



CHAPTER 5

WAN Edge—DMVPN Per VRF

DMVPN is used widely by enterprises to securely extend their private networks across public networks such as Internet. In a number of scenarios it provides backup to primary Layer 2 WAN connection. One of the ways that the existing DMVPN setup can be leveraged and expanded is by using it to extend virtualization to the branches. All the DMVPN functionality remains intact including bulk encryption and dynamic tunnel building.

Instead of the tunnel residing in the global space, it resides within the VRF. Thus for every VRF, you have to create a separate DMVPN cloud. DMVPN per VRF can create challenges, especially in terms of scale, management, and troubleshooting. So the overall recommendation is to implement this model only if the expectation is that total number of VRFs will remain three or less.

Since this is not new functionality, we focus on the implementation aspects of the solution, such as basic configuration, Multicast, QoS, and redundancy.

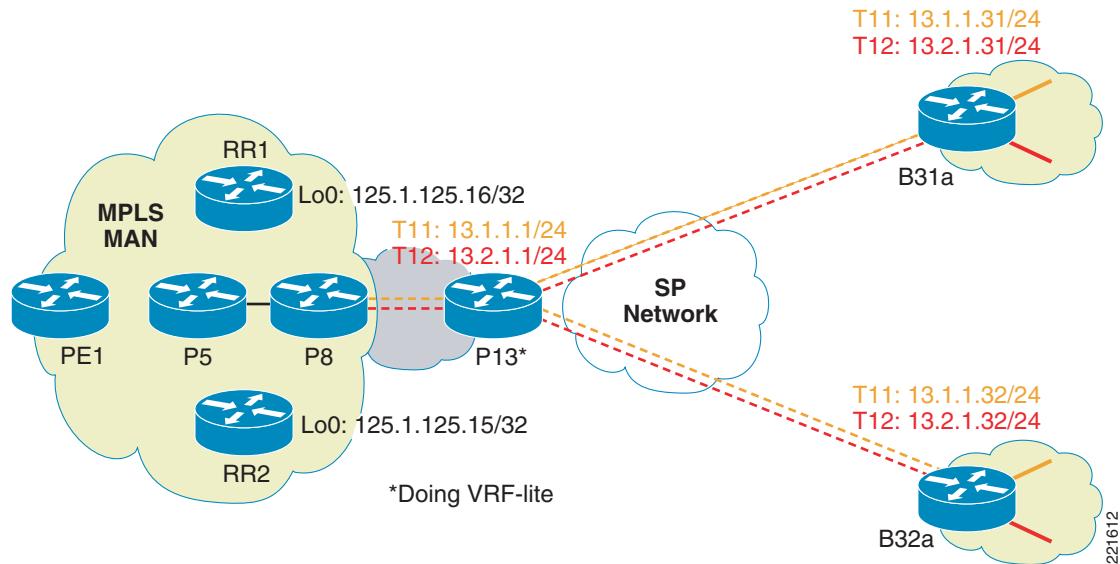
Platforms

The hub can be a 7200VXR (NPE-G1/G2) with encryption modules (VAM2/VAM2+/VSA) or a 7600 (Sup720-3BXL recommended) with encryption modules (VPNSM/VPN SPA). ISRs with hardware encryption accelerators (AIM II) are recommended as spoke routers. The lab tests were done with the following images:

- 7200VXR with NPE-G1/G2—12.4(11)T1
- 7600 with Sup720-3BXL—12.2(18)SXF
- ISRs (3825/2851)—12.4(11)T1

Example:

We discuss the basic implementation with an example. As shown in [Figure 5-1](#), we have two branches 31 (B31a) and 32 (B32a) connecting to hub PE13. The branch routers and the hub are running VRF-lite in this example. The hub is connected to a PE in the MPLS network. It can be a PE too, an example of which we will see later. The hub is not doing encryption in this example.

Figure 5-1 DMVPN per VRF Deployment**Hub PE13:**

```

ip vrf red-data
rd 10:103
!
ip vrf red-voice
rd 10:104
!
interface Tunnel11
ip vrf forwarding red-data
ip address 13.1.1.1 255.255.255.0
ip nhrp authentication spe
ip nhrp map multicast dynamic
ip nhrp network-id 11
ip ospf network broadcast
ip ospf priority 100
tunnel source GigabitEthernet0/2
tunnel mode gre multipoint
tunnel key 1111
!
interface Tunnel12
ip vrf forwarding red-voice
ip address 13.2.1.1 255.255.255.0
ip nhrp authentication spe
ip nhrp map multicast dynamic
ip nhrp network-id 12
ip ospf network broadcast
ip ospf priority 100
tunnel source GigabitEthernet0/2
tunnel mode gre multipoint
tunnel key 2222
!
interface GigabitEthernet0/2
ip address 135.0.16.2 255.255.255.252
!
interface GigabitEthernet0/3.1
encapsulation dot1Q 301
ip vrf forwarding red-data
ip address 125.1.108.2 255.255.255.252
!
```

```

interface GigabitEthernet0/3.2
  encapsulation dot1Q 302
  ip vrf forwarding red-voice
  ip address 125.1.108.2 255.255.255.252
!
router ospf 1 vrf red-data
  log-adjacency-changes
  capability vrf-lite
  network 13.1.1.0 0.0.0.255 area 0
  network 125.1.108.0 0.0.0.3 area 0
!
router ospf 2 vrf red-voice
  router-id 125.1.125.31
  log-adjacency-changes
  capability vrf-lite
  network 13.2.1.0 0.0.0.255 area 0
  network 125.1.108.0 0.0.0.3 area 0
!
router bgp 1
  bgp log-neighbor-changes
  neighbor 135.0.16.1 remote-as 2
  !
  address-family ipv4
    neighbor 135.0.16.1 activate
    neighbor 135.0.16.1 allowas-in
    no auto-summary
    no synchronization
    exit-address-family

```

Spoke B31a:

```

ip vrf red-data
  rd 10:103
!
ip vrf red-voice
  rd 10:104
!
interface Tunnel11
  ip vrf forwarding red-data
  ip address 13.1.1.31 255.255.255.0
  ip nhrp authentication spe
  ip nhrp network-id 11
  ip nhrp nhs 13.1.1.1
  ip ospf network broadcast
  ip ospf priority 0
  tunnel source FastEthernet1/1
  tunnel destination 135.0.16.2
  tunnel key 1111
!
interface Tunnel12
  ip vrf forwarding red-voice
  ip address 13.2.1.31 255.255.255.0
  ip nhrp authentication spe
  ip nhrp network-id 12
  ip nhrp nhs 13.2.1.1
  ip ospf network broadcast
  ip ospf priority 0
  tunnel source FastEthernet1/1
  tunnel destination 135.0.16.2
  tunnel key 2222
!
interface GigabitEthernet0/1.1
  encapsulation dot1Q 241

```

■ Building Redundancy

```

ip vrf forwarding red-data
ip address 125.1.18.1 255.255.255.0
!
interface GigabitEthernet0/1.2
encapsulation dot1Q 242
ip vrf forwarding red-voice
ip address 125.1.18.1 255.255.255.0
!
interface FastEthernet1/1
ip address 135.0.6.2 255.255.255.252
!
router ospf 1 vrf red-data
log-adjacency-changes
capability vrf-lite
passive-interface GigabitEthernet0/1.1
network 13.1.1.0 0.0.0.255 area 0
network 125.1.18.0 0.0.0.255 area 0
!
router ospf 2 vrf red-voice
log-adjacency-changes
capability vrf-lite
passive-interface GigabitEthernet0/1.2
network 13.2.1.0 0.0.0.255 area 0
network 125.1.18.0 0.0.0.255 area 0
!
router bgp 1
bgp log-neighbor-changes
neighbor 135.0.6.1 remote-as 2
!
address-family ipv4
neighbor 135.0.6.1 activate
neighbor 135.0.6.1 allowas-in
no auto-summary
no synchronization
exit-address-family

```

Configuration Notes:

- Every multipoint tunnel corresponds to a VRF. Each tunnel is placed in its own VRF.
- We are running OSPF within each VRF configured with “capability vrf-lite”.
- On the hub, the tunnel interfaces and the VLAN to the core MPLS PE are part of the corresponding OSPF process. On the spokes, the tunnel interface and optionally the LAN-facing VLAN (if there are other OSPF speakers on the LAN) are part of the OSPF process.

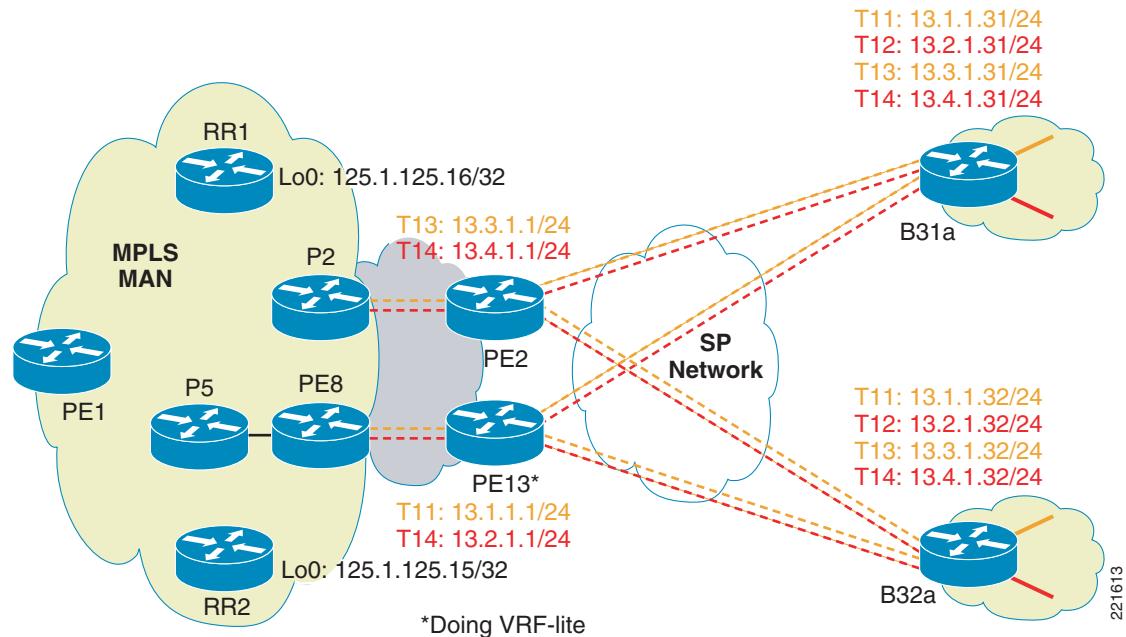
Building Redundancy

As in a normal DMVPN network, it is recommended to have multiple hubs. From the spoke perspective, it keeps connection to both the hubs but prefers one over the other. This can be done by changing the tunnel metric depending on the IGP—change delay for EIGRP and interface cost for OSPF. The return traffic from the headend network just picks the best path. The advantage of keeping such a arrangement is that it allows the hubs to be engineered to maintain a certain number of tunnels and level of traffic. With fast convergence mechanisms configured a tunnel failure would quickly switch the traffic to the backup tunnel.

Example:

As shown in [Figure 5-2](#), we introduce a second hub to our earlier setup. PE2 supports the same VRFs but is also a PE in the core MPLS network (connected to a P). We show the configuration from a 7600 hub PE with encryption enabled.

Figure 5-2 DMVPN per VRF Redundancy



Hub PE2:

```

ip vrf red-data
  rd 10:1032
  route-target export 10:103
  route-target import 10:103
!
ip vrf red-voice
  rd 10:1042
  route-target export 10:104
  route-target import 10:104
!
mls mpls tunnel-recir
!
crypto isakmp policy 1
  encr 3des
  authentication pre-share
  crypto isakmp key Cisco12345 address 0.0.0.0 0.0.0.0
  crypto isakmp keepalive 10 5
!
crypto ipsec transform-set T1 esp-3des esp-sha-hmac
  mode transport
!
crypto ipsec profile P1
  set transform-set T1
!
crypto engine mode vrf
!
interface Loopback0
  ip address 125.1.125.6 255.255.255.255
!
interface Loopback1

```

Building Redundancy

```

ip address 135.10.1.1 255.255.255.255
crypto engine slot 6
!
interface Loopback2
ip address 135.10.1.2 255.255.255.255
crypto engine slot 6
!
interface Tunnel13
ip vrf forwarding red-data
ip address 13.3.1.1 255.255.255.0
no ip redirects
ip nhrp authentication spe
ip nhrp map multicast dynamic
ip nhrp network-id 13
ip ospf network broadcast
ip ospf priority 100
tunnel source Loopback1
tunnel mode gre multipoint
tunnel protection ipsec profile P1
crypto engine slot 6
!
interface Tunnel14
ip vrf forwarding red-voice
ip address 13.4.1.1 255.255.255.0
no ip redirects
ip nhrp authentication spe
ip nhrp map multicast dynamic
ip nhrp network-id 14
ip ospf network broadcast
ip ospf priority 100
tunnel source Loopback2
tunnel mode gre multipoint
tunnel protection ipsec profile P1
crypto engine slot 6
!
interface GigabitEthernet2/9
description To P2
ip address 125.1.100.78 255.255.255.252
tag-switching ip
mls qos trust dscp
!
interface GigabitEthernet2/10
description To SP
ip address 135.0.8.2 255.255.255.252
mls qos trust dscp
crypto engine slot 6
!
interface GigabitEthernet6/1
switchport
switchport trunk encapsulation dot1q
switchport trunk allowed vlan 1,1002-1005
switchport mode trunk
mtu 4500
no ip address
load-interval 30
flowcontrol receive on
flowcontrol send off
spanning-tree portfast trunk
!
interface GigabitEthernet6/2
switchport
switchport trunk encapsulation dot1q
switchport trunk allowed vlan 1,1002-1005
switchport mode trunk

```

```

mtu 4500
no ip address
load-interval 30
flowcontrol receive on
flowcontrol send off
spanning-tree portfast trunk
!
router ospf 2 vrf red-voice
log-adjacency-changes
redistribute bgp 1 subnets
network 13.4.1.0 0.0.0.255 area 0
network 125.1.101.12 0.0.0.3 area 0
!
router ospf 1 vrf red-data
log-adjacency-changes
redistribute connected subnets
redistribute bgp 1 subnets
network 13.3.1.0 0.0.0.255 area 0
network 125.1.101.8 0.0.0.3 area 0
!
router ospf 10
log-adjacency-changes
network 125.0.0.0 0.255.255.255 area 0
maximum-paths 8
!
router bgp 1
no synchronization
bgp log-neighbor-changes
network 135.10.1.1 mask 255.255.255.255
network 135.10.1.2 mask 255.255.255.255
neighbor 125.1.125.15 remote-as 1
neighbor 125.1.125.15 update-source Loopback0
neighbor 125.1.125.16 remote-as 1
neighbor 125.1.125.16 update-source Loopback0
neighbor 135.0.8.1 remote-as 2
no auto-summary
!
address-family vpnv4
neighbor 125.1.125.15 activate
neighbor 125.1.125.15 send-community extended
neighbor 125.1.125.16 activate
neighbor 125.1.125.16 send-community extended
exit-address-family
!
address-family ipv4 vrf red-voice
redistribute ospf 2 vrf red-voice match internal external 1 external 2
maximum-paths ibgp unequal-cost 8
no auto-summary
no synchronization
exit-address-family
!
address-family ipv4 vrf red-data
redistribute ospf 1 vrf red-data match internal external 1 external 2
maximum-paths ibgp unequal-cost 8
no auto-summary
no synchronization
exit-address-family

```

DMVPN per VRF Configuration Notes:

- On a 7600, when using tunnel keys, the packets get process switched. Thus use of tunnel keys should be disabled on the 7600 as well as the connecting branch routers.

- On a 7600, GRE tunnels in different VRFs cannot share the same tunnel source. Hence in our example we are using two different loopback interfaces as sources for the VRF red-data (loopback 1) and red-voice (loopback 2) which are advertised to the SP via BGP.
- On a 7600, the “crypto engine <slot>” commands needs to be configured on the tunnel source interface (loopbacks in this case) and on the tunnels themselves.
- On a 7600, “mls mpls tunnel-recir” needs to be configured when implementing VRF-aware DMVPN.
- On a 7200, “tunnel protection ... shared” must be configured on the tunnel interface if the tunnel is sourced from a single address and tunnel key is used to distinguish different tunnels.
- Since this hub is a full PE, it is configured with route import/export configuration under the VRF. It is also peered with couple of Route Reflectors (125.1.125.15 and 125.1.125.16) in the core MPLS network.
- OSPF is used over the tunnel interface. So a process is created for each VRF and the hub peers with every branch router to which it is connected. The routes are mutually distributed with MP-BGP to establish the connectivity between the core MPLS network and the virtualized branches.
- The spoke configuration is similar to the ones shown in the earlier section, with additional tunnels configured to connect to the 7600. On spokes, for two hub connections, we need two tunnels configured per VRF. Thus, the number of tunnels start multiplying as the number of VRFs increases (four tunnels for two VRFs and two hubs in our example).

Implementing Multicast

Multicast is implemented as a combination of MVPN and non-VRF configuration. The RPs would still reside within the core MPLS network, but need to be reachable by the branches from within the VRF. The branches have multicast capabilities enabled for each VRF. The hub if it is configured as VRF-lite then would only require the appropriate PIM mode to be enabled on the interfaces, RP reachability configured, and multicast enabled for each VRF, similar to the spokes. If the hub is a PE for the core MPLS network, then it will have a full MVPN configuration which includes default and data MDTs for the core network.

Example:

Continuing with our dual hub example, hub PE13 is configured for VRF-lite and hub PE2 is configured for MVPN since its a full fledged PE. Only the additions to earlier configuration are shown here.

Hub PE13:

```
ip multicast-routing
ip multicast-routing vrf red-data
ip multicast-routing vrf red-voice
!
interface Tunnel11
 ip pim nbma-mode
 ip pim sparse-mode
!
interface Tunnel12
 ip pim nbma-mode
 ip pim sparse-mode
!
ip pim vrf red-data rp-address 3.3.3.11
ip pim vrf red-voice rp-address 4.4.4.11
```

Hub PE2:

```

ip vrf red-data
  mdt default 239.232.10.3
  mdt data 239.232.20.32 0.0.0.15 threshold 1
!
ip vrf red-voice
  mdt default 239.232.10.4
  mdt data 239.232.20.48 0.0.0.15 threshold 1
!
ip multicast-routing
  ip multicast-routing vrf red-data
  ip multicast-routing vrf red-voice
!
interface Tunnel13
  ip pim nbma-mode
  ip pim sparse-mode
!
interface Tunnel14
  ip pim nbma-mode
  ip pim sparse-mode
!
ip pim ssm range 1
ip pim vrf red-data rp-address 3.3.3.11
ip pim vrf red-voice rp-address 4.4.4.11
!
access-list 1 permit 239.232.20.0 0.0.0.255

```

The spoke configurations are similar to the VRF-lite case (hub PE13).

Configuration Notes:

- PIM sparse mode is configured on all the interfaces including the core facing (not shown).
- PIM NBMA mode is configured on the multipoint GRE tunnels. This creates the spoke specific entries in the Multicast Output Interface List (OIL).
- In the MPLS network, PIM SSM is used for the data MDTs in the core.



Note Multicast is not supported in DMVPN per VRF on Cat6500 and c7600 routers. It is not recommended to use either of these platforms as a DMVPN mGRE hub if multicast must be implemented.

Implementing QoS

QoS configurations and recommendations do not need to change with DMVPN per VRF. Policies used for existing DMVPN setup are applicable as well. One exception is the DMVPN Hub Support by QoS (http://www.cisco.com/en/US/products/ps6558/prod_bulletin0900aecf803f93d6.html) is not supported).

Example:

The following example shows the configuration on the 7200 hub (PE13) with a sub-rate GE connection to the provider. Hence we will apply a Hierarchical QoS policy to the outgoing interface. It shows a 8-class model with dual LLQ for voice and interactive video traffic.

```

class-map match-all Bulk-Data
  match ip dscp af11  af12
class-map match-any Network-Control
  match ip dscp cs6
  match ip dscp cs2
class-map match-all Critical-Data
  match ip dscp af21  af22

```

■ Scale Considerations

```

class-map match-any Call-Signaling
match ip dscp cs3
match ip dscp af31
class-map match-any Video
match ip dscp af41
match ip dscp af42
class-map match-all Voice
match ip dscp ef
class-map match-all Scavenger
match ip dscp cs1
!
policy-map WAN-EDGE-child
class Voice
  priority percent 18
class Call-Signaling
  bandwidth percent 5
class Network-Control
  bandwidth percent 5
class Critical-Data
  bandwidth percent 27
  random-detect dscp-based
class Bulk-Data
  bandwidth percent 4
  random-detect dscp-based
class Scavenger
  bandwidth percent 1
class Video
  priority percent 15
class class-default
  bandwidth percent 25
  random-detect
!
policy-map WAN-EDGE-parent
class class-default
  shape peak 500000000
  service-policy WAN-EDGE-child
!
interface GigabitEthernet0/2
  ip address 135.0.16.2 255.255.255.252
  service-policy output WAN-EDGE-parent

```

Scale Considerations

Traditional non-VRF DMVPN deployment scales have been limited by the number of IGP peers that can be supported per DMVPN cloud by the headend. For example, tests conducted by ESE/NSITE suggest that a single DMVPN domain can typically support 350-500 EIGRP peers on a 7200 with NPE-G1/VAM2. A platform itself can support two such domains before CPU becomes the limiting factor. Better performance can be expected with NPE-G2. One of the workarounds for that includes splitting the cloud into multiple headends. DMVPN per VRF has the same scale limitations, but now they are multiplied by the number of VRFs.

One approach is to use a similar concept as non-VRF DMVPN deployments and split the DMVPN clouds to multiple headends as the headend reaches the peer scale limits.

Another approach may be to terminate the different VRFs at different hubs. Depending on the requirements, this could mean a subset of VRFs are terminated on a particular hub. While this does not help with the peer scale limits, it can be useful in cases where DMVPN per VRF is used for large number of VRFs but small number of sites per VRF.

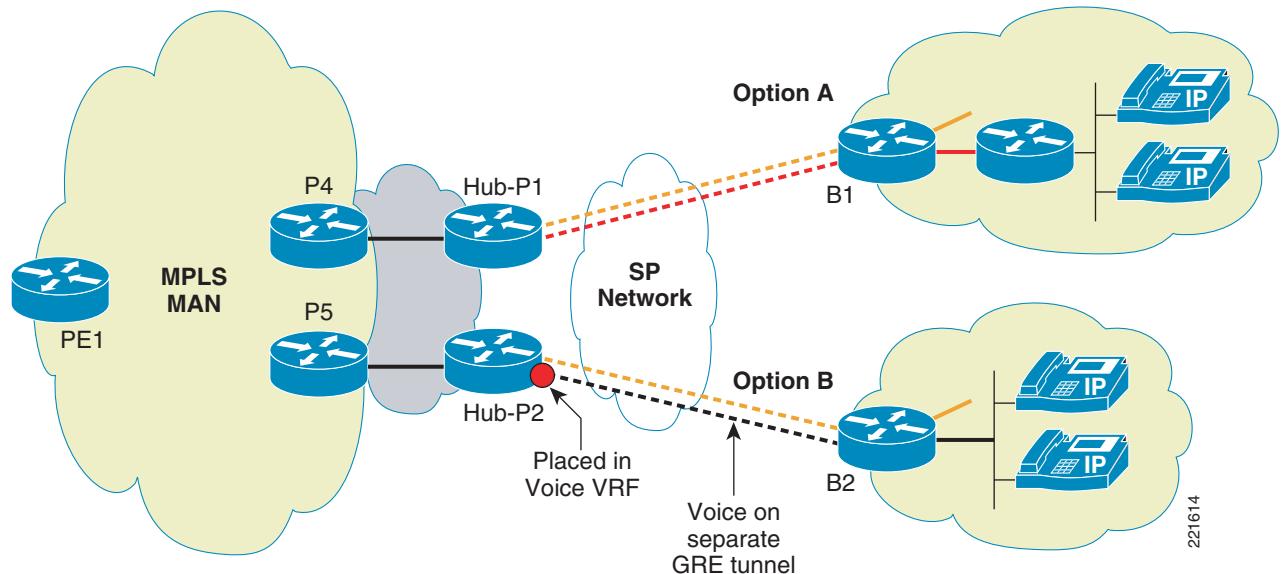
Voice and VRFs

Typically voice traffic has no dependency on the network type since they are just transported as IP packets and require correct QoS behavior applied to them. An exception is when routers are used as gateways for voice services because a lot of voice features and protocols deployed at the branches are not VRF aware (for example, SRST, CME, etc.). Thus just getting the voice traffic in a VRF could be a challenge. This is apart from larger issues of having the voice in a VRF; while you can have the IP phones within a VRF, other services such as softphones VT advantage may be in a different VRF. There are challenges in implementing Inter-VRF IP communications. These are not discussed here as it is part of the larger virtualization architecture issue. The current recommendation is to keep voice within the global space especially at the branches. At the hub they could remain in the global space or would have to be placed within its own VRF. We look at both the options, getting the voice in the VRF at the branch as well keeping it in the global table at the branch.

Voice in a VRF at the Branch

If we need to put the voice in the VRF and still want to use voice features such as CME, then the only way to currently do this is by having two separate routers at the branch. The branch edge router still has a voice VRF configured but treats it like any other VRF. It has a second router (such as a low end ISR) connected to its voice VRF VLAN. The second router, as shown in [Figure 5-3](#) (option A) has all the phones attached to it. Since it requires two routers at every such branch, this can be a expensive proposition

Figure 5-3 DMVPN per VRF—Voice and VRFs



Voice Global at the Branch

If we choose to keep the voice global at the branch then a single router would suffice. The voice VLAN is connected to the branch router but remains in the global space. It is carried across the a global GRE tunnel as normal IPv4 traffic. The DMVPN tunnels per VRF would co-exist with the global tunnel. At

Voice and VRFs

the hub, the DMVPN tunnel carrying voice traffic is placed in the voice VRF from where on it is treated like other VRF traffic ([Figure 5-3](#) option B). Another option at the hub could be to keep the tunnel in the global table for scenarios that keep the voice traffic in the global space even within the hub network.