



# IPsec Direct Encapsulation VPN Design Guide

---

This design guide provides guidelines and best practices for customer deployments of IP Security (IPsec) direct encapsulation VPNs. It is assumed that the reader has a basic understanding of IPsec.

## Contents

Introduction	3
Design Overview	4
Design Components	5
Best Practices and Known Limitations	6
Best Practices Summary	6
Known Limitations Summary	7
Design and Implementation	8
IPsec Direct Encapsulation Deployment	8
Dead Peer Detection	10
Reverse Route Injection	10
Dynamic Crypto Maps	10
Tunnel Initiation	11
VPN High Availability	11
Configuration and Implementation	12
ISAKMP Policy Configuration	12
Dead Peer Detection	13
Reverse Route Injection	14
Static Route Redistribution	14
VPN High Availability (IPsec Failover)	15
HA Design Example	15
Hot Standby Router Protocol	16



---

**Corporate Headquarters:**  
**Cisco Systems, Inc., 170 West Tasman Drive, San Jose, CA 95134-1706 USA**

Copyright © 2006 Cisco Systems, Inc. All rights reserved.

- Stateless Failover without HSRP 16
- Stateful Failover 17
- Stateless Failover with HSRP Configuration 17
- Quality of Service 18
- IP Multicast 19
- Interactions with Other Networking Functions 19
  - Network Address Translation and Port Address Translation 19
  - Dynamic Host Configuration Protocol 19
  - Firewall Considerations 19
- Common Configuration Errors 21
  - Crypto Peer Address Matching Using PSK 21
  - Transform Set Matches 21
  - ISAKMP Policy Matching 21
- Scalability Considerations 21
  - General Scalability Considerations 22
    - IPsec Encryption Throughput 22
    - Packets Per Second—The Most Important Factor 22
    - Tunnel Quantity Affects Throughput 23
  - Headend Scalability 23
    - Sizing the Headend 23
    - Tunnel Aggregation Scalability 24
    - Aggregation Scalability 24
    - Customer Requirement Aggregation Scalability Case Studies 24
  - Branch Office Scalability 26
- Scalability Test Results (Unicast Only) 27
  - Scalability Test Methodology 27
    - Overview 27
    - Headend Scalability Test Results 29
    - Branch Office Scalability Test Results 30
    - Scalability Test Results (AES Compared to 3DES) 30
    - Failover and Convergence Testing 31
  - Software Releases Evaluated 32
- Scalability Test Bed Configuration Files 33
  - Cisco 7200VXR Headend Configuration 33
    - Cisco 7200VXR Headend Configuration 33
    - Cisco 7600 Headend Configuration 34
  - ISR Branch Configuration 36
- Appendix A—Scalability Test Results for Other Cisco Products 37
  - Cisco Headend VPN Routers (Legacy) 37

Other Cisco Products for the Headend	37
Cisco Branch Office VPN Routers (Legacy)	38
Appendix B—References	38
Appendix C—Acronyms and Definitions	39

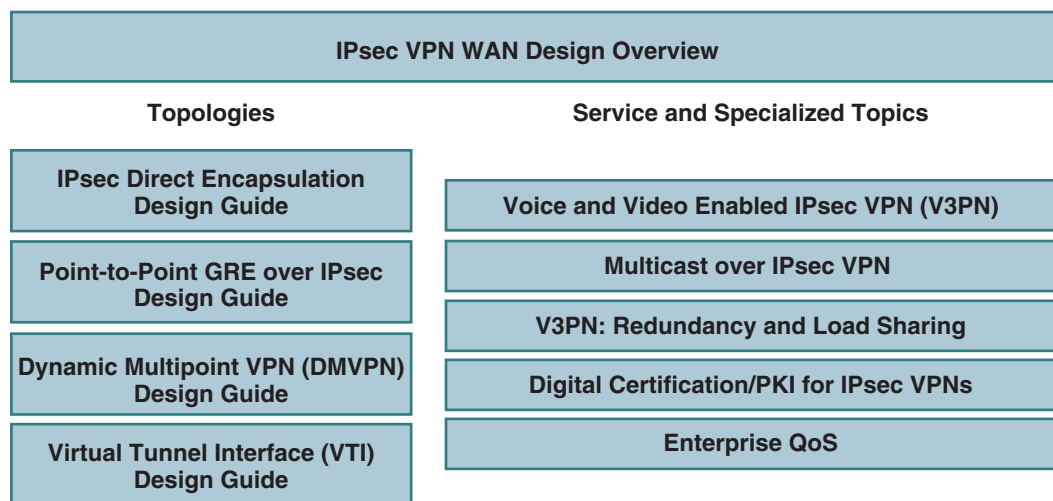
## Introduction

This design guide evaluates Cisco VPN product performance in scalable and resilient site-to-site VPN topologies, using Cisco VPN routers running Cisco IOS Software, with IPsec as the tunneling method. The concepts presented can also be applied to other Cisco products that do not run Cisco IOS software.

This design guide begins with an overview, followed by design recommendations and product selection and performance information. Finally, partial configuration examples are presented.

The chart in [Figure 1](#) shows the IPsec VPN WAN architecture documentation, which is divided into multiple design guides based on the technologies used. Each technology uses IPsec as the underlying transport mechanism for the VPNs.

**Figure 1** IPsec VPN WAN Design Overview



190897


The operation of IPsec is outlined in the *IPsec VPN WAN Design Overview* (<http://www.cisco.com/go/srnd>), which also outlines the criteria for selecting a specific IPsec VPN WAN technology. This document helps you to select the correct technology for the proposed network design. [Design and Implementation, page 8](#) provides more detail on the design considerations. [Scalability Considerations, page 21](#) presents Cisco product options for deploying the design.

This document addresses the following applications and implementations of IPsec direct encapsulation VPNs:

- Dead Peer Detection (DPD)
- Reverse Route Injection (RRI)
- VPN high availability using Hot Standby Router Protocol (HSRP) with stateless and stateful failover
- Data and VoIP converged traffic requirements
- Quality of service (QoS) features

The primary topology discussed in this document is a hub-and-spoke model. In this deployment, primary enterprise resources are located in a large central site, with a number of smaller sites or branch offices connected directly to the central site over a VPN. A high-level diagram of this topology is shown in [Figure 2](#).

**Figure 2** Hub-and-Spoke VPN



## Design Overview

This guide makes the following design assumptions and recommendations:

- The design supports a typical converged traffic profile for customers. See the [Scalability Considerations, page 21](#) for details about the traffic profile used during scalability testing.
- Built-in redundancy and failover with fast convergence are essential to help ensure high availability and resiliency. This is discussed further in [Design and Implementation, page 8](#).
- This design uses IPsec alone as the tunneling method, which is appropriate for enterprises that do not require an IGP routing protocol passing through the tunnel, IP multicast (IPmc) traffic, or multiprotocol traffic.

- Cisco devices should be maintained at reasonable CPU utilization levels. [Scalability Considerations, page 21](#) discusses this issue in detail, including recommendations for headend and branch devices and for software versions.
- The design recommendations assume that the customer deploys current VPN technologies, including hardware-accelerated encryption. Cost considerations have been taken into account in the proposed design, but not at the expense of necessary performance.
- Support for voice over IP (VoIP) and video are assumed to be requirements in the network design. Detailed design considerations for handling VoIP and other latency-sensitive traffic is not explicitly addressed in this design guide, but may be found in the *Voice and Video Enabled IPsec VPN (V3PN) Design Guide*, available at the following URL: <http://www.cisco.com/go/srnd>
- Recommendations are for enterprise-owned VPNs. However, the concepts and conclusions are valid regardless of the ownership of the edge tunneling equipment, so the recommendations are also useful for VPNs managed by service providers.

## Design Components

VPNs have the same requirements as traditional private WAN services, including multiprotocol support, high availability, scalability, and security. VPNs can often meet these requirements more cost-effectively and with greater flexibility than private WAN services.

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

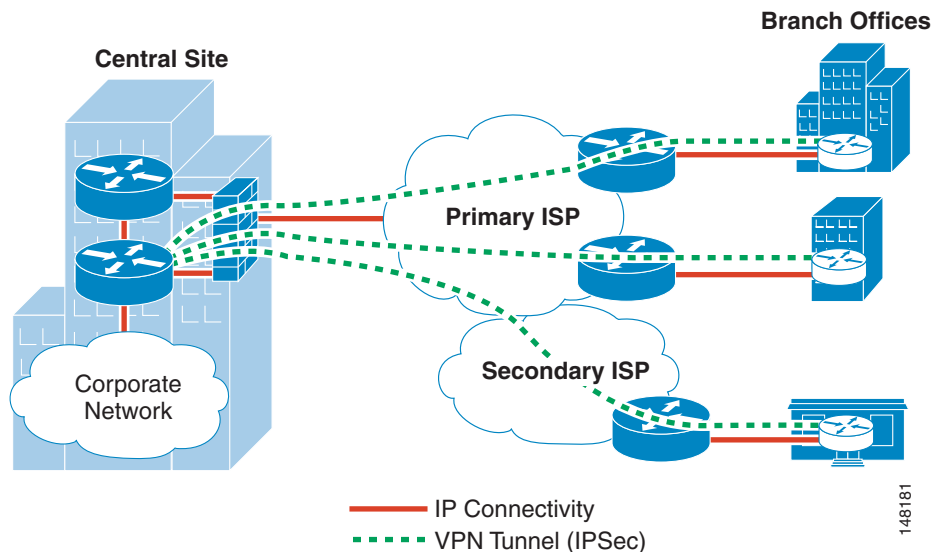
- Cisco high-end VPN routers serve as VPN headend termination devices at a central campus site.
- Cisco VPN access routers serve as VPN branch termination devices at branch office locations.
- IPsec direct encapsulation (with DPD, RRI, and HSRP) provides headend-to-branch interconnections.
- Internet services from a third-party ISP (or ISPs) provide 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 latency-sensitive traffic and resiliency. [Design and Implementation, page 8](#) describes how to select headend and branch devices.

The network topology of the hub-and-spoke design is shown in [Figure 3](#). The solution is a hub-and-spoke network with multiple headend devices for redundancy. Headends are high-end tunnel aggregation routers that service multiple IPsec tunnels for a prescribed number of branch office locations. In addition to terminating the VPN tunnels at the central site, headends can advertise routes to branch devices using RRI.

To ensure authentication and encryption, IPsec tunnels are provisioned to interconnect branch offices to the central site. The way that network resiliency is provided depends on the initial network requirements.

**Figure 3** VPN Hub-and-Spoke Network Topology



## 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 [Design and Implementation, page 8](#).

### Best Practices Summary

This section summarizes at a high level the best practices for an IPsec direct encapsulation VPN deployment.

#### General Best Practices

The following are general best practices:

- Use IPsec in tunnel mode for best performance.
- Configure Triple DES (3DES) or AES for encryption of transported data (exports of encryption algorithms to certain countries may be prohibited by law).
- Implement DPD to detect loss of communication between peers.
- Deploy hardware-acceleration for 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.
- Set up QoS service policies, as appropriate, on headend and branch router interfaces to help ensure performance of latency-sensitive applications. For more information, see the *Voice and Video Enabled IPsec VPN (V3PN) Design Guide* at the following URL: <http://www.cisco.com/go/srmd>.

- The QoS pre-classify feature is helpful in VPN designs where both QoS and IPsec occur on the same system. Alternatively, DSCP values in the ToS byte can be marked on the unencrypted packet at ingress and then matched on the encrypted packet on egress by the service policy.

## Headend Best Practices

The following are best practices for the headend device:

- Use RRI on headend routers for optimal routing between campus and remote sites.
- Configure dynamic crypto maps on headend routers to simplify configuration and provide touchless provisioning of new branches.
- If high-availability is a requirement, implement a design with redundancy for both headend equipment and WAN circuits.
- Select Cisco VPN router products at the headend based on considerations for the following:
  - Number of tunnels to be aggregated
  - Maximum throughput in terms of both pps and bps to be aggregated
  - Performance margin for resiliency and failover scenarios
  - Maintaining CPU utilization below design target

See [Headend Scalability, page 23](#) for more information.

## Branch Office Best Practices

The following are best practices for the branch office devices:

- Configure multiple crypto peers to provide headend redundancy
- 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 (for example, firewall, IPS, and NAT/PAT)
  - Maintaining CPU utilization below 65–80 percent

See [Branch Office Scalability, page 26](#) for more information.

## Known Limitations Summary

This section summarizes the known limitations for an IPsec direct encapsulation deployment.

### General Limitations

The following are general limitations for the recommended IPsec direct encapsulation design:

- Dynamic IGP routing protocols (for example, EIGRP and OSPF) are not supported, because dynamic routing protocols require IPmc support for forwarding hellos.
- IPmc traffic is not supported.
- Non-IP protocols, such as IPX or AppleTalk, are not supported.
- The network manager must verify the QoS service policies are matching packets as intended.
- IPsec direct encapsulation designs can be implemented only in a Single Tier Headend Architecture.

## Headend Limitations

The following are headend limitations for the recommended IPsec direct encapsulation design:

- Two versions of Stateful Failover (VPN High Availability) exist today, depending on the platform:
  - Cisco 7200VXR and ISR—Stateful Switchover (SSO)
  - Cisco Catalyst 6500 or 7600—State Synchronization Protocol (SSP)
- Eventually, all Cisco headend platforms will move to the SSO failover functionality.
- Digital certificates/PKI have not been verified with either SSO or SSP.
- QoS can be implemented only in a limited way in the headend-to-branch direction because it is not possible to configure a service policy at the tunnel/destination level.

## Branch Office Limitations

The following are branch office limitations for the recommended IPsec direct encapsulation design:

- The IPsec tunnel must be initiated by the remote branch in cases where remote routers acquire their address with a dynamically served IP address. The crypto headend cannot initiate the tunnel to the branch. As a result, interesting traffic must be present (for example, Cisco IP SLA) to keep the IPsec SA alive.
- There is no automatic failback when multiple crypto peers are configured. The IPsec Preferred Peer feature provides a limited means to influence the order in which multiple peers on a crypto map are tried
- 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 information about these recommendations is provided later in this document.

# Design and Implementation

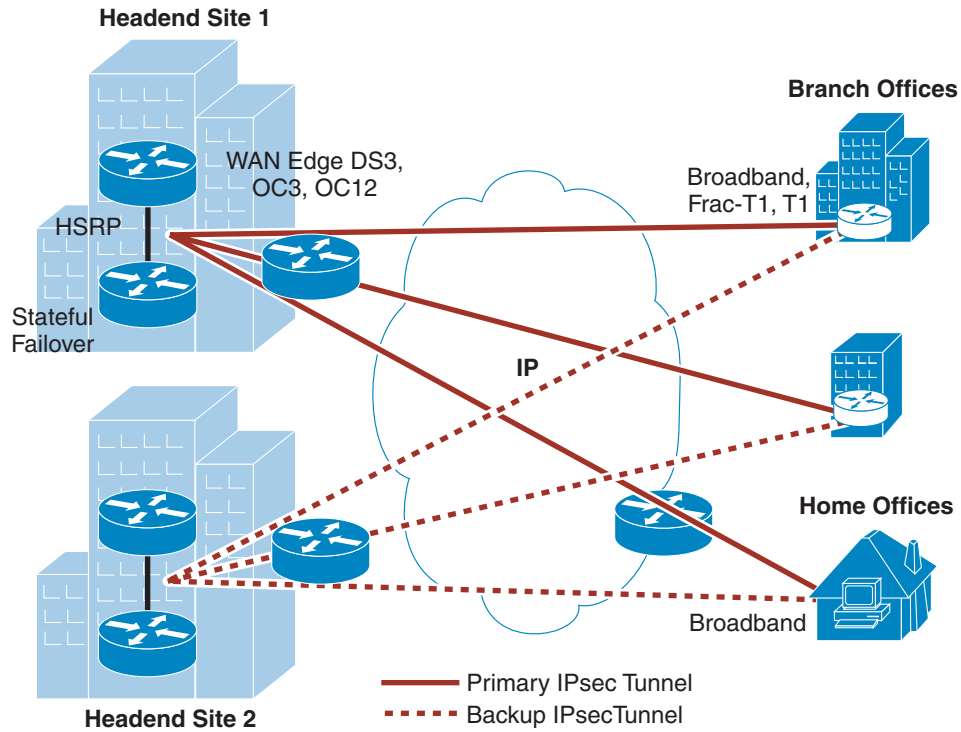
This section describes the recommended IPsec direct encapsulation deployment and discusses specific implementation issues.

## IPsec Direct Encapsulation Deployment

Figure 4 shows a typical IPsec direct encapsulation deployment.



**Figure 4 IPsec Direct Encapsulation Deployment**



Headend			Branch		
Routing Control Plane	Route Redistr.	RRI	Static Route		
IPsec Control Plane	Dynamic Crypto Map	DPD	Static Crypto Map	Peer List	DPD

148182

Headend sites are typically connected with DS3, OC3, or even OC12 bandwidth. Branch offices are typically connected by fractional T1, T1, T3, or fractional T3, and increasingly by broadband DSL or cable. Two possibilities are available for providing redundancy:

- Box-to-box redundancy with HSRP and Stateful Failover (VPN High Availability)
- Site-to-site stateless redundancy with geographically separated headend sites.

Typically, branch routers are configured with a list of possible headend crypto peers that are tried in succession until a tunnel is successfully established.

The IPsec control plane normally uses dynamic crypto maps at the headend to minimize configuration changes when new branches are added. Dynamic crypto maps are also used to support branches with a dynamic Internet addresses as their crypto peer. DPD automatically detects ISAKMP peer loss and tears down the IPsec SA (data tunnel) if the connection is lost completely.

The routing control plane generally uses static routes at the branch locations, with RRI at the headends to inject routes into the routing table for advertisement. IGP dynamic routing protocols are not exchanged over the VPN tunnel between headend and remote sites.

A routing protocol provides several vital features when deployed over a network. These include peer state detection, optimal routing, and the ability to facilitate alternate routes in the event of a link failure. IPsec VPNs implement this functionality without a routing protocol using DPD and RRI. The combined use of DPD and RRI is less network intensive than an actual routing protocol running over the VPN, but achieves a similar effect.

## Dead Peer Detection

Dead Peer Detection (DPD) is a relatively new Cisco IOS software feature that is an enhancement to the ISAKMP keepalives feature. DPD sends a hello message to a crypto peer from which it has not received traffic during a configurable period. If normal IPsec traffic is received from a crypto peer and decrypted correctly, the crypto peer is assumed alive, no hello message is sent, and the DPD counter for that crypto peer is reset. This produces lower CPU utilization than using ISAKMP keepalives.

If no traffic is received during the specified period, an ISAKMP R\_U\_THERE message is sent to the other crypto peer. If no response is received after the specified number of tries, the connection is assumed dead, and the IPsec tunnel is disconnected. This feature is vital to prevent blackholing traffic, in case the SA database on one peer is cleared manually or by rebooting the device. DPD is both a headend and branch technology and should be configured on both sides of each VPN tunnel.

## Reverse Route Injection

Another IPsec feature that has been added recently to Cisco IOS Software is Reverse Route Injection (RRI). RRI takes the information derived from the negotiated IPsec SAs and creates a static route to the networks identified in those SAs. Route redistribution then occurs between these static routes and whatever routing protocol is configured on the headend router. This makes the routes to the branch office networks available to networks behind the headend aggregation routers.

RRI is a headend technology that allows static routes to be automatically generated in the headend router IP routing table. These static routes are then redistributed using a routing protocol into the enterprise network. DPD works in conjunction with RRI. In the event that DPD detects the loss of a crypto peer connection (after the specified ISAKMP R\_U\_THERE retries have expired), DPD triggers the IPsec tunnel to be torn down. This causes RRI to remove the associated static route from the route table.

## Dynamic Crypto Maps

Dynamic crypto maps eliminate the need to statically predefine every crypto peer. Dynamic crypto maps allow an IPsec connection between two crypto peers when one of the crypto peers (usually the central site crypto peer) does not have the complete configuration necessary to complete the IPsec negotiation.

Dynamic crypto maps are required when the remote crypto peer has a dynamically assigned IP address, such as over a cable or ADSL connection. In this case, the remote peer cannot be preconfigured into the central site device because its IP address is unknown. The IKE authentication completes based on verification of identity through a pre-shared secret key or digital certificate. Information from the IPsec session is used to complete the current IP address of the remote branch router in the dynamic crypto map configuration on the headend.

## Tunnel Initiation

When dynamic crypto maps are used on the headend, the IPsec connection can be initiated only by the branch router. However, because the headend device uses dynamic crypto maps, it does not have all the information necessary to create an IPsec SA by itself. This is of concern when traffic forwarding is required from a central site to a remote site without the remote site initiating the connection.

Because an IPsec tunnel exists only when interesting traffic is transmitted, the tunnel may not be up when traffic arrives on the headend destined for the branch router. One way to work around this issue is to create a periodic traffic source that always keeps the tunnel active. Examples of this type of periodic traffic source include the following:

- Cisco IP SLA, formerly known as Service Assurance Agent (SAA)—This can be configured to send periodic probes
- Network Time Protocol (NTP)—Periodically polls the configured NTP servers
- Cisco CallManager—IP phones behind the branch router send periodic keepalive packets to a central primary and secondary Cisco CallManager

Of these options, the Cisco IP SLA feature offers the most robust capabilities.

When the headend must initiate the IPsec tunnel, static crypto maps must be used. After the IPsec SA is established, data traffic can flow in either direction, regardless of which side initiated the tunnel.

## VPN High Availability

Customer requirements determine the type of high availability required for the IPsec VPN design. The following failover topologies are discussed in this document:

- Stateless failover without HSRP
- Stateless failover with HSRP
- Stateful Failover using SSO (on Cisco 7200VXR and ISR platforms)
- Stateful Failover using SSP (on Cisco Catalyst 6500 or 7600 platforms)

Stateless failover (with or without HSRP) is an option when there is a primary and one or more secondary headend sites to which the remote site can establish a connection.

When there is no HSRP between the headends at the different geographic sites, if a connection cannot be achieved to the primary headend crypto peer, the remote site retries the next headend in the crypto peer list. In the case of stateless failover with HSRP, there is an HSRP virtual IP address that provides a single crypto peer for the branch router.

Stateful Failover using SSO or SSP is an option when two headend routers run HSRP and exchange IPsec state information. The remote points to a single HSRP address in its crypto peer list. If the active headend fails, the standby headend resumes the same IPsec tunnels to the branch locations, typically within one to three seconds.

Stateful HA failover can be used in a location with stateless failover without HSRP at the same time. This provides the highest level of availability with both box and site redundancy.

# Configuration and Implementation

This section describes how to configure and implement an IPsec direct encapsulation design.

## ISAKMP Policy Configuration

There must be at least one matching ISAKMP policy between any pair of crypto peers. The example configuration below shows a policy using pre-shared keys (PSK) with 3DES as the encryption algorithm. The default ISAKMP policy, which has the lowest priority, contains the default values for the encryption algorithm, hash method (HMAC), Diffie-Hellman group, authentication type, and ISAKMP SA lifetime parameters. This is the lowest priority ISAKMP policy.

The ISAKMP configuration must consider the tunnel authentication key method that will be chosen. The two most common options are pre-shared keys (PSK) and digital certificates. The use of digital certificates is more manageable and more secure than the use of pre-shared keys. More information about digital certificates is available in the *Digital Certification/PKI for IPsec VPN Design Guide* at the following URL: <http://www.cisco.com/go/srnd>.

If PSK is chosen, one pre-shared key must be assigned per remote crypto peer. Each pre-shared key is configured on a line by itself. An alternative to configuring pre-shared keys on the headend configuration is the use of IKE aggressive mode. This mode uses a RADIUS server to store the keys. IKE aggressive mode transmits the pre-shared key as a hashed but unencrypted string. If these packets are captured by a third party with a protocol analyzer, a dictionary attack can be executed to recover the hashed value. IKE main mode encrypts the hashed pre-shared key. This document focuses only on IKE main mode. In the following example, only one crypto peer with a single PSK is shown. This would be used with a static map on both the headend and branches.

- Headend #1 (stateless with HSRP):

```
interface GigabitEthernet0/1
ip address 192.168.251.2 255.255.255.0
standby ip 192.168.251.1
crypto isakmp policy 1
  encr 3des
  authentication pre-share
  group 2
crypto isakmp key bigsecret address 192.168.161.2
crypto is
akmp keepalive 10
```

- Headend #2 (stateless with HSRP):

```
interface GigabitEthernet0/1
ip address 192.168.251.3 255.255.255.0
standby ip 192.168.251.1
crypto isakmp policy 1
  encr 3des
  authentication pre-share
  group 2
crypto isakmp key bigsecret address 192.168.161.2
crypto isakmp keepalive 10
```

- Branch #1:

```
interface Serial0/0
ip address 192.168.161.2 255.255.255.0
!
crypto isakmp policy 1
  encr 3des
```

```

authentication pre-share
group 2
crypto isakmp key bigsecret address 192.168.251.1
crypto isakmp keepalive 10

```

## Dead Peer Detection

Dead Peer Detection is an enhancement to the ISAKMP Keepalive feature. The DPD **on-demand** option no longer automatically sends hello messages to the crypto peer if traffic has been received from that crypto peer within a “worry” interval. This option is triggered only if a packet is to be transmitted to that remote crypto peer. The **periodic** option sends ISAKMP keepalives to the crypto peer periodically, regardless of network traffic.

The first variable for the **ISAKMP keepalive** command is the number of seconds that the crypto peer waits for valid traffic from its IPsec neighbor. DPD **on demand** is the router default and is shown in the following configuration. This scheme helps conserve router CPU by not sending keepalive messages if a router has just received valid traffic.

- **Headend #1:**

```

interface GigabitEthernet0/1
ip address 192.168.251.2 255.255.255.0
standby ip 192.168.251.1
!
crypto isakmp policy 1
  encr 3des
  authentication pre-share
  group 2
crypto isakmp key bigsecret address 192.168.161.2
crypto isakmp keepalive 10

```

- **Headend #2:**

```

interface GigabitEthernet0/1
ip address 192.168.251.3 255.255.255.0
standby ip 192.168.251.1
!
crypto isakmp policy 1
  encr 3des
  authentication pre-share
  group 2
crypto isakmp key bigsecret address 192.168.161.2
crypto isakmp keepalive 10

```

- **Branch #1:**

```

interface Serial10/0
ip address 192.168.161.2 255.255.255.0
!
crypto isakmp policy 1
  encr 3des
  authentication pre-share
  group 2
crypto isakmp key bigsecret address 192.168.251.1
crypto isakmp keepalive 10

```

## Reverse Route Injection

RRI injects a static route into the routing table of the headend router for the network address referenced by the crypto ACL of the remote router. These static routes can be redistributed using a dynamic routing protocol.

RRI is implemented by the single command **reverse-route** under the crypto map of an IPsec configuration. RRI can be configured on a router with either a static or a dynamic crypto map. The static IP route is only present if that IPsec SA is active.

- Headend #1:

```
crypto ipsec transform-set vpn-test esp-3des esp-sha-hmac
!
crypto dynamic-map dmap 10
  set transform-set vpn-test
  reverse-route
!
!
crypto map dynamic-map local-address GigabitEthernet0/1
crypto map dynamic-map 10 ipsec-isakmp dynamic dmap
```

- Headend #2:

```
crypto ipsec transform-set vpn-test esp-3des esp-sha-hmac
!
crypto dynamic-map dmap 10
  set transform-set vpn-test
  reverse-route
!
!
crypto map dynamic-map local-address GigabitEthernet0/1
crypto map dynamic-map 10 ipsec-isakmp dynamic dmap
```

- Branch #1:

```
crypto ipsec transform-set vpn-test esp-3des esp-sha-hmac
!
crypto map static-map local-address Serial10/0
crypto map static-map 10 ipsec-isakmp
  set peer 192.168.251.1
  set transform-set vpn-test
  match address b000
!
ip access-list extended b000
  permit ip 10.60.0.0 0.0.0.255 10.0.0.0 0.255.255.255
```

RRI is not configured on the branch devices. The branch routers use a static default pointing to the upstream next hop.

## Static Route Redistribution

The redistribution of static routes inserted by RRI takes place using the normal route redistribution mechanisms already present in Cisco IOS software. Because of the IP routing “redistribution scan timer,” a change in the RRI static route may take up to a minute before being reflected in the distributed routing protocol. The following are examples of the ISAKMP headend configurations.

- Headend #1:

```
router eigrp 1
  redistribute static metric 1000 100 255 1 1500
```

```
network 10.0.0.0
no auto-summary
!
```

- Headend #2:

```
router eigrp 1
 redistribute static metric 1000 100 255 1 1500
 network 10.0.0.0
no auto-summary
!
```

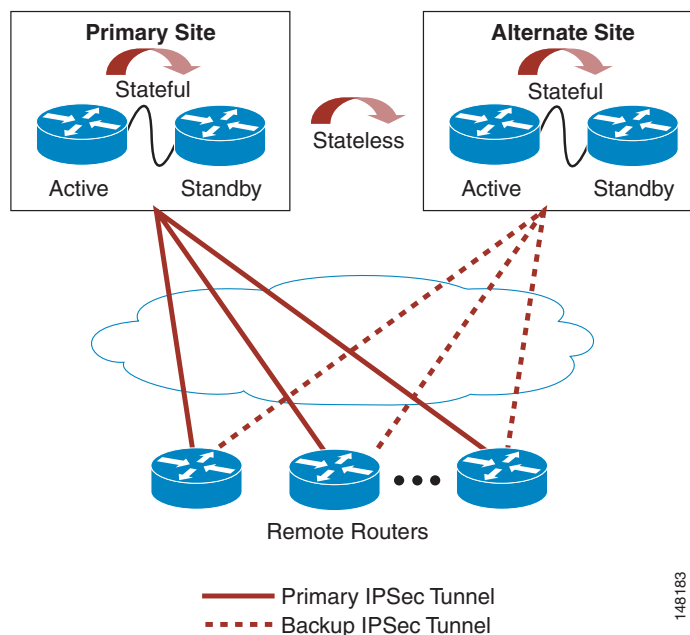
## VPN High Availability (IPsec Failover)

Network performance in the event of a failure is a primary concern for an IPsec VPN deployment. This section provides some recommendations for highly available IPsec VPNs. This design is covered in depth in the *High Availability for IPsec VPN Design Guide*.

### HA Design Example

High availability may be a customer requirement, but each customer is willing to invest for different levels of HA. Figure 5 is for a customer who has both a primary and a backup headend location that are geographically separated.

**Figure 5 High Availability VPN Design**



Within a site, redundant headend routers can be configured to run HSRP and also share IPsec state information (using SSP or SSO, depending on the platform). The alternate site can also be set up this way, if required.

Branch sites are configured with two (or more) crypto peer statements, with the primary site appearing first, followed by the alternate sites. Both crypto peer statements point to the appropriate HSRP address.

With this design, if a headend failure occurs at the primary site, HSRP triggers Stateful Failover to the standby headend router, typically within one to three seconds. The IPsec tunnel does not have to be re-established because the IPsec SA database is copied and active on the standby box. The branch router is unaware that the IPsec SA is now being serviced by the backup headend.

If the primary site fails, the branch routers are unable to receive traffic from the headend. After the configured DPD period, DPD begins sending ISAKMP R\_U\_THERE messages using an ISAKMP SA to the primary site headend. When no response is received for any of the retries, DPD declares the IPsec tunnel dead, removes the IPsec Security Associations (IPsec SAs), tears down the tunnel, and removes the corresponding RRI static route.

The branch router then tries to establish a new connection to the next headend crypto peer in the branch route crypto peer list. In this case, it is the alternate site headend HSRP address. Successful connection results in a new IPsec tunnel and traffic path. This process typically takes 30–45 seconds, depending on how aggressively DPD is configured, and depending on the routing protocol convergence in the core enterprise network. The redistributed route is removed or added from the IP routing table and the IGP neighbors are notified immediately.

The number of tunnels required for each headend device should be scaled to the overall size of the network. In addition, the normal load from a number of branch sites may be distributed across two or more headend devices, if stateless failover is used. To distribute the load, configure multiple standby groups, one group for each group of branch devices. By using HSRP this way, remote sites may be evenly divided among a number of headend devices for load sharing during normal operation. During a failure event, only the branch devices connected as primary to the failed HSRP group owner are subject to re-negotiation of the IPsec SAs. This results in enhanced failover performance.

If Stateful Failover is configured, you cannot distribute branch sites across different headend devices. With Stateful Failover, one headend router is active and terminates all ISAKMP and IPsec SAs. The other is completely dedicated to hot standby operation.

## Hot Standby Router Protocol

The Hot Standby Router Protocol (HSRP) lets IPsec use the standby group addresses for the crypto peer address. If the current owner of the HSRP group fails, the virtual IP address transfers over to the secondary standby router. HSRP is used between the active and standby crypto headend in either stateless or stateful mode.

## Stateless Failover without HSRP

A branch site is configured with two or more crypto peer statements, with the primary headend crypto peer appearing first, followed by the alternate crypto peers.

You can use DPD and Cisco IOS Software keepalive with multiple crypto peers in the crypto map to allow for stateless failover. DPD allows the router to detect a dead ISAKMP peer, and when the router detects the dead state, the router deletes the associated IPsec and ISAKMP SAs. If you configure multiple crypto peers, the router switches over to the next listed crypto peer for stateless failover. To control this option, use the **set peer** command with the following syntax:

```
set peer {host-name [dynamic] [default] | ip-address [default] }
```



## Stateful Failover

Stateful Failover lets headend routers share information in the SA database. If HSRP detects a headend device failure, the remote branch router continues to use the same IPsec SA to the backup headend without needing to create a new set of IPsec SAs. This process greatly reduces failover time and the amount of re-keying required in the event of a single headend system failure.

With Stateful Failover, only one crypto headend is active at one time. HSRP determines the crypto headend that receives the packets. A static IP route is manually inserted into the core router one hop above the crypto headends, pointing to the HSRP virtual IP as the next hop for the branch subnets. This static route is redistributed into the routing protocol in the core.

Cisco has developed various versions of Stateful Failover in conjunction with various platforms. The feature was initially released to work with State Synchronization Protocol (SSP) on the platforms listed in Table 1. Further information about SSP is available at the following website:

[http://www.cisco.com/en/US/partner/products/sw/iosswrel/ps5012/products\\_feature\\_guide09186a0080116d4c.html](http://www.cisco.com/en/US/partner/products/sw/iosswrel/ps5012/products_feature_guide09186a0080116d4c.html)

Further information about SSO is available at the following website:

[http://www.cisco.com/en/US/partner/products/sw/iosswrel/ps5207/products\\_feature\\_guide09186a00802d03f2.html](http://www.cisco.com/en/US/partner/products/sw/iosswrel/ps5207/products_feature_guide09186a00802d03f2.html)

**Table 1 Platform and Software Support for HA Features**

Platform	Software Release	HA Software
Cisco 7200VXR with NPE-400	Cisco IOS Software release 12.2(11)YX to 12.2(11)YX1	SSP
Cisco Catalyst 6500	Cisco IOS release 12.2(14)SY or greater	SSP
Cisco 7600		Cisco VPNM or Cisco VPN SPA modules
Cisco 7200VXR with NPE-G1	Cisco IOS Software release 12.3(11)T	Stateful Switchover (SSO) with HSRP
Platforms with the PA-ISA, SA-VAM, SA-VAMII, SA-VAMII+ encryption modules		
Cisco 3845 ISR router		



**Note**

The two versions of Stateful Failover (SSP and SSO) cannot be used together in a single chassis.

## Stateless Failover with HSRP Configuration

The HSRP configuration used on an interface with a crypto map is identical to normal HSRP use (see Table 6). The standby commands operate as they do without an IPsec configuration. The only difference between an IPsec configuration without HSRP and a configuration *with* HSRP is removing the **local crypto peer address** command when implementing PSK. When a crypto map is applied to an interface with the **redundancy** keyword, the IP address assigned to the standby group is automatically used as the local crypto peer address without any requirement for a **local crypto peer** statement with PSK.

In stateless failover mode, no IPsec or ISAKMP SA state information is transferred to the backup system. A remote crypto peer router configured with an HSRP group address as a crypto peer must renegotiate its ISAKMP SAs and IPsec SAs before traffic transmission. Stateless operation is supported with all platforms and ISAKMP authentication types.

Headend #1 in this example has a standby priority of 101. The default value is 100. A higher priority value is preferred within the standby group. Multiple standby groups can be configured, with the one headend with a higher priority value for one group, but the lowest priority value for a second group. Half the remote routers can use one IP address in their set peer statement, corresponding to the HSRP address of the first group, and the remaining routers can use the HSRP address for the second group. With this configuration, only half the remote routers need to failover in the event of a headend crypto failure. If this method is used, both HSRP configurations use the HSRP preempt command.

At the branch site, there are no special configurations required, because HSRP is configured only between the crypto headends.

- Headend #1:

```
interface GigabitEthernet0/1
  description GigabitEthernet0/1
  ip address 192.168.251.2 255.255.255.0
  duplex auto
  speed auto
  media-type gbic
  negotiation auto
  standby ip 192.168.251.1
  standby timers msec 50 1
  standby priority 101
  standby name group1
  standby track GigabitEthernet0/2
  crypto map dynamic-map redundancy group1
```

- Headend #2:

```
interface GigabitEthernet0/1
  description GigabitEthernet0/1
  ip address 192.168.251.3 255.255.255.0
  duplex auto
  speed auto
  media-type gbic
  negotiation auto
  standby ip 192.168.251.1
  standby timers msec 50 1
  standby name group1
  standby track GigabitEthernet0/2
  crypto map dynamic-map redundancy group1
```

## Quality of Service

You may need to configure QoS to support latency-sensitive traffic applications. QoS and IPsec have been integrated as part of the Cisco Voice and Video Enabled IPsec VPN (V3PN) technology. For more information, see the *Voice and Video Enabled IPsec VPN (V3PN) Design Guide* at the following URL: <http://www.cisco.com/go/srnd>.

## IP Multicast

IPsec direct encapsulation does not support IPmc traffic. It is therefore necessary to implement either a p2p GRE over IPsec, DMVPN, or Virtual Tunnel Interface (VTI) design to support IPmc. For more information, see one of the following design guides (<http://www.cisco.com/go/srnd>):

- *Point-to-Point GRE over IPsec VPN Design Guide*
- *Dynamic Multipoint VPN (DMVPN) Design Guide*
- *IPsec Virtual Tunnel Interface (VTI) Design Guide*

## Interactions with Other Networking Functions

Other networking functions such as PAT, DHCP, and firewall considerations apply to designing an IPsec direct encapsulation network. This section describes some of these functions.

### Network Address Translation and Port Address Translation

Although Network Address Translation (NAT) and Port Address Translation (PAT) can provide an added layer of security and conserve public addresses, they both present challenges when implementing an IPsec VPN. Internet Key Management Protocol (ISAKMP) relies on an individual IP address per crypto peer for proper operation. However, PAT works by representing multiple crypto peers with a single IP address.

The IPsec NAT Traversal feature (NAT-T) lets IPsec traffic travel through NAT or PAT devices by encapsulating both the IPsec SA and the ISAKMP traffic in a User Datagram Protocol (UDP) wrapper. NAT-T was first introduced in Cisco IOS software version 12.2(13)T, and is auto-detected by VPN devices. There are no configurations steps for a Cisco IOS Software router running this release or later because it is enabled by default as a global command. NAT-T detects a PAT device between the crypto peers and negotiates NAT-T if it is present. For further information about IPsec NAT Transparency, see the following URL:

<http://www.cisco.com/univercd/cc/td/doc/product/software/ios122/122newft/122t/122t13/ftipsnat.htm>

### Dynamic Host Configuration Protocol

For a host at a remote site to use a DHCP server over an IPsec tunnel at a central site, an IP helper address must be configured on the router interface associated with the host.

One drawback is that if connectivity to the central site is lost, a host at a remote site may not receive or renew an IP address. If the host cannot receive an IP address, the host is unable to communicate to the local network. For this reason, Cisco recommends using the remote branch router as a standalone DHCP server for branch offices without redundant links.

### Firewall Considerations

The section describes various firewall considerations when deploying an IPsec direct encapsulation design.

## Branch Considerations

This section describes the considerations that apply to the branch routers.

### Firewall Feature Set and Inbound ACL

Before Cisco IOS version 12.3(8)T, packets received on an interface with an inbound ACL and a crypto map were checked by the inbound ACL twice, before decryption, and as clear-text, following decryption. The Crypto Access Check on Clear-Text Packets feature removes the checking of clear-text packets that go through the IPsec tunnel just before encryption or just after decryption.

### Double ACL Check Behavior (before 12.3(8)T)

If the enterprise security policy does not permit the split-tunnel feature and the branch requires Internet access through the IPsec tunnel, the remote routers must also be configured to permit specific TCP and UDP traffic through the inbound access control list when the connection is initiated from within the remote router subnet.

To allow Internet access in configurations other than split tunnel, use Context Based Access Control (CBAC) in conjunction with the inbound access list. The following listing is an example of the configuration required:

```
ip inspect name CBAC tcp
ip inspect name CBAC udp
ip inspect name CBAC ftp
ip inspect name CBAC sip
interface Ethernet 0
description Inside
 ip address 10.81.7.1 255.255.255.248
interface Ethernet 1
description Outside
 ip address dhcp
 ip access-group INPUT_ACL in
 ip inspect CBAC out
ip access-list extended INPUT_ACL
 permit udp x.x.x.16 0.0.0.15 any eq isakmp
 permit udp x.x.x.16 0.0.0.15 any eq non500-isakmp
 permit esp x.x.x.16 0.0.0.15 any
 remark ! enterprise Address space
 permit ip 10.0.0.0 0.255.255.255 10.81.7.0 0.0.0.7
 permit udp any any eq bootpc
 permit udp x.x.x.40 0.0.0.1 eq ntp any
 permit tcp x.x.0.0 0.0.15.255 any eq 22
 permit icmp any any
 deny ip any any
```

### Crypto Access Check on Clear-Text Packets Feature (12.3(8)T and Later)

The Crypto Access Check on Clear-Text Packets feature removes the checking of inbound, decrypted clear-text packets against the outside interface inbound access list. When upgrading Cisco IOS software to a version that supports this feature, the following statement should be removed from the **ip access-list extended INPUT\_ACL**.

```
! enterprise Address space
 permit ip 10.0.0.0 0.255.255.255 10.81.7.0 0.0.0.7
```

The **ip inspect CBAC in** command can be removed from the Ethernet 0 interface.

If checking the decrypted clear-text packets against an access list is required, that function is now configured inside the crypto map global configuration. For more information about the Crypto Access Check on Clear-Text Packets feature, see the following URL:  
[http://www.cisco.com/en/US/partner/products/sw/iosswrel/ps5207/products\\_feature\\_guide09186a008022c2a5.html](http://www.cisco.com/en/US/partner/products/sw/iosswrel/ps5207/products_feature_guide09186a008022c2a5.html)

## Headend Considerations

Network location of the crypto headend in relation to headend firewalls affects the accessibility and performance of both systems. The network manager must make sure that all firewalls are properly configured to allow bi-directional tunnel traffic and that the crypto headend is accessible to the branch router.

## Common Configuration Errors

This section describes some common errors and problems encountered when configuring IPsec direct encapsulation.

### Crypto Peer Address Matching Using PSK

The IP address used as the crypto source address must match the address configured as the destination address on the crypto peer and vice-versa. Unless the address is configured specifically, the address of the outgoing interface is used as the crypto peer address, which causes the crypto peer to die during ISAKMP negotiation.

### Transform Set Matches

At least one matching IPsec transform set must be configured between each pair of crypto peers. When specifying a particular strength of encryption algorithm, both sides of the IPsec tunnel should match. Failure to do so could weaken the security of the entire solution.

### ISAKMP Policy Matching

The default ISAKMP policy for all Cisco IOS Software devices is as follows:

- Encryption—DES
- HMAC—SHA
- IKE authentication— RSA signature
- Diffie-Hellman—Group 1

If a stronger ISAKMP policy is required, both sides must support the policy. It is common but not required to use the same encryption level transform set and hash methods in the ISAKMP policy and the IPsec transform set.

## Scalability Considerations

This section presents steps for selecting Cisco products for a VPN solution, including sizing the headend and choosing the best Cisco product for the headend devices.

## General Scalability Considerations

This section describes general scalability considerations to help with design requirements.

### IPsec Encryption Throughput

Because each packet that is encrypted must traverse the encryption engine, you must consider the bidirectional throughput capacity of the IPsec encryption engine for both headend and branch devices. Several examples are shown in [Table 2](#) and [Table 3](#) for popular headend and branch connection speeds.

**Table 2**      *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**      *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 the router CPU also increases. However, with hardware-accelerated encryption available for all Cisco router products from the Cisco 871 through the Cisco 7600, processing is offloaded to the VPN hardware. However, main router CPU processing still occurs, so higher throughput still typically results in higher CPU consumption.

### Packets Per Second—The Most Important Factor

Bandwidth throughput capacity is important, but the packet rate for the connection speeds being terminated or aggregated is even more important.

In general, routers and encryption engines have upper boundaries for processing a given number of packets per second (pps). The size of the packets used for testing and throughput evaluations can understate or overstate true performance. For example, if a router with a VPN module can handle 20K pps, then 100-byte packets lead to 16 Mbps throughput. However, 1300-byte packets at the same packet rate yield 224 Mbps.

Because of the wide variance in throughput, pps is generally a better parameter for determining router forwarding potential than bits per second (bps). 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 affect the pps rate of that headend.

## Tunnel Quantity Affects Throughput

Generally speaking, as tunnel quantities increase, the overall throughput tends to decrease, although this is highly dependent on platform architecture. When a router receives a packet from a peer from which it has not recently decrypted a packet, it performs a lookup based on the security parameters index of the new packet. The new packet transform set information and negotiated key is then loaded into the hardware decryption engine for processing. Having traffic flowing on a larger numbers of SAs tends to negatively affect throughput performance.

Platforms with hardware-accelerated IPsec encryption are increasingly designed to offload tunnel processing overhead as well, which results 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 VPN Services Adapter for the Cisco 7200 VXR also implements tunnel processing overhead in hardware.

## Headend Scalability

This section describes considerations that are important for headend device scalability.

### Sizing the Headend

Headend devices are responsible for the following functions:

- Terminating ISAKMP and IPsec tunnels from the branch router
- Terminating ISAKMP keepalives to verify the state of the SAs to the branches
- Installing routes to the branch networks using RRI

It is important to size the headend correctly before choosing the device to deploy. This helps ensure that the overall network can support existing and future traffic profiles that the enterprise needs to run over the VPN.

The following are the critical factors that 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 HA requirement for the design?
- What is the expected performance margin or target CPU utilization?

All these factors must be considered together because any of them can become the limiting factor for sizing the headend.

The primary platform choices available today for headend routers are as follows:

- Cisco 7200VXR with NPE-G2 and VPN Services Adapter
- Cisco Catalyst 6500 or Cisco 7600 with Sup720, SIP400, and VPN SPA

The Cisco 7301 router, with performance nearly identical to the Cisco 7200VXR NPE-G1, is also an option. The main difference between the two platforms is that the Cisco 7200VXR is designed to be upgradeable to newer and faster processing engines and encryption modules. For example, a Cisco 7200VXR chassis can be upgraded from a NPE-G1 with a VAMII+ to a NPE-G2 with a VPN Services Adapter. The Cisco 7301 is a fixed-configuration platform, and is not upgradeable.

## Tunnel Aggregation Scalability

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

The number of IPsec tunnels that can be aggregated by a platform is a primary factor for recommending a platform, but the encryption pps rate is equally or more important.

## 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 each branch router IPsec tunnels are transported through at the headend? (For example, DS3, OC3, or OC12.)
- Branch connection speeds—What is the typical bandwidth at each branch office going to be? (For example, fractional-T1, T1, T3, broadband DSL, or cable.)
- Expected utilization—What is the maximum utilization of the WAN bandwidth under normal operation? The peak rate may also be important depending on customer requirements.

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

## Customer Requirement Aggregation Scalability Case Studies

This section includes examples that illustrate the factors affecting headend scalability.

### Customer Example with 300 Branches

A customer has the following design requirements:

- Number of branch offices—300
- Branch access speeds—128 kbps
- Headend access speed—DS3 (44.76 Mbps)
- Expected link utilization—80 percent

The calculation of aggregate bandwidth requirements is as follows:


- Typical case— $300 \times 128 \text{ kbps} \times 2(\text{bi-directional}) \times 80\% \text{ link utilization} = 61 \text{ Mbps}$
- Worst case— $300 \times 128 \text{ kbps} \times 2(\text{bi-directional}) \times 100\% \text{ link utilization} = 77 \text{ Mbps}$

In this example, the key requirements for the headend router platform are support for 300 tunnels and an aggregate encryption bandwidth of at least 61 Mbps. The traffic mix on the network determines the pps load on the processor. [Headend Scalability Test Results, page 29](#) includes tables that use a traffic mix commonly found on enterprise networks (e-mix) to determine pps based on bps. This bandwidth requirement is within the range of the headend access speed of a DS3 (90 Mbps, bi-directional).



A Cisco 7200VXR with NPE-G1 processor and SA-VAM2+ encryption accelerator supports these requirements. A Cisco 7200VXR with NPE-G2 and VPN Services Adapter supports even higher pps and Megabit per second rates, providing capacity for future growth. For platform-specific results, see [Headend Scalability Test Results, page 29](#). This design is shown in [Figure 6](#).

**Figure 6 Cisco 7200VXR-Based IPsec VPN Design Example**



### Customer Example with 1000 Branches

In this example, the customer has the following design requirements:

- Number of branch offices—1000
- Branch access speeds—384 kbps/1.5 Mbps cable/DSL
- Headend access speed—OC12 (622 Mbps)
- Expected link utilization—33 percent

The calculation for the aggregate bandwidth requirement is as follows:

- Typical case— $1000 \times (384 \text{ kbps} + 1.5 \text{ Mbps}) \times 33\% \text{ link utilization} = 628 \text{ Mbps}$
- Worst case— $1000 \times (384 \text{ kbps} + 1.5 \text{ Mbps}) \times 100\% \text{ link utilization} = 1.9 \text{ Gbps}$

In this example, the key requirements for the headend router platform are support for 1000 tunnels and an aggregate encryption bandwidth of at least 628 Mbps. This bandwidth requirement is within the range of the headend access speed of an OC12 (1.44 Gbps, bi-directional).

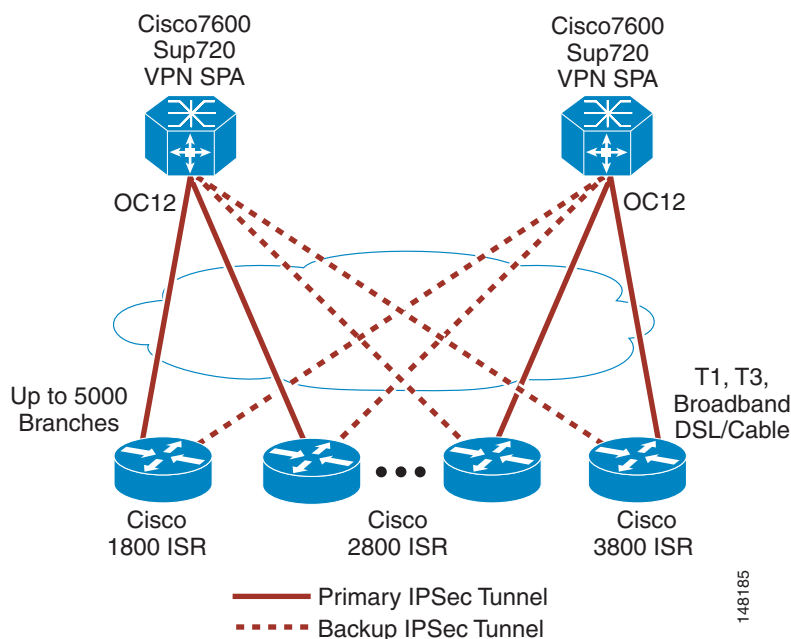
In this case, even though a Cisco 7200VXR with NPE-G1 processor and SA-VAM2+ encryption accelerator could support the number of tunnels required, the encryption bandwidth required would exceed performance of the platform, and produce a bottleneck at approximately 90–100Mbps.

The traffic mix on the network determines the pps load on the processor. [Headend Scalability Test Results, page 29](#) includes tables that use a traffic mix commonly found on enterprise networks (e-mix) to determine pps based on bps. Recommendations for this customer are to do one of the following:

- Divide the tunnels aggregated among multiple Cisco 7200VXR devices
- 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 Shared Port Adapter (SPA) can be recommended. Figure 7 illustrates this design. For platform-specific test results, see [Headend Scalability Test Results](#), page 29.

**Figure 7 Cisco 7600-Based IPsec VPN Design Example**



Headend aggregation designs based on the Cisco 7600 (or Catalyst 6500) and the VPN SPA can support many remote branches. The VPN SPA can support up to several thousand IPsec tunnels. For more details, see the corresponding data sheet at the following URL:

[http://www.cisco.com/en/US/partner/products/hw/routers/ps368/products\\_data\\_sheet0900aecd8027c9e.html](http://www.cisco.com/en/US/partner/products/hw/routers/ps368/products_data_sheet0900aecd8027c9e.html)

## Branch Office Scalability

The branch routers are responsible for the following functions:

- Terminating ISAKMP and IPsec SAs from the headend routers
- Terminating ISAKMP keepalives to verify the state of the SAs to the headend routers

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 (for example, fractional T1, T1, T3, or broadband cable/DSL)
- What other services is the branch router providing? (For example, DHCP, NAT/PAT, VoIP, IOS Firewall, IOS-IPS.)
- The pps rate (traffic size and traffic mix), which is the largest single factor affecting branch router scalability



---

The number of IPsec tunnels does not play a role in the branch device sizing because each branch router terminates a single tunnel in the recommended topology.

---

A primary concern is the amount of traffic throughput (pps and bps) along with the corresponding CPU utilization. The branch router must have sufficient CPU cycles to service periodic events, such as SNMP and Syslog activities, local CLI command processing, and establishing and re-keying ISAKMP and IPsec SAs. After initial deployment and testing, it may be possible to run branch routers at higher CPU utilization levels under normal operational conditions. However, the general Cisco recommendation and the specific recommendation of this design guide is to choose a branch router that can meet all processing demands with CPU utilization of 65 percent or less.

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

## Scalability Test Results (Unicast Only)

This section provides Cisco test results for design guidance on the scalability of various platforms in IPsec direct encapsulation VPN configurations. IPmc results are not included.

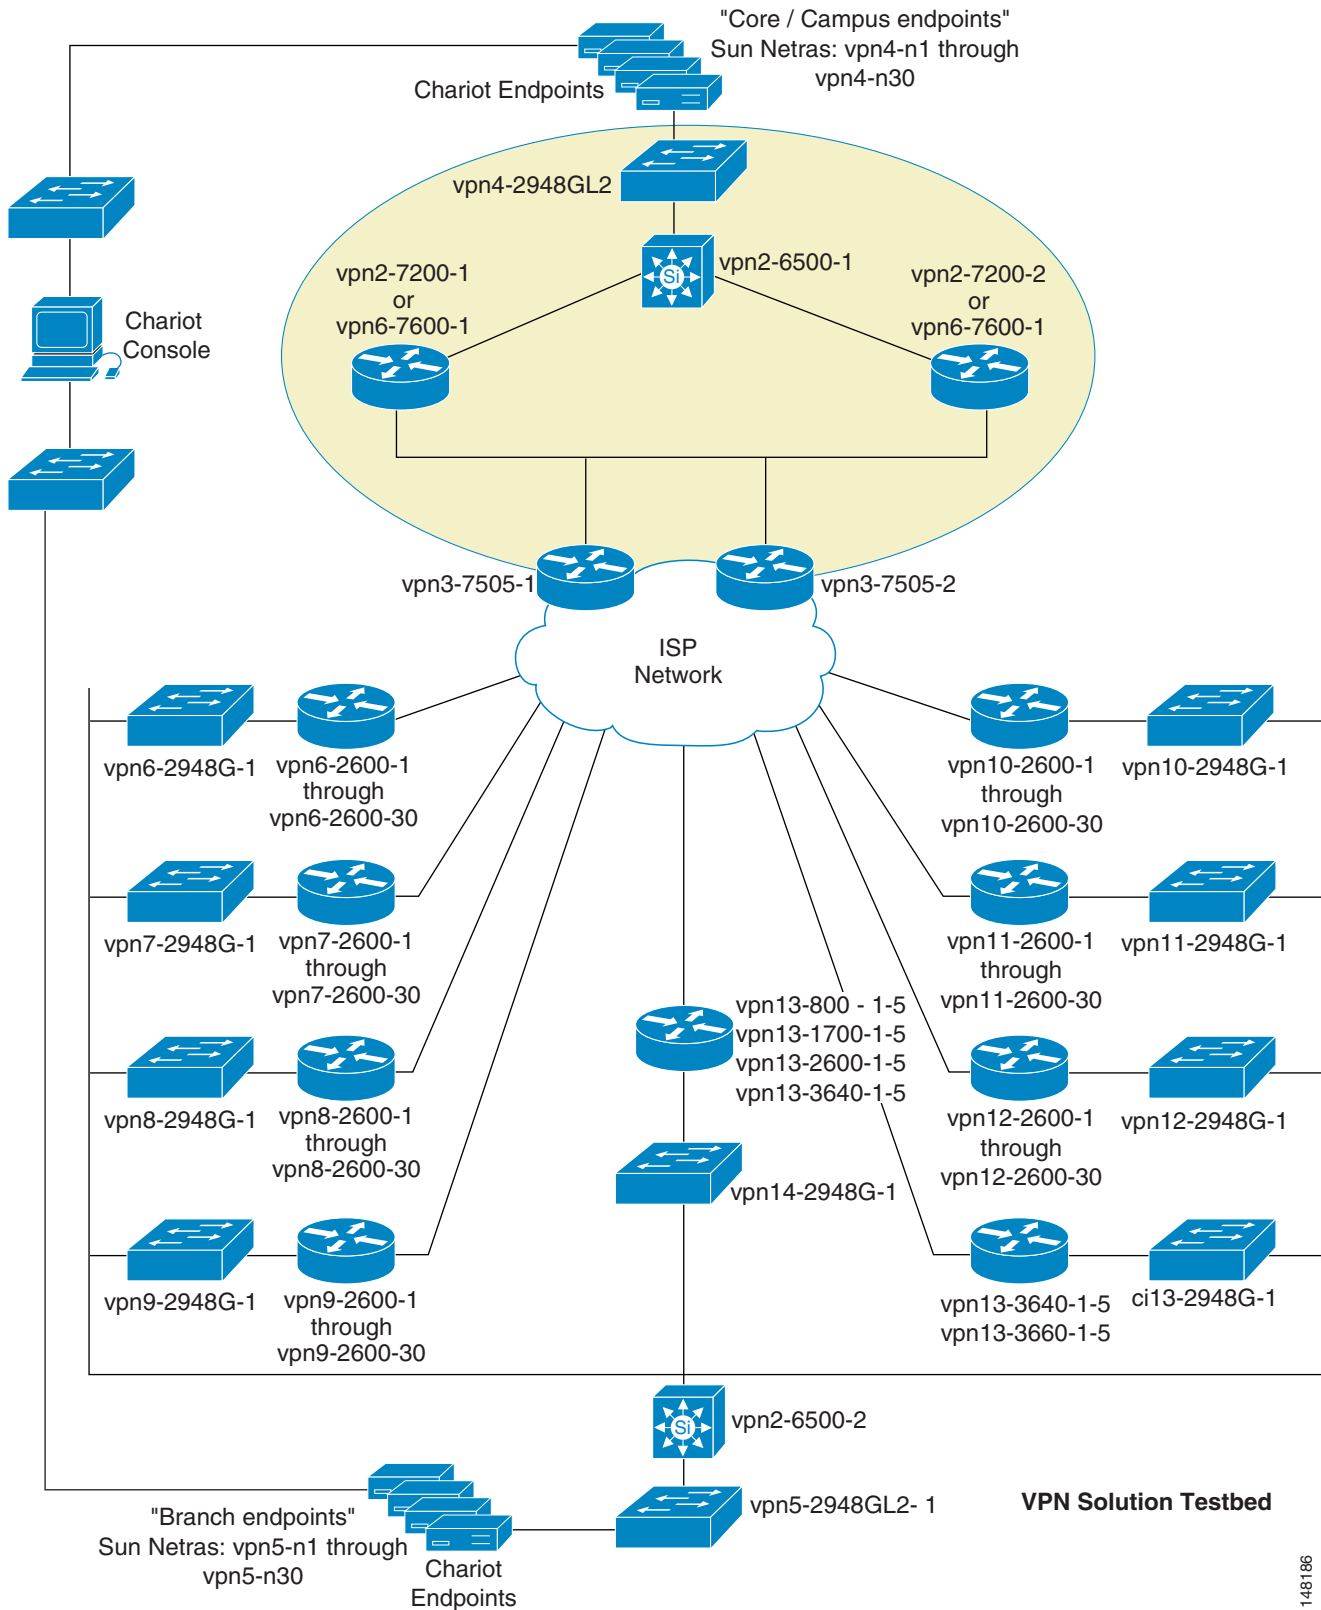
### Scalability Test Methodology

This section describes the methodology used for performing the scalability tests.

#### Overview

The headend scalability test bed consists of a variety of Cisco branch routers (including the Cisco 1700, 2600, 3600, 3700, 1800, 2800, and 3800 products) homed to various types of headends, as shown in [Figure 8](#).

Figure 8 Scalability Test Bed Network Diagram



148186

For most of the traffic, flows are established using the Ixia Chariot testing tool. The bps mix of traffic is approximately 35 percent UDP and 65 percent TCP. Application types represented in the mix include VoIP, FTP, DNS, HTTP, POP3, and TN3270.

The average packet size is 188 bytes, from headend to branch, and 144 bytes from branch to headend. This relatively small average packet size ensures that the scalability results presented support a converged network design, and tend to be fairly conservative. A network carrying data-only traffic, with a larger average packet size, may achieve better bps performance than indicated by the test results. However, the pps performance given a specific CPU value should be the same.

Some traffic is also generated by the IP SLA feature (formerly known as Cisco SAA) in Cisco IOS software, using the HTTP get script, with the branch routers making an HTTP get call to an HTTP server in the core. Testing was conducted without fragmentation occurring in the network by setting the MTU to 1300 bytes on the test endpoints.

## Headend Scalability Test Results

Table 4 show the results of testing with a configuration for IPsec direct encapsulation with DPD, RRI, and HSRP without enabling other Cisco IOS software features (such as IOS-FW, PAT, ACLs, IPS, or QoS):

**Table 4** Headend Scalability Test Results

Platform	# of Tunnels	Throughput (kpps)	Throughput (Mbps)	CPU %
Cisco 7200VXR NPE-G1 Dual SA-VAMII+	1040	49.4	109.0	81%
Cisco 7200VXR NPE-G2 Single SA-VAMII+ in Slot 3 <sup>1</sup>	1000	68	158	80%
Cisco 7200VXR NPE-G2 VPN Services Adapter (VSA)	1000	235	522	89%
Cisco Catalyst 6500 Sup2 VPNSM	1000	436.0	949.0	N/A
	4000	440.9	1004.0	N/A
Cisco 7600 Sup720 VPN SPA SSC-400	1000	601.6	1255.0	N/A
	5000	see <sup>2</sup>	see <sup>2</sup>	N/A

- Both the Cisco 7200VXR NPE-G1 and G2 were tested with single SA-VAMII+ and dual SA-VAMII+ in chassis slots as well as using the Port Adapter Jacket Card. With the conservative traffic profile used for these tests, that being predominately VoIP as a percent of total packets, the total throughput varied marginally. From that observation, it can be concluded that encryption performance is not the limiting factor, rather it is the main CPU packet per second switching rate. Traffic profiles that are predominately large data packets may differ and actually benefit from dual encryption adapters.
- The Cisco 7600 Sup720 VPN SPA SSC-400 has been tested supporting 5000 tunnels in a Single Tier with Route Server configuration. The Large-Scale IPsec Aggregation Network Design Guide describes this in detail.

Headend scalability testing did not include an exhaustive evaluation of the maximum number of tunnels that can be terminated to headend devices. In addition, scalability testing of branch routers was performed with two tunnels per branch. This did not include exhaustive testing of the number of tunnels that the various platforms can support.

## Branch Office Scalability Test Results

Table 5 shows the results of testing with a configuration for IPsec with DPD, RRI, and HSRP, using the same testing methodology described earlier.

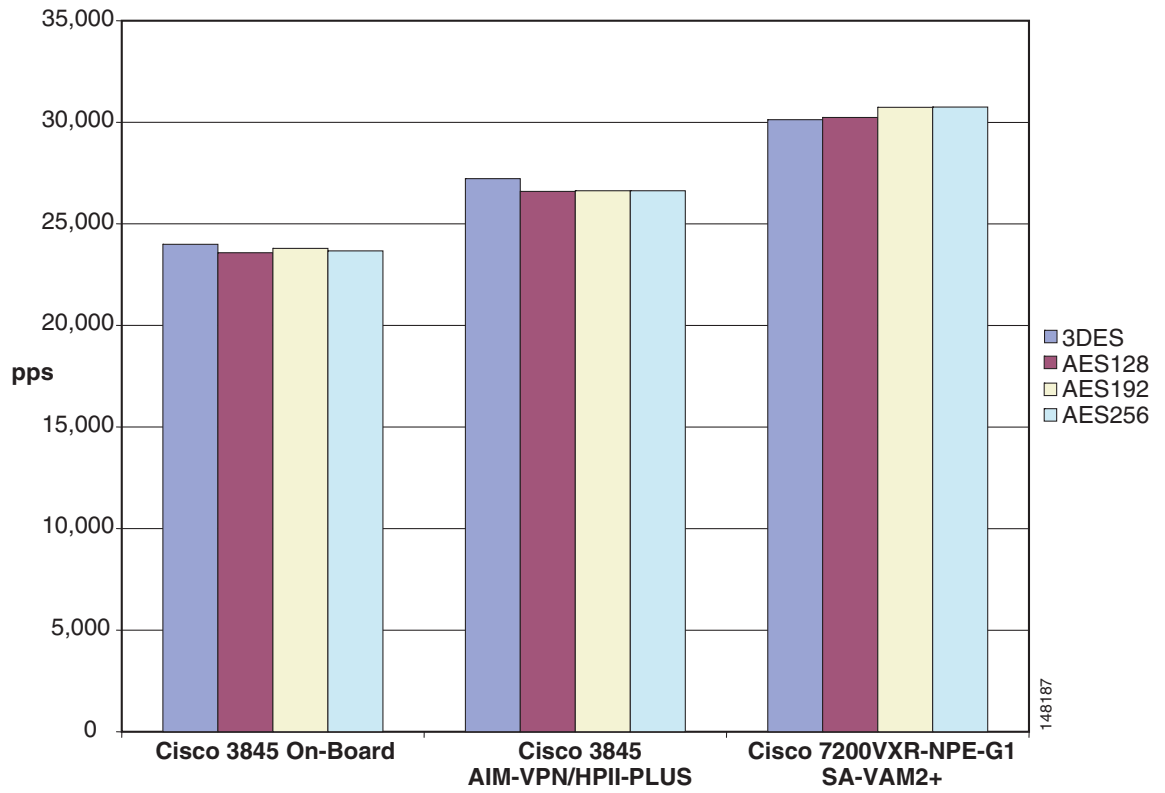
The results shown in [Table 5](#) are for a p2p GRE over IPsec design because test results were not available for an IPsec direct encapsulation topology using these platforms. Throughput and performance should not differ significantly, but are slightly better without p2p GRE. The results in [Table 5](#) include other integrated Cisco IOS Software features such as Cisco IOS-FW, PAT, and access lists, but not QoS or IPS.

**Table 5** *Branch Office Scalability Test Results (p2p GRE over IPsec Design)*

Platform	# of Tunnels	Throughput (kpps)	Throughput (Mbps)	CPU %
Cisco 3845 ISR	On-board	24.0	48.8	81%
	AIM-VPN/HPII-Plus	27.1	50.1	80%
Cisco 3825 ISR	On-board	18.2	36.6	81%
	AIM-VPN/EPII-Plus	20.1	42.8	79%
Cisco 2851 ISR	On-board	11.4	23.8	79%
	AIM-VPN/EPII-Plus	15.0	30.8	80%
Cisco 2821 ISR	On-board	6.0	13.6	53%
	AIM-VPN/EPII-Plus	12.3	25.9	78%
Cisco 2811 ISR	On-board	2.6	5.8	79%
	AIM-VPN/EPII-Plus	3.6	8.0	80%
Cisco 2801 ISR	On-board	2.6	5.8	83%
	AIM-VPN/EPII-Plus	3.9	8.4	79%
Cisco 1841 ISR	On-board	2.5	5.7	82%
	AIM-VPN/BPII-Plus	3.9	8.8	80%
Cisco 1811W with no BVI configured	On-board	7.6	16.0	81%
Cisco 1811W with BVI configured	On-board	4.3	9.3	82%
Cisco 871W with no BVI configured	On-board	2.0	4.4	85%
Cisco 871W with BVI configured	On-board	1.1	2.4	84%

## Scalability Test Results (AES Compared to 3DES)

Both 3DES and AES encryption are available for all the products shown here, including hardware accelerated IPsec. Not every test was executed with both 3DES and AES, but several snapshot tests were performed to compare performance. As can be seen in [Figure 9](#), results are fairly comparable, with little to no variation in performance, even for AES with wider key lengths.

**Figure 9 Comparison of 3DES and AES Performance**

## Failover and Convergence Testing

This section describes failover and convergence testing.

### Stateless Failover

With stateless failover, IPsec must re-negotiate ISAKMP and IPsec SAs with the standby router in addition to the time required for the HSRP process to discover that its primary crypto peer has failed. For a network with a large number of crypto peers, this process can take several minutes. The test results for stateless failover are shown in [Table 6](#).

**Table 6 Stateless Failover on Cisco 7200VXR**

Platform/Status	Headend 1	Headend 2
Cisco 7200VXR NPE-G1	Cisco IOS version 12.2(13)S	Cisco IOS version 12.2(13)S
Starting condition	81 Mbps 250 branches 64% CPU	0 Mbps 0 branches 0% CPU
During failover	Powered off	81 Mbps, 250 branches 64% CPU
Cisco 7200VXR NPE-G1	Cisco IOS version 12.2(13)S	Cisco IOS version 12.2(13)S

**Table 6 Stateless Failover on Cisco 7200VXR (continued)**

Platform/Status	Headend 1	Headend 2
Starting condition	79 Mbps 500 branches 68% CPU	0 Mbps 0 branches 0% CPU
During failover	Powered off	79 Mbps 500 branches 68% CPU

After completion of the test, the crypto peers renegotiated SAs with their primary headend via the HSRP preempt command. Both the failover and renegotiation processes took approximately three and one-half minutes to complete with the 250 tunnel scenario, and five and one-half minutes to complete with the 500 tunnel test.

### Stateful Failover

Stateful Failover was also tested, with two crypto headends (active and standby) running HSRP and exchanging IPsec SA database state information. The test results are shown in [Table 7](#).

**Table 7 Stateful Failover**

Platform/Status	Active Headend	Standby Headend
Cisco 7200VXR NPE-400 SA-VAM	Cisco IOS version 12.2(11)YX1	Cisco IOS version 12.2(11)YX1
Starting condition	65 Mbps 1040 branches 68% CPU	0 Mbps 0 branches 0% CPU
During failover	Powered off	65 Mbps 1040 branches 68% CPU
Cisco Catalyst 6500 Sup2 VPNISM	Cisco IOS version 12.2(14)SY1	Cisco IOS version 12.2(14)SY1
Starting condition	943 Mbps 1040 branches	0 Mbps 0 branches
During failover	Powered off	943 Mbps, 1040 branches

All tunnels failed over properly to the standby headend in 1–3 seconds.

## Software Releases Evaluated

[Table 8](#) lists the software releases used in the scalability testing.



**Table 8**      **Software Releases Evaluated**

Cisco 7600 VPN SPA	Cisco IOS 12.2(18)SXE2
Cisco Catalyst 6500 VPNSM	Cisco IOS 12.2(14)SY1 Cisco IOS 12.2(18)SXE2
Cisco headend routers (7200VXR, 7301)	Cisco IOS 12.2(13)S Cisco IOS 12.3(5) Cisco IOS 12.2(11)YX1
Cisco 7200 VXR G2 VPN Services Adapter	Cisco IOS 12.4(4)XD-0629
Cisco branch office routers (17xx, 26xx, 36xx, 37xx)	Cisco IOS 12.2(13)T Cisco IOS 12.3(8)T5
Cisco branch office ISRs (1841, 28xx, 38xx)	Cisco IOS 12.3(8)T5 Cisco IOS 12.3(11)T2
Cisco remote office routers (831, 871W, and 1811W)	831—Cisco IOS 12.3(8)T5 871W—Cisco IOS 12.3(8)Y1 1811W—Cisco IOS 12.3(14)YT1
Cisco PIX 535	PIXOS 6.3.1
Cisco VPN Concentrator 3080	SW version 4.0.0

Before selecting the Cisco IOS software, perform the appropriate research on Cisco.com to make sure you select the best release for your requirements. If you have technical questions, contact Cisco TAC.

## Scalability Test Bed Configuration Files

This section lists the configurations used for the central and branch sites. These configurations have been extracted from real configurations used in scalability testing and are provided as a reference only.

### Cisco 7200VXR Headend Configuration

Two headend devices are in the test bed, configured for dynamic crypto maps, DPD RRI, and HSRP. The listing for the first headend device is shown. No crypto peer statement or crypto access list is required. The ISAKMP pre-shared key is shown for one branch router.

### Cisco 7200VXR Headend Configuration

Headend #1:

```
ip cef
crypto isakmp policy 1
  encr 3des
  authentication pre-share
  group 2
crypto isakmp key bigsecret address 0.0.0.0 0.0.0.0
! wildcard preshared key on HE
crypto isakmp keepalive 10
crypto ipsec transform-set vpn-test esp-3des esp-sha-hmac
crypto dynamic-map dmap 10
  set transform-set vpn-test
```

```

reverse-route
crypto map dynamic-map local-address GigabitEthernet0/1
crypto map dynamic-map 10 ipsec-isakmp dynamic dmap
interface GigabitEthernet0/1
description GigabitEthernet0/1
ip address 192.168.251.2 255.255.255.0
duplex auto
speed auto
media-type gbic
negotiation auto
standby ip 192.168.251.1
standby timers msec 50 1
standby priority 101
standby preempt
standby name outside
standby track GigabitEthernet0/2
crypto map dynamic-map redundancy outside
interface GigabitEthernet0/2
description GigabitEthernet0/2
ip address 10.57.1.1 255.255.255.0
duplex auto
speed auto
media-type gbic
negotiation auto
router eigrp 1
redistribute static metric 1000 100 255 1 1500
network 10.0.0.0
no auto-summary
ip route 0.0.0.0 0.0.0.0 192.168.251.2

```

## Cisco 7600 Headend Configuration

### Headend #1:

```

no ip domain-lookup
ipv6 mfib hardware-switching replication-mode ingress
mls ip multicast flow-stat-timer 9
no mls flow ip
no mls flow ipv6
no mls acl tcam share-global
mls cef error action freeze
no scripting tcl init
no scripting tcl encdir
crypto isakmp policy 10
encr 3des
authentication pre-share
group 2
crypto isakmp key bigsecret address 0.0.0.0 0.0.0.0
! wildcard preshared key on HE
crypto isakmp keepalive 10
crypto ipsec transform-set vpn-test esp-3des esp-sha-hmac
crypto dynamic-map dmap 10
set transform-set vpn-test
crypto map dynamic-map local-address Vlan100
crypto map dynamic-map 10 ipsec-isakmp dynamic dmap
redundancy
mode sso
main-cpu
auto-sync running-config
spanning-tree mode pvst
no spanning-tree optimize bpdu transmission
spanning-tree extend system-id
power redundancy-mode combined

```

```

no diagnostic cns publish
no diagnostic cns subscribe
vlan internal allocation policy ascending
vlan access-log ratelimit 2000
interface GigabitEthernet3/1
  description GigabitEthernet3/1 Outside Interface
  no ip address
  load-interval 30
  crypto connect vlan 100
interface GigabitEthernet4/0/1
  description GigabitEthernet4/0/1
  switchport
  switchport trunk encapsulation dot1q
  switchport trunk allowed vlan 1,100,1002-1005
  switchport mode trunk
  mtu 9216
  no ip address
  load-interval 30
  flowcontrol receive on
  flowcontrol send off
  spanning-tree portfast trunk
interface GigabitEthernet4/0/2
  description GigabitEthernet4/0/2
  switchport
  switchport trunk encapsulation dot1q
  switchport trunk allowed vlan 1,1002-1005
  switchport mode trunk
  mtu 9216
  no ip address
  load-interval 30
  flowcontrol receive on
  flowcontrol send off
  spanning-tree portfast trunk
interface GigabitEthernet5/1
  description GigabitEthernet5/1 to vpn2-7200-1 GE0/1
  ip address 192.168.181.2 255.255.255.0 secondary
  ip address 192.168.161.2 255.255.255.0
  no ip redirects
  load-interval 30
interface GigabitEthernet5/2
  description GigabitEthernet5/2 to vpn2-7200-2 GE0/1
  ip address 192.168.191.2 255.255.255.0 secondary
  ip address 192.168.171.2 255.255.255.0
  no ip redirects
  load-interval 30
interface Vlan1
  description Vlan1
  no ip address
  load-interval 30
  shutdown
interface Vlan100
  description Vlan100
  ip address 192.168.241.1 255.255.255.0
  load-interval 30
  no mop enabled
  crypto map dynamic-map
  crypto engine subslot 4/0
ip classless
ip route 192.168.0.0 255.255.0.0 192.168.241.2

```

## ISR Branch Configuration

This section shows the configuration listing for one branch site router. The crypto peer is the HSRP address at the headend. This configuration shows QoS for VoIP flows (shaping and queueing) applied to the physical (outside) interface.

### Branch #1:

```
ip cef
class-map match-all VOICE
  match ip dscp ef
class-map match-any CALL-SETUP
  match ip dscp af31
  match ip dscp cs3
class-map match-any INTERNETWORK-CONTROL
  match ip dscp cs6
class-map match-all TRANSACTIONAL-DATA
  match ip dscp af21
policy-map 192kb
  class CALL-SETUP
    bandwidth percent 2
  class INTERNETWORK-CONTROL
    bandwidth percent 5
  class TRANSACTIONAL-DATA
    bandwidth percent 22
    queue-limit 16
  class VOICE
    priority 56
  class class-default
    fair-queue
    queue-limit 6
policy-map 192kb-shaper
  class class-default
    shape average 182400 1824 0
    service-policy 192kb
crypto isakmp policy 1
  encr 3des
  authentication pre-share
  group 2
crypto isakmp key bigsecret address 192.168.251.1
crypto isakmp keepalive 10
crypto ipsec transform-set vpn-test esp-3des esp-sha-hmac
crypto map static-map local-address Serial0/0
crypto map static-map 10 ipsec-isakmp
  set peer 192.168.251.1
  set transform-set vpn-test
  match address b000
  qos pre-classify
interface Loopback0
  description Loopback0
  ip address 10.61.138.254 255.255.255.255
interface Serial0/0
  description Serial0/0
  bandwidth 192
  ip address 192.168.90.6 255.255.255.252
  service-policy output 192kb-shaper
  crypto map static-map
interface FastEthernet0/1
  description FastEthernet0/1
  ip address 10.61.138.129 255.255.255.192 secondary
  ip address 10.61.138.1 255.255.255.128
  speed 100
  full-duplex
```

```
ip route 0.0.0.0 0.0.0.0 192.168.90.5
ip access-list extended b000
permit ip 10.61.138.0 0.0.0.255 10.0.0.0 0.255.255.255
```

## Appendix A—Scalability Test Results for Other Cisco Products

This section contains scalability and performance data for other Cisco products. The testing has been performed under the same conditions as outlined in [Scalability Test Results \(Unicast Only\)](#), page 27.

### Cisco Headend VPN Routers (Legacy)

[Table 9](#) shows the test results for older Cisco headend router platforms.

**Table 9** Cisco Headend Router Platform Throughput, IPsec with DPD, RRI and HSRP<sup>1</sup>

Headend Router Platform	Hardware Acceleration Option	# of Tunnels	Throughput	Kpps	Throughput
Cisco 7200VXR with NPE-G1	Dual SA-VAM	1040	48.1	106.7	81%
Cisco 7200VXR with NPE-400	SA-VAM	1040	31.7	71.7	88%
Cisco 3745	AIM-VPN/HPII	120	14.5	28.6	80%

1. Without access lists, PAT, IOS-FW, IPS, or QoS

### Other Cisco Products for the Headend

Several other Cisco products support IPsec VPN tunnel termination in a headend environment, such as the Cisco VPN 3000 Concentrator series and the Cisco PIX Firewall series. [Table 10](#) shows the test results for the Cisco PIX 535 with the VAC Plus and the Cisco 3080 with SEP and SEP-E.

**Table 10** Other Cisco Headend Platform Throughput

Headend Router Platform	Hardware Acceleration Option	# of Tunnels	Throughput	Throughput Mbps	CPU % Utilization
Cisco PIX 535	VAC Plus	500	61.8	122.9	80%
Cisco 3080	SEP	138	19.5	38.8	80%
	SEP-E	138	19.6	39.4	52%

Firewall rules require that traffic enters an interface on the firewall to exit that firewall through a different interface; in effect, passing all the way through the device. Cisco PIX OS before version 7 does not support branch-to-hub-to-branch communications through the Cisco PIX as a crypto headend. A packet cannot be received and transmitted out the same physical interface of a Cisco PIX system. Therefore, all the results shown for the Cisco PIX are branch-to-hub communications.

See the following URLs for more product information on the Cisco VPN 3000 and Cisco PIX series:

- <http://www.cisco.com/warp/public/cc/pd/hb/vp3000/>
- <http://www.cisco.com/warp/public/cc/pd/fw/sqfw500/>

## Cisco Branch Office VPN Routers (Legacy)

Table 11 shows the test results for Cisco branch router platforms configured for p2p GRE over IPsec.

**Table 11** Cisco Branch Router Platform Throughput with p2p GRE Over IPsec<sup>1</sup>

Branch Router Platform	Hardware Acceleration Option	Throughput Kpps	Throughput Mbps	CPU % Utilization
Cisco 3745	AIM-VPN/HPII	13.4	28.7	82%
Cisco 3725	AIM-VPN/EPII	6.8	15.5	81%
Cisco 3660	AIM-VPN/HPII	4.8	10.9	80%
Cisco 2691	AIM-VPN/EPII	5.1	11.4	81%
Cisco 2651XM	AIM-VPN/BPII	1.3	3.0	85%
Cisco 1760	MOD1700VPN	0.9	2.0	81%
Cisco 1711	On-board	0.8	2.0	81%
Cisco 831	On-board	0.4	1.0	74%

1. One tunnel with IOS-FW, PAT, and access lists, but no QoS or IPS

## Appendix B—References

This section includes the following references and readings for further information:

- Documents (<http://www.cisco.com/go/srnd>)
  - IPsec VPN WAN Design Overview
  - p2p GRE over IPsec Design Guide
  - Dynamic Multipoint VPN (DMVPN) Design Guide
  - Virtual Tunnel Interface (VTI) Design Guide
  - Voice and Video Enabled IPsec VPN (V3PN) Design Guide
  - Multicast over IPsec VPN Design Guide
  - IPsec VPN Redundancy and Load Sharing Design Guide
  - Digital Certification/PKI for IPsec VPN Design Guide
- Request For Comment (RFC) papers:
  - Security Architecture for the Internet Protocol—RFC 2401
  - IP Authentication Header—RFC 2402
  - The Use of HMAC-MD5-96 within ESP and AH—RFC 2403
  - The Use of HMAC-SHA-1-96 within ESP and AH—RFC 2404

- The ESP DES-CBC Cipher Algorithm With Explicit IV—RFC 2405
- IP Encapsulating Security Payload (ESP)—RFC 2406
- The Internet IP Security Domain of Interpretation for ISAKMP—RFC 2407
- Internet and Key Management Protocol (ISAKMP)—RFC 2408
- The Internet Key Exchange (IKE)—RFC 2409
- The NULL Encryption Algorithm and Its Use With IPsec—RFC 2410
- IP Security Document Roadmap—RFC 2411
- The OAKLEY Key Determination Protocol—RFC 2412

## Appendix C—Acronyms and Definitions

Term	Definition
3DES	Triple Data Encryption Standard
ACL	Access Control List
AES	Advanced Encryption Standard
AH	Authentication Header
AIM	Advanced Integration Module
ATM	Asynchronous Transfer Mode
CA	Certificate Authority
CAC	Call Admission Control
CAR	Committed Access Rate
CBWFQ	Class Based Weighted Fair Queuing
CEF	Cisco Express Forwarding
CPE	Customer Premises Equipment
cRTP	Compressed Real-Time Protocol
DES	Data Encryption Standard
DHCP	Dynamic Host Configuration Protocol
DMVPN	Dynamic Multipoint Virtual Private Network
DMZ	De-Militarized Zone
DNS	Domain Name Service
DPD	Dead Peer Detection
DSL	Digital Subscriber Line
EIGRP	Enhanced Interior Gateway Routing Protocol
ESP	Encapsulating Security Protocol
FIFO	First In First Out
FQDN	Fully Qualified Domain Name
FR	Frame Relay
FRTS	Frame Relay Traffic Shaping

<b>Term</b>	<b>Definition</b>
FTP	File Transfer Protocol
GRE	Generic Route Encapsulation
HSRP	Hot Standby Router Protocol
ICMP	Internet Control Message Protocol
IKE	Internet Key Exchange
IOS	Internetwork Operating System
IP	Internet Protocol
IPmc	IP Multicast
IPsec	IP Security
ISP	Internet Service Provider
Layer 2	OSI reference model Link Layer
Layer 3	OSI reference model Network Layer
Layer 4	OSI reference model Transport Layer
LFI	Link Fragmentation and Interleaving
LLQ	Low Latency Queuing
L2TP	Layer 2 Tunneling Protocol
mGRE	Multipoint Generic Route Encapsulation
MLPPP	Multi-link Point-to-point Protocol
MPLS	Multi-Protocol Label Switching
MTU	Maximum Transmission Unit
NAT	Network Address Translation
NetFlow	Cisco IOS component, collects and exports traffic statistics
NHRP	Next Hop Resolution Protocol
NHS	Next-Hop Server
ODR	On-Demand Routing
OSPF	Open Shortest Path First
p2p GRE	Point to Point Generic Route Encapsulation
PAT	Port Address Translation
PBR	Policy Based Routing
PE	Premises Equipment
PPTP	Point-to-Point Tunneling Protocol
PVC	Permanent Virtual Circuit
QoS	Quality of Service
RADIUS	Remote Authentication Dial In User System
RRI	Reverse Route Injection
RTP	Real-Time Protocol
SA	Security Association



<b>Term</b>	<b>Definition</b>
SAA	Service Assurance Agent
SHA-1	Secure Hash Algorithm One
SLA	Service Level Agreement
SNMP	Simple Network Management Protocol
SOHO	Small Office/Home Office
SPA	Shared Port Adapter
SRST	Survivable Remote Site Telephony
SSO	Stateful Switchover
SSP	State Synchronization Protocol
TCP	Transmission Control Protocol
TED	Tunnel Endpoint Discovery
ToS	Type of Service
UDP	User Datagram Protocol
VoIP	Voice over IP
V <sup>3</sup> PN	Voice and Video Enabled IPsec VPN
SA-VAM	VPN Acceleration Module
VPN	Virtual Private Network
VTI	Virtual Tunnel Interface
WAN	Wide Area Network
WRED	Weighted Random Early Detection