Tenant Configuration

- Tenant Configuration, on page 1

Tenant Configuration

Dynamic Fabric Automation provides touchless tenant provisioning on the desired leaf nodes using auto-configuration. This chapter describes a set of configuration commands that are required for provisioning a tenant. A tenant configuration involves setting up both sides of the network, the server side and the fabric side of the network. All the commands described here are part of the configuration profiles defined in the Profiles Database of the Cisco Prime DCNM. Refer to configuration profiles on the Cisco Prime DCNM for a predefined end-host auto-configuration profile, defaultUniversalNetworkEfProfile, vrf-common-universal-profile, and the vrf-tenant-profile which is pre-configured on the leaf switch.

Tenant Configuration (Server-side)

The server-side configuration at a high level involves:

Multi-tenancy lite version

1. Creating a VLAN for the servers that are attached to the tenant. The server can be either a virtual or physical server. This VLAN is local to the leaf nodes to which the server is connected and is in mode FabricPath.
2. Specifying a globally identifiable segment ID, which will be used to send traffic to these servers over the fabric.
3. Creating a server facing SVI to provide the pervasive first hop anycast gateway.
4. Choosing the right fabric mode for the IP traffic; that is, if an enhanced forwarding using the proxy-gateway is required or a traditional forwarding using the anycast-gateway option.

A sample configuration is as follows:

1. Configure terminal.
2. Configure the VLAN in mode FabricPath for attaching the workloads.
3. Specify that the workload is a part of a global segment ID, for example 10000.
4. Specify that the global segment ID is available over the FabricPath network.

5. Specify the distributed anycast gateway and the IP subnet that this tenant's workload belongs to.
6. Specify that this IP network that is created is a part of tenants VRF.
7. Specify the anycast gateway IP address.
8. Specify the anycast gateway IPv6 address.
9. Specify the mode for anycast distributed gateway.
10. Use this command if the distributed anycast gateway mode is a proxy gateway as there will be virtual machines under the same leaf, which are within the same subnet.
11. Administratively enable the interfaces.

```
configure terminal
 vlan 100
 vn-segment 10000
 mode fabricpath
 interface vlan 100
  vrf member xyz
  ip address 10.1.1.1/24
  ipv6 address 10:1:1::1/64
  fabric forwarding mode [anycast-gateway | proxy-gateway]
  no ip redirects
  no ipv6 redirects
  no shutdown
```

Multi-tenancy full version

1. Creating a bridge-domain for the servers that are attached to the tenant. The server can be either a virtual or physical server. This bridge-domain is local to the leaf nodes to which the server is connected and is in mode FabricPath.
2. Specifying a globally identifiable segment ID, which will be used to send traffic to these servers over the fabric.
3. Creating a server facing SVI to provide the pervasive first hop anycast gateway.
4. Choosing the right fabric mode for the IP traffic; that is, if an enhanced forwarding using the proxy-gateway is required or a traditional forwarding using the anycast-gateway option.

A sample configuration is as follows:

1. Configure terminal.
2. Configure the bridge-domain for attaching the workloads.
3. Specify that the workload is a part of a global VNI, for example 10000.
4. Specify that the global VNI is available over the FabricPath network.
5. Specify the distributed anycast gateway and the IP subnet that this tenant's workload belongs to.
6. Specify that this IP network that is created is a part of tenants VRF.
7. Specify the anycast gateway IP address.
8. Specify the anycast gateway IPv6 address.
9. Specify the mode for anycast distributed gateway.

10. Use this command if the distributed anycast gateway mode is a proxy gateway as there will be virtual machines under the same leaf, which are within the same subnet.
11. Administratively enable the interfaces.

```
configure terminal
 bridge-domain 100
  member vni 10000
 vni 10000
 interface bdi 100
  vrf member xyz
  ip address 10.1.1.1/24
  ipv6 address 10:1:1::1/64
  fabric forwarding mode [anycast-gateway | proxy-gateway]
  no ip redirects
  no ipv6 redirects
  no shutdown
```

Tenant Configuration (Fabric-side)

Multi-tenancy lite version

The fabric-side configuration at a high level involves:

1. Creating the required VRF for the tenant.
2. Setting the Layer-3 Segment or the VNI (as shown below) that is used for identifying the routed traffic of the tenant globally.
3. BGP route redistribution to the route reflector.

The following steps detail the configuration steps:

1. Creating the VRF for the tenant— When a Layer-3 segment is configured, configuration corresponding to the Layer-3 segment is performed using a specific configuration profile. This configuration profiles must be available as part of the POAP of the leaf node. IP forward and IPv6 forward in this configuration profile indicates Layer-3 IPv4/IPv6 forwarding is required on this SVI.

Note The keyword 'VNI' in the VRF context helps to trigger this specific config-profile.

A sample configuration is as follows:

```
Configure profile vrf-tenant-profile
   vlan $vrfVlanId
       vn-segment $vrfSegmentId
       mode fabricpath
  interface vlan $vrfVlanId
   vrf member $vrfName
   ip forward
   no ip redirects
   ipv6 forward
   no ipv6 redirects
   no shutdown
```

2. VRF context tenant— Specify a partition for this tenant. Each tenant will be placed in its own VRF.

3. VNI 100000— Specify the Layer-3 segment ID of the tenant, which is used for routing traffic over the fabric to the different leaf nodes.
4. rd auto— Specify the route distinguisher (RD) value. With fabric, the manual specification of RD becomes unnecessary. With this command a RD value is automatically generated by the BGP process for each VRF, using the fabric control segment IP address and the VRF ID of the VRF.
5. address family ipv4 unicast— Enable the IPv4 address family for BGP.
6. route-target both auto— Specify that the route-targets automatically generated are both exported and imported on this leaf.
7. address family ipv6 unicast— Enable the IPv6 address family for BGP.
8. route-target both auto— Specify that the route-targets automatically generated are both exported and imported on this leaf.
9. router BGP 100— Enable route distribution to the route reflector so that it is available in the other leaf nodes. Whatever is configured under BGP is relevant only for the end-host SVI and not applicable for the core-facing SVI.
10. vrf tenant— Specify the VRF name for which this distribution is configured.
11. address family IPv4 unicast— Enable the IPv4 address family for BGP.
12. redistribute hmm route-map redist-host— Any address resolution protocols that are learnt under an SVI in fabric forwarding mode can be distributed using BGP to the other leaf nodes.

Multi-tenancy full version

The fabric-side configuration at a high level involves:

1. Creating the required VRF for the tenant.
2. Setting the Layer-3 Segment or the VNI (as shown below) that is used for identifying the routed traffic of the tenant globally.
3. BGP route redistribution to the route reflector.

The following steps detail the configuration steps:

1. Creating the VRF for the tenant— When a Layer-3 segment is configured, configuration corresponding to the Layer-3 segment is performed using a specific configuration profile. This configuration profiles must be available as part of the POAP of the leaf node. IP forward and IPv6 forward in this configuration profile indicates Layer-3 IPv4/IPv6 forwarding is required on this BDI.

Note The keyword 'VNI' in the VRF context helps to trigger this specific config-profile.

A sample configuration of vrf-tenant-profile is as follows:

```
Configure profile vrf-tenant-profile
   bridge-domain $vrfbridge-domainId
     member vni $vrfSegmentId
  vni $vrfSegmentId
  interface bdi $vrfbridge-domainId
   vrf member $vrfName
```

```
   ip forward
   no ip redirects
   ipv6 forward
   no ipv6 redirects
   no shutdown
```

A sample configuration of vrf-common-profile is as follows:

```
vrf context $vrfName
vni $vrfSegmentId
rd auto
ip route 0.0.0.0/0 $include_serviceNodeIpAddress
  address-family ipv4 unicast
     route-target import $include_borderLeafRt
  route-target both auto
address-family ipv6 unicast
     route-target import $include_borderLeafRt
  route-target both auto
router bgp $asn
     vrf $vrfName
   address-family ipv4 unicast
    redistribute hmm route-map FABRIC-RMAP-REDIST-HOST
    redistribute direct route-map FABRIC-RMAP-REDIST-SUBNET
    maximum-paths ibgp 2
   address-family ipv6 unicast
       redistribute hmm route-map FABRIC-RMAP-REDIST-V6HOST
    redistribute direct route-map FABRIC-RMAP-REDIST-SUBNET
    maximum-paths ibgp 2
```

2. VRF context tenant— Specify a partition for this tenant. Each tenant will be placed in its own VRF.
3. VNI 100000— Specify the Layer-3 VNI of the tenant, which is used for routing traffic over the fabric to the different leaf nodes.
4. rd auto— Specify the Route Distinguisher (RD) value. With fabric, the manual specification of RD becomes unnecessary. With this command a RD value is automatically generated by the BGP process for each VRF, using the fabric control segment IP address and the VRF ID of the VRF.
5. address family ipv4 unicast— Enable the IPv4 address family for BGP.
6. route-target both auto— Specify that the route-targets automatically generated are both exported and imported on this leaf.
7. address family ipv6 unicast— Enable IPv6 address family for BGP.
8. route-target both auto— Specify that the route-targets automatically generated are both exported and imported on this leaf.
9. router BGP 100— Enable route distribution to the route reflector so that it is available in the other leaf nodes. Whatever is configured under BGP is relevant only for the end-host BDI and not applicable for the core-facing BDI.
10. vrf tenant— Specify the VRF name for which this distribution is configured.
11. address family IPv4 unicast— Enable the IPv4 address family for BGP.
12. redistribute hmm route-map redist-host— Any address resolution protocols that are learnt under an BDI in fabric forwarding mode can be distributed using BGP to the other leaf nodes.