Point-to-Point GRE over IPsec Design Overview

This chapter provides an overview of the VPN site-to-site design topology and characteristics. Chapter 2, "Point-to-Point GRE over IPsec Design and Implementation," provides more detail on the design considerations. Chapter 3, "Scalability Considerations," presents Cisco product options for deploying the design.

Starting Assumptions

The design approach presented in this design guide makes the following starting assumptions:

- The design supports a typical converged traffic profile for customers (see Chapter 4, "Scalability Test Results (Unicast Only)."
- It is assumed that the customer has a need for diverse traffic requirements, such as IP multicast, multiprotocol, and support for routing. The use of p2p GRE and a routing protocol are also discussed in more detail in Chapter 2, "Point-to-Point GRE over IPsec Design and Implementation."
- Cisco products should be maintained at reasonable CPU utilization levels. This is discussed in more detail in Chapter 3, "Scalability Considerations," including recommendations for both headend and branch routers, and software revisions.
- Although costs were certainly considered, the design recommendations assume that the customer will deploy current VPN technologies, including hardware-accelerated encryption.
- Voice over IP (VoIP) and video are assumed to be requirements in the network. Detailed design considerations for handling VoIP and other latency-sensitive traffic are not explicitly addressed in this design guide, but may be found in *Voice and Video Enabled IPsec VPN (V3PN)*, which is available at the following URL: http://www.cisco.com/en/US/docs/solutions/Enterprise/WAN_and_MAN/V3PN_SRND/V3PN_SRND.html.
- Finally, this design is targeted for deployment by enterprise-owned VPNs; however, the concepts and conclusions are valid regardless of the ownership of the edge tunneling equipment, and are therefore valuable for service provider-managed VPNs as well.

Quality of Service per p2p GRE Tunnel Interface

All headend scalability results were performed on various Cisco platforms. Results were obtained without a hierarchical class-based weighted fair queueing service policy applied per p2p GRE tunnel interface. This QoS configuration is a parent service policy that shapes packets in the logical tunnel interface and then queues packets within the shaped rate in a child service policy.

Applying a hierarchical class-based weighted fair queueing service policy to the p2p GRE tunnel interface may require packets to be process-switched while the shaper is active; thus, it is CPU intensive.

Note The *Dynamic Multipoint VPN (DMVPN) Design Guide (*http://www.cisco.com/en/US/docs/solutions/Enterprise/WAN_and_MAN/DMVPDG.html), provide performance guidance on per tunnel interface QoS.

Design Components

VPNs have many applications, including extending reachability of an enterprise WAN, or replacing classic WAN technologies such as leased lines, Frame Relay, and ATM. Site-to-site VPNs are primarily deployed to connect branch office locations to the central site (or sites) of an enterprise.

The requirements of enterprise customers for traditional private WAN services such as multiprotocol support, high availability, scalability, and security, are also requirements for VPNs. VPNs can often meet these requirements more cost-effectively and with greater flexibility than private WAN services.

The key components of this site-to-site VPN design are the following:

- Cisco high-end VPN routers serving as VPN headend termination devices at a central campus (headend devices)
- Cisco VPN access routers serving as VPN branch termination devices at branch office locations (branch devices)
- p2p GRE over IPsec to perform headend-to-branch interconnections
- Internet services procured from a third-party ISP (or ISPs) serving as the WAN interconnection medium

Cisco VPN routers are a good choice for site-to-site VPN deployments because they can accommodate any network requirement inherited from a Frame Relay or private line network, such as support for IP multicast and latency-sensitive traffic, routing for resiliency, and support for non-IP protocols such as IPX or SNA. See Chapter 3, “Scalability Considerations,” for a discussion on selection of headend and branch products.

Topology

In a p2p GRE over IPsec design, the following three topologies can be implemented:

- Hub-and-spoke
- Partial mesh
- Full mesh

The hub-and-spoke topology is discussed in this design guide because it is the most widely deployed.

Headend System Architectures

This section describes the two headend system architectures that can be implemented, depending on the scalability requirements.

Single Tier Headend Architecture

In a Single Tier Headend Architecture, both the p2p GRE and crypto functionally co-exist on the same router CPU. Figure 1-1 shows this hub-and-spoke topology.

Figure 1-1 Single Tier Headend Architecture

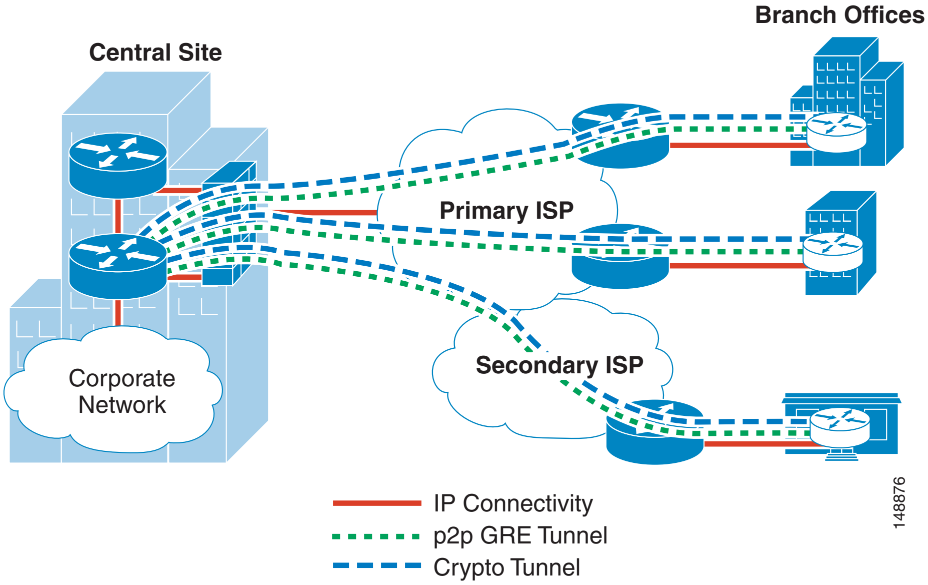

Figure 1-1 shows a hub-and-spoke network with multiple headend devices for redundancy. Headends are p2p GRE and crypto tunnel aggregation routers servicing multiple p2p GRE over IPsec tunnels for a prescribed number of branch office locations. In addition to terminating the VPN tunnels at the central site, headends can advertise branch routes using IP routing protocols such as EIGRP or OSPF.

Dual Tier Headend Architecture

In a Dual Tier Headend Architecture, the p2p GRE and crypto do not functionally co-exist on the same router CPU. Figure 1-2 shows this hub-and-spoke topology.

Figure 1-2 Dual Tier Headend Architecture

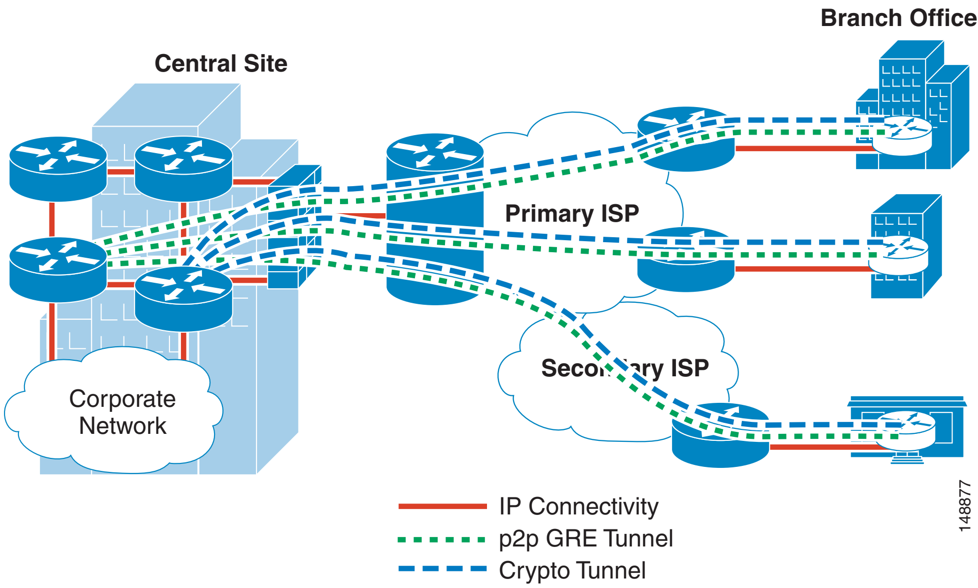

Figure 1-2 shows a hub-and-spoke network with multiple headend devices for redundancy. p2p GRE headends as well as crypto headends together service multiple p2p GRE over IPsec tunnels for a prescribed number of branch office locations. In addition to terminating the VPN tunnels at the central site, the p2p GRE headends can advertise branch routes using IP routing protocols such as EIGRP or OSPF.

Single Tier Headend Architecture versus Dual Tier Headend Architecture

The choice between the Single Tier Headend Architecture and the Dual Tier Headend Architecture depends on the following criteria:

- Topology
- Performance
- Value

In either architecture, the topology being designed is the first consideration. For this comparison, a hub-and-spoke topology is the only topology detailed because this discussion is focused on the headend router. Table 1-1 lists the technical limitations of the two architectures.

Table 1-1 Single Tier Headend versus Dual Tier Headend Architecture—Technical Limitations

Headend Architecture	Router	Crypto Configuration	Crypto IP Address	GRE Configuration	GRE IP Address	Tunnel Protection
Single Tier	Headend	Static or dynamic	Static	p2p GRE static	Static	Optional

Table 1-1 Single Tier Headend versus Dual Tier Headend Architecture—Technical Limitations

	Branch	Static	Static or dynamic	p2p GRE static	Static	Optional
Dual Tier	Headend	Static or dynamic	Static	p2p GRE static	Static	Not valid
	Branch	Static	Static or dynamic	p2p GRE static	Static	Not valid

Tunnel protection requires the same source and destination IP address for both the GRE tunnel and the crypto tunnel; therefore, it is not a valid option in a Dual Tier Headend Architecture. Both architectures support all of the options above as a hub-and-spoke topology.

When considering performance and value, the two should be considered together. Performance is based on the number of packets a router can forward in a given timeframe, or packets per second (pps). Value is the price for a specific router based on the pps rate. Given these two considerations, Table 1-2 shows both price and performance for the primary headend router choices currently available.

Table 1-2 Platform Price and Performance

Platform	pps (bi-directional)	Price (reference only)
Cisco 7200VXR with NPE-G1	p2p GRE only—200 Kpps	$20,000
Cisco 7200VXR with NPE-G1 and Dual SA-VAM2+	p2p GRE and 3DES—40 Kpps	$30,000
Cisco 7606 with Sup720, SSC400, and VPN SPA	p2p GRE and 3DES—520 Kpps	$100,000
Cisco 7606 with Sup720, SSC400, and VPN SPA	3DES only—600 Kpps	$100,000

For the purpose of this design guide, *these prices are not the actual price* of the specific platforms because prices change on a periodic basis. The prices are provided as a reference to show the value of one architecture over the other. The reader should obtain specific pricing per platform when comparing the platforms. However, the pps values have been verified in Cisco testing.

With this is mind, now consider the value for each of the architectures, given the performance numbers stated for a p2p GRE over IPsec design. The Single Tier Headend Architecture is the easiest to demonstrate because both the p2p GRE and encryption processes are housed on a single routing processor. Table 1-3 shows these results.

Table 1-3 Single Tier Headend Architecture Comparison

Platform	pps (bi-directional)	Quantity required	Aggregate pps (bi-directional)	Price (reference only)
Cisco 7200VXR with NPE-G1 and Dual SA-VAM2+	p2p GRE and 3DES—40 Kpps	3	120 Kpps (40 Kpps * 3)	$90,000 ($30,000 * 3)

Table 1-3 Single Tier Headend Architecture Comparison

Cisco 7200VXR with NPE-G1 and Dual SA-VAM2+	p2p GRE and 3DES—40 Kpps	13	520 Kpps (40 Kpps * 13)	$390,000 ($30,000 * 13)
Cisco 7606 with Sup720, SSC400, and VPN SPA	p2p GRE and 3DES—520 Kpps	1	520 Kpps	$100,000

Note from this comparison that the break-even point regarding value per performance is approximately 120 Kpps at $100,000. Below this price point, the clear choice is to deploy multiple Cisco 7200VXRs with NPE-G1 and Dual SA-VAM2+; however, above this point, the decision is much different. To obtain the same 520 Kpps with the Cisco 7200VXR with NPE-G1 and Dual SA-VAM2+ solution, it requires thirteen chasses at a cost of $390,000, versus a single Cisco 7606 with Sup720, SSC400, and VPN SPA at $100,000. The difference represents a substantial capital savings over the long term. Also, support contracts are increased from one to thirteen as well.

The Dual Tier Headend Architecture is slightly harder to demonstrate because the p2p GRE and encryption processes are housed on separate routing processors. Table 1-4 shows these results.

Table 1-4 Dual Tier Headend Architecture Comparison

Platform	pps (bi-directional)	Quantity required	Aggregate pps (bi-directional)	Price (reference only)
Cisco 7200VXR with NPE-G1 and Dual SA-VAM2+	p2p GRE and 3DES—40 Kpps	5	200 Kpps (40 Kpps *5)	$150,000 ($30,000 * 5)
Cisco 7200VXR with NPE-G1 and Dual SA-VAM2+	p2p GRE and 3DES—40 Kpps	15	600 Kpps (40 Kpps * 15)	$450,000 ($30,000 * 15)
Cisco 7606 with Sup720, SSC400, and VPN SPA for encryption only and Cisco 7200VXR with NPE-G1 for p2p GRE *only*	p2p GRE and 3DES—200 Kpps	1 Cisco 7606 and 1 Cisco 7200VXR	200 Kpps (200 Kpps p2p GRE only on the Cisco 7200VXR is the limitation)	$120,000 ($100,000 Cisco 7606 with Sup720, SSC400, and VPN SPA) + ($20,000 Cisco 7200VXR with NPE-G1)
Cisco 7606 with Sup720, SSC400, and VPN SPA for encryption only and Cisco 7200VXR with NPE-G1 for p2p GRE *only*	p2p GRE and 3DES—600 Kpps	1 Cisco 7606 and 3 Cisco 7200VXRs	600 Kpps (200 Kpps p2p GRE on each Cisco 7200VXR with 600 Kpps on a single VPN SPA)	$160,000 ($100,000 Cisco 7606 with Sup720, SSC400, and VPN SPA) + (3 * $20,000 Cisco 7200VXR with NPE-G1)

Note from this comparison that the break-even point regarding value per performance is approximately 200 Kpps at $150,000. Below this price point, the clear choice is to deploy multiple Cisco 7200VXRs with NPE-G1 and Dual SA-VAM2+; however, above this point, the decision is much different. To obtain the same 600 Kpps with the Cisco 7200VXR with NPE-G1 and Dual SA-VAM2+ solution, it requires fifteen chasses at a cost of $450,000, versus a single Cisco 7606 with Sup720, SSC400, and VPN SPA with three Cisco 7200VXRs with NPE-G1 at $160,000. The difference represents a substantial capital savings over the long term. Also, support contracts are increased from 4 to 15 as well.

Branch Router Considerations

Branches are typically access routers that provide p2p GRE over IPsec tunnel(s) from the branch office locations to the central site. In addition to terminating the VPN tunnels, the branch router often provides WAN access, and in some implementations may serve as a firewall.

Static p2p GRE over IPsec with a Branch Static Public IP Address

In this scenario, the public IP address of the branch router is a statically defined IP address. Both the p2p GRE and crypto tunnels are sourced from this statically defined public IP address.

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

In this scenario, the public IP address of the branch router is a dynamically assigned IP address. The p2p GRE tunnel is sourced from a loopback interface with an administratively assigned private IP address. The crypto tunnel is sourced from the dynamically assigned public IP address. A static host route is required in the headend router to ensure p2p GRE packets destined to the branch router loopback interface are encrypted.

High Availability

Network resiliency is provided differently depending on the initial network requirements. This design guide uses a dynamic IGP routing protocol across the VPN. Because IPsec does not provide the ability to run protocols requiring IP multicast (such as EIGRP), it is necessary to use p2p GRE in conjunction with IPsec. p2p GRE supports more diverse traffic across the VPN, including IP multicast and non-IP protocols.

For high availability in the case of a failure, each branch access router should have two p2p GRE over IPsec tunnels, a primary and secondary, provisioned to different headend tunnel aggregation routers. This is discussed further in High Availability, page 2-17.

Best Practices and Known Limitations

The following sections contain a summary of the best practices and limitations for the design. More detailed information is provided in Chapter 2, “Point-to-Point GRE over IPsec Design and Implementation.”

Best Practices Summary

The following list summarizes the best practices for a p2p GRE over IPsec design, supporting multiprotocol and/or IP multicast traffic including routing protocols:

- General best practices
  - Use IPsec in tunnel mode for best flexibility.
  - Configure Triple DES (3DES) or AES for encryption of transported data (exports of encryption algorithms to certain countries may be prohibited by law).
  - Implement Dead Peer Detection (DPD) to detect loss of communication between peers.
  - Deploy hardware-acceleration of IPsec to minimize router CPU overhead, to support traffic with low-latency/jitter requirements, and for the highest performance for cost.
  - Keep IPsec packet fragmentation to a minimum on the customer network by setting MTU size or using PMTU Discovery (PMTUD).
  - Use Digital Certificates/PKI for scalable tunnel authentication keys.
  - Set up QoS service policies as appropriate on headend and branch router interfaces to ensure performance of latency-sensitive applications (for more information, see the design guides at the following URL*:* http://www.cisco.com/en/US/netsol/ns742/networking_solutions_program_category_home.html.)
  - Configure a routing protocol such as EIGRP or OSPF with route summarization for dynamic routing.
- Headend best practices
  - Configure dynamic crypto maps on the headend to support dynamically addressed branches and to simplify provision of new branches.
  - If high availability is a requirement, implement a design with redundancy of headend equipment and WAN circuits.
  - Distribute branch office tunnels across a number of headend routers to balance loading and aggregation capacity of the hub(s).
  - Select Cisco VPN router products at the headend based on considerations for the following:
    - Number of tunnels to be aggregated
    - Maximum throughput in both pps and bps to be aggregated
    - Performance margin for resiliency and failover scenarios
    - Maintaining CPU utilization below design target
  - See Chapter 3, "Scalability Considerations," for more information.
- Branch office best practices
  - Configure multiple p2p GRE over IPsec tunnels to redundant headends.
  - Select Cisco VPN router products at the branch offices based on considerations for the following:
    - Maximum throughput in both pps and bps
    - Allowances for other integrated services that may be running on the router, such as firewall, IPS, and NAT/PAT
    - Maintaining CPU utilization below 65–80 percent

    - See Chapter 3, “Scalability Considerations,” for more information.
    - The QoS pre-classify feature is desirable in VPN designs where both QoS and IPsec occur on the same system. The network manager should verify correct operation.

Known Limitations Summary

The following lists at a high level the known limitations for a p2p GRE over IPsec VPN design:

- General limitations
    - p2p GRE acceleration is currently limited to 2047 tunnels with the VPN SPA and VPNSM, and 2000 tunnels with the Sup720. Other factors such as the number of sustainable routing peers may affect the maximum number of tunnels in a design.
    - Although IPsec can typically scale to thousands of tunnels on some platforms, a routed p2p GRE over IPsec design is generally limited by the routing protocol being used and the number of routing peers exchanging routing information, such as the following:
        - 500 for the Cisco 7200VXR with NPE-G1
        - 600 for the Cisco 7200VXR with NPE-G2
        - 1000 for the Cisco 7600 (or Catalyst 6500) with Sup720
    - There are significant scalability limitations for supporting IP multicast over p2p GRE over IPsec designs. For more information, see the *Multicast over IPsec VPN Design Guide* at the following URL:
http://www.cisco.com/en/US/docs/solutions/Enterprise/WAN_and_MAN/V3PNIPmc.html.
    - QoS pre-classify is not a supported feature in the Cisco Catalyst 6500 or on the Cisco 7600 platforms.
    - p2p GRE over IPsec designs do not readily support a “touchless” headend configuration concept, and therefore generally require configuration changes to the headend router(s) to provision new branch office tunnels.
    - GRE keepalives are not currently functional on the Cisco 7600 because of DDTS.
- Single Tier Headend Architecture limitations
    - It is possible to implement a QoS service policy at the tunnel/destination level to achieve per-VPN tunnel QoS and prevent hub-to-spoke overruns. However, the number of branches that can be supported in this configuration is limited if shaping is engaged concurrently for a large percentage of the configured tunnels.
- Dual Tier Headend Architecture limitations
    - Tunnel protection is not supported.
- Branch office limitations
    - The p2p GRE over IPsec tunnel must be initiated by the remote branch in cases where remote routers acquire their address via a dynamically served IP address. The crypto headend cannot initiate the tunnel to the branch. However, if a dynamic routing protocol is configured, the hello packets initiated by the branch router force the tunnel to be established when the branch router is started, and the tunnel is maintained by this control plane traffic.
    - In designs with QoS and IPsec, interaction between QoS and IPsec anti-replay can result in dropped packets if packets delayed by QoS fall outside the anti-replay sequence number window at the receiver.

Additional detailed information on these recommendations is discussed in the chapters that follow.