DMVPN to FlexVPN Soft Migration Configuration Example

Document ID: 116678

Contributed by Marcin Latosiewicz, Cisco TAC Engineer.
Feb 24, 2014

Contents

Introduction
Prerequisites
Requirements
Components Used
Background Information
Configure
Network Diagrams
Transport Network Diagram
Overlay Network Diagram
Configurations
Spoke Configuration
Hub Configuration
Verify
Pre−Migration Checks
Migration
EIGRP−to−EIGRP Migration
Post−Migration Checks
Additional Considerations
Existing Spoke−to−Spoke Tunnels
Communication between Migrated and Non−Migrated Spokes
Troubleshoot
Problems with Attempts to Establish Tunnels
Problems with Route Propagation
Known Caveats

Introduction

This document describes how to perform a *soft* migration where both Dynamic Multipoint VPN (DMVPN) and FlexVPN work on a device simultaneously without the need for a workaround and provides a configuration example.

Note: This document expands on the concepts described in the FlexVPN Migration: Hard Move from DMVPN to FlexVPN on Same Devices and FlexVPN Migration: Hard Move from DMVPN to FlexVPN on a Different Hub Cisco articles. Both of these documents describe *hard* migrations, which cause some disruption to traffic during migration. The limitations in these articles are due to a deficiency in Cisco IOS® software that is now rectified.

Prerequisites

Requirements

Cisco recommends that you have knowledge of these topics:

- DMVPN
- FlexVPN

Components Used

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

- Cisco Integrated Service Router (ISR) Versions 15.3(3)M or Later
- Cisco 1000 Series Aggregated Service Router (ASR1K) Releases 3.10 or Later

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

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.

Background Information

One of the advantages of the newer Cisco IOS platform and software is the ability to use Next Generation Cryptography. An example is the use of Advanced Encryption Standard (AES) in Galois/Counter Mode (GCM) for encryption in IPsec, as discussed in RFC 4106. AES GCM allows much faster encryption speeds on some hardware.

Note: For additional information about the use of and migration to Next Generation Cryptography, refer to the Next Generation Encryption Cisco article.

Configure

This configuration example focuses on a migration from a DMVPN Phase 3 configuration to a FlexVPN, because both designs work similarly.

	DMVPN Phase 2	***DMVPN Phase 3***	***FlexVPN***
Transport	GRE over IPsec	GRE over IPsec	GRE over IPsec, VTI
NHRP Usage	Registration and Resolution	Registration and Resolution	Resolution
Next Hop from Spoke	Other Spokes or Hub	Summary from Hub	Summary from Hub
NHRP Shortcut Switching	No	Yes	Yes (Optional)
NHRP Redirection	No	Yes	Yes
IKE and IPsec	IPsec Optional, IKEv1 Typical	IPsec Optional, IKEv1 Typical	IPsec, IKEv2

Network Diagrams

This section provides both transport and overlay network diagrams.

Transport Network Diagram

The transport network used in this example includes a single hub with two spokes connected. All of the devices are connected through a network that simulates the Internet.

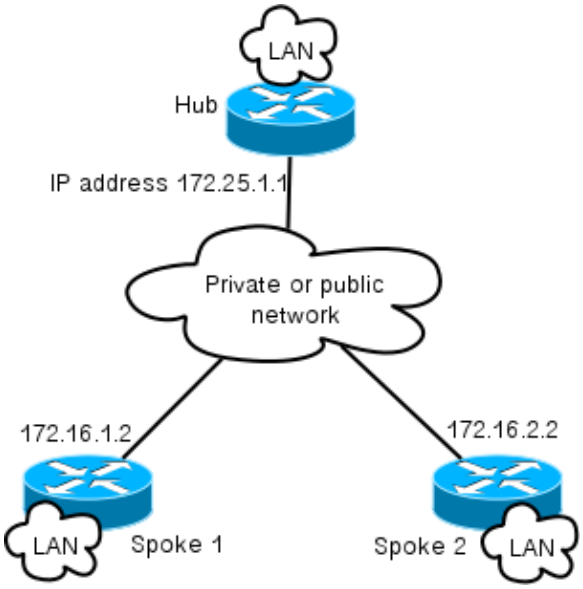

Overlay Network Diagram

The overlay network used in this example includes a single hub with two spokes connected. Remember that both DMVPN and FlexVPN are active simultaneously, but they use different IP address spaces.

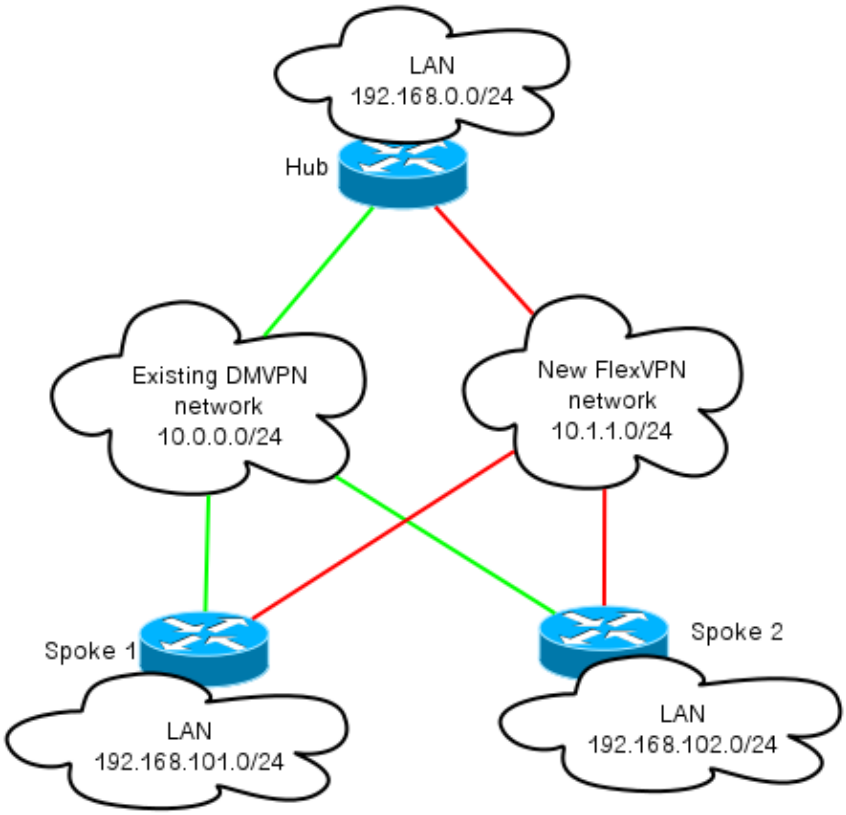

Configurations

This configuration migrates the most popular deployment of DMVPN Phase 3 via Enhanced Interior Gateway Routing Protocol (EIGRP) to FlexVPN with Border Gateway Protocol (BGP). Cisco recommends the use of BGP with FlexVPN, because it allows deployments to scale better.

Note: The hub terminates the IKEv1 (DMVPN) and IKEv2 (FlexVPN) sessions on the same IP address. This is possible only with recent Cisco IOS releases.

Spoke Configuration

This is a very basic configuration, with two notable exceptions that allow inter−operation of both IKEv1 and IKEv2, as well as two frameworks that use Generic Routing Encapsulation (GRE) over IPsec for transport in order to coexist.

Note: The relevant changes to the Internet Security Association and Key Management Protocol (ISAKMP) and IKEv2 configuration are highlighted in bold.

```
crypto keyring DMVPN_IKEv1
 pre−shared−key address 0.0.0.0 0.0.0.0 key cisco

crypto logging session

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

crypto isakmp policy 10
 encr aes
 authentication pre−share

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

crypto ipsec profile default
 set ikev2−profile Flex_IKEv2

interface Tunnel0
 desciption DMVPN tunnel
 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 key 0
 tunnel protection ipsec profile DMVPN_IKEv1 isakmp-profile DMVPN_IKEv1

interface Tunnel1
 description FlexVPN spoke-to-hub tunnel
 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
 tunnel source Ethernet0/0
 tunnel destination 172.25.1.1
 tunnel protection ipsec profile default ikev2-profile Flex_IKEv2

interface Virtual-Template1 type tunnel
 description FlexVPN spoke-to-spoke
 ip unnumbered Ethernet1/0
 ip mtu 1400
 ip nhrp network-id 2
 ip nhrp shortcut virtual-template 1
 ip nhrp redirect
 ip tcp adjust-mss 1360
 tunnel protection ipsec profile default ikev2-profile Flex_IKEv2
```

Cisco IOS Release 15.3 allows you to tie both IKEv2 and ISAKMP profiles together in a *tunnel protection* configuration. Along with some internal changes to the code, this allows IKEv1 and IKEv2 to operate on the same device simultaneously.

Because of the way Cisco IOS selects the profiles (IKEv1 or IKEv2) in releases earlier than 15.3, it led to some caveats, such as situations where IKEv1 is initiated to IKEv2 through the peer. The separation of IKE is now based on profile−level, not interface−level, which is achieved via the new CLI.

Another upgrade in the new Cisco IOS release is the addition of the *tunnel key*. This is needed because both the DMVPN and FlexVPN use the same source interface and the same destination IP address. With this in place, there is no way for the GRE tunnel to know which tunnel interface is used in order to decapsulate traffic. The tunnel key allows you to differentiate ***tunnel0*** and ***tunnel1*** with the addition of a small (4 byte) overhead. A different key can be configured on both interfaces, but you typically only need to differentiate one tunnel.

Note: The shared tunnel protection option is not required when DMVPN and FlexVPN share the same interface.

Thus, the spoke routing protocol configuration is basic. EIGRP and BGP work separately. EIGRP advertises only over the tunnel interface in order to avoid peering over spoke−to−spoke tunnels, which limits scalability. BGP maintains a relationship only with the hub router (***10.1.1.1***) in order to advertise the local network (***192.168.101.0/24***).

```
router eigrp 100
 network 10.0.0.0 0.0.0.255
```

```
 network 192.168.101.0
 passive-interface default
 no passive-interface Tunnel0

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

Hub Configuration

You must make similar changes on the hub-side configuration as those described in the ***Spoke Configuration*** section.

Note: The relevant changes to the ISAKMP and IKEV2 configuration are highlighted in bold.

```
crypto ikev2 authorization policy default
 pool FlexSpokes
 route set interface

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

crypto isakmp policy 10
 encr aes
 authentication pre-share

crypto isakmp key cisco address 0.0.0.0


crypto ipsec profile DMVPN_IKEv1
 set transform-set IKEv1

crypto ipsec profile default
 set ikev2-profile Flex_IKEv2

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 key 0
 tunnel protection ipsec profile DMVPN_IKEv1
```

```
interface Virtual-Template1 type tunnel
 ip unnumbered Loopback100
 ip mtu 1400
 ip nhrp network-id 2
 ip tcp adjust-mss 1360
 tunnel protection ipsec profile default
```

On the hub−side, the binding between the IKE profile and the IPsec profile occurs at the profile−level, unlike spoke configuration, where this is completed via the ***tunnel protection*** command. Both approaches are viable methods to complete this binding.

It is important to note that the Next Hop Resolution Protocol (NHRP) network IDs are different for DMVPN and FlexVPN in the cloud. In most cases, it is undesirable when NHRP creates a single domain over both frameworks.

The tunnel key differentiates DMVPN and FlexVPN tunnels at the GRE−level in order to achieve the same goal that is mentioned in the ***Spoke Configuration*** section.

The routing configuration on the hub is fairly basic. The hub device maintains two relationships with any given spoke, one that uses EIGRP and one that uses BGP. The BGP configuration uses listen−range in order to avoid a lengthy, per−spoke configuration.

The summary addresses are introduced twice. The EIGRP configuration sends a summary with use of the ***tunnel0*** configuration (IP summary−address EIGRP 100), and the BGP introduces a summary with use of the aggregate−address. The summaries are required in order to ensure that the NHRP redirection occurs, and in order to simplify the routing updates. You can send an NHRP redirect (much like an Internet Control Message Protocol (ICMP) redirect) that indicates whether a better hop exists for a given destination, which allows a spoke−to−spoke tunnel to be established. These summaries are also used in order to minimize the amount of routing updates that are sent between the hub and each spoke, which allows setups to scale better.

```
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

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

Verify

The verification for this configuration example is divided into several sections.

Pre−Migration Checks

Since both DMVPN/EIGRP and FlexVPN/BGP operate simultaneously, you must verify that the spoke maintains a relationship over IPsec with both IKEv1 and IKEv2, and that the appropriate prefixes are learned over EIGRP and BGP.

In this example, ***Spoke1*** shows that two sessions are maintained with the hub router; one uses IKEv1/***Tunnel0*** and one uses IKEv2/***Tunnel1***.

Note: Two IPsec Security Associations (SAs) (one inbound and one outbound) are maintained for each of the tunnels.

```
Spoke1#show cry sess
Crypto session current status

Interface: Tunnel0
Profile: DMVPN_IKEv1
Session status: UP-ACTIVE
Peer: 172.25.1.1 port 500
 Session ID: 0
 IKEv1 SA: local 172.16.1.2/500 remote 172.25.1.1/500 Active
 IPSEC FLOW: permit 47 host 172.16.1.2 host 172.25.1.1
 Active SAs: 2, origin: crypto map

Interface: Tunnel1
Profile: Flex_IKEv2
Session status: UP-ACTIVE
Peer: 172.25.1.1 port 500
 Session ID: 1
 IKEv2 SA: local 172.16.1.2/500 remote 172.25.1.1/500 Active
 IPSEC FLOW: permit 47 host 172.16.1.2 host 172.25.1.1
 Active SAs: 2, origin: crypto map
```

When you check the routing protocols, you must verify that a neighborship is formed, and that the correct prefixes are learned. This is first checked with the EIGRP. Verify that the hub is visible as a neighbor, and that the ***192.168.0.0/16*** address (the summary) is learned from the hub:

```
Spoke1#show ip eigrp neighbors
EIGRP-IPv4 Neighbors for AS(100)
H Address Interface Hold Uptime SRTT RTO Q Seq
 (sec) (ms) Cnt Num
0 10.0.0.1 Tu0 10 00:04:02 7 1398 0 13

Spoke1#show ip eigrp topology
EIGRP-IPv4 Topology Table for AS(100)/ID(192.168.101.1)
Codes: P - Passive, A - Active, U - Update, Q - Query, R - Reply,
 r - reply Status, s - sia Status

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 26880000
 via 10.0.0.1 (26880000/256), Tunnel0
P 10.0.0.0/24, 1 successors, FD is 26880000
 via Connected, Tunnel0
```

Next, verify the BGP:

```
Spoke1#show bgp summary
(...)

Neighbor V AS MsgRcvd MsgSent TblVer InQ OutQ Up/Down State/PfxRcd
10.1.1.1 4 65001 13 11 3 0 0 00:06:56 1
Spoke1#show bgp
BGP table version is 3, local router ID is 192.168.101.1
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
 r RIB-failure, S Stale, m multipath, b backup-path, f RT-Filter,
 x best-external, a additional-path, c RIB-compressed,
Origin codes: i - IGP, e - EGP, ? - incomplete
RPKI validation codes: V valid, I invalid, N Not found

Network Next Hop Metric LocPrf Weight Path
```

```
 r>i 192.168.0.0/16 10.1.1.1 0 100 0 i
 *> 192.168.101.0 0.0.0.0 0 32768 i
```

The output shows that the hub FlexVPN IP address (***10.1.1.1***) is a neighbor through which the spoke receives one prefix (***192.168.0.0/16***). Additionally, the BGP informs the administrator that a Routing Information Base (RIB) failure occurred for the ***192.168.0.0/16*** prefix. This failure occurs because there is a better route for that prefix that already exists in the routing table. This route is originated by EIGRP, and can be confirmed if you check the routing table.

```
Spoke1#show ip route 192.168.0.0 255.255.0.0
Routing entry for 192.168.0.0/16, supernet
 Known via "eigrp 100", distance 90, metric 26880000, type internal
 Redistributing via eigrp 100
 Last update from 10.0.0.1 on Tunnel0, 00:10:07 ago
 Routing Descriptor Blocks:
 * 10.0.0.1, from 10.0.0.1, 00:10:07 ago, via Tunnel0
 Route metric is 26880000, traffic share count is 1
 Total delay is 50000 microseconds, minimum bandwidth is 100 Kbit
 Reliability 255/255, minimum MTU 1400 bytes
 Loading 1/255, Hops 1
```

Migration

The previous section verified that both the IPsec and the routing protocols are configured and work as expected. One of the easiest ways to migrate from DMVPN to FlexVPN on the same device is to change the Administrative Distance (AD). In this example, the Internal BGP (iBGP) has an AD of ***200***, and the EIGRP has an AD of ***90***.

In order for traffic to flow through the FlexVPN properly, the BGP must have a better AD. In this example, the EIGRP AD is changed to ***230*** and ***240*** for internal and external routes, respectively. This makes the BGP AD (of ***200***) more preferable for the ***192.168.0.0/16*** prefix.

Another method that is used in order to acheive this is to decrease the BGP AD. However, the protocol that runs after the migration has non−default values, which can impact other parts of the deployment.

In this example, the ***debug ip routing*** command is used in order to verify operation on the spoke.

Note: If the information in this section is used on a production network, avoid the use of debug commands, and rely on the show commands listed in the next section. Also, the spoke EIGRP process must reestablish adjacency with the hub.

```
Spoke1#conf t
Enter configuration commands, one per line. End with CNTL/Z.
Spoke1(config)#router eigrp 100
Spoke1(config−router)# distance eigrp 230 240
Spoke1(config−router)#^Z
Spoke1#
*Oct 9 12:12:34.207: %SYS−5−CONFIG_I: Configured from console by console
*Oct 9 12:12:43.648: %DUAL−5−NBRCHANGE: EIGRP−IPv4 100: Neighbor 10.0.0.1
 (Tunnel0) is down: route configuration changed

*Oct 9 12:12:43.648: RT: delete route to 192.168.0.0 via 10.0.0.1,
 eigrp metric [90/26880000]
*Oct 9 12:12:43.648: RT: no routes to 192.168.0.0, delayed flush
*Oct 9 12:12:43.648: RT: delete network route to 192.168.0.0/16
*Oct 9 12:12:43.650: RT: updating bgp 192.168.0.0/16 (0x0) :
 via 10.1.1.1

*Oct 9 12:12:43.650: RT: add 192.168.0.0/16 via 10.1.1.1, bgp metric [200/0]
Spoke1#
```

```
*Oct 9 12:12:45.750: %DUAL-5-NBRCHANGE: EIGRP-IPv4 100: Neighbor 10.0.0.1
 (Tunnel0) is up: new adjacency
```

There are three important actions to notice in this output:

- The spoke notices that the AD changed, and disables the adjacency.
- In the routing table, the EIGRP prefix is retied, and the BGP is introduced.
- Adjacency to the hub over the EIGRP comes back online.

When you change the AD on a device, it only affects the path from the device to the other networks; it does not affect how other routers perform routing. For example, after the EIGRP distance is increased on ***Spoke1*** (and it uses FlexVPN on the cloud in order to route traffic), the hub maintains the configured (default) ADs. This means that it uses DMVPN in order to route traffic back to ***Spoke1***.

In certain scenarios, this can cause problems, such as when firewalls expect return traffic on the same interface. Therefore, you should change the AD on all spokes before you change it on the hub. Traffic is fully migrated by FlexVPN only once this is complete.

EIGRP-to-EIGRP Migration

A migration from DMVPN to FlexVPN that runs only EIGRP is not discussed in-depth in this document; however, it is mentioned here for completeness.

It is possible to add both DMVPN and EIGRP to the same EIGRP Autonomous System (AS) routing instance. With this in place, the routing adjacency is established over both types of clouds. This can cause load-balancing to occur, which is typically not recommended.

In order to ensure that either FlexVPN or DMVPN is chosen, an administrator can assign different ***Delay*** values on a per-interface basis. However, it is important to remember that no changes are possible on the virtual-template interfaces while corresponding virtual-access interfaces are present.

Post-Migration Checks

Similar to the process used in the ***Pre-Migration Checks*** section, the IPsec and routing protocol must be verified.

First, verify the IPsec:

```
Spoke1#show crypto session
Crypto session current status

Interface: Tunnel0
Profile: DMVPN_IKEv1
Session status: UP-ACTIVE
Peer: 172.25.1.1 port 500
 Session ID: 0
 IKEv1 SA: local 172.16.1.2/500 remote 172.25.1.1/500 Active
 IPSEC FLOW: permit 47 host 172.16.1.2 host 172.25.1.1
 Active SAs: 2, origin: crypto map

Interface: Tunnel1
Profile: Flex_IKEv2
Session status: UP-ACTIVE
Peer: 172.25.1.1 port 500
 Session ID: 1
 IKEv2 SA: local 172.16.1.2/500 remote 172.25.1.1/500 Active
 IPSEC FLOW: permit 47 host 172.16.1.2 host 172.25.1.1
 Active SAs: 2, origin: crypto map
```

As before, two sessions are seen, both of which have two active IPsec SAs.

On the spoke, the aggregate route (***192.168.0.0/16***) points from the hub and is learned over BGP.

```
Spoke1#show ip route 192.168.0.0 255.255.0.0
Routing entry for 192.168.0.0/16, supernet
 Known via "bgp 65001", distance 200, metric 0, type internal
 Last update from 10.1.1.1 00:14:07 ago
 Routing Descriptor Blocks:
 * 10.1.1.1, from 10.1.1.1, 00:14:07 ago
 Route metric is 0, traffic share count is 1
 AS Hops 0
 MPLS label: none
```

Similarly, the spoke LAN that is prefixed on the hub must be known via the EIGRP. In this example, the ***Spoke2*** LAN subnet is checked:

```
Hub#show ip route 192.168.102.0 255.255.255.0
Routing entry for 192.168.102.0/24
 Known via "bgp 65001", distance 200, metric 0, type internal
 Last update from 10.1.1.106 00:04:35 ago
 Routing Descriptor Blocks:
 * 10.1.1.106, from 10.1.1.106, 00:04:35 ago
 Route metric is 0, traffic share count is 1
 AS Hops 0
 MPLS label: none

Hub#show ip cef 192.168.102.100
192.168.102.0/24
 nexthop 10.1.1.106 Virtual-Access2
```

In the output, the forwarding path is updated properly and points out of a virtual−access interface.

Additional Considerations

This section describes some additional areas of importance that are relevant to this configuration example.

Existing Spoke−to−Spoke Tunnels

With a migration from EIGRP to BGP, the spoke−to−spoke tunnels are not impacted, because shortcut−switching is still in operation. Shortcut−switching on the spoke inserts a more specific NHRP route with an AD of 250.

Here is an example of such a route:

```
Spoke1#show ip route 192.168.102.100
Routing entry for 192.168.102.0/24
 Known via "nhrp", distance 250, metric 1
 Last update from 10.1.1.106 on Virtual-Access1, 00:00:42 ago
 Routing Descriptor Blocks:
 * 10.1.1.106, from 10.1.1.106, 00:00:42 ago, via Virtual-Access1
 Route metric is 1, traffic share count is 1
```

Communication between Migrated and Non−Migrated Spokes

If a spoke that is already on a FlexVPN/BGP wants to communicate with a device for which the migration process has not begun, the traffic always flows over the hub.

This is the process that occurs:

1. The spoke performs a route lookup for the destination, which points through a summary route that is advertised by the hub.
2. The packet is sent towards the hub.
3. The hub receives the packet and performs a route lookup for the destination, which points out of another interface that is part of a different NHRP domain.

Note: The NHRP network ID in the previous hub configuration is different for both FlexVPN and DMVPN.

Even if the NHRP network IDs are unified, a problem might occur where the migrated spoke routes objects over the FlexVPN network. This includes the directive used in order to configure shortcut switching. The non−migrated spoke attempts to run objects over the DMVPN network, with a specific goal to perform shortcut switching.

Troubleshoot

This section describes the two categories typically used in order to toubleshoot the migration.

Problems with Attempts to Establish Tunnels

Complete these steps if the IKE negotiation fails:

1. Verify the current state with these commands:

   - ***show crypto isakmp sa*** − This command reveals the amount, source, and destination of an IKEv1 session.
   - ***show crypto ipsec sa***− This command reveals the activity of IPsec SAs.

     Note: Unlike in IKEv1, in this output the Perfect Forward Secrecy (PFS) Diffie−Hellman (DH) Group value appears as ***PFS (Y/N): N, DH group: none*** during the first tunnel negotiation; however, after a rekey occurs, the correct values appear. This is not a bug, even though the behavior is described in CSCug67056. The difference between IKEv1 and IKEv2 is that in the latter, the Child SAs are created as a part of the ***AUTH*** exchange. The DH Group that is configured under the crypto map is used only during a rekey. For this reason, you see ***PFS (Y/N): N, DH group: none*** until the first rekey. With IKEv1, you see a different behavior because the Child SA creation occurs during Quick Mode, and the ***CREATE_CHILD_SA*** message has provisions for the transference of the Key Exchange payload that specifies the DH parameters in order to derive a new shared secret.
   - ***show crypto ikev2 sa*** − This command provides output similar to ISAKMP but is specific to IKEv2.
   - ***show crypto session*** − This command provides the summary output of the cryptographic sessions on this device.
   - ***show crypto socket*** − This command shows the status of crypto−sockets.
   - ***show crypto map*** − This command shows the mapping of IKE and IPsec profiles to the interfaces.
   - ***show ip nhrp*** − This command provides the NHRP informtion from the device. This is useful for spoke−to−spoke in FlexVPN setups, and for both spoke−to−spoke and spoke−to−hub bindings in DMVPN setups.
2. Use these commands in order to debug the tunnel establishment:

   - ***debug crypto ikev2***
   - ***debug crypto isakmp***

- ***debug crypto ipsec***
- ***debug crypto kmi***

Problems with Route Propagation

Here are some useful commands that you can use in order to troubleshoot the EIGRP and topology:

- ***show bgp summary*** − Use this command in order to verify the connected neighbors and their states.
- ***show ip eigrp neighbor*** − Use this command in order to show the neighbors that are connected via EIGRP.
- ***show bgp*** − Use this command in order to verify the prefixes learned over the BGP.
- ***show ip eigrp topology*** − Use this command in order to show the prefixes learned via EIGRP.

It is important to know that a learned prefix is different than a prefix that is installed in the routing table. For more information about this, reference the Route Selection in Cisco Routers Cisco article, or the Routing TCP/IP Cisco Press book.

Known Caveats

A limitation that parallels GRE tunnel handling exists on the ASR1K. This is tracked under Cisco bug ID CSCue00443. At this time, the limitation has a scheduled fix in Cisco IOS XE Software Release 3.12.

Monitor this bug if you desire a notification once the fix becomes available.

Updated: Feb 24, 2014 Document ID: 116678