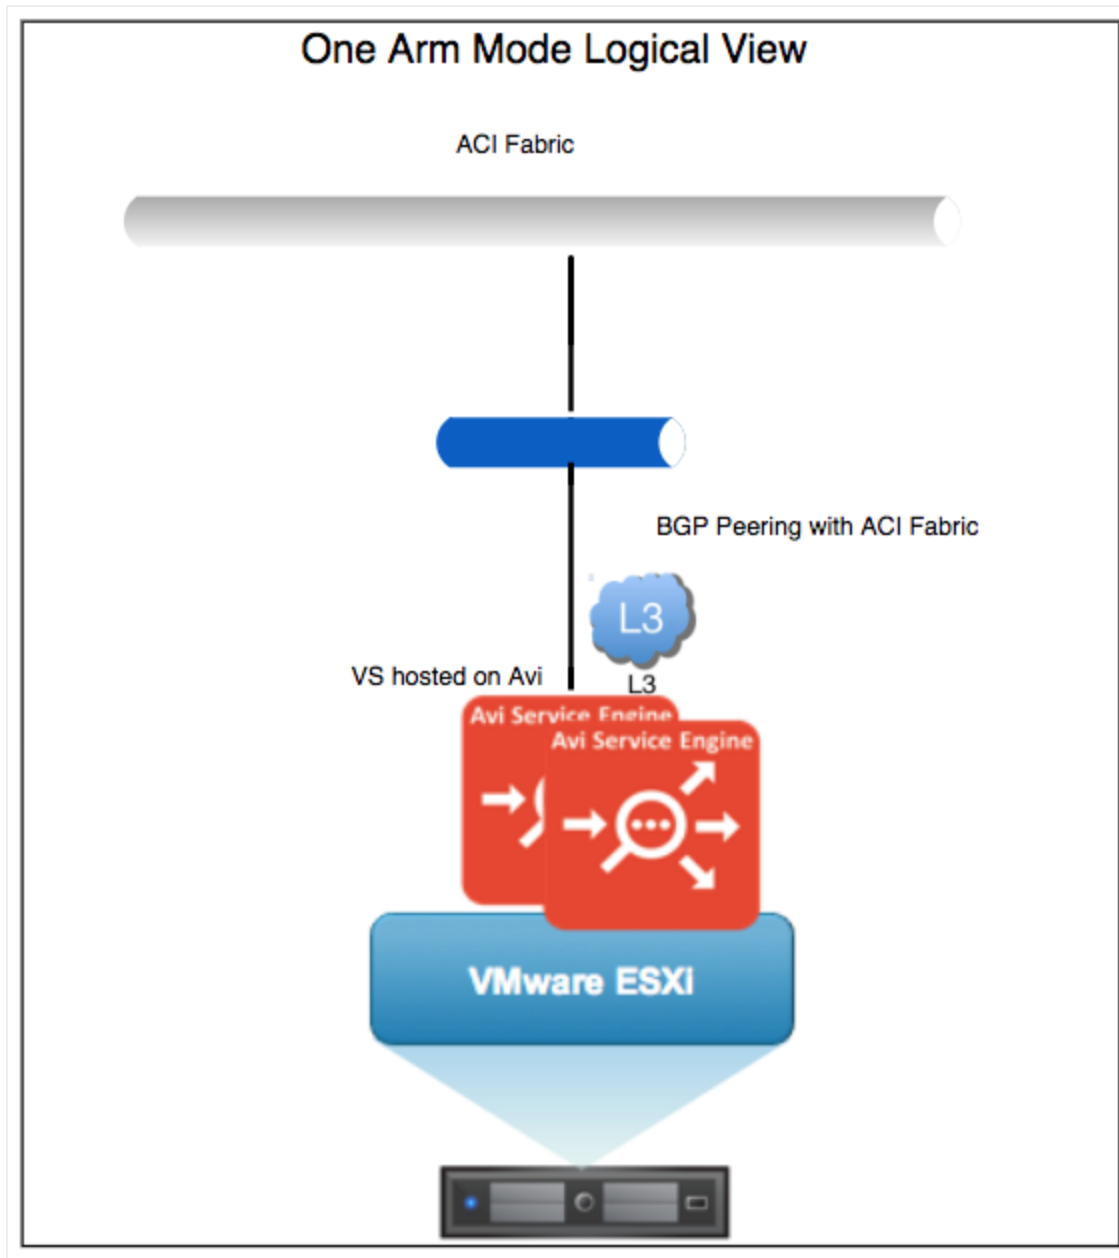
This option is best recommended for use cases where BGP scaleout for Service Engines is used instead of native L2 scaleout. This is particularly recommended for use cases with ECMP across the Service Engines.

For related information, refer to the following links: * [BGP Support for Scaling Virtual Services](#) * [Deploying Avi Vantage in Write Access Mode](#) * [Cisco ACI with Avi Vantage Deployment Guide](#)

Avi Service Engine in One Arm Mode

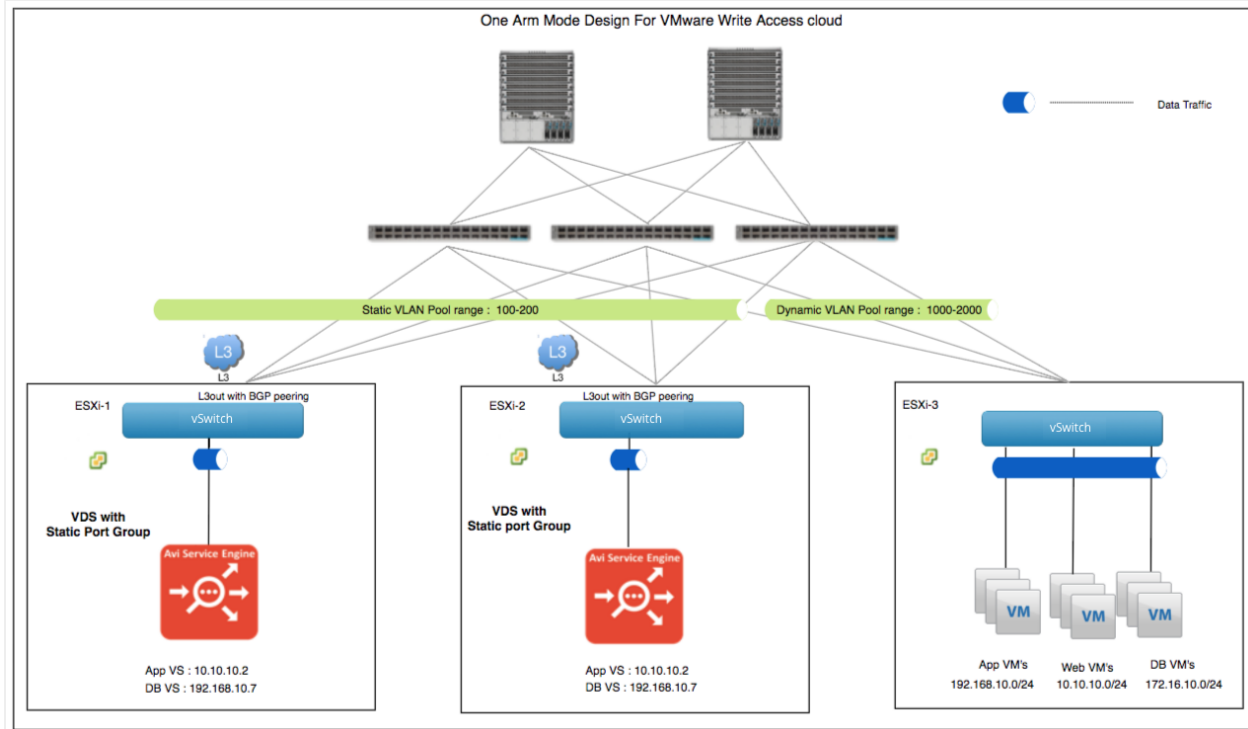
Logical Network Topology

The logical view of an one arm mode is as shown in the figure below.



The Service Engines are connected to a single port group on a virtual distributed switch and configured as BGP L3 out in APIC. The clients are sourced from ACI fabric and can access the virtual service hosted on the SE.

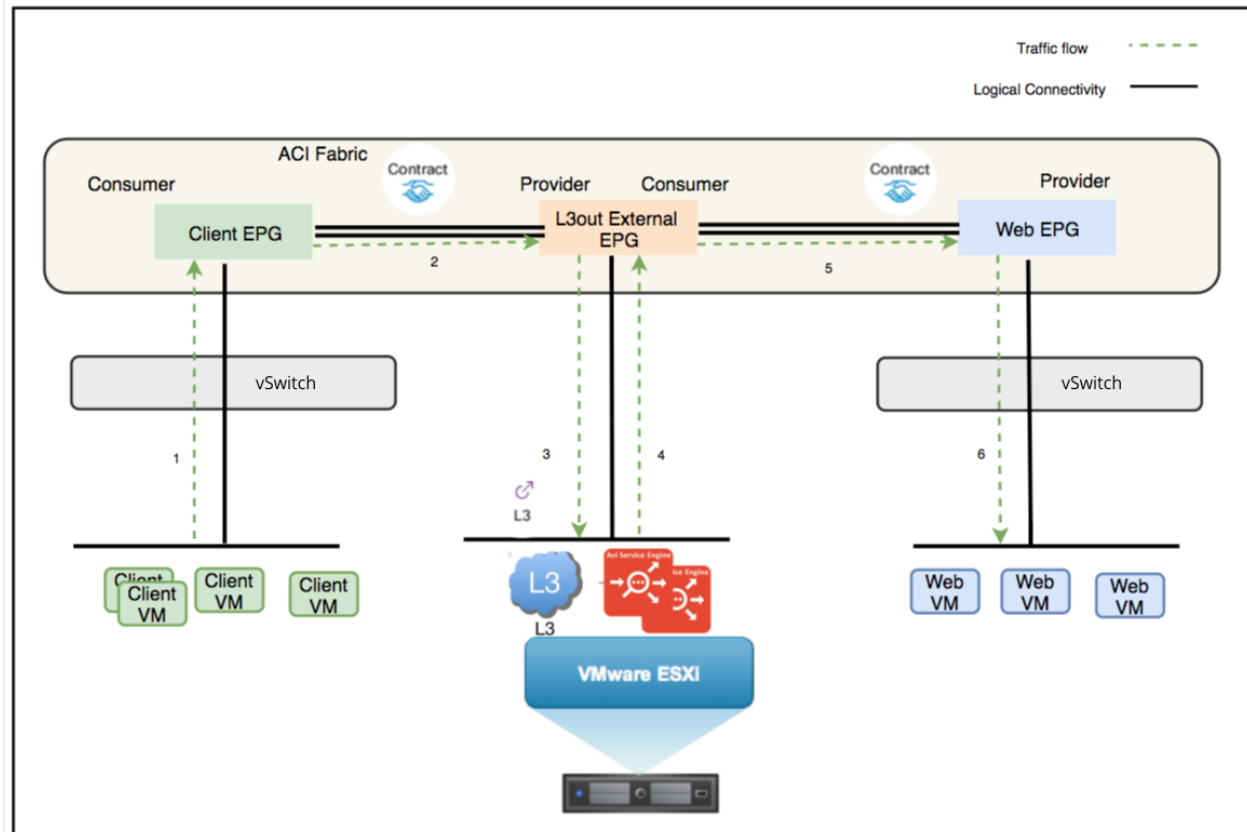
The following is the logical network topology of the one arm mode design with BGP peering for Avi Service Engines hosted on VMware write access mode.



Logical Traffic Flow

The traffic flow for client VM to access virtual service App hosted on Avi Service Engine is as shown below:

1. Client VM ? ACI Fabric ? Client EPG
2. Client EPG ? Contract ? L3out external EPG
3. L3out external EPG ? ACI fabric ? Avi SE
4. Avi SE ? load balancing to back-end servers ? ACI fabric ? L3out external EPG
5. L3out external EPG ? Contract ? Web EPG
6. Web EPG ? Web server VM
7. Return traffic follows the same path as the incoming traffic



Considerations

Virtual Distributed Switch

For Avia Service Engines hosted on VMware infrastructure with write access cloud configuration in an ACI fabric network, you can use the existing virtual distributed switch used by the payload VMs. Create a static port group on this switch to host the Service Engine data vNICs in this port group.

Else, create a new virtual distributed switch and use a port group in this switch for the Service Engines.

Service Engine Routing

Avia Service Engines can only publish the virtual service routes. The SEs can not learn any routes using BGP from any of the BGP peers.

For few cases SEs need to learn the routes from the BGP peer in order to send the return traffic to an appropriate next hop which forwards the traffic. In such cases, you can use the auto gateway option, to ensure that the return traffic is sent to the same MAC address from which it was received.

For more information, refer to [Auto Gateway](#).

High Availability

As BGP is used for exchanging routes, high availability is entirely dependent on BGP. By default, the Service Engines are in active/active state. For active/standby virtual service deployment, use the local preference option in ACI fabric. This allows you to choose one SE route as the most preferred one over other SEs.

For more information on local preference in ACI fabric, refer to [Cisco APIC Layer 3 Networking Configuration Guide](#).