



BGP Support on NFVIS

Table 1: Feature History

Feature Name	Release Information	Description
BGP Support on Remote Subnets Over IPsec.	NFVIS 4.4.1	This feature allows the NFVIS system to learn routes that are announced from the remote BGP neighbor and apply the learnt routes to the NFVIS system.
BGP Support Announcing Local Subnets (Route Distribution)	NFVIS 3.10.1	This feature allows you to announce or withdraw NFVIS local routes to the remote BGP neighbor using route distribution.

Border Gateway Protocol (BGP) is the dynamic routing protocol to exchange route information between BGP autonomous systems.

The NFVIS BGP feature works together with remote BGP router. This feature allows NFVIS system to learn routes announced from the remote BGP neighbor and apply the learnt routes to the NFVIS system. This feature also allows you to announce or withdraw NFVIS local routes from the remote BGP neighbor.

Starting from NFVIS 4.4.1 release, NFVIS BGP feature works with the secure overlay feature to learn routes from the BGP neighbor over a secure overlay tunnel. These learnt routes or subnets are added into the NFVIS routing table for the secure tunnel, which makes the routes accessible over the tunnel.

- [Configure BGP on NFVIS, on page 1](#)
- [Route Distribution, on page 4](#)
- [BGP Route Announcement over MPLS or IPsec, on page 5](#)

Configure BGP on NFVIS

A BGP neighbor can be configured using a neighbor IP address or a name string.

If a BGP neighbor is specified using a name string, it must be used in conjunction with the secure overlay `bgp-neighbor-name` field. A BGP session is established over the secure overlay tunnel. If the neighbor name matches the `bgp-neighbor-name` field configured in secure-overlay configuration, then NFVIS will determine the active remote system IP address used for IPsec connection and replace the neighbor name with that IP.

This will establish a BGP neighbor session with that IP address. For more information on how to configure secure overlay with BGP name, see [Secure Overlay and Single IP Configuration](#).

If a BGP neighbor is specified using an IP address which is headend VPN responder's tunnel IP address, which is the same as the IP address of a headend VPN responder tunnel, a BGP session is established over the secure overlay tunnel.

This example shows how to create or update BGP configuration for a neighbor with a specified name string:

```
config terminal
router bgp 200
  neighbor csrbgp remote-as 65000
  commit
```

This example shows how to create or update BGP configuration with a specified neighbor IP address:

```
config terminal
router bgp 200
  neighbor 166.34.121.112 remote-as 65000
  exit
  neighbor 166.35.121.112 remote-as 65000
  commit
```

This example shows how to delete BGP configurations:

```
no router bgp 200
commit
```

The following table provides the syntax description for each parameter in the commands mentioned in the examples above:

Property	Type	Description	Mandatory
as	Uint32	Local BGP AS number	Yes
router-id	IPv4	H.H.H.H: IPv4 address for local system	No
neighbor	list	Neighbor list	Yes
remote-ip	String	IPv4 address or Secure Overlay BGP neighbor name for BGP neighbor system	Yes
remote-as	Uint32	Remote BGP AS number	Yes
description	String	Description of neighbor	No

The following example displays the BGP session details:

```
nfvis# support show bgp

BIRD 1.6.8 ready.
name  proto  table  state  since  info
bgp1  BGP      bgptable  UP      23:53:18  Established
Preference: 100
Input filter: ACCEPT
```

```

Output filter: Accept
Import limit: 15
  Action: restart
Routes:      1 import, 0 exported, 1 preferred
Route change stats:  received  rejected  filtered  ignored  accepted
  Import updates:      1           0           0           0           1
  Import withdraws:    0           0           ---          0           0
  Export updates:      1           1           0           ---          0
  Export withdraws:    0           ---         ---         ---          0
BGP state:      Established
  Neighbour address: 166.34.121.112
  Neighbour AS:      65000
  Neighbour ID:      166.34.121.112
  Neighbour caps:    refresh enhanced-refresh AS4
  Session:           external multihop AS4
  Source Address:    112.112.112.1
  Route limit:       1/15
  Hold timer:        204/240
  Keepalive timer:   65/80

```

The following example displays the BGP routes learnt through BGP:

```

nfvis# support show bgp route

BIRD 1.6.8 ready.
91.91.91.0/24    dev ipsec0 [bgp1 23:53:18 from 166.34.121.112] (100) [AS65000?]

```



Note NFVIS can learn up to 15 prefixes.

BGP Neighbor Configuration Example

```

router bgp 65000
  bgp router-id 166.34.121.112
  bgp always-compare-med
  bgp log-neighbor-changes
  bgp deterministic-med
  bgp listen range 112.112.0.0/16 peer-group uCPEs
  bgp listen range 90.90.90.0/24 peer-group uCPEs
  bgp listen range 10.20.0.0/24 peer-group uCPEs
  bgp listen limit 255
  no bgp default ipv4-unicast
  !
  address-family ipv4 vrf private-vrf
    redistribute connected
    redistribute static
    neighbor uCPEs peer-group
    neighbor uCPEs remote-as 200
    neighbor uCPEs ebgp-multihop 10
    neighbor uCPEs timers 610 1835
    neighbor uCPEs prefix-list allow-list out
  exit-address-family

```

Route Distribution

The Route Distribution feature works together with a remote BGP router. It allows you to announce or withdraw specified routes to the remote BGP router.

You can use this feature to announce the route of int-mgmt-net subnet to a remote BGP router. A remote user, can access the VMs attached to int-mgmt-net through the VMs' IP address on int-mgmt-net-br through a BGP router, when the routes are successfully inserted on the remote BGP router.

To configure or update route distribution:

```
configure terminal
router bgp 172.25.221.17local-bridge wan-br local-as 45.45remote-as 65000 network-subnet
12.12.12.0/24
commit
```

Table 2: Property Description

Property	Type	Description	Mandatory
neighbor-address	IPv4	BGP neighbor IPv4 address. It is the key of the route distribution list.	Yes
local-address	IPv4	Local IPv4 address. This address must be configured as neighbor IP address on the remote BGP router. If not configured, local-address is set to local-bridge's IP address.	No
local-as		Local autonomous system number. It can be in following two formats: <decimal number, 1.0 .. 65535.65535><unsignedInt, 1 .. 4294967295>	Yes
local-bridge		Local bridge name for advertising routes (default wan-br).	No
remote-as		Remote autonomous system number. It can be in following two formats: <decimal number, 1.0 .. 65535.65535><unsignedInt, 1 .. 4294967295>	Yes
router-id	IPv4	Local router ID	No

Property	Type	Description	Mandatory
network-subnet		List of network subnet to be announced.	Yes
subnet	IPv4 prefix	Network subnet to be announced H.H.H.H/N	Yes
next-hop	IPv4	IPv4 address of next hop. Default local-address or IP address of local-bridge.	No

Use the **no router bgp** command to delete route distribution. To verify the route-distribution status use the **show router bgp** command.

Remote BGP Router Configuration Example

The NFVIS route distribution feature works together with the remote BGP router. The configuration on NFVIS and on remote BGP router must match.

This example shows the configuration on a remote BGP router.

```
router bgp 65000
  bgp log-neighbor-changes
  neighbor 172.25.221.106 remote-as 45.45
  neighbor 172.25.221.106 update-source GigabitEthernet2
```

BGP Route Announcement over MPLS or IPsec

Table 3: Feature History

Feature Name	Release Information	Description
BGP Route Announcement over MPLS or IPsec	NFVIS 4.5.1	This feature allows you to configure NFVIS to announce routes through BGP over MPLS. NFVIS allows the routes learnt through BGP available over IPsec tunnel over MPLS connection.

With this feature enhancement, the existing routes learnt through BGP over IPsec tunnel are now allowed over MPLS connection. Additionally, NFVIS can now announce routes through BGP, using the same **router bgp** command that is used for learning routes over BGP. For more information on this command, see the Cisco IOS XE [router bgp](#) command.

You can pair the secure overlay configurations to announce NFVIS routes over BGP through IPsec tunnel. The existing router bgp configurations can be updated to add the route announcement feature. Make sure that you remove the existing route distribution configurations before you configure the **router bgp** command.

The following example shows how to configure the announcement of 10.20.0.0/24 subnet over BGP.

```
router bgp 65000
 neighbor 172.25.221.17 remote-as 65001
 address-family ipv4 unicast
   network 10.20.0.0 mask 255.255.255.0
 neighbor 172.25.221.17 activate
```

The following example shows how to remove the announcement of 10.20.0.0/24 subnet from BGP.

```
router bgp 65000
 address-family ipv4 unicast
   no network 10.20.0.0 mask 255.255.255.0
```

The following example shows how to remove a neighbor from the IPv4 address family, and disable route announcements for the same neighbor.

```
router bgp 65000
 address-family ipv4 unicast
   no neighbor 172.25.221.17 activate
```

To view the local BGP status for BGP over MPLS use the **show bgp ipv4 unicast** command.

```
nfvis# show bgp ipv4 unicast
```

Family	Transmission	Router ID	Local AS Number
ipv4	unicast	10.20.0.1	65000

To view the BGP neighbor status for BGP over MPLS use the **show bgp ipv4 unicast summary** command.

```
nfvis# show bgp ipv4 unicast summary
```

Neighbor	IP Version	AS Number	Up/Down
172.25.221.17	4	65001	up

To view the BGP learned or announced routes for BGP over MPLS use the **show bgp ipv4 unicast route** command.

```
nfvis# show bgp ipv4 unicast route
```

Network	Next-Hop	Metric	LocPrf	Path
10.30.30.0/24	172.25.221.17	0	100	65001 ?
10.40.40.0/24	172.25.221.17	0	100	65001 ?
10.20.0.0/24	0.0.0.0			

To view the local BGP status for BGP over IPSec tunnel use the **show bgp vpnv4 unicast** command.

```
nfvis# show bgp vpnv4 unicast
```

Family	Transmission	Router ID	Local AS Number
vpnv4	unicast	10.20.0.1	200

To show BGP neighbor status for BGP over IPSec tunnel:

```
nfvis# show bgp vpnv4 unicast summary
```

Neighbor	IP Version	AS Number	Up/Down
10.90.90.1	4	65000	up

To show BGP learned/announced routes for BGP over IPSec tunnel:

```
nfvis# show bgp vpnv4 unicast route
```

Network	Next-Hop	Metric	LocPrf	Path
10.91.91.0/24	10.90.90.1	0	100	65000 ?
10.92.92.0/24	10.90.90.1	0	100	65000 ?
10.20.0.0/24	0.0.0.0			i



Note When you configure BGP route announcement over IPsec tunnel, ensure that you configure secure overlay to use the virtual IP address for the local tunnel IP address (**no local-system-ip-addr configured**).



Note When you configure BGP route announcement, the only configurable address-family or transmission combination is **ipv4 unicast** for both IPsec and MPLS. To view the BGP status, the configurable address-family or transmission for IPsec is **vpnv4 unicast** and for MPLS is **ipv4 unicast**.
