BFD Support on DMVPN

Bidirectional Forwarding Detection (BFD) support on DMVPN provides fast peer failure detection by sending rapid failure detection notices to the control protocols and reducing overall network convergence time.

Prerequisites for BFD Support on DMVPN

BFD for DMVPN supports both IPv4 and IPv6 overlay address and is agnostic to transport address family.

For more BFD prerequisites refer Prerequisites for Bidirectional Forwarding Detection

Restrictions for BFD Support on DMVPN

- NHRP currently acts only on BFD down events and not on up events.
- Both peers must configure BFD to get BFD support. If one of the peers is not configured with BFD, the other peer creates BFD sessions in down or unknown state.
- Before configuring BFD support on DMVPN, in case of point-to-point (P2P) tunnel, next hop server (NHS) must be configured.
- BFD intervals configured on the peers should be the same in the BFD echo mode for spoke to spoke refresh to work as expected.
- A single DMVPN hub with BFD can be scaled to a maximum of 4095 sessions on a Cisco Aggregation Service Router 1000 Series since the number of BFD sessions on these platforms is limited to 4095 currently. Regular methods of scaling DMVPN like clustering, Server Load Balancing (SLB), hierarchical designs, etc. still apply. This does not impact DMVPN scale without BFD.
- For hierarchical DMVPN deployment, BFD sessions cannot be established between spokes of two different regions if they are in disjoint networks.

Information About BFD Support on DMVPN

BFD Operation

BFD provides a low-overhead, short-duration method of detecting failures in the forwarding path between two adjacent routers, including the interfaces, data links, and forwarding planes.

BFD is a detection protocol that is enabled at the interface and protocol levels. Cisco supports BFD asynchronous mode, which depends on the sending of BFD control packets between two systems to activate and maintain BFD neighbor sessions between routers. Therefore, in order for a BFD session to be created, BFD must be configured on both systems (or BFD peers). Once BFD has been enabled on the interfaces and at the router level for the appropriate protocols (NHRP and the routing protocol on overlay), a BFD session is created, BFD timers are negotiated, and the BFD peers will begin to send BFD control packets to each other at the negotiated interval.

Benefits of BFD Support on DMVPN

- Faster detection of link failure.
- In non-crypto deployments, spoke can detect hub failure only after NHRP registration timeout but hub cannot detect a spoke failure until cache on hub expires (even though routing can re-converge much earlier). BFD allows for a very fast detection for such a failure.
- BFD validates the forwarding path between non authoritative sessions, for example, in scenarios where the hub is configured to respond on behalf of the spoke.
- BFD validates end-to-end data path including the tunnel unlike IKE keepalives/DPD that doesn't pass through the tunnel.
- BFD probes can be off-loaded.

There is no special NHRP configuration needed for BFD support on DMVPN, enabling BFD on an NHRP enabled interface suffices. For DMVPN configuration, refer How to Configure Dynamic Multipoint VPN.

How to Configure BFD Support on DMVPN

Configuring BFD Support on DMVPN

BFD intervals can be directly configured on tunnel interface as shown below:

```
enable
configure terminal
interface tunnel1
bfd interval 1000 min_rx 1000 multiplier 5
no echo
```

BFD intervals can also be configured by defining a template and attaching it to the tunnel interface as shown below

```
enable
configure terminal
bfd-template single-hop sample
interval min-tx 1000 min-rx 1000 multiplier 5
interface tunnel1
bfd template sample
```

Example: BFD Support on DMVPN

Example: BFD Support on DMVPN

The following is an example of configuring BFD support on DMVPN on hub.

```
bfd-template single-hop sample
 interval min-tx 1000 min-rx 1000 multiplier 5
!
interface Tunnel0
 ip address 10.0.0.1 255.255.255.0
 no ip redirects
 ip nhrp authentication cisco123
 ip nhrp network-id 5
 ip nhrp redirect
 ip mtu 1400
 ip tcp adjust-mss 1360
 bfd template sample
 tunnel source GigabitEthernet0/0/0
 tunnel mode gre multipoint
 tunnel key 6
!
interface GigabitEthernet0/0/0
 ip address 10.0.0.1 255.0.0.0
 negotiation auto
!
router eigrp 2
network 10.0.0.0 0.0.0.255
bfd all-interfaces
auto-summary
!
```

The following is an example of configuring BFD support on DMVPN on spoke.

```
bfd-template single-hop sample
 interval min-tx 1000 min-rx 1000 multiplier 5
!
interface Tunnel1
 ip address 10.0.0.10 255.255.255.0
 no ip redirects
 ip nhrp authentication cisco123
 ip nhrp network-id 5
 ip nhrp nhs 10.0.0.1 nbma 10.0.0.10 multicast
```

```
 bfd template sample
 tunnel source GigabitEthernet0/0/0
 tunnel mode gre multipoint
 tunnel key 6
!
interface GigabitEthernet0/0/0
 mtu 4000
 ip address 11.0.0.1 255.0.0.0
 media-type rj45
 negotiation auto
!
interface GigabitEthernet0/0/1
 mtu 6000
 ip address 111.0.0.1 255.255.255.0
 negotiation auto
!
router eigrp 2
 network 11.0.0.0 0.0.0.255
 network 111.0.0.0 0.0.0.255
 network 10.0.0.0 0.0.0.255
bfd all-interfaces
auto-summary
!
ip route 0.0.0.0 0.0.0.0 10.0.0.2
```

The following example outlines how to delete the tunnel entry details when BFD support is down by addition of ip nhrp bfd command. By default, the tunnel entry is not immediately deleted and is deleted after expiry of the entry.

```
!
interface Tunnel0
ip address 10.0.1.100 255.255.255.0
no ip redirects
ip nhrp authentication testing
ip nhrp summary-map 192.168.0.0/16 72.68.100.2
ip nhrp summary-map 77.77.0.0/16 72.68.100.2
ip nhrp network-id 100
ip nhrp bfd delete
ip nhrp redirect
bfd interval 1000 min_rx 1000 multiplier 5
tunnel source Ethernet0/0
tunnel mode gre multipoint
tunnel key 100
tunnel protection ipsec profile default
!
```

Note In this configuration, the tunnel entry is immediately deleted upon receiving a BFD down event. Without this configuration, the cache entry pertaining to the tunnel address of the peer is not deleted and performs its default behaviour.

The following is an example to illustrate faster convergence on spoke.

```
interface Tunnel1
ip address 18.0.0.10 255.255.255.0
no ip redirects
ip nhrp authentication cisco123
ip nhrp network-id 12
ip nhrp nhs 10.0.0.1 nbma 10.0.0.10 multicast
```

```
bfd template sample
tunnel source GigabitEthernet0/0/0
tunnel mode gre multipoint
tunnel key 18
tunnel protection ipsec profile MY_PROFILE
!
bfd-template single-hop sample
interval min-tx 1000 min-rx 1000 multiplier 3
echo
!
router eigrp 2
bfd interface Tunnel1    ------------------------> Specify the interface on which the routing
 protocol must act for BFD up/down events
network 11.0.0.0 0.0.0.255
network 111.0.0.0 0.0.0.255
```

With the above configuration, as soon as BFD is reported down (3 seconds to detect), EIGRP will remove the routes installed from RIB.

The following sample output shows a summary output on hub:

```
device#show dmvpn

Legend: Attrb --> S - Static, D - Dynamic, I - Incomplete
        N - NATed, L - Local, X - No Socket
        T1 - Route Installed, T2 - Nexthop-override
        C - CTS Capable
        # Ent --> Number of NHRP entries with same NBMA peer
        NHS Status: E --> Expecting Replies, R --> Responding, W --> Waiting
        UpDn Time --> Up or Down Time for a Tunnel
==========================================================================

Interface: Tunnel1, IPv4 NHRP Details
Type:Hub, NHRP Peers:2,

 # Ent  Peer NBMA Addr Peer Tunnel Add State  UpDn Tm Attrb
 ----- --------------- --------------- ----- -------- -----
     1 172.17.0.1            10.0.0.1    UP 00:00:14     D
     1 172.17.0.2            10.0.0.2   BFD 00:00:03     D
```

BFD is a new state which implies that while the session is UP as seen by lower layers (IKE, IPSec and NHRP), BFD sees the session as DOWN. As usual, the state is an indication of the lower most layer where the session is not UP. Also, this applies only to the parent cache entry. This could be because it was detected as DOWN by BFD or BFD is not configured on the other side.

The following sample output shows a summary output on spoke:

```
device#show dmvpn
Legend: Attrb --> S - Static, D - Dynamic, I - Incomplete
        N - NATed, L - Local, X - No Socket
        T1 - Route Installed, T2 - Nexthop-override
        C - CTS Capable
        # Ent --> Number of NHRP entries with same NBMA peer
        NHS Status: E --> Expecting Replies, R --> Responding, W --> Waiting
        UpDn Time --> Up or Down Time for a Tunnel
==========================================================================

Interface: Tunnel2, IPv4 NHRP Details
Type:Spoke, NHRP Peers:2,
```

```
 # Ent  Peer NBMA Addr Peer Tunnel Add State  UpDn Tm Attrb
 ----- --------------- --------------- ----- -------- -----
     2 172.17.0.2             10.0.0.2   BFD 00:00:02   DT1
                              10.0.0.2    UP 00:00:02   DT2
     1 172.17.0.11           10.0.0.11    UP 00:05:35     S
```

The following sample shows output for **show ip/ipv6 nhrp** command

```
device#show ip nhrp
10.0.0.2/32 via 10.0.0.2
   Tunnel2 created 00:00:15, expire 00:04:54
   Type: dynamic, Flags: router nhop rib bfd
   NBMA address: 172.17.0.2
10.0.0.11/32 via 10.0.0.11
   Tunnel2 created 00:09:04, never expire
   Type: static, Flags: used bfd
   NBMA address: 172.17.0.11
192.168.1.0/24 via 10.0.0.1
   Tunnel2 created 00:00:05, expire 00:04:54
   Type: dynamic, Flags: router unique local
   NBMA address: 172.17.0.1
    (no-socket)
192.168.2.0/24 via 10.0.0.2
   Tunnel2 created 00:00:05, expire 00:04:54
   Type: dynamic, Flags: router rib nho
   NBMA address: 172.17.0.2
```

BFD flag here implies that there is a BFD session for this peer. This marking is only for parent entries.

The following sample shows output for **show tunnel endpoints** command

```
device#show tunnel endpoints
 Tunnel2 running in multi-GRE/IP mode

 Endpoint transport 172.17.0.2 Refcount 3 Base 0x2ABF53ED09F0 Create Time 00:00:07
   overlay 10.0.0.2 Refcount 2 Parent 0x2ABF53ED09F0 Create Time 00:00:07
   Tunnel Subblocks:
      tunnel-nhrp-sb:
         NHRP subblock has 2 entries; BFD(0x2):U
 Endpoint transport 172.17.0.11 Refcount 3 Base 0x2ABF53ED0B80 Create Time 00:09:07
   overlay 10.0.0.11 Refcount 2 Parent 0x2ABF53ED0B80 Create Time 00:09:07
   Tunnel Subblocks:
      tunnel-nhrp-sb:
         NHRP subblock has 1 entries; BFD(0x1):U
```

For every tunnel endpoint, a new text "**BFD(handle):state**" is added. State here is UP(U), DOWN(D), NONE(N) or INVALID(I).

- In case, BFD is not configured on peer or a session is not UP for the first time, then the state will be N.

The following sample shows output for **show nhrp interfaces** command. This shows the configuration (and not operational) states on the interface or globally.

```
device#show nhrp interfaces
NHRP Config State
-----------------
```

```
Global:
   BFD: Registered

Tunnel1:
   BFD: Disabled

Tunnel2:
   BFD: Enabled
```

This is an internal and hidden command. This will currently display if NHRP is client of BFD and if BFD is enabled on the NHRP interface.

Additional References for BFD Support on DMVPN

Related Documents

Related Topic	Document Title
Dynamic Multipoint VPN Configuration Guide	Dynamic Multipoint VPN Configuration Guide
IP Routing: BFD Configuration Guide	IP Routing: BFD Configuration Guide

MIBs

MIB	MIBs Link
• CISCO-MIB • NHRP MIB • Cisco NHRP Extension MIB • BFD MIB • Tunnel MIB • IPSec MIBs	To locate and download MIBs for selected platforms, Cisco IOS releases, and feature sets, use Cisco MIB Locator found at the following URL: http://www.cisco.com/go/mibs

Technical Assistance

Description	Link
The Cisco Support and Documentation website provides online resources to download documentation, software, and tools. Use these resources to install and configure the software and to troubleshoot and resolve technical issues with Cisco products and technologies. Access to most tools on the Cisco Support and Documentation website requires a Cisco.com user ID and password.	http://www.cisco.com/cisco/web/support/index.html

Feature Information for BFD Support on DMVPN

The following table provides release information about the feature or features described in this module. This table lists only the software release that introduced support for a given feature in a given software release train. Unless noted otherwise, subsequent releases of that software release train also support that feature.

Use Cisco Feature Navigator to find information about platform support and Cisco software image support. To access Cisco Feature Navigator, go to www.cisco.com/go/cfn. An account on Cisco.com is not required.

Table 1: Feature Information for BFD Support on DMVPN

Feature Name	Releases	Feature Information
BFD Support on DMVPN	Cisco IOS Release 16.3	Bidirectional Forwarding Detection (BFD) support on DMVPN feature provides fast peer failure detection by sending rapid failure detection notices to the routing protocols and reducing overall network convergence time. The following commands were modified by this feature: **show dmvpn, show ip nhrp, show ipv6 nhrp, show tunnel endpoints, show nhrp interfaces**.