VMWARE NSBU FIELD ENGINEERING– **January 2022**

BEST PRACTICES FOR NSX ADVANCED LOAD BALANCER (AVI) INTEGRATION WITH NSX-T CLOUD

VMware Network & Security Business Unit

Table of Contents

About this Document

This technical document describes step-by-step process to integrate NSX Advanced Load Balancer (Avi) with NSX-T Cloud and the associated automation work-flow. It can also be used as a design guide for NSX-T integration with Avi with recommendations on best practices as applicable.

Contact Details

Company	VMware
NSBU Field Engineering Contact	Md Abdul Aziz <mdabdula@vmware.com> Ajay Mare <marea@vmware.com> Abhijith KS <kabhijith@vmware.com> Harsh Jaitly<hjaitly@vmware.com>

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 NSX Advanced Load Balancer (Avi) 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.

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://avinetworks.com/docs/20.1/roles-and-permissions-for-vcenter-nsx-t-users/

Applicability

- NSX-T versions 2.5 and above
- vCenter 6.7 and 7.0
- NSX Advanced Load Balancer (Avi Networks) – 20.1.4 and above

Caveats

Starting with Avi version 20.1.3, integration with multiple NSX-T clouds is supported with a maximum of 5 NSX-T clouds Also, multiple vCenters can be associated per NSX-T cloud.

Please note the following caveats with older versions of Avi lease note the following caveats.

- *Only 50 VRFs are supported per Avi controller cluster. Therefore, additional Avi controller clusters are required if if a customer has over 50 NSX-T T1s,*
- *Prefix of NSX-T cloud configuration cannot be modified.*
- *Maximum Virtual Services supported on a single NSX-T T1 is ~ 2000 VIPs This limitation is a consequence of the maximum number of static routes supported by a single NSX-T T1 Router)*

Creating an NSX-T Cloud

To create an NSX-T cloud, log in in to the Avi Controller and follow the steps given below:

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

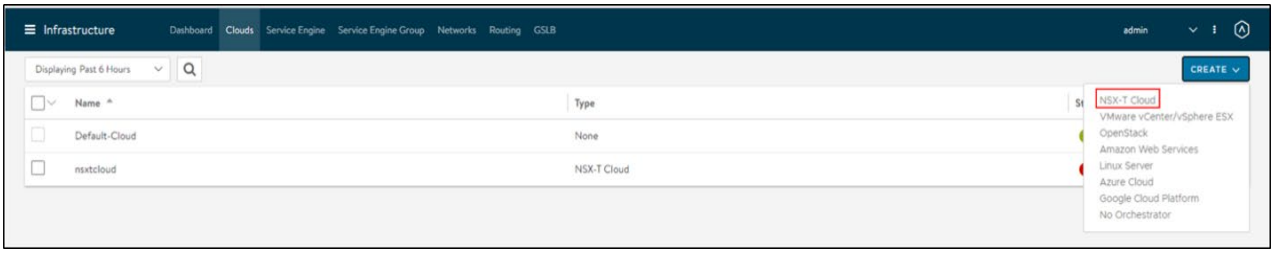

3. Enter the Name of the NSX-T cloud.
    A. Note: NSX-T Cloud is selected as the Cloud Type by default.
4. Check the DHCP option if SE management if segment has DHCP enabled.
    A. Note: The prefix string must only have letters, numbers, and underscore. This field cannot be changed once the cloud is configured.
5. Enter the NSX-T manager hostname or IP address as the **NSX-T Manager Address** and select the NSX-T Manager Credentials.
6. Click on **Connect** to authenticate with the NSX-T manager.

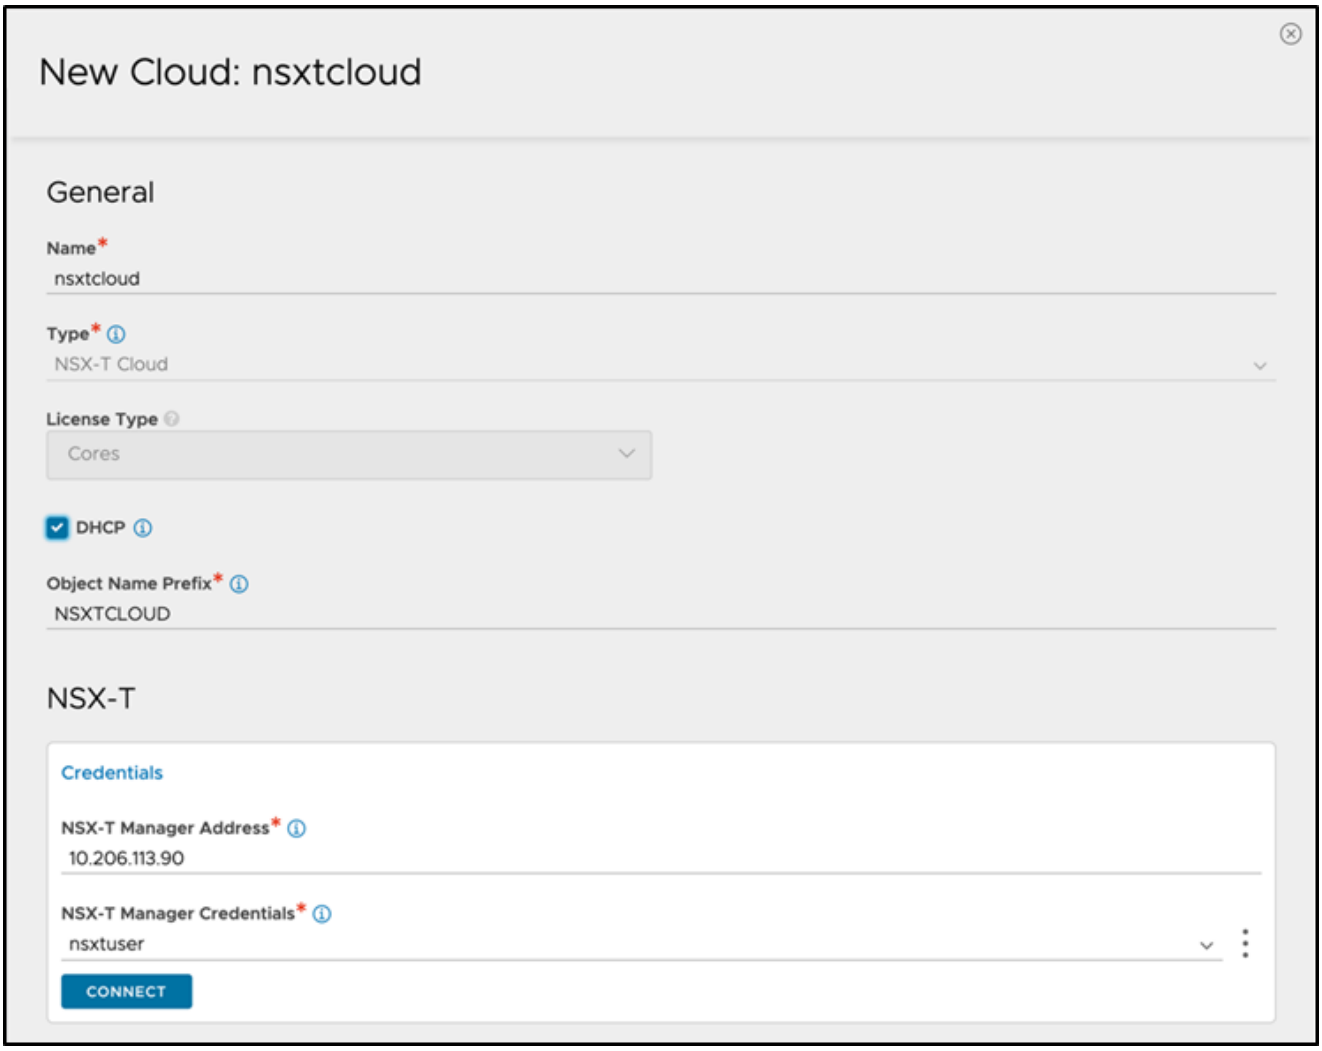

7. Select the **Transport Zone** required from the drop-down. (Only the Overlay type is supported)
    A. Only the Overlay type is supported on Avi release 20.1.1 to 20.1.4
    B. Overlay and VLAN for Service Engine Management is supported from Avi 20.1.5
    C. Overlay and VLAN for Service Engine Data is supported from Avi 20.1.6
8. In the Management Network Segment, select the Tier1 Logical Router ID and Segment ID.
   Note: Currently, only the Manual is supported as the Logical Segments Config Mode. Hence the option is greyed out. This requires the segment to be pre-created on NSX manager.
9. Select the tier-1 gateway and logical switch for VIP placement.
10. Click on **Add** to select one more tier-1 router and a connected logical segment for VIP placement

11. Under **vCenter Servers**, click on **Add**.
12. Enter the vCenter Server Name, and configure the credentials.
13. Click on **Connect**.
14. Select the Content Library and **DONE**.
    A. Select Existing content library on vCenter
    B. If there is no content library need to configure new content library to upload SE OVA file

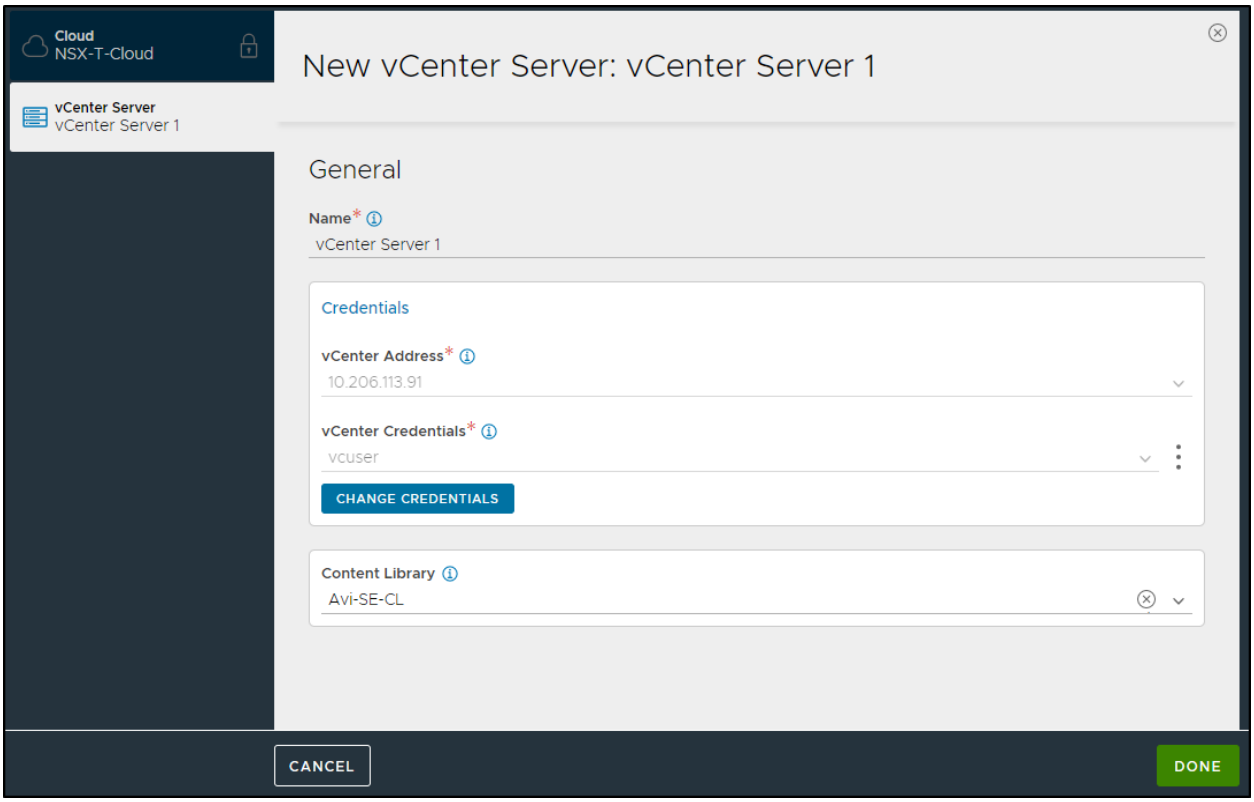

15. Select the IPAM/DNS Profile, as required.

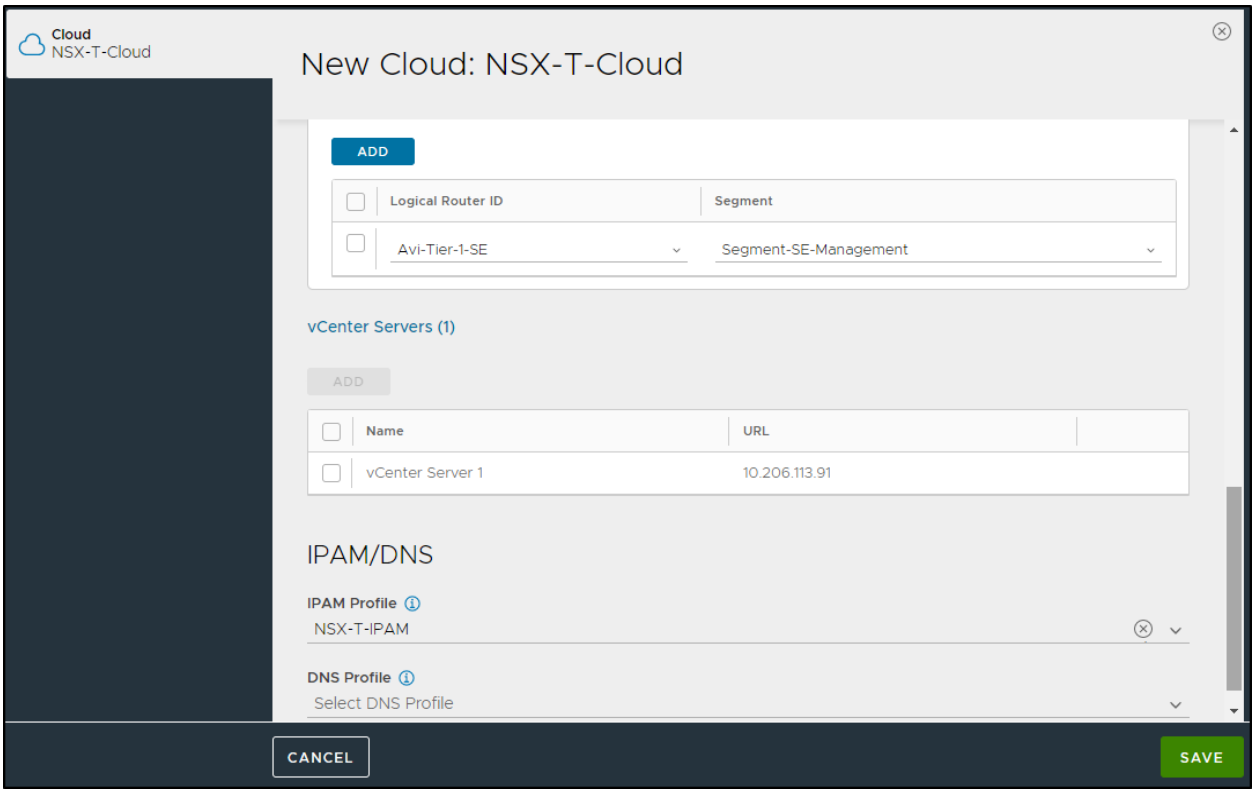

16. Click on **Save** to create the NSX-T cloud.

Multiple NSX-T Clouds

- *Starting with Avi version 20.1.3, multiple NSX-T clouds (maximum of 5) can be created.*
- Each NSX-T manager can be either created for the same NSX-T manager or different NSX-T manager. If different NSX-T managers are pointing to the same vCenter, then only one SE image per vCenter will be created.
- If there are multiple NSX-T managers pointing to different vCenters then the SE image will be created in the respective content libraries.
- Note: The clean-up of the SE image happens only after the last NSX-T cloud pointing to the SE image is removed.

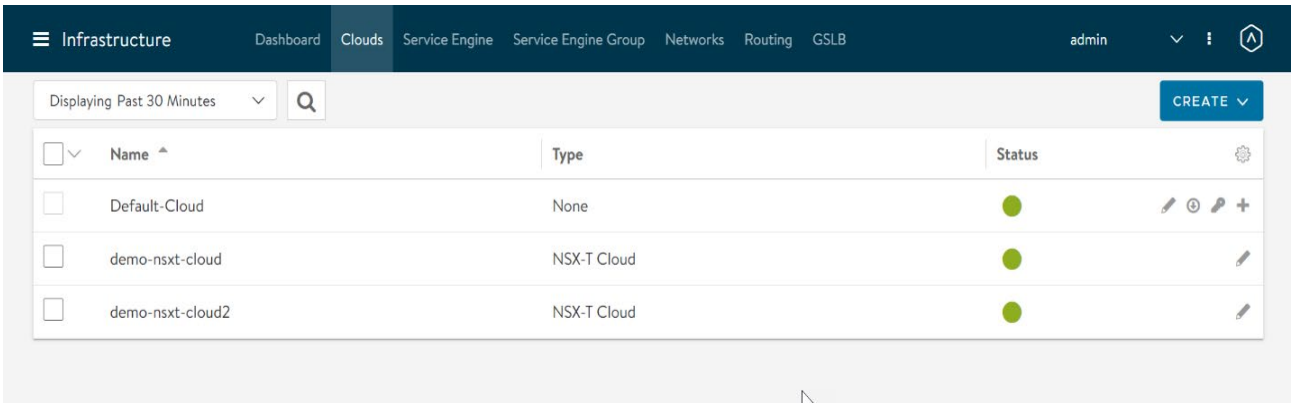

Cloud Configuration flow:

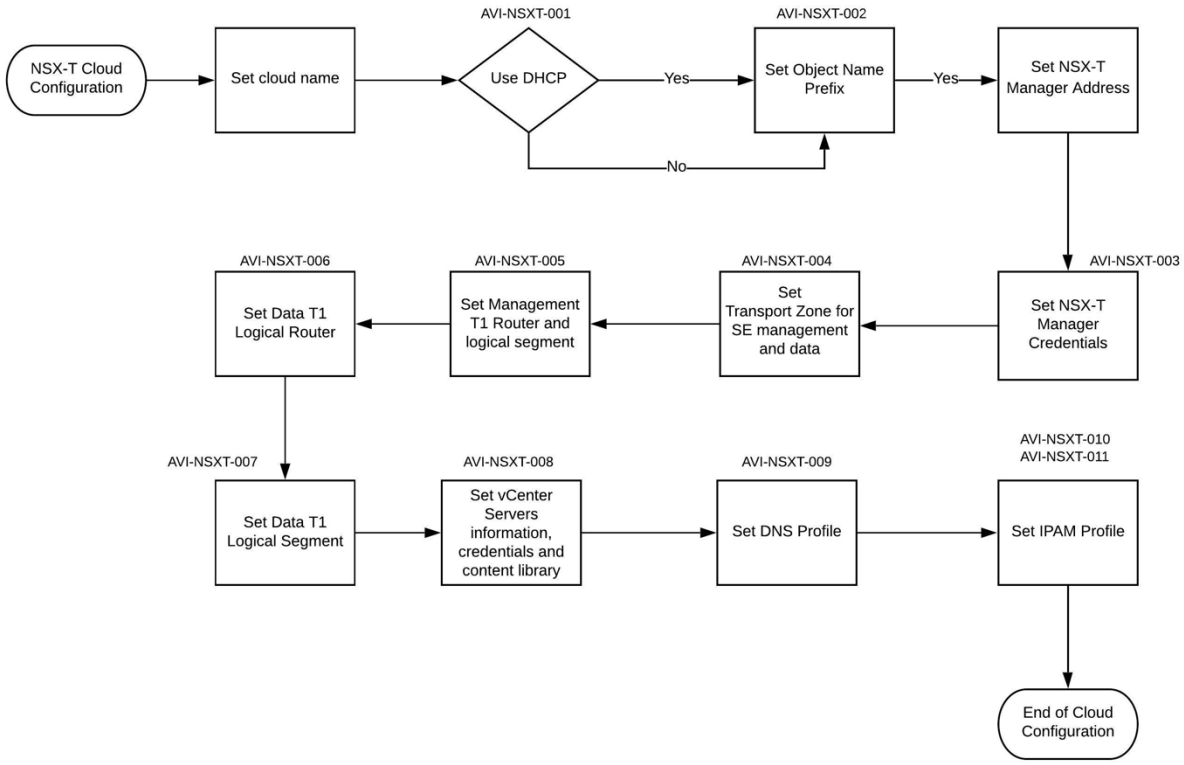

SE Group Configuration Flow:

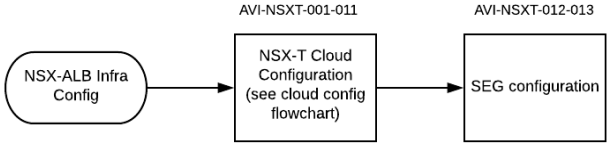

Virtual Service Configuration flow:

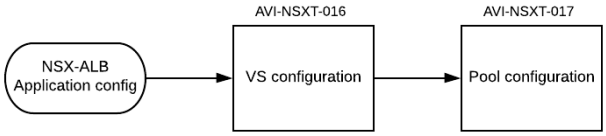

Cloud Configuration:

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 Decision ID - AVI-NSXT-012 Where need to defined static pool and enable Static IP Address Pool - Use Static IP Address for VIPs and SE or Use for Service Engine under network to allocate IP Addresses to SE interfaces.
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 LB services. It is recommended to configure a remote user using VIDM, LDAP, etc. on NSX-T with Network

			Engineer privileges to be used as a service account for Avi Controller to NSX-T Manager API communication. In cases where remote auth is not configured on NSX-T, use an NSX-T local user. Beginning NSX-T 3.1.1, 'guestuser1' user would be the recommended local user account. Prior releases would need to use the NSX-T admin account and assign "Network Engineer Role" (From NSX-T 3.1 onwards Network Engineer Role is renamed as Network Admin) Ref: https://avinetworks.com/docs/20.1/roles-and-permissions-for-vcenter-nsx-t-users/
AVI-NSXT-004	Selection of transport zone type for SE management	Overlay and VLAN (on version 20.1.5 or later) types are supported	Selected transport zone type is used for communication between Service engine and Controller
AVI-NSXT-005	Select a T1 router and segment for SE management	Segment must exist on NSX-T, Avi will use an existing one to place the SE VM management NIC on it.	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 existing ones to place the SE VM data NICs on them.	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 and IP Address 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 Static POOL for SE Data interface IP and Backend servers are in different subnet than data interface, we need to configure static route on SE to reach back servers. Because static Pool on controller doesn't 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 to allow clients to access load balancer services, as well for communication between the SE and backend servers.	When 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 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 on vCenter, 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 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://avinetworks.com/docs/20.1/roles-and-permissions-for-vcenter-nsx-t-users/#vcenter-roles
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 consumption when clients can reach services 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 specify the VIP and select the T1 router during the VS creation.	When not using an IPAM profile there is a risk of selecting an IP already in use by another VM on the segment, especially if the segment is shared with other workloads and not dedicated to VIP/Data allocation.
AVI-NSXT-012	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 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. And 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. 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 config (non-cloud)

AVI-NSXT-13	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-14	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-15	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 as indicated here: https://avinetworks.com/docs/20.1/setting-up-nsx-alb-basic-edition/ The Avi controller 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-NSXT-16	Deploying Avi on a multi-tenant 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 same subnet. VRF avoids issues when different tenants use the same subnet ranges.

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.

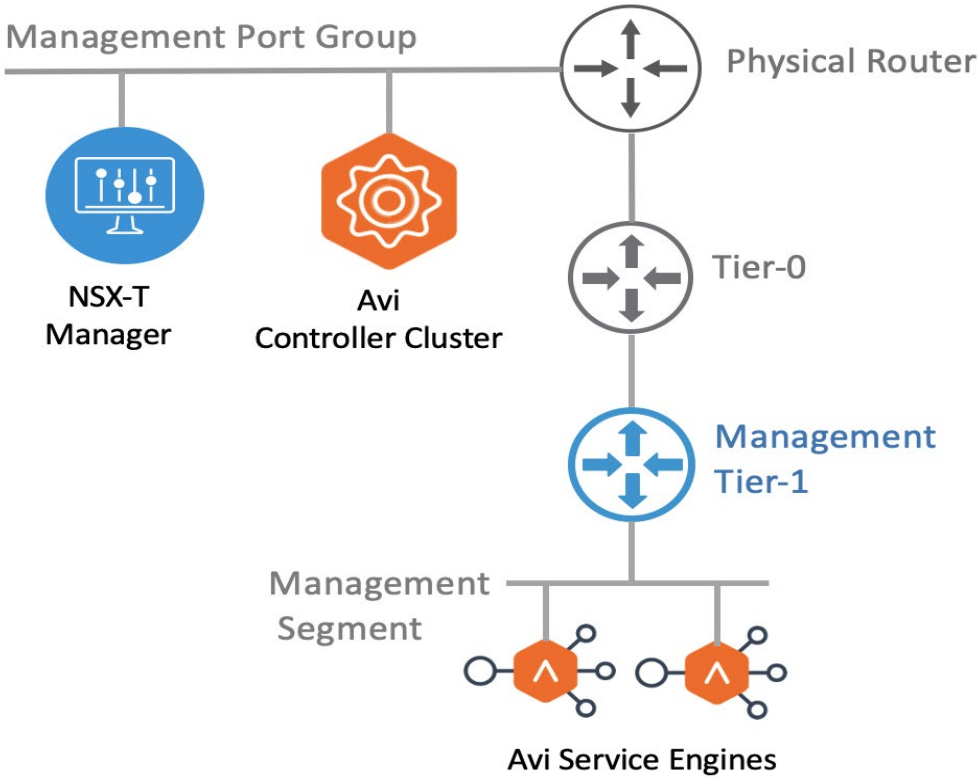

The network first network interface of all SE VMs are connected to the management segment. The management IP address of the SEs must be reachable from the Avi Controller. All Connected Segments & Service Ports" and "All Static Routes" to the tier-0. Tier-0 must advertise the learned routes to external router using BGP.

Load Balancer Topologies

- Avi currently supports load balancing only in an NSX-T transport zone of type overlay. The SE supports only one arm mode of deployment 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:

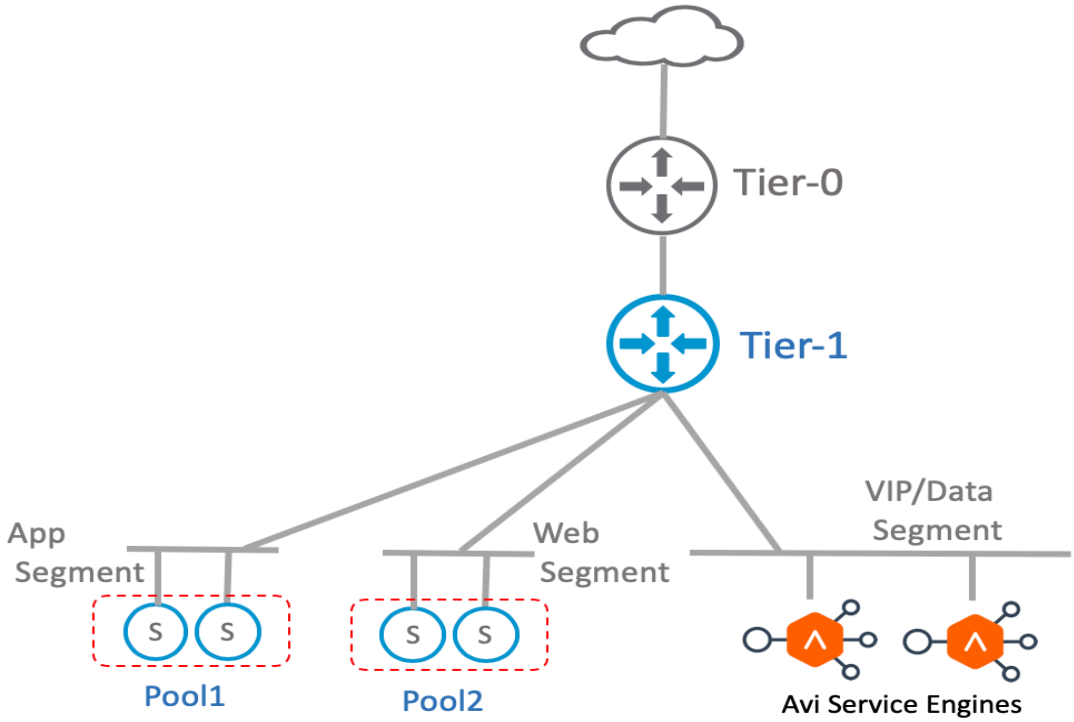

You need to create the VIP/data segment manually. The network adapter 1 of the Service Engine VM is reserved for management connectivity. You can connect only one of the remaining nine data interfaces (network adapter 2-10) of the Service Engine VM to the VIP/data segment. Rest of the interfaces must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

Decision ID	Design Decision	Design Justification	Design Implication
AVI-NSXT-017	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-018	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 Single 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.

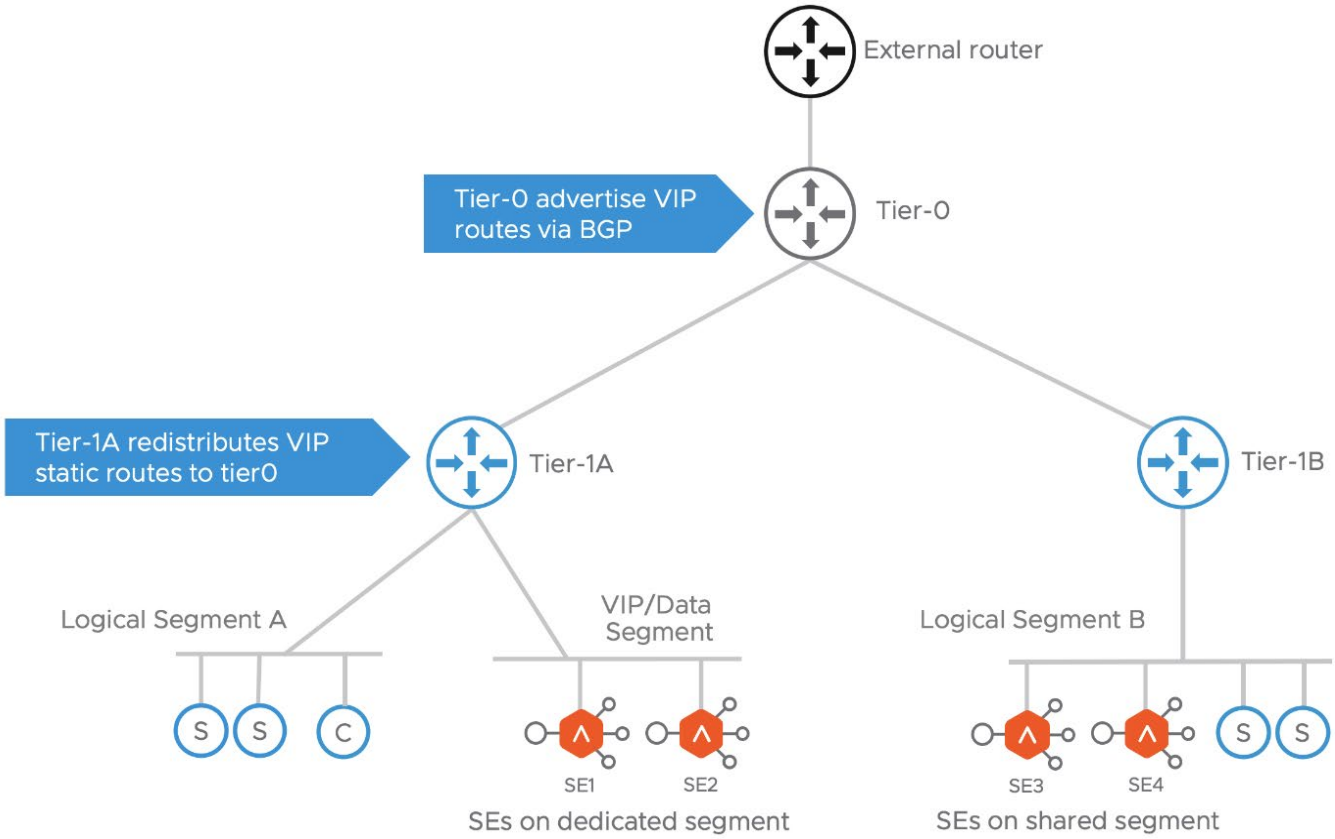

From shown in the diagram, two topologies are possible for SE deployment:

1. SEs on dedicated Logical Segment:
    - Allows to manage IP address assignment separately for SE interfaces
    - In the current version, this segment must be created on NSX-T prior to adding it to cloud configuration on Avi.

2. SEs on shared Logical Segment:
    - SE interfaces shares the same address space as the server VMs on same logical Segment.

Note: Only logical Segments connected to tier-1 router are supported. The cloud automation for NSX-T integration does not support placing SEs on logical segments directly connected to tier-0 routers.

Static route is configured on T1a.Therefore, only T1a AND T0 can do Proxy ARP. T1b does not have that static route so cannot proxy arp for the VIP.

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 Avi 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 should configure the tier-0 to advertise the VIP routes to external router via BGP.

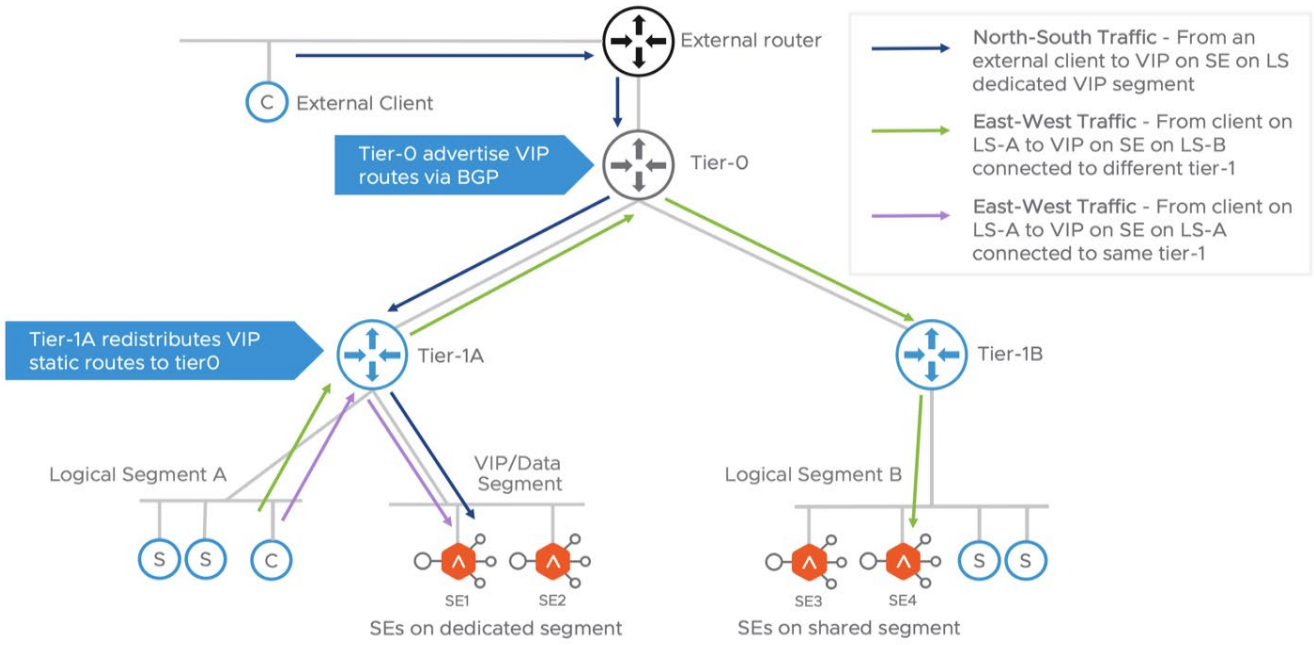

There are two traffic scenarios as discussed below:

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.

One Arm Mode with Multiple Overlay VIP Segment - Multiple Tier 1

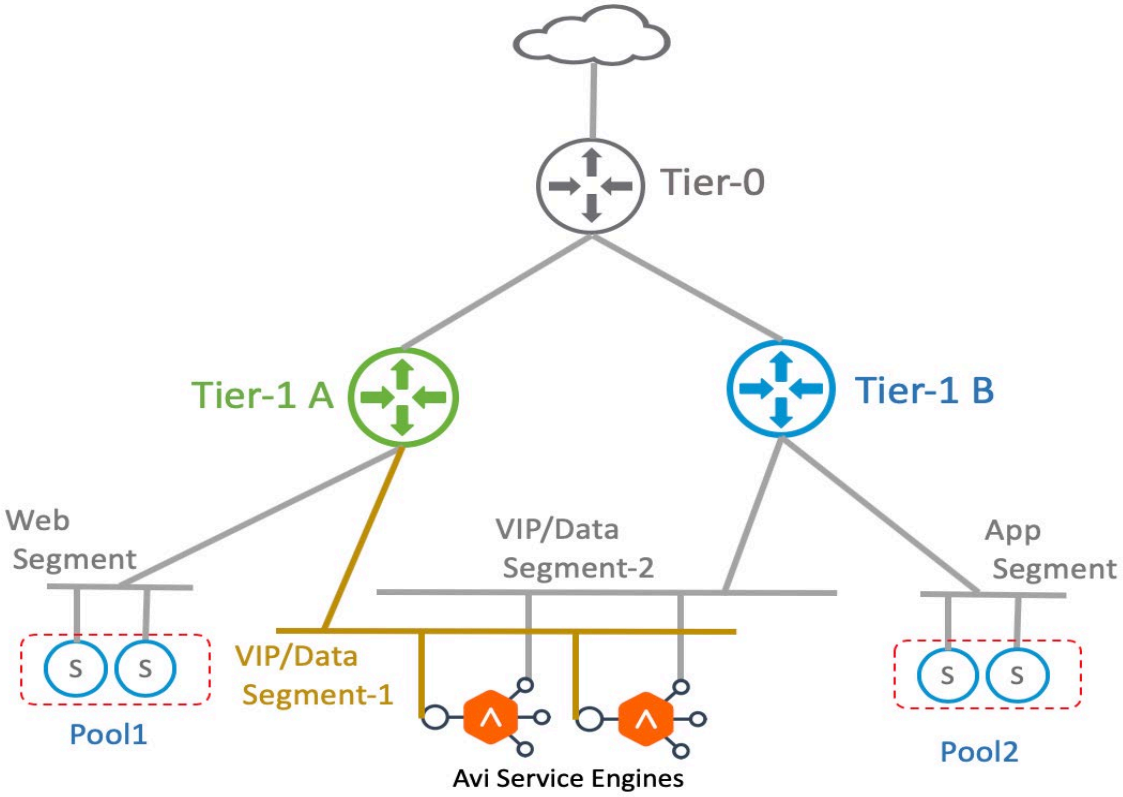

1. The network adapter 1 of the Service Engine VM is reserved for management connectivity
2. One data interface (network adapter 2) is connected to VIP/data segment-1
3. another data interface (network adapter 3) is connected to VIP/Data segment-2
4. Rest of the interfaces are kept disconnected.

- Admin needs to configure 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 respectively. Also, you should add SE interface connected to VIP/data segment-1 to VRF-A and add the interface connected to VIP/data segment-2 to VRF-B.

- While creating the virtual services for a pool, you should choose corresponding VRF. For instance, select VRF-A while creating a virtual service for Pool1 and VRF-B while creating virtual service for Pool2. This way the VIP of the virtual service managing Pool1 will be on VIP/data segment-1 and 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 does not have to hairpin to the Tier 0, and also because you can configure logical segments on different Tier 1 to have same subnet. Hence each Tier 1 traffic must be contained in its own VRF.

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, Avi 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 Avi controller adds equal cost next hops pointing to each SE where the the virtual service is placed. The tier-1 spreads out the incoming connections over the SEs, using Equal Cost Multi-Pathing (ECMP).

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

Prior to Avi version 20.1.3, if the client and the VIP are within the same segment, in the same subnet, ARP (Address Resolution Protocol) did not work because the client and VIP were in the same L2 Domain, resulting in no traffic flow.

Starting with Avi version 20.1.3, along with NSX-T 3.1.0, the proxy ARP functionality is available on both Tier-0 and Tier-1 gateways for Avi LB 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.

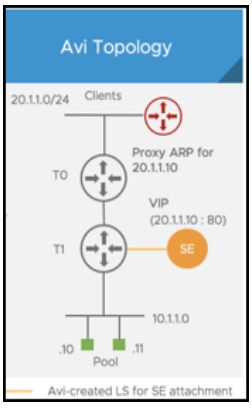

Proxy ARP on Tier-1 Gateway

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.

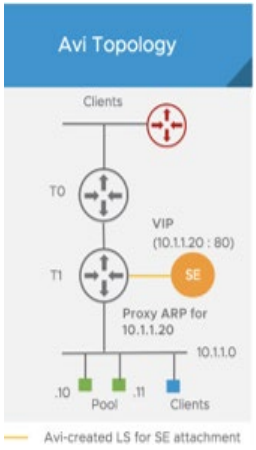

Decision ID	Design Decision	Design Justification	Design Implication
AVI-NSXT-019	Proxy ARP on T0 Gateway	External Client is reaching VIP over T0 Proxy ARP for VIP will be done by T0 Gateway	Starting with Avi version 20.1.3, along with NSX-T 3.1.0, the proxy ARP functionality is available on both Tier-0 and Tier-1 gateways for Avi LB VIPs

NSX Security Configuration

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

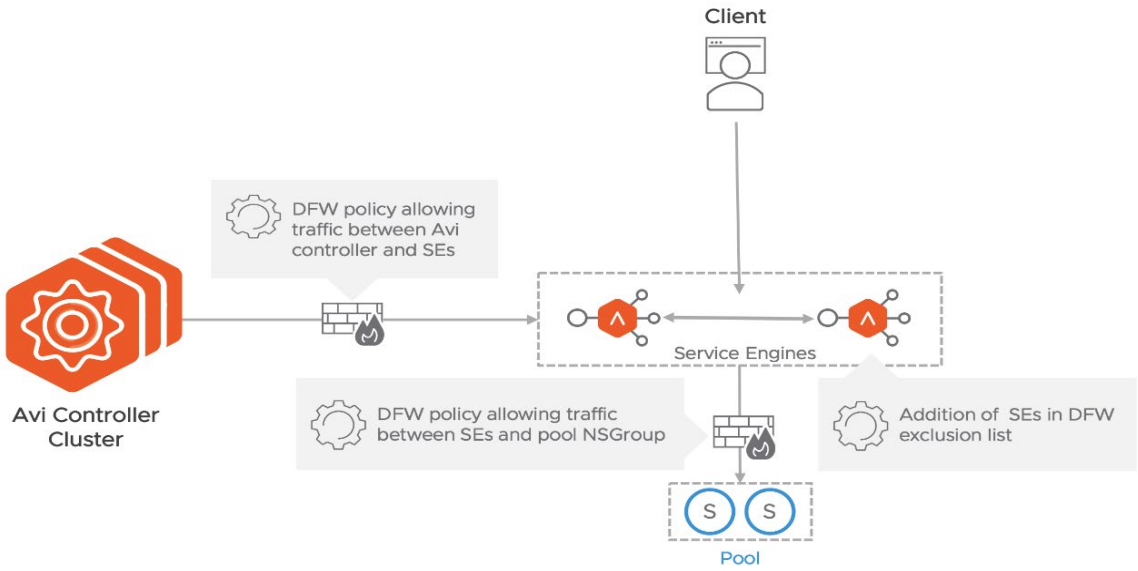

Creating NSGroups for SEs Avi Controller management IPs is automated by the NSX-T cloud.
Execute the following operations manually:

- 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 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 on Avi.
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 (TCP Strict Enable). Hence to ensure that the traffic is not dropped when a virtual service scales out, you should add the SE interfaces connected to the VIP/data segment 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 or TCP Strict Disable on NSX.
- 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 Controller and the SEs require certain protocols/ports to be allowed for management traffic as listed here. 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.

AVI-NSXT-020	Understand the creation of NSGroups	The NSX-T cloud connector creates NSGroups automatically for infrastructure components (SE, Controllers) as well as for VSs (VIPs, SE data IPs).	Those NSGroups are added to the exclusion list on the DFW. This allows the communication between the infrastructure components and the access to load balancer services provided by Avi (client to VIP and SE to pool members). For more information on the groups created please refer to https://avinetworks.com/docs/20.1/nsx-t-no-se-in-exclude-list/
AVI-NSXT-021	Secure Client to VIP traffic	Client to VIP traffic is allowed by the addition of an automatically created NSGroup to the exclusion list on the DFW.	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 on Avi. If the NSX-T gateway firewall is enabled, edge policies must be manually configured to allow VIP traffic form external clients.

Licensing

Customers who have active entitlements for NSX editions that have the load balancing capability are eligible to use the Basic edition. Any NSX serial key that has load balancing capability on the NSX platform can be imported onto the Avi Controller. Avi Controller would deposit an equivalent number of Basic Edition Service Core licenses based on the conversion table shown below.

NSX license type	Quantity	Avi Service Core
NSX cpuPackage	4	1
NSX vm	200	1
NSX concurrentUser	200	1
NSX core	48	1

In order to use the basic edition of the product, the Avi Controller needs be switch from the default Enterprise edition license to the Basic Edition by the following workflow from the UI:

- Administration > Settings > Licensing
- Click on the Settings icon
- Select Basic License
- Click on Save option

If regular enterprise licenses have been purchased there is no need to make any licensing changes on the controller, simply import the licenses.

It is important to note that the controller can only be in one license mode and cannot combine basic and enterprise features at the same time.

Feature difference between Basic License and enterprise license cab be found

https://avinetworks.com/docs/20.1/nsx-alb-basic-edition/#feature-comparison-datasheets

Recommended Actions

- Configure Unique Object Name Prefix while Create Cloud with NSX-T

  i. Avi Controller NSX-T Cloud Connector will create NSX Inventory resources (Services and Groups) with the configured 'Object Name Prefix' in the Cloud configuration on Avi.
  ii. Prefix can't be modified on cloud configuration
  iii. 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.

  iv. 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 load-balanced 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.

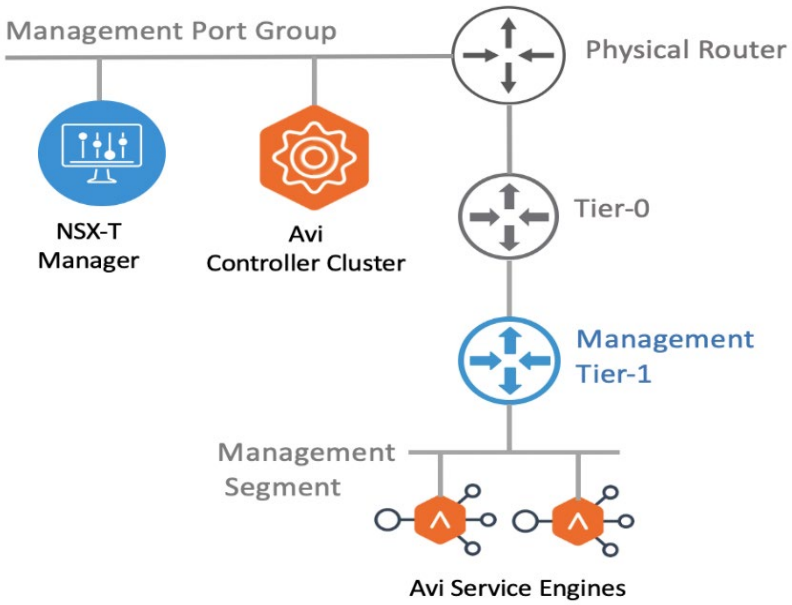

The network first network interface of all SE VMs are connected to the management segment. The management IP address of the SEs must be reachable from the Avi Controller

Note: Separate T1GW for redundancy, in case T1 GW goes down, recommend this or mention hosting the T1GW on a NSX Edge cluster.

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

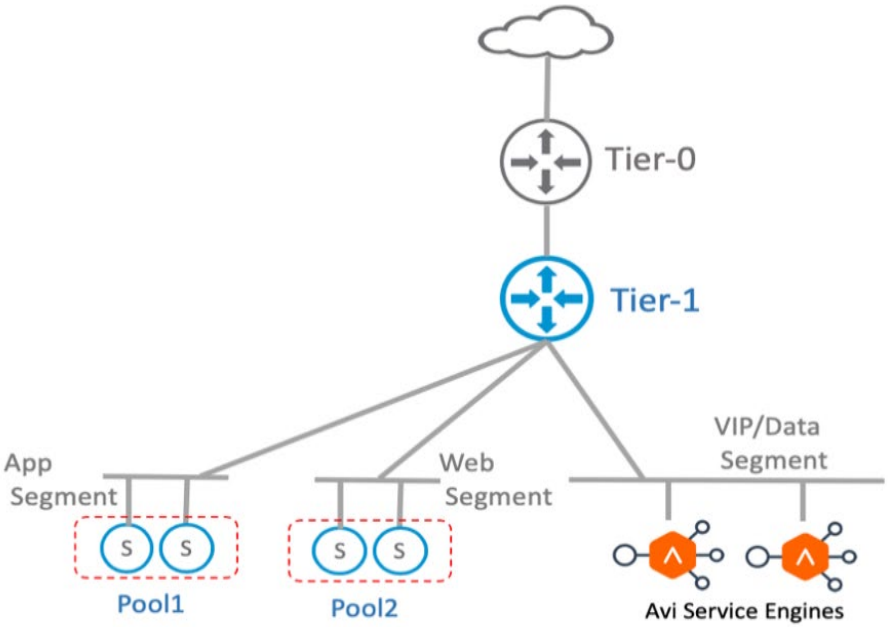

- Multiple T1 with Single VIP Segment please refer Design option - **One Arm Mode with Single Overlay VIP Segment - Multiple Tier 1**

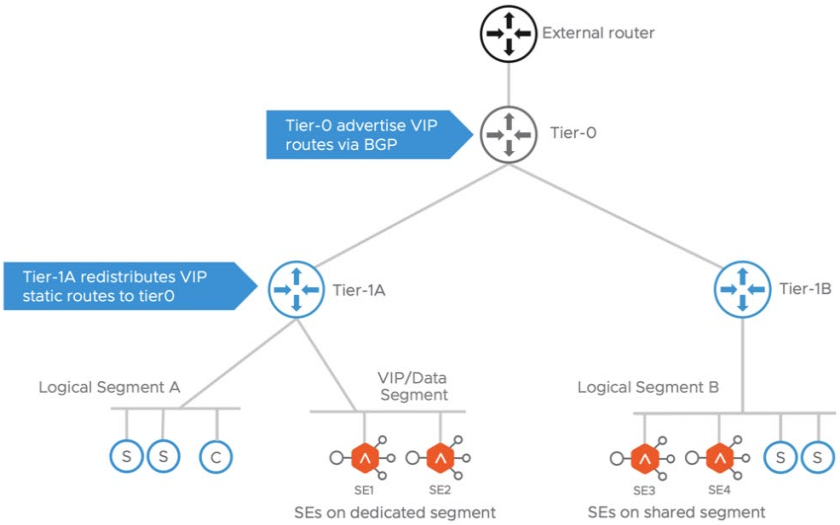

- Multiple T1 with Multiple VIP Segment please refer Design option - **One Arm Mode with Multiple Overlay VIP Segment - Multiple Tier 1**

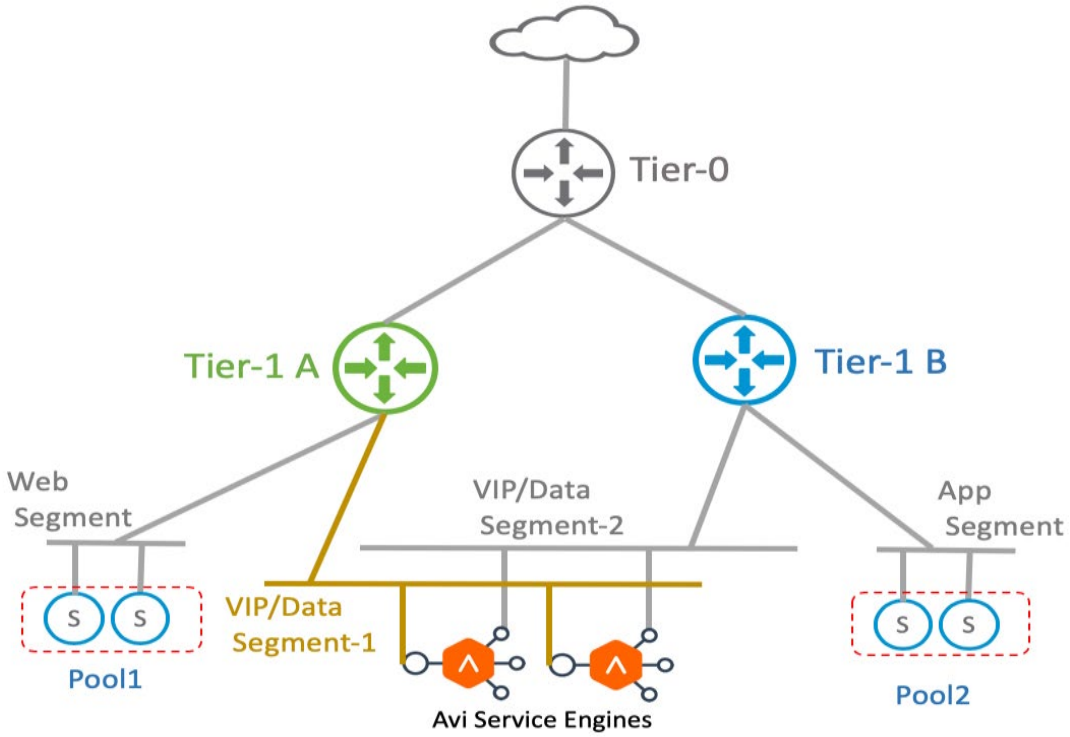

- SE management IP allocation can be DHCP (DHCP configuration required on the Tier-1) or Static, if using static needs a static pool configuration under networks (It will not provide default gateway to the service engine)

- Logical segments must be created manually, currently Avi Integration will not do create logical segments automatically.

- 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 T1 GW on the Service Engine is performed using VRF by T1.

- 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 incase SE interface and Backend servers are in different subnet

- 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:

    i. SE to Controller communication DFW policies, we can provide screenshots on how to do it, the KBs just mention it.
    ii. Add SE NSGroup to 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”, expand this and provide example.

    iii. DFW policy to allow traffic from SE NSGroup to the backend server NSGroup, is it enough from the exclusion list or do we have to manually go to each NSGroup configured as backend and set the DFW for it.

- To enable north-south connectivity, you should configure the following on the NSX-T Manager:

    i. Tier 1 to advertise static routes to Tier 0.

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