



Virtual Routing and Forwarding

- [Feature Summary and Revision History, on page 1](#)
- [Feature Description, on page 2](#)
- [How it Works, on page 2](#)
- [VRF Feature Configuration, on page 3](#)
- [OAM Support, on page 8](#)

Feature Summary and Revision History

Summary Data

Table 1: Summary Data

Applicable Product(s) or Functional Area	SMF
Applicable Platform(s)	SMI
Feature Default Setting	Disabled – Configuration Required
Related Changes in this Release	Not Applicable
Related Documentation	Not Applicable

Revision History

Table 2: Revision History

Revision Details	Release
Support overlapping addresses for AAA server for PAPN use case	2022.04.0
First introduced.	2020.02.5

Feature Description

In private APN deployments, the same SMF can support multiple PAPNs, requiring authentication and accounting with the enterprise AAA servers. As the AAA servers belong to different mobile virtual network operators (MVNOs), it is possible that their address ranges overlap. SMF uses Virtual Routing and Forwarding (VRF) functionality to support overlapping IP addresses for AAA servers.



Important Overlapping addresses for AAA client is currently not supported.

VRF is a technology that allows multiple instances of a routing table to coexist within the same router at the same time. As the routing instances are independent, the same or overlapping IP addresses can be used without conflicting with each other.

SMF enables configuration of VRF name in the IP pool configuration. The SMF sends IP address chunk details along with the configured VRF name to UPF. UPF maps the IP address to VRF configured within UPF.



Important This feature is not fully qualified in this release. For more information, contact your Cisco Account representative.

How it Works

This section describes how the SMF uses the VRF technology for PAPN use case.

The following are the prerequisites for PAPN support.

- Unique VRF is required for all RADIUS interface types.
- Enable VRF and configure required routes on leaf and spine.
- Each VRF has separate BGP or BFD peering.
- Two additional IPs and nonbonded interfaces are required for BGP peering.
- Unique VLAN or interface is required for RADIUS interface.
- Each logical SMF severing RADIUS interface should not share interface with other logical NF.

VRF Creation

To create the VRF, perform the following steps:

1. Create VRF with default route through CLI configuration and add the default route in VRF routing table. If the VRF is created without gateway, the default route will not be added in the routing table.
2. If BGP peering is associated with the created VRF, the routes will not be learned or published on that BGP peering.

3. If BGP peering is not associated with the VRF created without gateway, the routes will be learned or published on global BGP peering.

VRF Modification

To modify the device or gateway on the existing created VRF, use the following steps:

1. Disassociate the VRF to be modified from endpoint, router, and other applicable configuration.
2. Delete the VRF.



Important Make sure you delete the VRF in running system only.

3. Create VRF with modified device or gateway.
4. Associate newly created VRF with modified device or gateway with endpoint, router, and other configurations as required.

VRF Deletion



Important Make sure you delete the VRF in running system only.

If the VRF is shared across logical SMF, delete the VRF from all logical SMF to completely remove VRF configuration from interface,

Limitations

This feature has the following limitations:

- NAS-IP for authentication and accounting request must be the same as the interface CoA-NAS VIP-IP in the endpoint RADIUS interface CoA-NAS configuration.
- The VRF configured under RADIUS server-group and endpoint configuring VRF must be the same.
- For the DM request, client VRF and the server VRF (CoA-NAS VIP VRF) must match. If there is a mismatch, the DM request is discarded.

VRF Feature Configuration

Configuring VRF Name

Enable VRF name in the configuration of address pool. When “vrf-name” in the IP pool is configured with “@” character, the first part before “@” is treated as “VRF” name which is sent to UPF and the second part after “@” is sent as context name in N4 Association Update message.

When “vrf-name” in the IP pool is configured without “@” character, the entire name is sent as context name (existing behaviour) in N4 Association Update message.

The SMF sends the VRF name in N4 Session Establishment request if it is configured with “@” in the pool in “UE IP VRF” IE.

Use the following configuration to configure the VRF name:

```
config
  ipam instance instance_id
    address-pool pool_name
      vrf-name vrf_name
    end
```

NOTES:

- **ipam**: Enters the IPAM configuration mode.
- **address-pool** *pool_name*: Configures the address pool configuration. *pool_name* must be the name of the address pool.
- **vrf-name** *vrf_name* : Specifies the name of the VRF for the pool.

Configuring VRF for PAPN Support

To configure the VRF for PAPN, use the following sample configuration:

```
config
  vrf name vrf_name gateway ip_address device interface_name
end
```



Important VRF creation and deletion operations are supported. To modify the existing VRF, VRF must be deleted and then added again.

NOTES:

- **vrf name** *vrf_name* —Specify the virtual routing and forwarding (VRF) name of the pool.
- **gateway** *ip_address*—Specify the IP address of the gateway.
- **device** *interface_name*—Specify the interface name. The interface should be a bonded interface.

Configuration Example

The following are example configurations of VRFs with gateway:

```
vrf name papn_vrf_1 gateway 209.165.202.131 device bd2.radius.2161
vrf name papn_vrf_2 gateway 209.165.202.131 device bd2.radius.2162
```

The following are example configurations of VRFs without gateway:

```
vrf name papn_vrf_1 device bd2.radius.2161
vrf name papn_vrf_2 device bd2.radius.2162
```

Associating Endpoint with VRF

To associate the VRF with RADIUS endpoint,, use the following sample configuration:

```
config
  instance instance-id instance-id
    endpoint radius
      interface coa-nas
        vip-ip ip_address [ vip-interface interface_name | vip-port vip_port |
vrf vrf_name]
      end
```

NOTES:

- **vip-ip** *vip_ip*—Specify the IPv4 or IPv6 address for the RADIUS endpoint.
- **vip-port** *vip_port*—Specify the port number of RADIUS endpoint.
- **vip-interface** *vip_interface*—Specify the interface name. Note that this is a bonded interface which is associated with VRF.
- **vrf** *vrf_name*—Specify the VRF name of the pool in RADIUS endpoint interface CoA-NAS configuration.

Configuration Example

The following are example configurations:

```
instance instance-id 1 endpoint radius interface coa-nas vip-ip 209.165.202.131 vip-port
8112 vip-interface bd2.radius.2161 vrf papn_vrf_1

instance instance-id 1 endpoint radius interface coa-nas vip-ip 209.165.202.133 vip-port
8112 vip-interface bd2.radius.2162 vrf papn_vrf_2
```

Configuring VRF in RADIUS Server Group

To configure the VRF in RADIUS server group, use the following sample configuration:

```
config
  profile radius
    server-group group_name
      vrf vrf_name
      server { [ acct private_server_ip_address port ] | [ auth
private_server_ip_addressport] | [ range ] }
    end

  config
    profile radius-dynamic-author
      client client_ipaddress vrf vrf_name
      nas-identifier nas_identifier
      secret aes_encrypted_string
    end
```

NOTES:

- **acct** *private_server_ip_address* *port*—Specify the server used for Accounting requests. This server is private to the specific server-group.

Private servers in the server-group will be given priority over global servers that are associated to the group. If private servers are unreachable or dead, global servers will be selected to send authentication or accounting requests.

- **auth** *private_server_ip_addressport*—Specify the server used for Authentication or Authorization requests. This server is private to the specific server-group.
- **range**—Specify the IP address range.
- **client** *client_ip_address*—Specify the RADIUS Dynamic-authorization client configuration.
- **vrf** *vrf_name*—Specify the VRF name to be configured in AAA server group.

If VRF is configured in server-group, it is recommended to configure servers using server-private and not associate the global servers.

- **nas-identifier** *nas-identifierport*—Specify the Dynamic-authorization NAS Identifier.
- **secret** *aes_encrypted_string*—Specify the Dynamic-authorization server shared secret key.

Associating VRF with BGP Peering

To associate VRF with BGP peering, use the following sample configuration:

```
config
  router bgp bgp_name
    interface interface_name
      vrf vrf_name
    end
```

NOTES:

- **interface** *interface_name*—Specify the local BGP interface.
- **vrf** *vrf_name*—Specify the VRF details to be associated with BGP.

Configuration Example

The following is an example of BGP peering configuration with no VRF association:

```
interface enp94s0f0.3921
  bondingInterface enp216s0f0
  bondingInterface enp94s0f0
  neighbor 209.165.202.254 remote-as 65141 fail-over bfd
exit
interface enp94s0f1.3922
  bondingInterface enp216s0f1
  bondingInterface enp94s0f1
  neighbor 209.165.202.254 remote-as 65141 fail-over bfd
exit
```

The following is an example of BGP Peering configuration with association with *papn_vrf_1*.

```
interface enp94s0f0.3923 □ leaf1
  vrf papn_vrf_1
  bondingInterface enp216s0f0
  bondingInterface enp94s0f0
  neighbor 209.165.202.254 remote-as 65141 fail-over bfd
exit
```

```

interface enp94s0f1.3924 □ leaf2
  vrf papn_vrf_1
  bondingInterface enp216s0f1
  bondingInterface enp94s0f1
  neighbor 209.165.202.254 remote-as 65141 fail-over bfd
exit

```

The following is an example of BGP Peering configuration with association with papn_vrf_2.

```

interface enp94s0f0.3925
  vrf papn_vrf_2
  bondingInterface enp216s0f0
  bondingInterface enp94s0f0
  neighbor 209.165.202.254 remote-as 65141 fail-over bfd
exit
interface enp94s0f1.3926
  vrf papn_vrf_2
  bondingInterface enp216s0f1
  bondingInterface enp94s0f1
  neighbor 209.165.202.254 remote-as 65141 fail-over bfd
exit

```

Verifying VRF Configuration for BGP Peering

This section describes how to verify VRF configuration for BGP peering

To view the VRF information, use the **show vrf-info** command.

Following is a sample output of the **show vrf-info** command.

NAME ID	GATEWAY	STATE	TABLE		GATEWAY DEVICE	LINK DEVICE
			IPV6 POD	NAME		
npapn-vrf22 2527	102.1.16.174	true		bgpspeaker-pod-0	bd2.npv22.766	
npapn-vrf78 2588	102.1.18.142	true		bgpspeaker-pod-0	bd2.npv78.822	
npapn-vrf96 2608	102.1.19.46	true		bgpspeaker-pod-0	bd2.npv96.840	
npapn-vrf10 2502	102.1.16.78	true		bgpspeaker-pod-0	bd2.npv10.754	
npapn-vrf47 2554	102.1.17.134	true		bgpspeaker-pod-0	bd2.npv47.791	
npapn-vrf53 2561	102.1.17.182	true		bgpspeaker-pod-0	bd2.npv53.797	
npapn-vrf69 2578	102.1.18.70	true		bgpspeaker-pod-0	bd2.npv69.813	

To view the VRF route information, use the **show vrf-route-info** command.

Following is a sample output of the **show vrf-route-info** command.

```

-----bgpspeaker-pod-0 -----
Vrf      TableId  Route
npapn-vrf109  2512  default via 102.1.19.150 dev bd2.npv109.853 proto 217 metric 217
npapn-vrf109  2512  broadcast 102.1.19.144 dev bd2.npv109.853 proto kernel scope
link src 102.1.19.147
npapn-vrf109  2512  102.1.19.144/29 dev bd2.npv109.853 proto kernel scope link src
102.1.19.147
npapn-vrf109  2512  local 102.1.19.147 dev bd2.npv109.853 proto kernel scope host
src 102.1.19.147
npapn-vrf109  2512  broadcast 102.1.19.151 dev bd2.npv109.853 proto kernel scope

```

```

link src 102.1.19.147
npapn-vrf109      2512      anycast fe80:: dev bd2.npv109.853 proto kernel
npapn-vrf109      2512      local fe80::42a6:b7ff:fe48:5520 dev bd2.npv109.853 proto kernel
npapn-vrf109      2512      fe80::/64 dev bd2.npv109.853 proto kernel metric 256
npapn-vrf109      2512      ff00::/8 dev bd2.npv109.853 proto 3 metric 256
Vrf               TableId   Route
npapn-vrf82       2593      default via 102.1.18.174 dev bd2.npv82.826 proto 217 metric 217
npapn-vrf82       2593      broadcast 102.1.18.168 dev bd2.npv82.826 proto kernel scope link
src 102.1.18.171
npapn-vrf82       2593      102.1.18.168/29 dev bd2.npv82.826 proto kernel scope link src
102.1.18.171
npapn-vrf82       2593      local 102.1.18.171 dev bd2.npv82.826 proto kernel scope host src
102.1.18.171
npapn-vrf82       2593      broadcast 102.1.18.175 dev bd2.npv82.826 proto kernel scope link
src 102.1.18.171
npapn-vrf82       2593      anycast fe80:: dev bd2.npv82.826 proto kernel
npapn-vrf82       2593      local fe80::42a6:b7ff:fe48:5520 dev bd2.npv82.826 proto kernel
npapn-vrf82       2593      fe80::/64 dev bd2.npv82.826 proto kernel metric 256
npapn-vrf82       2593      ff00::/8 dev bd2.npv82.826 proto 3 metric 256
-----bgpspeaker-pod-1 -----
Vrf               TableId   Route
npapn-vrf109      2512      default via 102.1.19.150 dev bd2.npv109.853 proto 217 metric 217
npapn-vrf109      2512      broadcast 102.1.19.144 dev bd2.npv109.853 proto kernel scope
link src 102.1.19.147
npapn-vrf109      2512      102.1.19.144/29 dev bd2.npv109.853 proto kernel scope link src
102.1.19.147
npapn-vrf109      2512      local 102.1.19.147 dev bd2.npv109.853 proto kernel scope host
src 102.1.19.147
npapn-vrf109      2512      broadcast 102.1.19.151 dev bd2.npv109.853 proto kernel scope
link src 102.1.19.147
npapn-vrf109      2512      anycast fe80:: dev bd2.npv109.853 proto kernel
npapn-vrf109      2512      local fe80::42a6:b7ff:fe48:5520 dev bd2.npv109.853 proto kernel
npapn-vrf109      2512      fe80::/64 dev bd2.npv109.853 proto kernel metric 256
npapn-vrf109      2512      ff00::/8 dev bd2.npv109.853 proto 3 metric 256
Vrf               TableId   Route
npapn-vrf82       2593      default via 102.1.18.174 dev bd2.npv82.826 proto 217 metric 217
npapn-vrf82       2593      broadcast 102.1.18.168 dev bd2.npv82.826 proto kernel scope link
src 102.1.18.171
npapn-vrf82       2593      102.1.18.168/29 dev bd2.npv82.826 proto kernel scope link src
102.1.18.171
npapn-vrf82       2593      local 102.1.18.171 dev bd2.npv82.826 proto kernel scope host src
102.1.18.171
npapn-vrf82       2593      broadcast 102.1.18.175 dev bd2.npv82.826 proto kernel scope link
src 102.1.18.171
npapn-vrf82       2593      anycast fe80:: dev bd2.npv82.826 proto kernel
npapn-vrf82       2593      local fe80::42a6:b7ff:fe48:5520 dev bd2.npv82.826 proto kernel
npapn-vrf82       2593      fe80::/64 dev bd2.npv82.826 proto kernel metric 256
npapn-vrf82       2593      ff00::/8 dev bd2.npv82.826 proto 3 metric 25

```

OAM Support

Bulk Statistics Support

The following statistics are updated to support the VRF feature.

- `bgp_outgoing_routerequest_total` - This statistics includes "vrf" label to indicate the total count of successful BGP outgoing routes per VRF.

- `bgp_outgoing_failedrouterrequest_total` - This statistics includes "vrf" label to indicate the total count of failed BGP outgoing routes per VRF.
- `bgp_speaker_bfd_peer_status` - This statistics includes "vrf" label to indicate the BFD peer status.

