TECHNICAL GUIDE:
April 2024

VMware® Avi™ Load Balancer

Integration with NSX-T Cloud

vmware®
by Broadcom

Table of Contents

About this Document

This technical document describes the process to integrate VMware Avi Load Balancer with an NSX-T Cloud and the associated automation work-flows. It can also be used as a design guide for NSX-T integration with the Avi, with recommendations on best practices as applicable.

Solution Overview

- VMware NSX-T provides an agile software-defined infrastructure to build cloud-native application environments.
- NSX-T is focused on providing networking, security, automation, and operational simplicity for emerging application frameworks and architectures that have heterogeneous endpoint environments and technology stacks. NSX-T supports cloud-native applications, bare metal workloads, multi-hypervisor environments, public clouds, and multiple clouds.
- The solution comprises of the VMware Avi Load Balancer Controller which uses APIs to interface with the NSX-T manager and vCenter to discover the infrastructure. It also manages the lifecycle and network configuration of Service Engines (SE). Avi Controller provides the control plane and management console for users to configure the load balancing for their applications and the Service Engine provide a distributed and elastic load balancing fabric.

Applicability

- NSX-T versions 3.0 and above
- vCenter 6.7, 7.0 and 8.0
- VMware Avi Load Balancer – 20.1.3 and above

Prerequisites

- The integration requires the Avi Controller to authenticate with the NSX-T manager and the vCenter server(s).
- The user accounts configured on Avi Controller requires the following roles and permissions mentioned in the KB link for the integration to work successfully:
  https://docs.vmware.com/en/VMware-NSX-Advanced-Load-Balancer/22.1/Installation_Guide/GUID-4F8337D8-A21A-4C41-B1F4-370369CD2366.html

System Limits

This section provides the recommended configuration limits for Avi integration with NSX-T. While configuring, ensure that the mentioned limits are maintained.

- Starting with Avi 22.1.2, the maximum number of NSX-T Clouds per Avi Cluster is now 32. With versions prior to 22.1.2, that limit was 5.
- Starting with Avi 22.1.1, the maximum number of Tier-1 Routers supported and qualified across all NSX-T Clouds per Avi Cluster is 300, while on a single NSX-T Cloud is 128. With versions prior to 22.1.1, that limit was 50.
- The maximum number of Virtual Services supported per Tier-1 is 1000. This value is limited by "Routes per Distributed Router" as mentioned in NSX limits here: https://configmax.esp.vmware.com/guest?vmwareproduct=VMware%20NSX&release=NSX%204.0.0&categories=18-33
- The maximum number of Virtual Services supported and qualified across all Tier-1 Routers per a large sized Avi Cluster is 2000.

Configure NSX-T Cloud

Create User Credential Configuration

To create an NSX-T Cloud, we first need to create User Credential configuration Objects. Follow the steps below, to create the User Credentials for NSX-T Manager and vCenter Server:

1. Navigate to Administration > User Credentials
2. Click **Create**.
3. In the New User Credentials Pop-up window, provide the following parameters:
    a. **Name**: <Credential Object Name>
    b. **Credentials** Type: NSX-T
    c. **Username**: <NSX-T Manager Username>
    d. **Password:** <NSX-T Manager User Password>

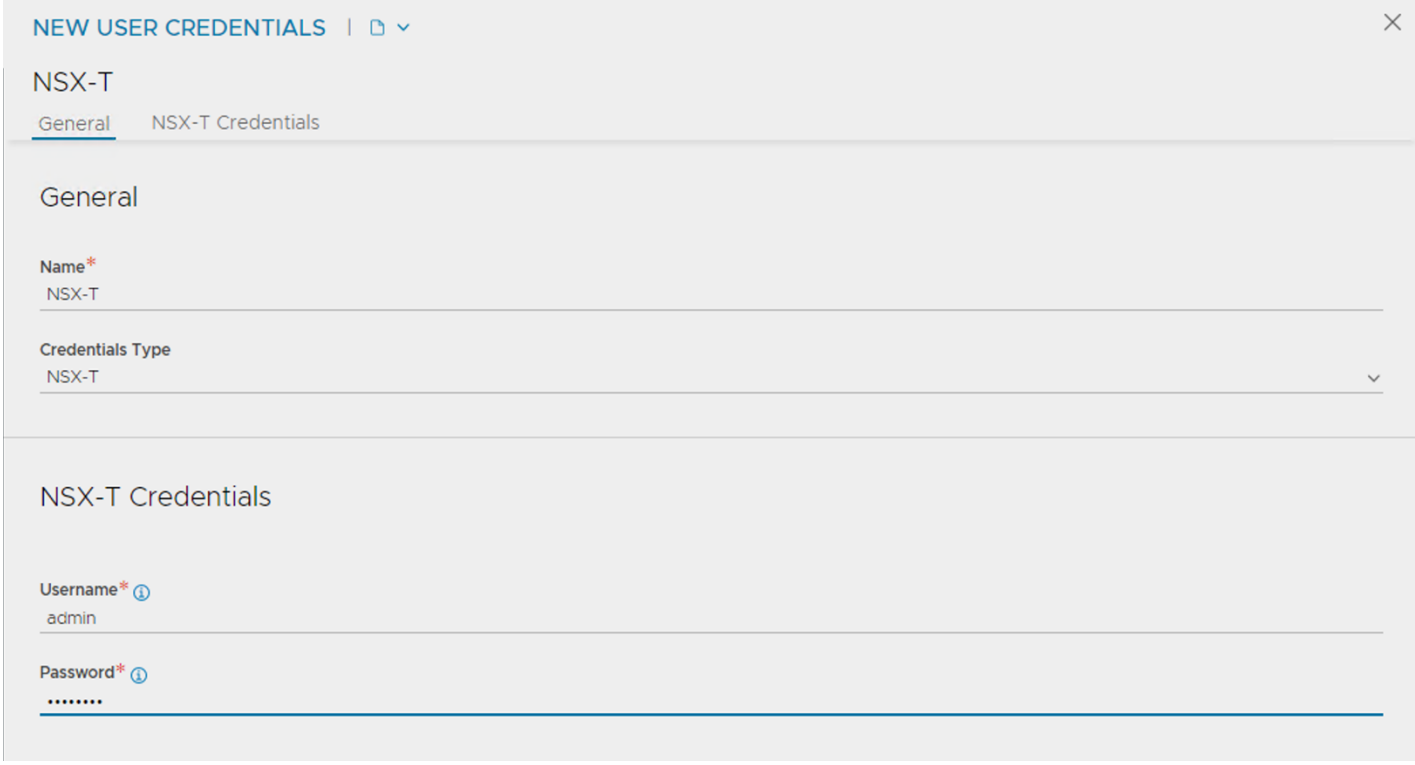

4. Click **Save**.
5. Repeat the steps above for the corresponding vCenter Server Credentials.

Create Cloud Configuration

To create an NSX-T cloud, log into the Avi Controller and complete the following steps:

1. Navigate to Infrastructure > Clouds.
2. Click on Create and select NSX-T Cloud.

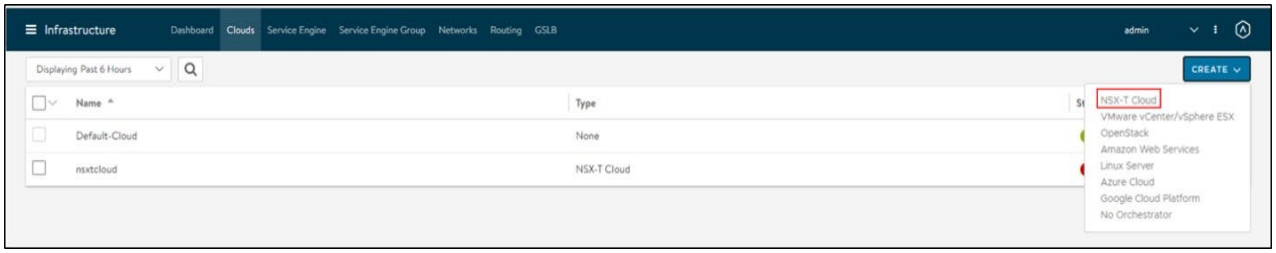

3. Enter the Name of the NSX-T cloud.
   Note: NSX-T Cloud is selected as the Cloud Type by default.
4. Check the DHCP option if SE management segment has DHCP enabled.

5. Enter a Prefix string for Objects created in the Avi and NSX environments.
   Note: The prefix string must only have letters, numbers, and underscore. This field cannot be changed once the cloud is configured.
6. Enter the NSX-T manager hostname or IP address as the **NSX-T Manager Address** and select the NSX-T Manager Credentials.
7. Click on **Connect** to authenticate with the NSX-T manager.

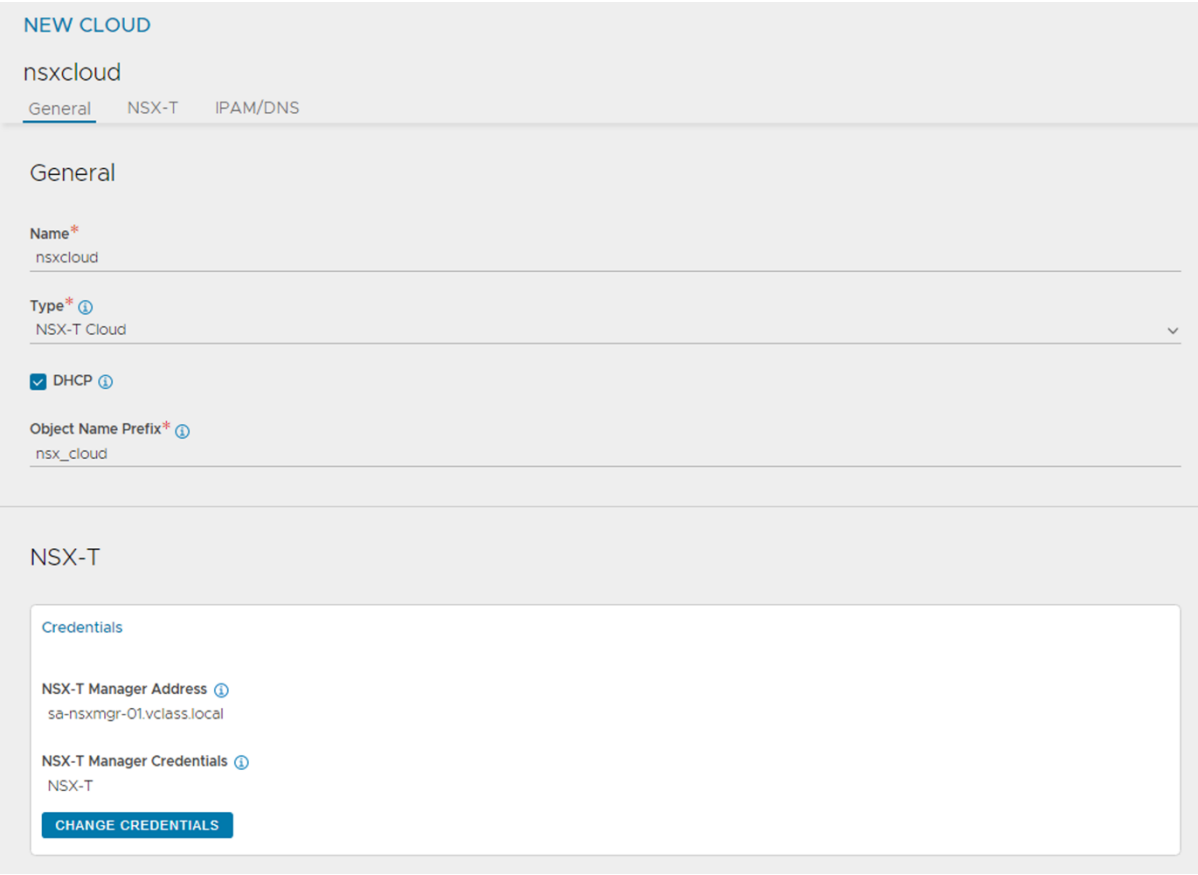

8. In the **Management Network** section, select the Transport Zone, and either the Tier1 Logical Router ID and Overlay Segment ID or the VLAN Segment ID.
9. Select the Transport Zone for VIP placement.
10. In the **Data Networks** section, select the Transport Zone.
    a. For Overlay Segments, click on Add to select each Tier1 Logical Router ID and Overlay Segment ID pair.
       Note: Only a single Segment ID for each T1 can be selected.
    b. For VLAN Segments, select all of the desired Segments from the drop down list.

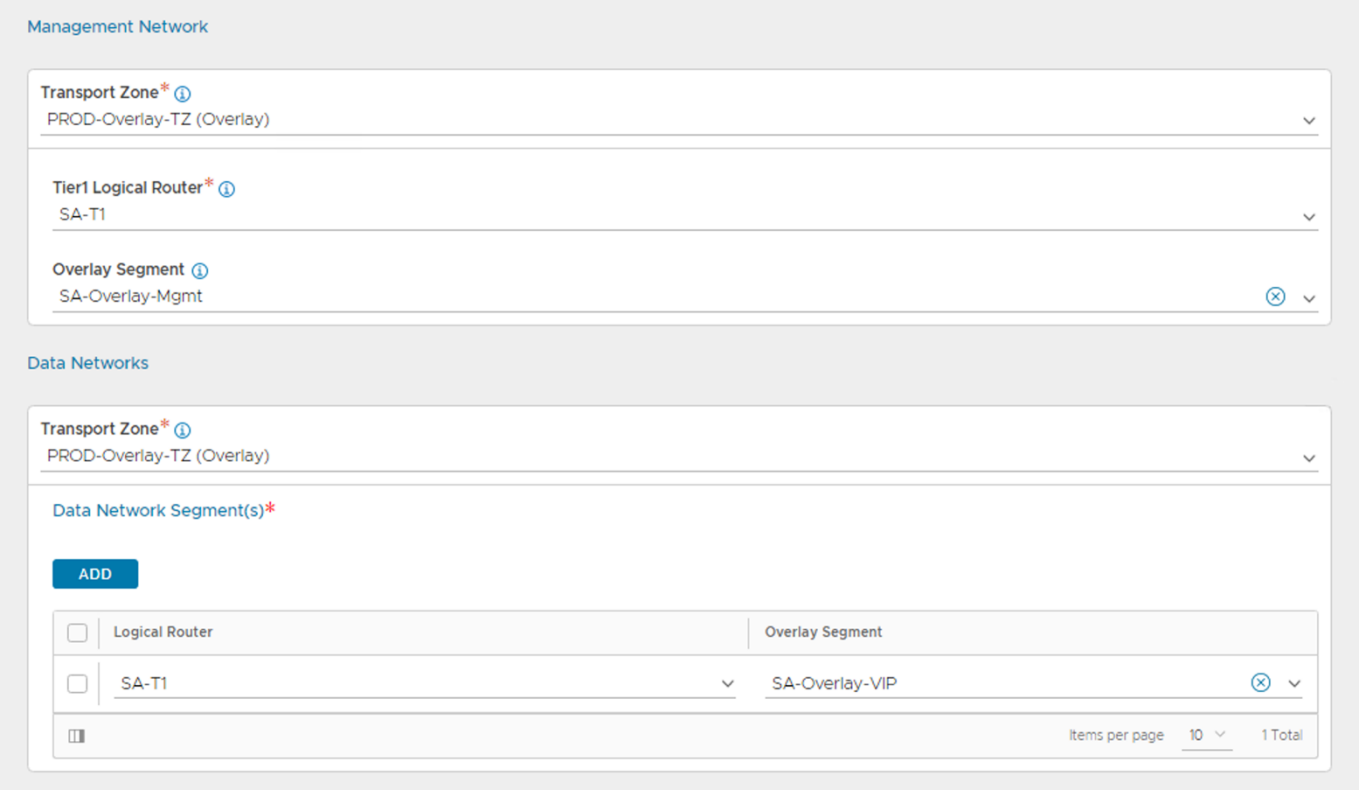

11. Under **vCenter Servers**, click on **Add**.
12. Provide a unique name for the vCenter Connection Object.
13. Select the vCenter Server from the dropdown, and configure the credentials.
14. Click on **Connect**.
15. Select the desired Content Library to store the Service Engine (SE) image.
16. Click **Done**.

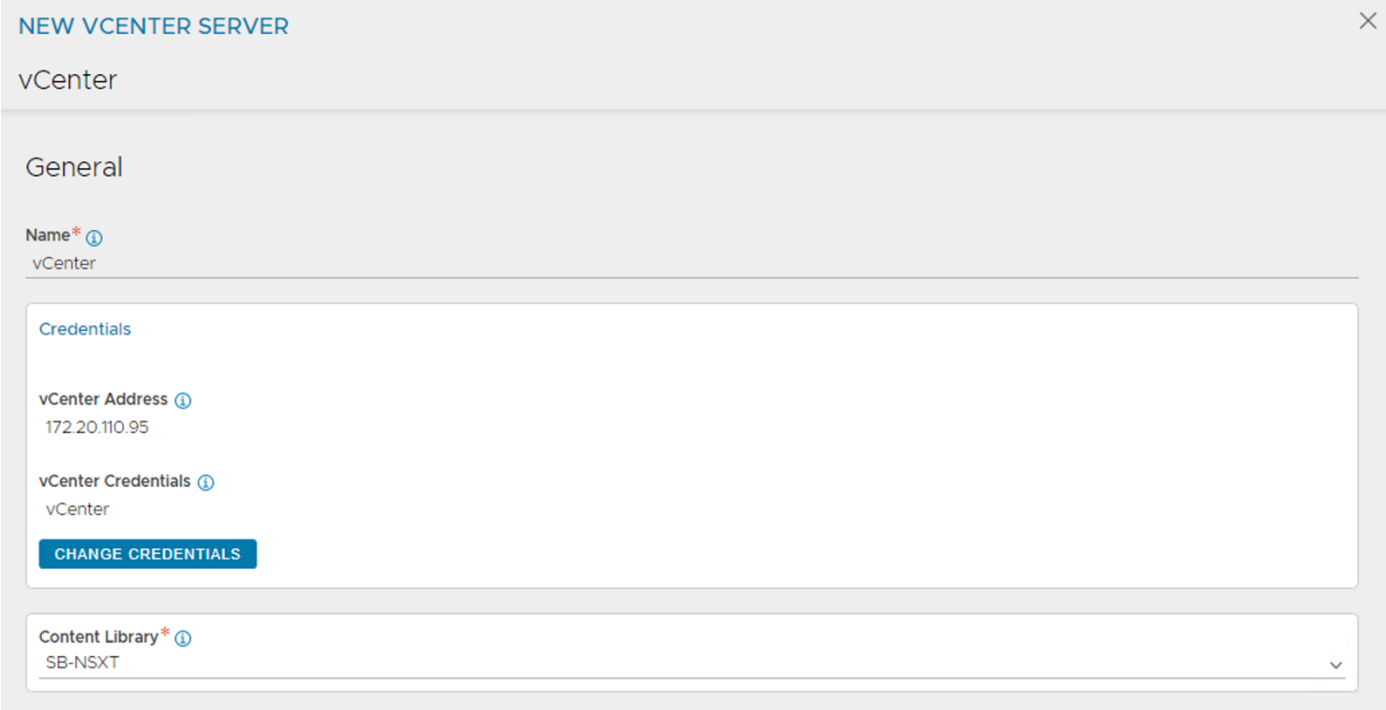

17. Select the IPAM/DNS Profile, as required.

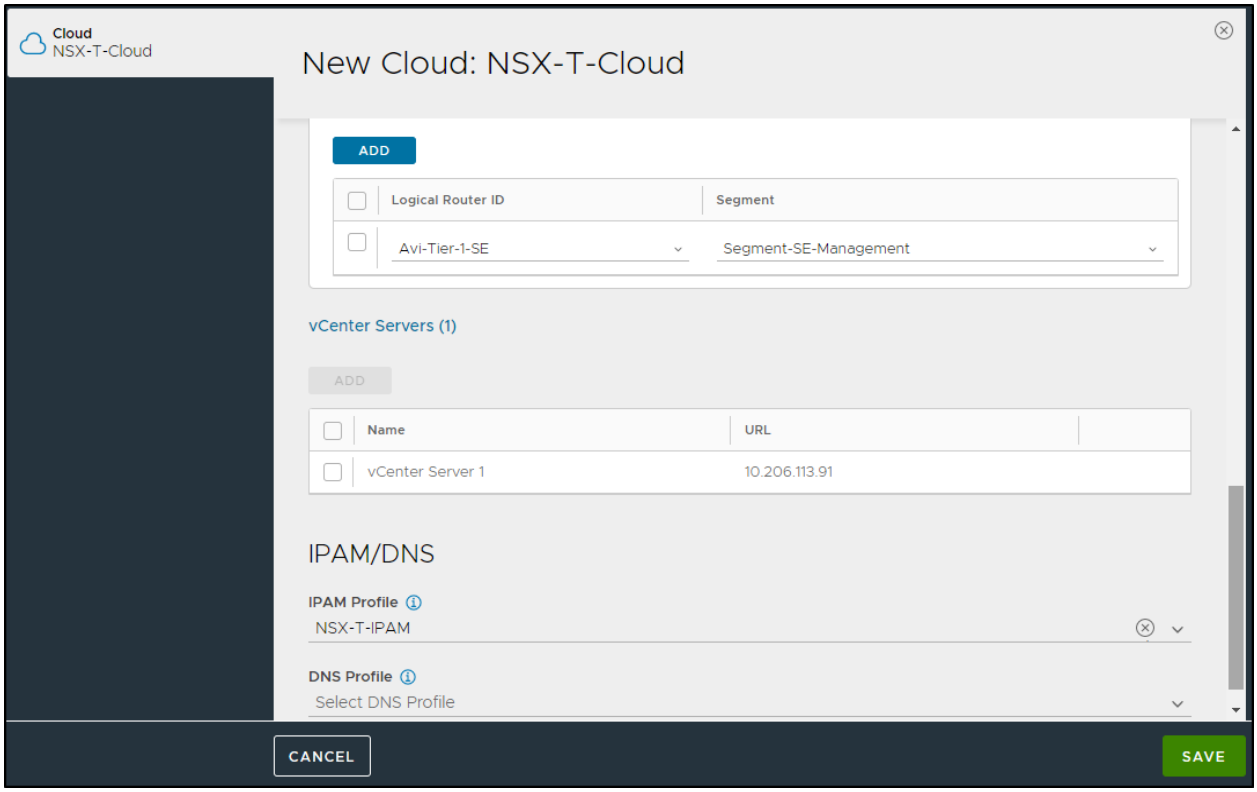

18. Click on **Save** to create the NSX-T cloud.

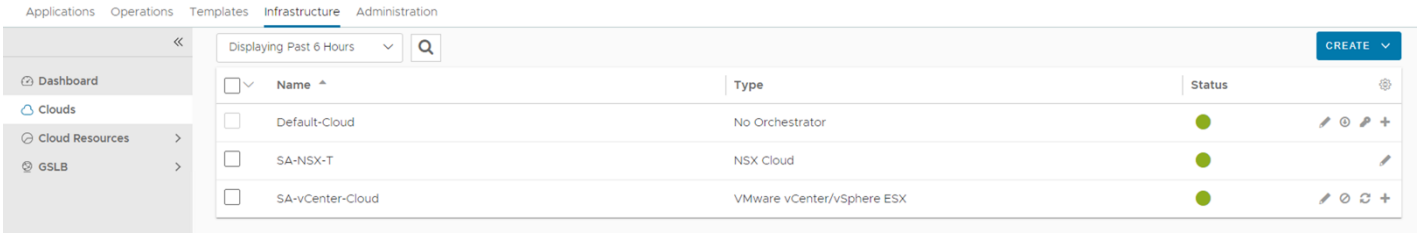

Cloud Configuration flow:

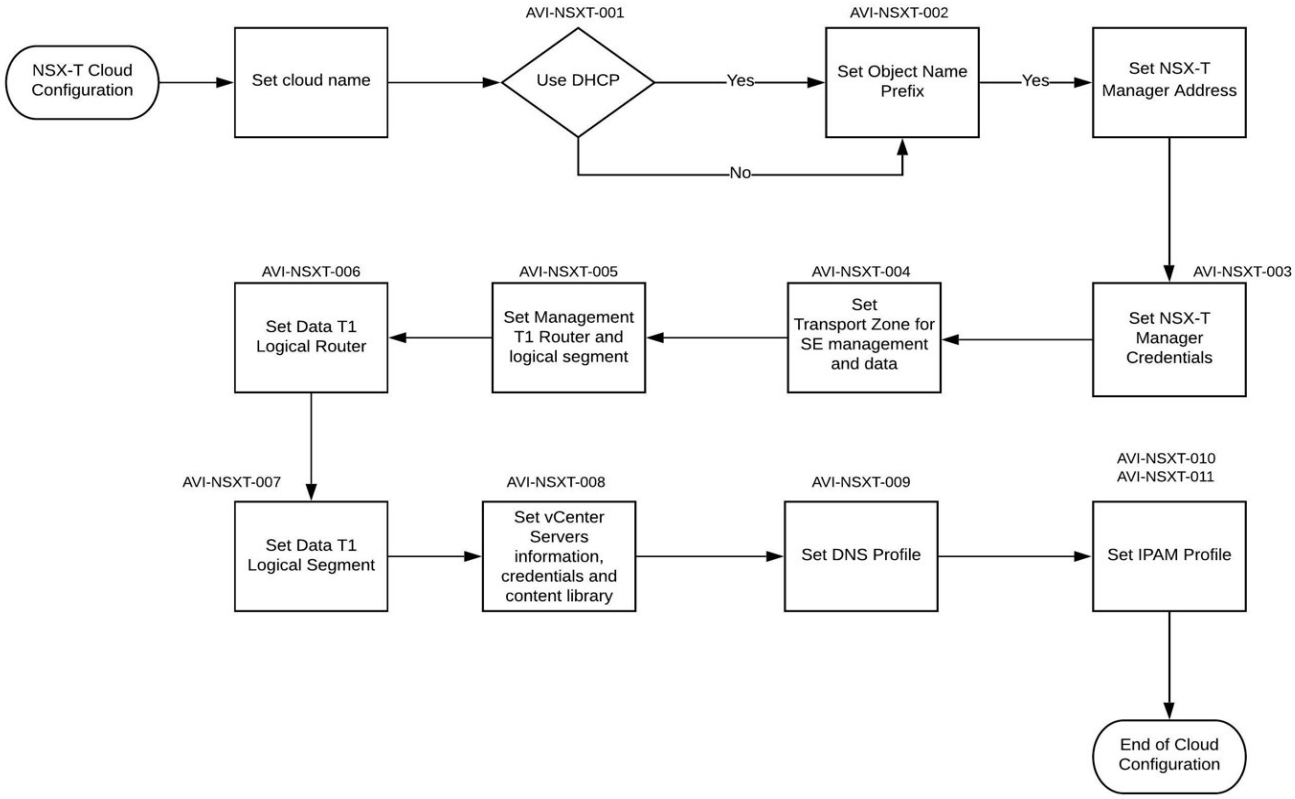

SE Group Configuration Flow:

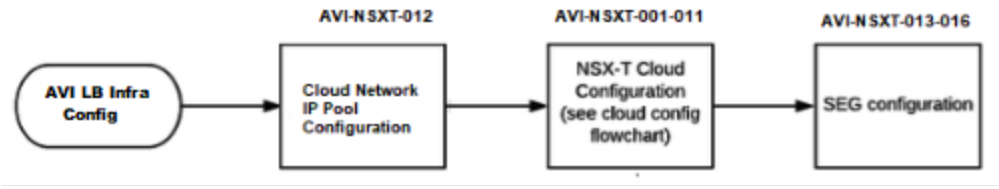

Virtual Service Configuration flow:

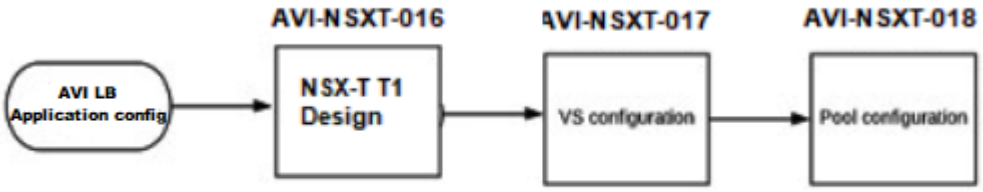

Cloud Configuration:

Decision ID	Design Decision	Design Justification	Design Implication
AVI-NSXT-001	Use DHCP option for SE management segment	Simplify the configuration of SEs automatically created by the Avi controller	Relying on the segment DHCP allocation simplifies the SE configuration and speeds up its deployment. Having a range of IPs also facilitates the DFW rule configuration required for SE to controller communication. External DHCP Server can provide Gateway information to the Service Engine. When DHCP Option is not available for SE interface refer to Decision ID - AVI-NSXT-012
AVI-NSXT-002	Choosing an Object Name Prefix	Use a meaningful name in order to identify the objects created on NSX-T.	NSGroups mapping SEs and Controller information use this prefix, therefore it will simplify the management of DFW rules and troubleshooting in general.
AVI-NSXT-003	Create or use an NSX-T Manager User/Role with the required privileges	This is required by the Avi controller and interact with the NSX-t Manager	Without the proper role assignment Avi controller won't be able to pull information from NSX-T and publish routes required for clients to reach Avi services. It is recommended to configure a remote user using VIDM or LDAP on NSX-T with Network
AVI-NSXT-005	Select a T1 router and segment for SE management	Segment must exist on NSX-T, Avi will use the segment to place the SE VM management NIC.	Dedicating a T1 router for SE management is optimal. SE management segment must have connectivity to the Avi controller on port 443, DFW rules need to be in place to allow this communication to happen.

AVI-NSXT-006	Select T1 router and segment(s) for Data Networks	Segments must exist on NSX-T, Avi will use the segment to place the SE VM data NICs.	Select multiple T1 router/Segment entries depending on your requirements. The SE needs a data interface on a segment belonging to the T1 router where it can reach backend servers. IP Addresses for the interface can be obtained using DHCP/Static Pool on controller. Avi is only supported on a one arm mode of deployment, meaning Client to VIP traffic and SE to backend server traffic both use the same SE data interface. If we are using a Static Pool for SE Data interface IP, and the Backend servers are in a different subnet than the data interface, we need to configure a static route on the SE to reach the backend servers. This is because the static Pool does not provide gateway details. There will be a VRF created per each T1 router added to the data networks, in order to provide isolation among different data interfaces and account for logical segments connected to different T1s to have the same subnet.
AVI-NSXT-007	Selection of a dedicated segment for SE data.	Using a dedicated segment for VIP/data placement simplifies the IP management and the configuration of DFW rules to allow clients to access load balancer services, as well as the communication between the SE and backend servers.	By not having a dedicated segment for VIP/data it is advised to reserve a block of static IPs on the segment for VIP allocation.

AVI-NSXT-008	Prepare vCenter requirements	Avi Controller nodes needs connectivity to vCenter(s) on port 443. Deployment and configuration of SEs happens through this integration. vCenter credentials need to be added for a user with the required level of access. Ref Decision ID: AVI-NSXT-009 A content library must be created beforehand, the controller uploads the SE OVA to it for later deployment when load balancer services are created.	vCenter objects must be configured on Avi for all vCenter compute managers added to NSX-T, in order to guarantee the successful creation and configuration of SEs.
AVI-NSXT-009	Create or use an vCenter User/Role with the required privileges	This is required by the Avi Controller nodes to interact with the vCenter	Create the following roles: AviRole-Global AviRole-Folder AviRole- Global: This role must apply Global Permissions. It allows the user to upload SE OVF to the content library, allocate space on datastore to create a virtual machine (VM) and assign networks to it. AviRole-Folder: This role must be applied to the folder where the admin wants the Avi service engine VMs to be created. It contains the permissions to create an SE folder, create SE VM from template, assign it to a resource pool, and perform operations on the VM like adding devices, powering it on/off, and connecting its vNICs to networks. This role restricts the VM operations only to the folder to which the role is applied. Ref: https://docs.vmware.com/en/VMware-NSX-Advanced-Load-Balancer/22.1/Installation_Guide/GUID-B68460EC-41B7-442E-ADCA-35A203A1027C.html
AVI-NSXT-010	Use of a DNS profile	By selecting a DNS profile every VS will be associated to a name, and a DNS record will be automatically created either on Avi or on a third-party integration.	Simplify load balancer services access by leveraging existing DNS infrastructure.

AVI-NSXT-011	Selection of an IPAM profile	By selecting an IPAM profile you simplify the allocation of VIPs for the VSs created on the NSX-T cloud. Otherwise, you must manually specify the VIP and select the T1 router during the VS creation.	By not using an IPAM profile there is a risk of selecting an IP already in use by another VS on the segment, especially if the segment is shared with other workloads and not dedicated to VIP/Data allocation.

Avi config (non-cloud):

Decision ID	Design Decision	Design Justification	Design Implication
AVI-NSXT-012	Configure Cloud IP Pool	If the MGMT and Data/VIP Networks are not utilizing DHCP, then IP Pools are required to automate the IP selection.	The network is created on the context of the NSX-T cloud and the VRF routing (mapping the different T1s selected for data segments). The network object is created automatically by the Avi Controller, but the subnet/mask value needs to be added manually to match the segment configuration on NSX-T. It is required to configure a static range for the IPAM profile to consume IPs. That range should be out of the DHCP segment scope if DHCP is enabled to avoid duplicate IPs.
AVI-NSXT-013	Configuration of IPAM profile	Configure a Network on Avi and associate it to an IPAM profile to simplify the allocation of VIPs for VSs created on NSX-T clouds.	The network on Avi can be used to allocate IPs for SE data NICs, VIPs or both. To have higher granularity and better control of the traffic (more specific DFW rules) you can either create two static ranges, one for SE NICs and one for VIPs, or one static range for VIPs only and let the SE NICs get the IPs from the segment's DHCP service.
AVI-NSXT-014	Selection of a Service Engine Folder under SEG config	Using this advanced setting on the Avi SEG configuration allows for organization of the SE VMs in vCenter.	As SEs are automatically created by the NSX-T cloud integration it helps to keep them organized and grouped on vCenter. The VM folder must be previously created on vCenter.
AVI-NSXT-015	Scope Datastore and Host, under SEG config	Using this advanced setting on the Avi SEG configuration allows for allocation of SEs on specific hosts or/and Datastore on the vCenter infrastructure.	As SEs are automatically created by the NSX-T cloud integration it helps to keep them organized or following infrastructure deployment rules by predetermining their deployment location in vCenter.
AVI-NSXT-016	Importing a license to Avi Controller	In order for VSs to be placed on SEs successfully, a valid license is required.	Any NSX-T client entitled for load balancer services can import licenses to Avi. However, it is required to change the Default License Tier to basic Edition. The Avi Controllers can only run on one license mode, either Enterprise or Basic. If Avi licenses have been purchased, import them on the controller and use the Default License mode.

Avi Controller Deployment and Management Connectivity

Avi Controller cluster VMs should be deployed adjacent to the NSX-T Manager, connected to the management port group. It is recommended to have a dedicated tier-1 gateway and logical segment for the Avi SE management.

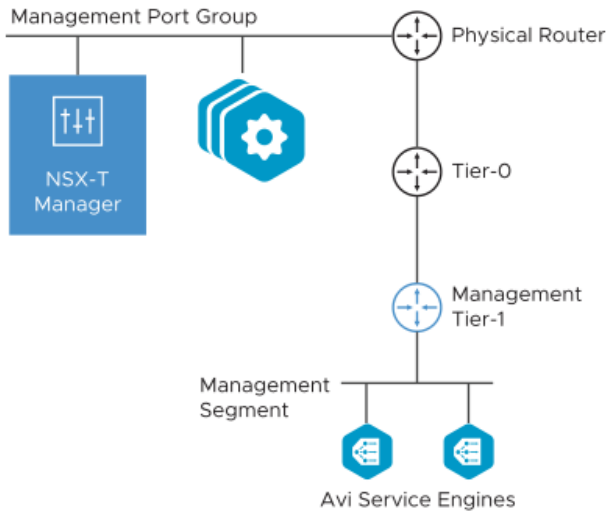

The network interface 1 of the SE VM is connected to the management segment. You should configure the management Tier 1 to redistribute the connected subnet routes to the Tier 0. Tier 0 must advertise the VIP to external peer using BGP.

Load Balancer Topologies

Avi supports only “one arm mode” deployments in NSX-T environments i.e. for a virtual service, the Client to VIP traffic and SE to backend server traffic both use the same SE data interface. An SE VM has nine data interfaces so it can connect to multiple logical segments but each one will be in a different VRF and hence will be isolated from all other interfaces.

The following are the recommended deployment modes for Avi on top of NSX-T managed infrastructure:

One Arm Mode with overlay VIP Segment - Single Tier 1

The diagrammatic representation of a typical Avi deployment on a simple NSX-T environment with all server segments connected to a single Tier 1 router is as follows:

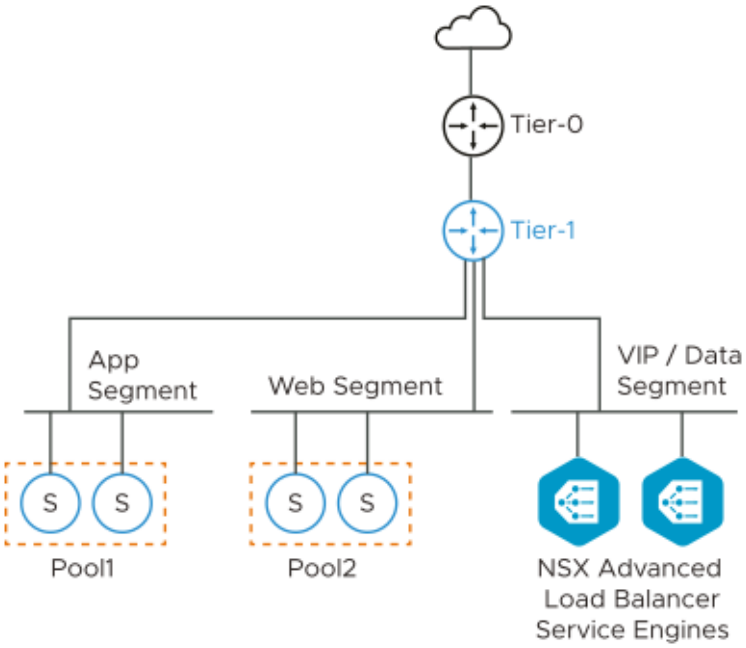

You need to create the VIP/data segment manually. The network adapter 1 of the Service Engine VM is reserved for management connectivity. You can connect only one of the remaining nine data interfaces (network adapter 2-10) of the Service Engine VM to the VIP/data segment. The rest of the interfaces must be left disconnected. The Service Engines are deployed in one-arm mode; i.e. the same interface is used for client and backend server traffic. The SE routes to backend servers through the Tier 1 router.

You can allocate the VIPs from the same subnet as that of the VIP/data interface of the Service Engine. You should reserve a range of static IP addresses to be used as VIPs in the subnet assigned to the VIP/data segment. Having a dedicated segment for VIPs makes managing the IP ranges independent of other subnets easier.

Optionally, you can place the data interface of the Service Engine on one of the server segments if you do not need a separate VIP segment. The subnet assigned to the server segment must have enough free static IP addresses to be used as VIPs.

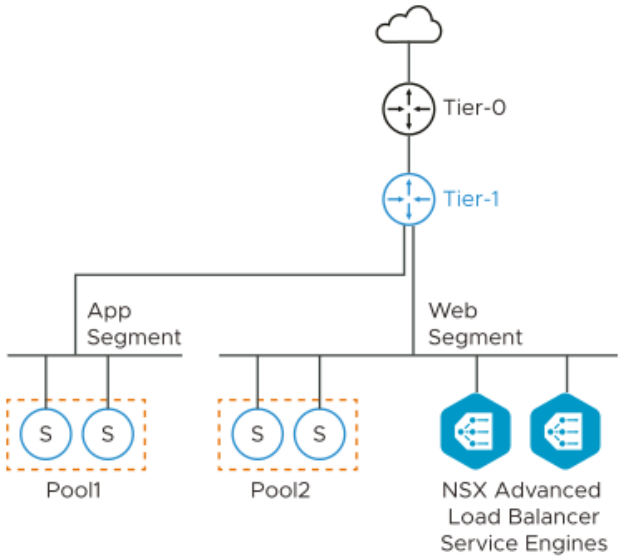

In both cases the SE to pool server traffic is handled by the Tier 1 DR (distributed router) which is present on all ESX transport nodes. Hence, there is no significant performance difference between using dedicated or shared segment for VIP/data.

Decision ID	Design Decision	Design Justification	Design Implication
AVI-NSXT-017	Deploying Avi on a multitenant environment	Dedicate a T1 router per tenant.	Avi automatically creates a VRF context for every tier-1 router selected during VIP network configuration. The logical Segments connected to different tier-1s can have the same subnet. VRF avoids issues when different tenants use the same subnet ranges.
AVI-NSXT-018	Selection of T1 router field when creating a Virtual Service.	Select the same T1 router that has the backend servers segment attached to it.	The Service Engine is deployed as one-arm mode, and its data interface is connected to the T1 specified on this configuration step, by selecting the same T1 as the backend servers avoids traffic going up to a T0 router and keeps all traffic east west within the same T1.
AVI-NSXT-019	Selection of T1 router field when creating a Pool.	As pointed out on the T1 selection for the Virtual Service, select the same T1 router that has been selected for the SE data interface.	If the VIP and the pool are connected to different T1, the traffic may pass through the T0 and hence through the NSX-T edge.

One Arm Mode with Overlay VIP Segment - Multiple Tier 1

In NSX-T environments, where web servers of different applications are connected to their individual Tier 1 routers, you need to create a VIP/data segment on each Tier 1.

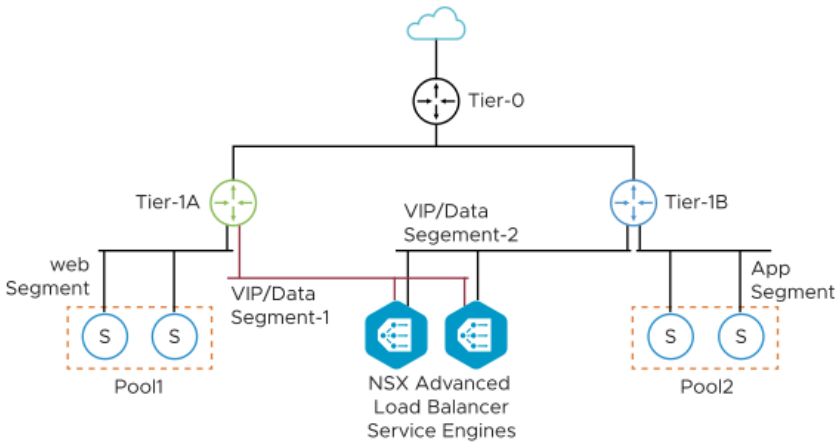

The network adapter 1 of the Service Engine VM is reserved for management connectivity. One data interface (network adapter 2) is connected to VIP/data segment-1, and one data interface (network adapter 3) is connected to VIP/Data segment-2. The rest of the interfaces are kept disconnected.

You can allocate the VIPs from the same subnet as that of the VIP/data interface of the SE. You should reserve a range of static IP addresses to be used as VIPs in the subnet assigned to each VIP/data segment. Having a dedicated segment for VIPs makes managing the IP ranges independent of other server subnets easier. You need to configure a separate VRF on Avi for each Tier 1 and add the data interfaces to the VRF corresponding to the Tier 1 segment it is connected to. For instance, in the above diagram, you should configure VRF-A and VRF-B for Tier 1 A and Tier 1 B. Also, you should add the SE interface connected to VIP/data segment-1 to VRF-A and the interface connected to VIP/data segment-2 to VRF-B. While creating the virtual services for a pool, you should choose the corresponding VRF.

For instance, select VRF-A while creating a virtual service for Pool1 and VRF-B while creating a virtual service for Pool2. This way, the VIP of the virtual service managing Pool1 will be on VIP/data segment-1, and the VIP of the virtual service managing Pool2 will be on VIP/data segment-2. This is required since you can route the SE to pool server traffic through Tier 1 DR and do not have to hairpin to Tier 0, and also because you can also configure logical segments on different Tier 1 to have the same subnet. Hence each Tier 1 traffic must be contained in its own VRF.

Optionally, if a separate VIP segment is not required, you can place the data interfaces of the Service Engine on one of the server segments. The subnet assigned to the server segment must have enough free static IP addresses to be used as VIPs.

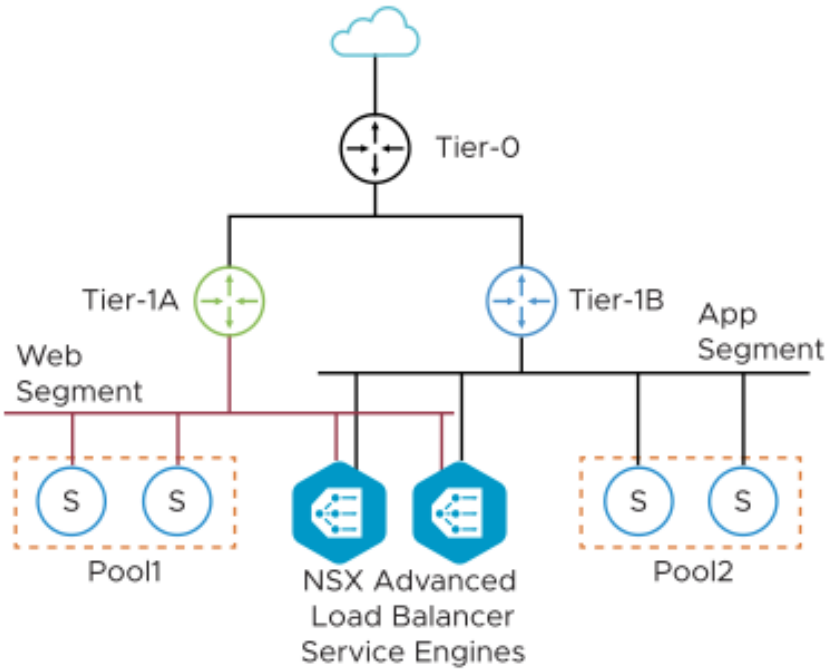

One Arm Mode with VLAN VIP Segment

Avi SEs can also be deployed in cases where the NSX-T logical segment is of VLAN type.

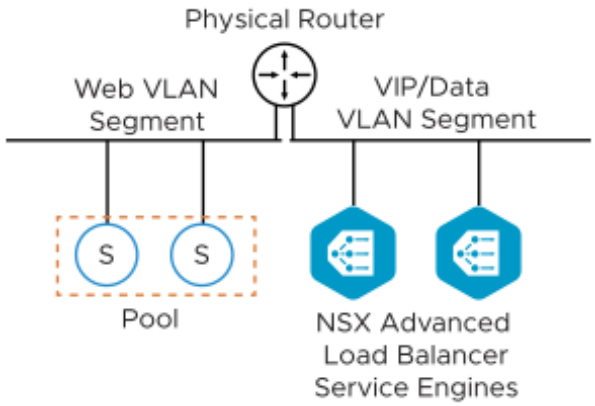

You need to create the VIP/data segment manually. The network adapter 1 of the Service Engine VM is reserved for management connectivity. You can connect only one of the remaining nine data interfaces (network adapter 2-10) of the SE VM to the VIP/data segment. Rest of the interfaces must be left disconnected. The SEs are deployed in one-arm mode; i.e. the same interface is used for client and backend server traffic. The SE routes to backend servers through the external physical router. In this mode, you can either allocate the VIPs from the same subnet as that of the VIP/data interface of the SE or allocate the VIP from a completely different subnet (not used anywhere else in the network). In the first case, you need to reserve a range of static IP addresses in the subnet assigned to each VIP/data segment to be used as VIPs.

In the second case, you can configure SEs to peer with the physical router by configuring BGP on Avi (provided the router supports BGP). The SEs will advertise VIP routes with the SE data interface IP as the next hop. The router does ECMP across the SEs if the VIP is scaled out. You can allocate the VIPs from the same subnet as that of the VIP/data interface of the SE. You should reserve a range of static IP addresses to be used as VIPs in the subnet assigned to the VIP/data segment. Having a dedicated segment for VIPs makes managing the IP ranges independent of other subnets easier.

Optionally, you can place the data interface of the SE on one of the server segments if you do not need a separate VIP segment. The subnet assigned to the server segment must have enough free static IP addresses to be used as VIPs.

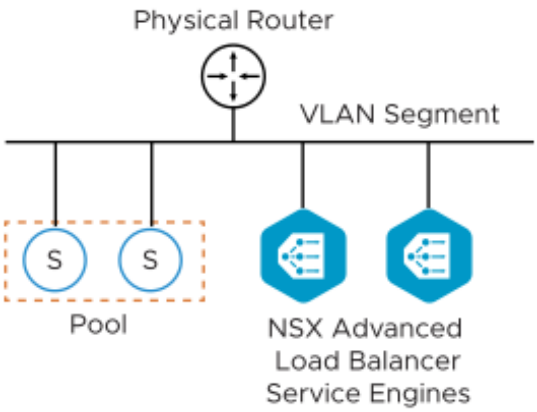

VIP Networking

For the virtual services placed on these SEs the VIP can belong to the subnet of the logical segment it is connected to or any other unused subnet. Once the virtual service is placed on the SE, the Controller updates the VIP static routes on the tier-1 router associated with the logical segment selected for the virtual service placement. The NSX admin is expected to configure the tier-1 router to redistribute these static routes with tier-0. For north-south reachability of the VIP, admin must configure the tier-0 to advertise the VIP routes to external router through BGP.

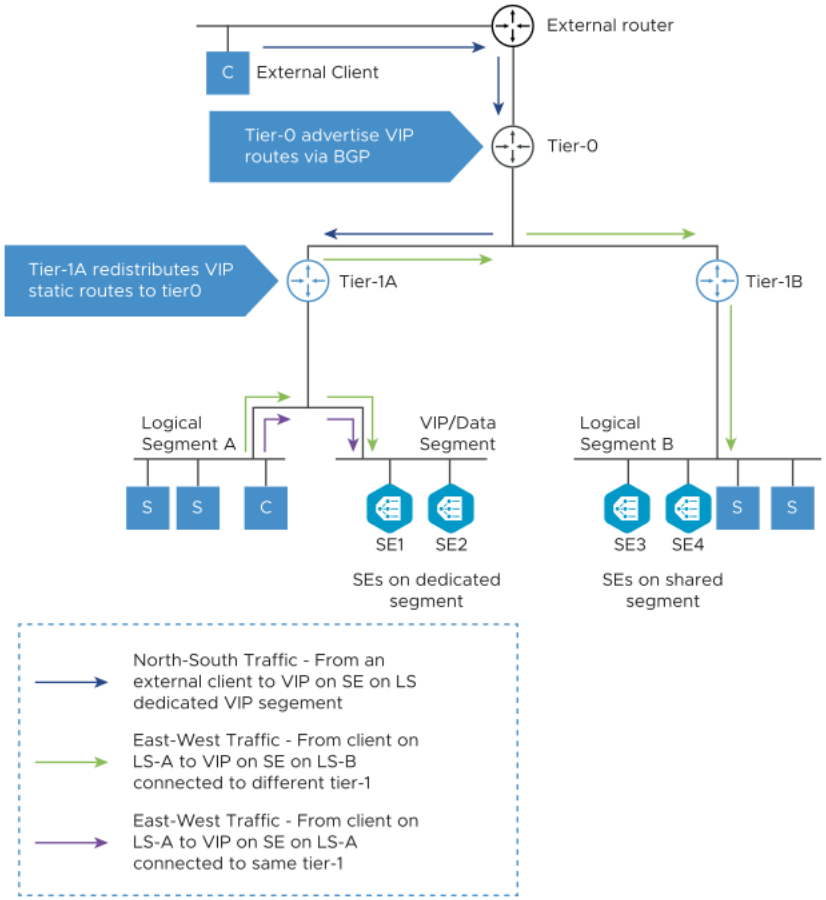

There are two traffic scenarios as discussed below:

North-South Traffic

As shown in the figure above, when an external client sends request to the VIP it gets routed from the external router to tier-0 which forwards it to the correct tier-1, which routes it to the VIP on the SE.

East-West Traffic

For a client on the NSX overlay trying to reach a VIP, the request is sent to its default gateway on the directly connected tier-1. Depending on where the VIP is placed there can be 2 sub-scenarios:

- **VIP on the SE connected to a different tier-1**: The traffic is routed to tier-0 which forwards the traffic to the correct tier-1 router. This then routes the traffic to the SE.
- **VIP on the SE connected to the same tier-1**: The traffic is routed to the SE on same tier-1.

HA Modes and Scale Out

All HA modes (Active-Active, M+N and Active-Standby) are supported in NSX-T environment. When a VIP is placed on an SE, the Controller adds a static route for it on the tier-1 router, with the SE's data interface as the next hop.

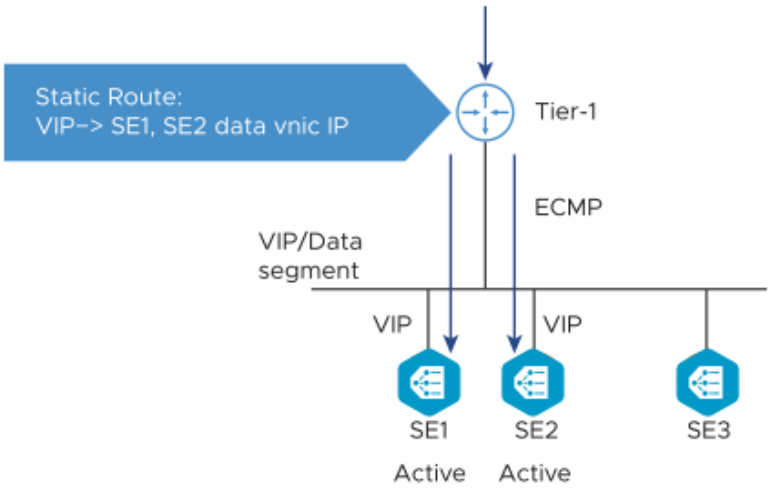

In the case of Active-Active and M+N HA modes, when the virtual service is scaled out, the Controller adds equal cost next hops pointing to each SE where the virtual service is placed. The tier-1 spreads out the incoming connections over the SEs, using Equal Cost Multi-Pathing (ECMP).

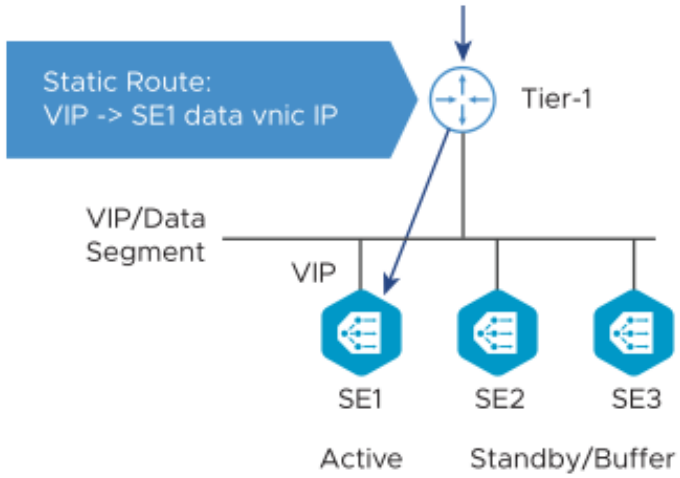

In case of Active-Standby where only one SE is Active or N+M HA mode with virtual service is not scaled out, the Avi controller programs route to the active SE only. ECMP is not required here.

Proxy Arp for VIP on Tier-1 and Tier-0

In Avi-NSX-T 3.1.0 integration, the proxy ARP functionality is available on both Tier-0 and Tier-1 gateways for Avi VIPs.

Proxy ARP on Tier-0 Gateway

The client and VIP are in the same segment, but the client is reaching the VIP through the tier 0. Proxying of the ARP or the VIP will be done by tier 0 to the external clients.

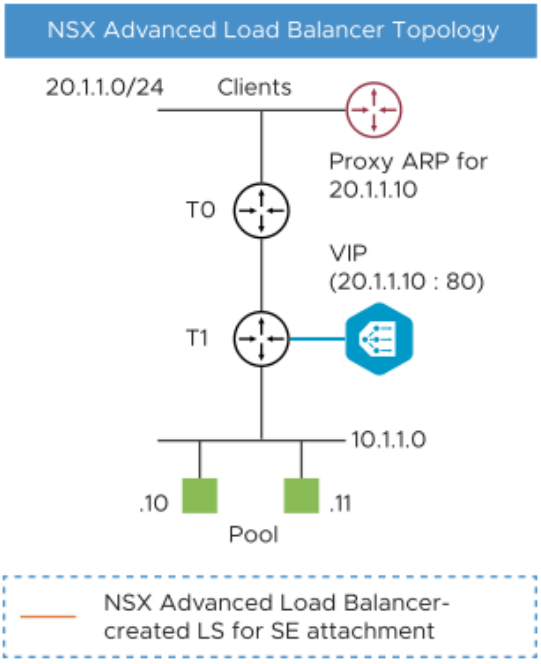

Proxy ARP on Tier-1 Gateway

If the client is local, (on the same tier 1), tier 1 does the proxy ARP for the VIP. Both the SE ad the tier 1 will respond, if they are attached.

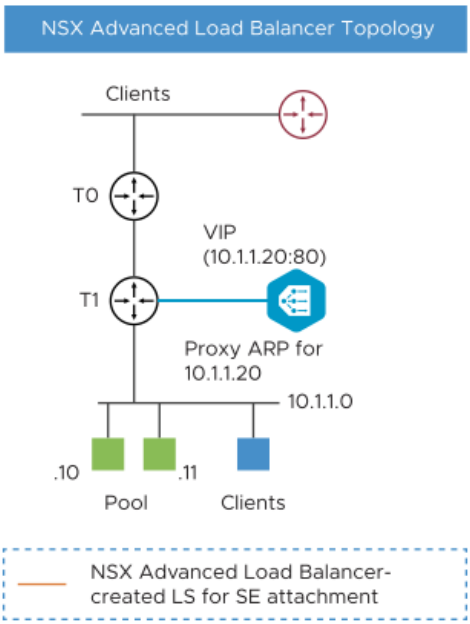

For Tier 0 to do the proxy ARP, enable the **"All LB VIP Routes"** for Tier-1. Starting with NSX-T 3.1, **"All LB SNAT IP Routes"** is also required:

1. From the NSX-T manager, navigate to **Networking > Tier-1 Gateway**.
2. Add a Tier Gateway or edit an existing one.
3. Under Route Advertisement, enable **"All LB VIP Routes"** and **"All LB SNAT IP Routes"**.

4. Click **Save**.

NSX Security Configuration

Security Automation

Creating NSGroups for SEs and Avi Controller management IPs is automated by the NSX-T cloud.

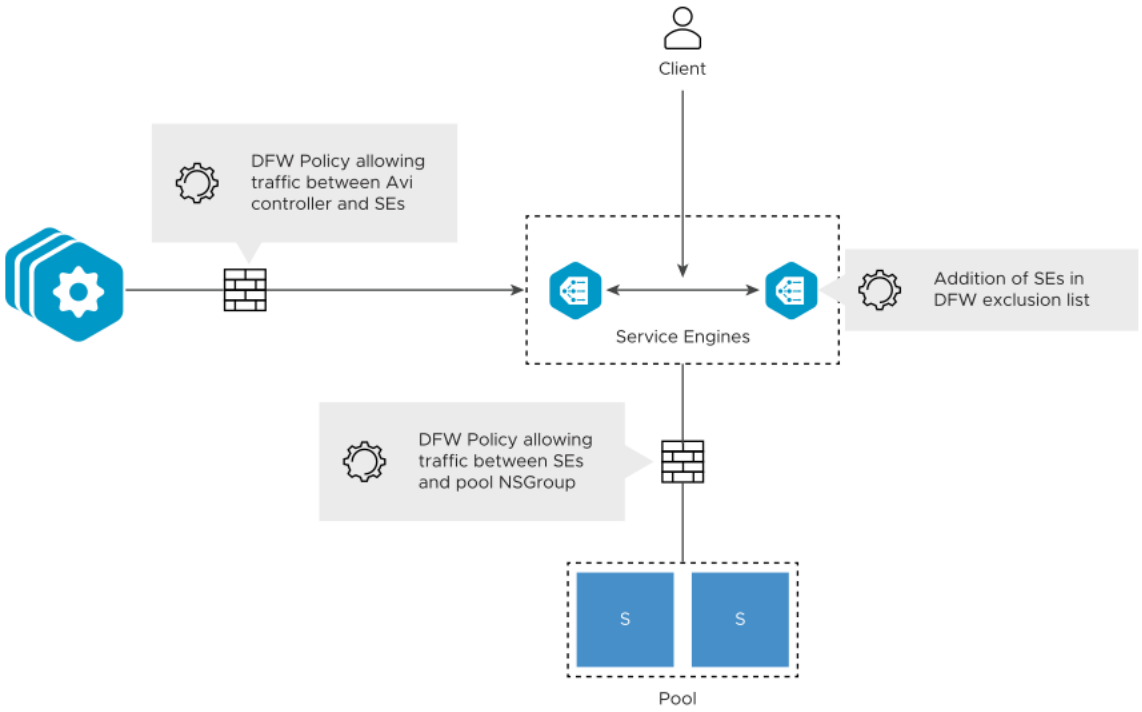

Run the following operations manually:

- Add the SE NSGroup to the exclusion list. This is required to allow cross SE traffic and prevent packet drop due to stateful DFW when asymmetric routing of application traffic happens.
- Create DFW policy to allow management traffic from SE NSGroup to Avi Controller NSGroup.

  Note: The SE initiates the TCP connection to the Avi Controller management IP
- For every Virtual Service configured on Avi, create a DFW policy to allow traffic from SE NSGroup to the NSGroup/IP group configured as pool.

Note: The NSX-T cloud connector creates and manages the NSGroups for different Avi objects. But the DFW rule creation is not supported currently. Add the service engine NSGroup to exclude the list before virtual service creation.

Since the SEs are in the exclusion list, DFW cannot be enforced on the Client to VIP traffic. This can be secured by configuring the network security policies on the virtual service. If the NSX-T gateway firewall is enabled, edge policies must be manually configured to allow VIP traffic form external clients.

Exclusion List

Avi SE redirects traffic from the primary SE to secondary SEs when using L2 scale out mode. This leads to asymmetric traffic which can get blocked by the Distributed Firewall because of its stateful nature. Hence to ensure that the traffic is not dropped when a virtual service scales out, you should add the SE NSGroup to the exclusion list.

This can be done by creating an NSG on NSX-T and adding the VIP/data segment as member. You can then add this NSG to the exclusion list. This way if a new SE is deployed its VIP/data interface will dynamically get added to Exclusion list. In case the SEs are connected to server segment, adding the segment to Exclude list is not an option as that will put all servers in the list too. You need to add individual SE VMs as members to the NSG.

Distributed Firewall

Avi Controllers and SEs require certain protocols/ports to be allowed for management traffic as listed in Ports and Protocols section. If the distributed firewall is enabled with default rule as block/reject all, create the following allow rules on DFW:

Controller UI Access

Source — Any (can be changed to restrict the UI access)
Destination — Avi Controller management IPs and Cluster IP
Service — TCP(80,443)
Action — Allow

Note: This rule is required only if Avi Controller is connected to NSX-T managed segment.

Controller Cluster Communication

Source — Avi Controller management IPs
Destination — Avi Controller management IPs Service — TCP(22, 8443) Action — Allow

Note: This rule is required only if Avi Controller is connected to NSX-T managed segment.

SE to Controller Secure Channel

Source — Avi SE management IPs
Destination — Avi Controller management IPs
Service — TCP(22, 8443), UDP(123)
Action — Allow

Note: SE initiates TCP connection for the secure channel to the Controller IP.

SE to Backend

Source — Avi SE data IPs
Destination — Backend server IPs
Service — Any (can be restricted to service port, for instance, TCP 80)
Action — Allow

Note: Client to VIP traffic does not require a DFW rule as the VIP interface is in Exclusion list. The front-end security can be enforced for each VIP using network security policies on the virtual service.

North-South VIP

If a certain VIP is required to have north-south access to allow external clients to reach the application, additional configuration is required on NSX-T manager.

- Tier-1 to advertise static routes to Tier-0
- Tier-0 to re-distribute Tier-1 advertised static routes to external peer.

Note: If all VIPs are required to be north-south or for simplicity of configuration, you can configure Tier 1 to redistribute the entire VIP range to Tier 0. You can configure Tier 0 to advertise all learned routes in that range to external peer. This way whenever a new VIP is created, it will be automatically advertised to the external peer.

Recommended Actions

- Configure Unique Object Name Prefix while Create Cloud with NSX-T
    1. The Avi NSX-T Cloud Connector will create NSX Inventory resources (Services and Groups) with the configured 'Object Name Prefix' in the Cloud configuration on Avi.
    2. Prefix can't be modified on cloud configuration .
    3. During cloud creation the following NSGroup(s)/NSService(s) will be created:

Object	Naming Convention	Description
Group	<prefix>-ControllerCluster	Contains all the Avi Controller Management IPs
Group	<prefix>-ServiceEngineMgmtIPs	Contains all the Avi Service Engine IPs
Group	<prefix>-ServiceEngines	Contains all the Service Engines as VMs
Service	<prefix>-ControllerCluster	Contains protocols/ports for the Controller. Allows TCP ports 22, 8443 and UDP 123.

    4. During load-balanced application creation the following NSGroup(s)/NSService(s) would be created:

Object	Naming Convention	Description
Group	<prefix>-<VS-Name>	Contains all the data vNIC IPs of all the Avi Service Engines servicing traffic for this load-balanced application (vs)
Group	<prefix>-<VS-Name>VsServiceEngines	Contains all the Service Engine VMs servicing traffic for this loadbalanced application (vs)
Service	<prefix>-<VS-Name>	Contains protocols/ports for the load-balanced application (vs)
Service	<prefix>-<Pool-Name>	Contains protocols/ports for the backend servers (pool)

• Dedicated Tier-1 GW and segment for SE management.

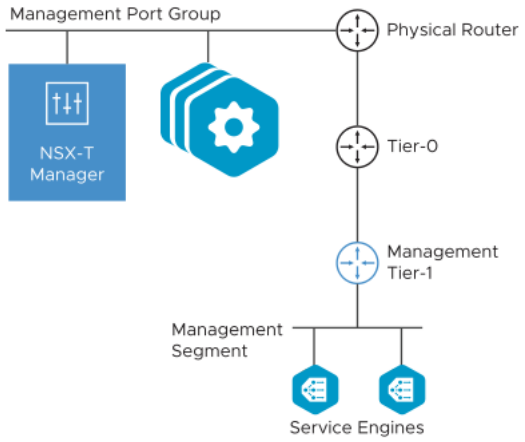

The network interface 1 of the SE VM is connected to the management segment. You should configure the management Tier 1 to redistribute the connected subnet routes to the Tier 0. Tier 0 must advertise the VIP to external peer using BGP.

• Single T1 with VIP Segment(s) please refer Design option - **One Arm Mode with overlay VIP Segment - Single Tier 1**

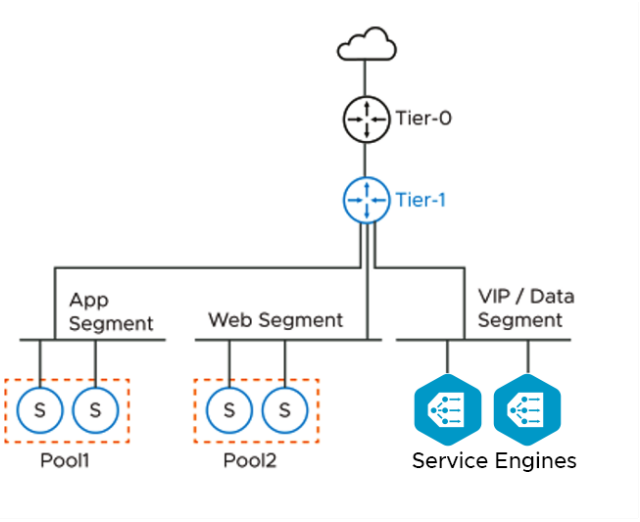

• Multiple T1 with VIP Segment(s) please refer Design option - **One Arm Mode with Overlay VIP Segment - Multiple Tier 1**

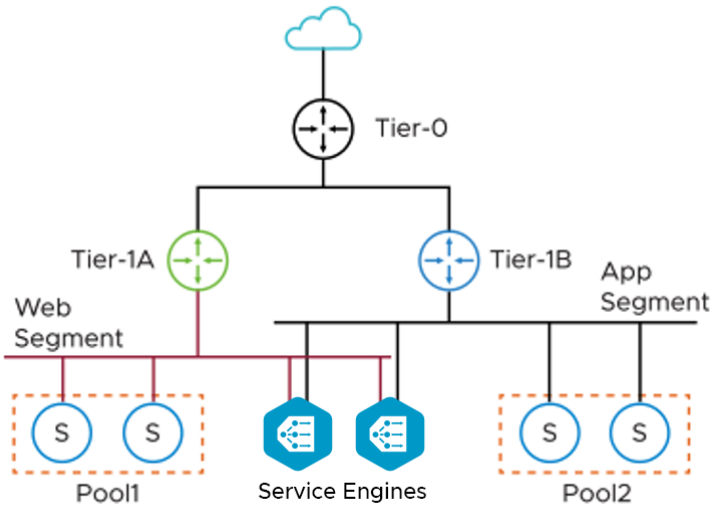

- VLAN VIP Segment(s) please refer Design option - **One Arm Mode with VLAN VIP Segment**

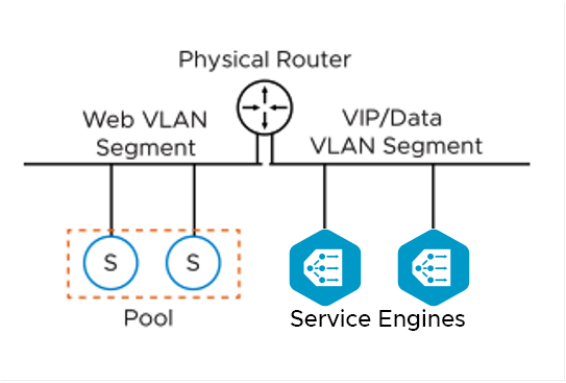

- SE management IP allocation can be DHCP (DHCP configuration required on the Tier-1) or Static, if using static needs a static pool configuration under networks (It will not provide default gateway to the service engine)
- Currently, only Manual is supported as the Logical Segments Config Mode. Hence the option is greyed out. This requires the segment to be pre-created in the NSX manager.
- Organize SEs using the "Service Engine Folder" advanced option on the Service Engine Group
- Configuration, folder must exist on vCenter, Avi Integration will not create the SE folder automatically.
- Leverage Host and Data Store Scope, under Service Engine Group advanced options, use cases is to make sure we have service Engines are deployed under specified cluster / Host / Datastore.
- Isolation for different Tier-1 GW on the Service Engine is performed using VRF by Tier-1.
- All deployments are one arm (client to VIP traffic and SE to backend server use the same SE interface) Need to add static route to reach backend server in case SE interface and Backend servers are in different subnets.

- For pool configuration, group application VMs by NSGroups, these objects can be imported by the Avi controller during VS creation, NSGroup changes will be synced automatically (5 min default).
- Manual operations:

1. Create DFW policy to allow management traffic from SE NSGroup to VMware Avi Load Balancer Controller NSGroup.

   Note: The SE initiates the TCP connection to the Controller management IP

2. Avi SE redirects traffic from the primary SE to secondary SEs when using L2 scale out mode. This leads to asymmetric traffic which can get blocked by the Distributed Firewall because of its stateful nature. Therefore, you should add the SE NSGroup to the exclusion list.
3. For every Virtual Service configured on Avi SE, create a DFW policy to allow traffic from SE NSGroup to the NSGroup/IP group configured as pool.
4. To enable North-South connectivity, you should configure the following on the NSX-T Manager:
   a. Tier-1 to advertise static routes to Tier-0.
   b. Tier-0 to re-distribute Tier-1 advertised static routes to external peer.

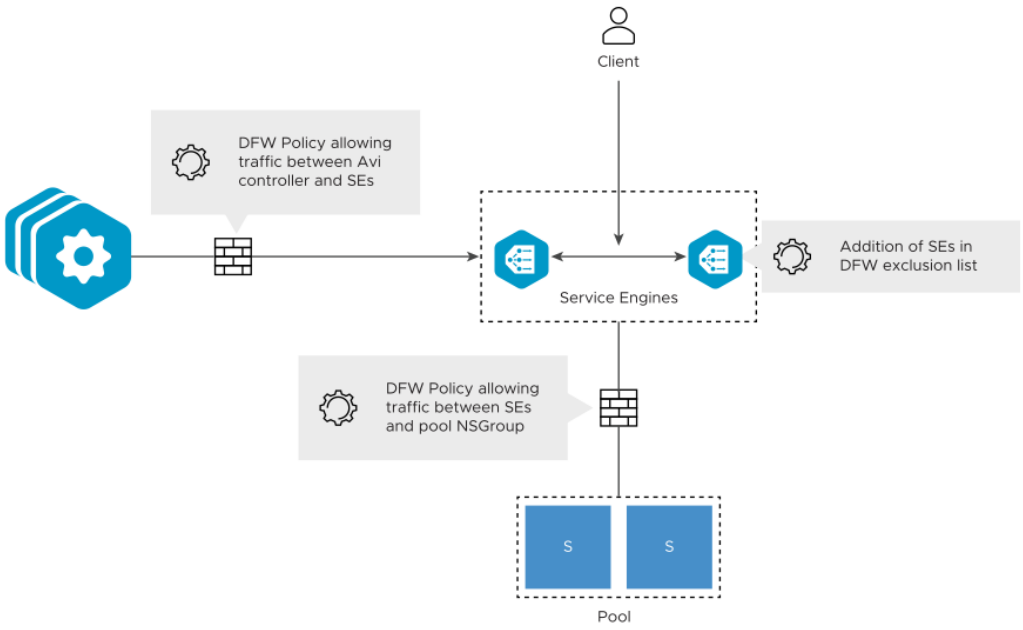

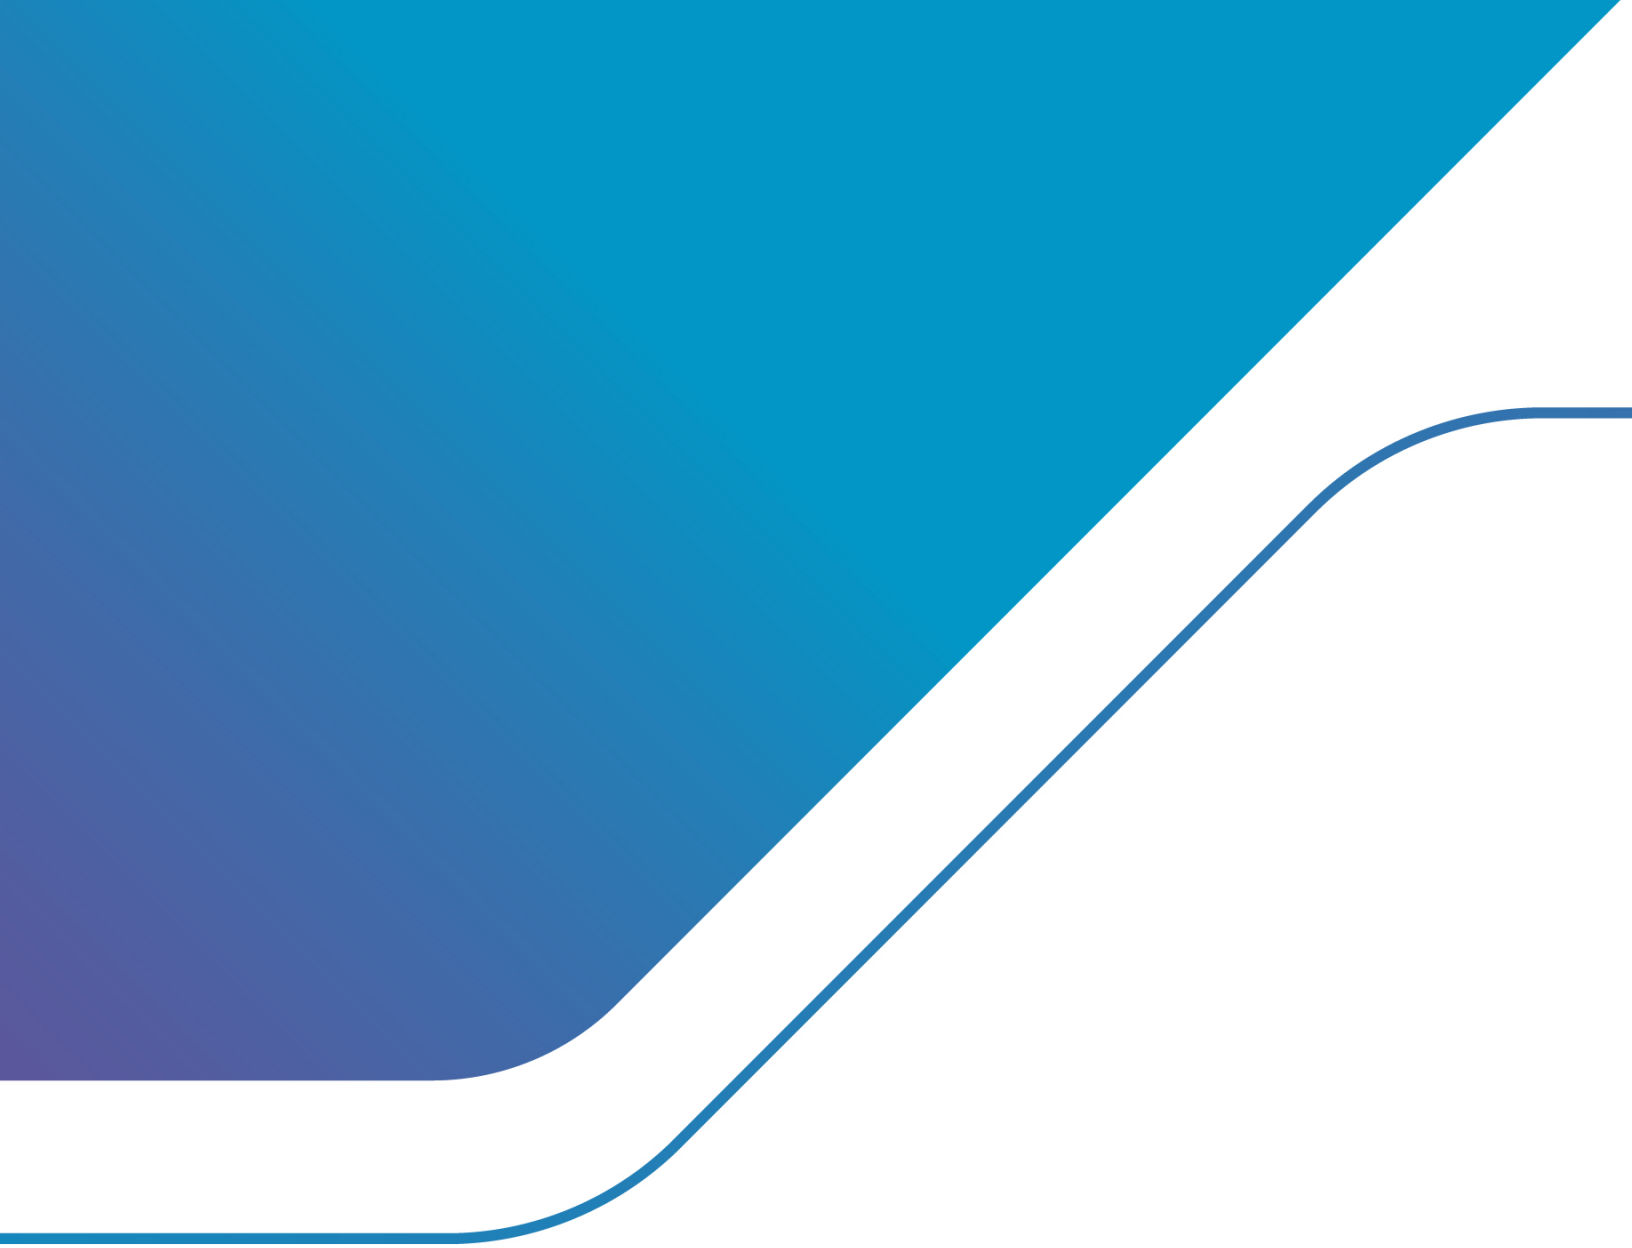

vmware® by Broadcom

Item No: vmw-bc-wp-tech-temp-uslet-word-2024 1/24