CHAPTER 3

Scalability Considerations

This chapter presents the steps to selecting Cisco products for a VPN solution, starting with sizing the headend, and then choosing Cisco products that can be deployed for headend devices. This chapter concludes with product sizing and selection information for branch devices.

General Scalability Considerations

This section provides general scalability considerations to assist with design requirements.

IPsec Encryption Throughput

The throughput capacity of the IPsec encryption engine in each platform (headend or branch) must be considered for scalable designs, because each packet that is encrypted must traverse through the encryption engine.

Encryption throughput must therefore consider bi-directional speeds. Several examples are shown in Table 3-1 and Table 3-2 for popular headend and branch connection speeds.

Table 3-1 Headend Connection Speeds

Connection Type	Speed (in Mbps)	Encryption throughput required (in Mbps)
T3/DS3	44.7	90.0
OC3	155.0	310.0
OC12	622.0	1250.0

Table 3-2 Branch Connection Speeds

Connection Type	Speed (in Mbps)	Encryption Throughput Required (in Mbps)
T1	1.5	3.0
2 x T1	3.0	6.0
T3/DS3	44.7	90.0
Broadband cable/DSL	384 Kbps uplink/ 2 Mbps downlink	2.4

In general, as throughput increases, the burden on router CPU also increases. However, with hardware-accelerated encryption available for all Cisco router products from the 871 through the 7600, impact on the main CPU is offloaded to the VPN hardware. Main router CPU processing still occurs, however, so higher throughput typically results in higher CPU utilization.

Packets Per Second—Most Important Factor

Although bandwidth throughput capacity must be considered, even more important is the packet rate for the connection speeds being terminated or aggregated.

In general, routers and encryption engines have upper boundaries for processing a given number of packets per second (pps). Size of packets used for testing and throughput evaluations can understate or overstate true performance. For example, if a router with a VPN module can handle 20 Kpps, then 100-byte packets lead to 16 Mbps throughput, while 1400-byte packets at the same packet rate lead to 224 Mbps.

Because of such a wide variance in throughput, pps is generally a better parameter than bits per second (bps) to determine router forwarding potential. Scalability of the headend is the aggregate forwarding potential of all branches that terminate a tunnel to that headend. Therefore, the aggregate pps from all branches impacts the pps rate of that headend.

Tunnel Quantity Affects Throughput

Although it is highly dependent on platform architecture, the overall throughput generally tends to decrease as tunnel quantities are increased. When a router receives a packet from a different peer than the one whose packet was just decrypted, a lookup based on the security parameters index of the new packet must be performed. The transform set information and negotiated key of the new packet is then loaded into the hardware decryption engine for processing. Traffic flowing on larger numbers of SAs tends to negatively affect throughput performance.

Increasingly, platforms with hardware-accelerated IPsec encryption are designed to offload tunnel processing overhead as well, resulting in more linear performance regardless of the number of tunnels. For example, the VPN SPA blade for the Cisco 7600 has relatively linear throughput regardless of whether the traffic load is offered on a few tunnels or several thousand. The Cisco VPN Services Adapter (VSA) for Cisco 7200 Series also incorporates these features and as such, performance is improved accordingly over a VAM2+ adapter.

GRE Encapsulation Affects Throughput

Router encryption throughput is affected by the configuration of GRE. In addition to the headers that are added to the beginning of each packet, these headers must also be encrypted. The GRE encapsulation process, when not hardware-accelerated, increases total CPU utilization. Total throughput in a p2p GRE over IPsec design results in a lower throughput than that of an IPsec Direct Encapsulation design.

Routing Protocols Affect CPU Overhead

CPU overhead is affected by running a routing protocol. Router processing of keepalives and maintenance of a routing table uses a finite amount of CPU time, which varies with the number of routing peers and the size of the routing table. The network manager should design the routing protocol based on well-known and accepted practices.

Headend Scalability

Headend devices are primarily responsible for the following:

- Terminating p2pGRE over IPsec tunnels from the branch routers
- Running a routing protocol inside the p2p GRE tunnels to advertise internal routes to the branches
- Providing redundancy to eliminate the possibility of a single point of failure

It is important to size the headend correctly before choosing the devices to deploy. This ensures that the overall network can support the intended (and possibly future) traffic profiles that the enterprise wants to run over the VPN.

The following critical factors must be considered when sizing the headend:

- How many branch offices need to be connected to the headend? This information provides the number of primary tunnels requiring aggregation.
- What is the expected traffic profile, including the average pps and bps throughput rates for each branch office? This information provides the aggregated data throughput required across the VPN.
- What is the headend connection speed?
- What is the high availability requirement for the design?
- What is the expected performance margin or target CPU utilization?

In general, it is recommended that headend devices be chosen so that CPU utilization does not exceed the target value. This recommendation is to ensure that the device has enough performance remaining to deal with various events that take place during the course of normal network operations, including network re-convergence in the event of a failure, re-keying IPsec SAs, and bursts of traffic.

However, keep in mind that the target value is just a guideline and your customer requirements should determine where to set the limits of operation. In addition, for some platforms such as the Catalyst 6500 and Cisco 7600 router, CPU is not an accurate reflection of performance limits. Other factors may limit the performance, such as backplane and packet switching speeds.

To provide an idea of design scalability and limits, several different design topologies and headend platforms have been evaluated under typical customer traffic profiles. These scalability test results are presented in Chapter 4, “Scalability Test Results (Unicast Only).”

The primary platform choices available today for headend routers are the following:

- Cisco 7200VXR with NPE-G1 and SA-VAM2+ encryption module
- Cisco 7200VXR with NPE-G2 and either a SA-VAM2+ or VPN Services Adapter (VSA)
- Cisco Catalyst 6500 or Cisco 7600 with Sup720, SSC400, and VPN SPA

The Cisco 7301 router is also an option, with performance nearly identical to the Cisco 7200VXR NPE-G1. The main difference is the Cisco 7200VXR will be upgradeable at some point in the future to newer and faster processing engines and encryption modules. The Cisco 7301 is a fixed-configuration platform, and is not upgradeable.

Tunnel Aggregation Scalability

The maximum number of IPsec tunnels that a headend can terminate must be considered. Tunnel scalability is a function of the number of branch routers that are terminated to the headend aggregation point. This number needs to include both the primary tunnels as well as any alternate tunnels for which each headend may be responsible in the event of a failover situation.

The number of IPsec tunnels that can be aggregated by a platform is used as the primary determining factor in recommending a platform. Equally or more important is the encryption pps rate.

Aggregation Scalability

Aside from the number of tunnels that a headend terminates, the aggregated pps must be considered. Requirements are influenced by several factors, including the following:

- Headend connection speed—What is the speed of the WAN link on which the IPsec tunnels of the branch routers are transported through at the headend? (DS3, OC3, OC12, other?)
- Branch connection speeds—What is the typical bandwidth at each branch office going to be? (Fractional-T1, T1, T3, broadband DSL/cable, other?)
- Expected utilization—What is the maximum utilization of the WAN bandwidth under normal operation (or perhaps peak, depending on customer requirements)?

The pps rate (traffic size and traffic mix) is the largest single factor to branch router scalability.

Customer Requirement Aggregation Scalability Case Studies

This section includes examples to illustrate headend scalability factors.

Customer Example with 300–500 Branches

A customer has the design requirements shown in Table 3-3:

Table 3-3 Customer Requirements

Customer Requirement	Value
Number of branch offices	300
Branch access speeds	256 Kbps FR PVCs
Headend access speed	DS3 (90 Mbps bi-directional))
Expected utilization	75%

The calculation of aggregate bandwidth requirements is as follows:

- Typical case—300 x 256 Kbps x 2 (bi-directional) x 80% utilization = 122 Mbps
- Worst case—300 x 256 Kbps x 2 (bi-directional) x 100% utilization = 155 Mbps

Even though the worst case aggregation option calculated is 155 Mbps, the total headend connection speed of DS3 (90 Mbps bi-directional) is the constraining factor. In this example, the key factor to consider is a headend router platform that can support 300 tunnels and provide an aggregate encryption bandwidth of at least 61 Mbps. The traffic mix on the network determines the pps load on the processor. Chapter 4, “Scalability Test Results (Unicast Only),” includes tables that use a traffic mix commonly found on enterprise networks (e-mix) to determine pps based on bps.

At a minimum, a Cisco 7200VXR with NPE-G1 processor and SA-VAM2+ encryption accelerator supports these requirements. For additional future growth capacity, a Cisco 7200VXR with NPE-2 processor and VSA could also be positioned. For platform-specific results, see Headend Scalability Test Results—p2p GRE over IPsec, page 4-3. This design is shown in Figure 3-1.

Figure 3-1 Cisco 7200VXR-Based p2p GRE over IPsec VPN Design

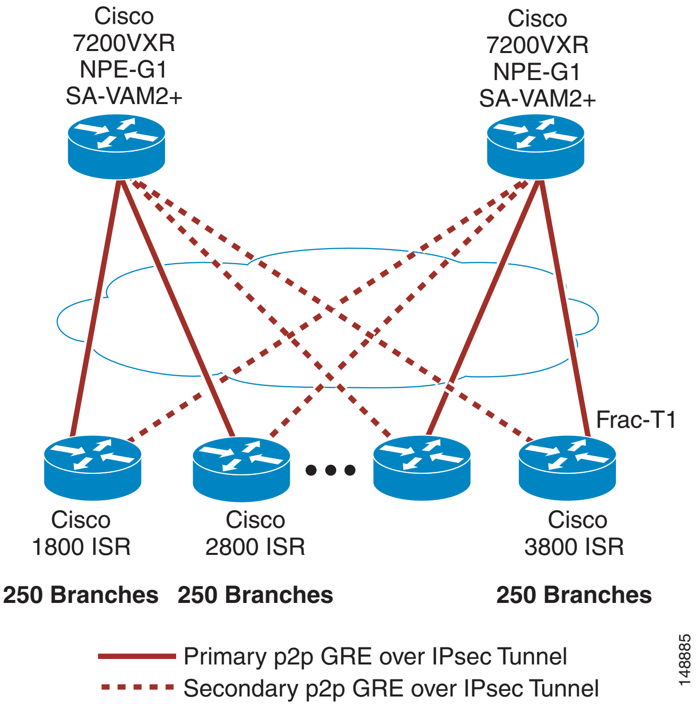

In this design, each Cisco 7200VXR router with NPE-G1 and SA-VAM2+ is handling half of the branch routers (250) and traffic load (45 Mbps) under normal circumstances. Both p2p GRE and IPsec are being terminated directly on each Cisco 7200VXR.

If either of the headends experiences a failure, the surviving headend can handle the total number of branches (500) and traffic load (90 Mbps) during the failure.

The headends can reside in the same or separate geographic locations. If in separate locations, it is likely that each location has independent DS3 (in this example) connections. Separate DS3s should be factored into the maximum aggregation bandwidth requirements to make sure the platforms recommended can handle the load.

Customer Example with 1000 Branches

Next, consider another customer example with many branch locations, each with relatively small traffic requirements, such as a point-of-sale terminal model, as shown in Table 3-4.

Table 3-4 Customer Requirements

Customer Requirement	Value
Number of branch offices	1000
Branch access speeds	128 Kbps FR PVCs
Headend access speed	DS3 (90 Mbps bi-directional))
Expected utilization	25%

The calculation of aggregate bandwidth requirements is as follows:

- Typical case—2000 x 128 Kbps x 2 (bi-directional) x 25% utilization = 64 Mbps
- Worst case—2000 x 128 Kbps x 2 (bi-directional) x 100% utilization = 256 Mbps

In this example, the main factor to consider is a headend router platform that can support at least 1000 tunnels. Although the worst case aggregation option calculated is 256 Mbps, the aggregate headend connection speed of DS3 (90 Mbps) is the upper constraining factor. Normal utilization is expected to be in the range of 64 Mbps.

In this case, even though a Cisco 7200VXR with NPE-G1 processor and SA-VAM2+ encryption accelerator can support the number of tunnels required, the encryption bandwidth required exceeds performance of the platform and produces a bottleneck at approximately 90–100 Mbps. Recommendations for this customer are to do one of the following:

- Divide the tunnels aggregated across multiple Cisco 7200VXRs
- Use a platform with higher encryption performance, such as the Cisco 7600 with VPN Shared Port Adapter (SPA)

A design based on the Cisco 7600 with a VPN SPA is recommended. For platform-specific results, see Headend Scalability Test Results—p2p GRE over IPsec, page 4-3.

Consider the first option. The "building block" of 500 branches to a pair of Cisco 7200VXRs is duplicated to support the total number of branch offices. In the customer example above, this requires two pairs of Cisco 7200VXRs for a total of 1000 branches. There is no limit to the number of branches that can be aggregated with this approach. Figure 3-2 shows this solution.

Figure 3-2 Stacked Cisco 7200VXR Design

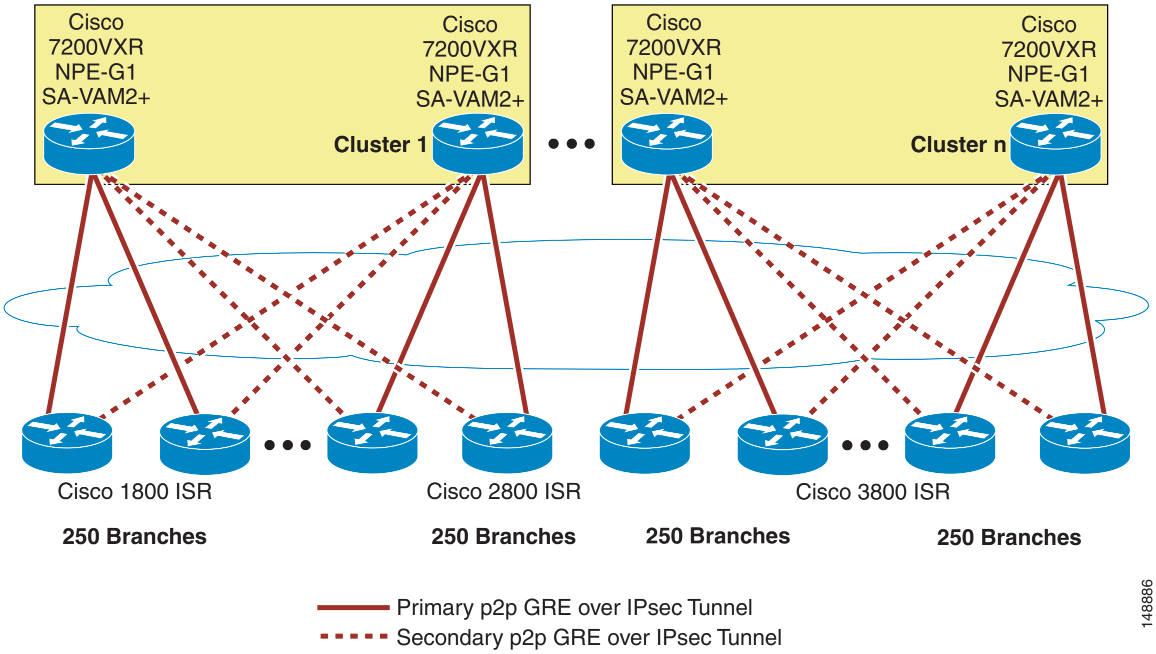
Cisco 7200VXR NPE-G1 SA-VAM2+
Cisco 7200VXR NPE-G1 SA-VAM2+
Cluster 1
Cisco 7200VXR NPE-G1 SA-VAM2+
Cisco 7200VXR NPE-G1 SA-VAM2+
Cluster n
Cisco 1800 ISR
Cisco 2800 ISR
Cisco 3800 ISR
250 Branches
250 Branches
250 Branches
250 Branches
Primary p2p GRE over IPsec Tunnel
Secondary p2p GRE over IPsec Tunnel
148886

Another way to address these same customer requirements is to position a larger aggregation platform in the design. The design alternative is to implement a pair of Cisco 7600 routers, each with Sup720 and VPN SPA. Each router can support 1000 GRE tunnels. (See Figure 3-3.)

Figure 3-3 Cisco 7600-Based p2p GRE over IPsec VPN Design

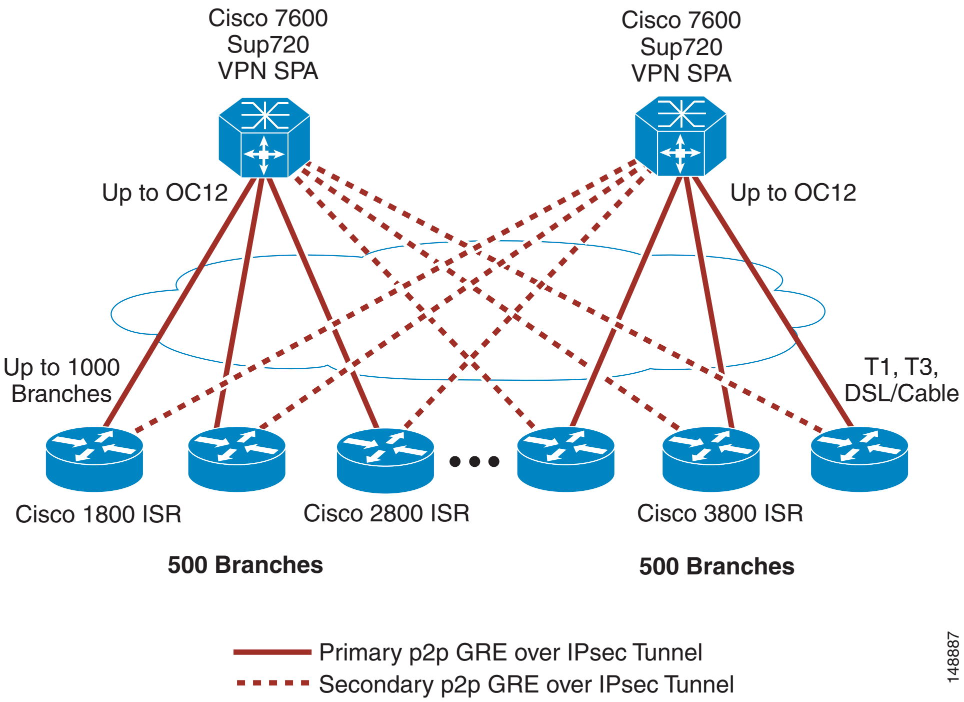

This design is an example of Load Sharing with Failover Headend Resiliency Design, page 2-21. Each branch office has a primary p2p GRE over IPsec tunnel to one Cisco 7600, and an alternate tunnel to a second Cisco 7600. Figure 3-3 can support up to 1000 branch offices, with each Cisco 7600 having 500 primary and 500 secondary neighbors. Routing neighbors tends to be the limiting factor. On the Cisco 7200VXR platform, routing neighbors tend to be limited to 500–700. On the Cisco 7600 with Sup720, up to 1000 routing neighbors have been proven to work in the Cisco scalability test lab.

The VPN SPA can hardware accelerate up to 8000 IPsec tunnels, but only 2047 p2p GRE tunnels. Encryption throughput is not a concern, because the VPN SPA can support line rates up to OC12.

Another design consideration is to determine where to terminate the p2p GRE tunnels. IPsec is terminated on the VPN SPA. The p2p GRE tunnels can either be terminated on the VPN SPA or on the Sup720. Both options were tested in the Cisco scalability lab, with the Sup720 option providing approximately 10–20 percent additional pps performance for the overall design. The design chosen depends on other factors, such as what other functions the Sup720 may be performing.

Customer Example with 1000–5000 Branches

There are cases where customers require aggregation of several thousand branch offices over a VPN. Consider the case of 5000 branch offices, each with a tunnel to a primary and alternate hub location, such as a point-of-sale terminal model. (See Table 3-5.)

Table 3-5 Customer Requirements

Customer Requirement	Value
Number of branch offices	5000
Branch access speeds	128 Kbps/1 Mbps DSL
Headend access speed	OC12 (1.24 Gbps bi-directional)
Expected utilization	25%

The calculation of aggregate bandwidth requirements is as follows:

- Typical case—5000 x (128 Kbps + 1 Mbps) x 22% utilization = 1.24 Gbps
- Worst case—5000 x (128 Kbps + 1 Mbps) x 100% utilization = 5.64 Gbps

One design decision that needs to be made is where to terminate the p2p GRE tunnels. IPsec is terminated on the VPN SPA. The p2p GRE tunnels can either be terminated on the VPN SPA or on the Sup720. Both options were tested in the Cisco scalability lab, with the Sup720 option providing approximately 10–20 percent additional pps performance for the overall design. The design chosen depends on other factors, such as what other functions the Sup720 may be performing.

Although the worst case aggregation option calculated is 5.64 Gbps, the total headend connection speed of OC12 (1.24 Gbps bi-directional bandwidth) is the upper constraining factor; thus, oversubscribing the OC12 circuit. Normal utilization is expected to be in the range of 1.24 Gbps.

Currently, no Cisco platform can aggregate 5000 p2p GRE over IPsec tunnels on a single platform. Factors such as the routing protocol limitations cause scalability issues. 5000 IGP RP neighbors do not scale on a single platform, which contradicts the widely known accepted practices, regardless of VPN technologies. Options for such large designs include the following:

- Duplicating a smaller-scale design, such as either the Cisco 7200VXR-based design for 500 branches, or the Cisco 7600-based design for 1000 branches.
- Implementing a Dual Tier Headend Architecture, using the Cisco 7200VXR platform to terminate the p2p GRE tunnels, and the Cisco 7600 platform for high-capacity IPSec encryption.

The Dual Tier Headend Architecture is shown in Figure 3-4.

Figure 3-4 Dual Tier Headend Architecture with p2p GRE over IPsec

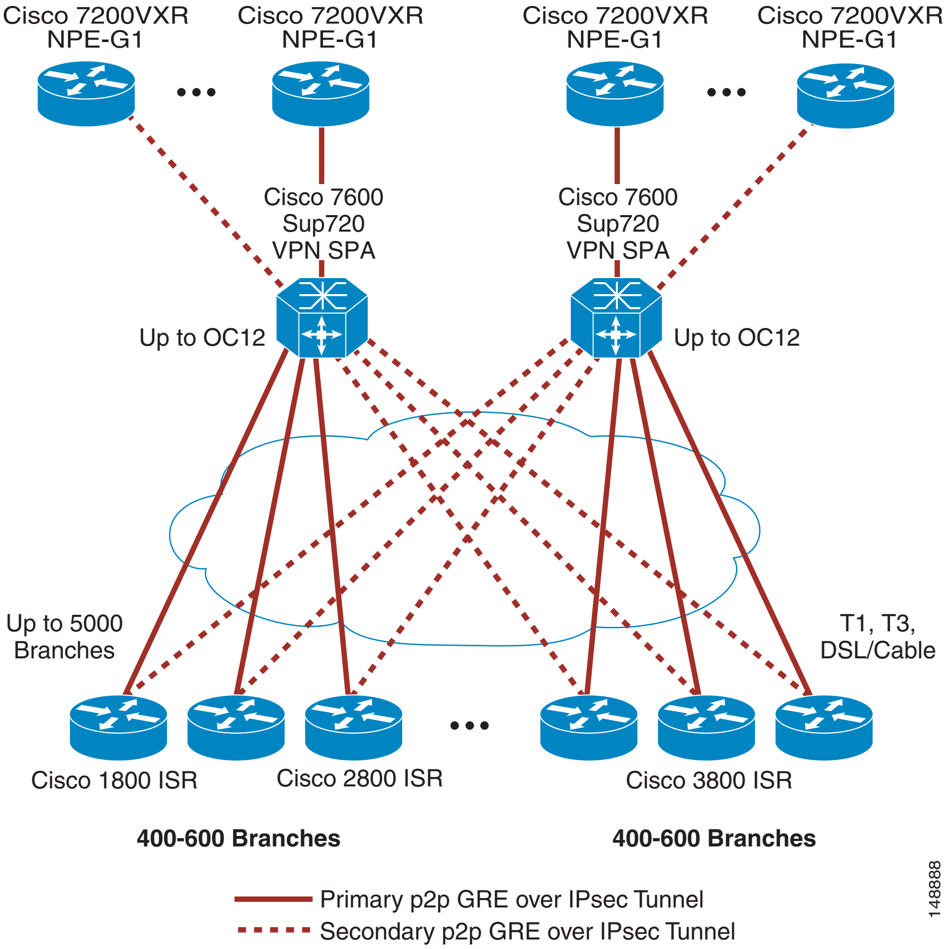

The Cisco 7200VXR platforms terminate the p2p GRE tunnels. Because there are no IPsec encryption requirements in this "tier" of the design, no SA-VAM2+ is required, and also these platforms can typically handle more spokes than if the router was performing both p2p GRE and IPsec.

In the other "tier" of the design, the Cisco 7600 with Sup720 and VPN SPA performs IPsec encryption services, which enables a single Cisco 7600, providing up to OC12 encryption speed, to perform as the IPsec tunnel aggregation point for up to 5000 tunnels.

A very important limitation of this design approach is that IP multicast limits the total number of tunnels that can be terminated through the VPN SPA. For designs requiring IP multicast, 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.

Branch Office Scalability

The branch routers are primarily responsible for the following:

- Terminating p2p GRE over IPsec tunnels from the headend routers

- Running a routing protocol inside of the p2p GRE tunnels to advertise internal routes

The most important factors to consider when choosing a product for the branch office include the following:

- Branch access speed, and expected traffic throughput to the headend (Fractional T1, T1, T3, broadband cable/DSL, other?)
- What other services is the branch router providing (for example, DHCP, NAT/PAT, VoIP, Cisco IOS firewall, IOS-IPS, and so forth)?

The pps rate (traffic size and traffic mix) is the largest single factor to branch router scalability.

The number of p2p GRE over IPsec tunnels does not play a large role in the branch sizing, because each branch router must be able to terminate a single tunnel for this design topology.

A primary concern is the amount of traffic throughput (pps and bps) along with the corresponding CPU utilization. Cisco recommends that branch routers be chosen so that CPU utilization does not exceed 65 percent under normal operational conditions. The branch router must have sufficient CPU cycles to service periodic events that require processing. Examples include ISAKMP and IPsec SA establishing and re-keying, SNMP, and Syslog activities, as well as local CLI exec processing. The average CPU busy on the branch router can be higher than the crypto headend because the headend is responsible for termination of all branches, not only the requirements of the branch being serviced by the remote router.

After initial deployment and testing, it may be possible to run branch routers at CPU utilization levels higher than 65 percent under normal operational conditions. However, this design guide conservatively recommends staying at or below 65 percent.

The Cisco Integrated Services Router (ISR) 1840, 2800, and 3800 Series of products have higher CPU performance than the products they replace. The ISR has an encryption module on the motherboard, or can be upgraded to an AIM series of encryption module for increased crypto performance.