CHAPTER 15

IPsec High Availability Options

Redundancy is typically found at various spots throughout networks. Because any path or component has the potential to fail, an alternate solution ensures that data continues to flow from one point to another. However, redundancy does come with a price. The configuration of additional paths could imply that such paths must be procured from the provider, and such paths are not free. To avoid a hardware failure, additional hardware must be procured and installed. To avoid a cable cut catastrophe, alternate physical paths should be planned.

As the amount of equipment and the number of options increase, so do the complexity and cost of the network. This chapter explores how to offer high availability options to IPsec VPNs. Many of the best practice concepts for general redundancy and high availability apply to the IPsec VPN world. Such opportunities are highlighted when appropriate.

This chapter only discusses the CLI methodology for offering high availability solutions. Such configurations appear on the ISCW certification test.

"Do I Know This Already?" Quiz

The purpose of the "Do I Know This Already?" quiz is to help you decide whether you really need to read the entire chapter. If you already intend to read the entire chapter, you do not necessarily need to answer these questions now.

The 13-question quiz, derived from the major sections in the "Foundation Topics" portion of the chapter, helps you to determine how to spend your limited study time.

Table 15-1 outlines the major topics discussed in this chapter and the "Do I Know This Already?" quiz questions that correspond to those topics.

Table 15-1 *"Do I Know This Already?" Foundation Topics Section-to-Question Mapping*

Foundation Topics Section	Questions Covered in This Section	Score
Sources of Failures	1	
Failure Mitigation	2	
Failover Strategies	3–12	
WAN Backed Up by an IPsec VPN	13	
Total Score		

CAUTION The goal of self-assessment is to gauge your mastery of the topics in this chapter. If you do not know the answer to a question or are only partially sure of the answer, you should mark this question wrong for purposes of self-assessment. Giving yourself credit for an answer that you correctly guess skews your self-assessment results and might provide you with a false sense of security.

1. Which of the following are the primary failure sources (select all that apply)?
   a. Remote peer failure
   b. Local peer failure
   c. Access link failure
   d. Device failure
   e. Path failure
   f. Cable failure
2. Multiple interfaces or devices are a way to mitigate which of the following failures (select all that apply)?
   a. Access link failure
   b. Remote peer failure
   c. Local peer failure
   d. Device failure
   e. Path failure
3. Which of the following failover means are considered stateless (select all that apply)?
   a. DPD
   b. HSRP
   c. SSO
   d. IPC
   e. IGP in GRE over IPsec

4. Periodic DPD mode has which of the following characteristics (select all that apply)?
   a. It is the default mode in Cisco IOS devices.
   b. DPD keepalive messages are used only when there is a lull in tunnel traffic.
   c. DPD keepalive messages are never sent during idle tunnel moments.
   d. DPD keepalive messages are periodically sent between IPsec VPN peers.
   e. DPD keepalive messages are sent in addition to normal IPsec rekey messages.
5. How is the primary DPD peer configured in an IOS device?
   a. **set peer** *ip-address*
   b. **set peer** *ip-address* **default**
   c. **set dpd peer** *ip-address*
   d. **set dpd peer** *ip-address* **default**
   e. **set peer** *ip-address* **dpd default**
6. What IOS command determines the frequency of DPD keepalive messages?
   a. **dpd keepalive** *seconds* [*retries*]
   b. **set peer** *ip-address* **dpd-keepalive** *seconds*
   c. **set dpd keepalive** *seconds* [*retries*]
   d. **crypto isakmp keepalive** *seconds* [*retries*]
   e. **crypto isakmp dpd keepalive** *seconds* [*retries*]
7. Which dynamic routing protocols should be used within GRE over IPsec tunnels due to fast convergence (select all that apply)?
   a. RIPv2
   b. OSPF
   c. EIGRP
   d. BGP
   e. IS-IS
8. Which of the following terms best describes an HSRP active router?
   a. The router is properly configured and ready to participate in HSRP elections.
   b. The router is actively working on electing a primary HSRP router for the group.
   c. The router is the current primary forwarding router for the group.
   d. The router helps elect the next primary router if the current one fails.
   e. The router is alive but not the primary router in the group.

9. In HSRP, what is the purpose of the standby router?
   a. The router is next in line to be the active router of the group.
   b. The router is awaiting proper configuration.
   c. The router is forwarding traffic for only a small percentage of the group.
   d. The router is waiting for the group to become active.
   e. The router is waiting for permission from the primary router to become active.
10. What does the **preempt** command do in an HSRP router?
   a. Forces this router to be the active router in any condition
   b. Forces this router to be the active router only if it has the highest HSRP priority of the group
   c. Forces this router to be the active router only if it has the lowest HSRP priority of the group
   d. Prevents this router from becoming the active router of the group
   e. Prevents any other router from ever becoming active
11. What happens to an IPsec VPN that terminates on an HSRP group IP address when the active router fails?
   a. Nothing; the standby router assumes active status.
   b. The IPsec VPN is dynamically migrated to the standby router.
   c. The IPsec VPN drops, and is reestablished with the group IP address.
   d. The IPsec VPN drops, and is reestablished with the standby router IP address.
   e. Nothing, because all routers in the HSRP group terminated the IPsec VPN initially.
12. Which protocols are used to provide IPsec stateful failover (select all that apply)?
   a. OSPF
   b. HSRP
   c. EIGRP
   d. DPD
   e. SSO

13. How can an IPsec VPN be used to back up a typical WAN connection (select all that apply)?

    a. IPsec VPNs are not used for this type of backup.

    b. Use a floating static route with a high AD that becomes active only when the primary path fails.

    c. Use a floating static route with a low AD that becomes active only when the primary path fails.

    d. Use an IGP over the IPsec VPN, and make the VPN dynamic routes less favorable.

    e. Use an IGP over the GRE over IPsec VPN, and make the VPN dynamic routes less favorable.

The answers to the "Do I Know This Already?" quiz are found in Appendix A, "Answers to the 'Do I Know This Already?' Quizzes and Q&A Sections." The suggested choices for your next step are as follows:

- **8 or fewer overall score**—Read the entire chapter. This includes the "Foundation Topics," "Foundation Summary," and "Q&A" sections.
- **9 or 11 overall score**—Begin with the "Foundation Summary" section, and then go to the "Q&A" section.
- **12 or more overall score**—If you want more review on these topics, skip to the "Foundation Summary" section, and then go to the "Q&A" section. Otherwise, move to the next chapter.

Foundation Topics

Sources of Failures

The network has a number of possible points vulnerable to failure. Remember that an IPsec VPN is an end-to-end connection. It typically travels across untrusted networks (such as the Internet), and through many different network devices. The loss of any one of these components can cause the IPsec VPN to fail. Such potential failure points include

- Access link failure
- Remote peer failure
- Device failure
- Path failure

An access link failure could include the failure of a physical interface on any transit network device (although *the* access link is typically seen at your end of the IPsec VPN), a module that contains many interfaces, or the "cable" (electrical, optical, or wireless) that provides transport.

Failure of the remote peer is typically attributed to "the other guy." Unless you have some network management reachability into the remote site, it is difficult to determine what the exact cause of the failure is.

A device failure is typically a failure of any device between, and including, the source and destination of the IPsec VPN. In many cases, these devices are beyond your administrative control, and the reason for failure cannot be determined.

A path failure could be a routing or circuit issue in a network between the two IPsec VPN endpoints. The failure is typically outside of your administrative reach, and cannot be easily determined.

The IPsec VPN design must consider all facets of potential network failure and implement redundancy accordingly to ensure that the secure traffic continues to flow from one site to another.

Failure Mitigation

Each of the failure sources mentioned earlier can be mitigated by employing one or more redundancy mechanisms. It is important to remember that the greater the level of high availability

in the network, the greater the implementation cost. The primary failure points and some preventive solutions are as follows:

- **Access link failure**—To overcome the loss of an access link, multiple interfaces and devices can be used. A single IPsec VPN endpoint could have multiple interfaces, multiple interface cards, or multiple endpoint devices.
- **Remote peer failure**—Failure of the remote peer is mitigated in a similar manner to the way in which access link failure is mitigated. Multiple interfaces and devices can be used to survive a failure. Such duplication allows multiple IPsec VPN tunnels to securely connect the two sites, and each uses a different infrastructure.
- **Device failure**—Device failure comes in many flavors. As described for both the access link and remote peer, duplicate interfaces and devices can help overcome a local failure. However, a device failure outside of your administrative control is a challenge to correct. So rather than fix someone else's equipment, simply avoid it. Ensure that you have multiple diverse paths between endpoints in case a problem arises in the untrusted network.
- **Path failure**—A path failure is typically beyond your control. Path redundancy can be used to circumvent a path failure in an untrusted network.

It is important to consider what is truly required to achieve path redundancy. Any single point of failure should be removed from the path. Within your network, this would mean duplicate equipment and wiring. It would also imply separate and diverse paths into and out of the building. Many costly redundancy plans have been knocked out with a single swipe of a backhoe cutting the single physical path into the building.

The use of different ISPs ensures that the traffic starts in different pieces of the Internet. But it is difficult to ensure that a common circuit (from an upstream ISP) is not used "somewhere" between the source and destination points.

Failover Strategies

The best redundancy plans cannot be executed if the failure state cannot be recognized. There are two ways that IPsec failover can be executed:

- **Stateless**—In a stateless environment, redundant logical connections (IPsec VPN tunnels) are used to provide primary and backup paths. The use of the paths is determined by message exchanges between the peers, or a determination by the end devices on which path to use. The *state* of the IPsec VPN tunnels is not known. Traffic is sent across the backup tunnel if the end-to-end path has failed.

- **Stateful**—To provide a stateful failover, redundant equipment is employed. The devices used to provide stateful failover are typically identical (configuration, interfaces, operating system, and so on). These devices also communicate with each other to determine which one is the current best device.

Most redundancy plans react to a failure and send traffic on an alternate path. The overall ability to provide timely redundancy begins with the detection of a failure.

IPsec Stateless Failover

There are three primary stateless means to detect and react to a fault. The ideal reaction to a detected fault is to automatically send traffic a different way. The three failure detection methods are as follows:

- Dead peer detection (DPD)
- An IGP within GRE over IPsec
- Hot Standby Routing Protocol (HSRP) (or one of the related protocols)

The sections that follow discuss each of these methods in greater detail.

Dead Peer Detection

Dead peer detection is a configuration option during the IPsec VPN setup. DPD also offers a stateless failover from one VPN tunnel to another. This means that the routers are not keeping track of which VPN tunnels are currently alive. Instead, traffic flows through the primary tunnel until it fails, at which time a secondary tunnel is selected.

DPD has two operational modes: periodic mode and on-demand mode.

DPD periodic mode has the following characteristics:

- DPD keepalive messages are periodically sent between IPsec VPN peers.
- DPD keepalive messages are in addition to the normal IPsec rekey messages that also regularly traverse the tunnel.
- DPD keepalive messages are not sent if user data is transmitted through the VPN tunnel.
- DPD keepalive messages are used only when there is a lull in tunnel traffic.

One negative consequence of periodic DPD mode is the potentially excessive tunnel overhead. IKE already has a regular set of keepalive messages that pass through the tunnel. This keepalive mechanism is the IPsec SA rekeying messages that occur as the IPsec lifetime nears expiration.

Use of an IPsec VPN tunnel normally means that packets are encrypted at one end and decrypted at the other. The addition of DPD keepalive messages adds more encryption/decryption overhead to the VPN endpoints. However, the addition of these DPD keepalive messages provides more timely failure detection.

In contrast, DPD on-demand mode has the following attributes:

- It is the default DPD mode in a Cisco IOS device.
- DPD keepalive messages are sent only if the liveliness of the remote peer is in question. If traffic is sent to the peer, a response is expected. If such a response does not arrive, a DPD keepalive message is sent.
- DPD keepalive messages are never sent during otherwise idle tunnel moments.
- It is possible that a router might not discover a dead peer until the IKE or IPsec security association (SA) rekey is attempted.

The use of on-demand mode reduces the additional tunnel overhead that normal mode introduced. However, an alternate IPsec VPN tunnel might not be used immediately upon the failure of the primary one. This is not as bad as it may sound. If there is no traffic traveling through an IPsec VPN, and the VPN fails, there truly is no need to change to the alternate tunnel until user data arrives.

The configuration of DPD in a Cisco IOS device is simply a modification of an existing IPsec VPN setup. As already discussed, DPD uses keepalive messages to determine if the primary peer has failed, and then swaps over to a backup peer. Figure 15-1 shows a sample DPD configuration and topology.

Figure 15-1 shows how a remote site is configured with redundant IPsec VPN tunnels back to a central office using DPD. The two Cisco IOS commands that enable DPD are

```
crypto isakmp keepalive seconds [retries] [periodic ¦ on-demand]
set peer ip-address [default]
```

The **crypto isakmp keepalive** IOS command determines the mode and frequency of DPD. Remember that **periodic** mode sends DPD keepalive messages, which are continually sent to verify that the remote VPN peer is still alive. The default DPD mode is **on-demand**, which sends DPD messages only if the remote peer is believed to be dead. Default options do not appear in the configuration.

The **crypto isakmp keepalive** command has two timer options. The *seconds* option defines how often DPD keepalive messages are sent in periodic mode. The *retries* option defines how long to wait to resend DPD messages after the previous one has failed.

Figure 15-1 *DPD Configuration*

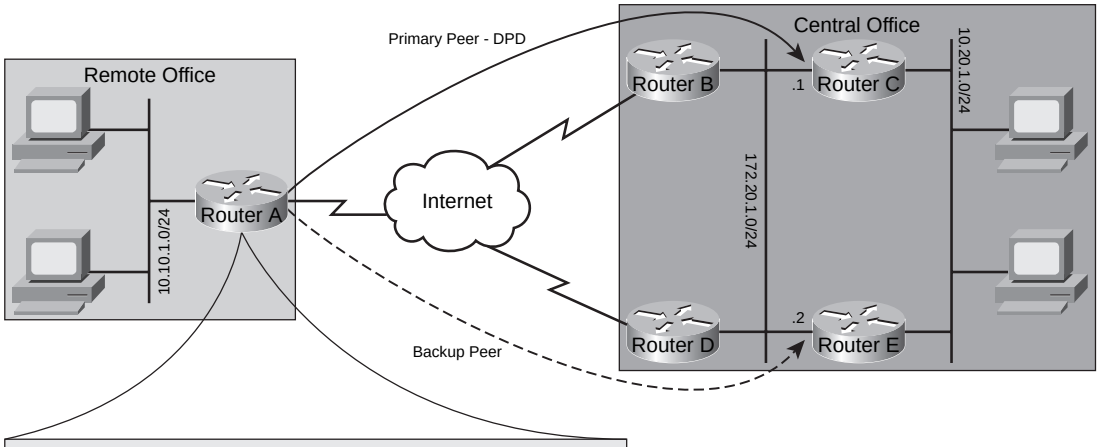


```
crypto isakmp keepalive 10 3
!
crypto ipsec transform-set to-central esp-3des esp-sha-hmac
!
crypto map central-office 10 ipsec-isakmp
 set peer 172.20.1.1 default
 set peer 172.20.1.2
 set transform-set to-central
 match address 120
!
access-list 120 permit ip 10.10.1.0 0.0.0.255 10.20.1.0 0.0.0.255
```

Figure 15-1 shows two peer configurations (with the set peer IOS commands). The primary peer (172.20.1.1) is indicated with the **default** option. This is the peer that is initially used between the remote and central offices. The secondary peer (172.20.1.2—the one without the **default** option) is not used until DPD determines that the primary peer has failed.

IGP Within a GRE over IPsec Tunnel

Chapter 14, "GRE Tunneling over IPsec," covered the use of GRE over IPsec. Remember that a normal IPsec VPN cannot transport dynamic routing protocols. A GRE tunnel is created for the routing protocol traffic, and then sent through the IPsec VPN for confidentiality and integrity.

OSPF and EIGRP have very fast convergence around failed links. The use of a backup GRE over IPsec VPN tunnel does provide redundancy, at the cost of additional IGP overhead in the VPN tunnel.

If two sites are connected with two or more GRE over IPsec tunnels, the IGP that runs across the tunnels can make very rapid routing decisions on alternative paths. Of course, it is important to create the tunnels such that there is no single point of failure in the paths. For example, if all the tunnels start on one router and end on a different router, the failure of either router eliminates all the tunnels.

HSRP

Most hosts are configured with a single gateway, or default, router. The address of this default router is typically delivered to the host during address acquisition via DHCP. However, if the gateway router fails, then all hosts that use it become isolated.

A good network design attempts to remove any single points of failure; however, such design options come at a price. The addition of a second gateway router not only costs money, but adds complexity to the network. The simple configuration of a second default gateway in the end hosts does not ensure a timely failover to the secondary gateway when needed.

It is possible to have the end hosts actually discover the gateways, or run routing protocols with the gateways. However, neither of these options is desirable for a number of reasons (administrative and processing overhead, feature support for some platforms, network security concerns).

HSRP offers the capability to use more than one router as a default gateway for end hosts. A group of routers form a logical gateway. This gateway IP address is used by the end hosts as their default gateway. A virtual MAC address is also used when the hosts broadcast (use ARP) for their default gateway.

Normally, the actual gateway IP address is configured on a single router. However, the HSRP group handles traffic destined for the logical gateway IP address. Within the group, the active router handles all packets destined for the logical IP address (and MAC address). A standby router exists to forward packets only if the active router fails.

Any number of routers can be in an HSRP group (although a large number quickly becomes impractical). There is only one active router per group (per gateway IP address). The remaining routers in the HSRP group elect the standby router. The active and standby routers periodically communicate with each other, which is how the standby router determines if the active router has failed. If the active router fails, the standby router takes control of the group and forwards traffic sent to the virtual group IP address. At this time, the remaining routers in the HSRP group elect a new standby router. Although the HSRP routers communicate with each other, this is still considered stateless VPN failover because the state of the IPsec VPN tunnels is unknown.

It is possible for one physical LAN to be home for multiple IP subnets. As such, each subnet would typically need a gateway router. With HSRP, each subnet would use a virtual standby group, where each standby group emulates a physical gateway router. HSRP groups can coexist and overlap on the same physical router. For example, one router could be the active router for one group and the standby router for another. In such a case, the router forwards traffic only for the active group. Another router forwards traffic for the other HSRP group.

Figure 15-2 shows a sample HSRP configuration and topology for the remote office. This actually shows the ultimate in redundancy, because there are two connections to the central office, and each uses a separate ISP. Because there are two physical connections, there are two different IPsec VPNs configured also. Not all remote sites are as fortunate.

Figure 15-2 *HSRP Configuration at the Remote Office*

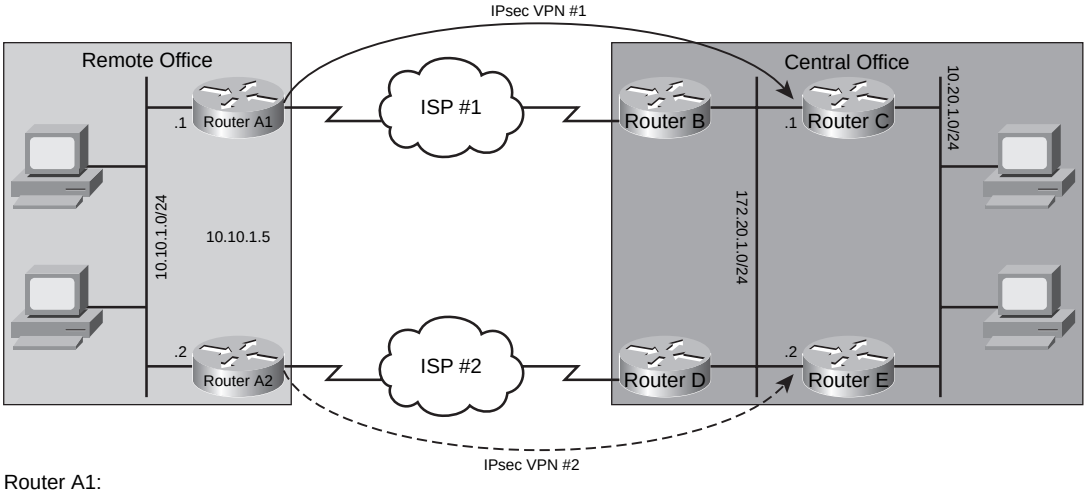

Router A1:

```
interface fastethernet 0/1
 ip address 10.10.1.1 255.255.255.0
 standby 1 ip 10.1.1.5
 standby 1 priority 150
 standby 1 preempt
```

Router A2:

```
interface fastethernet 0/1
 ip address 10.10.1.2 255.255.255.0
 standby 1 ip 10.1.1.5
```

The hosts at the remote site would use 10.10.1.5 as their default gateway. This is the HSRP group IP address (virtual IP address) between Routers A1 and A2. Router A1 is configured with a higher HSRP priority (the default is 100), which means that it will initially be the active router. The **preempt** command says that if it has a higher priority (and it does), it will regain active HSRP status if it ever fails and comes back to life.

The HSRP service provided to end hosts does not interact with the IPsec VPN configuration. For the hosts, and thus at the remote site, HSRP simply selects the active default gateway.

Figure 15-3 shows how HSRP can be used at the central office to terminate IPsec VPN connections from remote offices.

Figure 15-3 *HSRP Configuration at the Central Office*

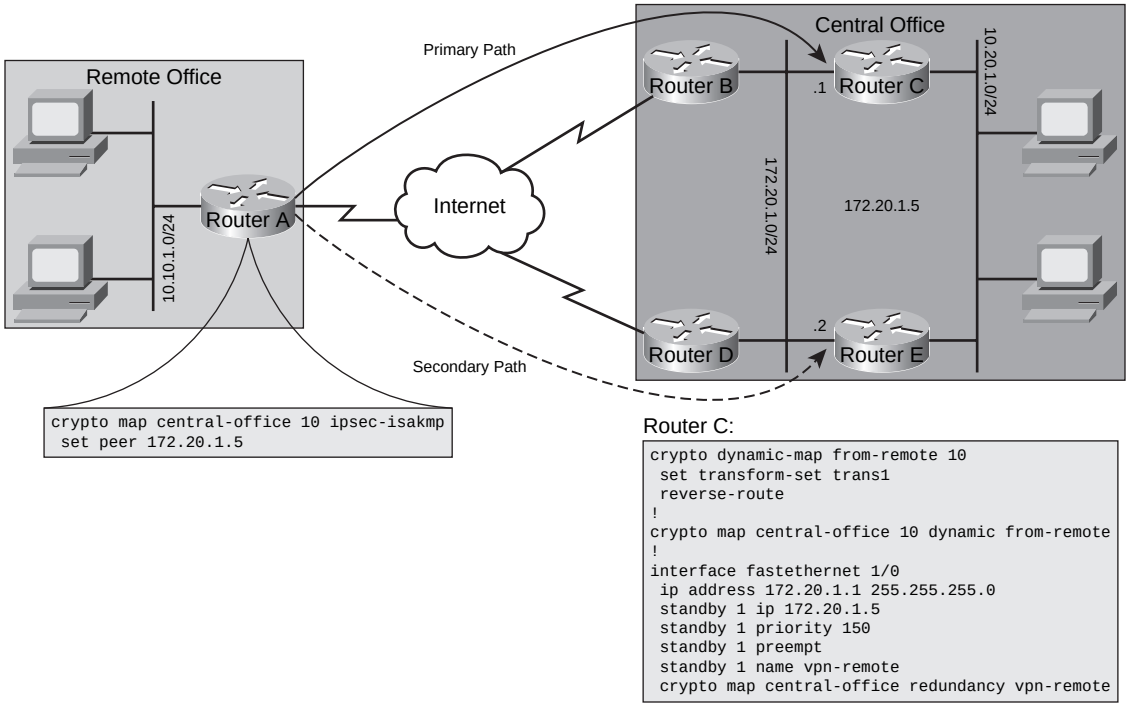

In this example, HSRP is configured between Routers C and E for the benefit of incoming IPsec VPN connections—not the hosts shown at the central office. These two routers represent the IPsec VPN headend for all remote sites. The 172.20.1.0/24 LAN is globally reachable. The remote site is configured to terminate its VPN connection to 172.20.1.5. At the central office, this IP address is actually a virtual group IP address between Routers C and E. In this example, the remote site does not benefit from as much redundancy as it does in Figure 15-2.

Figure 15-3 shows the HSRP configuration for Router C. The HSRP configuration for Router E would be very similar. A separate HSRP group can be configured between Router C and Router E to offer the hosts at the central office a redundant gateway. Such a configuration would be similar to the one shown in Figure 15-2.

The interface **crypto map** statement indicates that the HSRP group **vpn-remote** provides redundancy. This HSRP group name is defined on the interface. The central office is also configured with a dynamic crypto map. This means that any remote office (source IP address) can initiate a VPN connection with the central office. It is possible that remote offices that use DSL or cable connectivity to the Internet do not have fixed external IP addresses, and thus cannot be statically configured at the central office.

It is important to remember that if Router C is active and fails, the IPsec VPN to it will also drop. The remote site will reestablish an IPsec VPN to the same remote IP address (the HSRP group IP address—172.20.1.5), which is then handled by Router E. When Router C comes back to life, the IPsec VPN again drops (because Router C becomes active and preempts Router E) and is reestablished to Router C.

IPsec Stateful Failover

IPsec stateful failover typically requires a set of identical equipment so that failover can occur, and requires some continuous exchange of data between the devices to track the state of the IPsec VPNs (SA information). This also implies that there are multiple active IPsec VPN tunnels. Thus, the failure of one path can immediately switch the traffic to an alternate and operational IPsec VPN.

As described in the previous section on IPsec stateless failover, failover typically involves the creation of a new IPsec VPN tunnel when the first tunnel fails or becomes unreachable. Thus, there is a period of time during which secure connectivity does not exist. A stateful environment eliminates the temporary inability to communicate securely.

Stateful failover is accomplished through active (primary) and backup (secondary) devices. This concept is similar to how HSRP operates; however, SA information is also being maintained. The backup router automatically forwards traffic upon the failure (planned or unplanned) of the primary path. The switch from the primary to the backup is transparent to both the users and the remote IPsec VPN peer.

IPsec stateful failover uses two protocols for proper and continual operation:

- **HSRP**—Monitors both the inside and outside interfaces. If either goes down, the entire router is deemed unworthy and ownership of the IKE and IPsec SA processes is passed to the standby router. When this transition occurs, the standby router becomes the active HSRP router.
- **Stateful Switchover (SSO)**—Shares the IKE and IPsec SA information between the active and backup routers. At any time, either router knows enough to be the active IPsec VPN router.

There are some limitations/restrictions that exist when IPsec stateful failover is deployed. Some of the more important points to understand are as follows:

- Both the active and standby devices must run an identical Cisco IOS release.
- The active and standby devices must be connected via LAN ports, either directly or through a switch. WAN interfaces are not supported.

- Both the inside and outside interfaces must be connected via LAN ports.
- Only "box-to-box" failover is supported. Intrachassis (card-to-card) failover is not currently supported.
- Load balancing is not supported. Only one device in a redundancy group can be active at any time.
- IKE keepalive messages are not supported. DPD and periodic DPD are supported.
- Stateful failover of Layer 2 Tunneling Protocol (L2TP) is not supported.
- IPsec idle timers are not supported.

Because IPsec stateful failover uses HSRP and SSO, both protocols must be properly configured. Figure 15-4 shows the configuration necessary at the central office for the topology illustrated.

Figure 15-4 *IPsec Stateful Failover*

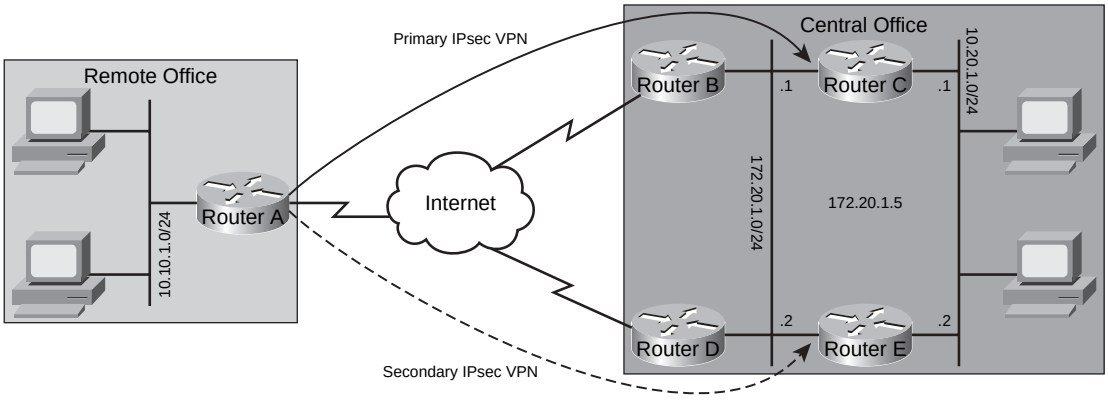

Router C:

```
crypto dynamic-map from-remote 10
 set transform-set trans1
 reverse-route
!
crypto map central-office 10 ipsec-isakmp dynamic from-remote
!
interface fastethernet 1/0
 ip address 172.20.1.1 255.255.255.0
 standby 1 ip 172.20.1.5
 standby 1 priority 150
 standby 1 preempt
 standby 1 name vpn-remote
 crypto map central-office redundancy vpn-remote stateful
```

```
!
redundancy inter-device
 scheme standby vpn-remote
!
ipc zone default
 association 1
 protocol sctp
  local-port 12321
 local-ip 10.20.1.1
 retransmit-timeout 300 10000
 path-retransmit 10
 assoc-retransmit 20
  remote-port 12321
 remote-ip 10.20.1.2
```

The crypto map and interface configurations for Router C in Figure 15-4 are nearly identical to those from Figure 15-3. One minor addition is the term **stateful** to the crypto map on the interface. This permits the use of SSO to perform stateful failover. The HSRP configuration is the same as

before. Router E would have a similar configuration as Router C to complete the stateful configuration.

The follow-on configuration box shows the IOS commands needed to enable SSO. The **redundancy inter-device** command configures redundancy and enters inter-device configuration mode. Currently, the only scheme supported is **standby**. Note that the name of the standby, *vpn-remote*, must match the standby group name defined with the crypto map on the interface.

The next block of commands configures the inter-device communication protocol (IPC) between the two gateways. The **ipc zone default** command initiates the communication link between active and standby routers. The subcommand **association** creates an association between the active and standby routers and uses the Stream Control Transmission Protocol (SCTP) as the transport protocol.

Within SCTP, the local and remote SCTP ports and IP addresses are defined. The *local-port* defined on this router must match the *remote-port* configured on the peer router. Also, the *local-ip* and *remote-ip* addresses should point to physical interface IP addresses and *not* to virtual IP addresses.

The **path-retransmit** command defines the number of SCTP retries before an attempt to create an SCTP session fails, and the **retransmit-timeout** command defines the maximum amount of time that SCTP waits before retransmitting data.

WAN Backed Up by an IPsec VPN

This chapter has focused on how to ensure that the loss of one IPsec VPN can be easily recovered by a second. Both stateful and stateless methods were examined. IPsec VPN tunnels can also be used to back up "normal" WAN connections.

Most of Part III, "IPsec VPNs," of this book deals with IPsec VPNs, which offer confidentiality to data as it passes from one site to another. A "normal" WAN connection is simply a PVC, such as a Frame Relay or ATM link between sites. No confidentiality or integrity is offered for such connections. However, if such a connection should fail, there is no reason that the traffic that does not expect protection cannot travel through the IPsec VPN.

The assumption is that both a "normal" WAN connection and an IPsec VPN link exist between two sites. The WAN connection is some sort of provider-based PVC, while the IPsec VPN travels across the untrusted Internet. As already explained in Chapter 13, an IPsec VPN can be statically configured to know which traffic is permitted to travel through it (interesting traffic). It has also been shown how to configure dynamic routing protocols across the IPsec VPN through the use of GRE over IPsec (refer to Chapter 14).

The “normal” WAN connection exchanges dynamic routing updates via OSPF or EIGRP. When this link fails, both sides realize the loss very quickly, due to the fast convergence time of both OSPF and EIGRP. There are two ways that routers on either end can decide to forward traffic over the IPsec VPN link.

The first solution is to ensure that the same dynamic routing protocol is also configured to run across the IPsec VPN, which is accomplished with GRE over IPsec. The IPsec VPN connection should be used only after the “normal” WAN connection fails. To ensure this, the EIGRP interface delay or OSPF cost can be adjusted to make the dynamic IPsec VPN routes less favorable than the “normal” WAN ones.

A second way to route traffic through the IPsec VPN upon WAN failure is to use floating static routes. A floating static route is a manually configured route with a high administrative distance (AD). Due to the high AD, the static route is not chosen as the best available path until the dynamic routes (with lower ADs) have evaporated. The loss of such dynamic routes occurs as a result of either path failure to the prefix or failure of the prefix itself.

With either of these solutions, the IPsec VPN is used primarily for specific traffic. Upon failure of the WAN connection, all traffic is permitted to temporarily travel through the VPN. When the primary WAN path has been reestablished, the normal WAN traffic returns to its desired connection.

Foundation Summary

Potential network failure points and some of the ways to mitigate them include:

- **Access link**—Use multiple interfaces and devices.
- **Remote peer**—Use multiple interfaces and devices.
- **Device failure**—Use duplicate interfaces and devices to help overcome a local failure. Having multiple diverse paths between endpoints helps avoid misbehaving devices beyond your administrative control.
- **Path failure**—Use path redundancy to circumvent a path failure in an untrusted network.

Two ways that IPsec failover can be executed are as follows:

- **Stateless**—In a stateless environment, redundant logical connections (IPsec VPN tunnels) are used to provide primary and backup paths. The use of the paths is determined by message exchanges between the peers, or a determination by the end devices on which path to use. The "state" of the IPsec VPN tunnels is not known. Traffic is sent across the backup tunnel if the end-to-end path has failed.
- **Stateful**—In a stateful environment, redundant equipment is employed. The devices used to provide stateful failover are typically identical (configuration, interfaces, operating system, and so on). These devices also communicate with each other to determine which one is the current best device.

Three primary stateless means to detect and react to a fault are as follows:

- Dead peer detection (DPD)
- An IGP within GRE over IPsec
- HSRP (or one of the related protocols)

DPD has two operational modes:

- DPD periodic mode, which has the following characteristics:
  - — DPD keepalive messages are periodically sent between IPsec VPN peers.
  - — DPD keepalive messages are in addition to the normal IKE keepalive messages that also regularly traverse the tunnel.

- DPD keepalive messages are not sent if user data is transmitted through the VPN tunnel.
- DPD keepalive messages are used only when there is a lull in tunnel traffic.

- DPD on-demand mode, which has the following attributes:
  - It is the default DPD mode in a Cisco IOS device.
  - DPD keepalive messages are sent only if the liveliness of the remote peer is in question. If traffic is sent to the peer, a response is expected. If one does not arrive, then a DPD keepalive is sent.
  - DPD keepalive messages are never sent during otherwise idle tunnel moments.
  - It is possible that a router might not discover a dead peer until the IKE or IPsec SA rekey is attempted.

The two Cisco IOS commands that enable DPD are

```
crypto isakmp keepalive seconds [retries] [periodic | on-demand]
set peer ip-address [default]
```

OSPF and EIGRP have very fast convergence around failed links. The use of a backup GRE over IPsec VPN tunnel does provide redundancy, but at the cost of additional IGP overhead in the VPN tunnel.

HSRP uses virtual MAC and IP addresses as default gateway addresses for end hosts.

An HSRP group consists of two or more routers. Each HSRP group is intended for one IP subnet.

One router can participate in more than one HSRP group.

In an HSRP group, there is only one active router and one standby router. Only the HSRP active router forwards traffic.

Typical host-based HSRP interface commands include

- **standby** *group* **ip** *virtual-IP-address*—Defines the HSRP group ID and virtual IP address, which is the same for all group members.
- **standby** *group* **priority** *priority-#*—Defines the HSRP priority for this router (the default is 100).
- **standby** *group* **preempt**—Causes this router to regain active status if it has the highest priority.

Stateless IPsec VPN HSRP interface commands include

- **standby** *group* **name** *group-name*—Defines a name for the HSRP group that can be added to the crypto map.
- **crypto map** *map-name* **redundancy** *group-name*—Defines the HSRP group that provides redundancy for this crypto map.

IPsec stateful failover uses two protocols for proper and continual operation: HSRP and SSO.

Stateful IPsec VPN interface commands include

- **standby** *group* **ip** *virtual-IP-address*—Defines the HSRP group ID and virtual IP address, which is the same for all group members.
- **standby** *group* **priority** *priority-#*—Defines the HSRP priority for this router (the default is 100).
- **standby** *group* **preempt**—Causes this router to regain active status if it has the highest priority.
- **standby** *group* **name** *group-name*—Defines a name for the HSRP group that can be added to the crypto map.
- **crypto map** *map-name* **redundancy** *group-name* **stateful**—Defines the HSRP group that provides stateful redundancy for this crypto map.

SSO global commands include

- **redundancy inter-device**—Enables SSO.
- **scheme standby** *group-name*—Maps an HSRP group to the stateful failover.
- **ipc zone default**—Defines the inter-device communications protocol parameters for coordination between the active and standby routers. Local and remote ports and local and remote IP addresses must be defined on both routers.

There are two ways that an IPsec VPN link can be used to back up a typical WAN link: IGP via GRE over IPsec and floating static routes.

Q&A

The questions and scenarios in this book are designed to be challenging and to make sure that you know the answer. Rather than allowing you to derive the answers from clues hidden inside the questions themselves, the questions challenge your understanding and recall of the subject.

Hopefully, mastering these questions will help you limit the number of exam questions on which you narrow your choices to two options, and then guess.

You can find the answers to these questions in Appendix A. For more practice with exam-like question formats, use the exam engine on the CD-ROM.

1. What are the potential failure points in a network?
2. What are some of the ways to overcome an access link failure?
3. What are the three forms of stateless IPsec failover?
4. Which DPD mode is the default in a Cisco IOS device?
5. What is a negative consequence of periodic DPD mode?
6. What IOS command enables DPD?
7. Which routing protocols should be used within the GRE over IPsec tunnels to permit fast convergence around failed links?
8. How do the HSRP active and standby routers work together?
9. If an IPsec VPN terminates on an HSRP virtual IP address, and the active router fails, what happens to the VPN?
10. What two protocols are used to provide IPsec stateful failover?
11. If dynamic routing is used to permit an IPsec VPN to back up a normal WAN connection, what must be done?
12. What is a floating static route?

Exam Topic List

This chapter covers the following topics that you need to master for the CCNP ISCW exam:

- **Cisco Easy VPN Components**—Describes the constituent elements of the Easy VPN solution
- **Easy VPN Connection Establishment**—Describes the process of connecting to another site with Easy VPN
- **Easy VPN Server Configuration**—Describes the Easy VPN Server configuration process
- **Monitoring the Easy VPN Server**—Describes possible options available for connection monitoring with Easy VPN Server
- **Troubleshooting the Easy VPN Server**—Describes the basic process and options available in troubleshooting Easy VPN Server