

- **Border Leaf**—This leaf node is typically connected to L3Outs. L3Outs can serve as a path into the WAN, or into the core of a legacy network.
- **Compute Leaf**—This leaf node is typically connected to compute resources, whether the resources are physical or virtualized servers.
- **Services Leaf**—Services within ACI are typically those given by Layer 4 to Layer 7 services. Services include firewalls, load balancers, and intrusion prevention systems. Services do not need to be integrated into ACI through a service graph template to be considered a service; that is a definition from the applications point of view.
- **Storage Leaf**—This leaf node is typically connected to storage devices for compute resources. This can include iSCSI, NFS, or other Ethernet medium storage devices.

Leaf nodes do not need to be delegated to only one category. Depending on the design, the categories can overlap. For example, a leaf node serving as a border leaf node can also provide compute resources.

Additional References for Leaf Node Categorization

For additional information on border leaf switches:

http://www.cisco.com/c/en/us/solutions/collateral/data-center-virtualization/application-centric-infrastructure/white-paper-c07-732033.html#_Toc395143551

Fabric Provisioning

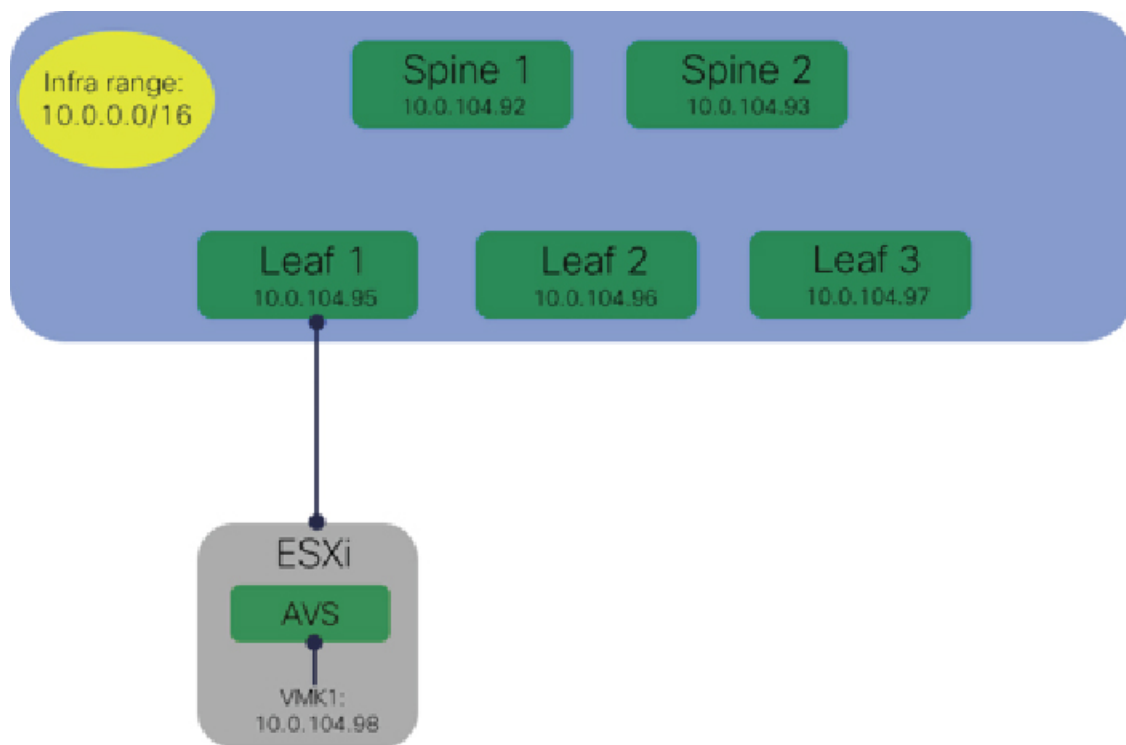
About Fabric Provisioning

Fabric Infrastructure IP Range Recommendations

When provisioning an Application Policy Infrastructure Controller (APIC), one of the required data points during the setup stage is an IP address range for infrastructure addressing inside of the fabric. This is primarily for the purposes of allocating tunnel endpoint (TEP) addresses. The default value for this range is 10.0.0/16. Although technically you can select a range that overlaps with other subnets in the network, you should choose a unique range for this infrastructure range.

Frequently, the infrastructure IP address range must be extended beyond the Cisco Application Centric Infrastructure (ACI) fabric. For example, when the Application Virtual Switch (AVS) is used, a VMK interface is automatically created that uses an address from the infrastructure range as shown in the following figure:

Figure 1: Extending the Infrastructure IP Range Beyond the ACI Fabric



500801

If the infrastructure range overlaps with other subnets elsewhere in the network, routing problems might occur.

The minimum supported subnet size in the recommended three APIC scenario is /22. The number of addresses required depends on a variety of factors, including the number of APICs in your fabric, the number of leaf and spine nodes, the number of AVS instances, and the number of virtual port channels required. To avoid issues with address exhaustion, you should consider allocating a /16 or /17 range if possible.



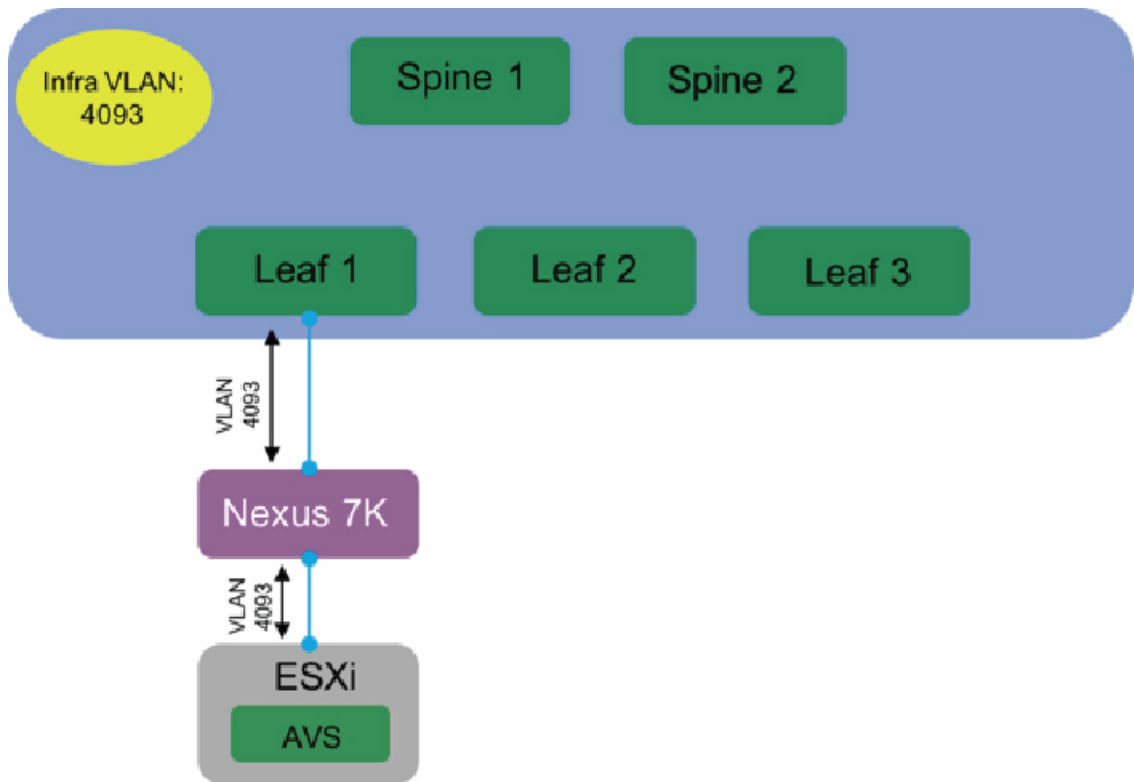
Note When considering the preceding requirements, remember that changing either the infrastructure IP address range or the VLAN after initial provisioning is not possible without rebuilding the fabric.

Fabric Infrastructure VLAN Recommendations

During fabric provisioning, the system requires a VLAN number to be used as the infrastructure VLAN. This VLAN is used for control communication as a reserved overlay VLAN between the fabric nodes (leaf, spine, and APIC controllers) to bring up the fabric. This VLAN is hard coded on the fabric nodes.

If possible, this VLAN number should be unique within the network. In a scenario where the infrastructure VLAN is extended outside of the ACI fabric (for example, if using Cisco AVS or OpenStack integration with Opflex), this VLAN might need to traverse other (non-ACI) devices. In that case, be sure that the infrastructure VLAN does not fall within a range that is prohibited on the non-ACI device. The following figure shows an example of the reserved VLAN range within a Cisco Nexus 7000:

Figure 2: Reserved VLAN Range



500802

In many cases, VLAN 3967 is a good choice for the ACI infrastructure VLAN to avoid the issue outlined in the preceding section.

For more information about fabric infrastructure VLAN recommendations, see the *Cisco APIC Getting Started Guide* at the following URL:

<http://www.cisco.com/c/en/us/support/cloud-systems-management/application-policy-infrastructure-controller-apic/tsd-products-support-series-home.html>

Fabric Node ID Recommendation

The fabric node ID is used to form the fabric membership during fabric initialization. It is also used to configure underlay physical policies, such as access policies and fabric policies, within the fabric. Having a good node ID structure is important to ease the management and operation for the ACI fabric.

Below are general guidelines for configuring fabric node IDs:

- Plan the node ID wisely to allow for further growth and expansion.
- Use different node ID ranges for spine switches and leaf switches. For example, the 100 range for spine switches and the 200 range for leaf switches.
- Using different node ID ranges for leaf switches depends on the use case. For example, if leaf switches are categorized in a different functionality, consider using a different range based on the use. For example, the 200 range for border leaf switches and service leaf switches, the 300 range for compute leaf switches, and the 400 range for storage leaf switches.



Note Node IDs 1 through 29 are reserved for APICs, which cannot be changed.

When APIC redundancy is configured, you should use IDs 1 to 19 for active APICs and IDs 20 to 29 for standby APICs. This allows for expansion of the fabric.

- When a pair of switches is used for the server uplink connectivity using either vPC or active/standby, consider using sequential numbers for the leaf node ID for those switch pairs. For example, node ID 201 for the vPC side A connectivity and node ID 202 for side B. That way, it is easier to configure and easier to manage an upgrade when using maintenance groups.
- If only one ToR switch is deployed, reserve the even leaf ID for future use.



Note Once the fabric node ID is assigned, the ID is difficult to change unless the fabric nodes (spine and leaf) are decommissioned from the fabric and cleanly rebooted.
