Tech Talk Series

Tim McConnaughy
CCIE #58615 R/S
Enterprise Route/Switch NCE
tmcconna@cisco.com

About Me

Tim
Level 8 Engineer

STR: 10
DEX: 12
CON: 13
INT: 16
WIS: 12
CHA: 18

Attacks:

Winded Explanation (Sonic): Enemies are afflicted by Confusion for 1d6 rounds.

About Me

Josephine (9)
Level 1 Fairy Princess

Fiona (4)
Level 1 Moana

About Me

Maria
Level 29+ Wife

DMVPN Overview

*Marketing Jargon Optional

- Point-to-Multipoint Layer 3 VPN
- Hub/Spoke Topology
- Supports Multiple Passenger Protocols
- Scalable Connectivity

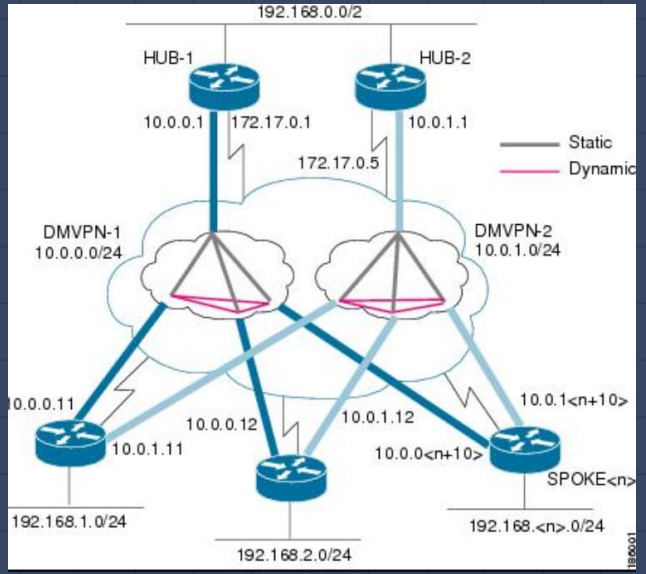

DMVPN Components

- Generic Routing Encapsulation (Specifically mGRE)
- Next-Hop Resolution Protocol
- Dynamic Routing Protocols
- Crypto IPSec (Optional but recommended through insecure transport)

DMVPN Topology

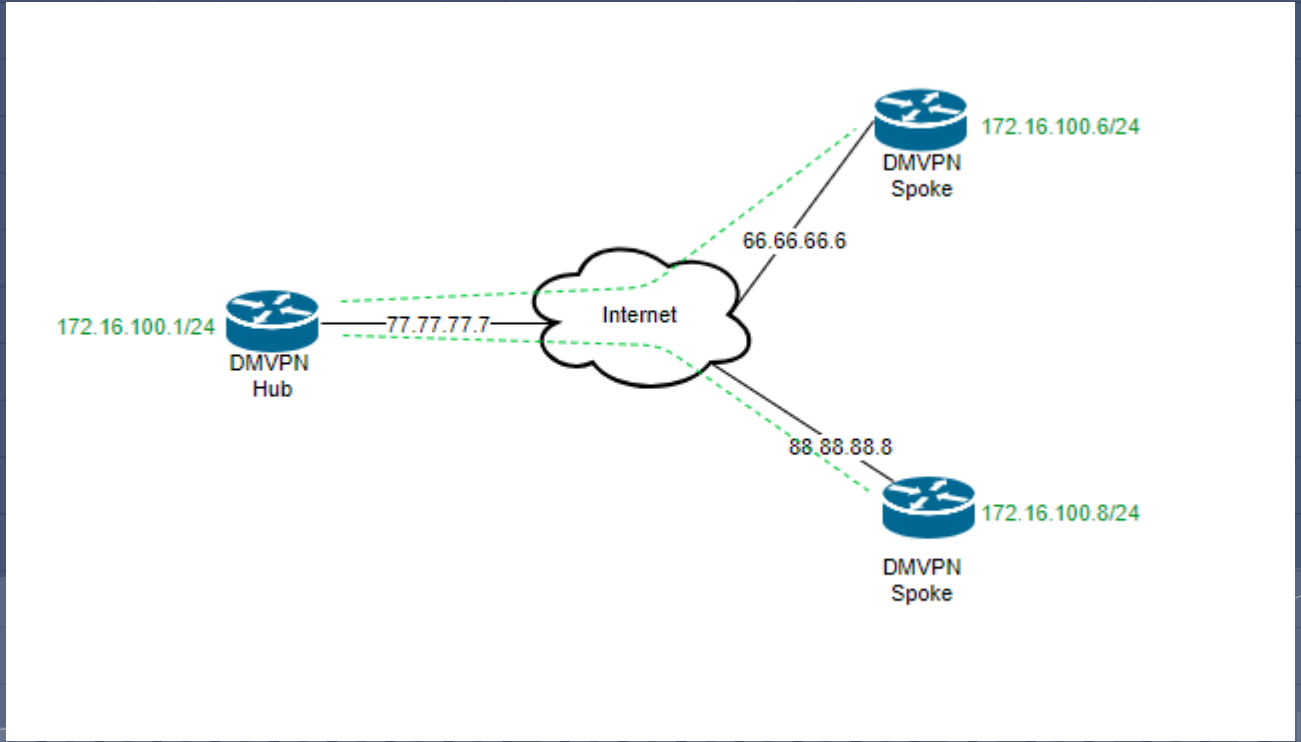

Generic Routing Encapsulation

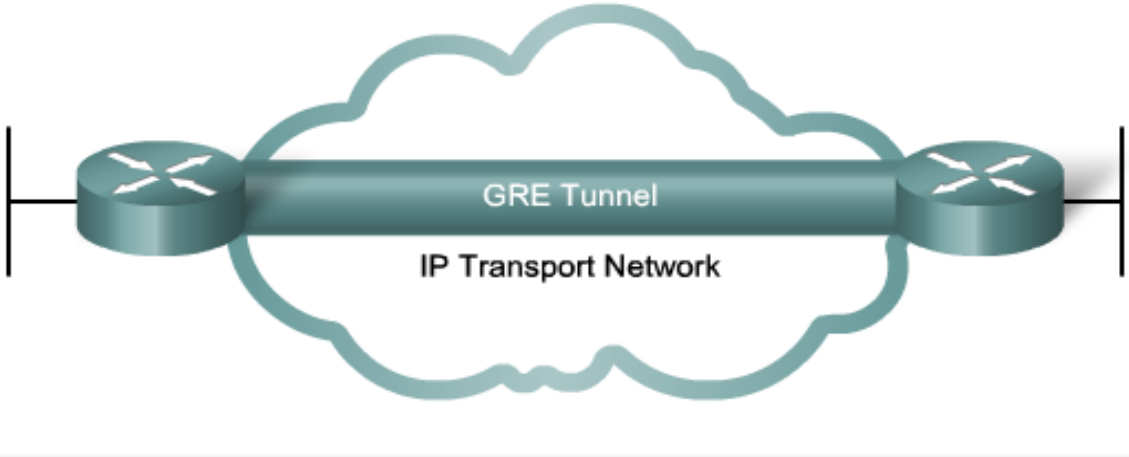

GRE Tunneling Protocol

- Encapsulates a wide variety of protocol packet types inside IP tunnels.
- Creates a virtual point-to-point link to Cisco routers at remote points over an IP internetwork.
- Uses IP for transport.
- Uses an additional header to support IP multicasting and any other OSI Layer 3 protocol as payload.

GRE: What is it Good For?

- Transport the untransportable – IPv6 over IPv4 transport or vice versa as an example
- Virtual network adjacency – Supporting some legacy application protocols and solutions
- Routing manipulation – Path preference and traffic engineering

GRE: How Does it Work?

- A logical GRE tunnel interface is created on two device endpoints. This logical tunnel requires that an interface be defined as the source of the GRE tunnel. This can be a physical or logical interface, but can not be itself.
- A tunnel destination is also configured on the GRE Tunnel interface*. This destination IP address must be routable in order for the tunnel to go from a down to up status and allow traffic to be routed through the tunnel.
- When each tunnel source and destination is routable from the other, the GRE tunnel comes up and traffic can be directed into the tunnel by any method, be it IGP/EGP or static routing.

* Multipoint GRE does not need a tunnel destination configured, and will be discussed later

GRE: How Does it Work?

- When the next-hop of a packet is through a GRE tunnel, the router encapsulates the packet in a GRE header which includes the entire original packet, preserving the source/destination of that original packet. While that GRE-encapsulated packet traverses the infrastructure to the tunnel destination, the original packet is not consulted for routing decisions.
- When the GRE-encapsulated packet reaches the tunnel destination, the GRE header is stripped and the remote endpoint makes a routing decision based on the original packet destination.

GRE: Flying The Friendly Skies

GRE Tunneling is like being a passenger on an airplane

Passengers fly over the infrastructure from hop to hop, or from airport to airport, without being aware of the supporting infrastructure or conditions outside

GRE: Underlay vs. Overlay

- Overlay network refers to the network virtualization that is achieved by abstracting the details of the underlay network from the passenger protocol in a GRE tunnel.

- Underlay network refers to the true network infrastructure and layout which serves as transport for the overlay network endpoints to reach each other.

GRE Tunnel Checklist

- ✓ Tunnel interface created – Number is locally significant
- ✓ Tunnel IP addressing configured for overlay network
- ✓ Tunnel source selected – IP or interface
- ✓ Tunnel destination selected – IP only*
- ✓ (Optional) Tunnel Key – Needed if sourcing multiple GRE tunnels from same interface
- ✓ Verify reachability between tunnel source and destination

* Except in case of GRE multipoint

```
interface Tunnel100
```

```
ip address 172.16.100.8 255.255.255.128
```

```
tunnel source Ethernet0/0
```

```
tunnel destination 77.77.77.7
```

```
tunnel key 100
```

```
DMVPN-SPOKE2#ping 77.77.77.7 source Eth0/0
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 77.77.77.7, timeout is 2 seconds:
Packet sent with a source address of 88.88.88.8
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/1 ms
```

Anatomy of a GRE Packet

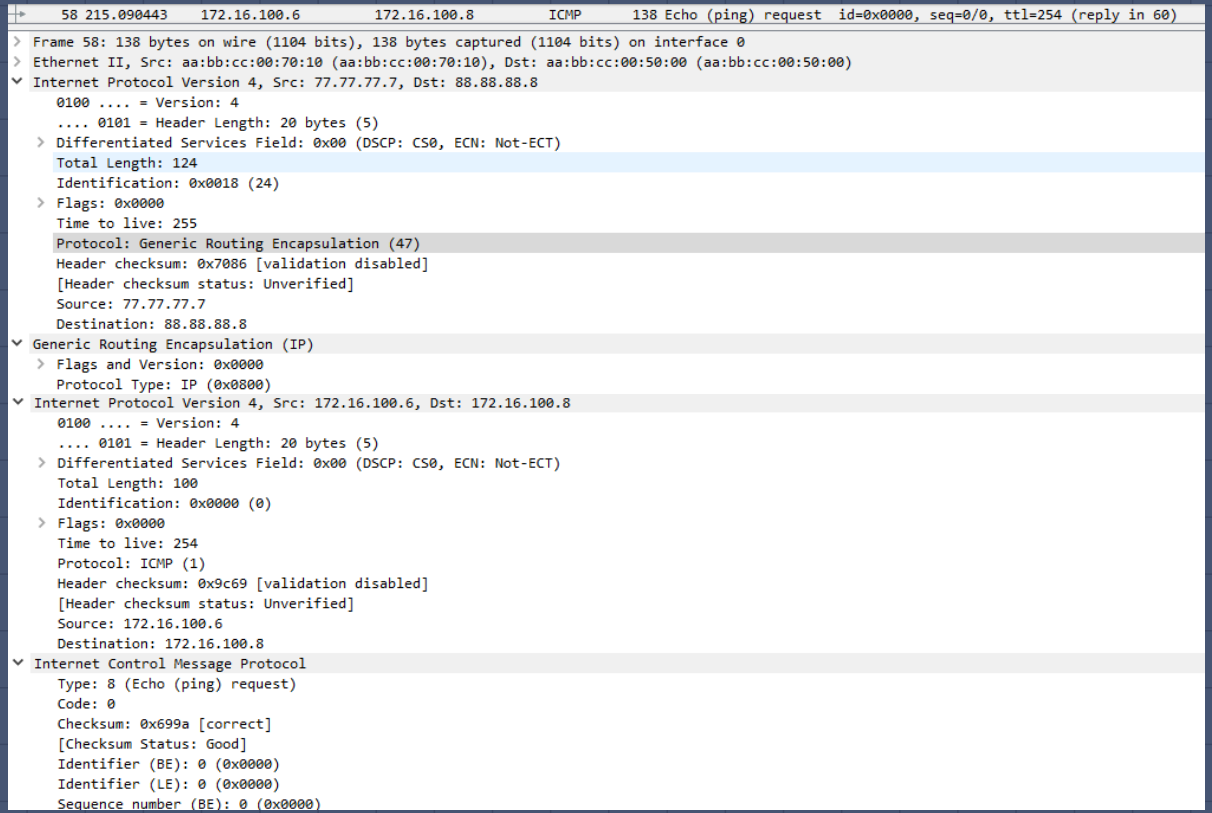

```
    58 215.090443    172.16.100.6         172.16.100.8         ICMP      138 Echo (ping) request  id=0x0000, seq=0/0, ttl=254 (reply in 60)
> Frame 58: 138 bytes on wire (1104 bits), 138 bytes captured (1104 bits) on interface 0
> Ethernet II, Src: aa:bb:cc:00:70:10 (aa:bb:cc:00:70:10), Dst: aa:bb:cc:00:50:00 (aa:bb:cc:00:50:00)
v Internet Protocol Version 4, Src: 77.77.77.7, Dst: 88.88.88.8
    0100 .... = Version: 4
    .... 0101 = Header Length: 20 bytes (5)
  > Differentiated Services Field: 0x00 (DSCP: CS0, ECN: Not-ECT)
    Total Length: 124
    Identification: 0x0018 (24)
  > Flags: 0x0000
    Time to live: 255
    Protocol: Generic Routing Encapsulation (47)
    Header checksum: 0x7086 [validation disabled]
    [Header checksum status: Unverified]
    Source: 77.77.77.7
    Destination: 88.88.88.8
v Generic Routing Encapsulation (IP)
  > Flags and Version: 0x0000
    Protocol Type: IP (0x0800)
v Internet Protocol Version 4, Src: 172.16.100.6, Dst: 172.16.100.8
    0100 .... = Version: 4
    .... 0101 = Header Length: 20 bytes (5)
  > Differentiated Services Field: 0x00 (DSCP: CS0, ECN: Not-ECT)
    Total Length: 100
    Identification: 0x0000 (0)
  > Flags: 0x0000
    Time to live: 254
    Protocol: ICMP (1)
    Header checksum: 0x9c69 [validation disabled]
    [Header checksum status: Unverified]
    Source: 172.16.100.6
    Destination: 172.16.100.8
v Internet Control Message Protocol
    Type: 8 (Echo (ping) request)
    Code: 0
    Checksum: 0x699a [correct]
    [Checksum Status: Good]
    Identifier (BE): 0 (0x0000)
    Identifier (LE): 0 (0x0000)
    Sequence number (BE): 0 (0x0000)
```

Multipoint GRE

```
interface Tunnel100
 ip address 172.16.100.8 255.255.255.128
 no ip redirects
 tunnel source Ethernet0/0
 tunnel mode gre multipoint
```

- Multipoint GRE works the same as traditional GRE with some exceptions
- With multipoint GRE, there is no tunnel destination configured because the same tunnel can have multiple destinations
- Multipoint GRE relies on another mechanism to determine tunnel destinations such as NHRP in the DMVPN model
- With multipoint GRE, the same logical tunnel can have multiple endpoints instead of needing a separate tunnel interface per endpoint
- The use of multipoint GRE is central to the scalability and dynamic nature of DMVPN

NHRP and DMVPN

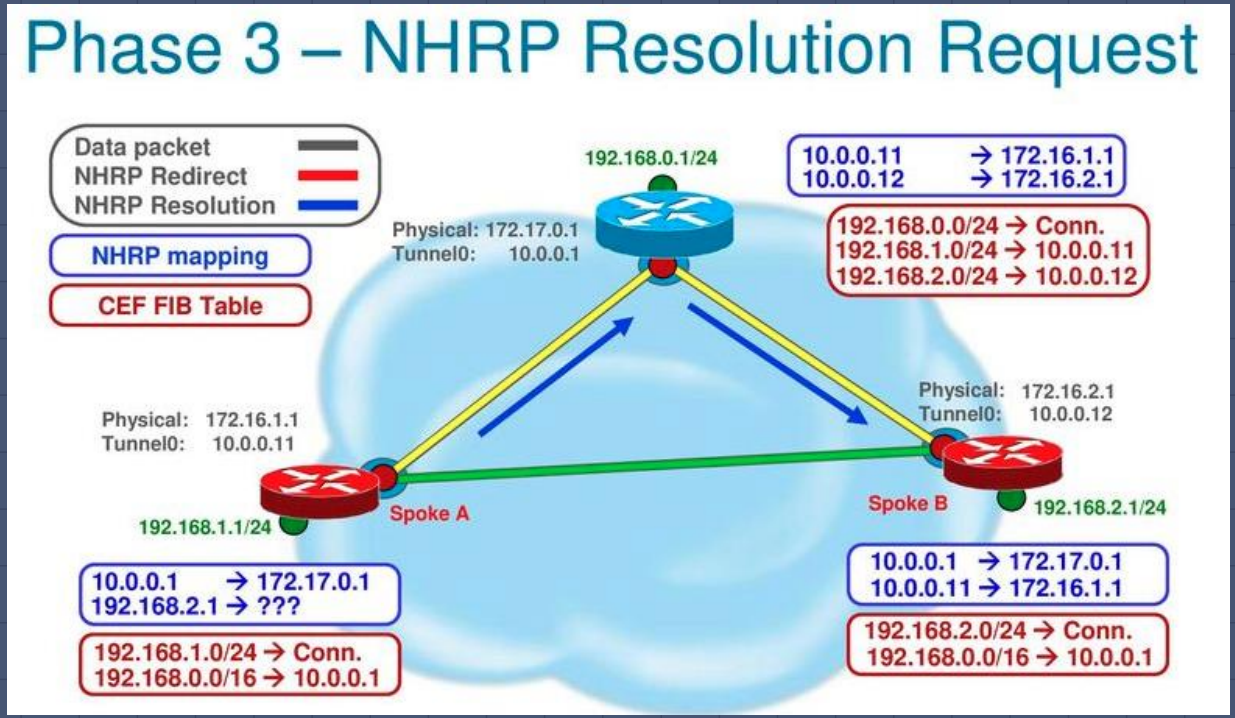

Next-Hop Resolution Protocol

- NHRP is a protocol which provides a function similar to ARP in broadcast networks, but across a non-broadcast medium
- By having hub devices in the DMVPN topology configured as Next-Hop Servers, and spoke devices configured as Next-Hop Clients, a mechanism exists to dynamically learn underlay addresses

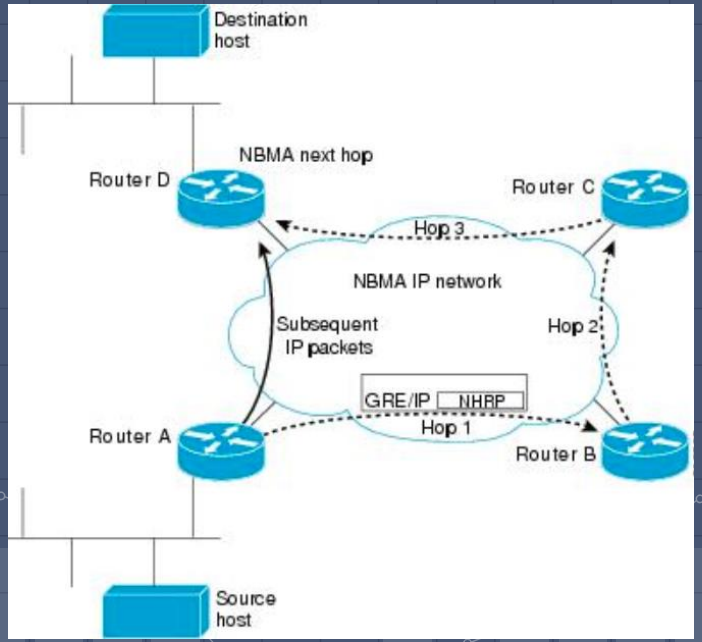

NHRP: What is it Good For?

- Facilitate direct spoke-to-spoke traffic flows to avoid hairpin or suboptimal routing
- Extends scalability of DMVPN by ensuring that hub devices do not need separate configuration for each spoke device
- Provides the glue that allows multipoint GRE tunnels to be built

NHRP: How Does it Work?

- Typically, at least one router is configured as the Next-Hop Server, which will register the underlay/overlay addresses of the Next-Hop Clients and keep a table of registrations. In the case of DMVPN, the hub is configured as the server and spokes as clients.
- Spoke routers are configured with the underlay and overlay addresses of the NHS. There can be multiple NHS configured, as well as NHS priority and other NHRP features to aid in NHS selection.
- Spoke routers will attempt to communicate with the configured NHS and register their underlay/overlay addresses to that NHS. If the spoke router needs to build a GRE multipoint tunnel to an another spoke, it will request the underlay information from the NHS which matches the overlay address needed to route traffic.

NHRP: Let's Book a Flight

NHRP is similar to booking a flight on a travel website

Airlines register their flights and destinations with the travel website, and when a customer wants to go to a destination, they query the travel website for the best flights to get there

Since the airlines registered their flights, the travel website can return best matches to the customer for trip planning

NHS Checklist

- ✓ NHRP Multicast Mapping – Replicated unicast to support pseudo-multicast

```
ip nhrp map multicast dynamic
```

- ✓ NHRP Network ID – Differentiates NHRP NBMA networks

```
ip nhrp network-id 100
```

- ✓ NHRP Authentication (Optional) – Authenticates spokes/hubs

```
ip nhrp authentication securepw
```

- ✓ NHRP Redirect (Optional) – Similar to IP redirect, facilitates spoke-to-spoke traffic flows

```
ip nhrp redirect
```

Configuration required on the hub is minimal by design in order to be as dynamic as possible

NHC Checklist

- ✓ NHRP Multicast Mapping – Map multicast traffic to NHS underlay address

```
ip nhrp map multicast 77.77.77.7
```

- ✓ NHRP Network ID – Differentiates NHRP NBMA networks, must match NHS

```
ip nhrp network-id 100
```

- ✓ NHRP NHS Mapping – Static mapping of NHS underlay IP address to overlay IP address

```
ip nhrp map 172.16.100.1 77.77.77.7
ip nhrp nhs 172.16.100.1
```

- ✓ NHRP Authentication (Optional) – Authenticates spokes/hubs

```
ip nhrp authentication securepw
```

- ✓ NHRP Shortcut (Optional) – Works in conjunction with NHRP Redirect to facilitate spoke-to-spoke traffic flows

```
ip nhrp shortcut
```

NHS Registration Dissected

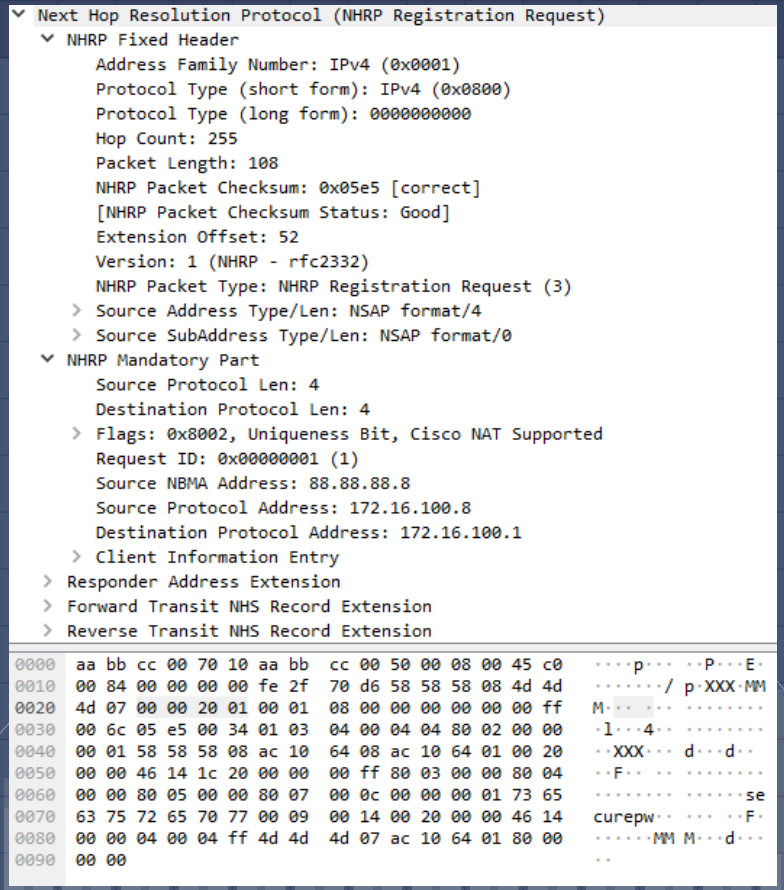

```
▼ Next Hop Resolution Protocol (NHRP Registration Request)
  ▼ NHRP Fixed Header
      Address Family Number: IPv4 (0x0001)
      Protocol Type (short form): IPv4 (0x0800)
      Protocol Type (long form): 0000000000
      Hop Count: 255
      Packet Length: 108
      NHRP Packet Checksum: 0x05e5 [correct]
      [NHRP Packet Checksum Status: Good]
      Extension Offset: 52
      Version: 1 (NHRP - rfc2332)
      NHRP Packet Type: NHRP Registration Request (3)
    > Source Address Type/Len: NSAP format/4
    > Source SubAddress Type/Len: NSAP format/0
  ▼ NHRP Mandatory Part
      Source Protocol Len: 4
      Destination Protocol Len: 4
    > Flags: 0x8002, Uniqueness Bit, Cisco NAT Supported
      Request ID: 0x00000001 (1)
      Source NBMA Address: 88.88.88.8
      Source Protocol Address: 172.16.100.8
      Destination Protocol Address: 172.16.100.1
    > Client Information Entry
  > Responder Address Extension
  > Forward Transit NHS Record Extension
  > Reverse Transit NHS Record Extension
```

```
0000  aa bb cc 00 70 10 aa bb  cc 00 50 00 08 00 45 c0   ····p··· ··P···E·
0010  00 84 00 00 00 00 fe 2f  70 d6 58 58 58 08 4d 4d   ·······/ p·XXX·MM
0020  4d 07 00 00 20 01 00 01  08 00 00 00 00 00 00 ff   M··· ···· ········
0030  00 6c 05 e5 00 34 01 03  04 00 04 04 80 02 00 00   ·l···4·· ········
0040  00 01 58 58 58 08 ac 10  64 08 ac 10 64 01 00 20   ··XXX··· d···d··
0050  00 00 46 14 1c 20 00 00  00 ff 80 03 00 00 80 04   ··F·· ·· ········
0060  00 00 80 05 00 00 80 07  00 0c 00 00 00 01 73 65   ········ ······se
0070  63 75 72 65 70 77 00 09  00 14 00 20 00 00 00 46 14   curepw·· ··· ··F·
0080  00 00 04 00 04 ff 4d 4d  4d 07 ac 10 64 01 80 00   ·····MM M···d···
0090  00 00                                              ··
```

NHS Registration Dissected

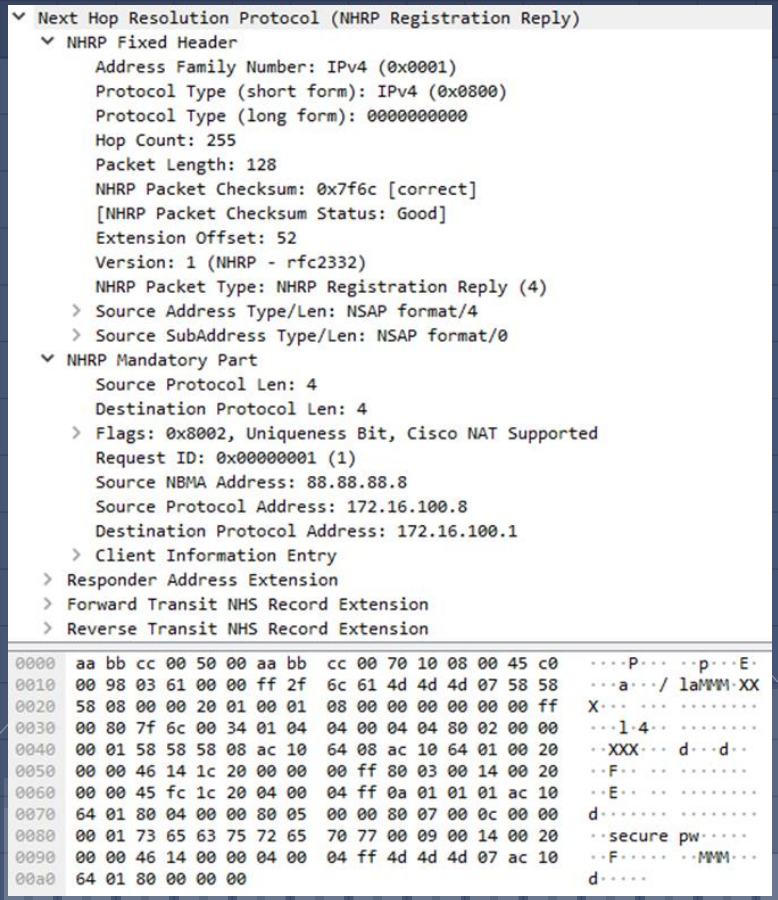

```
Next Hop Resolution Protocol (NHRP Registration Reply)
  NHRP Fixed Header
      Address Family Number: IPv4 (0x0001)
      Protocol Type (short form): IPv4 (0x0800)
      Protocol Type (long form): 0000000000
      Hop Count: 255
      Packet Length: 128
      NHRP Packet Checksum: 0x7f6c [correct]
      [NHRP Packet Checksum Status: Good]
      Extension Offset: 52
      Version: 1 (NHRP - rfc2332)
      NHRP Packet Type: NHRP Registration Reply (4)
    > Source Address Type/Len: NSAP format/4
    > Source SubAddress Type/Len: NSAP format/0
  NHRP Mandatory Part
      Source Protocol Len: 4
      Destination Protocol Len: 4
    > Flags: 0x8002, Uniqueness Bit, Cisco NAT Supported
      Request ID: 0x00000001 (1)
      Source NBMA Address: 88.88.88.8
      Source Protocol Address: 172.16.100.8
      Destination Protocol Address: 172.16.100.1
    > Client Information Entry
  > Responder Address Extension
  > Forward Transit NHS Record Extension
  > Reverse Transit NHS Record Extension
```

```
0000  aa bb cc 00 50 00 aa bb  cc 00 70 10 08 00 45 c0   ····P··· ··p···E·
0010  00 98 03 61 00 00 ff 2f  6c 61 4d 4d 4d 07 58 58   ···a···/ laMMM·XX
0020  58 08 00 00 20 01 00 01  08 00 00 00 00 00 00 ff   X··· ···· ········
0030  00 80 7f 6c 00 34 01 04  04 00 04 04 80 02 00 00   ···l·4·· ········
0040  00 01 58 58 58 08 ac 10  64 08 ac 10 64 01 00 20   ··XXX··· d···d··
0050  00 00 46 14 1c 20 00 00  00 ff 80 03 00 14 00 20   ··F·· ·· ········
0060  00 00 45 fc 1c 20 04 00  04 ff 0a 01 01 01 ac 10   ··E·· ·· ········
0070  64 01 80 04 00 00 80 05  00 00 80 07 00 0c 00 00   d··· ···· ········
0080  00 01 73 65 63 75 72 65  70 77 00 09 00 14 00 20   ··secure pw·····
0090  00 00 46 14 00 00 04 00  04 ff 4d 4d 4d 07 ac 10   ··F····· ··MMM···
00a0  64 01 80 00 00 00                                  d·····
```

Routing and DMVPN

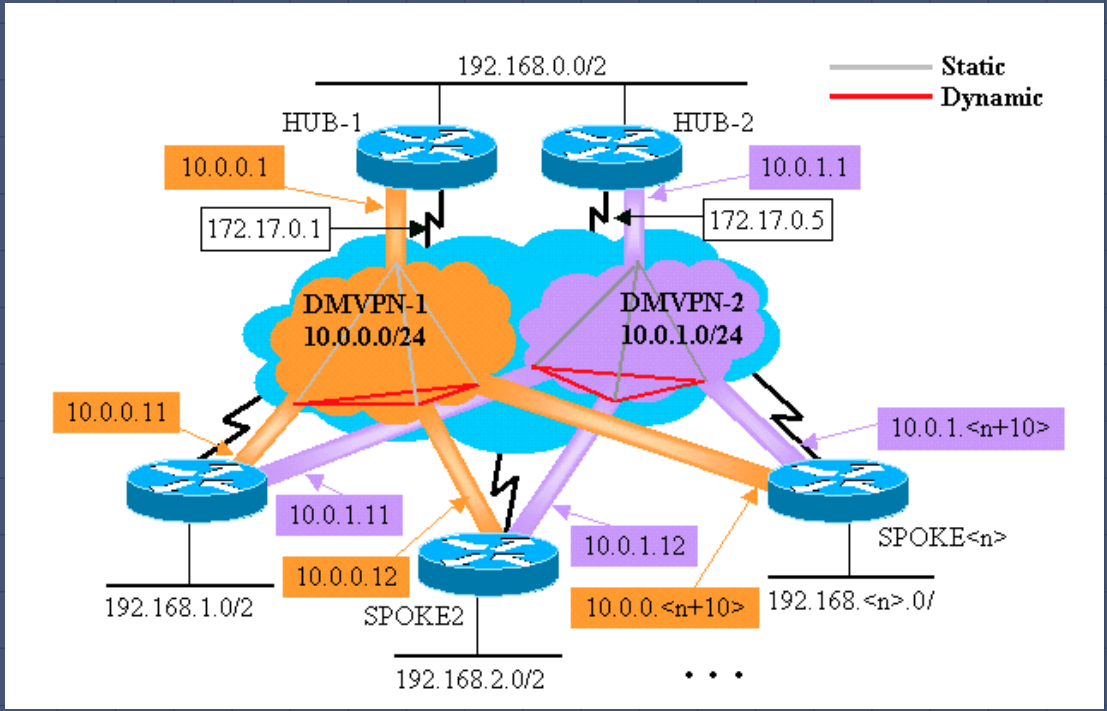

Routing With DMVPN

Routing With DMVPN carries specific design considerations depending on the routing protocol used.

- Will the routing protocol be using unicast or multicast to form neighbors?
- Does the routing protocol support different next-hop addresses?
- What per-neighbor metrics (if any) are supported?
- On the NHS/Hub routers, how will neighbors scale from a control plane perspective?

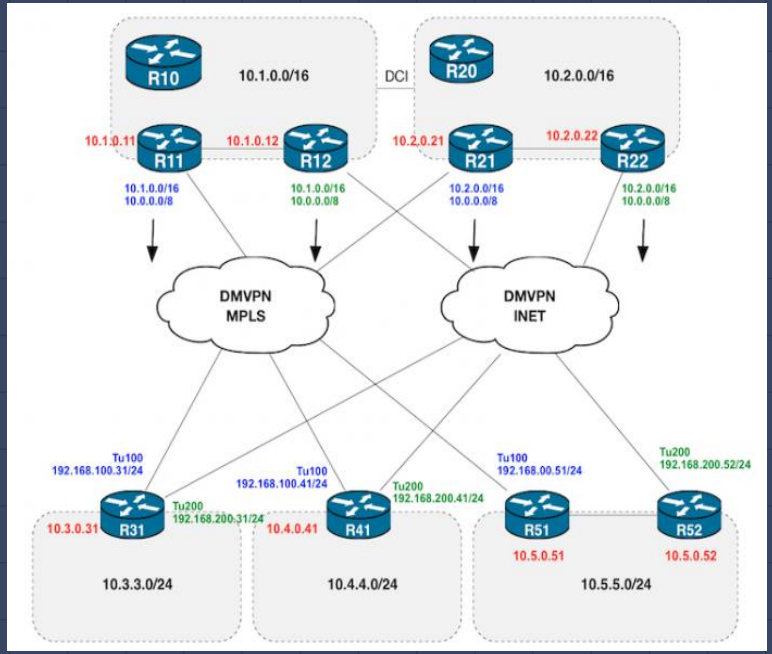

DMVPN Routing: How Does It Work?

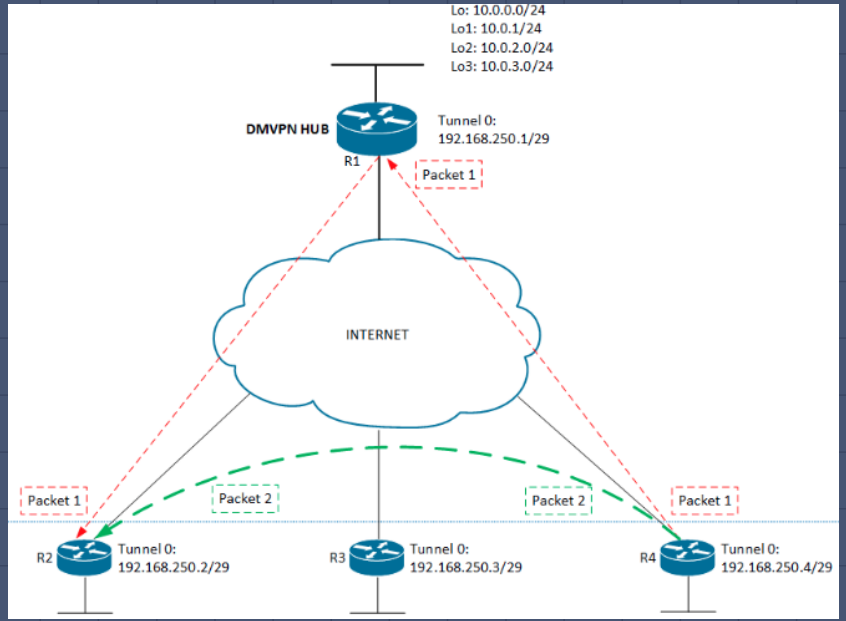

- DMVPN routing is a hub-and-spoke flow in the control plane
- Spokes will form routing adjacencies with hubs but not other spokes, and will learn route advertisements from the hub routers
- In the data plane, traffic will be initially sent to the hub, but subsequent flows will be sent directly from spoke to spoke via two methods:
  - NHRP Redirect/Shortcut will be used to trigger spoke-to-spoke traffic flow
  - Alternatively, the IP next-hop of advertised routes can be preserved from spoke to hub to spoke, in order to trigger spoke-to-spoke traffic flows

DMVPN Routing with EIGRP

- Hellos will be sent to the hub because of NHRP configuration. Because a hub may have thousands of neighbors, some timer tuning is needed to scale the control plane effectively
- Scaling the EIGRP Query Domain with DMVPN is important – This is usually accomplished with a summary-address toward spokes on the hub router and also by settings spokes as EIGRP stub
- From the hub perspective, metrics cannot be changed on a per-spoke basis without moving the spoke to an entirely different DMVPN tunnel interface. Offset-lists and delay can be applied only to the tunnel interface, which affects all spokes
- Disable split-horizon in order to allow route updates to propagate between spokes if not using summarization
- Consider using a different EIGRP AS for the DMVPN network for traffic engineering purposes

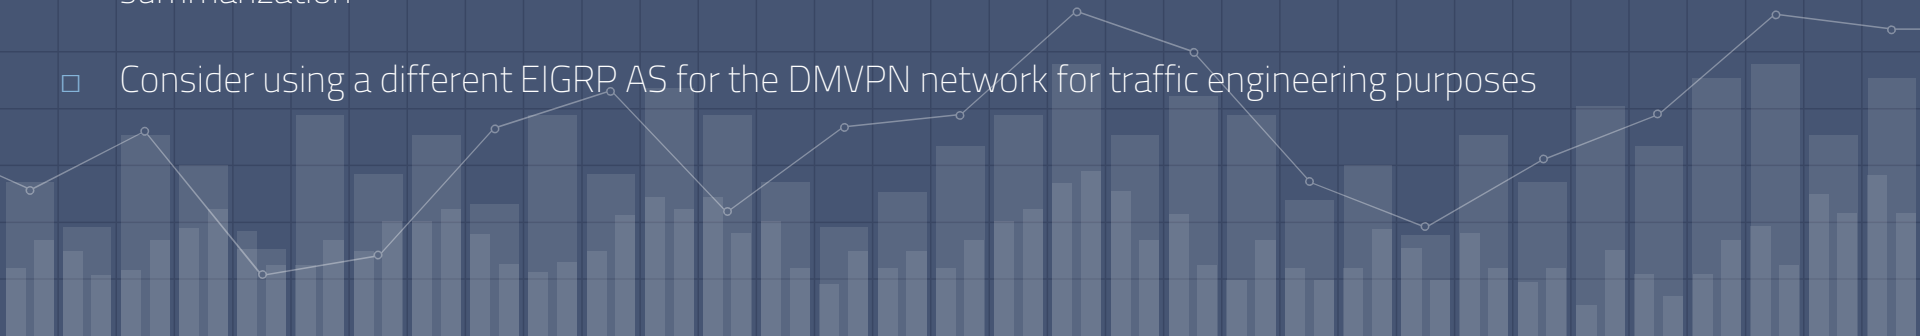

DMVPN Routing with OSPF

- Hellos will be sent to the hub because of NHRP configuration. Because a hub may have thousands of neighbors, some timer tuning is needed to scale the control plane effectively

- Under normal circumstances, all OSPF routers within an area share an identical copy of the link-state database, and with DMVPN, all routers are in the same area due to sharing the same IP subnet

- Because of this, scaling is a challenge. The DMVPN spoke routers should be totally stubby or totally not-so-stubby if possible to aid scaling of the LSDB

- OSPF network type is integral to DMVPN design
  - ***Point to Multipoint***: Increases routing table with host addressing but is preferred due to simple configuration
  - ***Point to Point / P2MP***: Spokes run P2P and Hubs run P2MP; OSPF timers must be adjusted
  - ***Broadcast / Non-Broadcast***: Not preferred, only needed if NHRP Redirect/Shortcut isn't supported
  - All network types except Broadcast support per-neighbor metric manipulation on the hub

- Routing changes on spoke routers can trigger an SPF recalculation for all DMVPN routers, for this reason OSPF is not preferred for large-scale DMVPN deployments

DMVPN Routing with BGP

- Largest routing decision to make with using BGP over DMVPN is whether to use iBGP or eBGP

- eBGP:
  - Uses BGP's built-in loop-prevention mechanism
  - AS-Path typically used to help path preference across DMVPN
  - Requires far more BGP configuration on hub per spoke
  - Each hub and spoke requires its own BGP AS, or configuration on each hub to break loop prevention in order to advertise spoke routes to other spokes

- iBGP:
  - Better scaling for neighbors (dynamic BGP neighbors possible on hub, simple configuration)
  - Local Preference can be used to set path preference across the DMVPN
  - Hub will typically be a route-reflector and spokes will be clients
  - All spokes peer with hub over BGP using next-hop-self to modify next-hop
  - Hub doesn't require next-hop-self, but if configured, it can still trigger spoke/spoke tunnels using NHRP Redirect/Shortcut

Routing: Air Traffic Control

In the DMVPN model, routing protocols act as air traffic control in the control plane

Planes do not have to reach their destinations through the ATC towers, but will ask ATC on which flight plans are available

ATC directs the planes on the best paths to take to reach the destination

Crypto IPSEC and DMVPN

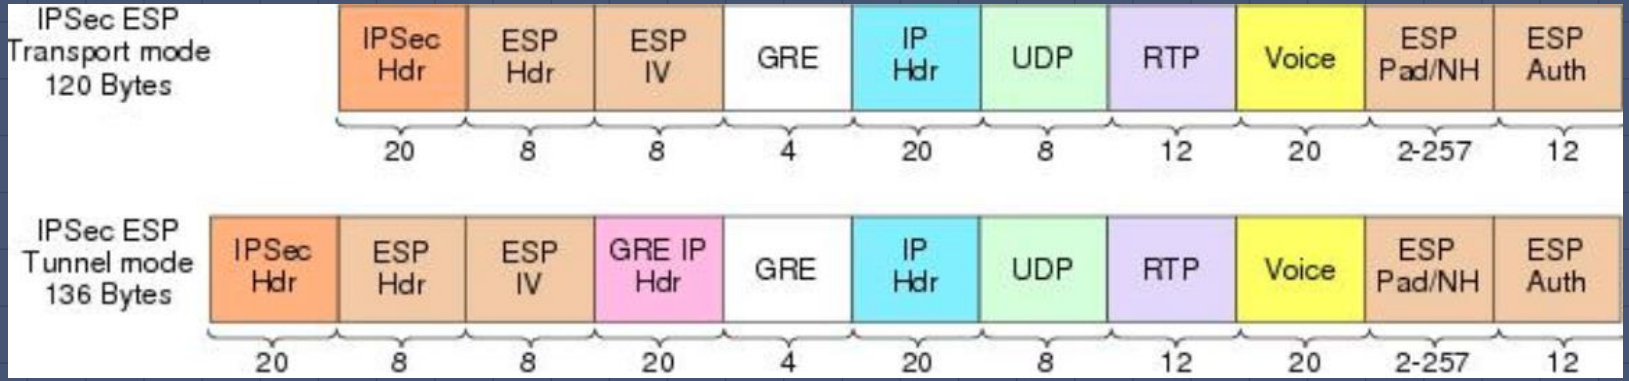

IPSEC With DMVPN

Because the IP transport for DMVPN is often over a third-party unsecure network (ie, the Internet) there is a need to encrypt traffic. For this reason, Internet Key Exchange and IP Security were created

- IPSEC transport mode preferred to tunnel mode in DMVPN because GRE is already providing tunneling
- No need for static crypto maps with DMVPN because crypto is based on traffic entering the GRE tunnel
- IPSEC tunnel protection profiles simply encryption configuration

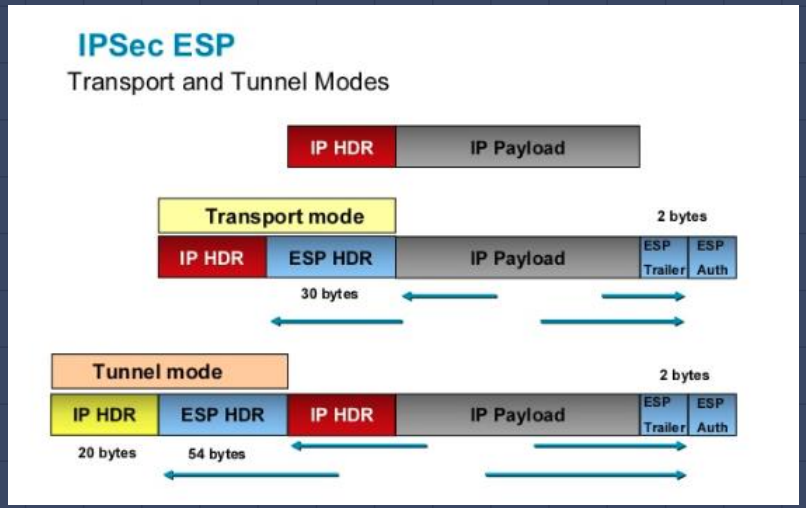

IPSEC: How Does It Work?

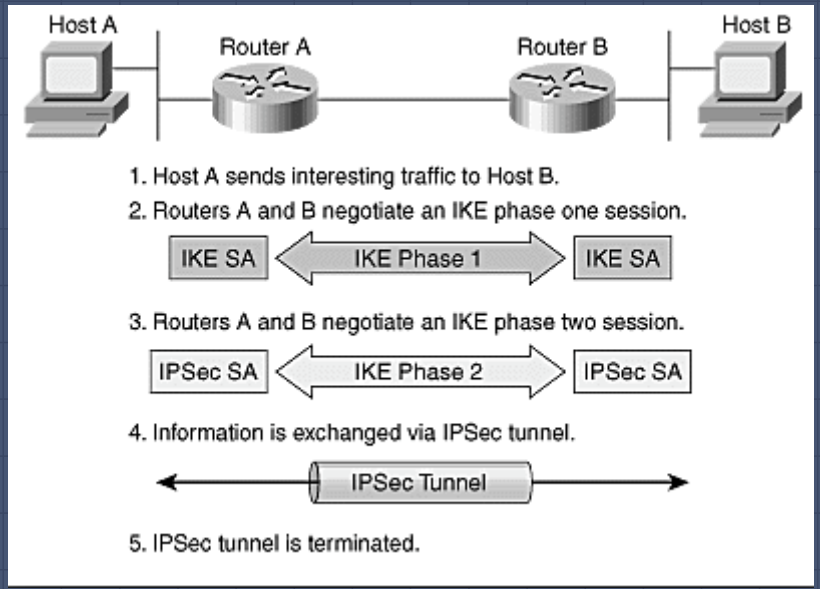

- IPSEC encrypts the entire payload of a packet in order to provide confidentiality and integrity to the data within
- IPSEC parameters must be configured on each router in order to choose common encryption methods
- Authentication can be provided by pre-shared keys or digital certificates
- Because the establishment of an IPSEC tunnel starts with unencrypted data, there is first an IKE Phase 1 negotiation to establish the secure channel to negotiate the IPSEC tunnel, followed by an IKE Phase 2 negotiation to create the tunnel itself

IPSEC in Action: Setup

No.	Time	Source	Destination	Protocol	Length	Info
1	0.000000	aa:bb:cc:00:50:10	aa:bb:cc:00:50:10	LOOP	60	Reply
2	4.710130	aa:bb:cc:00:50:10	CDP/VTP/DTP/PAgP/UD…	CDP	386	Device ID: ISP Port ID: Ethernet0/1
3	6.674096	aa:bb:cc:00:60:00	aa:bb:cc:00:60:00	LOOP	60	Reply
4	8.334172	aa:bb:cc:00:60:00	Broadcast	ARP	60	Who has 66.66.66.5? Tell 66.66.66.6
5	8.334556	aa:bb:cc:00:50:10	aa:bb:cc:00:60:00	ARP	60	66.66.66.5 is at aa:bb:cc:00:50:10
6	10.005085	aa:bb:cc:00:50:10	aa:bb:cc:00:50:10	LOOP	60	Reply
7	16.682775	aa:bb:cc:00:60:00	aa:bb:cc:00:60:00	LOOP	60	Reply
8	18.341843	66.66.66.6	77.77.77.7	ISAKMP	210	Identity Protection (Main Mode)
9	18.343160	77.77.77.7	66.66.66.6	ISAKMP	150	Identity Protection (Main Mode)
10	18.349362	66.66.66.6	77.77.77.7	ISAKMP	326	Identity Protection (Main Mode)
11	18.385749	77.77.77.7	66.66.66.6	ISAKMP	346	Identity Protection (Main Mode)
12	18.393759	66.66.66.6	77.77.77.7	ISAKMP	154	Identity Protection (Main Mode)
13	18.395101	77.77.77.7	66.66.66.6	ISAKMP	122	Identity Protection (Main Mode)
14	18.399955	66.66.66.6	77.77.77.7	ISAKMP	250	Quick Mode
15	18.400991	77.77.77.7	66.66.66.6	ISAKMP	250	Quick Mode
16	18.401705	66.66.66.6	77.77.77.7	ISAKMP	106	Quick Mode
17	18.410915	66.66.66.6	77.77.77.7	ESP	206	ESP (SPI=0xd1c9c3c7)
18	18.411862	77.77.77.7	66.66.66.6	ESP	222	ESP (SPI=0x1d873980)
19	20.013881	aa:bb:cc:00:50:10	aa:bb:cc:00:50:10	LOOP	60	Reply
20	20.044853	66.66.66.6	77.77.77.7	ESP	158	ESP (SPI=0xd1c9c3c7)
21	20.055072	77.77.77.7	66.66.66.6	ESP	158	ESP (SPI=0x1d873980)
22	20.055108	77.77.77.7	66.66.66.6	ESP	126	ESP (SPI=0x1d873980)

```
    Reserved: 00
    Payload length: 60
    Domain of interpretation: IPSEC (1)
  > Situation: 00000001
  v Payload: Proposal (2) # 1
      Next payload: NONE / No Next Payload  (0)
      Reserved: 00
      Payload length: 48
      Proposal number: 1
      Protocol ID: ISAKMP (1)
      SPI Size: 0
      Proposal transforms: 1
    v Payload: Transform (3) # 1
        Next payload: NONE / No Next Payload  (0)
        Reserved: 00
        Payload length: 40
        Transform number: 1
        Transform ID: KEY_IKE (1)
        Reserved: 0000
      > IKE Attribute (t=1,l=2): Encryption-Algorithm: AES-CBC
      > IKE Attribute (t=14,l=2): Key-Length: 128
      > IKE Attribute (t=2,l=2): Hash-Algorithm: SHA
      > IKE Attribute (t=4,l=2): Group-Description: Alternate 1024-bit MODP group
      > IKE Attribute (t=3,l=2): Authentication-Method: Pre-shared key
      > IKE Attribute (t=11,l=2): Life-Type: Seconds
      > IKE Attribute (t=12,l=4): Life-Duration: 86400
v Payload: Vendor ID (13) : RFC 3947 Negotiation of NAT-Traversal in the IKE
```

- Tunnel Setup: Traffic to be encrypted is forwarded to the GRE tunnel
- IKE Phase 1: Authenticate peer and negotiate IKE security associations, set up a secure channel for IKE Phase 2
- IKE Phase 2: Negotiate SA parameters, set up matching IPSEC SA for peer
- Data transfer, encryption based on IPSEC parameters negotiated and active SA
- Tunnel Termination: SA terminates through deletion or timeout after traffic flow completes

IPSEC in Action: Ping Between Spokes

No.	Time	Source	Destination	Protocol	Length	Info
1	0.000000	172.16.100.6	224.0.0.10	EIGRP	98	Hello
2	0.136214	aa:bb:cc:00:50:10	aa:bb:cc:00:50:10	LOOP	60	Reply
3	3.402760	172.16.100.1	224.0.0.10	EIGRP	98	Hello
4	3.626647	aa:bb:cc:00:60:00	aa:bb:cc:00:60:00	LOOP	60	Reply
5	3.636046	aa:bb:cc:00:60:00	CDP/VTP/DTP/PAgP/UD…	CDP	380	Device ID: DMVPN-SPOKE Port ID: Ethernet0/0
6	4.626853	172.16.100.6	224.0.0.10	EIGRP	98	Hello
7	8.186524	172.16.100.1	224.0.0.10	EIGRP	98	Hello
8	9.472363	172.16.100.6	224.0.0.10	EIGRP	98	Hello
9	10.137585	aa:bb:cc:00:50:10	aa:bb:cc:00:50:10	LOOP	60	Reply
10	10.523122	1.1.1.1	2.2.2.2	ICMP	138	Echo (ping) request id=0x0002, seq=0/0, ttl=255 (reply in 12)
11	10.524485	77.77.77.7	66.66.66.6	NHRP	138	NHRP Traffic Indication
12	10.525277	2.2.2.2	1.1.1.1	ICMP	138	Echo (ping) reply id=0x0002, seq=0/0, ttl=254 (request in 10)
13	10.525626	1.1.1.1	2.2.2.2	ICMP	138	Echo (ping) request id=0x0002, seq=1/256, ttl=255 (reply in 14)
14	10.526979	2.2.2.2	1.1.1.1	ICMP	138	Echo (ping) reply id=0x0002, seq=1/256, ttl=254 (request in 13)
15	10.527170	1.1.1.1	2.2.2.2	ICMP	138	Echo (ping) request id=0x0002, seq=2/512, ttl=255 (reply in 16)
16	10.528429	2.2.2.2	1.1.1.1	ICMP	138	Echo (ping) reply id=0x0002, seq=2/512, ttl=254 (request in 15)
17	10.528613	1.1.1.1	2.2.2.2	ICMP	138	Echo (ping) request id=0x0002, seq=3/768, ttl=255 (reply in 18)
18	10.529783	2.2.2.2	1.1.1.1	ICMP	138	Echo (ping) reply id=0x0002, seq=3/768, ttl=254 (request in 17)
19	10.529958	1.1.1.1	2.2.2.2	ICMP	138	Echo (ping) request id=0x0002, seq=4/1024, ttl=255 (reply in 20)
20	10.531141	2.2.2.2	1.1.1.1	ICMP	138	Echo (ping) reply id=0x0002, seq=4/1024, ttl=254 (request in 19)
21	10.538144	66.66.66.6	77.77.77.7	NHRP	126	NHRP Resolution Request, ID=4
22	10.539809	88.88.88.8	66.66.66.6	NHRP	174	NHRP Resolution Reply, ID=4, Code=Success
23	10.539843	77.77.77.7	66.66.66.6	NHRP	146	NHRP Resolution Request, ID=4
24	10.550943	66.66.66.6	88.88.88.8	NHRP	174	NHRP Resolution Reply, ID=4, Code=Success
25	12.593670	172.16.100.1	224.0.0.10	EIGRP	98	Hello

```
> Frame 1: 98 bytes on wire (784 bits), 98 bytes captured (784 bits) on interface 0
> Ethernet II, Src: aa:bb:cc:00:60:00 (aa:bb:cc:00:60:00), Dst: aa:bb:cc:00:50:10 (aa:bb:cc:00:50:10)
> Internet Protocol Version 4, Src: 66.66.66.6, Dst: 77.77.77.7
> Generic Routing Encapsulation (IP)
> Internet Protocol Version 4, Src: 172.16.100.6, Dst: 224.0.0.10
> Cisco EIGRP
```

```
0000  aa bb cc 00 50 10 aa bb  cc 00 60 00 08 00 45 c0   ····P··· ··`···E·
0010  00 54 07 dc 00 00 ff 2f  94 42 42 42 42 06 4d 4d   ·T·····/ ·BBBB·MM
0020  4d 07 00 00 08 00 45 c0  00 3c 07 c6 00 00 01 58   M·····E· ·<·····X
```

No.	Time	Source	Destination	Protocol	Length	Info
1	0.000000	aa:bb:cc:00:60:00	aa:bb:cc:00:60:00	LOOP	60	Reply
2	1.237295	77.77.77.7	66.66.66.6	ESP	158	ESP (SPI=0x6bafc4fd)
3	1.961940	66.66.66.6	77.77.77.7	ESP	158	ESP (SPI=0xbdc9acce)
4	5.836577	77.77.77.7	66.66.66.6	ESP	158	ESP (SPI=0x6bafc4fd)
5	6.391852	aa:bb:cc:00:50:10	aa:bb:cc:00:50:10	LOOP	60	Reply
6	6.431113	66.66.66.6	77.77.77.7	ESP	158	ESP (SPI=0xbdc9acce)
7	10.008526	aa:bb:cc:00:60:00	aa:bb:cc:00:60:00	LOOP	60	Reply
8	10.240782	66.66.66.6	88.88.88.8	ESP	182	ESP (SPI=0x55130274)
9	10.242004	88.88.88.8	66.66.66.6	ESP	182	ESP (SPI=0xaff87615)
10	10.246987	66.66.66.6	88.88.88.8	ESP	182	ESP (SPI=0x55130274)
11	10.248051	88.88.88.8	66.66.66.6	ESP	182	ESP (SPI=0xaff87615)
12	10.252566	66.66.66.6	88.88.88.8	ESP	182	ESP (SPI=0x55130274)
13	10.259555	88.88.88.8	66.66.66.6	ESP	182	ESP (SPI=0xaff87615)
14	10.264848	66.66.66.6	88.88.88.8	ESP	182	ESP (SPI=0x55130274)
15	10.267217	88.88.88.8	66.66.66.6	ESP	182	ESP (SPI=0xaff87615)
16	10.272671	66.66.66.6	88.88.88.8	ESP	182	ESP (SPI=0x55130274)
17	10.273700	88.88.88.8	66.66.66.6	ESP	182	ESP (SPI=0xaff87615)
18	10.694113	66.66.66.6	77.77.77.7	ESP	158	ESP (SPI=0xbdc9acce)
19	10.822637	77.77.77.7	66.66.66.6	ESP	158	ESP (SPI=0x6bafc4fd)

```
> Frame 1: 60 bytes on wire (480 bits), 60 bytes captured (480 bits) on interface 0
> Ethernet II, Src: aa:bb:cc:00:60:00 (aa:bb:cc:00:60:00), Dst: aa:bb:cc:00:60:00 (aa:bb:cc:00:60:00)
> Configuration Test Protocol (loopback)
> Data (40 bytes)
```

```
0000  aa bb cc 00 60 00 aa bb  cc 00 60 00 90 00 00 00   ····`··· ··`·····
0010  01 00 00 00 00 00 00 00  00 00 00 00 00 00 00 00   ········ ········
0020  00 00 00 00 00 00 00 00  00 00 00 00 00 00 00 00   ········ ········
```

IPSEC: Engage Stealth Mode

No aircraft is completely invisible, just as packets are not invisible on the wire

Stealth aircraft rely on low radar profile and active measures to minimize detection, just as encrypted packets rely on a strong encryption algorithm to protect the data inside

DMVPN Design Challenges

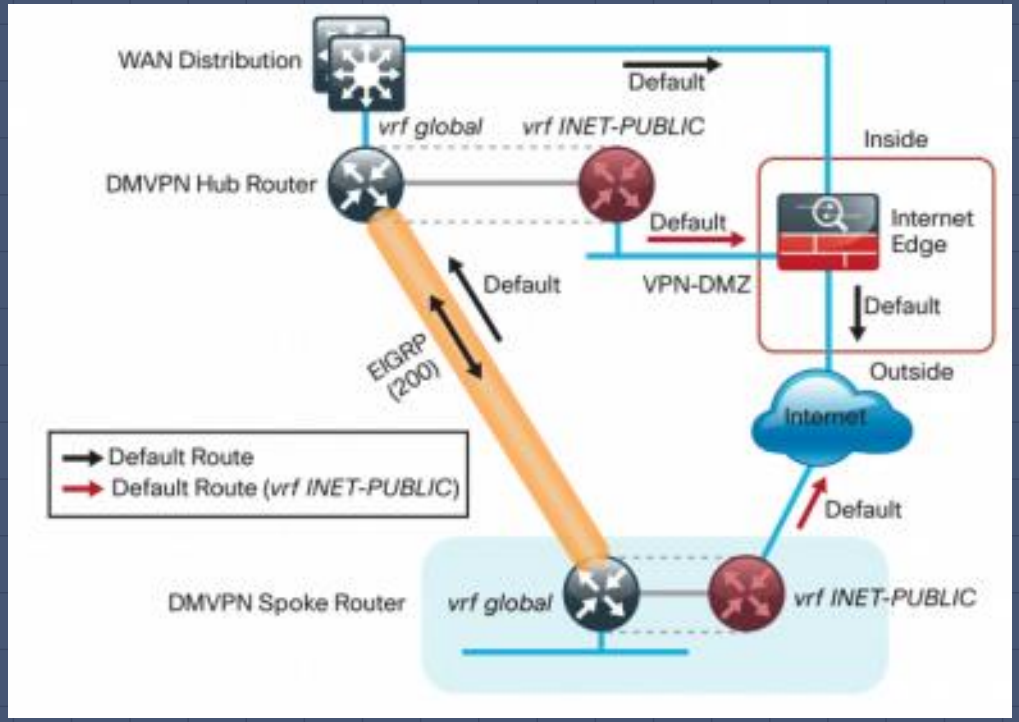

Solution: Front Door VRF

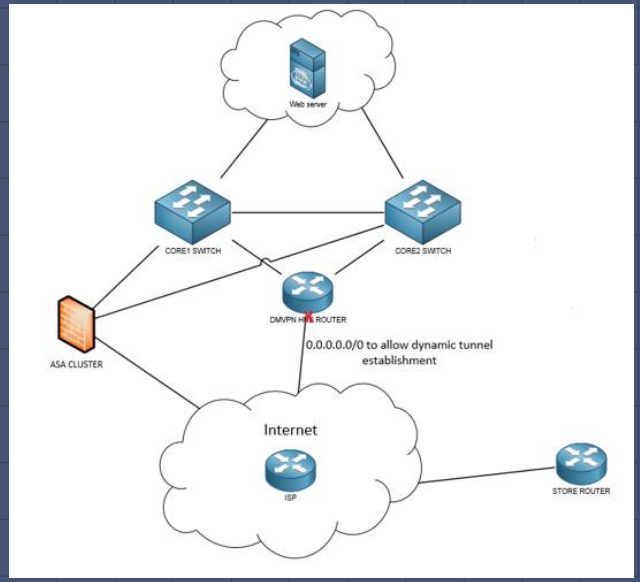

Problem: Underlay/Overlay Route Table Fusion

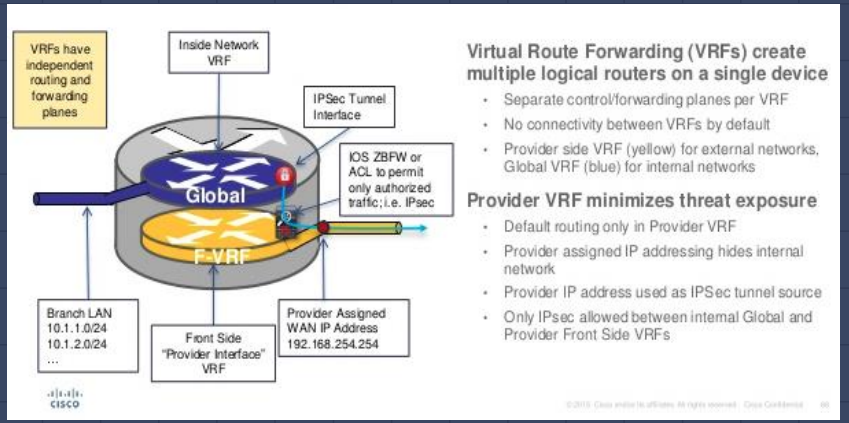

Front Door VRF: What Is It Good For?

- Needed to separate the routing table of the underlay and overlay networks
- Allows default routing on hub for dynamic tunnel establishment without overriding a preferred default route for spokes
- Allows underlay network traffic flow to be different than overlay network traffic flow if desired

Front Door VRF: Why Do We Need It?

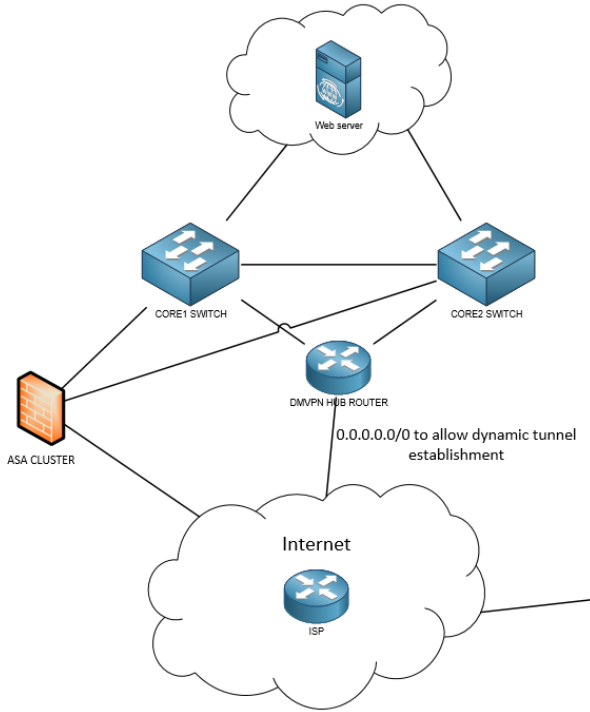

Default route sent from Data Center forces store to use Data Center (with firewall) for Internet transport

Store traffic lands on DMVPN router, which has default route statically set to Internet – hairpins on DMVPN router to Internet, does not reach firewall to be inspected/NAT – Store can not reach internet

Front Door VRF: Why Do We Need It?

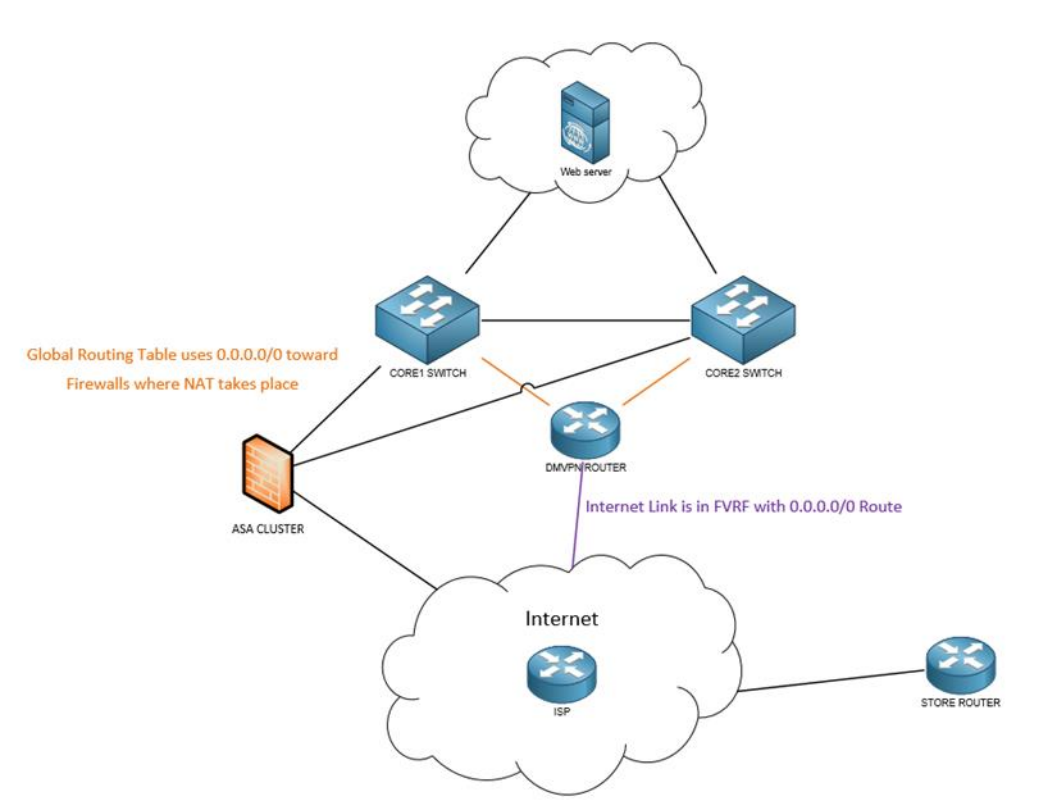

Problem: Route Recursion Causes Tunnel to Bounce

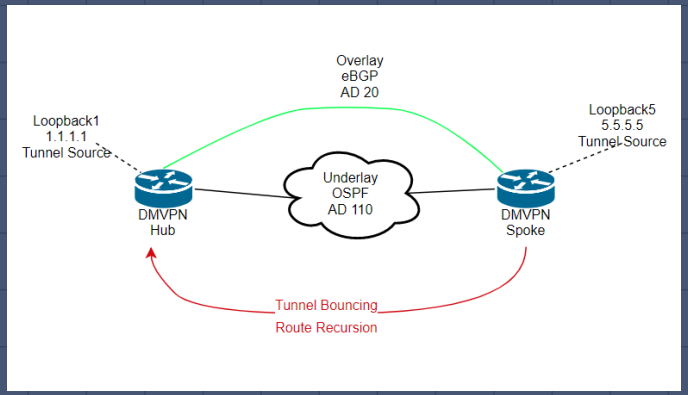

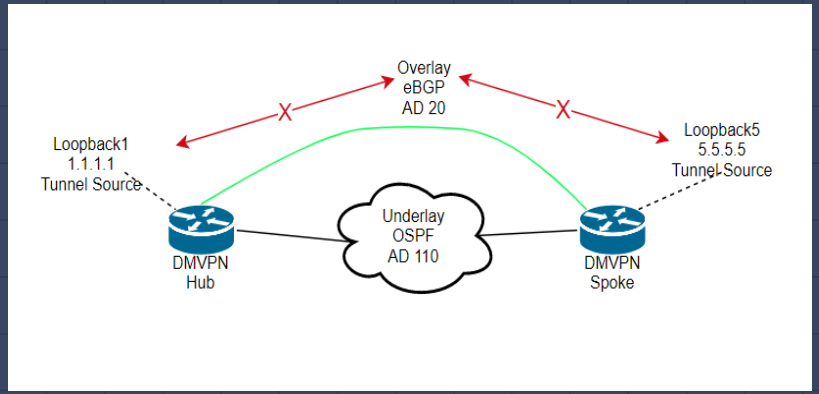

Solution: Selective Advertisement Through Tunnel

Route Recursion in DMVPN Explained

- In this example, DMVPN tunnels are built using Loopback as tunnel source
- Underlay IGP provides reachability between tunnel endpoints
- Overlay eBGP advertises all networks (including tunnel endpoints)
- Routers learn path to tunnel endpoints with better AD through the tunnel and update RIB

Route Recursion in DMVPN Explained

- Issue is that the underlay address of the tunnel endpoint is now thought to be best reachable through overlay tunnel
- When CEF does route-lookup for GRE-encapsulated packet, exit interface will be the GRE tunnel itself
- Tunnel will be torn down because underlay reachability is lost, until route convergence causes the best path to tunnel endpoint to be through the underlay again
- When tunnel is rebuilt and routing adjacencies recovered, the issue will repeat until the tunnel endpoints are no longer advertised through the tunnel, longest match is through underlay only, or AD is changed to prefer underlay instead

Problem: One or More Devices Using NAT

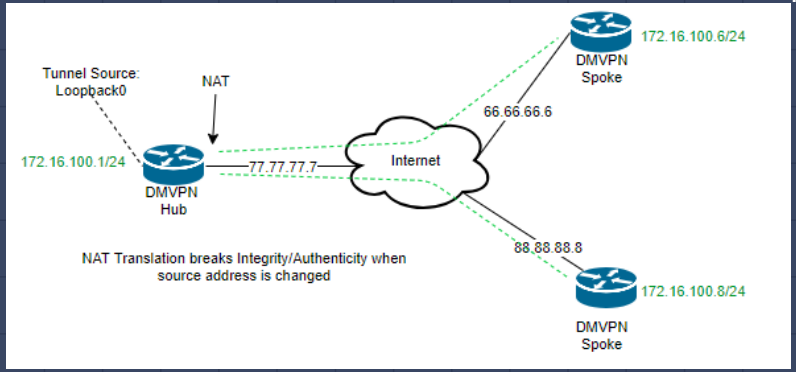

Solution: NAT Traversal (NAT-T)

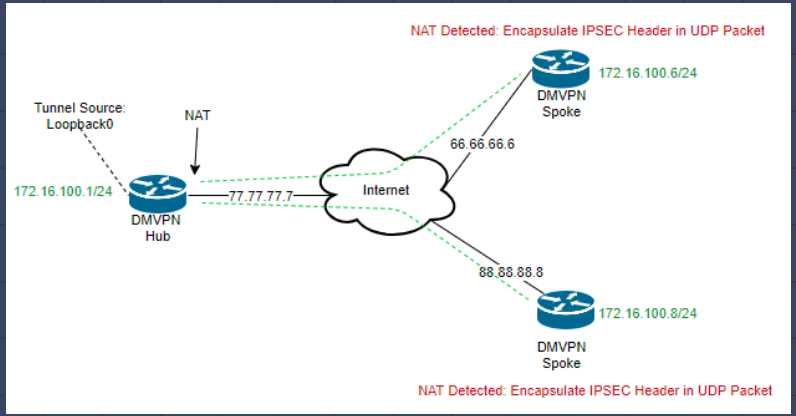

IPSEC NAT Traversal in DMVPN Explained

- Since NAT changes the source address of a packet, it breaks the integrity/authenticity principles of IPSEC
- DMVPN routers can detect if one or more endpoints are behind a NAT when starting the crypto setup between them by checking the result of a NAT-T probe, which includes a hash of the source/destination IP/port.
- If the NAT-T probe's hashes do not match on both ends, it is understood that NAT took place somewhere along the path and NAT Traversal is needed
- If NAT-T is to be used, IPSEC packets are encapsulated in a UDP packet using port 4500, which allows the NAT device to process the packet without touching the IPSEC header within and thereby retaining authenticity/integrity

DMVPN Designs

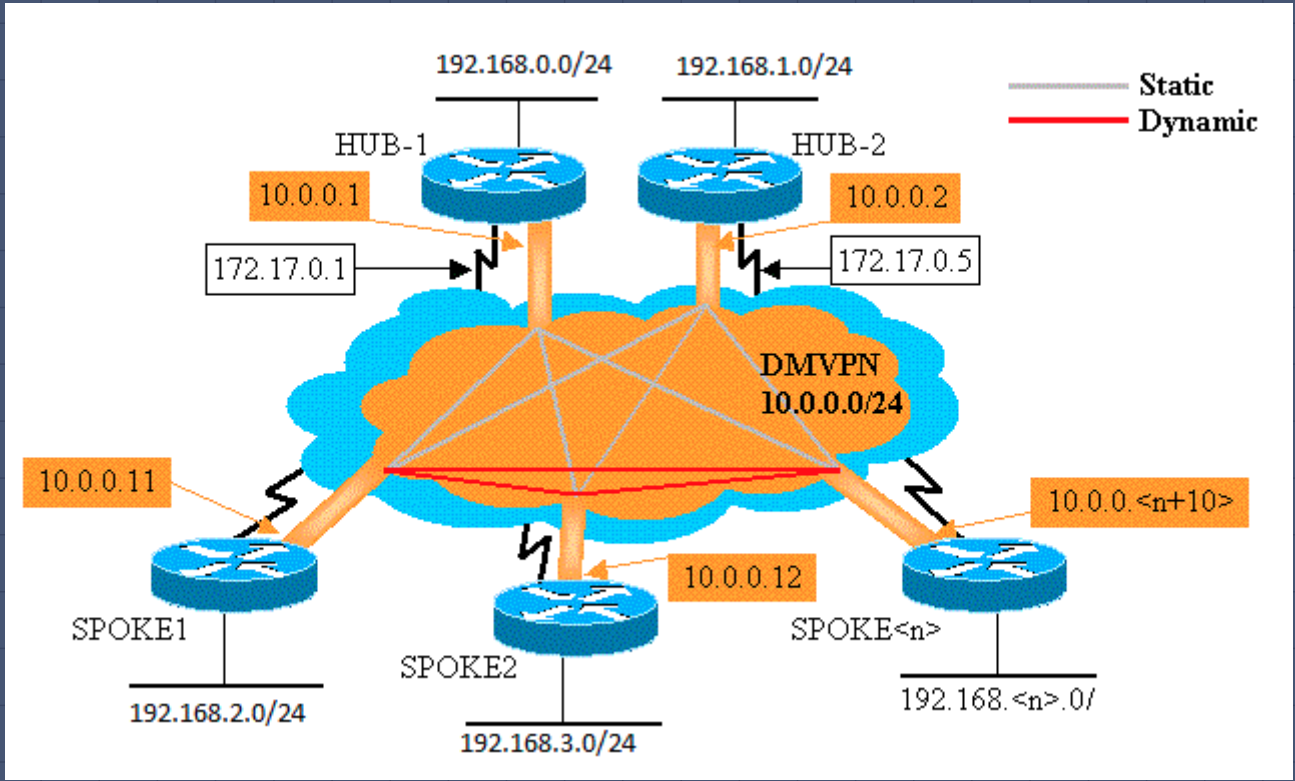

Single Hub DMVPN

- Benefits:
- Simple DMVPN Configuration
- Simple routing configuration

- Weaknesses:
- Single point of failure
- One hub must have all control plane adjacencies
- Any traffic destined behind the hub has only a single path to take, could lead to congestion if spokes all need resources behind hub at same time

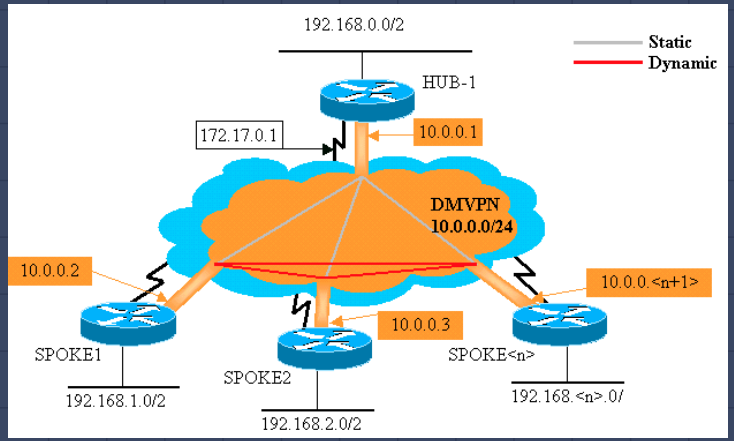

Single Hub Dual Cloud DMVPN

- Benefits:
- Redundant paths
- Easier to set path preference per-DMVPN
- Can load balance traffic destined from/to network behind hub
- Failover time is based on routing convergence only

- Weaknesses:
- Routing design can be complicated, loops possible
- More tunnel interfaces per spoke needed
- More IP addressing required, different subnets per-DMVPN

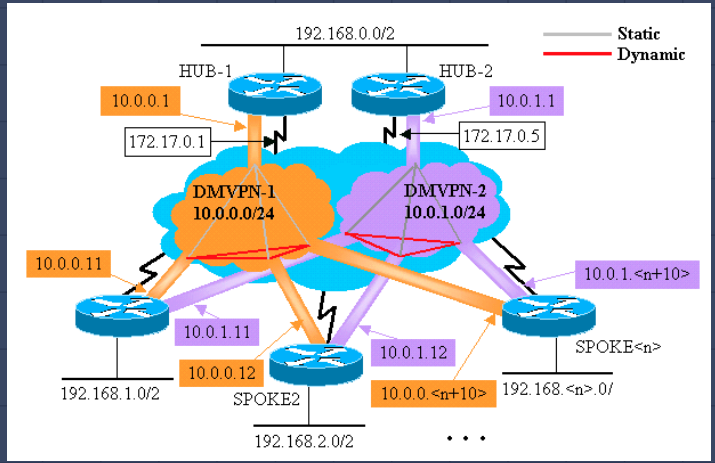

Dual Hub Single Cloud DMVPN

- Benefits:
  - Redundant DMVPN hubs
  - Simpler routing configuration
  - Can load balance traffic destined from/to network behind hubs
  - Less IP addressing required, fewer tunnel interfaces

- Weaknesses:
  - Path preference, if desired, affects all spokes because metric must be adjusted per hub
  - More complicated hub and spoke configuration

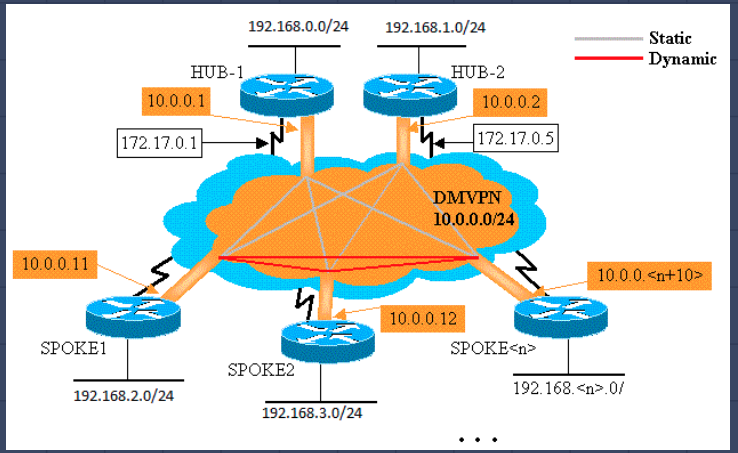

Hierarchical DMVPN

- Benefits:
- Excellent scale for control plane and spokes
- Regional hubs can connect with central hub instead of full mesh of hubs
- Potential to add multiple hubs per region or centrally to increase redundancy further

- Weaknesses:
- Very complicated design, only needed to scale to very high number of spokes
- More complicated hub configuration
- More complicated routing design

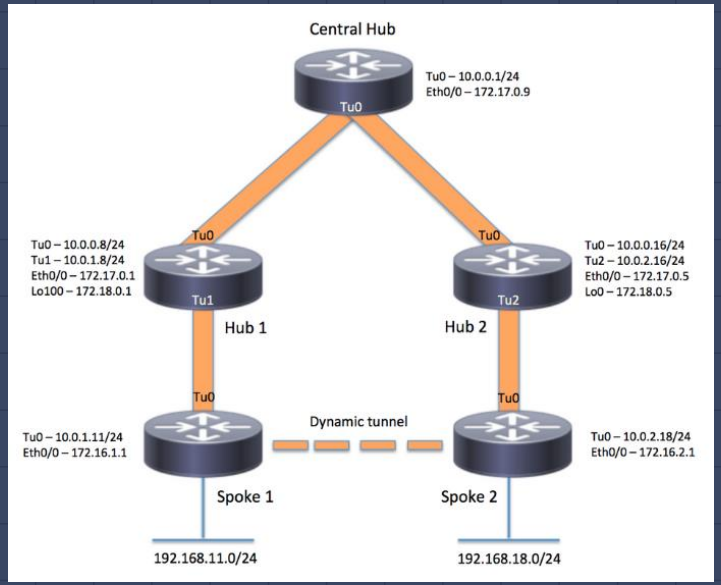

Where Do I Go From Here?

- GRE Resources:
  - Anatomy of GRE Tunnels
  - GRE Tunnel Configuration Guide
- NHRP Resources:
  - NHRP Configuration Guide
- DMVPN Solution Resources:
  - DMVPN White Paper
  - IPSec VPN WAN Design Overview (CVD)

Credits

- Graphics:
- GRE Tunnel Configuration Guide
- DMVPN Configuration Guide
- iWAN Solutions Guide: Cisco DocWiki
- United Airlines
- Wikipedia.org


I KNOW
DMVPN
memegenerator.net