CHAPTER 7

Migration Strategy and Integrating Non-VRF Sites

Either during migration or otherwise, there will be sites that do not have any VRFs configured since the branches are not virtualized. Thus while migrating from a totally non-virtualized WAN to a virtualized WAN it is recommended to have a parallel headend with VRFs enabled. Thus individual branches can be moved from hub1 (non virtualized) to hub2 as each of them is virtualized.

A similar setup is needed to support non-VRF sites as well. There could be scenarios where the enterprise may choose to implement virtualization in only some of the branches (larger ones) and choose to keep the existing setup in the smaller sites. But at the hub, the network is already virtualized with each segment in its own VRF. In such a case, while there is no separation of routes/traffic at the smaller branches, the routes/traffic need to be separated at the hub and placed in their own VRF.

Figure 7-1 demonstrates one such method. This uses DMVPN as a example but would have relevance in the other cases as well. The DMVPN hub terminates non-virtualized spokes. This would mean that any existing WAN setup can remain un changed. Physically, the only change required is the creation of subinterfaces that connect the hub to the MPLS PE upstream. The PE has the VRFs configured while the hub has everything in the global routing table.

Figure 7-1 Integrating non-VRF Sites

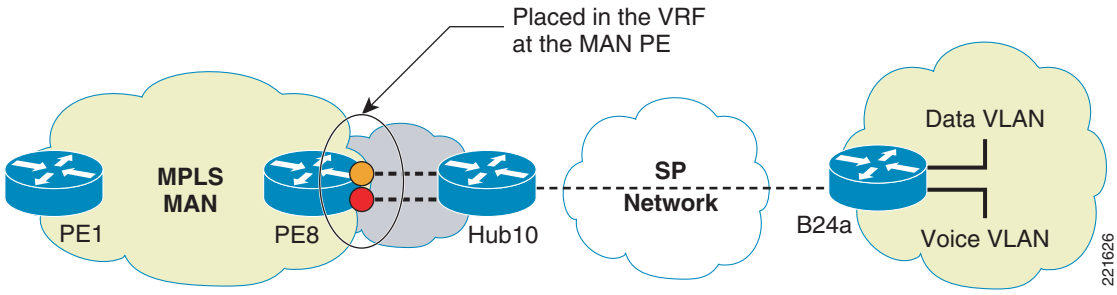

Successful implementation requires that the routes advertisement is controlled between the MPLS PE and the DMVPN hub, so that the traffic flows appropriately as well. Traffic from spoke to hub do not have any issues as the MPLS PE advertises the VRF routes over the matching VLAN to the DMVPN hub, thus creating separate paths for the segmented traffic. The traffic from hub to spoke can get a little tricky. It requires the DMVPN hub to control the route advertisement into the MPLS PE. Spoke routes that need to be visible in VRF A need to be advertised over VLAN A and routes that need to be visible in VRF B need to be advertised over VLAN B. This requires two key abilities in the way the networks have been addressed:

- The spoke addresses need to be easily identifiable such that they can be classified into VRFs at the hub.
- These identifiable addresses need to be summarizable for ease of configuration.

Note It is recommended to use BGP as the routing protocol between MPLS PE and DMVPN hub. It allows for most flexibility in filtering routes.

Note Overlapping addresses cannot exist between the VRFs (unless they are NATed) since the spokes and the DMVPN hub would have no way of differentiating between them since they would be in the global table.

Example:

In the following example the DMVPN hub (hub10) supporting spokes (B24a shown here), connected to MPLS PE (PE8). PE8 has two VRFs (red-data and red-voice) configured on the two VLANs connecting to hub10. The spoke has identifiable subnets that correspond to the two VRFs although they co-exist in the global table till they reach PE8. OSPF is running over DMVPN between the hub and the spoke. BGP is running between PE8 and the hub—1 session per VRF to be supported. Routes are mutually redistributed between OSPF and BGP but filtered in BGP when advertised out from the hub to PE8. B24a and PE8 are configured normally. Relevant portions of hub10 configuration are shown below.

Hub10:

```
interface Tunnel11
 ip address 12.2.1.1 255.255.255.0
 no ip redirects
 ip nhrp authentication spoke
 ip nhrp map multicast dynamic
 ip nhrp network-id 1
 ip ospf network point-to-multipoint
 ip ospf priority 100
 tunnel source POS1/0
 tunnel mode gre multipoint
 tunnel key 111
!
interface GigabitEthernet0/3
 description To PE8
!
interface GigabitEthernet0/3.1
 description VLAN for red-data
 encapsulation dot1Q 201
 ip address 125.1.141.2 255.255.255.252
 no snmp trap link-status
!
interface GigabitEthernet0/3.2
 description VLAN for red-voice
 encapsulation dot1Q 202
 ip address 125.1.141.6 255.255.255.252
 no snmp trap link-status
!
router ospf 10
 log-adjacency-changes
 redistribute bgp 3 subnets
 network 12.2.1.0 0.0.0.255 area 0
!
router bgp 3
 no synchronization
 bgp log-neighbor-changes
 redistribute ospf 10 match internal external 1 external 2
```

```
 neighbor 125.1.141.1 remote-as 1
 neighbor 125.1.141.1 update-source GigabitEthernet0/3.1
 neighbor 125.1.141.1 distribute-list data out
 neighbor 125.1.141.5 remote-as 1
 neighbor 125.1.141.5 update-source GigabitEthernet0/3.2
 neighbor 125.1.141.5 distribute-list voice out
 no auto-summary
!
ip access-list standard data
 permit 125.1.17.0 0.0.0.63
ip access-list standard voice
 permit 125.1.17.64 0.0.0.63
```