# Case Studies

The following two case studies are provided as reference material for implementing p2p GRE over IPsec designs.

# Static p2p GRE over IPsec with a Branch Dynamic Public IP Address Case Study

This case study explains how to encrypt p2p GRE tunnels with a dynamic crypto map to support deployments where the remote IP address is not statically defined. This technique can also be used when there is static IP addressing on the branch office where the goal of the network administrator is to simplify the headend configuration.

## Overview

To fully understand the mechanics of this configuration, is it helpful to visualize the topology as a tunnel within a tunnel. The IP addresses of the crypto endpoints are routable over the network. The Internet is the most common deployment. The headend crypto IP address is a static IP address. The IP address of the outside interface of the branch router is also routable over the Internet but is dynamically assigned by the broadband or Internet service provider, either through DHCP or PPPoE. The branch router always initiates the p2p GRE and crypto tunnel to the crypto headend IP address. The interesting traffic (traffic matching the crypto ACL) is either a GRE keepalive or hello packet of whatever routing protocol is configured on the branch router.

The crypto headend router learns the dynamically assigned IP address of the outside interface of the branch router because it is the source IP address in the ISAKMP packets. Following successful negotiation of IKE Phase 1 (main mode) and subsequently IKE Phase 2 (quick mode), the outer (crypto) tunnel is established providing transport for the p2p GRE tunnel.

The p2p GRE tunnel is transported within the crypto tunnel. The p2p GRE tunnel endpoint IP addresses are only locally significant. The tunnel source is a connected interface, and the tunnel destination IP address must be routed out the interface with the crypto map configured. After the encrypting router has identified the p2p GRE packet as requiring encryption, and is encapsulated in IPsec, the packet is routed on destination IP address in the IPsec header and not the destination address in the p2p GRE header.

As the encrypted packet is forwarded over the network, the p2p GRE destination IP address is simply part of the IPsec payload. All routing decisions are based on the IPsec header. When the IPsec packet is decrypted by the receiving router, the packet forwarding decision is then based on the p2p GRE destination IP address. This p2p GRE destination address is a connected interface on the decrypting router.

# Sample Topology

A typical customer topology includes two p2p GRE tunnels from the branch router to separate crypto headend routers. One tunnel is the primary tunnel and the second is the backup tunnel providing redundancy in the event the primary headend or network connection failure. To simplify the explanation, only one branch router, headend, and tunnel is shown in this section and in Figure 5-1.

*Figure 5-1*        *Sample Topology*
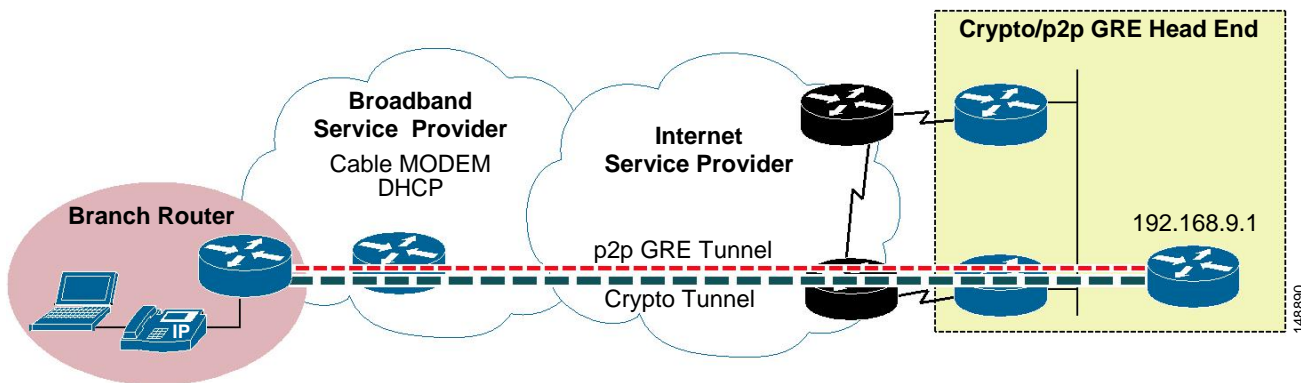


# Addressing and Naming Conventions

In this section, assume that 10.0.0.0/8 address space is the enterprise IP address space and 192.168.0.0/16 represents routable Internet address space.

The addressing scheme and naming conventions are defined as follows:

- Branch

    - Routers in this example are assigned address space within a /22 prefix.

    - This address space is advertised as a summary route to the headend(s).

    - Loopback 0 is used for p2p GRE tunnel source IP address.

    - Loopback 1 is used for Network Management.

    - Tunnel 1 is the p2p GRE tunnel to the primary headend.

    - Tunnel 0 is the p2p GRE tunnel to the secondary headend (if present).

    - The last octet of the IP address of Loopback 0 corresponds to the branch site number (17,18, and so on).

- Headend (primary)

    – Loopback 0 IP address is Internet routable represented by 192.168.9.1.

    – The p2p GRE tunnel interface instance number represents the branch router supported by this interface. For example, branch_17 is Tunnel17, and branch_18 is Tunnel18, and so on.

    – A static or dynamically learned route for 10.63.0.0/16 is present to switch packets out the interface with the dynamic crypto map configured.

- Headend (secondary—configuration not shown)

    – Loopback 0 IP address is Internet routable represented by 192.168.8.1.

    – *All primary headend assumptions also apply to this secondary headend.*

It is a best practice for a network manager to document and define an address plan before implementation. Table 5-1 shows an example. The configuration for Branch 17 is also shown in the configuration details. The data for Branch 18 is shown to illustrate how the addressing scheme progresses for subsequent branch routers. The configuration for Branch 18 or the secondary headend router is not shown.

*Table 5-1      Address Plan*

|  |  |  |  |  | Primary | Headend |
|---|---|---|---|---|---|---|
| Branch 17 |  | Loopback0 | 10.63.0.17/32 | 10.63.0.17 | 192.168.9.1 | Loop0 |
|  | 10.0.4.0/22 | Branch network | Summary | ? |  |  |
|  |  | FastEth0/0 | 10.0.4.0/24 |  |  |  |
|  |  | Loopback1 | 10.0.6.1/32 | 10.0.6.1 |  |  |
|  |  | Tunnel 1 | 10.0.7.4/30 | 10.0.7.5 | 10.0.7.6 | Tunnel 17 |
|  |  | Tunnel 0 | 10.0.7.0/30 | 10.0.7.1 | 10.0.7.2 |  |
| Branch 18 |  | Loopback0 | 10.63.0.18/32 | 10.63.0.18 | 192.168.9.1 | Loop0 |
|  | 10.0.8.0/22 | Branch network | Summary | ? |  |  |
|  |  | FastEth0/0 | 10.0.8.0/24 |  |  |  |
|  |  | Loopback1 | 10.0.10.1/32 | 10.0.10.1 |  |  |
|  |  | Tunnel 1 | 10.0.11.4/30 | 10.0.11.5 | 10.0.11.6 | Tunnel 18 |
|  |  | Tunnel 0 | 10.0.11.0/30 | 10.0.11.1 | 10.0.11.2 |  |

# Configuration Examples

The following sections provide configuration examples.

## p2p GRE Tunnel and Interface Addressing

Based on Table 5-1, a partial configuration example for the branch router (branch_17) is shown.

Items to note include the following:

- The branch router has a static route for 192.168.8.1/32 and 192.168.9.1/32 hosts with the DHCP learned default gateway as the next hop. This network contains the IP address of both the primary and secondary headend addresses 192.168.8.1 and 192.168.9.1.

- A default route is learned through the dynamic routing protocol running in the p2p GRE tunnel following crypto tunnel establishment and routing protocol neighbors are formed on the p2p GRE tunnel and routes are exchanged. The 192.168.8.1/32 and 192.168.9.1/32 routes are a best practice to avoid recursive routing.

```
!
hostname branch_17
!
interface Loopback0
 description For Tunnel Termination
 ip address 10.63.0.17 255.255.255.255
!
interface Loopback1
 description For Network Management ip
 address 10.0.6.1 255.255.255.255
!
interface FastEthernet 0/0
 description Inside (LAN) interface
 ip address 10.0.4.1 255.255.255.0
!
interface Ethernet1/0
 description Outside to CABLE MODEM / Broadband SP / Internet
 ip address dhcp
 crypto map GRE
!
interface Tunnel1
 description Tunnel to (Primary Headend)
 ip address 10.0.7.5 255.255.255.252
 ip summary-address eigrp 44 10.0.4.0 255.255.252.0 5
 tunnel source Loopback0
 tunnel destination 192.168.9.1
!
router eigrp 44
 network 10.0.4.0 0.0.3.255
!
ip route 192.168.8.1 255.255.255.255 dhcp
ip route 192.168.9.1 255.255.255.255 dhcp
!
end
```

The primary headend outside tunnel (for branch 17) and loopback interfaces are shown. The inside interface as well as tunnel interfaces for other branches are not shown.

Note that both the primary and secondary headends are configured using EIGRP as the IGP; however, to emphasis the routing requirements, two static routes are shown. The first is a default (0/0) route using the enterprise campus Internet gateway as a next hop. The second route is a route with a /16 prefix for network 10.63.0.0, also to the enterprise campus Internet gateway as a next hop. This route is not

required, because the default (0/0) route accomplishes that same purpose. It is included to draw attention to the fact the tunnel source IP addresses of the branch routers are allocated from the 10.63.0.0 address space as individual /32 networks. This /16 prefix route illustrates that the encrypting router selects the p2p GRE packets for encryption based on the fact the p2p GRE encapsulated packet is routed out the interface (outside) that contains the crypto map configuration.

```
hostname primary_headend
!
interface Loopback0
 ip address 192.168.9.1 255.255.255.0
!
interface Ethernet1/0
 description Outside Interface
 ip address 10.254.1.49
 255.255.255.0 crypto map DYNO-MAP
!
interface Tunnel17
 ip address 10.0.7.6 255.255.255.252
 tunnel source Loopback0
 tunnel destination 10.63.0.17
!
ip route 0.0.0.0 0.0.0.0 10.254.1.1 name To_INTERNET_GateWay
ip route 10.63.0.0 255.255.0.0 10.254.1.1 name
Branch_Tunnel_Endpoints end
```

## Crypto Map Configurations (Crypto Tunnel)

The crypto map of the branch router is simply a static crypto map matching the p2p GRE tunnel endpoint IP addresses. Only the primary headend is shown.

```
!
hostname branch_17
!
interface Loopback0
 ip address 10.63.0.17 255.255.255.255
!
crypto map GRE 10 ipsec-
 isakmp set peer 192.168.9.1
 set transform-set AES_SHA_TUNNEL
 3DES_SHA_TUNNEL match address GRE_to_NINE
!
ip access-list extended GRE_to_NINE
 permit gre host 10.63.0.17 host 192.168.9.1
!
```

The primary headend is configured with a dynamic crypto map. Note there is no **match address** configured. It is learned dynamically from the branch router.

The **match address** ACL on the crypto headend crypto map is optional. If a remote site has not established its crypto tunnel, the headend router sends p2p GRE packets out the crypto map interface unencrypted. However, this a minimal security exposure for two reasons: the destination IP address of the unencrypted p2p GRE packet in this example is an RFC1918 IP address and is not routed over the Internet by the ISP; and the contents of these packets are either GRE keepalives or routing protocol hello packets. Data packets are not routed until the p2p GRE tunnel is up and the routing protocol has formed a neighbor relationship.

```
hostname primary_headend
!
crypto dynamic-map DYNO-TEMPLATE
 10 description dynamic crypto map
```

```
 set transform-set AES_SHA_TUNNEL 3DES_SHA_TUNNEL
!
crypto map DYNO-MAP local-address Loopback0
crypto map DYNO-MAP 10 ipsec-isakmp dynamic DYNO-TEMPLATE
```

Note that crypto maps are not required on tunnel interface beginning in 12.2(13)T.

## Headend EIGRP Configuration

The headend routers determine what networks are learned by the branch router through the tunnel. This is effectively a policy push from the headend to all branch routers. The branch router learns a default route if the outside IP address is obtained by DHCP, and this default route is inserted in the routing table of the branch router by default with an administrative distance of 254 or least preferred.

The remote router is configured to encrypt p2p GRE packets. The networks advertised by the headend through the tunnel interface determine what gets encrypted or what is routed to the default route to the Internet. If the headend router advertises a default (0/0) route through the p2p GRE tunnel, split-tunnel is not implemented on the branch router. If the headend router advertises, for example, 10.0.0.0/8, split-tunnel is enabled and all packets not destined for 10.0.0.0/8 are routed to the Internet. Of course, the network manager should include an inbound access list as one component in securing the branch router, and this along with Context-Based Access Control (CBAC, the Cisco IOS firewall feature) controls if the return packets are permitted into the branch network.

The following is an example of the headend configuration advertising only a default route to branch 17.

```
hostname primary_headend
!
router eigrp 44
 network 10.0.0.0
 distribute-list Quad_ZERO_to_BRANCH out Tunnel17
!
ip access-list standard Quad_ZERO_to_BRANCH
 permit 0.0.0.0
!
```

It is assumed the network manager has either a default route learned from an EIGRP AS 44 neighbor or the static route to the default network is redistributed into this configuration.

# Verification

The branch_17 router has an outside IP address of 192.168.31.2. Assume this is an Internet routable IP address and was learned by DHCP.

As an illustration, the primary headend dynamic crypto map is shown:

```
primary_headend#show crypto map
Crypto Map: "DYNO-MAP" idb: Loopback0 local address: 192.168.9.1

Crypto Map "DYNO-MAP" 10 ipsec-isakmp
        Dynamic map template tag: DYNO-TEMPLATE

Crypto Map "DYNO-MAP" 65536 ipsec-isakmp
        Peer = 192.168.31.2
        Extended IP access list
            access-list  permit gre host 192.168.9.1 host 10.63.0.17
            dynamic (created from dynamic map DYNO-TEMPLATE/10)
        Current peer: 192.168.31.2
```

Note in the above display that the dynamic crypto map entry has been updated with the crypto map ACL (source and destination IP address are reversed accordingly) configured on branch_17.

The ISAKMP SA is displayed for branch_17, showing that the source IP address of the ISAKMP session is the address learned by DHCP.

```
branch_17# show crypto isakmp sa
dst             src            state          conn-id slot
192.168.9.1     192.168.31.2   QM_IDLE             37   0
```

# Summary

Using the configuration guide in this section, network managers can implement p2p GRE tunnels in network topologies where the IP address of the branch router is learned dynamically. Static crypto maps are configured on the headend router; as branch routers are brought online, no change to the headend crypto configuration is required. The headend router needs only a corresponding p2p GRE tunnel interface for the new branch.

The primary advantage for this configuration compared to a DMVPN configuration is a p2p GRE interface as opposed to an mGRE interface. P2p GRE interfaces support multiprotocol (IPX, Appletalk, and so on) routing as well as interface (and peer) specific configurations such as a QoS service policy.

# Moose Widgets Case Study

The key objective of this case study is to provide a reference example for a site-to-site VPN design. It provides an example of how these design principles can be applied in a real-world scenario.

The details of the service provider backbone and WAN connectivity are not addressed in the case study, because the focus is on VPN deployment on the enterprise customer side.

## Customer Overview

Moose Widgets has been developing products at their Portland, Oregon headquarters (HQ) for several years. In addition, Moose has a single distribution center and 20 retail outlets across the United States (US).

Currently, Moose Widgets uses a traditional Frame Relay (FR) WAN service to connect its headquarters to its distribution center. There is currently no connectivity to retail outlets, with the exception of a few outlets that use a personal computer (PC) to dial-up to the corporate HQ. The current network topology is shown in Figure 5-2.

*Figure 5-2        Moose Widgets Case Study—Current Topology*
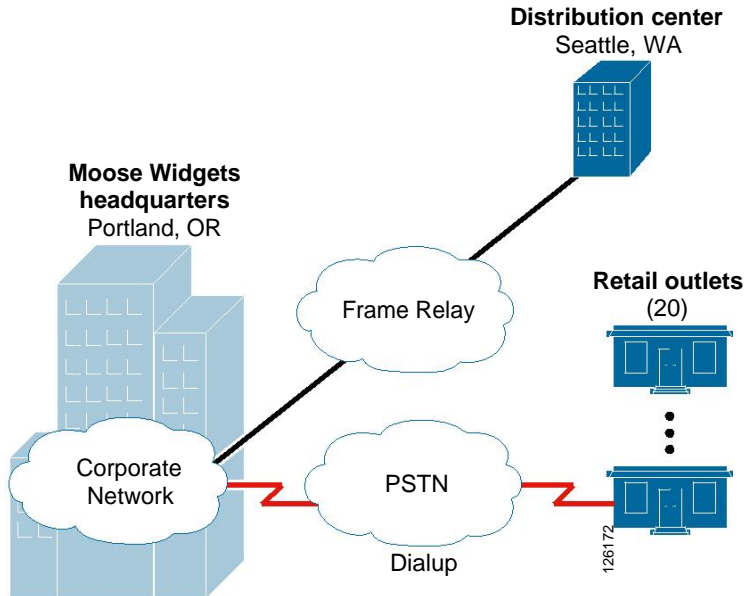


Moose has recently acquired two companies; one in San Jose, California and the other in Great Falls, Montana. Moose wants to connect its newly acquired companies and its retail outlets to its corporate network as an intranet. In addition, Moose plans to expand its retail outlets to 40–50 over the next year and sees already that it will most likely need additional distribution centers on the East and West coasts of the US.

As part of a corporate initiative, Moose is implementing a centralized inventory tracking system to better manage inventory in its growing distribution centers and retail outlets to significantly lower costs. The existing dial-in access does not provide adequate bandwidth to support the new applications. Further, Moose is concerned about escalating dial-in charges as each retail outlet relies more on corporate resources. As a result, Moose is looking to transition to a dedicated connection for each of its retail outlets using the Internet and VPN technology.

Moose is concerned about the costs of adding the connections. They are also concerned about the ability to quickly get retail outlets up and running. Moose indicates they are primarily concerned about data traffic today, but there is some degree of interest in adding voice services in the future.

Moose estimates their traffic requirements for their various site locations as shown in Table 5-2.

*Table 5-2        Moose Widgets Case Study—Bandwidth*

| Location | Connectivity Requirements |
|---|---|
| Distribution center (one today, potentially three in the future) | T1 (3 Mbps bi-directional) |
| San Jose | 2 x T1 (6 Mbps bi-directional) |
| Great Falls | 2 x T1 (6 Mbps bi-directional) |
| Retail outlets (up to 50) | 384 Kbps/1.5 Mbps broadband DSL and cable |

# Design Considerations

A site-to-site IPsec VPN will be deployed with the Moose corporate HQ serving as the headend, and all other locations treated as branch sites. This allows a branch office to subscribe to a local ISP, get authenticated, and be inside the corporate intranet.

At the same time, end-to-end encryption is attained using IPsec tunneling. Switching to VPN offers Moose significant cost savings over dial-up solutions and the ability to outsource to a service provider who has VPN service as a core competency, providing more efficiency with cost and scalability.

Following the design practice outlined in Chapter 2, "Point-to-Point GRE over IPsec Design and Implementation," and Chapter 3, "Scalability Considerations," there are four main design steps to perform:

- Preliminary design considerations assessing customer requirements and expectations
- Sizing of the headend devices and product selection
- Sizing of the branch site devices and product selection
- Tunnel aggregation and load distribution planning

## Preliminary Design Considerations

The design is straightforward and offers flexibility. As new retail locations are put into service, Moose can purchase Internet connectivity from the local ISP, deploy a Cisco VPN router at the branch site, configure the IPsec tunnels to the headend devices at the corporate headquarters, and be up and running in a short amount of time.

Using the questions from Design Considerations, page 2-1, Table 5-3 summarizes the preliminary design considerations.

*Table 5-3*     *Preliminary Design Considerations*

| Question | Answer | Comments |
|---|---|---|
| What applications does the customer expect to run over the VPN? | Data | Interested in future voice services |
| Is multiprotocol support required? | Yes, IP and IP multicast | GRE tunnels will enable multi-protocol traffic transport. |
| How much packet fragmentation does the customer expect on their network? | Minimal | Path MTU discovery enabled |
| How many branches does the customer expect to aggregate to the headend? | 55 sites | |
| What is customer expected traffic throughput to/from branch offices? | See Table 5-2 | |
| What are customer expectations for resiliency? | Resiliency is required | 1 primary, 1 backup tunnel |
| What encryption level is required? | 3DES | |

*Table 5-3*          *Preliminary Design Considerations (continued)*

| | | |
|---|---|---|
| What type of IKE authentication method will be used? | The use of pre-shared keys is selected because of the relatively small number of sites to manage. | Migration to digital certificates should be considered if number of branches increases beyond 50 in the future. |
| What other services will be run on the branch VPN routers? | None | |

EIGRP is recommended as the routing protocol, with route summarization.

# Sizing the Headend

Although the traffic loads involved do not exceed the recommended capacity of a single headend device, Moose has indicated they would like redundancy built-in at the central location. The tunnels from the remote ends will be allocated to each of the headend devices to balance the traffic load. Secondary tunnels will also be configured and allocated so that, in the event of a headend failure, traffic will be transitioned over to the partner headend device.

Applying the sizing algorithm defined in section Headend Scalability, page 3-3, the calculation of headend sizing based on number of GRE tunnels is as follows:

$N = 55$

$T = N \times 2 = 110$

$C(t) = (T / 500)$ rounded up $+ 1 = 110/500$ rounded $+1 = 1 + 1 = 2$ headends

Next, using the throughput estimates from Table 5-4, the calculation of headend sizing based on branch traffic throughput is as follows:

$A = (3*3 \text{ Mbps}) + 6 \text{ Mbps} + 6 \text{ Mbps} + (50*1.9 \text{Mbps}) = 115 \text{ Mbps} * 50\%$ utilization $= 58$ Mbps

$H = 109$ Mbps (for Cisco 7206VXR NPE-G1 with SA-VAM2)

$C(a) = A/H$, rounded up $+ 1 = 58/109$ rounded up $+ 1 = 1+1 = 2$ headends

Comparing the number of headend devices calculated based on number of tunnels, $C(t)$, to the number based on aggregate throughput, $C(a)$, the outcomes match. Therefore it is appropriate to deploy two headend devices.

Presented with the headend product options, the customer selects to deploy two Cisco 7206VXR NPE-G1s, each equipped with an SA-VAM2 hardware encryption adapter.

# Sizing the Branch Sites

The primary consideration for sizing of branch office sites is expected traffic throughput. Accordingly, starting with Table 5-2, and applying the concepts presented in Branch Office Scalability, page 3-9, the branch products selected are summarized in Table 5-4.

*Table 5-4*          *Moose Widgets Case Study—Branch Devices*

| Location | Estimated Throughput | Branch Office Platform Selected |
|---|---|---|
| Distribution centers | 3 Mbps | Cisco 2851 ISR |

***Table 5-4       Moose Widgets Case Study—Branch Devices***

| San Jose | 6 Mbps | Cisco 3845 ISR |
|---|---|---|
| Great Falls | 6 Mbps | Cisco 3845 ISR |
| Retail outlets (typical) | 384 Kbps/1.5 Mbps DSL/cable | Cisco 2801/1841 ISR |

At each of the acquired company locations, a Cisco 3845 ISR is deployed. The choice of the Cisco 3845 platform is based on the assumption that the acquired companies are large offices with a substantial number of employees. Future VoIP expansion is also a factor.

At each of the distribution centers, a Cisco 2851 ISR is deployed. Finally, at each of the retail locations, a combination of Cisco 2801 ISR and 1841 ISR are deployed, depending on the size of the retail outlets.

## Tunnel Aggregation and Load Distribution

Given 55 branch sites, the total number of tunnels that need to be aggregated is 110 (primary and secondary). Therefore, the first headend device is allocated 27 primary and 28 backup tunnels, while the second headend device is allocated 28 primary and 27 backup tunnels.

# Network Layout

The new network topology is shown in .

***Figure 5-3       Moose Widgets Case Study—VPN Topology***