

Cisco Virtual Office—Advanced Deployment

Contents

Scope of Document	1
Hardware Platforms and Software Images	2
Cisco Virtual Office: Per-Tunnel QoS Deployment	2
Introduction	2
Per-Tunnel QoS Configuration	3
Per-Tunnel QoS Diagnostic Commands	5
Cisco Virtual Office: BGP Routing Between Hub and Spoke	6
Introduction	6
Dynamic Update Peer-Groups Configuration	6
Cisco Virtual Office: Next-Hop Server Clustering	7
Introduction	7
Configuration	8
Sample Output	9
Resources	10

Scope of Document

This document addresses some of the advanced features configured for the Cisco[®] Virtual Office deployment. Although the default configuration in the Cisco Virtual Office Deployment guide would suffice for configuring a functional VPN network, this guide provides information about extra features and control over the network.

Three features are discussed in detail in this document:

1. Per-Dynamic Multipoint VPN (DMVPN) tunnel quality of service (QoS)
2. Border Gateway Protocol (BGP) routing between hub and spoke (in lieu of Enhanced IGRP [EIGRP])
3. Next-hop server (NHS) clustering

Per-DMVPN tunnel QoS enables more granular control over the network on a spoke-per-spoke basis.

BGP routing between hub and spoke provides an alternative to routing using EIGRP, and potentially increases scalability.

NHS clustering allows more deployment flexibility when multiple VPN headend devices are present.

This document details the recommended settings of each of the features for Cisco Virtual Office. It is assumed that a functional Cisco Virtual Office network is operational before these settings are attempted.

For more information about Cisco Virtual Office, please refer to <http://www.cisco.com/go/cvo>.

Hardware Platforms and Software Images

This guide is based on Cisco 881 Integrated Services Routers as spokes with wireless running Cisco IOS® Software Release 15.1(2)T1. Hub routers are based on Cisco 3945 Integrated Service Routers. For other Cisco router platforms, the sample configurations may need minor modifications. For a list of all supported hardware and software, please refer to the Cisco Virtual Office Hardware and Software list.

Cisco Virtual Office: Per-Tunnel QoS Deployment

Introduction

Per-tunnel QoS allows traffic from a hub to spokes to be regulated on a per-spoke basis. It allows QoS to be configured on tunnel interfaces used for DMVPN and IPsec, whereas previously QoS was restricted to physical interface support.

Per-tunnel QoS solves two problems commonly found in hub-and-spoke topologies:

1. Prevents lower-end spoke routers from being overrun by higher-end hubs
2. Prevents some spokes from hogging hub resources and starving other spokes.

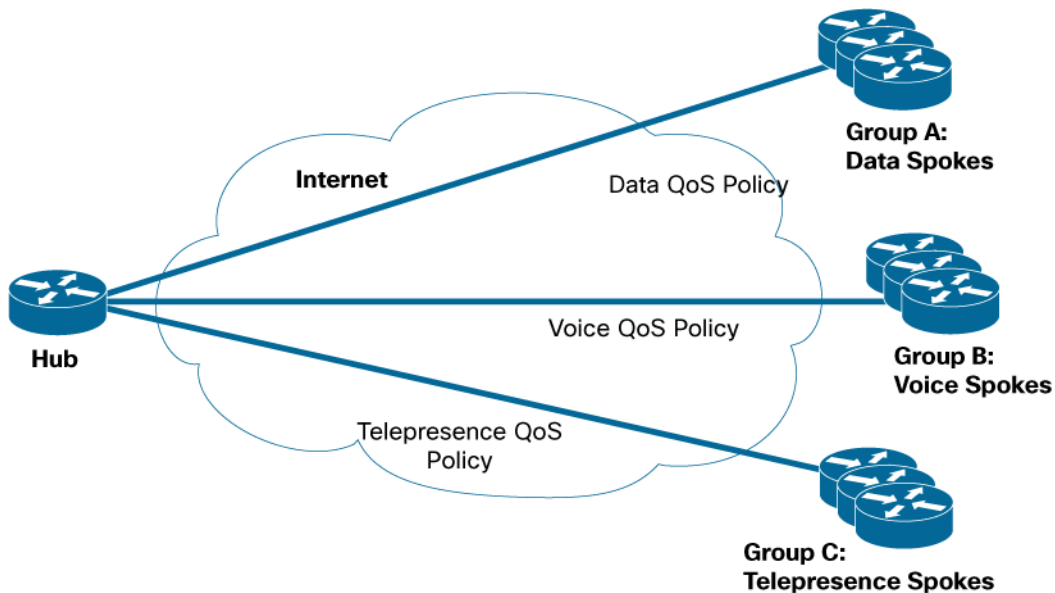
Case 1 occurs in many DMVPN networks, because the hub is usually a high-end router, whereas spokes are often lower-end routers. In this case, if the link between the hub and spoke is unregulated, the hub may send traffic at a higher rate than the spoke can handle, causing congestion and packet loss. In case 2, one spoke may be performing traffic-intensive operations, such as VoIP calls. Because that spoke is requesting a lot of traffic, it may overrun the link between the hub and itself, not allowing traffic to flow from the hub to other spokes.

Previously, QoS could be configured on spokes, but the hub is not aware of the policies and cannot prevent the problem indicated in case 2. Alternatively, past configurations of QoS on the hub cannot adjust for multiple classes of service for the spokes.

Per-tunnel QoS is most advantageous when different levels of service are desired for separate subsets of spokes, or when the spoke routers themselves vary in capability. One of the first design considerations is to divide the spokes into groups, which have different QoS policies.

In the Cisco Virtual Office configuration, there are three different class levels to which spokes belong: data, voice, and telepresence. Figure 1 shows the Cisco Virtual Office hub-and-spoke topology and the spoke groups.

Figure 1. Cisco Virtual Office DMVPN Hub-and-Spoke Topology and Per-Tunnel QoS Setup



In Figure 1, group A comprises spokes with data traffic being the primary traffic flowing from hub to spoke. Users in group B make frequent calls and so require a QoS policy that can guarantee good voice and video quality for the calls, whereas users in group C are executives who use Cisco TelePresence[®] conferencing applications and need guaranteed bandwidth to ensure adequate telepresence quality. With the deployment of per-tunnel QoS, different QoS policies can be applied to each tunnel, depending on which group the spoke belongs to, helping ensure that users in each group receive adequate bandwidth for their applications without overstressing the hub or their spokes.

Per-Tunnel QoS Configuration

The bulk of per-tunnel QoS configuration is on the hub, which uses Cisco Policy Language to set up a hierarchy of class maps and policy maps. On the DMVPN hub interface, these policy mappings are tied together through the use of an **nhrp-group** command. Similarly on the spoke side, the **nhrp-group** command specifying the policy name is configured on the spoke router DMVPN interface. The hub receives the **nhrp-group** string from the spoke to the hub in the periodic NHRP registration requests. The **nhrp-group** string is then mapped to the QoS policy defined on the hub, and the policy is applied to the tunnel from the hub to the spoke.

The following is a sample Cisco Virtual Office configuration for per-tunnel QoS:

```

!!! Hub configuration !!!
! The QoS configuration on the hub side follows a parent-child hierarchy in which a
! child policy is defined and then can be called by the parent policy.
! This configuration creates 3 classes to support data; voice and video;
! and Telepresence QoS

! Define parent class 'data,' which provides default QoS policy for spokes
! requiring
! just data traffic
policy-map data

```

```

class class-default
  shape average 1000000

! Configure the default behavior. Parent classes of 'data,' 'voice,' and
! 'Telepresence' will all call the default class as it is assumed data traffic
will be
! common to all the groups.
class class-default
  fair-queue
  random-detect

! Define parent class 'voice,' which provides voice and video QoS policy
policy-map voice
  class class-default
    shape average 2000000
    service-policy voice_and_video

! Child policy 'voice_and_video' is called by parent policy 'voice' above
policy-map voice_and_video
  class voice
    priority 384
  class class-default
    fair-queue
    random-detect

! Define parent class 'tp,' which provides QoS policy for spokes needing
Telepresence
! support
policy-map tp
  class TelePresence
    priority 10000
  class class-default
    fair-queue
    random-detect

! Child class 'TelePresence' is called by parent policy 'tp' above
class TelePresence
  priority 500000
  class class-default
    fair-queue
    random-detect

! Policies are attached to the DMVPN tunnel interface using the 'nhrp map group'
! command
interface Tunnel300
  ip nhrp map group persa_data service-policy output data
  ip nhrp map group persa_voice service-policy output voice

```

```

ip nhrp map group persa_tp service-policy output tp

!!! Spoke configuration !!!

! Spokes requiring data-only QoS policy should configure the following on their
DMVPN
! tunnel interface
interface Tunnel300
  ip nhrp group data

! Spokes requiring voice and video QoS policies should configure the following on
! their DMVPN tunnel interface
interface Tunnel300
  ip nhrp group voice

! Spokes requiring Telepresence QoS policy should configure the following on
their
! DMVPN tunnel interface
interface Tunnel300
  ip nhrp group tp

```

Per-Tunnel QoS Diagnostic Commands

Table 1 lists the per-tunnel QoS diagnostic commands.

Table 1. Per-Tunnel QoS Diagnostic Commands

Hub-Side Diagnostics	
show ip nhrp	Displays NHRP group received from spoke
show dmvpn detailed	Displays NHRP group received from spoke and QoS policy applied to the spoke tunnel
show ip nhrp group-map	Displays group-to-policy maps and tunnels to which QoS policy is applied
show policy-map multipoint tunnel <ip addr> output	Displays QoS policy applied to multipoint tunnels
show tunnel endpoints	Displays QoS policy applied
debug nhrp group	Debugs nhrp groups received from spokes, group-to-policy map existence, and policy status
debug tunnel qos	Debugs QoS policy application and removal
debug nhrp extension	Debugs receipt of nhrp group from spoke
debug cce dp target	Debugs data plane Common Classification Engine (CCE) packet classification and QoS policy for packet
Spoke-Side Diagnostics	
debug nhrp extension	Debugs receipt of nhrp group from hub

Cisco Virtual Office: BGP Routing Between Hub and Spoke

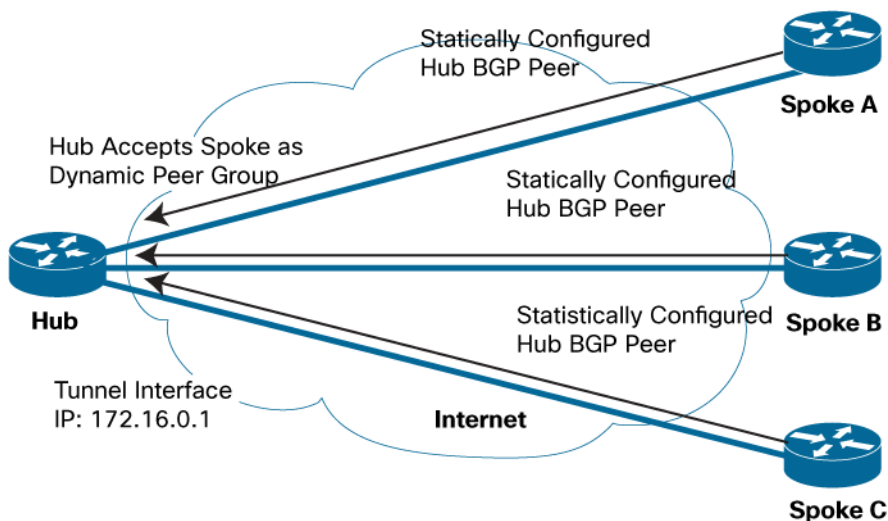
Introduction

In most Cisco Virtual Office deployments, EIGRP is used as the underlining routing protocol between the hub and the spoke. Although EIGRP offers simple configurations and fast convergence times, there are cases when using a Cisco proprietary routing protocol is undesirable. In those cases, BGP could be used as an alternative.

Using the Dynamic Update Peer-Groups feature in the Cisco IOS® Software, configuring BGP could become as simple as configuring EIGRP. Besides, the number of spokes supported per hub using BGP as the routing protocol might increase as well.

Dynamic Update Peer-Groups Configuration

Figure 2. Cisco Virtual Office DMVPN Hub-and-Spoke Topology and BGP Dynamic Peer Groups.



The Cisco Virtual Office spoke will have the Cisco Virtual Office hub configured as a static peer, with a hardcoded IP address. On the Cisco Virtual Office hub, however, the spokes are accepted as dynamic peers under a dynamic peer group. Spokes can dynamically request peering with the hub when they come online, and it is not necessary to configure each spoke IP address on the hub.

The configuration on the spoke follows (arbitrarily assigning 1628 as the AS number on the spoke and 65159 as the AS number on the hub):

```
router bgp 1628
  bgp log-neighbor-changes
  neighbor 172.16.0.1 remote-as 65159
  ! 172.16.0.1 is the DMVPN tunnel address on the hub
  !
  address-family ipv4
    network 10.32.242.88 mask 255.255.255.248
    !10.32.242.88 is the spoke subnet
    neighbor 172.16.0.1 activate
  exit-address-family
```

!

The configuration on the hub follows:

```
router bgp 65159
  bgp listen range 172.16.0.0/16 peer-group dynamic
  neighbor dynamic peer-group
  neighbor dynamic remote-as 1628
!
address-family ipv4
  aggregate-address 10.32.224.0 255.255.224.0 summary-only
  neighbor dynamic activate
  no auto-summary
exit-address-family
```

In the configuration, a peer group called “dynamic” is created in the BGP on the hub. The hub will listen to incoming peering requests from remote AS 1628, originating from source IP in the range of 172.16.0.0/16 (specified in the listen range). The group “dynamic” must be activated using the

```
neighbor dynamic activate
```

command under address-family ipv4.

Regular route redistribution can be configured normally on the BGP 65159 instance on the hub.

Note: Dynamic Update peer group does not presently support IPv6.

Cisco Virtual Office: Next-Hop Server Clustering

Introduction

In situations when more than one DMVPN hub router is deployed, it is possible to logically group multiple hub routers into a cluster in configuration. Connections are made to a cluster instead of to an individual router. The number of connections made to each cluster can be altered, so various combination of connections to hub routers are possible.

Clustering is particularly useful in the following two scenarios:

1. Configure an inactive backup hub router: A hub router can be part of the cluster but not activated, and connection is made only when the primary connection to other hub router(s) fails.
2. Multiple hub routers: If there is a farm of multiple hub routers that the spoke can connect to, hub routers with different priorities can be assigned to a single cluster, and the spoke will attempt connection according to the priority settings dynamically. A limit can be set on each hub router for the number of connections accepted at a given time; excess connections will not establish, and the spoke router will try the router with next priority in the cluster.

Figure 3. Scenario 1: Configure an Inactive Backup Hub Router

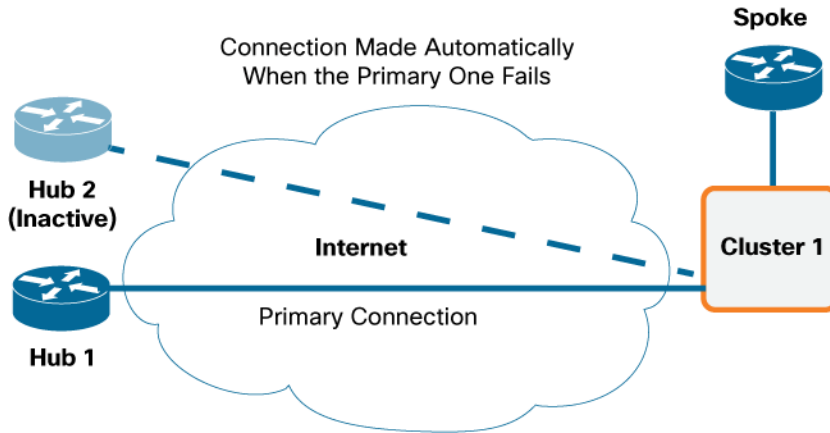
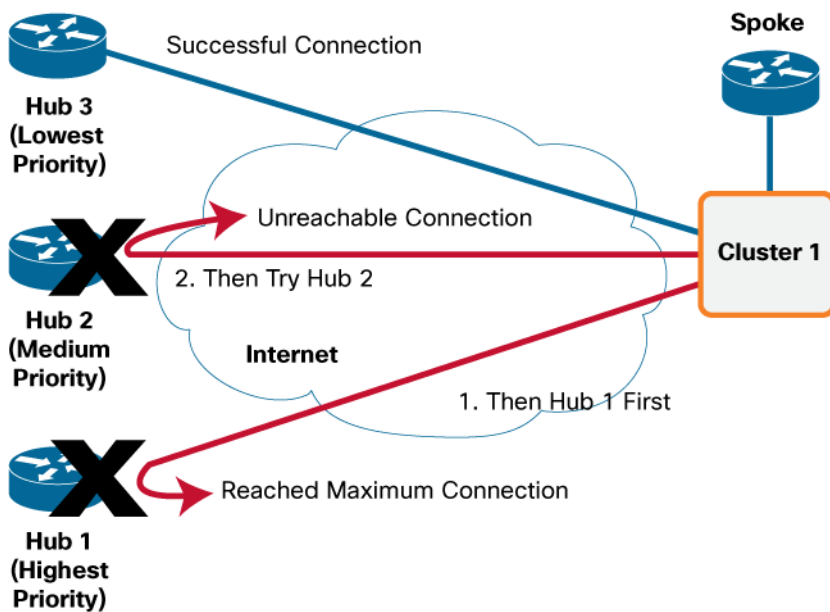


Figure 4. Scenario 2: Multiple Hub Routers



Configuration

Clustering is configured on the remote spoke router only. The hub router has no knowledge of the cluster configuration on the spoke. The hub involved in a cluster can be either geographically co-located, or dispersed. More than one cluster can be configured on the spoke.

The configuration is put on the DMVPN tunnel interface on the remote spoke router:

```
interface tunnel 0
 ip nhrp nhs cluster $cluster_number max-connections $max_connection
 ip nhrp nhs $nhs_address priority $nhs_priority cluster $cluster_number
 ip nhrp nhs $nhs_address2 priority $nhs_priority2 cluster $cluster_number
```

\$cluster_number is a locally significant number assigned to the cluster for identification.

\$max_connection is the maximum number of connections to establish to hubs in the cluster.

\$nhs_address/ \$nhs_address2 are the outside addresses of the DMVPN hub routers.

\$nhs_priority/ \$nhs_priority2 are the priorities assigned to each hub; the **lower** the number the **higher** the priority.

A sample of an actual configuration follows:

```
interface tunnel 0
 ip nhrp nhs cluster 1 max-connections 1
 ip nhrp nhs 1.1.1.1 priority 10 cluster 1
 ip nhrp nhs 2.2.2.2 priority 20 cluster 1
```

In the configuration, only one DMVPN connection will be established (max-connection 1), and the spoke will try 1.1.1.1 first (lower-priority number) before trying to connect to 2.2.2.2.

Note: The spoke router will periodically try to establish a tunnel to 1.1.1.1, and it will try indefinitely.

To configure the maximum connections allowed to a hub router (e.g., to accomplish scenario 2), this command is used on the hub router:

```
call admission limit $num_connection
```

\$num_connection is the maximum number of connections allowed to the hub router, and it should be dependent on the capability of the hub router. For scenario 2 deployment, you should randomize the priority order between each of the spoke routers so as not to overwhelm a single hub router.

Sample Output

The output from a spoke router with two hub routers in the cluster, a maximum connection of one, and the highest-priority hub failing should look like this:

```
spoke#show ip nhrp nhs redundancy
```

Legend: E=Expecting replies, R=Responding, W=Waiting

No.	Interface	Cluster	NHS	Priority	Cur-State	Cur-Queue	Prev-State
1	Tunnel0	1	2.2.2.2	10	RE	Running	E
2	Tunnel0	1	1.1.1.1	9	E	Running	E

No.	Interface	Cluster	Status	Max-Con	Total-NHS	Registering/UP	Expecting
1	Tunnel0	1	Enable	1	2	1	1
0	0						

A hub set with call admission to restrict the number of incoming connections to three is configured as follows:

```
hub#show call admission statistics
Total call admission charges: 48, limit 3
Total calls rejected 5971, accepted 0
Load metric: charge 48, unscaled 48%
```

Note: Clusters can be created only within a tunnel interface, with a single source interface. It is not possible to fail to a redundant WAN interface using NHS cluster.

Resources

Dynamic Update Peer Groups:

http://www.cisco.com/en/US/docs/ios/iproute_bgp/configuration/guide/irg_neighbor_ps10592_TSD_Products_Configuration_Guide_Chapter.html#wp1055389

Next-Hop Server (NHS) Clustering:

http://www.cisco.com/en/US/docs/ios/sec_secure_connectivity/configuration/guide/sec_dmvpn_backup_nhs.html#wp1059394



Americas Headquarters
Cisco Systems, Inc.
San Jose, CA

Asia Pacific Headquarters
Cisco Systems (USA) Pte. Ltd.
Singapore

Europe Headquarters
Cisco Systems International BV Amsterdam,
The Netherlands

Cisco has more than 200 offices worldwide. Addresses, phone numbers, and fax numbers are listed on the Cisco Website at www.cisco.com/go/offices.

Cisco and the Cisco Logo are trademarks of Cisco Systems, Inc. and/or its affiliates in the U.S. and other countries. A listing of Cisco's trademarks can be found at www.cisco.com/go/trademarks. Third party trademarks mentioned are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (1005R)