



GRE and DM VPNS

You can configure Generic Routing Encapsulation (GRE) and Dynamic Multipoint (DM) VPNs that include GRE mode configurations. You can configure IPsec GRE VPNs for hub-and-spoke, point-to-point, and full mesh VPN topologies. DMVPN is available for hub-and-spoke topologies only.

This chapter contains the following topics:

- [Understanding the GRE Modes Page](#) , on page 1
- [GRE and Dynamic GRE VPNs](#) , on page 2
- [Dynamic Multipoint VPNs \(DMVPN\)](#) , on page 10

Understanding the GRE Modes Page

Use the GRE Modes page to define the routing and tunnel parameters for IPsec tunneling with GRE, GRE Dynamic IP, and DMVPN policies.

The content of the policy differs depending on how you access it:

- ([Site-to-Site VPN Manager Window](#)) When you select a GRE VPN or DMVPN, the GRE Modes policy contains the properties related to the technology and technology type used in the VPN.
- (Policy view) When you select **Site-to-Site VPN > GRE Modes**, and create a new policy or select an existing policy, there is an additional field in the policy called **GRE Method**. From the GRE Method list, you must select the VPN technology and technology type for which you are defining the policy: IPsec/GRE, GRE Dynamic IP, DMVPN, or Large Scale DMVPN. This option controls which fields are displayed in the policy. You cannot change the GRE Method after you save the policy.

When you assign a shared GRE Modes policy to a VPN, the GRE Method and the VPN's technology and type must match or the policy cannot be selected. For example, you cannot assign a shared DMVPN GRE Modes policy to an IPsec/GRE VPN.

The following topics describe the GRE Modes policy in detail based on the selected GRE Methods:

- IPsec/GRE or GRE Dynamic IP—See [Configuring GRE Modes for GRE or GRE Dynamic IP VPNs](#) , on page 6.
- DMVPN or Large Scale DMVPN—See [Configuring GRE Modes for DMVPN](#) , on page 13.



Note When configuring an IPsec/GRE, GRE Dynamic IP, or DMVPN routing policy, Security Manager adds a routing protocol to all the devices in the secured IGP, on deployment. If you want to maintain this secured IGP, you must create a router platform policy (on each member device) using the same routing protocol and autonomous system (or process ID) number as defined in the GRE Modes policy.

Related Topics

- [Understanding GRE , on page 2](#)
- [Understanding GRE Configuration for Dynamically Addressed Spokes , on page 5](#)
- [Understanding DMVPN , on page 10](#)
- [Understanding IPsec Technologies and Policies](#)

GRE and Dynamic GRE VPNs

You can use Generic Routing Encapsulation (GRE) to create VPNs using Cisco IOS security routers and Catalyst 6500/7600 devices in hub-and-spoke, point-to-point, and full mesh VPN topologies.

This section contains the following topics:

- [Understanding GRE , on page 2](#)
- [Configuring IPsec GRE VPNs , on page 5](#)
- [Configuring GRE Modes for GRE or GRE Dynamic IP VPNs , on page 6](#)

Understanding GRE

Generic Routing Encapsulation (GRE) is a tunneling protocol that encapsulates a variety of protocol packet types inside IP tunnels, creating a virtual point-to-point connection to devices at remote points over an IP network. With this technology, GRE encapsulates the entire original packet with a standard IP header and GRE header before the IPsec process. Then, IPsec views the GRE packet as an unremarkable IP packet and performs encryption and authentication services, as dictated by the IKE negotiated parameters. Because GRE can carry multicast and broadcast traffic, it is possible to configure a routing protocol for virtual GRE tunnels. The routing protocol detects loss of connectivity and reroutes packets to the backup GRE tunnel, thus providing high resiliency.

For VPN resilience, a spoke must be configured with two GRE tunnels, one to the primary hub and the other to the backup hub. Both GRE tunnels are secured with IPsec: each one has its own IKE security association (SA) and a pair of IPsec SAs. An associated routing protocol automates the failover mechanism, transferring to the backup tunnel if virtual link loss is detected.



Note GRE can be configured on Cisco IOS security routers and Catalyst 6500/7600 devices in hub-and-spoke, point-to-point, and full mesh VPN topologies.

This section contains the following topics:

- [Advantages of IPsec Tunneling with GRE](#) , on page 3
- [How Does Security Manager Implement GRE?](#) , on page 3
- [Prerequisites for Successful Configuration of GRE](#) , on page 3
- [Understanding GRE Configuration for Dynamically Addressed Spokes](#) , on page 5

Advantages of IPsec Tunneling with GRE

The main advantages of IPsec tunneling with GRE are the following:

- GRE uses a routing protocol by which every IPsec peer knows the status of every other peer at all times.
- GRE provides higher resiliency than IKE keepalive.
- Spoke-to-spoke connectivity is supported when you use GRE.
- GRE supports multicast and broadcast transmissions.



Note GRE does not support the use of dynamic cryptographic tunnels.

How Does Security Manager Implement GRE?

Security Manager implements an additional Interior Gateway Protocol (IGP) solution for GRE. An IGP refers to a group of devices that receive routing updates from one another by a routing protocol, EIGRP, OSPF, or RIP. Each “routing group” is identified by a logical number. For general routing purposes, the interfaces on the routers in your networks belong to an IGP. Security Manager adds an additional IGP that is dedicated for IPsec and GRE-secured communication. This additional IGP is the secured IGP. The existing IGP (unsecured IGP), is used for routing traffic that does not require encryption.

For a GRE tunnel to be established, Security Manager configures a virtual interface on each device. These virtual interfaces are the endpoints of the GRE tunnel. Each virtual interface is unique. The GRE tunnel interface has an IP address (inside tunnel IP address) which is taken from an interface that Security Manager creates. The GRE tunnel points to the source and destination IP addresses of either the physical or loopback interfaces on each device. The GRE virtual interfaces belong to the secured IGP, as do the inside interfaces. Routing updates within the secured IGP are GRE encapsulated and IPsec is applied. A flow whose destination is a secured interface (according to the routing updates of the secured IGP) is directed through the GRE interface where it is GRE encapsulated and then evaluated against the crypto ACL. If it matches the crypto ACL, it is routed through the GRE and VPN tunnels.

Prerequisites for Successful Configuration of GRE

Consider the following prerequisites before using GRE in your network:

- You must identify the inside interfaces on your devices—the physical interfaces on the device that connect the device with its internal subnets and networks.
- You must select a routing protocol (known as an IGP) or a static route, whenever you enable GRE.

Security Manager supports the EIGRP, OSPF, and RIPv2 dynamic routing protocols, and GRE static routes.

- EIGRP—Enhanced Interior Gateway Routing Protocol enables the exchange of routing information within an autonomous system and addresses some of the more difficult issues associated with routing in large, heterogeneous networks. Compared to other protocols, EIGRP provides superior convergence properties and operating efficiency, and combines the advantages of several different protocols. For more information, see [EIGRP Routing on Cisco IOS Routers](#).
- OSPF—Open Shortest Path First is a link-state, hierarchical protocol that features least-cost routing, multipath routing, and load balancing.

Using OSPF, a host that obtains a change to a routing table or detects a change in the network immediately multicasts the information to all other hosts in the network, so that all will have the same routing table information. For more information, see [OSPF Routing on Cisco IOS Routers](#).

- RIPv2—Routing Information Protocol is a distance-vector protocol that sends routing-update messages at regular intervals and whenever the network topology changes.

Using RIPv2, a gateway host (with a router) sends its entire routing table to its closest neighbor host every 30 seconds, which in turn passes the information on to its next neighbor, and so on, until all hosts within the network have the same knowledge of routing paths. RIPv2 uses a hop count to determine network distance. Each host with a router in the network uses the routing table information to determine the next host to route a packet to for a specified destination.

RIP is considered an effective solution for small homogeneous networks. For larger, more complicated networks, RIP's transmission of the entire routing table every 30 seconds may put a heavy amount of extra traffic in the network. For more information, see [RIP Routing on Cisco IOS Routers](#).

- Static route—Use a static routing policy to provide a robust, stable IPsec-protected GRE tunnel if there is a fixed, unchanging route between two devices. For each device subnet, a static route is created on the device pointing to the corresponding tunnel interface. For more information, see [Static Routing on Cisco IOS Routers](#).
- You must specify an IGP process number. The IGP process number identifies the IGP process to which the inside interface on the device belongs. When GRE is implemented, this will be the secured IGP. For secure communication, the inside interfaces on the devices in your VPN must use the same IGP process. The IGP process number must be within a specified range. If you have an existing IGP process on the device that is within this range, but is different from the IGP process number specified in your GRE settings, Security Manager removes the existing IGP process. If the existing IGP process matches the one specified in your GRE settings, any networks included in the existing IGP process that do not match the specified inside interfaces are removed.
- If the inside interfaces on your devices are configured to use an IGP process other than the IGP process specified in your GRE settings (meaning that the interfaces belong to an unsecured IGP):
 - For spokes: Manually remove the inside interfaces from the unsecured IGP through the device CLI before configuring GRE.
 - For hubs: If the hub inside interface is used as a network access point for Security Manager, then on deployment, the interface is advertised in both secured and unsecured IGPs. To ensure that the spoke peers use only the secured IGP, manually add the auto-summary command for the unsecured IGP or remove the unsecured IGP for that inside interface.

- You must provide a subnet that is unique yet it can be non-globally-routable for loopback. This subnet must only be used to support the implementation of loopback for GRE. The loopback interfaces are created, maintained, and used only by Security Manager. You should not use them for any other purpose.
- If you are using static routes, not unsecured IGP, make sure you configure static routes on the spokes through to the hub inside interfaces.



Note You can configure the above settings in the GRE Modes page when IPsec/GRE is the selected IPsec technology.

Understanding GRE Configuration for Dynamically Addressed Spokes

When a spoke has a dynamic IP address, there is no fixed GRE tunnel source address (to be used by the GRE tunnel on the spoke side) or destination address (to be used by the GRE tunnel on the hub side). Therefore, Security Manager creates additional loopback interfaces on the hub and the spoke, to be used as the GRE tunnel endpoints. You must specify a subnet from which Security Manager can allocate an IP address for the loopback interfaces.



Note GRE Dynamic IP can only be configured on Cisco IOS routers and Catalyst 6500/7600 devices in hub-and-spoke VPN topologies.

Security Manager uses the Cisco Configuration Engine to retrieve device IP addresses and other information from dynamically addressed devices. Devices that have dynamic IP addresses connect to the Configuration Engine manager at periodic intervals to upgrade device configuration files and to pass device and status information.

For more information, see [Adding, Editing, or Deleting Auto Update Servers or Configuration Engines](#).



Note You can configure the GRE Dynamic IP settings in the GRE Modes page when GRE Dynamic IP is the selected IPsec technology.

Related Topics

- [Understanding GRE](#) , on page 2
- [Configuring GRE Modes for DMVPN](#) , on page 13

Configuring IPsec GRE VPNs

To configure an IPsec GRE (generic routing encapsulation) VPN, use the Create VPN wizard as described in [Creating or Editing Extranet VPNs](#). You can also edit the membership of the VPN, or some of its policies, using the described procedures. If you are creating a hub-and-spoke VPN with dynamically addressed spokes, also see [Understanding GRE Configuration for Dynamically Addressed Spokes](#) , on page 5.

If you need to make changes to other policies and settings, open the policies from the Site-to-Site Manager page, as follows:

- For ISAKMP and IPsec settings, select **VPN Global Settings**. See [Configuring VPN Global Settings](#).
- For IKE proposal policies, select **IKE Proposal**. See [Configuring an IKE Proposal](#).
- For IPsec proposals, select **IPsec Proposal**. See [Configuring IPsec Proposals in Site-to-Site VPNs](#).
- For preshared key policies, select **IKEv1 Preshared Key**. See [Configuring IKEv1 Preshared Key Policies](#).
- For public key (PKI) policies, select **Public Key Infrastructure**. See [Configuring IKEv1 Public Key Infrastructure Policies in Site-to-Site VPNs](#).
- For Generic Routing Encapsulation configuration, select **GRE Modes**. See [Configuring GRE Modes for DMVPN , on page 13](#).

Related Topics

- [Understanding IKE](#)
- [Understanding GRE , on page 2](#)
- [Prerequisites for Successful Configuration of GRE , on page 3](#)
- [Advantages of IPsec Tunneling with GRE , on page 3](#)

Configuring GRE Modes for GRE or GRE Dynamic IP VPNs

Use the GRE Modes policy to define the routing and tunnel parameters for IPsec tunneling in a GRE or GRE Dynamic IP VPN.

To open the GRE Modes policy:

- ([Site-to-Site VPN Manager Window](#)) Select an IPsec/GRE or GRE Dynamic IP topology, then select **GRE Modes** from the policies list.
- (Policy view) Select **Site-to-Site VPN > GRE Modes**, and create a new policy or select an existing policy. Then, select either IPsec/GRE or Dynamic GRE from the **GRE Method** list.

The following table describes the elements on the GRE Modes page for configuring IPsec tunneling with GRE or GRE Dynamic IP.



Note When configuring a GRE routing policy, Security Manager adds a routing protocol to all the devices in the secured IGP, on deployment. If you want to maintain this secured IGP, you must create a router platform policy (on each member device) using the same routing protocol and autonomous system (or process ID) number as defined in the GRE Modes policy.

Table 1: GRE Modes Page for GRE or GRE Dynamic IP VPNs

| Element | Description |
|------------------------|-------------|
| Routing Parameters Tab | |

| Element | Description |
|-------------------------------------|---|
| Routing Protocol | <p>Select the required dynamic routing protocol (EIGRP, OSPF, or RIPv2,) or static route to be used for GRE or GRE Dynamic IP.</p> <p>The default routing protocol is EIGRP.</p> <p>For more information about configuring these protocols, see Prerequisites for Successful Configuration of GRE , on page 3.</p> |
| AS Number (EIGRP only.) | <p>The number that is used to identify the autonomous system (AS) area to which the EIGRP packet belongs. The range is 1-65535. The default is 110.</p> <p>An autonomous system (AS) is a collection of networks that share a common routing strategy. An AS can be divided into a number of areas, which are groups of contiguous networks and attached hosts. Routers with multiple interfaces can participate in multiple areas. An AS ID identifies the area to which the packet belongs. All EIGRP packets are associated with a single area, so all devices must have the same AS number.</p> |
| Hello Interval (EIGRP only.) | The interval between hello packets sent on the interface, between 1 and 65535 seconds. The default is 5 seconds. |
| Hold Time (EIGRP only.) | The number of seconds the router will wait to receive a hello message before invalidating the connection. The range is between 1 and 65535. The default hold time is 15 seconds (three times the hello interval). |
| Delay (EIGRP only.) | The throughput delay for the primary route interface, in microseconds. The range of the tunnel delay time is 1-16777215. The default is 1000. |
| Failover Delay (EIGRP only.) | The throughput delay for the failover route interface, in microseconds. The range of the tunnel delay time is 1-16777215. The default is 1500. |
| Bandwidth (EIGRP only.) | <p>The amount of bandwidth available to the primary route interface for the EIGRP packets. You should enter a value that gives priority to the primary route over other routes.</p> <p>You can enter a value in the range 1 to 10000000 kb. The default is 1000 kb.</p> <p>Note By default, the cost of sending a packet on an interface is calculated based on the bandwidth—the higher the bandwidth, the lower the cost.</p> |
| Failover Bandwidth (EIGRP only.) | <p>The amount of bandwidth available to the failover route interface for the EIGRP packets.</p> <p>Enter a value in the range 1 to 10000000 kb. The default is 1000 kb.</p> |

| Element | Description |
|---|--|
| Process Number (OSPF only.) | <p>The routing process ID number that will be used to identify the secured IGP that Security Manager adds when configuring GRE.</p> <p>The range is between 1 and 65535. The default is 110.</p> <p>Security Manager adds an additional Interior Gateway Protocol (IGP) that is dedicated for IPsec and GRE secured communication. An IGP refers to a group of devices that receive routing updates from one another by means of a routing protocol. Each “routing group” is identified by the process number.</p> <p>For more information, see Understanding GRE , on page 2.</p> |
| Hub Network Area ID (OSPF only.) | The ID number of the area in which the hub’s protected networks will be advertised, including the tunnel subnet. You can specify any number. The default is 0. |
| Spoke Protected Network Area ID (OSPF only.) | The ID number of the area in which the remote protected networks will be advertised, including the tunnel subnet. You can specify any number. The default is 1. |
| Authentication (OSPF or RIPv2 only.) | A string that specifies the OSPF or RIPv2 authentication key. The string can be up to eight characters long. |
| Cost (OSPF or RIPv2 only.) | <p>The cost of sending a packet on the primary route interface.</p> <p>If the selected protocol is OSPF, enter a value in the range 1-65535; the default is 100.</p> <p>If the selected protocol is RIPv2, enter a value in the range 1-15; the default is 1.</p> |
| Failover Cost (OSPF or RIPv2 only.) | <p>The cost of sending a packet on the secondary (failover) route interface.</p> <p>You can enter a value in the range 1-65535 for OSPF (the default is 125), or in the range 1-15 for RIPv2 (the default is 2).</p> |
| Filter Dynamic Updates on Spokes | When selected, enables the creation of a redistribution list that filters all dynamic routing updates on the spokes. This forces the spoke devices to advertise (populate on the hub device) only their own protected subnets and not other IP addresses. |
| Tunnel Parameters Tab | |

| Element | Description |
|--|--|
| Tunnel IP | <p>Select the required option to specify the GRE or GRE Dynamic IP tunnel interface IP address.</p> <ul style="list-style-type: none"> • Use Physical Interface—When selected, uses the private IP address of the tunnel taken from the protected network. • Use Subnet—When selected, uses the tunnel IP address taken from an IP range. This is the default. <p>In the Subnet field, enter the private IP address including the unique subnet mask (default is 1.1.1.0/24).</p> <p>If you are also configuring a dial backup interface, enter its subnet in the Dial Backup Subnet field provided (default is 1.1.2.0/24).</p> <p>Note In most cases, when you use a subnet to specify a GRE tunnel interface IP address, Security Manager creates a loopback interface on the device which is used for the tunnel IP address. If the device belongs to a VPN topology whose configurations were discovered by Security Manager, and you configure an IP address directly on the device's GRE tunnel, Security Manager keeps that configuration and does not create a loopback interface on the device. However, a loopback is always configured on a hub in a VPN topology; in a hub-and-spoke VPN topology with multiple hubs, a loopback interface is also configured on the spokes.</p> <ul style="list-style-type: none"> • Use Loopback Interface—When selected, uses the tunnel IP address taken from an existing loopback interface. In the Role field, enter the name of the interface role object that defines the loopback interface name, or click Select to select it from a list or to create a new object. <p>Note To view the new GRE tunnel or loopback interfaces in the Router Interfaces page, you must rediscover the device inventory details after successfully deploying the VPN to the device.</p> |
| Configure Unique Tunnel Source for each Tunnel | <p>When enabled, each GRE tunnel interface in the VPN is assigned a unique tunnel source. In the Tunnel Source IP Range field, enter a subnet IP to be used as tunnel sources.</p> <p>Note When enabled, this feature is set for all GRE tunnel interfaces in the VPN. If you want to assign a specific tunnel source for an interface, use the Peers policy to configure the endpoints for the desired devices; see Defining the Endpoints and Protected Networks.</p> |

| Element | Description |
|--|--|
| Tunnel Source IP Range (GRE Dynamic IP only.) | <p>The private IP address including the unique subnet mask that supports the loopback for GRE. The GRE tunnel interface has an IP address (inside tunnel IP address) which is taken from a loopback interface that Security Manager creates specifically for this purpose.</p> <p>When a spoke has a dynamic IP address, there is no fixed GRE tunnel source address (to be used by the GRE tunnel on the spoke side) or destination address (to be used by the GRE tunnel on the hub side). Therefore, Security Manager creates additional loopback interfaces on the hub and the spoke to use as the GRE tunnel endpoints. You must specify a subnet from which Security Manager can allocate an IP address for the loopback interfaces.</p> |
| Enable IP Multicast | When selected, enables multicast transmissions across your GRE tunnels. IP multicast delivers application source traffic to multiple receivers without burdening the source or the receivers, while using a minimum of network bandwidth. |
| Rendezvous Point | <p>Only available if you selected the Enable IP Multicast check box.</p> <p>If required, you can enter the IP address of the interface that will serve as the rendezvous point (RP) for multicast transmission. Sources send their traffic to the RP. This traffic is then forwarded to receivers down a shared distribution tree.</p> |

Dynamic Multipoint VPNs (DMVPN)

Dynamic Multipoint VPN (DMVPN) is a hub-and-spoke VPN technology that enables better scaling of large and small IPsec VPNs by combining generic routing encapsulation (GRE) tunnels, IP Security (IPsec) encryption, and Next Hop Resolution Protocol (NHRP) routing.

This section contains the following topics:

- [Understanding DMVPN](#) , on page 10
- [Configuring DMVPN](#) , on page 12
- [Configuring GRE Modes for DMVPN](#) , on page 13
- [Configuring Large Scale DMVPNs](#) , on page 17
- [Configuring Server Load Balancing in Large Scale DMVPN](#) , on page 18

Understanding DMVPN

Dynamic Multipoint VPN (DMVPN) enables better scaling of large and small IPsec VPNs by combining generic routing encapsulation (GRE) tunnels, IP Security (IPsec) encryption, and Next Hop Resolution Protocol (NHRP) routing. (For information about large scale DMVPNs, see [Configuring Large Scale DMVPNs](#) , on page 17.)

Security Manager supports DMVPN using the EIGRP, OSPF, and RIPv2 dynamic routing protocols, and GRE static routes. In addition, On-Demand Routing (ODR) is supported. ODR is not a routing protocol. It may be used in a hub-and-spoke VPN topology when the spoke routers do not connect to any router other than the hub. If you are running dynamic protocols, ODR is not suitable for your network environment.

You can use DMVPN on a hub-and-spoke VPN topology only with devices running Cisco IOS Software release 12.3T devices and later, or ASRs running Cisco IOS XE Software 2.x or later (known as 12.2(33)XNA+ in Security Manager). DMVPN is not supported on Catalyst VPN Services Module devices or on High Availability (HA) groups. If your device does not support DMVPN, use GRE dynamic IP to configure GRE for dynamically addressed spokes. See [Understanding GRE Configuration for Dynamically Addressed Spokes](#), on page 5.

The following topics provide more overview information on DMVPN:

- [Enabling Spoke-to-Spoke Connections in DMVPN Topologies](#), on page 11
- [Advantages of DMVPN with GRE](#), on page 12

The following documents on Cisco.com explain DMVPN in further detail:

- *Cisco Dynamic Multipoint VPN: Simple and Secure Branch-to-Branch Communications*—Explains DMVPN technology and where and why you would use it. This data sheet explains the technologies used with DMVPN and the benefits derived from those technologies.
- *Migrating from Dynamic Multipoint VPN Phase 2 to Phase 3*—Explains the difference between phase 2 and phase 3 spoke-to-spoke connections. Creating spoke-to-spoke connections is a configuration option with DMVPN. Phase 3 uses shortcut switching enhancements to increase network performance and scalability.
- Additional white papers and presentations are available at http://www.cisco.com/en/US/products/ps6658/prod_literature.html.

Enabling Spoke-to-Spoke Connections in DMVPN Topologies

You can use DMVPN to essentially create a full-mesh VPN, in which traditional hub-and-spoke connectivity is supplemented by dynamically-created IPsec tunnels directly between the spokes. With direct spoke-to-spoke tunnels, traffic between remote sites does not need to traverse the hub; this eliminates additional delays and conserves WAN bandwidth. Spoke-to-spoke capability is supported in a single-hub or multihub environment. Multihub deployments provide increased spoke-to-spoke resiliency and redundancy.

You can use the 80:20 traffic rule to determine whether to use a pure hub-and-spoke topology or to allow direct spoke-to-spoke connections:

- If 80 percent or more of the traffic from the spokes are directed into the hub network itself, deploy the hub-and-spoke model.
- If more than 20 percent of the traffic is meant for other spokes, consider the spoke-to-spoke model.

For networks with a high volume of IP Multicast traffic, the hub-and-spoke model is usually preferred.

When you configure the GRE Modes policy for a DMVPN, you can elect to allow spokes to create these direct connections. You must select the DMVPN phase to use for these connections:

- **Phase 2**—Spoke to spoke connections go through regional hubs and routing protocol updates from hubs to spokes are not summarized.
- **Phase 3 (Default)**—Spokes can create direct connections with each other and routing updates from hubs to spokes are summarized. This option allows the greatest scalability and reduces latency. Devices must run IOS Software release 12.4(6)T or later; ASRs must run IOS XE Software release 2.4 (called 12.2(33)XND) or later. Security Manager automatically creates a phase 2 configuration for devices running a lower OS version.

For more information on configuring the GRE Modes policy, see [Configuring GRE Modes for DMVPN](#), on page 13.

Related Topics

- [Understanding DMVPN](#), on page 10
- *Cisco Dynamic Multipoint VPN: Simple and Secure Branch-to-Branch Communications*
- *Migrating from Dynamic Multipoint VPN Phase 2 to Phase 3*

Advantages of DMVPN with GRE

Using DMVPN with GRE provides the following advantages:

- **Simplified GRE configuration on the hub**

With GRE, a tunnel is configured on the hub for each connected spoke. With GRE + DMVPN, only one tunnel is configured for all the connected spokes.

- **Support for dynamically addressed spokes**

When using GRE, the physical interface IP address of the spoke routers must be configured as the GRE tunnel destination address, when configuring the hub router. DMVPN enables spoke routers to have dynamic external interface IP addresses, and provides robust configuration that does not have to be redeployed to the device even if the external interface IP address changes. When the spoke comes online, it sends to the hub registration packets that contain the physical interface IP address of the spoke.

- **Dynamic tunnel creation for direct spoke-to-spoke communication**

NHRP enables spoke routers to dynamically learn the external interface IP address of the routers in the VPN network. Using NHRP, the hub maintains an NHRP database of the public interface addresses of all the spokes (the clients). Each spoke registers its real address with the hub when it boots.

When a spoke wants to transmit a packet to another spoke, it can use NHRP to dynamically determine the required destination address of the destination spoke. The hub acts as the NHRP server, handling the request for the source spoke. This enables the dynamic creation of an IPsec+GRE tunnel directly between spoke routers, without having to go through a hub router, thus reducing the delay of multiple encryption and decryption actions on the hub.

Configuring DMVPN

To configure a hub-and-spoke Dynamic Multipoint VPN, use the Create VPN wizard as described in [Creating or Editing VPN Topologies](#). You can also edit the membership of the VPN, or some of its policies, using the described procedures. If you are creating a Large Scale DMVPN, also see [Configuring Large Scale DMVPNs](#), on page 17.

If you need to make changes to other policies and settings, open the policies from the Site-to-Site Manager page, as follows:

- For ISAKMP and IPsec settings, select **VPN Global Settings**. See [Configuring VPN Global Settings](#).
- For IKE proposal policies, select **IKE Proposal**. See [Configuring an IKE Proposal](#).
- For IPsec proposals, select **IPsec Proposal**. See [Configuring IPsec Proposals in Site-to-Site VPNs](#).

- For preshared key policies, select **IKEv1 Preshared Key**. See [Configuring IKEv1 Preshared Key Policies](#).
- For public key (PKI) policies, select **Public Key Infrastructure**. See [Configuring IKEv1 Public Key Infrastructure Policies in Site-to-Site VPNs](#).
- For Generic Routing Encapsulation configuration, including the selection of phase 2 or 3 connections between spokes, select **GRE Modes**. See [Configuring GRE Modes for DMVPN](#), on page 13.
- For server load balancing policies that are used with Large Scale DMVPN, select **Server Load Balance**. See [Configuring Server Load Balancing in Large Scale DMVPN](#), on page 18.

Related Topics

- [Understanding IKE](#)
- [Understanding DMVPN](#), on page 10
- [Enabling Spoke-to-Spoke Connections in DMVPN Topologies](#), on page 11
- [Advantages of DMVPN with GRE](#), on page 12

Configuring GRE Modes for DMVPN

Use the GRE Modes policy to define the routing and tunnel parameters for IPsec tunneling in a DMVPN.

To open the GRE Modes policy:

- ([Site-to-Site VPN Manager Window](#)) Select a DMVPN or Large Scale DMVPN topology, then select **GRE Modes** from the policies list.
- (Policy view) Select **Site-to-Site VPN > GRE Modes**, and create a new policy or select an existing policy. Then, select either DMVPN or Large Scale DMVPN from the **GRE Method** list.

The following table describes the elements on the GRE Modes page for configuring a DMVPN.



Note

When configuring a DMVPN routing policy, Security Manager adds a routing protocol to all the devices in the secured IGP, on deployment. If you want to maintain this secured IGP, you must create a router platform policy (on each member device) using the same routing protocol and autonomous system (or process ID) number as defined in the GRE Modes policy.

Table 2: GRE Modes Page for DMVPN

| Element | Description |
|------------------------|-------------|
| Routing Parameters Tab | |

| Element | Description |
|-------------------------------------|---|
| Routing Protocol | <p>Select the required dynamic routing protocol, or static route, to be used in the DMVPN tunnel.</p> <p>Options include the EIGRP, OSPF, and RIPv2 dynamic routing protocols, and GRE static routes. On-Demand Routing (ODR) is also supported. On-Demand Routing is not a routing protocol. It can be used in a hub-and-spoke VPN topology when the spoke routers connect to no other router other than the hub. If you are running dynamic protocols, On-Demand Routing is not suitable for your network environment.</p> <p>For more information, see Understanding GRE, on page 2.</p> |
| AS Number (EIGRP only.) | <p>The number that is used to identify the autonomous system (AS) area to which the EIGRP packet belongs. The range is 1-65535. The default is 110.</p> <p>An autonomous system (AS) is a collection of networks that share a common routing strategy. An AS can be divided into a number of areas, which are groups of contiguous networks and attached hosts. Routers with multiple interfaces can participate in multiple areas. An AS ID identifies the area to which the packet belongs. All EIGRP packets are associated with a single area, so all devices must have the same AS number.</p> |
| Hello Interval (EIGRP only.) | The interval between hello packets sent on the interface, from 1 to 65535 seconds. The default is 5 seconds. |
| Hold Time (EIGRP only.) | The number of seconds the router will wait to receive a hello message before invalidating the connection. The range is 1-65535. The default hold time is 15 seconds (three times the hello interval) |
| Delay (EIGRP only.) | The throughput delay for the primary route interface, in microseconds. The range of the tunnel delay time is 1-16777215. The default is 1000. |
| Bandwidth (EIGRP only.) | The bandwidth for the primary route interface, in kilobits. The range of bandwidth is 1 to 10000000. The default is 1000. |
| Bandwidth (EIGRP only.) | <p>The amount of bandwidth available to the primary route interface for the EIGRP packets. You should enter a value that gives priority to the primary route over other routes.</p> <p>You can enter a value in the range 1 to 10000000 kb. The default is 1000 kb.</p> <p>Note By default, the cost of sending a packet on an interface is calculated based on the bandwidth—the higher the bandwidth, the lower the cost.</p> |
| Process Number (OSPF only.) | <p>The routing process ID number that will be used to identify the secured IGP that Security Manager adds when configuring DMVPN.</p> <p>The valid range for either protocol is 1-65535. The default is 110.</p> |
| Hub Network Area ID (OSPF only.) | The ID number of the area in which the hub's protected networks will be advertised, including the tunnel subnet. You can enter any number. The default is 0. |

| Element | Description |
|---|--|
| Spoke Protected Network Area ID (OSPF only.) | The ID number of the area in which the remote protected networks will be advertised, including the tunnel subnet. You can enter any number. The default is 1. |
| Authentication Key (OSPF and RIPv2.) | A string that indicates the OSPF or RIPv2 authentication key. The string can be up to eight characters long. |
| Cost (OSPF and RIPv2.) | The cost of sending a packet on the primary route interface. If the selected protocol is OSPF, enter a value in the range 1-65535; the default is 100. If the selected protocol is RIPv2, enter a value in the range 1-15; the default is 1. |
| Allow Direct Spoke to Spoke Connectivity | Whether to enable direct communication between spokes without going through the hub. Select the DMVPN phase you want to use, which determines the types of connections that spokes can make: <ul style="list-style-type: none"> • Phase 2—Spoke to spoke connections go through regional hubs and routing protocol updates from hubs to spokes are not summarized. • Phase 3 (Default)—Spokes can create direct connections with each other and routing updates from hubs to spokes are summarized. This option allows the greatest scalability and reduces latency. Devices must run IOS Software release 12.4(6)T or later; ASRs must run IOS XE Software release 2.4 (called 12.2(33)XND) or later. Security Manager automatically creates a phase 2 configuration for devices running a lower OS version. <p>For detailed information on how phase 2 and 3 differ, see “Migrating from Dynamic Multipoint VPN Phase 2 to Phase 3” on Cisco.com.</p> <p>Note With direct spoke-to-spoke communication, you must use the Main Mode Address option for preshared key negotiation. For more information, see Understanding IKEv1 Preshared Key Policies in Site-to-Site VPNs.</p> |
| Filter Dynamic Updates On Spokes | Unavailable if you are using On-Demand Routing or a static route for your DMVPN tunnel. When selected, enables the creation of a redistribution list that filters all dynamic routing updates (EIGRP, OSPF, and RIPv2) on spokes. This forces the spoke devices to advertise (populate on the hub device) only their own protected subnets and not other IP addresses. |
| Tunnel Parameters Tab | |
| Tunnel IP Range | The IP address range of the inside tunnel interface IP address, including the unique subnet mask. This field defines a subnet, such as 10.1.1.0/24. Note If Security Manager detects that a tunnel interface IP address already exists on the device, and its IP address matches the tunnel’s IP subnet field, it will use that interface as the GRE tunnel. |

| Element | Description |
|-----------------------------|--|
| Dial Backup Tunnel IP Range | If you are configuring a dial backup interface, enter its inside tunnel interface IP address range, including the unique subnet mask. This field defines a subnet. |
| Server Load Balance | When selected, enables the configuration of load balancing on a Cisco IOS router that serves as a hub in a multiple hubs configuration. Server load balancing optimizes performance in a multiple hubs configuration, by sharing the workload. In this configuration, the DMVPN server hubs share the same tunnel IP and source IP addresses, presenting the appearance of a single device to the spokes in a VPN topology. |
| Enable IP Multicast | When selected, enables multicast transmissions across your GRE tunnels. IP multicast delivers application source traffic to multiple receivers without burdening the source or the receivers, while using a minimum of network bandwidth. |
| Rendezvous Point | Only available if you selected the Enable IP Multicast check box. If required, you can enter the IP address of the interface that will serve as the rendezvous point (RP) for multicast transmission. Sources send their traffic to the RP. This traffic is then forwarded to receivers down a shared distribution tree. |
| Tunnel Key | A number that identifies the tunnel key. The default is 1. The tunnel key differentiates between different multipoint GRE (mGRE) tunnel Non Broadcast Multiple Access (NBMA) networks. All mGRE interfaces in the same NBMA network must use the same tunnel key value. If there are two mGRE interfaces on the same router, they must have different tunnel key values. Note To view the newly created tunnel interfaces in the Router Interfaces page for routers that are members of the VPN, you must rediscover the device inventory details after successfully deploying the VPN to the device. |
| NHRP Parameters | |
| Network ID | All Next Hop Resolution Protocol (NHRP) stations within one logical Non-Broadcast Multi-Access (NBMA) network must be configured with the same network identifier. Enter a globally unique, 32-bit network identifier within the range of 1 to 4294967295. |
| Hold time | The time, in seconds, that routers will keep information provided in authoritative NHRP responses. The cached IP-to-NBMA address mapping entries are discarded after the hold time expires. The default is 300 seconds. |
| Authentication | An authentication string that controls whether the source and destination NHRP stations allow intercommunication. All routers within the same network using NHRP must share the same authentication string. The string can be up to eight characters long. |

Configuring Large Scale DMVPNs

You can configure DMVPN for large scale deployments that might comprise thousands of spokes. In large scale DMVPN topologies, IPsec Terminators, also referred to as Server Load Balance (SLB) devices, reside between the spokes and the hubs. The hubs must be directly connected to the IPsec Terminator—there can be no other device between them.

The IPsec Terminator, which is a Catalyst 6500/7600 device, performs encryption and decryption while the hubs handle all tasks related to Next Hop Resolution Protocol (NHRP) and multipoint generic routing encapsulation (mGRE). The IPsec Terminator is configured to specifically load balance GRE traffic to the hubs, and is configured with dynamic crypto to accept any spokes with any proxies. When using tunnel protection on spokes, these proxies are automatically set to match GRE traffic. One GRE tunnel is configured on the spokes. All hubs connecting to the same IPsec Terminator will use the same Tunnel IP address, and the tunnel source is the Virtual IP address of the IPsec Terminator.

In Security Manager, you configure a Large Scale DMVPN during the creation of a new hub-and-spoke VPN topology as described in [Creating or Editing VPN Topologies](#). You cannot edit an existing standard DMVPN and convert it to a Large Scale DMVPN. When you create the Large Scale DMVPN, keep the following points in mind:

- When you define the technology of the VPN, select **DMVPN** as the technology, and **Large Scale with IPsec Terminator** as the type. For the procedure, see [Defining the Name and IPsec Technology of a VPN Topology](#).
- When you select the devices for the VPN, select the required IPsec Terminators (Catalyst 6500/7600 devices), the hubs and all the spokes. For the procedure, see [Selecting Devices for Your VPN Topology](#).

There must be direct connectivity between the IPsec Terminators and the hubs.

- When you configure the endpoints, as described [Defining the Endpoints and Protected Networks](#), configure the following in the Edit Endpoints dialog box:
 - For each hub device, in the Hub Interface tab, select the interface that is connected to the IPsec Terminator. Each hub can be connected to only one IPsec Terminator. Also, identify the protected networks. Each hub in the Large Scale DMVPN must identify itself and its protected networks.
 - For each IPsec Terminator in the Large Scale DMVPN, specify a VPN external interface, the crypto engine slot, and the Inside VLAN. No protected networks are configured on an IPsec Terminator.

After you create the Large Scale DMVPN topology, a Server Load Balance policy is configured on the IPsec Terminators with all the required parameters, which you can edit if required. Initially, all hubs are given the same priority and number of VPN connections. For information on configuring the Server Load Balance policy, see [Configuring Server Load Balancing in Large Scale DMVPN](#), on page 18.



Note VRF-Aware IPsec cannot be configured in a Large Scale DMVPN.

Related Topics

- [Understanding DMVPN](#), on page 10
- [Configuring DMVPN](#), on page 12

Configuring Server Load Balancing in Large Scale DMVPN

Use the Server Load Balance page to view or edit the server load balance policy configured on the IPsec Terminators in a large scale DMVPN. Server load balancing optimizes performance in multiple hub-and-spoke VPN topologies by sharing the workload among a group of hubs. In large scale DMVPN configurations, the IPsec Terminators perform the traffic load balancing. For more information, see [Configuring Large Scale DMVPNs , on page 17](#).

A weighted round robin (WRR) scheduling algorithm is used to control the bandwidth allocated to output transmission queues. Weighting is based on the amount of bandwidth used by each transmit queue on an interface. Packets from queues with higher capacity are transmitted more often than those from queues with less capacity.

To open the Server Load Balance policy, in the [Site-to-Site VPN Manager Window](#), select an existing Large Scale DMVPN topology, then select **Server Load Balance** from the Policies list.

The table displays the hubs in the VPN, the hub's weight relative to other hubs connected to the same IPsec Terminator, and the maximum number of active connections allowed for the hub. To change the weight or maximum connections, select the hub and click the Edit (pencil) button beneath the table to open the [Edit Load Balancing Parameters Dialog Box , on page 18](#).

Related Topics

- [Configuring Large Scale DMVPNs , on page 17](#)
- [Filtering Tables](#)

Edit Load Balancing Parameters Dialog Box

Use the Edit Load Balancing Parameters dialog box to change the server load balance parameters configured on a hub that is connected to an IPsec Terminator in a large scale DMVPN.

Navigation Path

From the Server Load Balance policy, select a hub and click the **Edit (pencil)** button below the table. For information on opening the Server Load Balance policy, see [Configuring Server Load Balancing in Large Scale DMVPN , on page 18](#).

Related Topics

- [Configuring Large Scale DMVPNs , on page 17](#)

Field Reference

Table 3: Edit Load Balancing Parameters Dialog Box

| Element | Description |
|---------|---|
| Weight | The capacity of the hub relative to other hubs connected to the IPsec Terminator, based on the weighted round robin (WRR) scheduling algorithm. You can enter a value between 1 and 255. The default is 1. |

| Element | Description |
|-----------------|---|
| Max Connections | The maximum number of active connections to the IPsec Terminator that are permitted to the hub. You can enter a value between 1 and 65535. The default is 500. |

