Migration Overview

This chapter contains the following sections:

- Prerequisites, page 1
- Existing FabricPath Topology, page 2
- Cisco Dynamic Fabric Automation Topology, page 3
- Traffic Flow Before and After Migration, page 6

Prerequisites

To prepare for migration to the Cisco Dynamic Fabric Automation (DFA) solution, you must meet the following prerequisites.

- Install and configure Cisco Data Center Network Manager 7.0
  - Perform tasks specified in the DCNM 7.0 OVA Installation Guide
  - Perform tasks specified in the DCNM 7.0 Fundamentals Guide
- FabricPath on Spine-Leaf Topology
  - Nexus 7000 spine switches with NX-OS 6.2.(2) images
  - Nexus 6000 border leaf switches with NX-OS 6.02.N2 images
  - Nexus 6000 leaf switches with NX-OS 6.02.N2 images

Note All non-Nexus 6000 boxes must be physically replaced with Nexus 6000 boxes with NX-OS 7.0(0)N1(1).

  - Nexus 1000v Series virtual switches at the virtual machine access layer

Existing FabricPath Topology

The existing FabricPath topolgy from which you are migrating includes:

- An access layer with FabricPath-enabled VPCpath peers (VPC+)
- Layer 3 aggregation layer-only connection to Spine layers
- Two peers of Layer 3 boxes
- Switched Virtual Interfaces (SVI) on only one set of VPC+ peers
- HSRP running in local Layer 3 VLANS

Figure 1: Figure: Pre-migration Fabric Topology

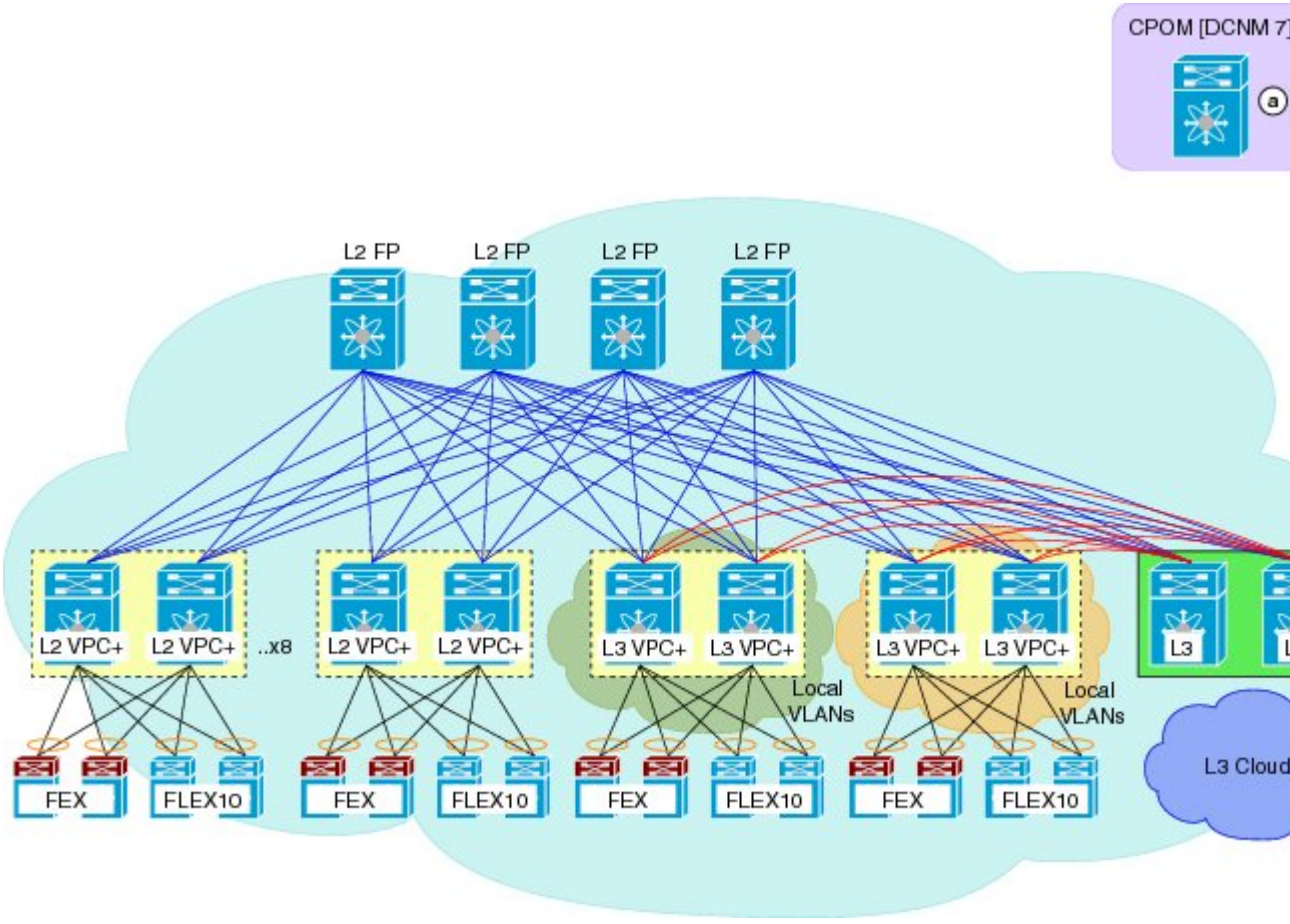

Cisco Dynamic Fabric Automation Topology

An illustration of the Cisco Dynamic Fabric Automation (DFA) topology is shown in the following figure.

Figure 2: Cisco DFA topology

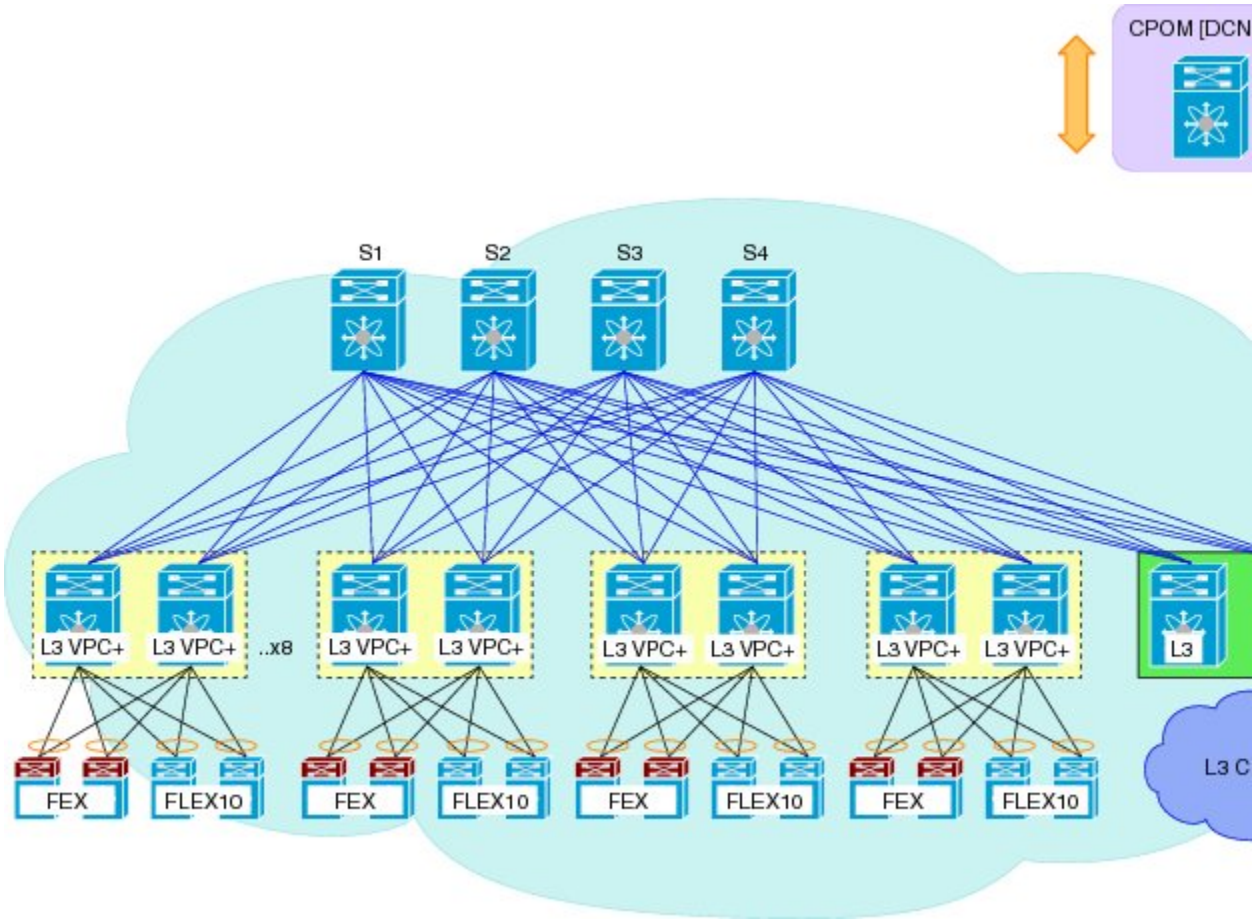

You can structure your Cisco Dynamic Fabric Automation (DFA) topology with two distinct fabrics:

- Fabric with a mix of Nexus 5000 and Nexus 6000 leaves

- Fabric with only Nexus 6000 leaves

Figure 3: DFA Fabric with a mix of Nexus 5000 and Nexus 6000 leaves

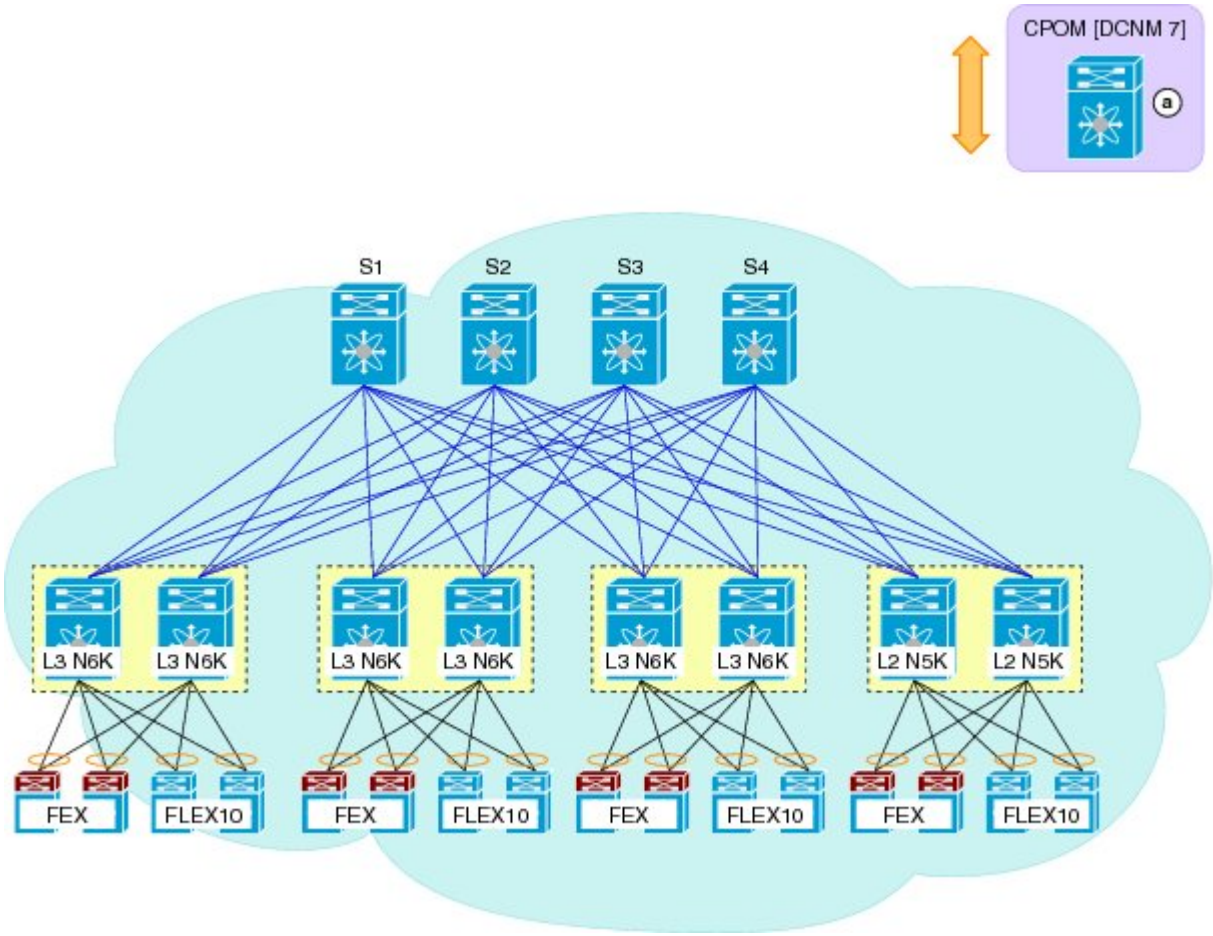

Figure 4: DFA Fabric with only Nexus 6000 leaves

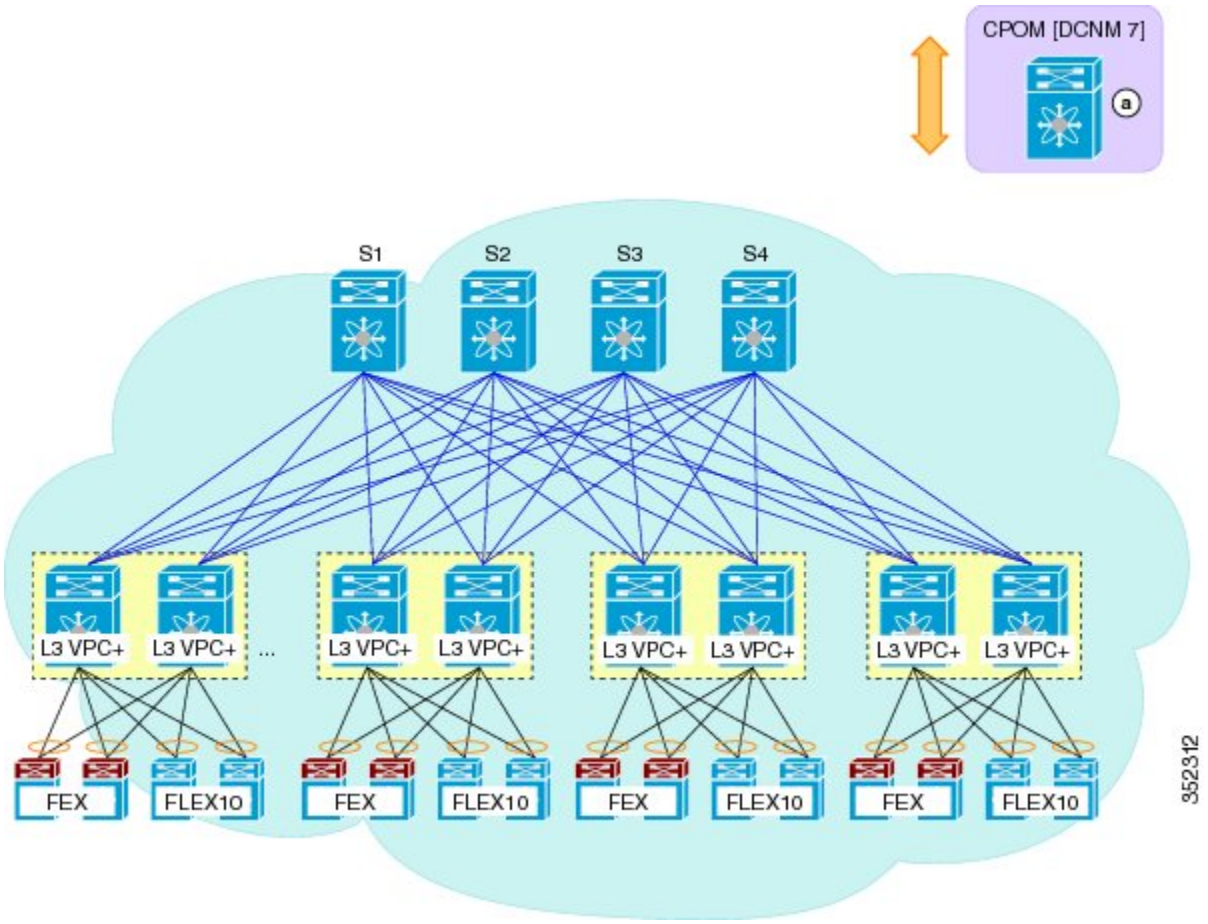

The Cisco DFA fabric with both Nexus 5000 and 6000 leaves includes the following:

- Nexus 5000 remains as Layer 2
- Spine switches that can forward both 1q and 2q traffic, encapsulated in a FabricPath header
- VLAN/SVI distinctions:
  - On a Nexus 5000, the VLAN/SVI is non-Segment ID-enabled across all Cisco DFA leaves running anycast gateway mode on Nexus 6000 leaves. Border leaf runs HSRP/VRRP as well as anycast gateway
  - On a Nexus 6000, the VLAN/SVI is Segment ID-enabled. The forwarding mode can be either proxy or anycast gateway.
  - Multicast will continue to run in the legacy multicast mode. Cisco DFA multicast should not be turned on.

The DFA fabric with only Nexus 6000 leaves includes the following:

- Nexus 6000 leaves running either Anycast Gateway mode or Proxy Gateway mode
- Spine switches that can forward both 1q and 2q traffic, encapsulated in a FabricPath header
- VLANS that can be Segment ID-enabled

Traffic Flow Before and After Migration

As a result of changes to the topology and configuration of switches, traffic flow is optimized after the migration. Differences in traffic flow are shown in the following set of figures:

Prior to migration, Inter-VLAN traffic from Host 1 on VLan10 goes through single Layer 3 hops up through the spine to get to host 2 on Vlan11.

Figure 5: Figure: pre-migration inter-vlan single hop

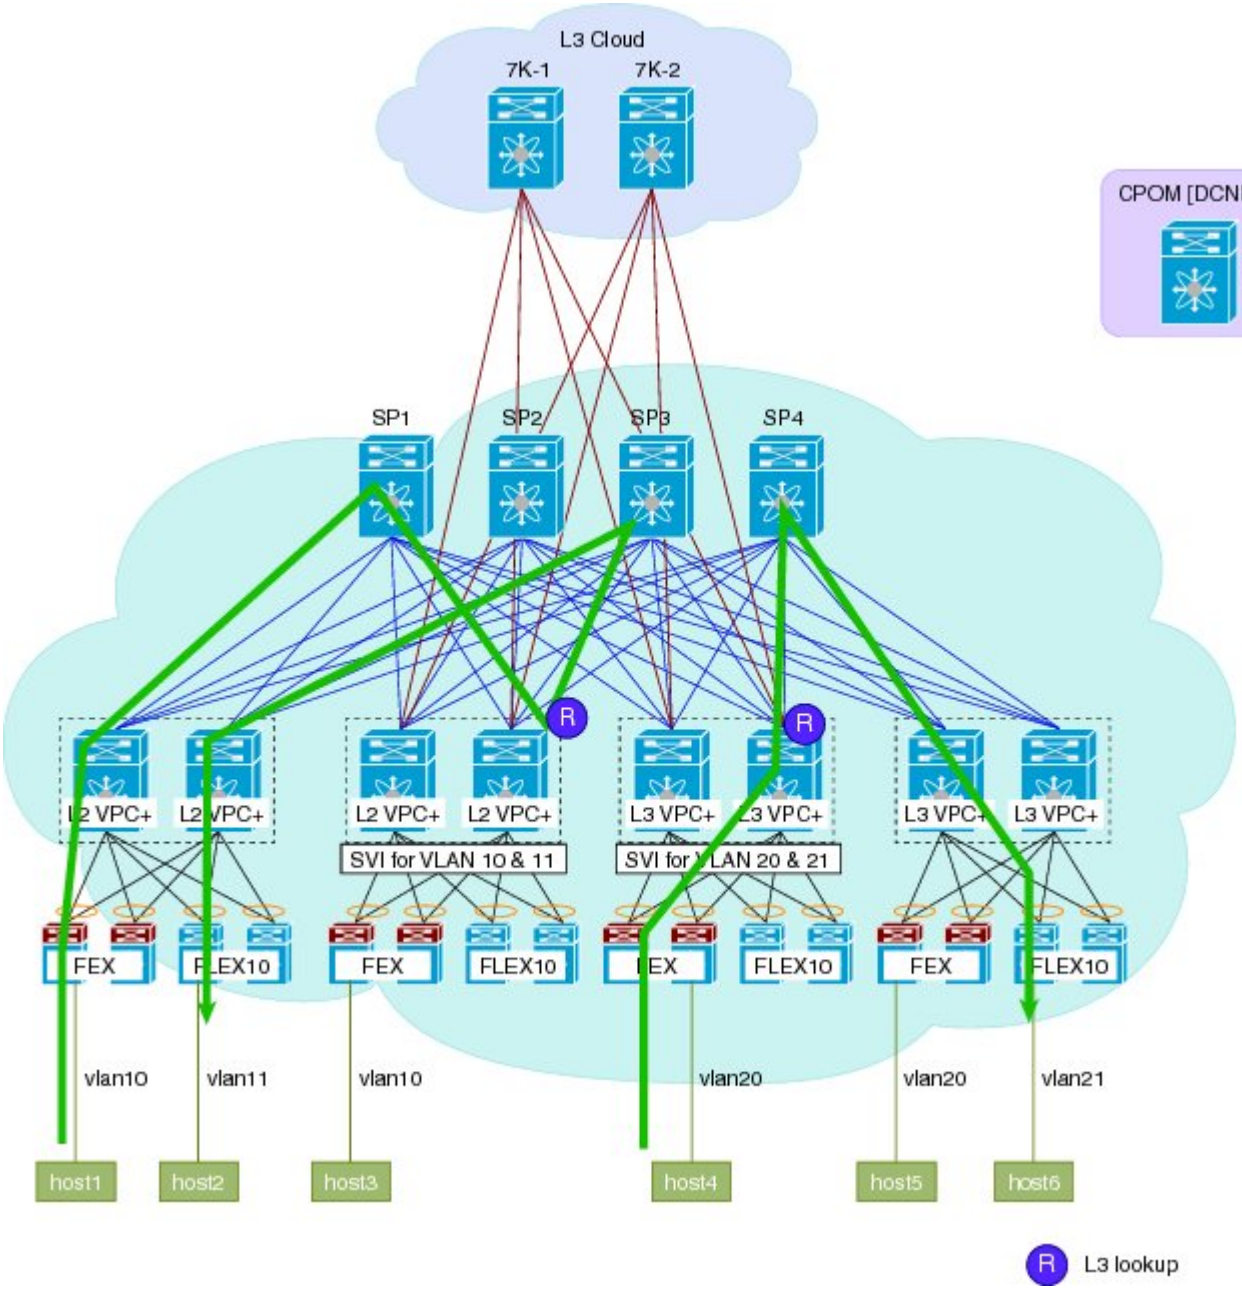

After migration to the Cisco DFA fabric, inter-Vlan traffic from Host 1 on Vlan 10 takes a single hop through a single leaf node, where a Layer 3 lookup is performed and traffic is routed to host 2 on Vlan 11. Border Leafs start to respond to address resolution protocol (ARP) with anycast gateway media access control (MAC).

Figure 6: Figure: post-migration inter-vlan single hop

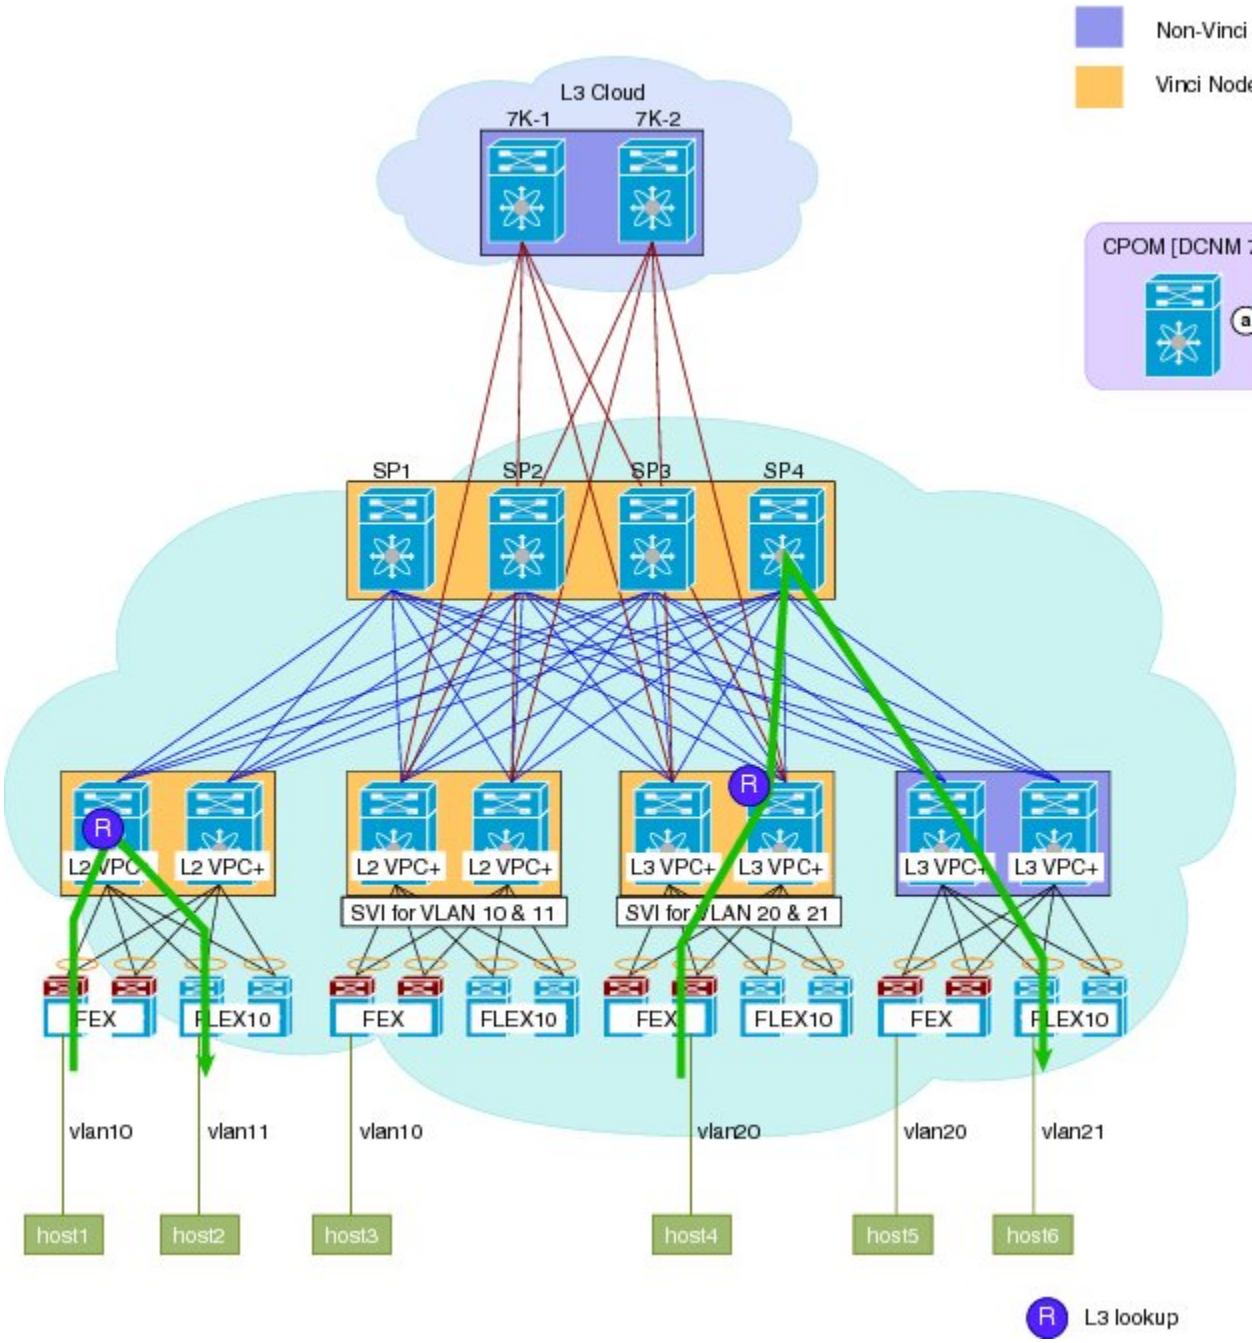

Prior to migration, traffic going from host1 on vlan10 to host 5 on vlan20 takes multiple Layer 3 hops up to the Nexus 7000 Layer 3 and a series of Layer 3 lookups.

Figure 7: Figure: inter-vlan trafic multiple l3 hops x

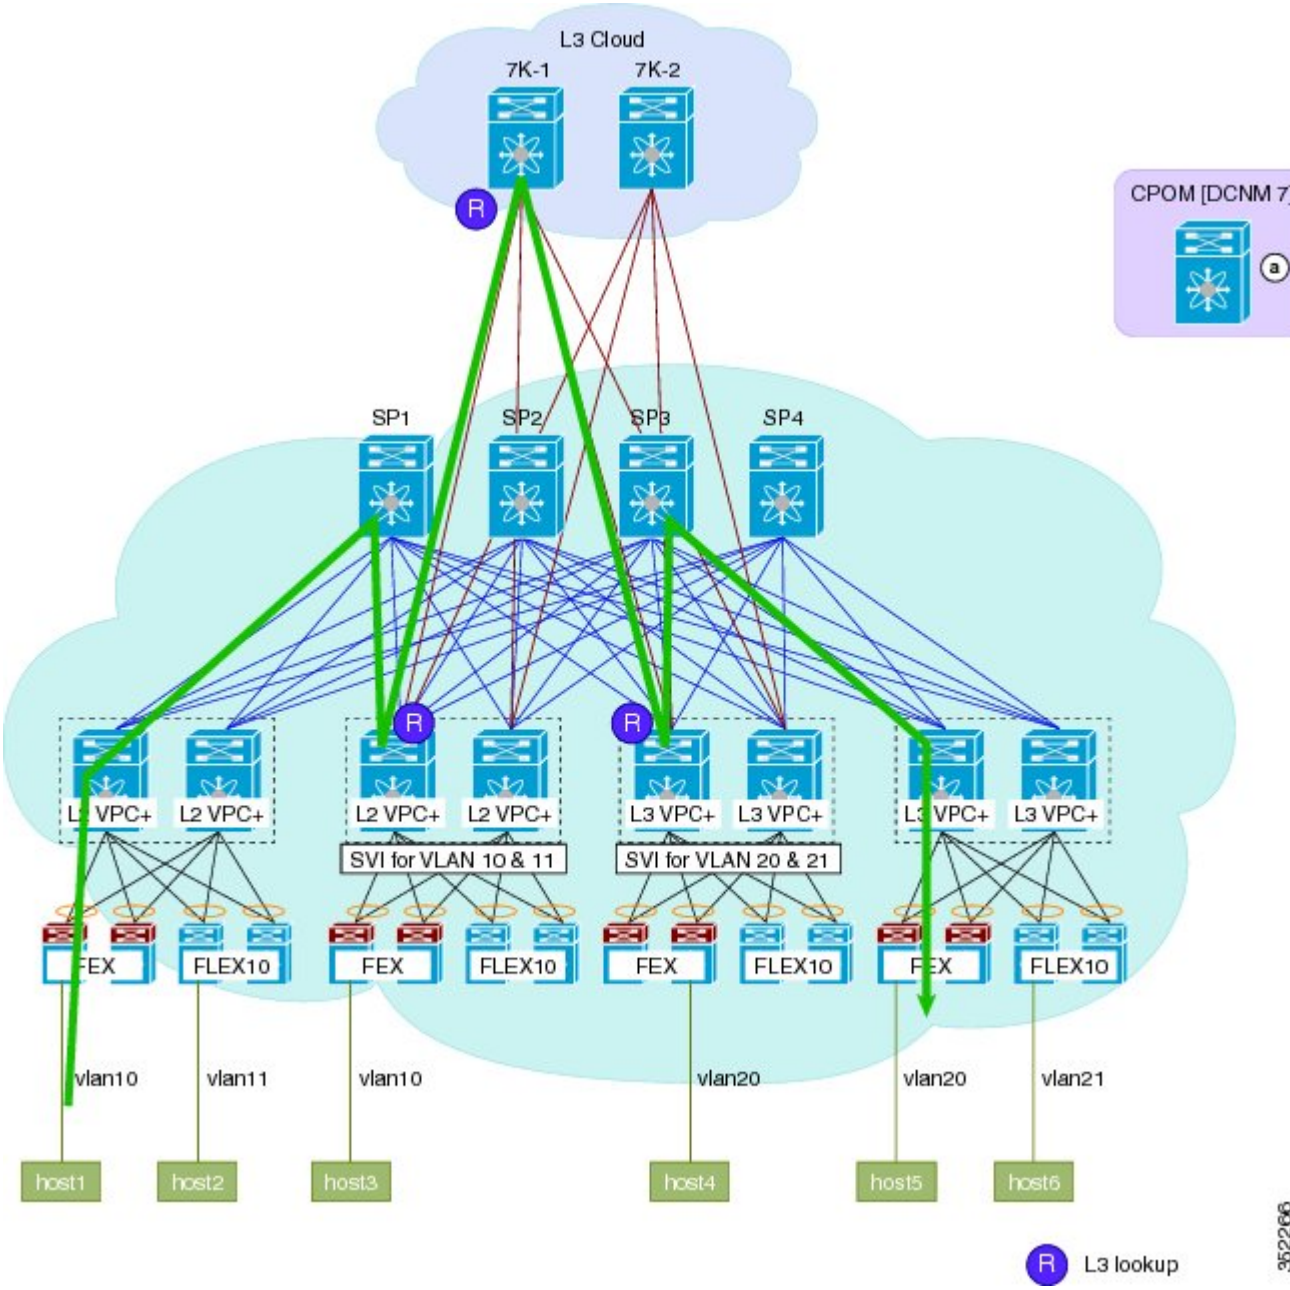

After migration, unicast traffic going from host 1 on vlan 10 to host 4 on vlan 20 takes fewer Layer 3 lookups at the leaf-level, and direct forwarding occurs between border leaf pairs through the spine without going to the Nexus 7000.

Figure 8: Figure: post-migration Unicast traffic flow

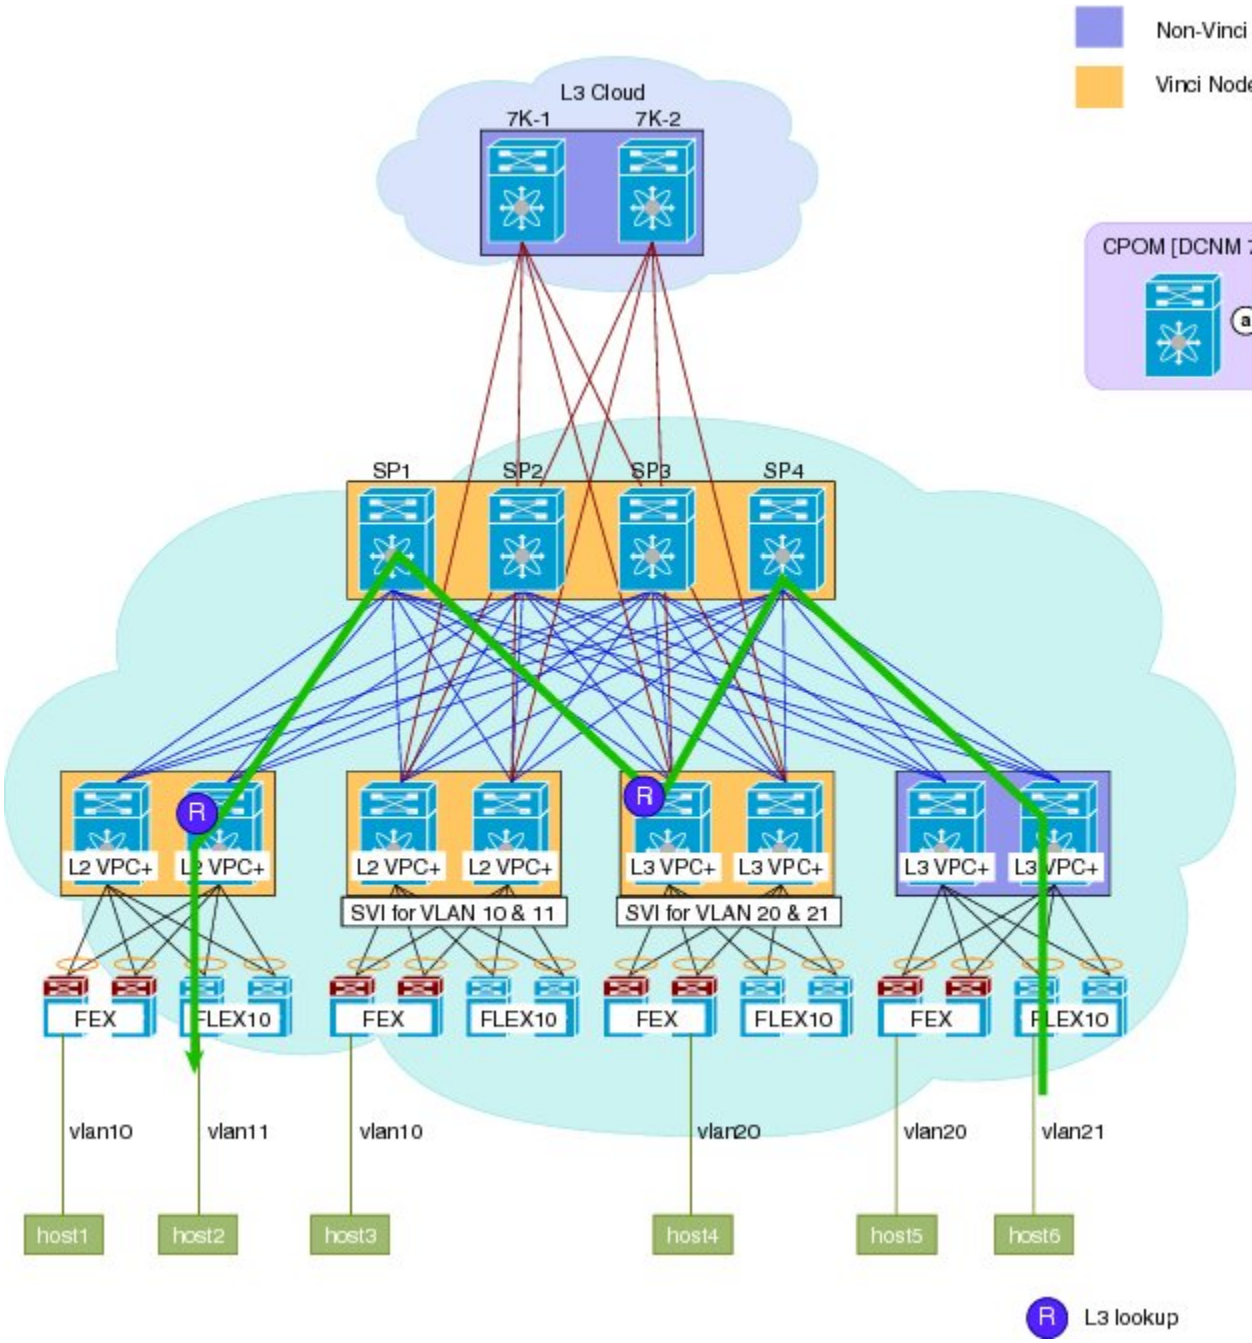

Another illustration of post-migration unicast traffic flow.

Figure 9: Figure: post-migration unicast traffic flow

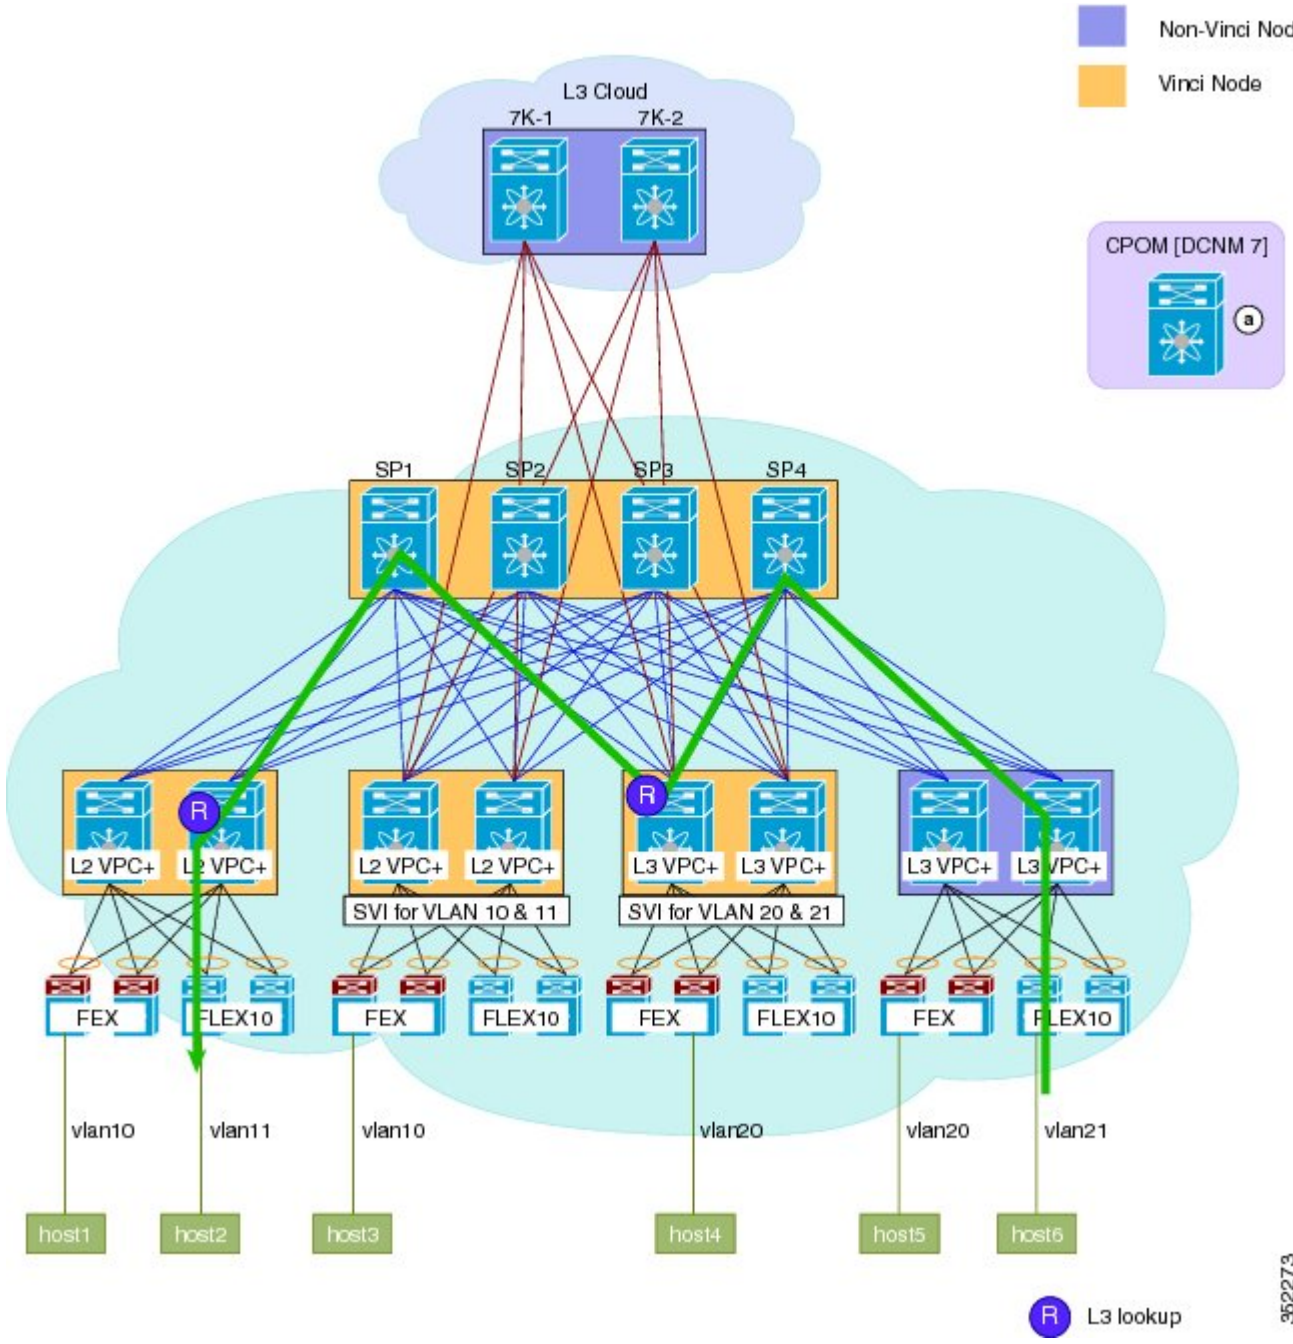

North-South traffic remains unchanged after the migration and requires two Layer 3 lookups before reaching the Layer 3 cloud

Figure 10: Figure: North South Traffic Flow

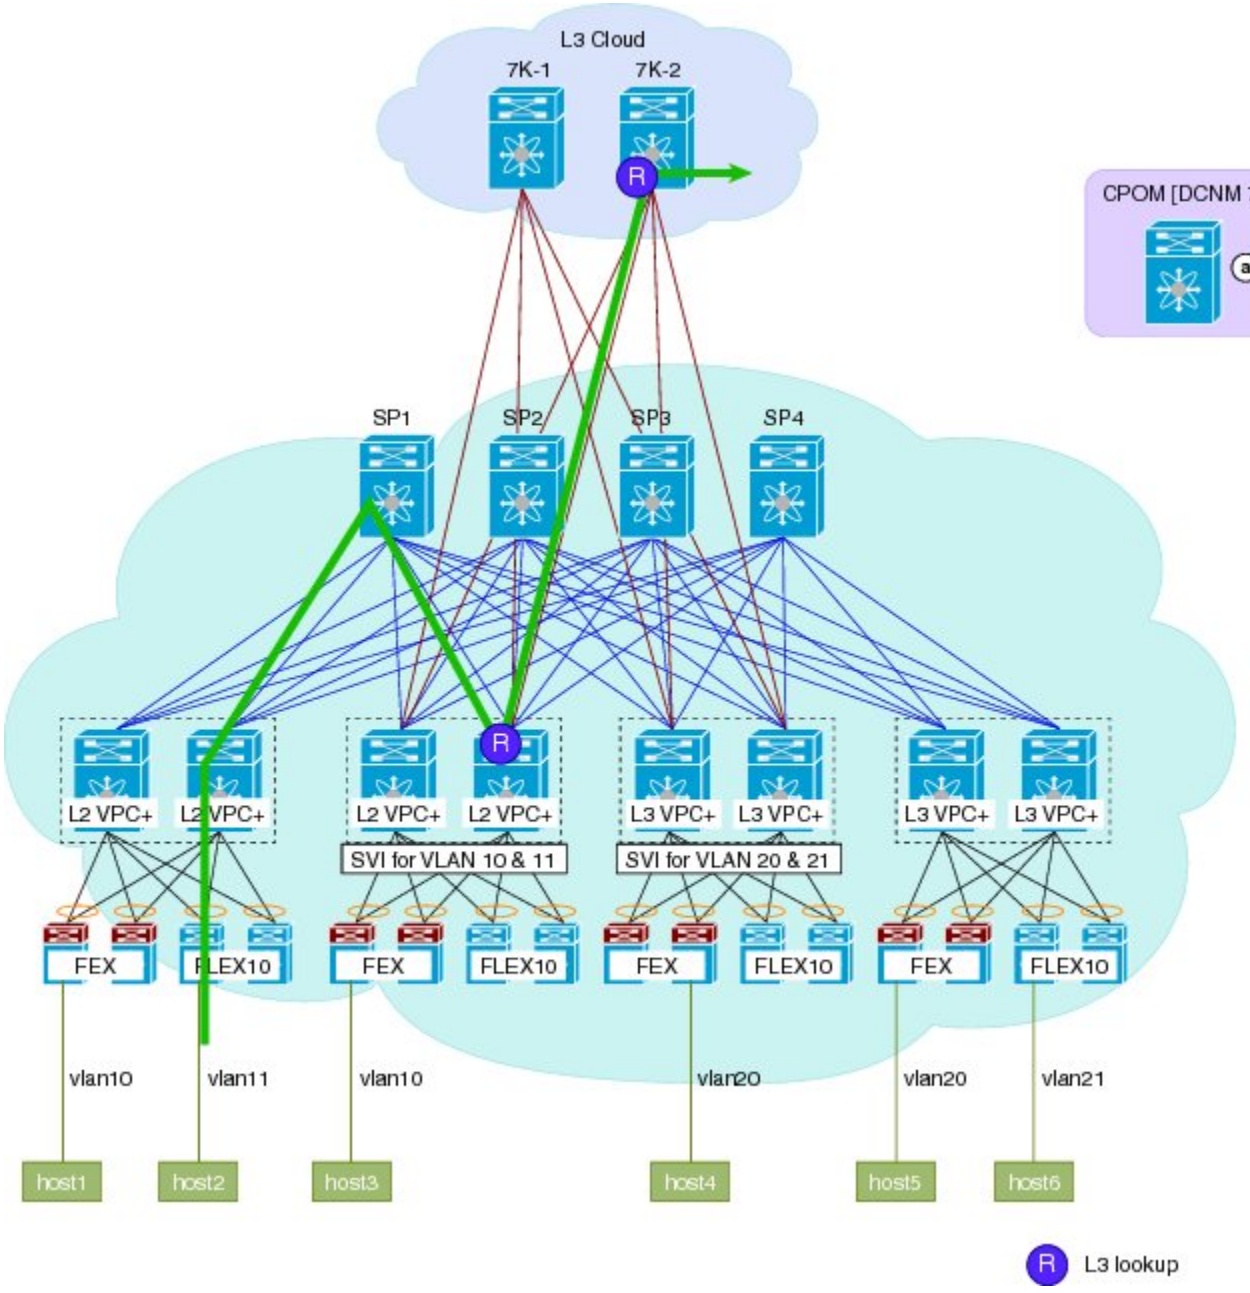

PIM-SM and multicast replication behavior is the same as a non-FabricPath topology. Layer 2 multicast forwarding follows a pruned FabricPath tree. Internet Group Management Protocol (IGMP) is propagated to all FabricPath nodes via Intermediate-system to intermediate-system (ISIS).

Figure 11: Figure: Pre-migration Multicast Traffic Flow

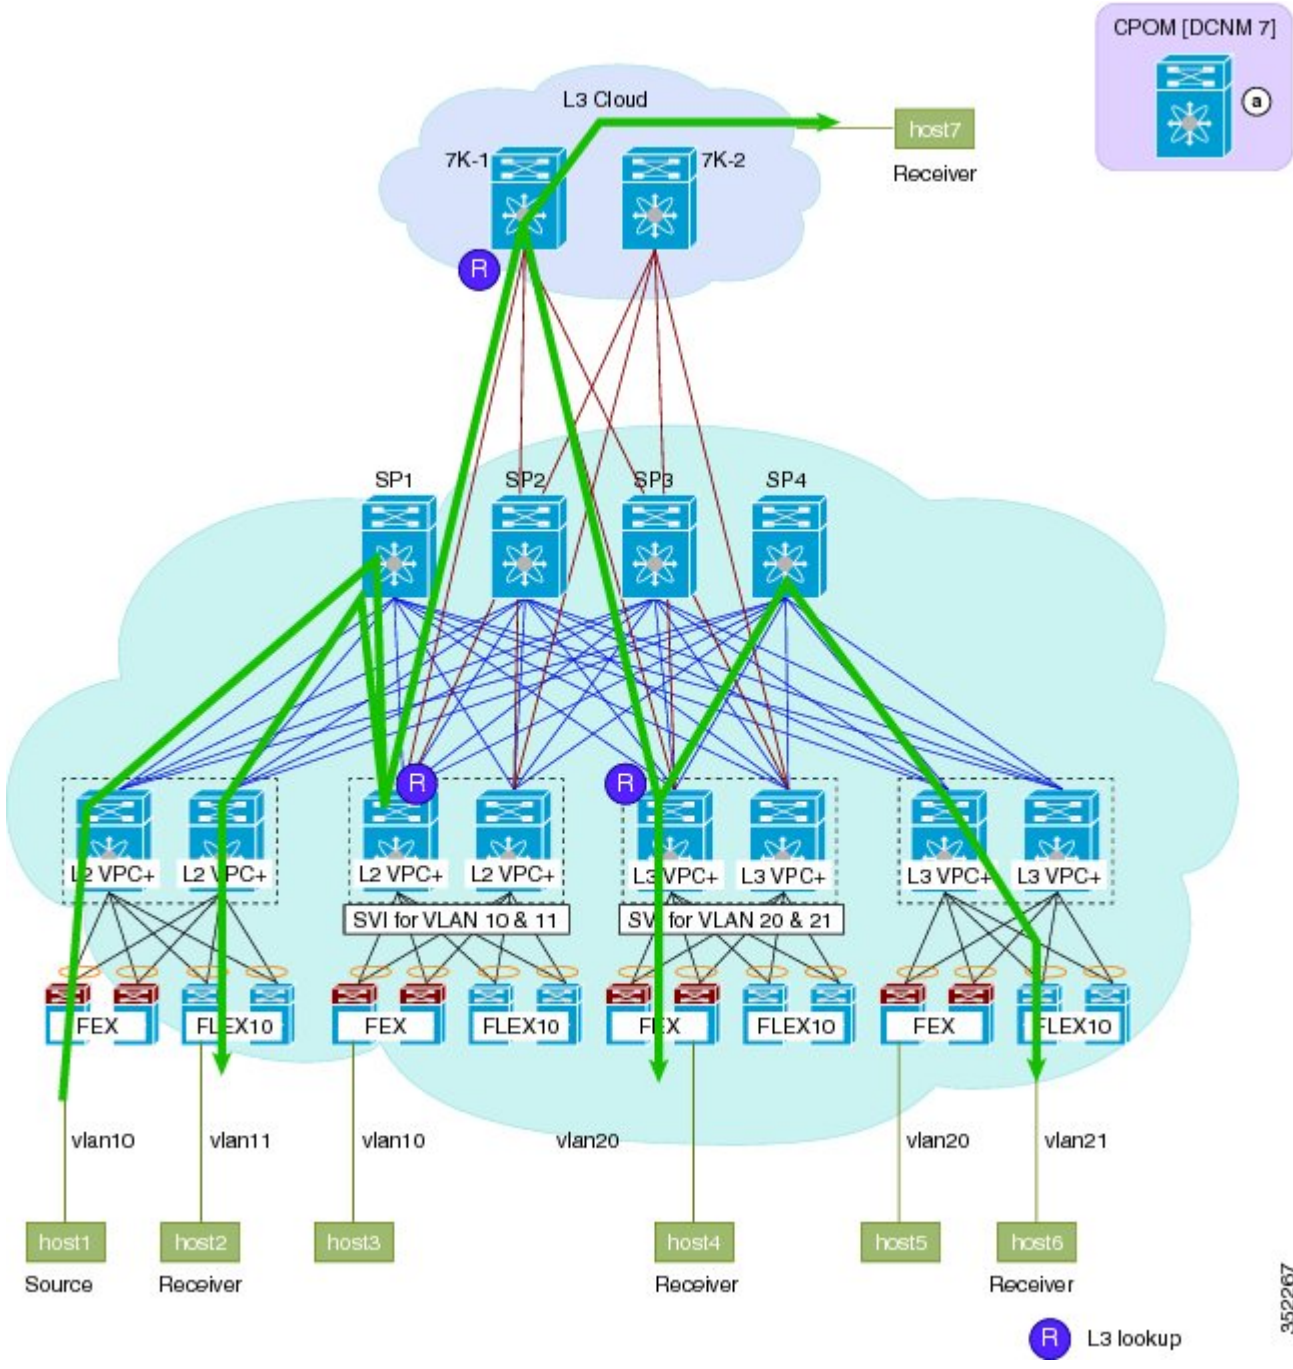

IF there is a Nexus 5000 in the topology, legacy multicast will continue to run.

Figure 12: Figure: post-migration multicast traffic flow