FlexVPN Migration: Hard Move from DMVPN to FlexVPN on a Different Hub

TAC **Document ID: 115727**

Contributed by Marcin Latosiewicz and Atri Basu, Cisco TAC Engineers.
Jan 09, 2015

Contents

Introduction
Prerequisites
Requirements
Components Used
Migration Procedure
Hard Migration Between Two Different Hubs
Custom Approach
Network Topology
Transport Network Topology
Overlay Network Topology
Configuration
DMVPN Configuration
Spoke DMVPN Configuration
Hub DMVPN Configuration
FlexVPN Configuration
Spoke FlexVPN Configuration
FlexVPN Hub Configuration
Traffic Migration
Migrate to BGP as the Overlay Routing Protocol [Recommended]
Spoke BGP Configuration
Hub BGP Configuration
Migrate Traffic to BGP/FlexVPN
Migrate to New Tunnels with EIGRP
Updated Spoke Configuration
Updated FlexVPN Hub Configuration
DMVPN Hub − Updated BGP Configuration
FlexVPN Hub − Updated BGP Configuration
Migrate Traffic to FlexVPN
Verification Steps
Additional Considerations
Spoke−to−Spoke Tunnels that Already Exist
Clear NHRP Entries
Known Caveats
Related Information

Introduction

This document provides information about how to migrate from a Dynamic Multipoint VPN (DMVPN) network that currently exists to FlexVPN on different hub devices. The configuations for both frameworks coexist on the devices. In this document, only the most common scenario is shown − DMVPN with the use of

the preshared key for authentication and Enhanced Interior Gateway Routing Protocol (EIGRP) as the routing protocol. In this document, migration to Border Gateway Protocol (BGP), which is the recommended routing protocol, and the less−desirable EIGRP is demonstrated.

Prerequisites

Requirements

Cisco recommends that you have basic knowledge of these topics:

- DMVPN
- FlexVPN

Components Used

Note: Not all software and hardware supports Internet Key Exchange version 2 (IKEv2). Refer to the Cisco Feature Navigator for more information.

The information in this document is based on these software and hardware versions:

- Cisco Integrated Service Router (ISR) Version 15.2(4)M1 or newer
- Cisco Aggregation Services Router 1000 Series (ASR1K) 3.6.2 Release 15.2(2)S2 or newer

One the advantages of a newer platform and software is the ability to use Next Generation Cryptography, such as Advanced Encryption Standard (AES) Galois/Counter Mode (GCM) for encryption in Internet Protocol Security (IPsec), as discussed in Request for Comments (RFC) 4106. AES GCM allows you to reach a much faster encryption speed on some hardware. In order to see Cisco recommendations on use of and migration to Next Generation Cryptography, refer to the Next Generation Encryption article.

The information in this document was created from the devices in a specific lab environment. All of the devices used in this document started with a cleared (default) configuration. If your network is live, make sure that you understand the potential impact of any command.

Migration Procedure

Currently, the recommended method to migrate from DMVPN to FlexVPN is for the two frameworks to not operate at the same time. This limitation is scheduled to be removed due to new migration features to be introduced in the ASR 3.10 Release, tracked under multilple enhancement requests on the Cisco side, which include Cisco bug ID CSCuc08066. Those features should be available in late June 2013.

A migration where both frameworks coexist and operate at the same time on the same devices is referred to as a ***soft migration***, which indicates the minimal impact and smooth failover from one framework to another. A migration where configurations for both frameworks coexist, but do not operate at the same time is referred to as a ***hard migration***. This indicates that a switchover from one framework to another means a lack of communication over the VPN, even if minimal.

Hard Migration Between Two Different Hubs

In this document, migration from the DMVPN hub that is currently used to a new FlexVPN hub is discussed. This migration allows inter−communication between spokes migrated already to FlexVPN, and those that still run on DMVPN and can be performed in multiple phases, on each spoke separately.

Provided that routing information is properly populated, communication between migrated and nonmigrated spokes should remain possible. However, additional latency can be observed because migrated and nonmigrated spokes do not build spoke−to−spoke tunnels between each other. At the same time, migrated spokes should be able to establish direct spoke−to−spoke tunnels between themselves. The same applies to nonmigrated spokes.

Until this new migration feature is available, complete these steps in order to perform migrations with a different hub from DMVPN and FlexVPN:

1. Verify connectivity over DMVPN.
2. Add the FlexVPN configuration, and shut down the tunnel that belongs to the new configuration.
3. (During a maintenance window) On each spoke, one by one, shut down the DMVPN tunnel.
4. On the same spoke as in Step 3, unshut the FlexVPN tunnel interfaces.
5. Verify spoke−to−hub connectivity.
6. Verify spoke−to−spoke connectivity within FlexVPN.
7. Verify spoke−to−spoke connectivity with DMVPN from FlexVPN.
8. Repeat Steps 3 through 7 for each spoke separately.
9. If you encounter any problems with the verifications described in Steps 5, 6, or 7, shut down the FlexVPN interface, and unshut the DMVPN interfaces in order to revert to DMVPN.
10. Verify spoke−to−hub communication over the backed−up DMVPN.
11. Verify spoke−to−spoke communication over the backed−up DMVPN.

Custom Approach

If the previous approach might not be the best solution for you due to your network or routing complexities, start a discussion with your Cisco representative before you migrate. The best person with which to discuss a custom migration process is your System Engineer or Advanced Services Engineer.

Network Topology

Transport Network Topology

This diagram shows the typical connection topology of hosts on the Internet. The hub's IP address of ***loopback0*** (***172.25.1.1***) is used in order to terminate the DMVPN IPsec session. The IP address on the new hub (***172.25.2.1***) is used for FlexVPN.

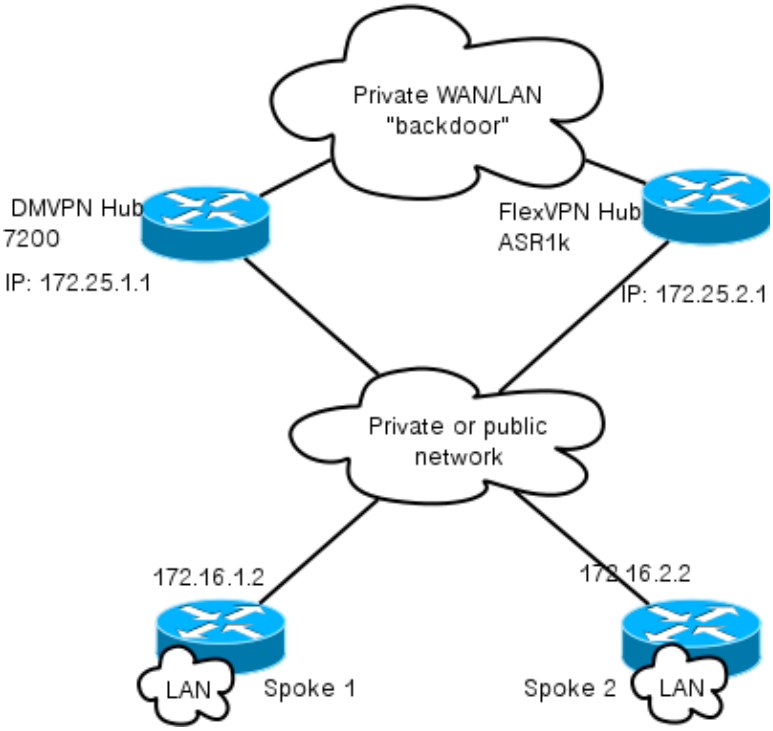

Notice the link between the two hubs. This link is crucial in order to allow connectivity between the FlexVPN and DMVPN clouds during migration. It allows spokes already migrated to FlexVPN to communicate with DMVPN networks and vice versa.

Overlay Network Topology

This topology diagram shows two separate clouds used for overlay: DMVPN (green connections) and FlexVPN (red connections). LAN prefixes are shown for corresponding sites. The ***10.1.1.0/24*** subnet does not represent an actual subnet in terms of interface addressing, but represents a chunk of IP space dedicated to the FlexVPN cloud. The rationale behind this is discussed later in the ***FlexVPN Configuration*** section.

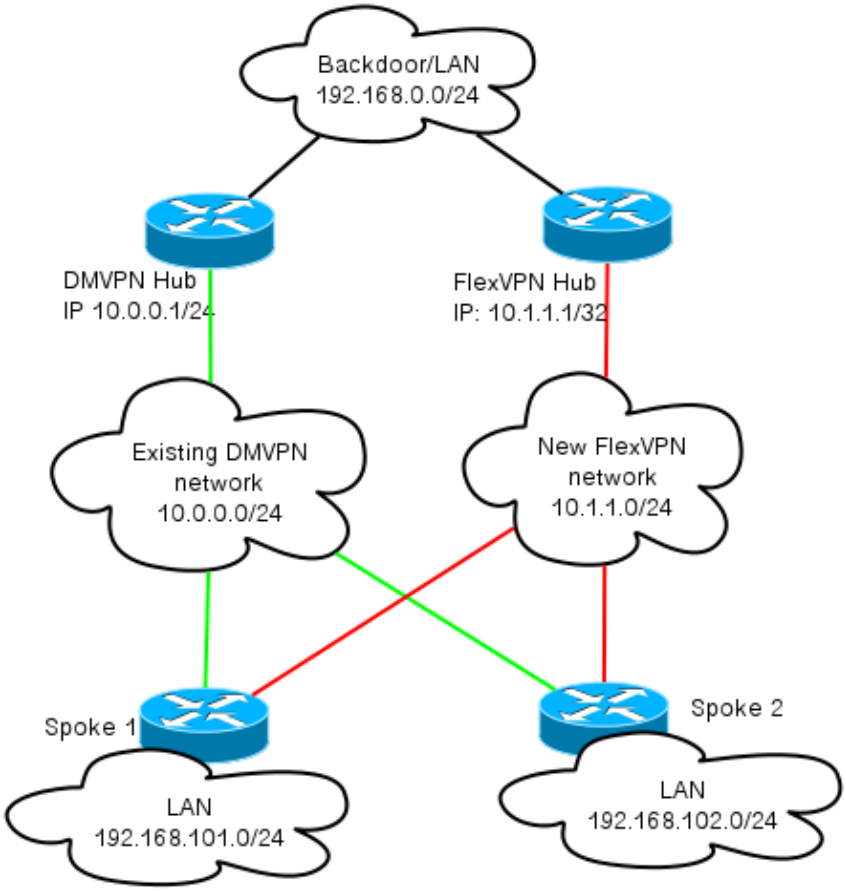

Configuration

This section describes the DMVPN and the FlexVPN configurations.

DMVPN Configuration

This section describes the basic configuration for the DMVPN hub and spoke.

The Pre−Shared Key (PSK) is used for IKEv1 authentication. Once IPsec is established, Next Hop Resolution Protocol (NHRP) registration from spoke−to−hub is performed so that the hub can learn the spokes' Nonbroadcast Multiaccess (NBMA) addressing dynamically.

When NHRP performs registration on the spoke and the hub, routing adjacancy can establish, and routes can be exchanged. In this example, EIGRP is used as a basic routing protocol for the overlay network.

Spoke DMVPN Configuration

Here you can find a basic example configuration of DMVPN with PSK authentication and EIGRP as the routing protocol.

```
crypto isakmp policy 10
  encr aes
  authentication pre−share

crypto isakmp key cisco address 0.0.0.0

crypto isakmp keepalive 30 5

crypto isakmp profile DMVPN_IKEv1
```

```
  keyring DMVPN_IKEv1
  match identity address 0.0.0.0

crypto ipsec transform-set IKEv1 esp-aes esp-sha-hmac
  mode transport

crypto ipsec profile DMVPN_IKEv1
  set transform-set IKEv1
   set isakmp-profile DMVPN_IKEv1

interface Tunnel0

ip address 10.0.0.101 255.255.255.0
 no ip redirects
 ip mtu 1400
 ip nhrp map 10.0.0.1 172.25.1.1
 ip nhrp map multicast 172.25.1.1
 ip nhrp network-id 1
 ip nhrp holdtime 900
 ip nhrp nhs 10.0.0.1
 ip nhrp shortcut
 ip tcp adjust-mss 1360
 tunnel source Ethernet0/0
 tunnel mode gre multipoint
 tunnel protection ipsec profile DMVPN_IKEv1

router eigrp 100
 network 10.0.0.0 0.0.0.255
 network 192.168.102.0
 passive-interface default
 no passive-interface Tunnel0
```

Hub DMVPN Configuration

In the hub configuration, the tunnel is sourced from ***loopback0*** with an IP address of ***172.25.1.1***. The rest is a standard deployment of a DMVPN hub with EIGRP as the routing protocol.

```
crypto isakmp policy 10
 encr aes
 authentication pre-share

crypto isakmp key cisco address 0.0.0.0

crypto ipsec transform-set IKEv1 esp-aes esp-sha-hmac
 mode transport
crypto ipsec profile DMVPN_IKEv1
 set transform-set IKEv1

interface Tunnel0
 ip address 10.0.0.1 255.255.255.0
 no ip redirects
 ip mtu 1400
 ip nhrp map multicast dynamic
 ip nhrp network-id 1
 ip nhrp holdtime 900
 ip nhrp server-only
 ip nhrp redirect
 ip summary-address eigrp 100 192.168.0.0 255.255.0.0
 ip tcp adjust-mss 1360
 tunnel source Loopback0
 tunnel mode gre multipoint
 tunnel protection ipsec profile DMVPN_IKEv1

router eigrp 100
 network 10.0.0.0 0.0.0.255
```

```
network 192.168.0.0 0.0.255.255
passive-interface default
no passive-interface Tunnel0
```

FlexVPN Configuration

FlexVPN is based on these same fundamental technologies:

- ***IPsec***: Unlike the default in DMVPN, IKEv2 is used instead of IKEv1 in order to negotiate IPsec Security Associations (SAs). IKEv2 offers improvements over IKEv1, such as resiliency and the number of messages that are needed in order to establish a protected data channel.

- ***GRE***: Unlike DMVPN, static and dynamic point−to−point interfaces are used, and not only one static multpoint GRE interface. This configuration allows added flexibility, especially for per−spoke/per−hub behavior.

- ***NHRP***: In FlexVPN, NHRP is primarily used in order to establish spoke−to−spoke communication. Spokes do not register to the hub.

- ***Routing***: Because spokes do not perform NHRP registration to the hub, you must rely on other mechanisms in order to make sure the hub and spokes can communicate bidirectionally. Simliar to DMVPN, dynamic routing protocols can be used. However, FlexVPN allows you to use IPsec in order to introduce routing information. The default is to introduce as ***/32*** route for the IP address on the other side of the tunnel, which allows spoke−to−hub direct communication.

In a hard migration from DMVPN to FlexVPN, the two framemworks do not work at the same time on the same devices. However, it is recommended to keep them separate.

Separate them on several levels:

- NHRP − Use a different NHRP network ID (recommended).
- Routing − Use separate routing processes (recommended).
- Virtual Routing and Forwarding (VRF) − VRF separation allows added flexibility but is not discussed here (optional).

Spoke FlexVPN Configuration

One of the differences in the spoke configuration in FlexVPN as compared to DMVPN is that you potentially have two interfaces. There is a required tunnel for spoke−to−hub communication and an optional tunnel for spoke−to−spoke tunnels. If you choose not to have dynamic spoke−to−spoke tunneling and would prefer that everything goes through the hub device, you can remove the virtual template interface, and remove the NHRP shortcut switching from the tunnel interface.

Notice that the static tunnel interface receives an IP address based on negotiation. This allows the hub to provide the tunnel interface IP address to the spoke dynamically without the need to create static addressing in the FlexVPN cloud.

```
aaa new-model
aaa authorization network default local
aaa session-id common

crypto ikev2 profile Flex_IKEv2
 match identity remote fqdn domain cisco.com
 local identity fqdn spoke.cisco.com
 authentication remote rsa-sig
 authentication local rsa-sig
```

```
 aaa authorization group cert list default default
 virtual-template 1
crypto ikev2 dpd 30 5 on-demand
```

Note: By default, the local identity is set in order to use the IP address. So the corresponding match statement on the peer must match based on the address as well. If the requirement is to match based on the Distinguished Name (DN) in the certifcate, then the match must be done with the use of a certificate map.

Cisco recommends that you use AES GCM with hardware that supports it.

```
crypto ipsec transform-set IKEv2 esp-gcm
 mode transport

crypto ipsec profile default
 set ikev2-profile Flex_IKEv2
! set transform-set IKEv2

interface Tunnel1
 ip address negotiated
 ip mtu 1400
 ip nhrp network-id 2
 ip nhrp shortcut virtual-template 1
 ip nhrp redirect
 ip tcp adjust-mss 1360
 shutdown
 tunnel source Ethernet0/0
 tunnel destination 172.25.2.1
 tunnel path-mtu-discovery
 tunnel protection ipsec profile default

interface Virtual-Template1 type tunnel
 ip unnumbered Tunnel1
 ip mtu 1400
 ip nhrp network-id 2
 ip nhrp shortcut virtual-template 1
 ip nhrp redirect
 ip tcp adjust-mss 1360
 tunnel path-mtu-discovery
 tunnel protection ipsec profile default
```

Public Key Infrastructure (PKI) is the recommended method to perform large scale authentication in IKEv2. However, you can still use PSK as long as you are aware of its limitations.

Here is an example configuration that uses ***cisco*** as the PSK.

```
crypto ikev2 keyring Flex_key
 peer Spokes
 address 0.0.0.0 0.0.0.0
 pre-shared-key local cisco
 pre-shared-key remote cisco
crypto ikev2 profile Flex_IKEv2
 match identity remote address 0.0.0.0
 authentication remote pre-share
 authentication local pre-share
 keyring local Flex_key
 aaa authorization group psk list default default
 virtual-template 1
crypto ikev2 dpd 30 5 on-demand
```

FlexVPN Hub Configuration

Typically, a hub only terminates dynamic spoke−to−hub tunnels. This is why you do not find a static tunnel interface for FlexVPN in the hub configuration. Instead, a virtual template interface is used.

Note: On the hub side, you must indicate the pool addresses to be assigned to spokes.

Addresses from this pool are added later in the routing table as ***/32*** routes for each spoke.

```
aaa new−model
aaa authorization network default local
aaa session−id common

crypto ikev2 authorization policy default
 pool FlexSpokes
crypto ikev2 profile Flex_IKEv2
 match identity remote fqdn domain cisco.com
 local identity fqdn hub.cisco.com
 authentication remote rsa−sig
 authentication local rsa−sig
 aaa authorization group cert list default default
 virtual−template 1
crypto ikev2 dpd 30 5 on−demand
```

Cisco recommends that you use AES GCM with hardware that supports it.

```
crypto ipsec transform−set IKEv2 esp−gcm
 mode transport
```

Note: In this configuration, the AES GCM operation has been commented out.

```
crypto ipsec profile default
 set ikev2−profile Flex_IKEv2
! set transform−set IKEv2

interface Loopback0
 description DMVPN termination
 ip address 172.25.2.1 255.255.255.255
interface Loopback100
 ip address 10.1.1.1 255.255.255.255
interface Virtual−Template1 type tunnel
 ip unnumbered Loopback100
 ip nhrp network−id 2
 ip nhrp redirect
 tunnel path−mtu−discovery
 tunnel protection ipsec profile default
ip local pool FlexSpokes 10.1.1.100 10.1.1.254
```

With authentication in IKEv2, the same principle applies on the hub as on the spoke. For scalability and flexibility, use certificates. However, you can reuse the same configuration for PSK as on the spoke.

Note: IKEv2 offers flexibility in terms of authentication. One side can authenticate with PSK while the other side uses Rivest−Shamir−Adleman Signature (RSA−SIG).

If the requirement is to use preshared keys for authentication, then the configuration changes are similar to those described for the spoke router here.

Inter−Hub BGP Connection

Make sure that the hubs know where particular prefixes are located. This becomes increasingly important because some spokes were migrated to FlexVPN while some other spokes remain on DMVPN.

Here is the inter−hub BGP connection based on the DMVPN hub configuration:

```
router bgp 65001
 network 192.168.0.0
 neighbor 192.168.0.2 remote−as 65001
```

Traffic Migration

Migrate to BGP as the Overlay Routing Protocol [Recommended]

BGP is a routing protocol that is based on unicast exchange. Due to its characteristics, it is the best scaling protocol in DMVPN networks.

In this example, Internal BGP (iBGP) is used.

Spoke BGP Configuration

Spoke migration consists of two parts. First, enable BGP as dynamic routing:

```
router bgp 65001
 bgp log−neighbor−changes
 network 192.168.101.0
 neighbor 10.1.1.1 remote−as 65001
```

After the BGP neighbor comes up (see the next section) and new prefixes over BGP are learned, you can swing traffic from the current DMVPN cloud to a new FlexVPN cloud.

Hub BGP Configuration

FlexVPN Hub − Full BGP Configuration

On the hub, in order to avoid keeping the neighborship configuration for each spoke separately, configure dynamic listeners. In this setup, BGP does not initiate new connections, but accepts connections from the provided pool of IP addresses. In this case, the said pool is ***10.1.1.0/24***, which is all of the addresses in the new FlexVPN cloud.

Two points to note:

- The FlexVPN hub advertises specific prefixes to the DMVPN hub; thus the unsupress map is being used.
- Either advertise the FlexVPN subnet of ***10.1.1.0/24*** to the routing table, or make sure that the DMVPN hub sees the FlexVPN hub as the next hop.

This document shows the latter approach.

```
access−list 1 permit any
route−map ALL permit 10
 match ip address 1

route−map SET_NEXT_HOP permit 10
 set ip next−hop 192.168.0.2
```

```
router bgp 65001
 network 192.168.0.0
 bgp log-neighbor-changes
 bgp listen range 10.1.1.0/24 peer-group Spokes
 aggregate-address 192.168.0.0 255.255.0.0 summary-only
 neighbor Spokes peer-group
 neighbor Spokes remote-as 65001

 neighbor 192.168.0.1 remote-as 65001
 neighbor 192.168.0.1 route-reflector-client
 neighbor 192.168.0.1 unsuppress-map ALL
 neighbor 192.168.0.1 route-map SET_NEXT_HOP out
```

DMVPN Hub − Full BGP and EIGRP Configuration

The configuration on the DMVPN hub is basic, because it only receives specific prefixes from the FlexVPN hub and advertises prefixes it learns from EIGRP.

```
router bgp 65001
 bgp log-neighbor-changes
 redistribute eigrp 100
 neighbor 192.168.0.2 remote-as 65001
```

Migrate Traffic to BGP/FlexVPN

As discussed before, you must shut down DMVPN functionality and bring FlexVPN up in order to perform migration.

This procedure guarantees minimal impact:

1. On each spoke, separately, enter this:

   ```
   interface tunnel 0
     shut
   ```

   At this point, make sure there are no IKEv1 sessions established to this spoke. This can be verified if you check the output of the ***show crypto isakmp sa*** command and monitor syslog messages generated by the ***crypto logging session*** command. Once this is confirmed, you can proceed to bring up FlexVPN.

2. On the same spoke, enter this:

   ```
   interface tunnel 1
      no shut
   ```

Verification Steps

IPsec Stability

The best way to evaluate IPsec stability is to monitor sylogs with the ***crypto logging session*** configuration command enabled. If you see sessions that go up and down, this can indicate a problem on the IKEv2/FlexVPN level that must be corrected before migration can begin.

BGP Information Populated

If IPsec is stable, make sure that the BGP table is populated with entries from the spokes (on the hub) and

summary from the hub (on the spokes). In the case of BGP, this can be viewed with these commands:

```
show bgp
! or
show bgp ipv4 unicast
! or
show ip bgp summary
```

Here is an example of correct information from the FlexVPN hub:

```
BGP router identifier 172.25.2.1, local AS number 65001
(...omitted...)

Neighbor V AS MsgRcvd MsgSent TblVer InQ OutQ Up/Down State/PfxRcd
*10.1.1.100 4 65001 112 123 16 0 0 01:35:58 1
192.168.0.1 4 65001 97 99 16 0 0 01:24:12 4
```

The output shows that the hub has learned one prefix from each of the spokes, and both spokes are dynamic and marked with an asterisk (*) sign. It also shows that a total of four prefixes from the inter−hub connection is received.

Here is an example of similar information from the spoke:

```
show ip bgp summary
 BGP router identifier 192.168.101.1, local AS number 65001
(...omitted...)

Neighbor V AS MsgRcvd MsgSent TblVer InQ OutQ Up/Down State/PfxRcd
10.1.1.1 4 65001 120 109 57 0 0 01:33:23 2
```

The spoke has received two prefixes from the hub. In the case of this setup, one prefix should be the summary advertised on the FlexVPN hub. The other is DMVPN ***10.0.0.0/24*** network redistributed on the DMVPN spoke into BGP.

Migrate to New Tunnels with EIGRP

EIGRP is a popular choice in DMVPN networks due to its relatively simple deployment and fast convergence. However, it scales worse than BGP, and does not offer many advanced mechanisms that can be used by BGP straight out of the box. The next section describes one of the ways to move to FlexVPN with a new EIGRP process.

Updated Spoke Configuration

A new Autonomous System (AS) is added with a separate EIGRP process:

```
router eigrp 200
 network 10.1.1.0 0.0.0.255
 network 192.168.101.0
 passive−interface default
 no passive−interface Tunnel1
```

Note: It is best not to establish routing protocol adjacency over spoke−to−spoke tunnels. Therefore, only make the interface of ***tunnel1*** (spoke−to−hub) not passive.

Updated FlexVPN Hub Configuration

Similarly, for the FlexVPN hub, prepare the routing protocol in the appopriate AS, matching one configured on the spokes.

```
router eigrp 200
 network 10.1.1.0 0.0.0.255
```

There are two methods that are used in order to provide summary back towards the spoke.

- Redistribute a static route that points to ***null0*** (preferred option).

  ```
  ip route 192.168.0.0 255.255.0.0 null 0
  ip route 10.1.1.0 255.255.255.0 null 0

  ip prefix−list EIGRP_SUMMARY_ONLY seq 5 permit 192.168.0.0/16
  ip prefix−list EIGRP_SUMMARY_ONLY seq 10 permit 10.1.1.0/24

  route−map EIGRP_SUMMARY permit 20
   match ip address prefix−list EIGRP_SUMMARY_ONLY

  router eigrp 200
   distribute−list route−map EIGRP_SUMMARY out Virtual−Template1
   redistribute static metric 1500 10 10 1 1500 route−map EIGRP_SUMMARY
  ```

  This option allows control over summary and redistribution without modifications to the hub's Virtualization Technology (VT) configuration. This is important, because the hub's VT configuration cannot be modified if there is active virtual access associated with it.

- Set up a DMVPN−style summary address on a virtual template.

  This configuration is *not recommended*, because of internal processing and replication of said summary to each virtual access. It is shown here for reference.

  ```
  interface Virtual−Template1 type tunnel
   ip summary−address eigrp 200 192.168.0.0 255.255.0.0
  ```

  Another aspect to account for is the inter−hub routing exchange. This can be done if you redistribute EIGRP instances to iBGP.

DMVPN Hub − Updated BGP Configuration

The configuration remains basic. You must redistribute specific prefixes from EIGRP to BGP:

```
router bgp 65001

 redistribute eigrp 100

 neighbor 192.168.0.2 remote−as 65001
```

FlexVPN Hub − Updated BGP Configuration

Similar to the DMVPN hub, in FlexVPN, you must redistribute the new EIGRP process' prefixes to BGP:

```
router bgp 65001

 redistribute eigrp 200 redistribute static

 neighbor 192.168.0.1 remote−as 65001
```

Migrate Traffic to FlexVPN

You must shut down DMVPN functionality and bring FlexVPN up on each spoke, one at a time, in order to perform migration. This procedure guarantees minimum impact:

1. On each spoke, separately, enter this:

   ```
   interface tunnel 0
     shut
   ```

   At this point, make sure there are no IKEv1 sessions established on this spoke. This can be verified if you check the output of the ***show crypto isakmp sa*** command and monitor syslog messages generated by the ***crypto logging session*** command. Once this is confirmed, you can proceed to bring up FlexVPN.
2. On the same spoke, enter this:

   ```
   interface tunnel 1
      no shut
   ```

Verification Steps

IPsec Stability

As in the case of BGP, you must evaluate if IPsec is stable. The best way to do so is to monitor sylogs with the ***crypto logging session*** configuration command enabled. If you see sessions go up and down, this can indicate a problem on the IKEv2/FlexVPN level that must be corrected before migration can begin.

EIGRP Information in Topology Table

Make sure that your EIGRP topology table is populated with spoke LAN entries on the hub and summary on the spokes. This can be verified if you enter this command on the hub(s) and the spoke(s):

show ip eigrp [AS_NUMBER] topology

Here is an example of output from the spoke:

```
Spoke1#show ip eigrp 200 topology
EIGRP−IPv4 Topology Table for AS(200)/ID(192.168.101.1)
Codes: P − Passive, A − Active, U − Update, Q − Query, R − Reply,
 r − reply Status, s − sia Status

P 10.1.1.1/32, 1 successors, FD is 26112000
 via Rstatic (26112000/0)
 via 10.1.1.1 (26240000/128256), Tunnel1

P 192.168.101.0/24, 1 successors, FD is 281600
 via Connected, Ethernet1/0
```

```
P 192.168.0.0/16, 1 successors, FD is 26114560
 via 10.1.1.1 (26114560/2562560), Tunnel1

P 10.1.1.100/32, 1 successors, FD is 26112000
 via Connected, Tunnel1

P 10.1.1.0/24, 1 successors, FD is 26114560
 via 10.1.1.1 (26114560/2562560), Tunnel1
```

The output shows that the spoke knows about its LAN subnet (in *italic*) and the summaries for those (in ***bold***).

Here is an example of output from the hub:

```
hub2# show ip eigrp 200 topology
 EIGRP-IPv4 Topology Table for AS(200)/ID(172.25.2.1)
Codes: P - Passive, A - Active, U - Update, Q - Query, R - Reply,
 r - reply Status, s - sia Status

P 10.1.1.1/32, 1 successors, FD is 128256
 via Connected, Loopback200

P 192.168.101.0/24, 1 successors, FD is 26905600
 via 10.1.1.100 (26905600/281600), Virtual-Access1

P 192.168.0.0/16, 1 successors, FD is 2562560
 via Rstatic (2562560/0)

P 10.1.1.0/24, 1 successors, FD is 2562560
 via Rstatic (2562560/0)
```

The output shows that the hub knows about the spokes' LAN subnets (in *italic*), the summary prefix it advertises (in ***bold***), and each spoke's assigned IP address via negotiation.

Additional Considerations

Spoke−to−Spoke Tunnels that Already Exist

Because a shutdown of the DMVPN tunnel interface causes NHRP entries to be removed, spoke−to−spoke tunnels that already exist will be torn down.

Clear NHRP Entries

A FlexVPN hub does not rely on the NHRP registration process from the spoke in order to know how to route traffic back. However, dynamic spoke−to−spoke tunnels rely on NHRP entries.

In DMVPN, if NHRP on the hub is cleared, it can result in short−lived connectivity problems. In FlexVPN, clearing NHRP on the spokes will cause the FlexVPN IPsec session, related to spoke−to−spoke tunnels, to be torn down. Clearing NHRP on the hub has no effect on the FlexVPN session.

This is because, in FlexVPN by default:

- Spokes do not register to hubs.
- Hubs work only as NHRP redirectors, and do not install NHRP entries.
- NHRP shortcut entries are installed on spokes for spoke−to−spoke tunnels and are dynamic.

Known Caveats

Spoke−to−spoke traffic might be affected by Cisco bug ID CSCub07382 .

Related Information

- ***DMVPN to FlexVPN Soft Migration Configuration Example***
- ***Technical Support & Documentation − Cisco Systems***

Updated: Jan 09, 2015 Document ID: 115727