Campus Fabric

This chapter contains the following sections:

Campus Fabric

Overview of Campus Fabric

The Campus Fabric feature provides the basic infrastructure for building virtual networks based on policy-based segmentation constructs. Fabric overlay provides services such as host mobility and enhanced security, which are in addition to normal switching and routing capabilities.

This feature enables a LISP-based Control Plane for VXLAN Fabric. This feature is supported only on the M3 module. The Cisco Nexus 7700 Series with M3 Module acts as a fabric border which connects traditional Layer 3 networks or different fabric domains to the local fabric domain, and translates reachability and policy information from one domain to another. However inter-fabric connectivity on a same switch is not supported.

Cisco Nexus 7700 is positioned as a fabric border node in the Campus Fabric architecture.

Figure 1: Campus Fabric Architecture

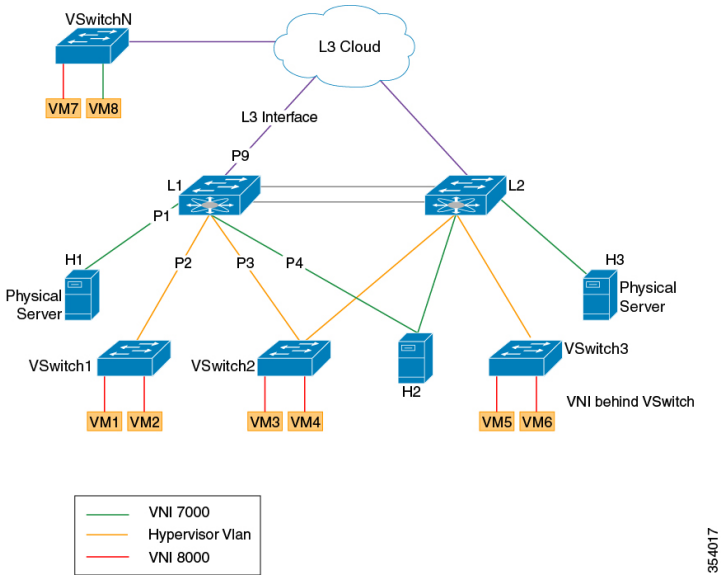

The key elements of the Campus fabric architecture are explained below.

Campus Fabric : The Campus Fabric is an instance of a "Network Fabric". A Network Fabric describes a network topology where data traffic is passed through interconnecting switches, while providing the abstraction of a single Layer-2 and/or Layer-3 device. This provides seamless connectivity, independent of physical topology, with policy application and enforcement at the edges of the fabric. Enterprise fabric uses IP overlay, which makes the network appear like a single virtual router/switch without the use of clustering technologies. This logical view is independent of the control plane used to distribute information to the distributed routers or switches.

Fabric Edge Node : Fabric edge nodes are responsible for admitting, encapsulating/decapsulating and forwarding traffic to and from endpoints connected to the fabric edge. Fabric edge nodes lie at the perimeter of the fabric and are the first points for attachment of the policy. It is to be noted that the endpoints need not be directly attached to the fabric edge node. They could be indirectly attached to a fabric edge node via an intermediate Layer-2 network that lies outside the fabric domain.

Traditional Layer-2 networks, wireless access points or end-hosts are connected to Fabric Edge nodes.

Fabric Intermediate Node: Fabric intermediate nodes provide the Layer-3 underlay transport service to fabric traffic. These nodes are pure layer-3 forwarders that connect the Fabric Edge and Fabric Border nodes.

In addition, Fabric intermediate nodes may be capable of inspecting the fabric metadata and could apply policies based on the fabric metadata (not mandatory). However, typically, all policy enforcement is at the Fabric Edge and Border nodes.

Fabric Border Node : Fabric border nodes connect traditional Layer-3 networks or different fabric domains to the Campus Fabric domain.

If there are multiple Fabric domains, the Fabric border nodes connect a fabric domain to one or more fabric domains, which could be of the same or different type. Fabric border nodes are responsible for translation of context from one fabric domain to another. When the encapsulation is the same across different fabric domains,

the translation of fabric context is generally 1:1. The Fabric Border Node is also the device where the fabric control planes of two domains exchange reachability and policy information.

APIC-EM Controller : This is the SDN controller developed by the Enterprise Networking Group. This controller serves both Brownfield and Greenfield deployments. Campus Fabric service will be developed on the APIC-EM controller. This service will drive the management and orchestration of the Campus Fabric, as well as the provision of policies for attached users and devices.

Fabric Header: Fabric header is the VXLAN header which carries the segment ID(VNI) and user group(SGT). SGT is encoded in the reserved bits of the VXLAN header.

Cisco Catalyst 3000 is positioned as the fabric edge and Cisco Nexus 7700 is positioned as the fabric border in this architecture. LISP is the control plane in the campus fabric architecture and it programs the VXLAN routes. LISP is enhanced to support VXLAN routes for Campus Fabric architecture.

The following features are supported on Cisco Nexus 7700:

- LISP control plane pushing VXLAN routes
- VXLAN L3GW (VRF-lite Hand-off)
- Optimal Tenant L3 Multicast (ASM/Bidir/SSM) based on LISP Multicast (Ingress Replication over unicast core)
- IS-IS as underlay
- TTL propagation

VXLAN Encapsulation for Layer-3 LISP Configuration

This section summarizes only the steps that are used for configuring LISP for a hand-off from VXLAN on the border spine or border leaf switch.

Figure 2: Overall Topology of Campus Fabric

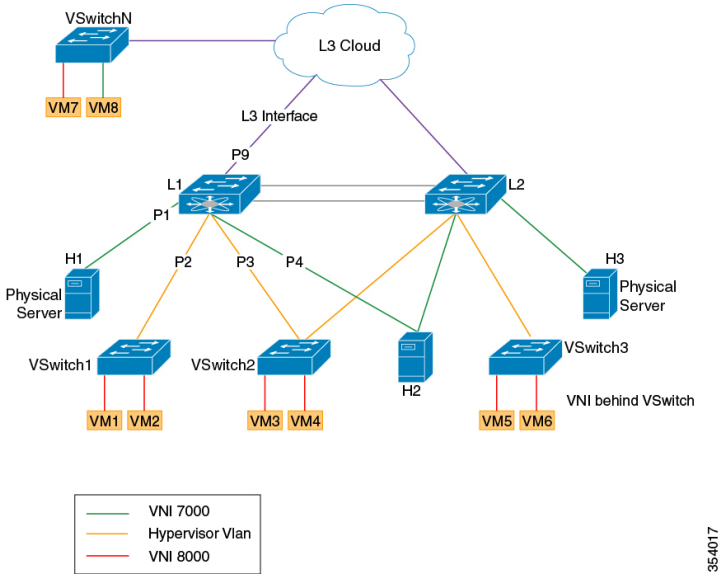


```
feature-set fabric
hostname PxTR1

feature telnet
feature bgp
feature pim
feature isis
feature lisp
feature interface-vlan
system bridge-domain 100
feature nv overlay
feature vni
vni 5000

vlan 1
bridge-domain 100

route-map LISP-RMAP permit 10
bridge-domain 100

/* UNDERLAY VRF*/
vrf context core
  description "UNDERLAY VRF"
  ipv6 lisp itr-etr
  ip lisp itr-etr
  ipv6 lisp itr map-resolver 9.9.9.9
  ip lisp itr map-resolver 9.9.9.9
  ip lisp etr map-server 9.9.9.9 key 3 a97b0defe7b8ff70
  ip lisp multicast
  lisp encapsulation vxlan

/*  OVERLAY VRF */
vrf context vrf5000
```

```
  description "OVERLAY VRF "
  vni 5000
  ip pim rp-address 200.1.2.1 group-list 225.0.0.0/24
  ip lisp proxy-itr 10.1.1.1
  ip lisp proxy-etr
  lisp instance-id 5000
  ip lisp locator-vrf core
  ip lisp map-cache 100.0.1.0/24 map-request
  ip lisp multicast
  lisp encapsulation vxlan
  address-family ipv4 unicast
    route-target import 3:3
    route-target export 1:1

bridge-domain 100
  member vni 5000

interface Bdi100
  description "BDI in OVERLAY vrf"
  no shutdown
  vrf member vrf5000
  no ip redirects
  ip forward
  ip pim sparse-mode

interface nve1
  no shutdown
  source-interface loopback0
  host-reachability protocol lisp
  member vni 5000 associate-vrf

interface Ethernet1/5
  description PxTR1 to SPINE1 link(UNDERLAY VRF)
  vrf member core
  ip address 10.1.1.1/24
  isis circuit-type level-1-2
  ip router isis 100
  ip pim sparse-mode
  no shutdown

interface Ethernet5/1
  no shutdown

interface Ethernet5/1.1
  description PxTR1 to CORE vrf vrf5000(OVERLAY VRF)
  encapsulation dot1q 2
  vrf member vrf5000
  ip address 80.1.1.1/24
  ip pim sparse-mode
  no shutdown

interface loopback0
  description "Source Locator loopback"
  vrf member core
  ip address 1.1.1.1/32
  isis circuit-type level-1-2
  ip router isis 100
  ip pim sparse-mode

interface loopback100
  Description "OVERLAY VRF loopback"
  vrf member vrf5000
  ip address 111.1.1.1/32
  ip pim sparse-mode

/*  IGP on the  UNDERLAY VRF */
router isis 100
  net 49.0001.1111.1111.1111.00
  vrf core
    net 49.0001.1111.1111.1111.00
  vrf vrf5000

/* BGP neighbor towards the CORE */
```

```
router bgp 100
  description "
  router-id 12.12.12.13
  vrf vrf5000
    address-family ipv4 unicast
      redistribute lisp route-map LISP-RMAP
      redistribute direct route-map LISP-RMAP
    neighbor 80.1.1.2 remote-as 100
      Description "BGP neigbor to the CORE Router"
      address-family ipv4 unicast
      address-family ipv6 unicast
```

SGT Propagation, Termination, and Generation

The security group tag (SGT) allows the network to enforce the access control policy by enabling the endpoint device to act upon the SGT to filter traffic.

At the ingress point, traffic from the source is tagged with an SGT containing the security group number of the source entity. The SGT is propagated with the traffic across the domain. At the egress point, an egress device uses the source SGT and the security group number of the destination entity to determine which access policy to apply from the security group access control lists (SGACL) policy matrix.

The least significant 16 bits in the reserved field of the VXLAN header is used to carry the SGT information.

For traffic ingressing the site from internet a mechanism is needed to classify the packets as internet packets and drive SGT based on the classification. This SGT is used in the reserved field of the VXLAN header during VXLAN encapsulation.

For traffic egressing the site the SGT field should be used from the reserved field during VXLAN decapsulation and policy enforcement can be done based on the sg tag. This is where M3 module acts as a PETR. This is enabled using the**ip lisp sgt** command.

Multicast Head-end Replication

Head-end replication for LISP multicast over a unicast core is supported on M3 modules.

Head-end replication accomplishes the need of a multicast transport for Overlay Transport Virtualization (OTV) control plane communications. Multicast transport is used to let a single OTV update or packets to reach all other OTV devices using a specific multicast group address across domains.

LISP Multicast configuration on an ETR or ITR is covered in the "VXLAN Encapsulation for Layer-3 LISP Configuration" section described above.

TTL Propagation

TTL (Time-to-Live) is a setting for each DNS record that specifies how long a resolver should cache the DNS query before the query expires and a new query needs to be made.

TTL propagation from the inner header to the outer header during VXLAN encapsulation is done based on a CLI. On enabling this CLI, the TTL propagation will be disabled from the inner header to the outer header during encapsulation. This is enabled using the **lisp disable-ttl-propagate** command.

Feature History for Campus Fabric

This table lists the release history for this feature.

Table 1: Feature History for Campus Fabric

Feature Name	Releases	Feature Information	
Campus Fabric	7.3(1)D1(1)	This feature was introduced.	