CHAPTER 6

Enterprise Internet Edge

The Internet edge is the network infrastructure that provides connectivity to the Internet and that acts as the gateway for the enterprise to the rest of the cyberspace. The Internet edge serves other building blocks—referred to by Cisco as *Places-in-the-network* (PINs)—that are present in a typical enterprise network. This modular building-block approach enables flexibility and customization in network design to meet the needs of customers and business models of differing sizes and requirements.

Figure 6-1 shows the Internet edge infrastructure as part of an enterprise network. The Internet edge infrastructure serves most areas of the enterprise network, including the data center, campus, and remote branches. The proper design and implementation of the Internet edge infrastructure is crucial to ensure the availability of Internet services to all enterprise users. The Internet edge infrastructure includes the following functional elements:

- *Service Provider (SP) Edge*

  This border part of the Internet edge infrastructure consists of routers that interface directly to the Internet. Internet-facing border routers peer directly to the Internet SP. Careful consideration must be made to routing design, redundancy, and security of these border routers.

- *Corporate Access and DMZ*

  One of the major functions of the Internet edge is to allow for safe and secure Internet access by corporate users while providing services to the general public. The firewalls in this module secure these functions through implementation enforcement of stateful firewall rules and application-level inspection. Users at the campuses may access email, instant messaging, web browsing, and other common services through the Internet edge firewalls. Optionally, the same infrastructure may serve users at the branches that are mandated to access the Internet over a centralized connection. Public-facing services, such as File Transfer Protocol (FTP) servers and websites, can be provided by implementing a de-militarized zone (DMZ) within this network domain. The web application firewall is another appliance that protects web servers from application-layer attacks (such as XML). The web application firewall also resides in the DMZ infrastructure and provides primary security for Hypertext Transfer Protocol (HTTP)-based and E-commerce applications.

- *Remote Access*

  The remote access infrastructure that provides corporate access to remote users through protocols such as Secure Socket Layer (SSL) Virtual Private Networking (VPN) and Easy VPN.

- *Edge Distribution*

  The edge distribution infrastructure provides the interface for the Internet edge network devices to the rest of the enterprise network. Appliances, such as the Web Security Appliances (WSA), reside in this part of the network. Within the edge distribution infrastructure, you can also implement an Intrusion Prevention Appliance (IPS) to guard against worms, viruses, denial-of-service (DoS) traffic, and directed attacks.

- *Branch Backup*

  Some branches may adopt an Internet connection to provide a backup link to a WAN network. This backup functionality may be performed by using dedicated appliances, such as a Cisco ASR 1000 Series router. The branch backup functionality is implemented within at Internet WAN edge block in the Enterprise WAN edge module.

The Internet edge module provides many of the essential Internet-based services used in enterprise networking environments (see Figure 6-1). Providing these services in a secure manner is essential to business continuity and availability. The best practices for securing these services in the context of Internet edge are presented in this chapter.

Figure 6-1 Internet Edge Infrastructure as part of an Enterprise Network

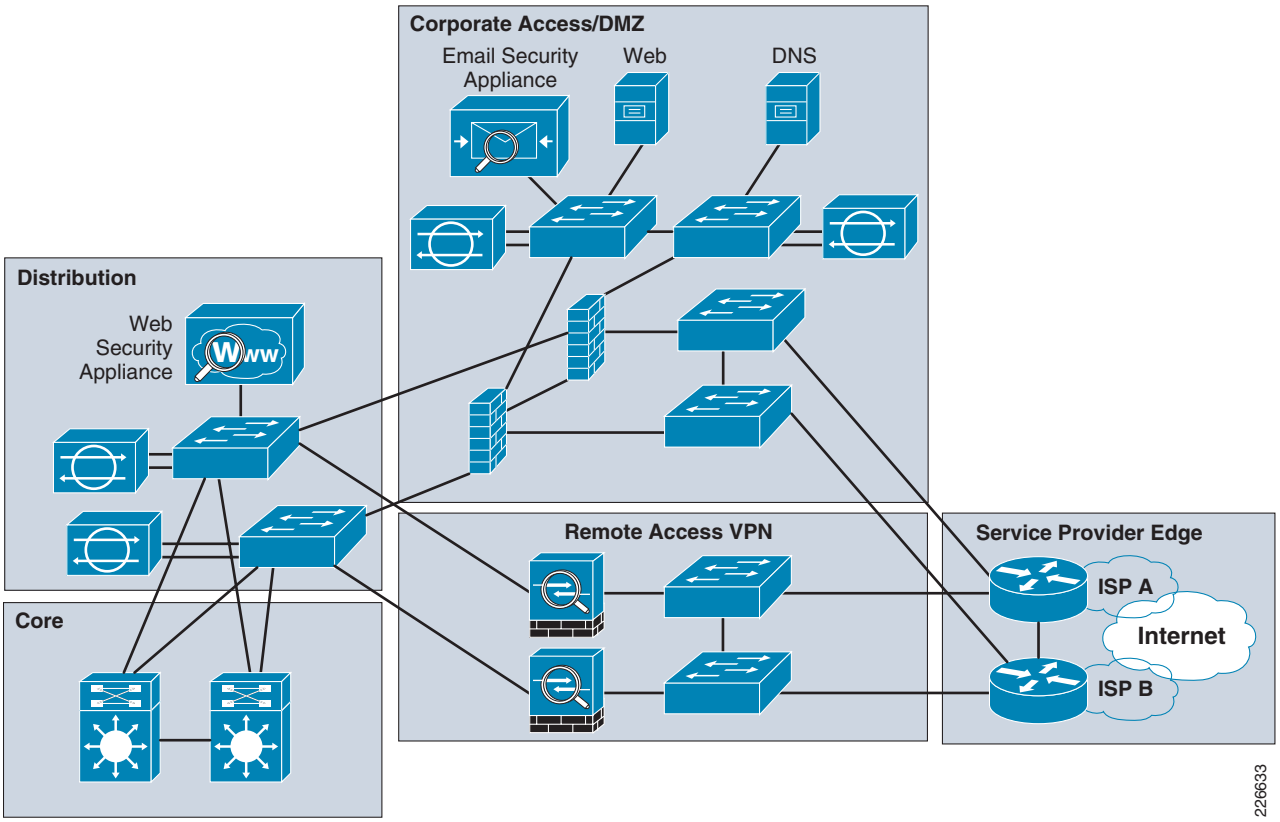

Key Threats in Internet Edge

The Internet edge is a public-facing network infrastructure and is particularly exposed to large array of external threats. Some of the expected threats are as follows:

- Denial-of-service (DoS), distributed DoS (DDoS)
- Spyware, malware, and adware
- Network intrusion, takeover, and unauthorized network access
- E-mail spam and viruses
- Web-based phishing, viruses, and spyware
- Application-layer attacks (XML attacks, cross scripting, and so on)
- Identity theft, fraud, and data leakage

Design Guidelines for the Enterprise Internet Edge

This section focuses on the overall design of the Internet edge module in the SAFE design. The Internet edge network can be divided into several functional blocks. Each functional block has its own design and security considerations:

- *Service Provider Edge*—This block is composed by the Internet-facing border routers. The primary function of the border routers is to route traffic between the organization's network and the Internet. These routers also act as the first line of defense against external attacks. The SAFE design accommodates a redundant Internet connection. To that end, the two border routers at the edge connect to the Internet through dual Internet Service Providers (ISP)—ISP-A and ISP-B, as shown in Figure 6-1. The two border routers provide redundancy and run external Border Gateway Protocol (eBGP). The eBGP allows efficient policy-based routing and prevents leakage of routes from one ISP to another. The border routers also run Performance Routing (PfR) to optimize traffic flows and improve performance. With PfR enabled, load balancing between border routers is possible without the need for storing the full Internet routing on the border routers. For outgoing traffic to the Internet, PfR also enables the routers to make intelligent decisions in choosing an optimized path based on the traffic characteristics of each ISP. This improves the overall performance of outgoing Internet traffic. For PfR implementation details, refer to *Transport Diversity: Performance Routing (PfR)*. The border routers should be secured following the device hardening best practices outlined in Chapter 2, "Network Foundation Protection."
- *Corporate Access and DMZ* —A pair of firewalls provide stateful access control and deep packet inspection. These firewalls are deployed to protect the organization's internal resources and data from external threats by preventing incoming access from the Internet; to protect public resources served by the DMZ by restricting incoming access to the public services and by limiting outbound access from DMZ resources out to the Internet; and to control user's Internet-bound traffic. To that end, firewalls are configured to enforce access policies, keep track of connection status, and inspect packet payloads. The firewalls are configured in active/standby mode for redundancy purposes. The DMZ hosts services such as the E-mail Security Appliance, HTTP, Domain Name System (DNS), and FTP. The web application firewall also resides in the DMZ. The web application firewall provides perimeter security for application-based attacks that the firewall cannot guard against and can provide safeguards for key applications, such as business-to-business transactions. In most cases the data center implements its own web application firewall. The web application firewall on the DMZ can provide the first line of defense for commerce applications and protect any web servers on the DMZ from application-layer attacks. IronPort Email Security Appliance (ESA) may be deployed at the DMZ to protect email communications.

- *Remote Access VPN*—One essential function of the Internet-edge module is to provide secure access to remote workers. Many different approaches can be taken, depending on particular requirements and policies within the enterprise. Access for remote clients can be implemented using SSL VPN with thin clients. Clients in this case are only allowed access to specific HTTP services within the enterprise. This is in contrast to full client remote access in which clients have full access to all services within the enterprise and experience the same level of service as internal corporate users. It is recommended that two separate firewalls are used to provide remote access functionality. Although a single pair of firewalls could be leveraged for both remote access and corporate access, it is a good practice to keep them separate. Remote access may be further secured by requiring user authentication and authorization, and enforcing granular per user or per group access control policies.
- *Edge Distribution*—The Enterprise Internet Edge module implements a layer of distribution switches that aggregate services common to the various blocks present in the module. The distribution switches reside in the inside network of the firewalls and connect to the core switches, making it accessible to other parts of the enterprise network. URL filtering, content inspection, and intrusion prevention are examples of services that may be aggregated at the distribution layer. Ironport Web Security Appliance (WSA) may be deployed at this layer to enforce acceptable use policies for Web access, content inspection, and to block malware, spyware and other threats. The WSA is a web proxy and sits logically in the path between corporate users and the Internet edge firewalls. Proper placement and configuration of these appliances provides secure web access, content security and threat mitigation for web services. Cisco Intrusion Prevention Systems (IPS) may be implemented in this part of the Internet edge infrastructure. Logically, the IPSs are placed between the firewall and core routers, protecting the enterprise from threats originating from campus and remote users.

Edge Distribution Layer

The position of the edge distribution layer within the Internet edge network is shown Figure 6-2.

Figure 6-2 Edge Distribution in Internet Edge

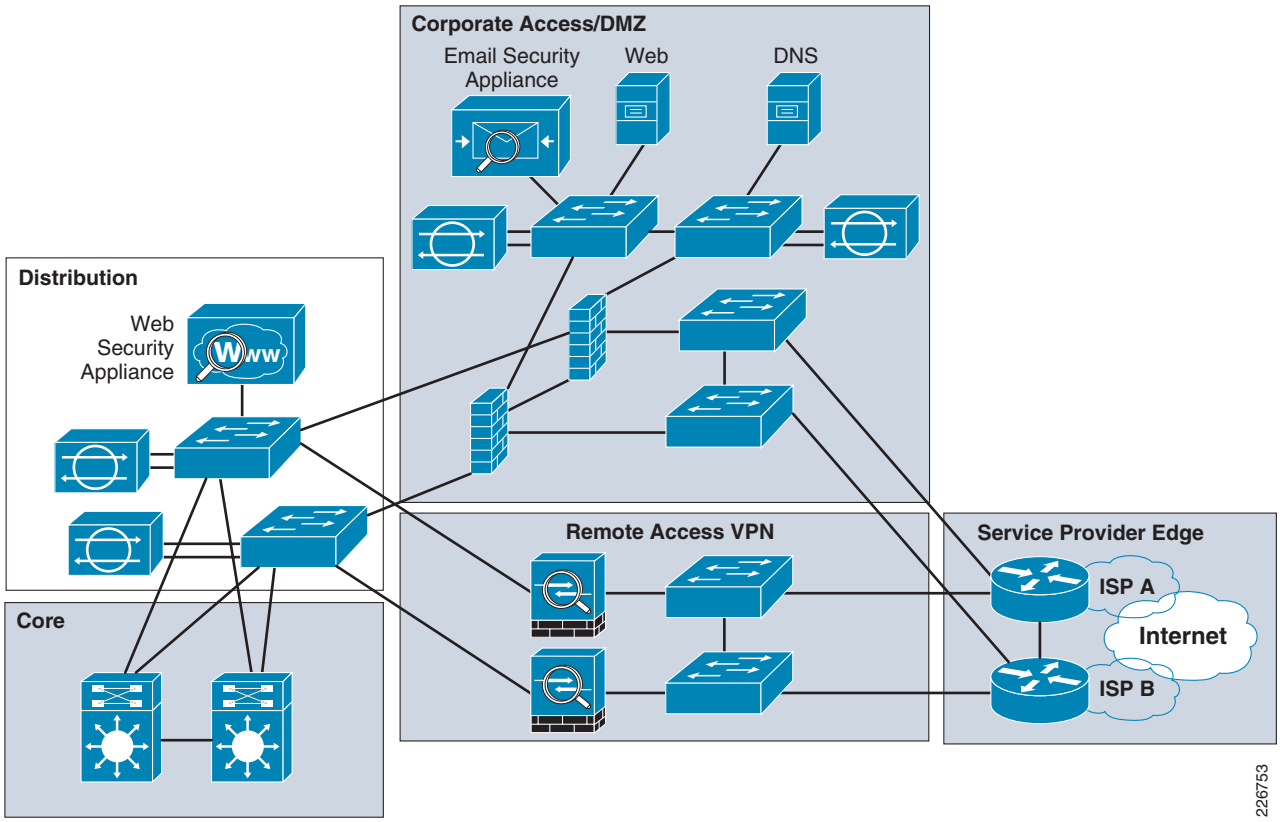

This section addresses the following edge distribution topics:

- Design Guidelines and Best Practices
  - Infrastructure Protection Best Practices
  - Internet Edge Cisco IPS Design Best Practices

Design Guidelines and Best Practices

The Internet edge *distribution layer* refers to the part of the network that aggregates common services used by the various blocks in the Internet Edge module, that resides within the inside network, and that is adjacent to the core network. Common services include web security with the WSA, and intrusion protection with Cisco IPS. It is a good practice to deploy WSA at the distribution layer, as close to the clients as possible. Per contrary, the ESA should be deployed as close to the Internet as possible with a reasonable level of firewall protection (i.e., within the DMZ). The deployment of WSA and ESA are discussed in detail in the "E-mail and Web Security" section on page 6-15. Other functions covered in this section include connectivity and routing to and from the core and implementation of the Cisco IPS appliances.

Infrastructure Protection Best Practices

Infrastructure protection plays an important role in the Internet edge distribution block. The following best practices are recommended:

- All infrastructure protection hardening, such as management access control lists (ACL), authentication, control plane policing, or Layer-2 hardening, must be implemented on the *inner* switches.
- Routing protocols between switches and Cisco ASAs and core routers must be authenticated.
- Use separate interfaces for management of the WSA.
- Disable unnecessary services, such as Telnet, HTTP, and the like on the data interfaces for the WSA in order to prevent even inside corporate users from taking advantage of open ports.

Internet Edge Cisco IPS Design Best Practices

The Cisco SAFE Internet edge design leverages the Cisco IPS to provide protection against threats originating from both inside and outside the enterprise. When implemented in inline mode, the Cisco IPS inspects all transit traffic and automatically blocks all malicious packets. The Cisco IPS can also be deployed in promiscuous mode, in which the sensor does not reside in the traffic path. In promiscuous mode, the Cisco IPS is able to identify and generate alarms whenever malicious traffic is seen, but the sensor cannot block the attack in real-time by itself. To ensure adequate threat visibility and detection, Cisco IPS sensors must be maintained with the latest signature database. This can be automated by using CSM.

When deployed in inline mode, the Cisco IPS sensor is configured to bridge traffic between interface or VLAN pairs. The sensor inspects the traffic as it is bridged between the interfaces or VLANs. Figure 6-3 shows the placement of Cisco IPS in the context of Internet edge infrastructure.

Figure 6-3 Internet Edge with Integrated Cisco IPS

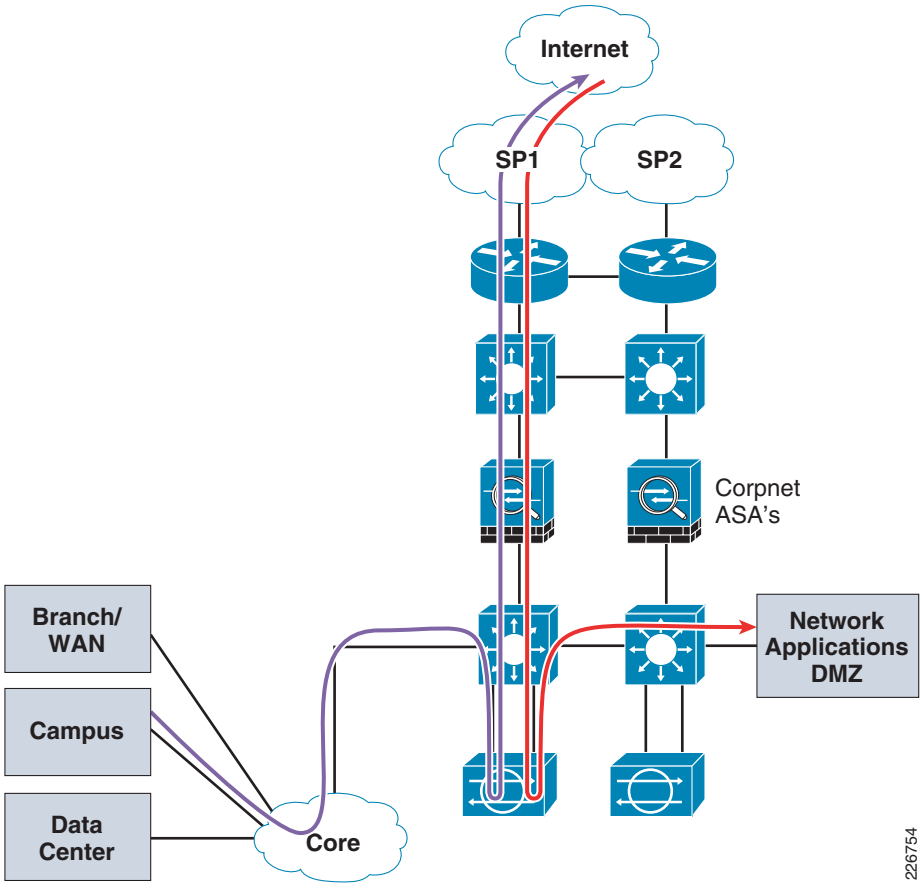

By implementing the Cisco IPS at the distribution layer, the sensors can inspect traffic from all sources, whether from the Internet, remote-access clients, or corporate users. In addition, traffic destined to the DMZ or corporate users can be monitored. Figure 6-3 shows how the Cisco IPS can inspect traffic from corporate users or from users accessing public-facing services at the DMZ.

The deployment framework for the Cisco IPS is as follows:

- The Cisco IPS is logically located between the edge firewalls and the distribution switches.
- Cisco IPS is configured to enable it to send alarms to a CS-MARS appliance. CSM and Cisco Intrusion Detection System Device Manager (IDM) GUI interfaces can be used to monitor the Intrusion Detection System (IDS). CSM and CS-MARS are located in the out-of-band management network.
- Two Cisco IPS devices are used for redundancy.
- Because the Cisco IPS loses visibility into the traffic flows with asymmetrical data paths, it is a good practice to ensure symmetrical paths by fine tuning spanning tree parameters and by implementing the firewalls in active/standby mode. With correct tuning of spanning tree and by firewalls implemented in standby/active mode, a single Cisco IPS is actively inspecting traffic at any point in time while the other is idle.

Redundancy can be achieved very easily when using Cisco ASAs in active/standby mode. In this case, only a single Cisco IPS is actively monitoring traffic for both directions at any time. If a link fails, the other Cisco IPS inspects traffic.

Corporate Access/DMZ Block

The location of the corporate access/DMZ network infrastructure within the Internet edge network is shown in Figure 6-4.

Figure 6-4 Corporate Access/DMZ in Internet Edge

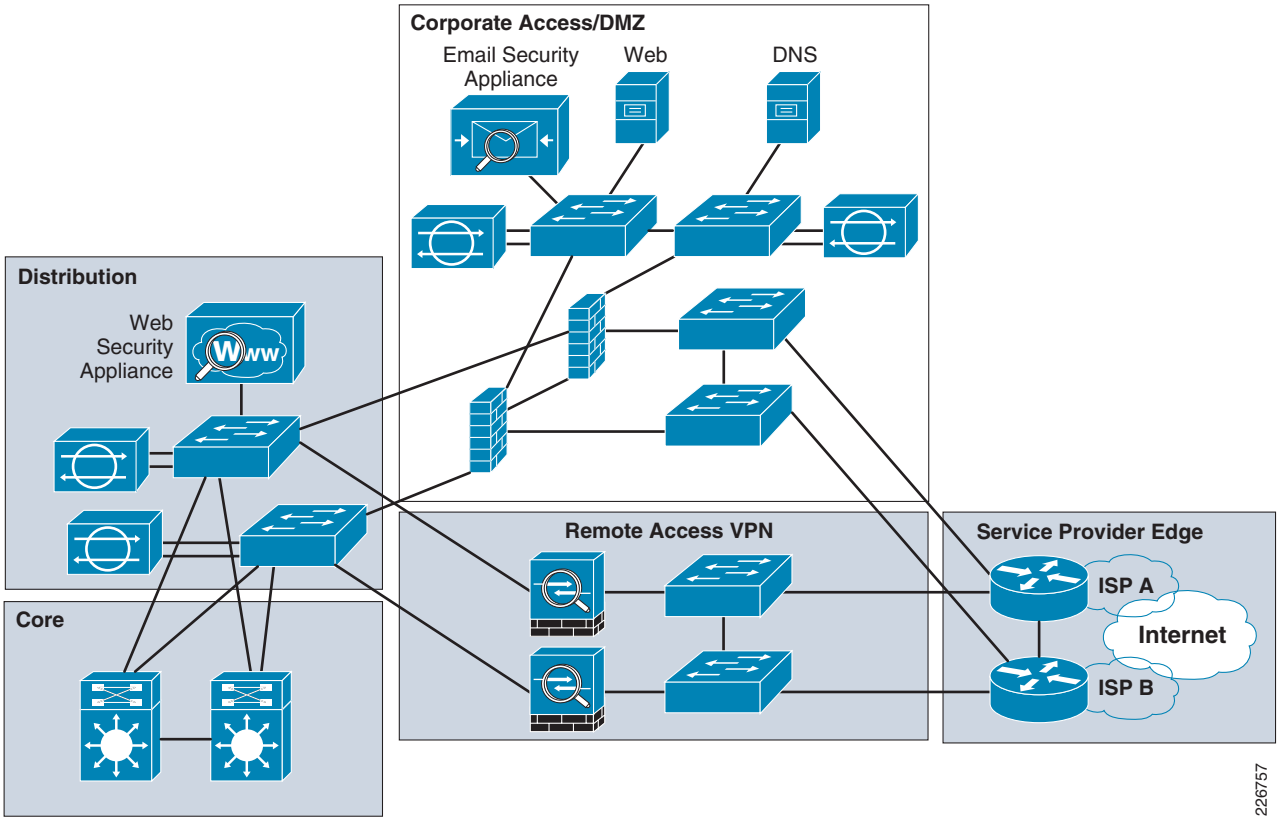

Corporate access/DMZ design is an essential aspect of the overall Internet edge design. Most enterprise customers must provide Internet access for all employees. However, this comes with security challenges. The challenge is to provide Internet access, but at the same time block malicious traffic from entering the corporate network. The first step is to deploy a firewall with proper policies configured.

The design considerations for the Cisco Application Control Engine (ACE) Web Application Firewall appliance is covered in this section. The Cisco ACE Web Application Firewall provides perimeter security functionality and protection for public-facing, web-based services located within the DMZ.

This section addresses three key topics:

- How to provide corporate access
- How to implement the firewall policy
- How to implement the Cisco ACE Web Application Firewall at the DMZ

Design Guidelines for Corporate Access/DMZ Block

The corporate access policies are enforced by Internet edge firewalls. Two Cisco ASAs are used in order to provide redundancy. They are used in active/standby mode. This simplifies Cisco IPS deployment and ensures that no traffic loss occurs in the event of a failover.

The key objectives of firewall requirements are as follows:

- All corporate users must be able to access the Internet.
- All HTTP/HTTPS traffic must pass through the WSA.
- Only web, E-mail, and some Internet Control Message Protocol (ICMP) traffic are allowed into the network.
- Cisco ASAs should be hardened.
- Cisco ASAs should be configured for redundancy.
- The Cisco ACE Web Application Firewall serves all web servers on the DMZ and all public addresses of the web servers must point to the Cisco ACE Web Application Firewall.
- Secure device access by limiting accessible ports, authentication for access, specifying policy for permitable action for different groups of people, and proper logging of events.
- Disable Telnet and HTTP; allow only secure shell (SSH) and HTTPS.
- Secure firewall routing protocols by implementing Message Digest 5 (MD5) authentication.
- Enable firewall network telemetry functionality by using features such as Network Time Protocol (NTP), logging, and NetFlow.

Figure 6-5 illustrates traffic flow through a firewall in a corporate access environment.

Figure 6-5 Traffic Flow for Typical Corporate Access

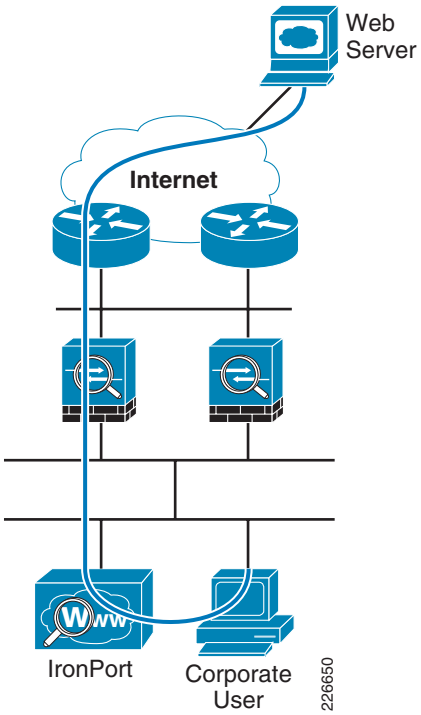

As shown in Figure 6-5, all the corporate users should pass through the WSA to reach the Internet. The Cisco ASA should not allow any web traffic to go out that does not originate from the WSA, with the exception of the ESA, Cisco Security MARS, and CSM that need to access the Internet for updates. The different logical interfaces on the Cisco ASA can be used to separate the DMZ, SP-facing interfaces, and the inside corporate infrastructure. See Figure 6-6.

Figure 6-6 Cisco ASA Physical Interface Layout

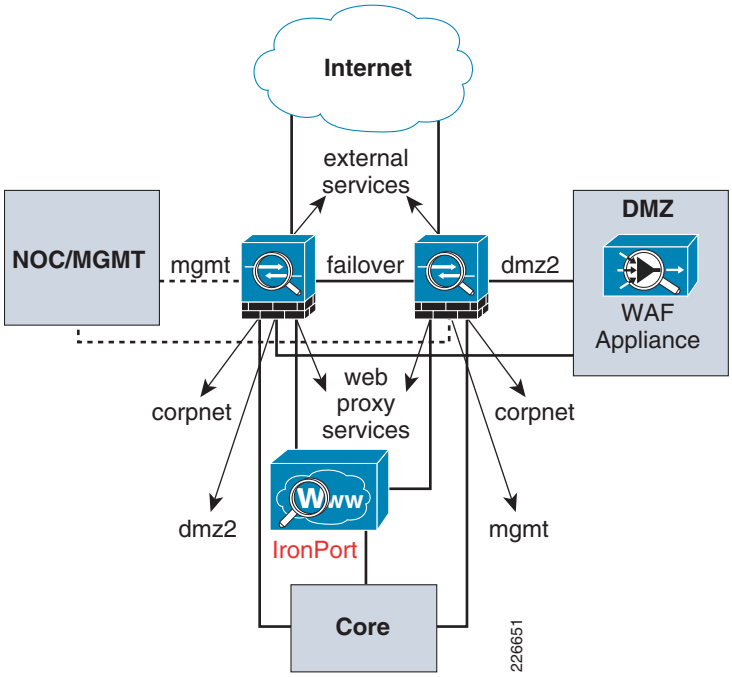

The following lists the importance of each particular interface:

- *management*—This interface is used for management traffic, including AAA, HTTPS, and so on.
- *dmz2*—This interface is used to host the web servers and web application firewall.
- *emailservices*—This interface is used to host the IronPort ESA.
- *corpnet*—This is the gateway interface for all the corporate users.
- *Failover Interface*—This is the interface used to facilitate communication and status between standby/active firewalls.
- *externalservices*—This is the interface connected to the outside world (which in this scenario is connected to the border routers).

The Cisco ASDM management tool can be used to verify firewall rules, monitor events, and configure the Cisco ASAs. Figure 6-7 and Figure 6-8 illustrate information available through the Cisco ASDM.

Figure 6-7 Cisco ASDM Screen Capture—Device Information and Status Page

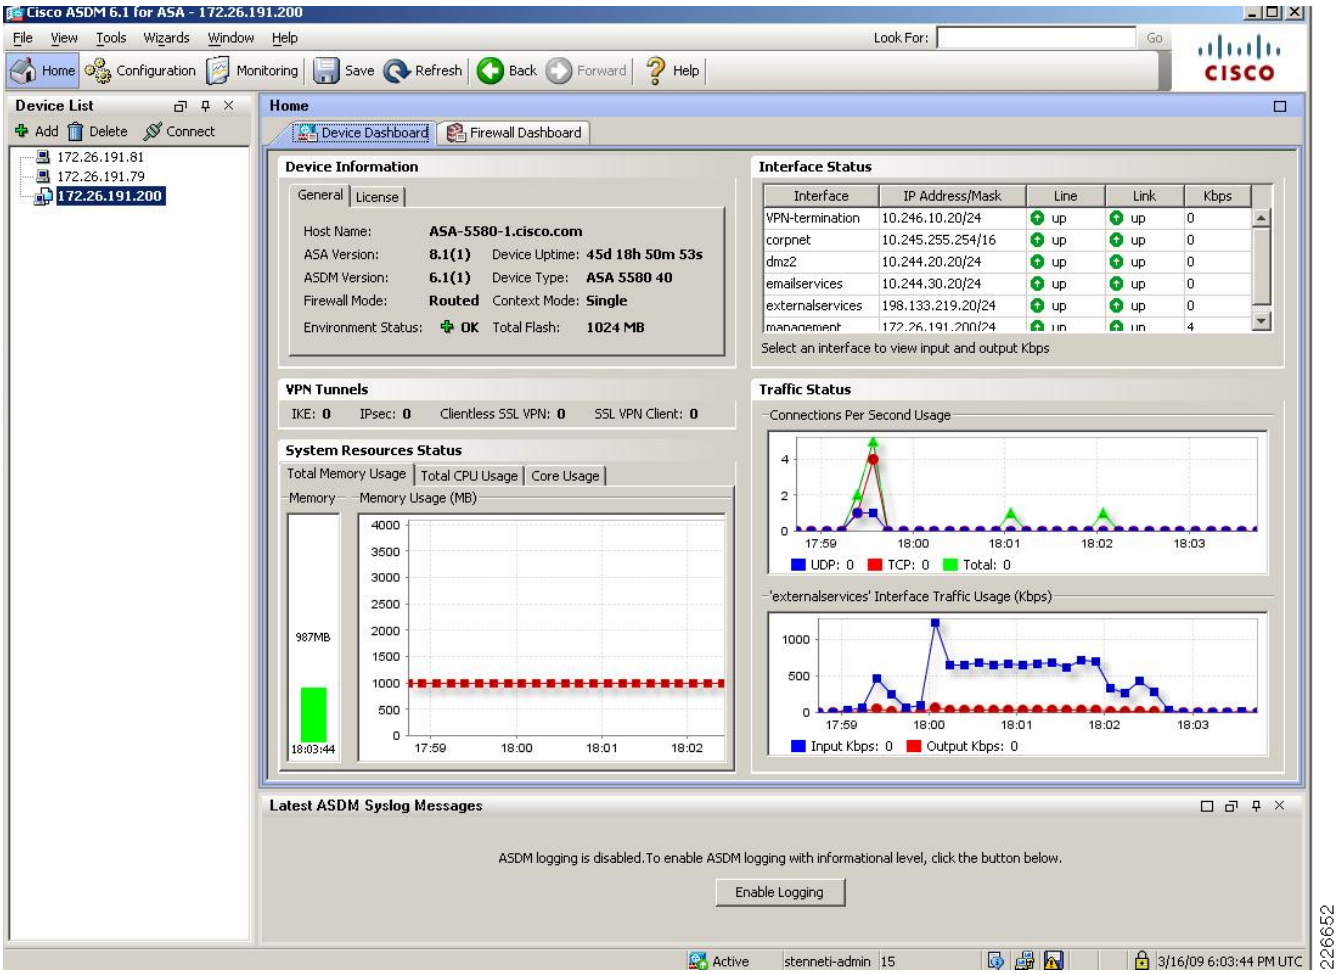

Figure 6-8 Cisco ASDM Screen Capture—Firewall Access Rules Page

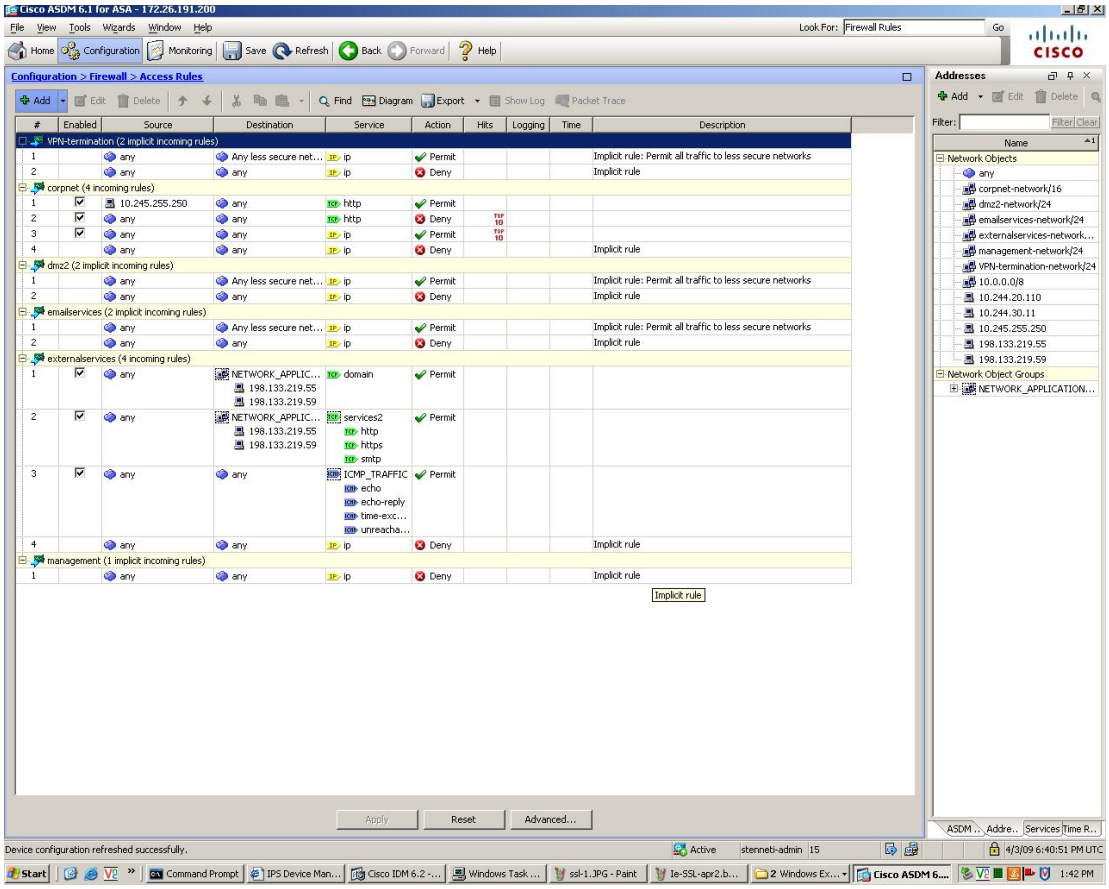

As shown in Figure 6-7 and Figure 6-8, the Cisco ASDM can be used to configure firewall rules and monitor a variety of statistics and system parameters. The following configuration steps illustrate the process necessary to implement the security best practices for the Internet edge firewalls.

Note The 64.104.0.0/16 and 198.133.219.0/24 address blocks used in the examples provided below are reserved for the exclusive use of Cisco Systems, Inc.

Step 1 Define the inter-interface and intra-interface security policy. The configuration that follows allows traffic to flow between the interfaces and within an interface of same security-level. This is required if two or more interfaces on the firewall are configured with the same security level.

```
same-security-traffic permit inter-interface
same-security-traffic permit intra-interface
```

Step 2 Define the object groups. The configuration that follows allows objects—such as IP hosts or networks, protocols, ports, and ICMP types—to be collected into object groups. This simplifies deployment and makes it easier to deploy future servers or hosts without modifying the ACLs.

```
object-group network NETWORK_APPLICATION_HOSTS
 network-object 198.133.219.55 255.255.255.255 <-- This is Iron Port E-mail server
 network-object 198.133.219.59 255.255.255.255 <-- This is web application firewall
object-group protocol NETWORK_APPLICATION_PROTOCOL
 protocol-object tcp
 protocol-object udp
```

```
object-group service services1 tcp-udp
 description DNS Group
 port-object eq domain
object-group service services2 tcp
 port-object eq www
 port-object eq https
 port-object eq smtp
object-group icmp-type ICMP_TRAFFIC
 icmp-object echo-reply
 icmp-object time-exceeded
 icmp-object unreachable
 icmp-object echo
object-group service ICMP_TRAFFIC_1
 description (Generated by Cisco SM from Object "ICMP_TRAFFIC")
 service-object icmp echo
 service-object icmp unreachable
 service-object icmp time-exceeded
 service-object icmp echo-reply
```

Step 3 Define the key services that are to be visible to the outside world. In the design presented here, the web application firewall appliance and IronPort E-mail servers are visible to outside. As a result, static NAT translations for these services must be defined. The following example commands illustrate this configuration.

```
static (dmz2,externalservices) tcp 198.133.219.59 www 10.244.20.110 www netmask
255.255.255.255
static (emailservices,externalservices) 198.133.219.55 10.244.30.11 netmask
255.255.255.255
```

Step 4 Define the Protocol Address Translation (PAT) and NAT pool for corporate access as illustrated in the following configuration.

```
global (externalservices) 20 198.133.219.129-198.133.219.254 netmask 255.255.255.128
global (externalservices) 20 198.133.219.128 netmask 255.255.255.255
nat (corpnet) 20 access-list corp-net
nat (VPN-termination) 2 10.246.10.0 255.255.255.0
```

Step 5 Define the ACL for allowing external access as illustrated in the following configuration.

```
access-list OUTSIDE_IN extended permit tcp any object-group NETWORK_APPLICATION_HOSTS eq
domain
access-list OUTSIDE_IN extended permit tcp any object-group NETWORK_APPLICATION_HOSTS
object-group services2
```

Note Defining object groups greatly simplifies deploying the firewall policy.

Step 6 Define the ACL to prevent the inside users from trying to access the Internet without going to the IronPort appliance as illustrated in the following configuration.

```
access-list WEB_ACCESS extended permit tcp host 10.245.255.250 any eq www
access-list WEB_ACCESS extended permit tcp host 10.242.50.99 any eq www
access-list WEB_ACCESS extended permit tcp host 10.242.50.96 any eq www
access-list WEB_ACCESS extended deny tcp any any eq www
access-list WEB_ACCESS extended permit ip any any
```

Step 7 Apply the ACL WEB_ACCESS to *corpnet* and apply the ACL OUTSIDE_IN to *externalservices* as illustrated in the following configuration.

```
access-group OUTSIDE_IN in interface externalservices
access-group WEB_ACCESS in interface corpnet
```

Web Application Firewall

The web application firewall acts as a reverse proxy for the web servers that it is configured to protect. The *virtual web application* is used to create a virtual URL that will be used to intercept incoming client connections. You can configure one more virtual web applications based on the protocol and port, as well as the policy you want to be applied. Covering every aspect of web application firewall configuration and policy management is beyond the scope of this publication. Only basic implementation steps as they pertain to the Internet edge architecture are addressed. For more details, refer to the web application firewall reference guide listed in Appendix A, "Reference Documents."

Basic configuration of network services is handled through the console port or a keyboard. The policy configuration and management are done through a GUI via HTTPS. The web application firewall can be configured with a virtual address that acts as the IP address of a web server. The web application firewalls can then point to the actual web server and inspect all traffic destined for the web server.

The logical placement and deployment model of web application firewall is shown in Figure 6-9.

Figure 6-9 Cisco Application Control Engine (ACE) Web Application Firewall Logical Placement

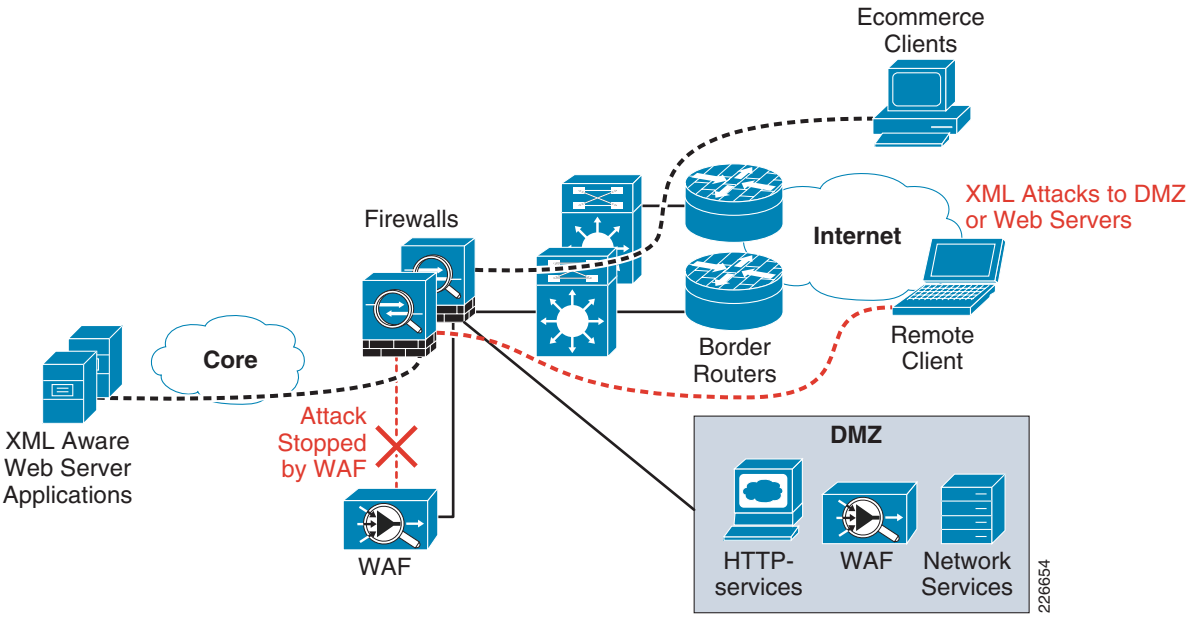

The following are some of the best practices that should be used when implementing web application firewall:

- The web application firewall is implemented as one-armed design with the single interface connecting to the DMZ.
- Configure the web application firewall to retain the source IP address if the traffic is directed to appliances in the data center.
- It is recommended that HTTPS traffic directed to the data center, not be encrypted as the Cisco ACE
- module in data center will perform the load-balancing and decryption while also providing higher performance.
- The web application firewall in the Internet edge and the web application firewall in data center to be configured in the same cluster.

For more information on best practices, configuration steps and threat control and monitoring capability of the web application firewall, refer to the web application firewall reference guide listed in Appendix A, "Reference Documents."

Examples of the GUI interface used to view rules and monitor events are shown in Figure 6-10 and Figure 6-11.

Figure 6-10 Cisco ACE Web Application Firewall—Viewing Rules and Signatures

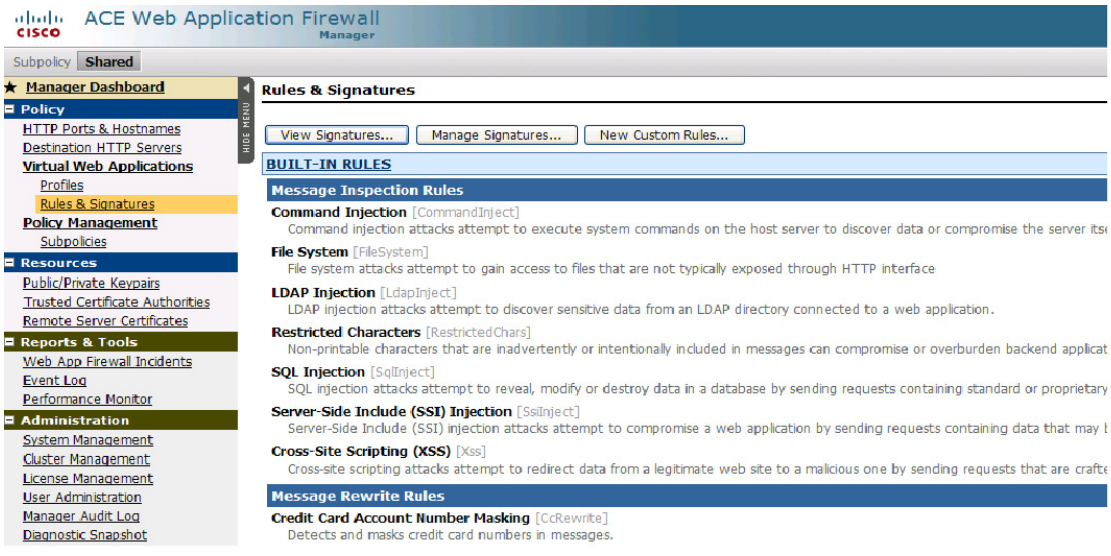

Figure 6-11 Cisco ACE Web Application Firewall—Viewing Event Log Viewer

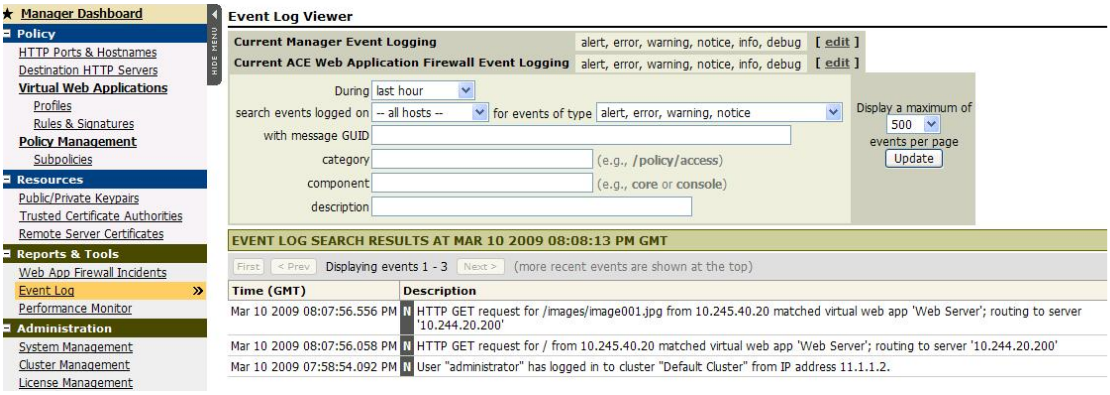

E-mail and Web Security

To implement the best practices for the ESA and WSA, a good understanding of the SenderBase SensorBase network is required. The ESA and WSA use the information gathered by the SenderBase SensorBase network to make decisions about threat level of websites and senders of received E-mails. The following section summarizes the operation and advantages of the SenderBase SensorBase network.

IronPort SensorBase

The IronPort ESA and WSA use the SensorBase network to gain a real-time view into security threats and stop E-mail spam and E-mails from malicious sites. The SensorBase network is an extensive network that monitors global E-mail and web traffic for anomalies, viruses, malware and other and abnormal behavior. It queries a significant percentage of all global E-mail and web traffic and uses tens of parameters to determine spam E-mail sites and malicious or compromised websites.

SensorBase examines more than 90 different parameters about E-mail traffic and 20 different parameters about web traffic. Parameters tracked include global sending volume, complaint levels, "spamtrap" accounts, whether a sender's DNS resolves properly and accepts return mail, country of origin, blacklist information, probability that URLs are appearing as part of a spam or virus attack, open proxy status, use of hijacked IP space, valid and invalid recipients, and other parameters. By using sophisticated algorithms, SensorBase creates a reputation score for domains and websites that ranges from -10 to +10. This score is analogous to credit scores for individuals and is used to determine risk. Every ESA implemented at the enterprise can dynamically lookup reputation scores for domains of each E-mail it receives, or each website to which it is connected. The appliance can use preconfigured policies to drop, monitor, or quarantine E-mails from suspect mail sites-and to drop connections to malicious websites.

Figure 6-12 depicts the operation of the IronPort SensorBase network.

Figure 6-12 IronPort SensorBase Network

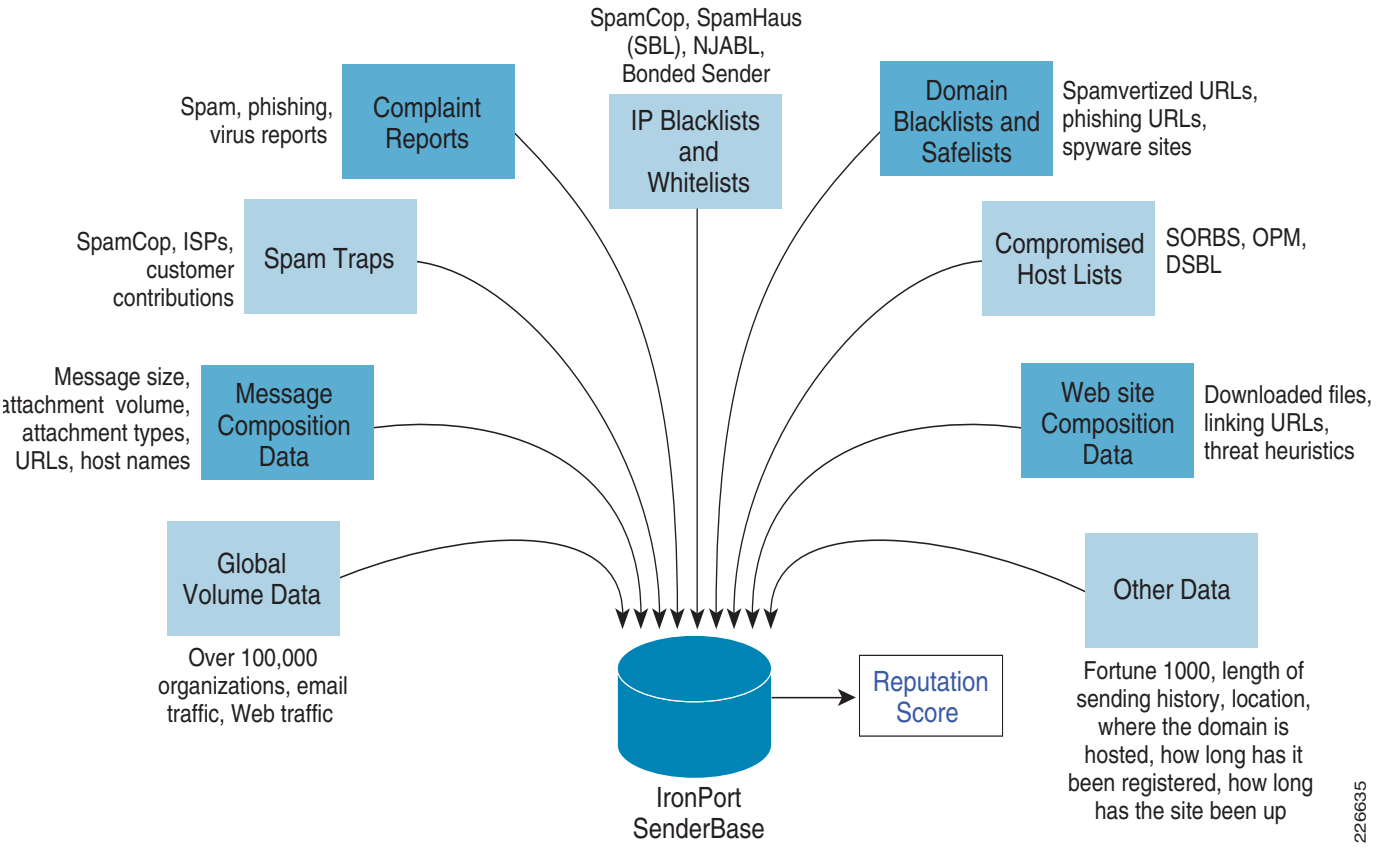

Web Security Appliance Best Practices

The function of the WSA is to monitor and mitigate any abnormal web activity between corporate users and the outside world. The WSA is logically located in the path between corporate web users and the Internet. In effect, the WSA acts as a web proxy for the corporate users residing inside the network. This logical placement of the WSA implies proper configuration of the browser. There are three different ways clients may interact with WSA:

- *Explicit mode without use of Proxy Auto Configuration (PAC) files*—This requires manual configuration of the browser to point to the WSA as its proxy. This choice does not support redundancy, does not work with multiple WSAs, and requires changes to every browser in the enterprise network. This is the preferred method to test and verify proper operation of the WSA.
- *Explicit mode with use of PAC files*—In this mode, the proxy information is stored in a file that can be downloaded automatically or the file's location can be referenced manually. The advantage of this mode is that more than one proxy can be referenced in the files and used by the browser. This allows for load balancing and redundancy of WSAs. You can use Dynamic Host Configuration Protocol (DHCP) or DNS to download the files automatically to the browser. This eliminates the need to manually configure each browser separately.
- *Transparent mode with Web Cache Communications Protocol (WCCP)*—In this mode, the web traffic is transparently directed to the WSA using WCCP redirection and does not require any adjustments to the browser. This mode requires the configuration of a WCCP-enabled firewall, router or Layer-3 switch to direct client traffic to the appliance. Care should be taken when asymmetrical traffic flows exist in the network. This is a good method for load sharing and redundancy.

It is recommended that explicit mode be initially implemented. You may use this mode with the use of PAC files for initial testing-and then transition to WCCP for final implementation. As mentioned in the preceding description, with PAC files, you may achieve load balancing and redundancy between multiple WSAs. Alternatively, WCCP-based transparent mode may be used if you require weighted load-balancing or source and destination hashing. More sophisticated load-balancing is also possible with the use of a Layer-4 load balancer, such as Cisco Application Control Engine (ACE). Figure 6-13 illustrates the manual proxy configuration in Microsoft Internet Explorer.

Figure 6-13 Example Browser Configuration Window for Setting up WSA Reference

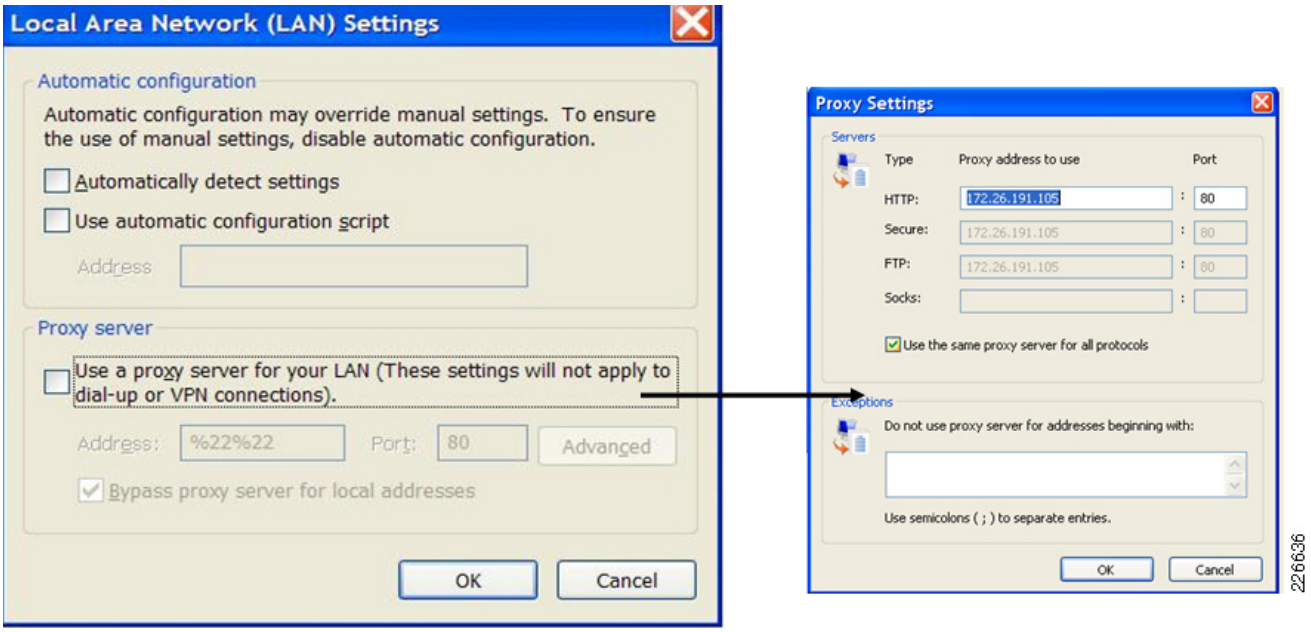

To implement explicit mode without using PAC files, use the **proxy server** configuration setting shown in Figure 6-13 and manually enter the IP address of the WSA. The **use automatic configuration script** configuration is used to indicate the location of the PAC file used by the browser; with WCCP redirection, you do not configure anything. Similar configuration options are available for other popular browsers.

Other recommendations and best practices for WSA deployment are as follows:

- The edge firewalls should be configured to allow only outgoing HTTP or Hypertext Transfer Protocol over SSL (HTTPS) connections sourced from the WSA and other devices requiring such access—such as ESA and Cisco Security Monitoring, Analysis and Response System (Cisco Security MARS), or Cisco Security Manager (CSM). This would prevent users from bypassing the WSA in order to directly connect to the Internet.
- Determine the IP address of the WSA to which all browsers will point as a web proxy.
- Configure all browsers in the organization to point to the WSA either through PAC or manual configuration. Once the location of the PAC files are configured, no other changes to the browser are necessary. If using WCCP, end-station configuration is not needed.
- If an upstream proxy is present, configure the WSA to point to the upstream proxy.
- Determine policies for handling HTTPS traffic and configure WSA to enforce those policies.
- Configure the WSA policies and actions to be taken for the different ranges in the Web Reputation Score. Based on the reputation score, the WSA can pass, monitor, or drop web traffic.
- Configure enforcement policies for your organization through the URL filter available on the WSA. URL filters allow blocking or monitoring of users based on the website categories users visit.
- In creating policies for encrypted traffic. It is recommended that a white list of well-known sites be created through which traffic to those sites is allowed to pass. This saves resources on the WSA. It is also good practice to monitor traffic for other sites that use encryption.

- If a no split-tunneling policy is enforced at the branches, then the browsers on all branches should point to the WSA. This practice will ensure that all Internet traffic flows through the corporate network and is passed through and monitored by the WSA.
- Use a separate interface to connect to the management network.

You can use the WSA reporting tools to monitor web activity and to look for any malicious activity. The screen shots in Figure 6-14 through Figure 6-16 illustrate the monitoring capabilities available.

Figure 6-14 WSA Reporting Window—Web Site Activity

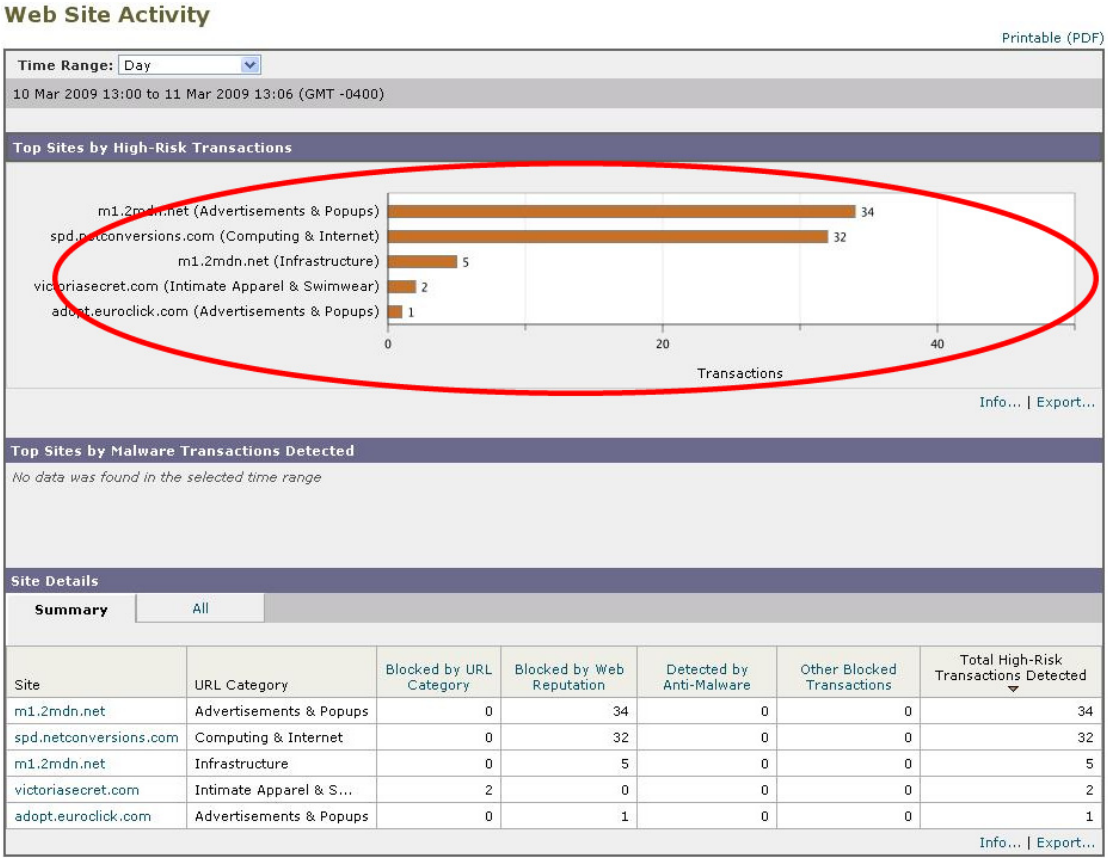

Figure 6-15 WSA Reporting Window—URL Categories

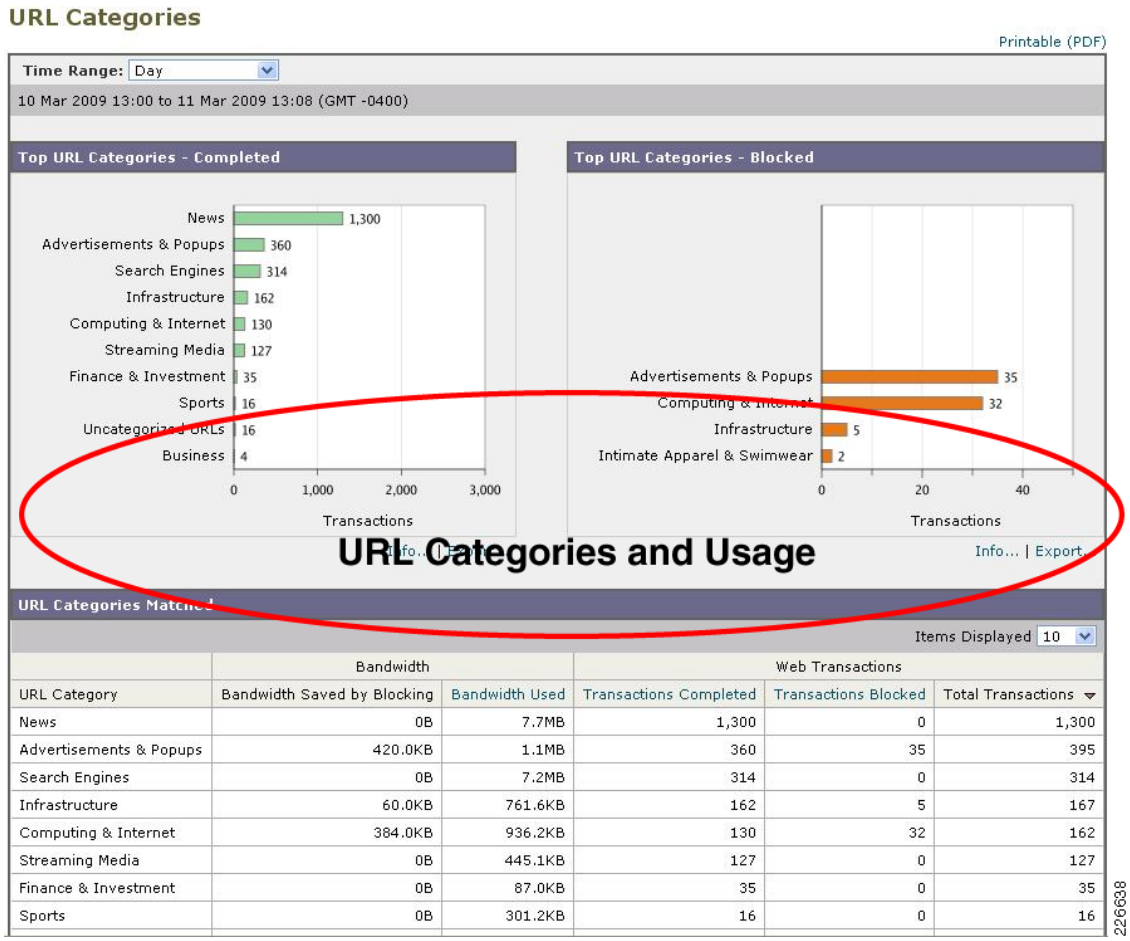

The *Web-Site activity* screen allows the administrator to determine what websites were blocked from the user and the reason for blocking access. A website can be blocked because of a bad reputation score, because spyware or malware was detected by anti-malware, or due to URL filtering. The URL filtering window categorizes all the visited websites and shows the amount of traffic and number of blocked transactions for each category. The client website window in Figure 6-16 shows the website activity for each client that can be used for enforcing acceptable-use policies.

Figure 6-16 WSA Reporting Window—Client Web Activity

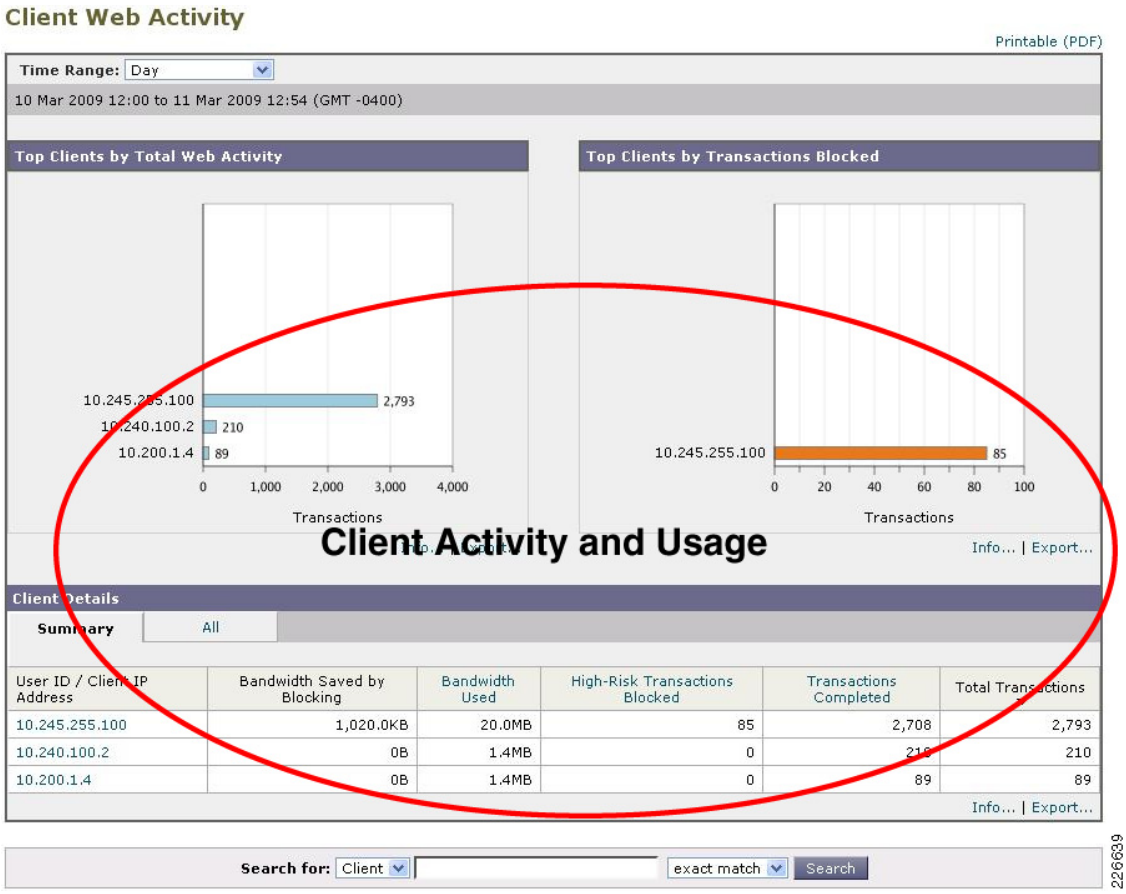

The E-mail Security Appliance

E-mail is a medium through which spyware and viruses can be propagated. In addition to outside threats, E-mail spam and malicious malware can reduce employee productivity. The ESA is a type of firewall and threat monitoring appliance for Simple Mail Transfer Protocol (SMTP) traffic (TCP port 25). Logically speaking, the ESA acts as a Mail Transfer Agent (MTA) within the E-mail delivery chain. There are multiple deployment approaches for the security appliance depending on the number of interfaces used. ESA may be deployed with a single physical interface to transfer emails to and from both the Internet and the internal mail servers. If desired, two physical interfaces may be used, one for email transfer to and from the Internet, and another one for email communications to the internal servers. In the former approach, the ESA would reside on the DMZ, while in the later, the ESA would have an interface connecting to the DMZ and the other one connecting to the inside network. In this case, the DMZ-based interface would send and receive E-mail to and from the Internet. The inside network interface would be used to deliver E-mail to the internal mail server.

This guide follows the single-interface model as it is the simplest and most commonly deployed. Figure 6-17 shows both deployment models.

Figure 6-17 IronPort ESA Deployment Models

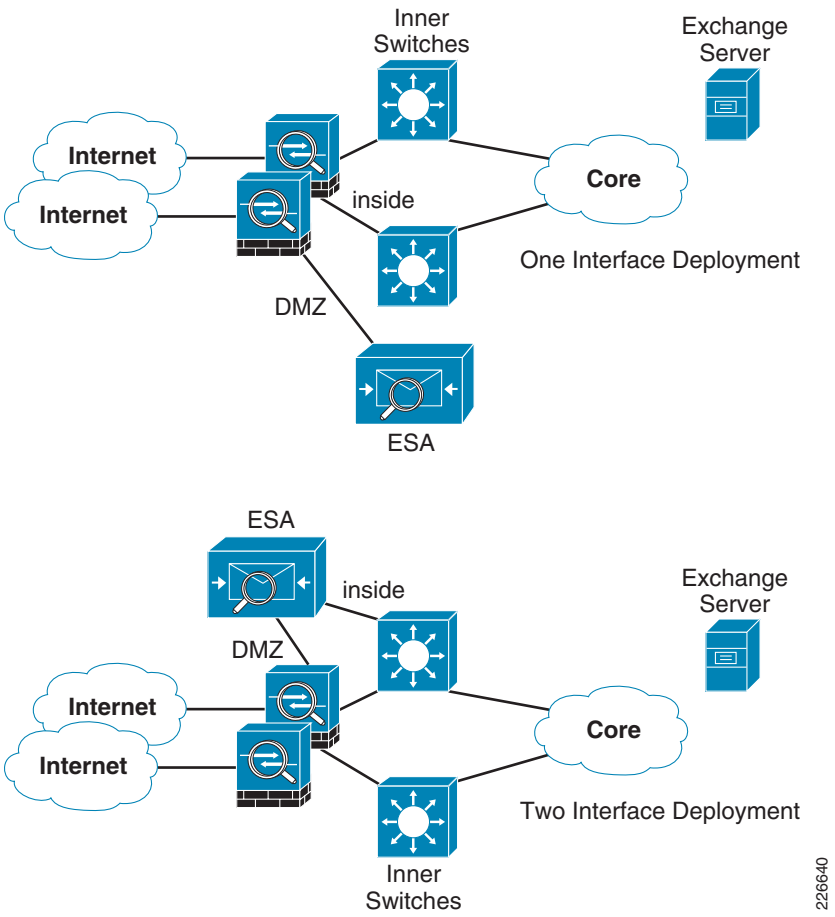

E-mail Data Flow

Consider a sender somewhere in the Internet first sending an E-mail to a mail server. The E-mail server resolves the E-mail domain name to the public IP address of the domain and sends the E-mail using SMTP to the corresponding IP address. Upon receiving the E-mail, the enterprise firewall translates the public IP address of the ESA to the DMZ IP address and forwards traffic to the ESA. The ESA then does a DNS query on the sender domain name, compares the IP address of the sender to its own SensorBase database, and determines the reputation score of the sender. It rejects the E-mail if it falls within a pre-configured reputation score. A typical dataflow for inbound E-mail traffic is shown in Figure 6-18.

Figure 6-18 Typical Data Flow for Inbound E-mail Traffic

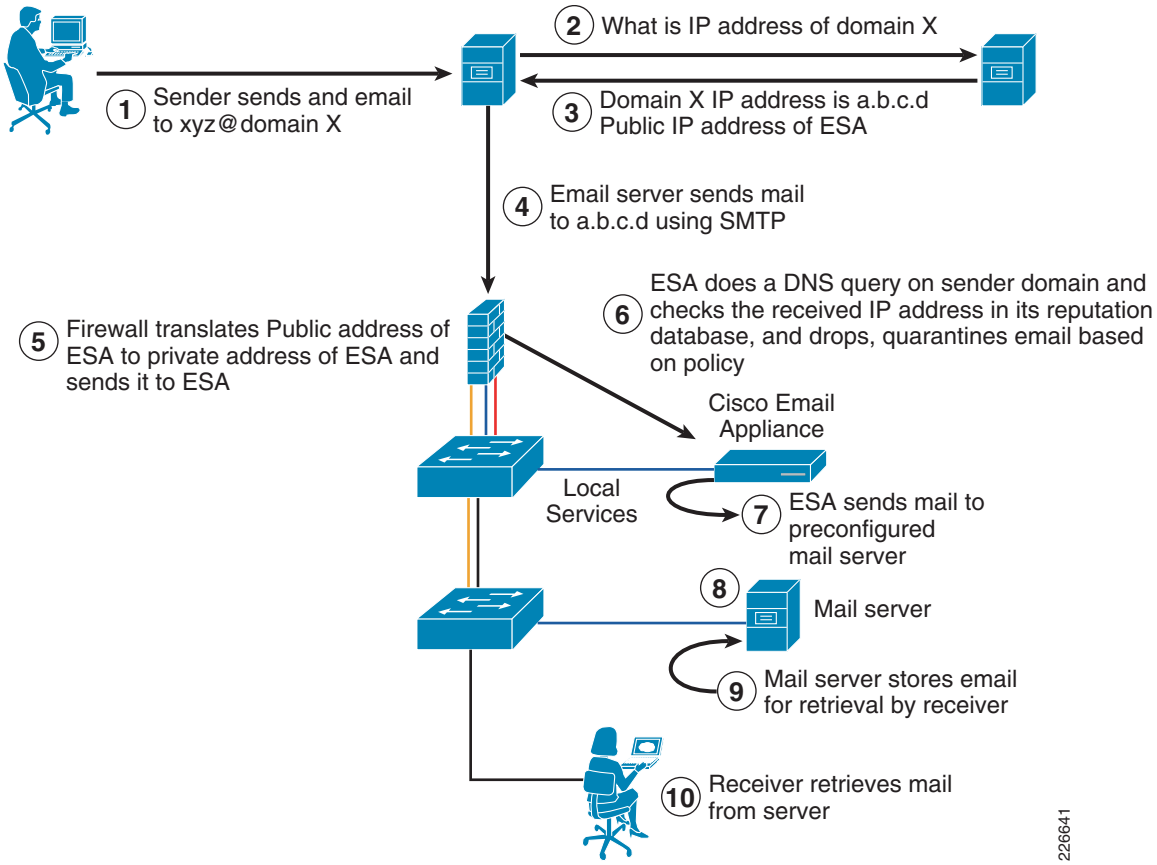

Redundancy and Load Balancing of an E-mail Security Appliance

Redundancy is often a requirement, a failure of an ESA can cause mail service outage. There are multiple ways to configure redundancy; the simplest one is to add the second appliance with an equal cost secondary MX record, as shown in Figure 6-19. In this method, traffic will be shared across two or more ESAs. A more advanced design would include a load-balancing platform for balancing traffic among multiple appliances.

Figure 6-19 IronPort E-mail Appliance High Availability Environment

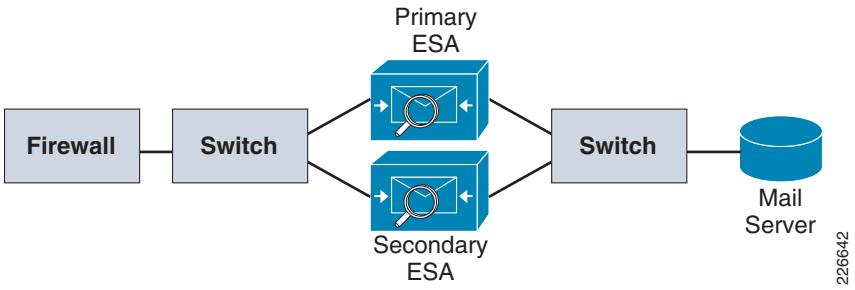

Best Practices and Configuration Guidelines for ESA Implementation

The first task when implementing an ESA in the enterprise is to define firewall rules. Important considerations when defining firewall rules are as follows:

- A static address must be defined on the firewall to translate a publicly accessible IP address for the E-mail server to a private IP address used by the ESA.
- It is recommended that the ESA be configured to access a DNS in the outside network, rather than the internal DNS. This means that the firewall must allow ESA do perform DNS queries and receive DNS replies.
- The ESA downloads the latest SensorBase information, virus updates, and so on through HTTP/HTTPS connections. Again firewall rules must allow HTTP/HTTPS traffic from the ESA.
- SMTP routes must be set to point to inside E-mail servers.
- Either the same interface or a separate interface can be used for incoming or outgoing mail. If the same interface is used, you will need to relay mail on the interface.
- Use a separate interface to connect to the management network.
- Use separate subnets for different organizational domains. This simplifies the configuration of policies in the WSA for different groups of users.

IronPort has a very intuitive and powerful configuration web interface. The network, interface, and SMTP routing information can be configured using the wizard. Firewall rules and Network Address Translation (NAT) are configured on the Cisco Adaptive Security Appliances (ASA). IronPort ESA is a functionally rich appliance. The following guidelines give the implementation framework and the actions necessary to implement an ESA on the network. A more detailed discussion of the IronPort ESA can be found at the following URL: http://www.ironport.com/resources/whitepapers.html

Step 1 Determine IP addresses, domain names, and networks with which the ESA will be configured.

Step 2 Obtain a public address for the ESA.

Step 3 Create SMTP routes to the private E-mail servers.

Step 4 Create firewall rules to allow in TCP port 25 (SMTP), UDP, and TCP port 53 (DNS).

Step 5 Create firewall rules to allow HTTP/HTTPS so that the ESA can contact SensorBase and get virus protection updates.

Step 6 Configure the ESA with DNS and the default route.

Step 7 Configure the management interface.

Step 8 Configure incoming and outgoing E-mail policies and content filters to match the requirements of the enterprise organization.

From a security perspective, you can use the monitoring functionality available in the ESA's GUI to manage and react to threats. The screen shots for some of the monitoring tools are presented in Figure 6-20 and Figure 6-21.

Figure 6-20 ESA Monitoring Screen—Virus Outbreaks

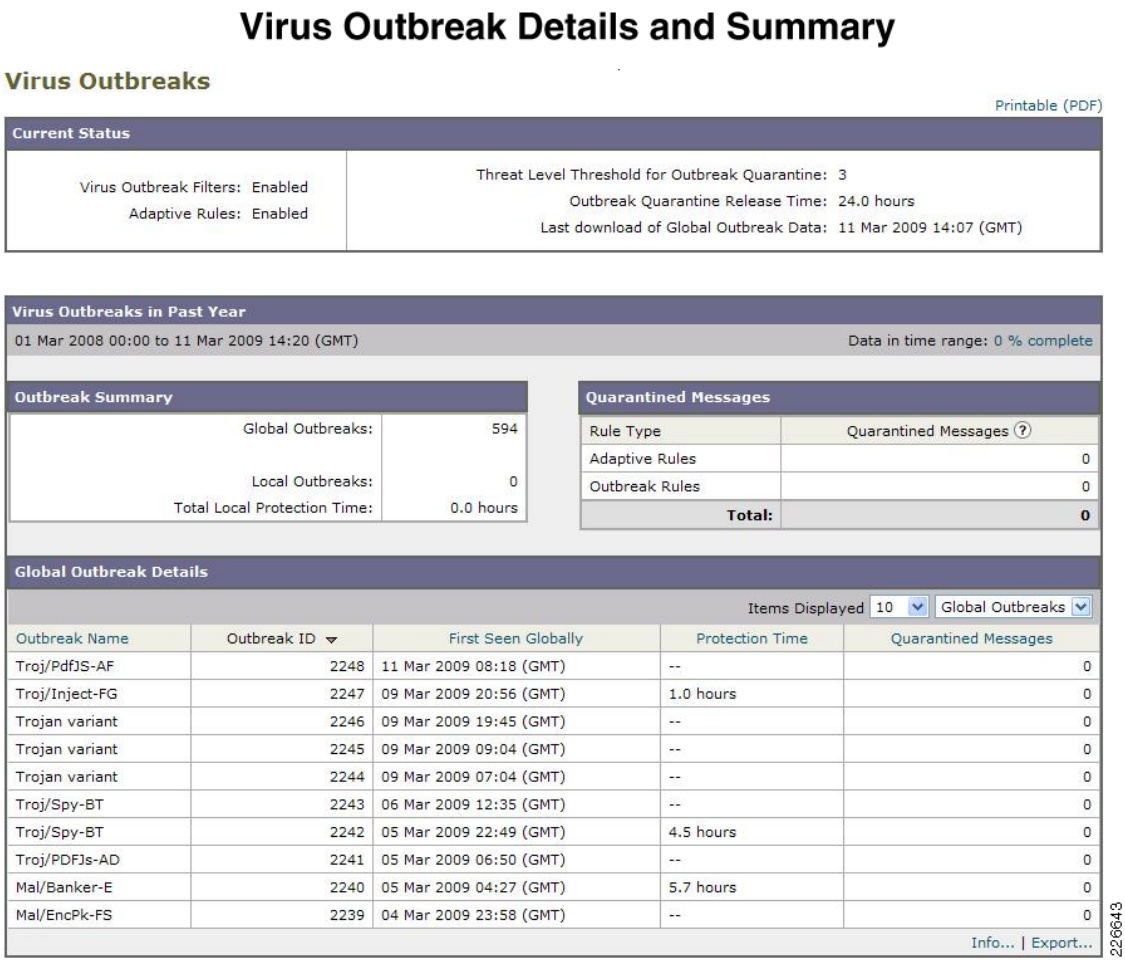

The virus outbreak screen (Figure 6-20) shows the different viruses detected, action taken, and total number of outbreaks. The message analysis screen (Figure 6-21) categorizes different types of threats that were blocked and provides statistical analysis of total threats received.

Figure 6-21 ESA Monitoring Screen—Message Analysis

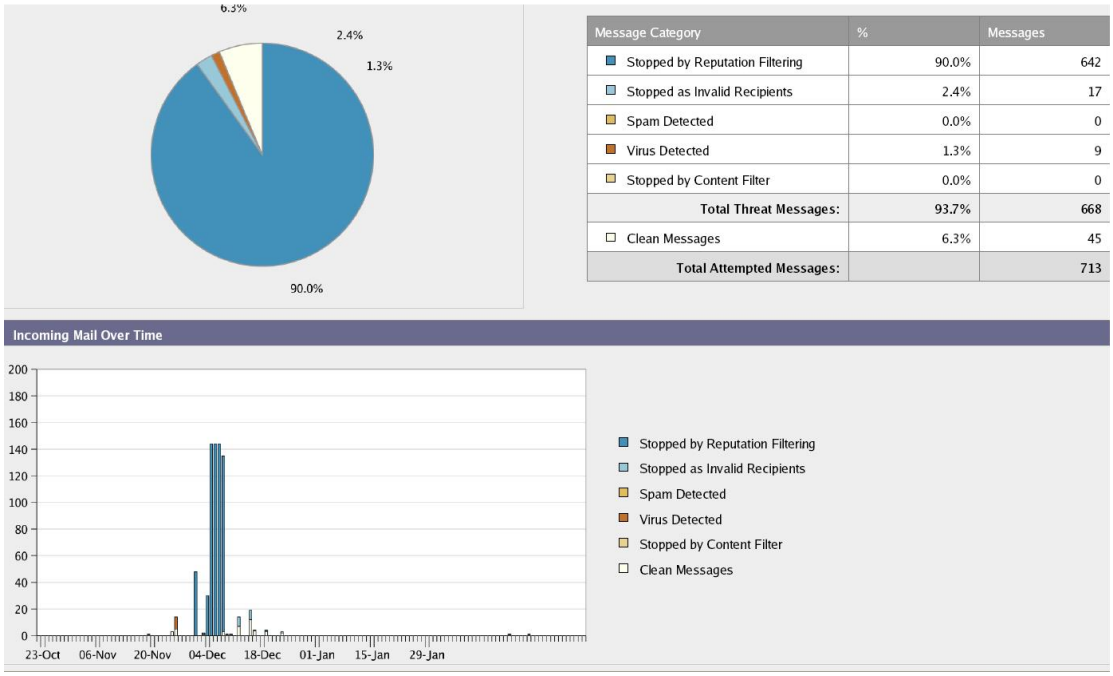

Service Provider Block

The Service Provider (SP) edge block is a critical part of the Internet edge because it provides the interface to the public Internet infrastructure. The following topics are covered in this section:

- Design Guidelines and Best Practices for the SP Edge Block, page 6-28
- Security Features for BGP, page 6-29
- Infrastructure ACL Implementation, page 6-33

Figure 6-22 illustrates the SP edge block topology.

Figure 6-22 Service Provider Block Topology

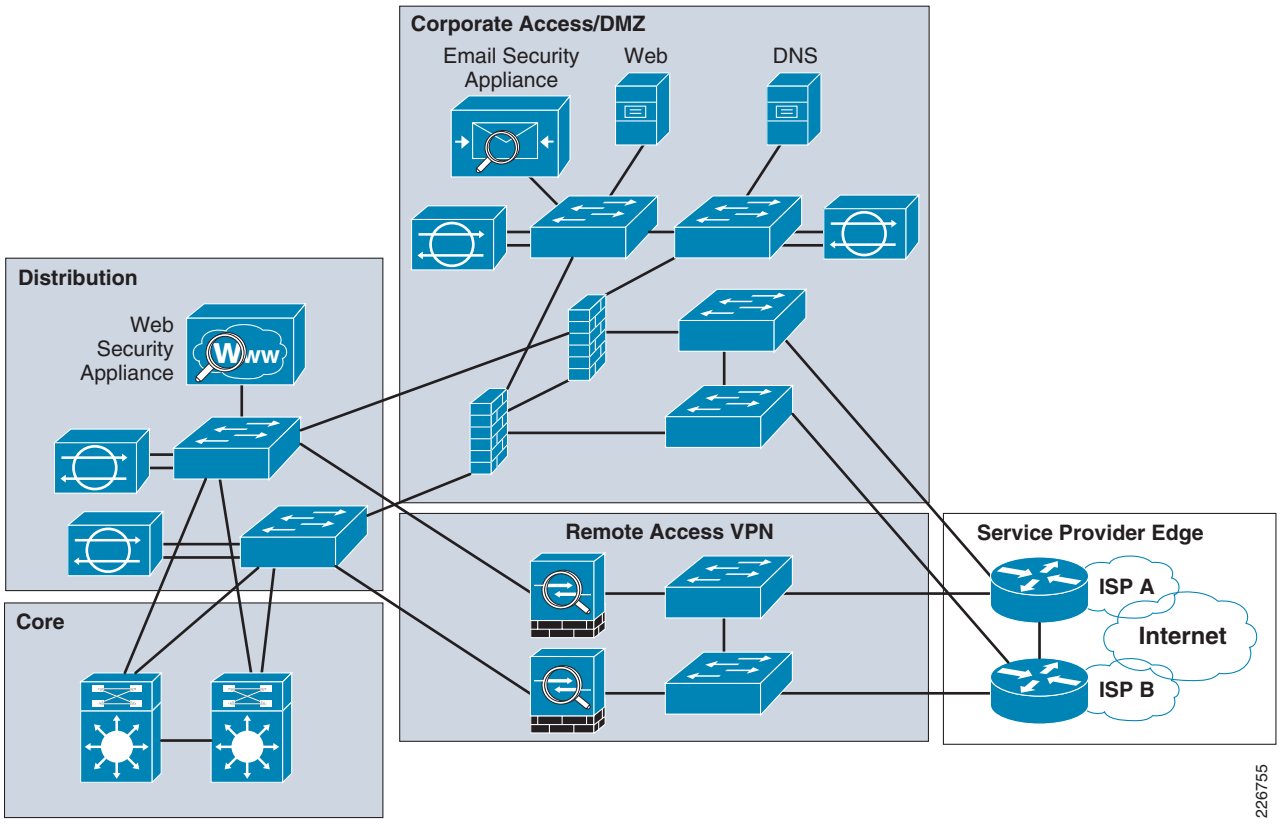

Figure 6-23 illustrates an example of the interface with the SP environment via Border Gateway Protocol (BGP).

Figure 6-23 BGP Design

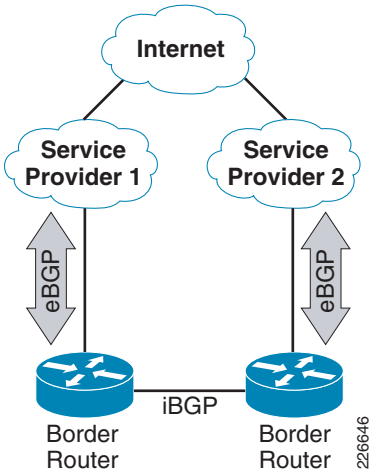

Design Guidelines and Best Practices for the SP Edge Block

Figure 6-24 illustrates the topology used to implement a BGP-based, SP-edge block environment. The configuration examples presented in the subsequent descriptions depict best practices for this scenario and are taken from this example topology.

Note The 64.104.0.0/16 and 198.133.219.0/24 address blocks used in the examples provided below are reserved for the exclusive use of Cisco Systems, Inc.

Figure 6-24 BGP Topology Diagram

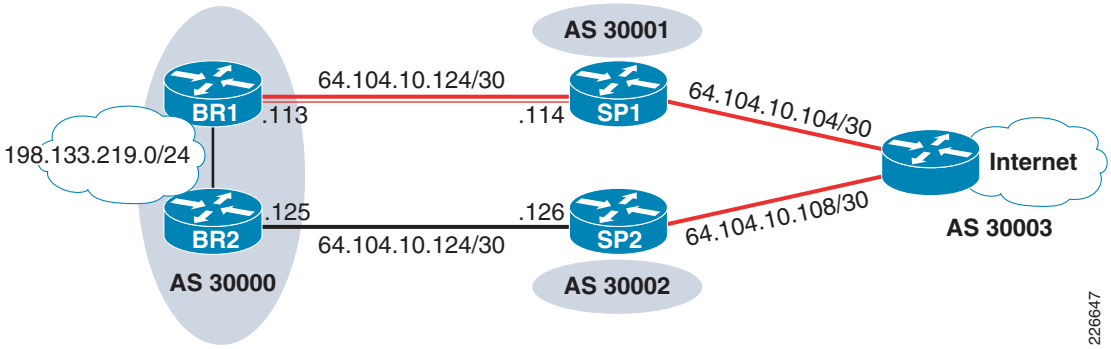

The following are the design recommendations for the SP edge:

- Use BGP as the routing protocol for all dynamic routing—both between the border routers and between the border routers and SP.
- Have an independent autonomous system number. This will give the flexibility of advertising the Internet prefix to different SPs.

- Use PfR as path-optimization mechanism. This will ensure that the optimal path is selected between the SPs—thereby increasing the application performance.

Harden the SP edge infrastructure by following the best practices described in Chapter 2, "Network Foundation Protection."

Security Features for BGP

The BGP support for the (time-to-live) TTL security check feature introduces a lightweight security mechanism to protect eBGP peering sessions from CPU utilization-based attacks. These types of attacks are typically brute force DoS attacks that attempt to disable the network by flooding the network with IP packets that contain forged source and destination IP addresses.

TTL security check allows the configuration of a minimum acceptable TTL value for the packets exchanged between two eBGP peers. When enabled, both peering routers transmit all BGP packets with a TTL value of 255. An eBGP router will establish a peering session with another router only if that other is an eBGP peer that sends packets with a TTL equal-to-or-greater-than an expected TTL value for the peering session. The expected TTL value is calculated by subtracting the hop count configured for the session to 255. All packets received with TTL values less than the expected value are silently discarded.

Although it is possible to forge the TTL field in an IP packet header, accurately forging the TTL count to match the TTL count from a trusted peer is impossible unless the network to which the trusted peer belongs has been compromised. For more information, refer to the following URL:

http://www.cisco.com/en/US/docs/ios/12_3t/12_3t7/feature/guide/gt_btsh.html#wp1027184

Note The 64.104.0.0/16 and 198.133.219.0/24 address blocks used in the examples provided below are reserved for the exclusive use of Cisco Systems, Inc.

The following configuration command example illustrates the command required to enable TTL security on the SP edge router:

```
neighbor 64.104.10.114 ttl-security hops 2
```

Note This feature must be enabled on both sides of a connection (enterprise border router and the SP router).

The following **show** command verifies whether the TTL security feature is properly configured on the router (the relevant line is highlighted):

```
IE-7200-3# show ip bgp neighbors

BGP neighbor is 64.104.10.114,  remote AS 30001, external link
  BGP version 4, remote router ID 172.26.191.176
  BGP state = Established, up for 1w6d
  Last read 00:00:53, last write 00:00:42, hold time is 180, keepalive interval is 60
seconds
  Neighbor capabilities:
    Route refresh: advertised and received(new)
    Address family IPv4 Unicast: advertised and received
  Message statistics:
    InQ depth is 0
    OutQ depth is 0

                         Sent       Rcvd
    Opens:                  1          1
    Notifications:          0          0
    Updates:                1          1
```

```
    Keepalives:         19943      20081
    Route Refresh:          0          0
    Total:              19945      20083
  Default minimum time between advertisement runs is 30 seconds

 For address family: IPv4 Unicast
  BGP table version 37589, neighbor version 37589/0
  Output queue size : 0
  Index 2, Offset 0, Mask 0x4
  2 update-group member
  Outbound path policy configured
  Route map for outgoing advertisements is my_routes
                               Sent       Rcvd
  Prefix activity:             ----       ----
    Prefixes Current:             1          1 (Consumes 416 bytes)
    Prefixes Total:               1          1
    Implicit Withdraw:            0          0
    Explicit Withdraw:            0          0
    Used as bestpath:           n/a          8
    Used as multipath:          n/a          0


                                 Outbound    Inbound
  Local Policy Denied Prefixes:  --------    -------
    route-map:                      18758          0
    Total:                          18758          0
  Number of NLRIs in the update sent: max 1, min 1

  Address tracking is enabled, the RIB does have a route to 64.104.10.114
  Connections established 1; dropped 0
  Last reset never

  External BGP neighbor may be up to 2 hops away.
  Transport(tcp) path-mtu-discovery is enabled
Connection state is ESTAB, I/O status: 1, unread input bytes: 0
Connection is ECN Disabled, Mininum incoming TTL 253, Outgoing TTL 255
Local host: 64.104.10.113, Local port: 179
Foreign host: 64.104.10.114, Foreign port: 36929
Connection tableid (VRF): 0

Enqueued packets for retransmit: 0, input: 0  mis-ordered: 0 (0 bytes)

Event Timers (current time is 0x47D127AC):
Timer          Starts    Wakeups            Next
Retrans         19946          1             0x0
TimeWait            0          0             0x0
AckHold         20081      19750             0x0
SendWnd             0          0             0x0
KeepAlive           0          0             0x0
GiveUp              0          0             0x0
PmtuAger            0          0             0x0
DeadWait            0          0             0x0
Linger              0          0             0x0
ProcessQ            0          0             0x0

iss: 2105715872  snduna: 2106094887  sndnxt: 2106094887     sndwnd:  15700
irs: 2928015827  rcvnxt: 2928397480  rcvwnd:      15947  delrcvwnd:    437

SRTT: 300 ms, RTTO: 303 ms, RTV: 3 ms, KRTT: 0 ms
minRTT: 0 ms, maxRTT: 300 ms, ACK hold: 200 ms
Status Flags: passive open, gen tcbs
Option Flags: nagle, path mtu capable, md5
IP Precedence value : 6
```

```
Datagrams (max data segment is 1440 bytes):
Rcvd: 39763 (out of order: 0), with data: 20084, total data bytes: 381652
Sent: 39962 (retransmit: 1, fastretransmit: 0, partialack: 0, Second Congestion: 0), with
data: 19946, total data bytes: 379033
 Packets received in fast path: 0, fast processed: 0, slow path: 0
 fast lock acquisition failures: 0, slow path: 0
```

The routes learned from SP 1 should not be leaked to SP 2 and vice versa. To prevent the routes from leaking, an **as-path** access list and the **route-map** command are used. The following commands are required to implement this filtering:

- **as-path** filtering command

```
ip as-path access-list 20 permit ^$
ip as-path access-list 20 deny .*
```

- **route-map** command to match the **as-path** command

```
route-map my_routes permit 10
 match as-path 20
```

- **route-map** command applied to the external peers

```
neighbor 64.104.10.114 route-map my_routes out
```

The following **show** command output presents the number of prefixes that are denied; this information verifies that leakage is not happening between the border routers (output of interest is highlighted):

```
IE-7200-3# show ip bgp neighbors 64.104.10.114
BGP neighbor is 64.104.10.114,  remote AS 30001, external link
  BGP version 4, remote router ID 172.26.191.176
  BGP state = Established, up for 1w6d
  Last read 00:00:05, last write 00:00:44, hold time is 180, keepalive interval is 60
seconds
  Neighbor capabilities:
    Route refresh: advertised and received(new)
    Address family IPv4 Unicast: advertised and received
  Message statistics:
    InQ depth is 0
    OutQ depth is 0

                         Sent       Rcvd
    Opens:                  1          1
    Notifications:          0          0
    Updates:                1          1
    Keepalives:         19968      20107
    Route Refresh:          0          0
    Total:              19970      20109
  Default minimum time between advertisement runs is 30 seconds

 For address family: IPv4 Unicast
  BGP table version 37623, neighbor version 37623/0
  Output queue size : 0
  Index 2, Offset 0, Mask 0x4
  2 update-group member
  Outbound path policy configured
  Route map for outgoing advertisements is my_routes
                               Sent       Rcvd
  Prefix activity:             ----       ----
    Prefixes Current:             1          1 (Consumes 312 bytes)
    Prefixes Total:               1          1
    Implicit Withdraw:            0          0
    Explicit Withdraw:            0          0
    Used as bestpath:           n/a          6
    Used as multipath:          n/a          0
```

```
                                  Outbound    Inbound
  Local Policy Denied Prefixes:   --------    -------
    route-map:                       18773          0
    Total:                           18773          0
  Number of NLRIs in the update sent: max 1, min 1

  Address tracking is enabled, the RIB does have a route to 64.104.10.114
  Connections established 1; dropped 0
  Last reset never
  External BGP neighbor may be up to 2 hops away.
  Transport(tcp) path-mtu-discovery is enabled
Connection state is ESTAB, I/O status: 1, unread input bytes: 0
Connection is ECN Disabled, Mininum incoming TTL 253, Outgoing TTL 255
Local host: 64.104.10.113, Local port: 179
Foreign host: 64.104.10.114, Foreign port: 36929
Connection tableid (VRF): 0

Enqueued packets for retransmit: 0, input: 0  mis-ordered: 0 (0 bytes)

Event Timers (current time is 0x47E81AEC):
Timer          Starts    Wakeups            Next
Retrans         19971          1             0x0
TimeWait            0          0             0x0
AckHold         20106      19775             0x0
SendWnd             0          0             0x0
KeepAlive           0          0             0x0
GiveUp              0          0             0x0
PmtuAger            0          0             0x0
DeadWait            0          0             0x0
Linger              0          0             0x0
ProcessQ            0          0             0x0

iss: 2105715872  snduna: 2106095362  sndnxt: 2106095362     sndwnd:  15225
irs: 2928015827  rcvnxt: 2928397955  rcvwnd:      15472  delrcvwnd:    912

SRTT: 300 ms, RTTO: 303 ms, RTV: 3 ms, KRTT: 0 ms
minRTT: 0 ms, maxRTT: 300 ms, ACK hold: 200 ms
Status Flags: passive open, gen tcbs
Option Flags: nagle, path mtu capable, md5
IP Precedence value : 6

Datagrams (max data segment is 1440 bytes):
Rcvd: 39812 (out of order: 0), with data: 20109, total data bytes: 382127
Sent: 40011 (retransmit: 1, fastretransmit: 0, partialack: 0, Second Congestion: 0), with
data: 19970, total data bytes: 379489
 Packets received in fast path: 0, fast processed: 0, slow path: 0
 fast lock acquisition failures: 0, slow path: 0
IE-7200-3#

IE-7200-3# show ip as-path-access-list
AS path access list 20
    permit ^$
    deny .*

IE-7200-3# show route-map my_routes
route-map my_routes, permit, sequence 10
  Match clauses:
    as-path (as-path filter): 20
  Set clauses:
  Policy routing matches: 0 packets, 0 bytes
IE-7200-3#

IE-7200-3# show ip bgp route-map my_routes
```

```
BGP table version is 37619, local router ID is 198.133.219.2
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop            Metric LocPrf Weight Path
* i198.133.219.0    64.104.20.4              0    100      0 i
*>                  0.0.0.0                  0         32768 i
IE-7200-3#
```

The BGP updates between the SPs and the border routers should be authenticated using passwords. The following is a example of the required command:

```
neighbor 64.104.10.114 password 7 045802150C2E
```

The following is the BGP configuration for the border router.

```
router bgp 30000
 bgp log-neighbor-changes
 neighbor 64.104.10.114 remote-as 30001 ! <----- This is connection to SP1
 neighbor 64.104.10.114 ttl-security hops 2 ! <---- TTL -security feature
 neighbor 64.104.10.114 password 7 045802150C2E ! <---- Password protection
 neighbor 64.104.20.4 remote-as 30000 ! <---- iBGP connection to the other Border router
 maximum-paths ibgp  3 ! <--- Maximum mumber of paths to be allowed.
 !
 address-family ipv4
  neighbor 64.104.10.114 activate
  neighbor 64.104.10.114 route-map my_routes out
  neighbor 64.104.20.4 activate
  neighbor 64.104.20.4 next-hop-self
  maximum-paths ibgp  3
  no auto-summary
  no synchronization
  network 198.133.219.0

route-map my_routes permit 10
 match as-path 20
!
ip as-path access-list 20 permit ^$ ! <-- Permit only if there is no as-path prepend
ip as-path access-list 20 deny .* ! <-- Deny if there is as-path prepend.
```

Infrastructure ACL Implementation

The Infrastructure ACL (iACL) forms the first layer of defense to the Internet edge module. The iACL should be constructed following the best practices outlined in the *Network Security Baseline* document (see Appendix A, "Reference Documents."). The ACL example that follows protects the PIN from several unwanted sources and illustrates how an iACL protects the border routers.

Note The 64.104.0.0/16 and 198.133.219.0/24 address blocks used in the examples provided below are reserved for the exclusive use of Cisco Systems, Inc.

```
! The global address for IE pin is 198.133.219.0. The first three lines prevent fragments
from any source to this address space.
access-list 110 deny   tcp any 198.133.219.0 0.0.0.255 fragments
access-list 110 deny   udp any 198.133.219.0 0.0.0.255 fragments
access-list 110 deny   icmp any 198.133.219.0 0.0.0.255 fragments

access-list 110 deny   ip host 0.0.0.0 any ! prevent traffic from default route
access-list 110 deny   ip 127.0.0.0 0.255.255.255 any ! prevent traffic from host address.
```

```
access-list 110 deny   ip 192.0.2.0 0.0.0.255 any
access-list 110 deny   ip 224.0.0.0 31.255.255.255 any ! prevent traffic from special
multicast address space.
access-list 110 deny   ip 10.0.0.0 0.255.255.255 any ! prevent spoofing traffic from
10.x.x.x address space.
access-list 110 deny   ip 192.168.0.0 0.0.255.255 any ! prevent spoofing traffic from
192.168.x.x address space.
access-list 110 permit tcp host 64.104.10.114 host 64.104.10.113 eq bgp ! permit bgp
traffic only with the known service provider
access-list 110 permit tcp host 64.104.10.114 eq bgp host 64.104.10.113 ! same comment as
above.
access-list 110 deny   ip 198.133.219.0 0.0.0.255 any ! prevent spoofing traffic, which is
traffic orginate from outside using the IE address space.
access-list 110 deny   ip 10.240.0.0 0.15.255.255 any ! deny spoofing management traffic,
10.240.0.0 is the internal management address space.
access-list 110 permit ip any any
```

Remote Access Block

The position of the remote-access network infrastructure within the Internet edge network is shown in Figure 6-25.

Figure 6-25 Remote Access Block in Internet Edge Network

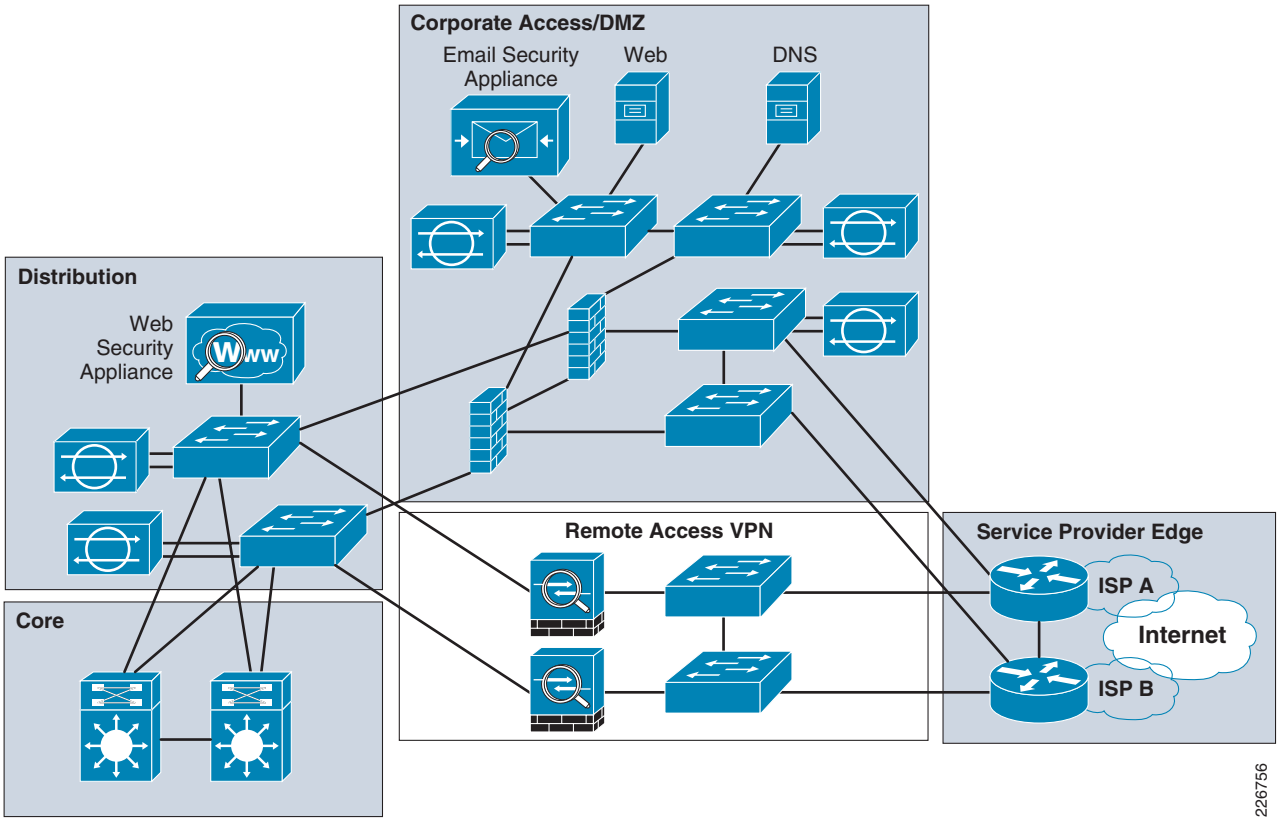

In the Internet edge module, remote access provides access primarily to users connecting to network resources from external locations, such as Internet hot spots, public access, and so on. Remote access does not refer to access from remote offices or teleworkers. To provide access for remote users, there are several technologies available, including Easy VPN, SSL VPN, and Virtual Tunnel Interfaces (VTI). The principal function of remote access is to provide access to internal resources and applications. Common examples are product information pages, blogs, FTP sites, and so on. Remote access functionality is provided by using a separate pair of Cisco ASAs. Figure 6-26 shows the placement of the remote-access firewalls in the context of Internet edge.

Figure 6-26 Placement of Remote Access Firewall

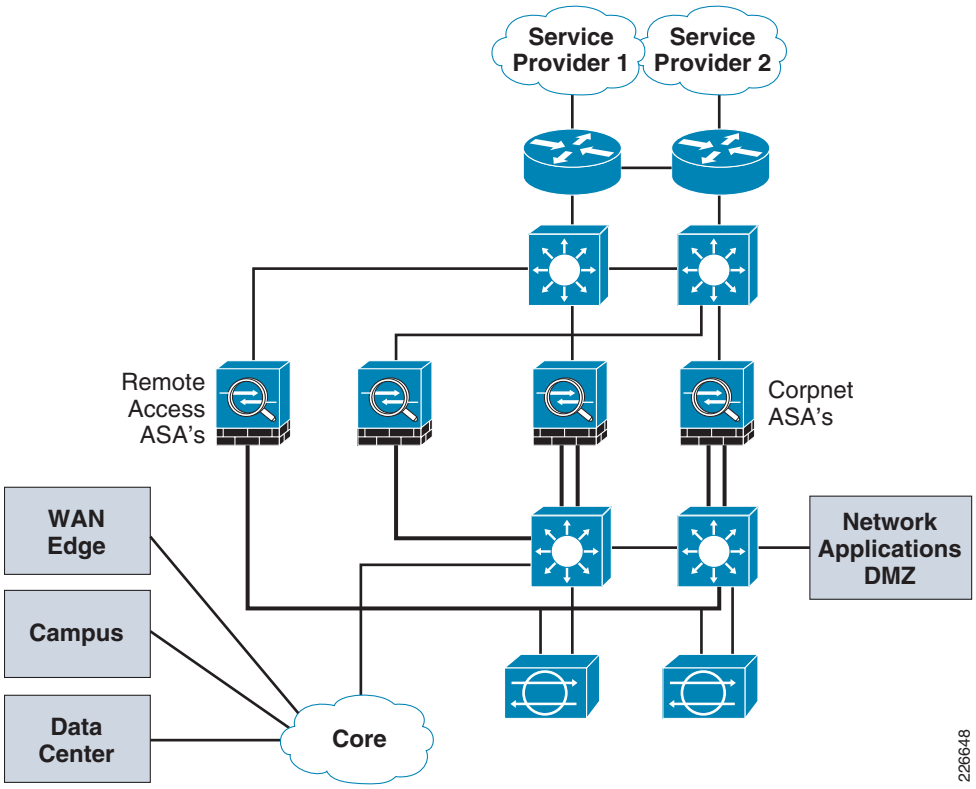

Design Guidelines for the Remote Access Block

This description focuses on an SSL VPN-based implementation. To implement SSL VPN, there are several factors and best practices that are recommended. These can be summarized as follows:

- In simple deployments, the Cisco ASA can issue its own certificate. In a more complex enterprise system, you can use a certificate issued and verified by a third-party vendor.
- Use redundant Cisco ASAs for reliability. In this design, an active/standby scenario is featured.
- It is recommended that the Cisco IPS be used to inspect traffic to or from remote users. Cisco IPS sensors are placed at the distribution block, allowing the inspection of traffic after it is decrypted.
- Use Authentication, Authorization, and Accounting (AAA) for authentication of remote users.

The following configuration steps illustrate some of the practices to implement remote access using SSL VPN.

Note The 64.104.0.0/16 and 198.133.219.0/24 address blocks used in the examples provided below are reserved for the exclusive use of Cisco Systems, Inc.

Step 1 Enable the HTTP server on the Cisco ASA.

```
http server enable
```

Step 2 Configure a different port for management purposes. This is required because WebVPN listens by default on 443. As a result, a separate port is required for management.

```
http redirect management 445
```

Step 3 Enable WebVPN on outside interface.

```
webvpn
 enable VPN-termination
```

Step 4 (Optional) Configure DNS.

```
dns -lookup inside
dns server-group DefaultDNS
 name-server 10.244.30.10
 domain-name cisco.com
```

Step 5 Define a group policy. The following example illustrates creating a group policy named *executive*.

```
group-policy executive internal
group-policy executive attributes
 vpn-simultaneous-logins 25
 vpn-tunnel-protocol webvpn
 default-domain value cisco.com
```

Step 6 Define a tunnel policy. The following configuration illustrates creating a tunnel-policy named *executive-tunnel*.

```
tunnel-group executive-tunnel type remote-access
tunnel-group executive-tunnel general-attributes
 default-group-policy executive
tunnel-group executive-tunnel webvpn-attributes
 group-alias executive enable
```

Step 7 Configure certificates. The SSL gateway uses a certificate as its identity for remote users. The gateway can issue its own certificate and use it as its identity or use a certificate issued by a third-party vendor. For a simple deployment, the gateway can use its own certificate. The following configuration example illustrates configuration of a locally signed certificate:

```
crypto ca trustpoint LOCAL-TP
 revocation-check crl none
 enrollment self
 fqdn IE-SSL-1.cisco.com
 subject-name CN=198.133.219.40
 serial-number
 ip-address 198.133.219.40
 crl configure

route-map my_routes permit 10
 match as-path 20
!
ip as-path access-list 20 permit ^$ !<-- Permit only if there is no as-path prepend
ip as-path access-list 20 deny .* ! <-- Deny if there is as-path prepend.
```

You can use the Cisco Adaptive Security Device Manager (ASDM) tool to configure and monitor the remote-access Cisco ASAs. With Cisco ASDM, you can monitor traffic statistics, look an interface status and monitor events. An example of the Cisco ASDM monitoring capabilities is given in Figure 6-27.

Figure 6-27 ASDM Example Management and Monitoring Screen

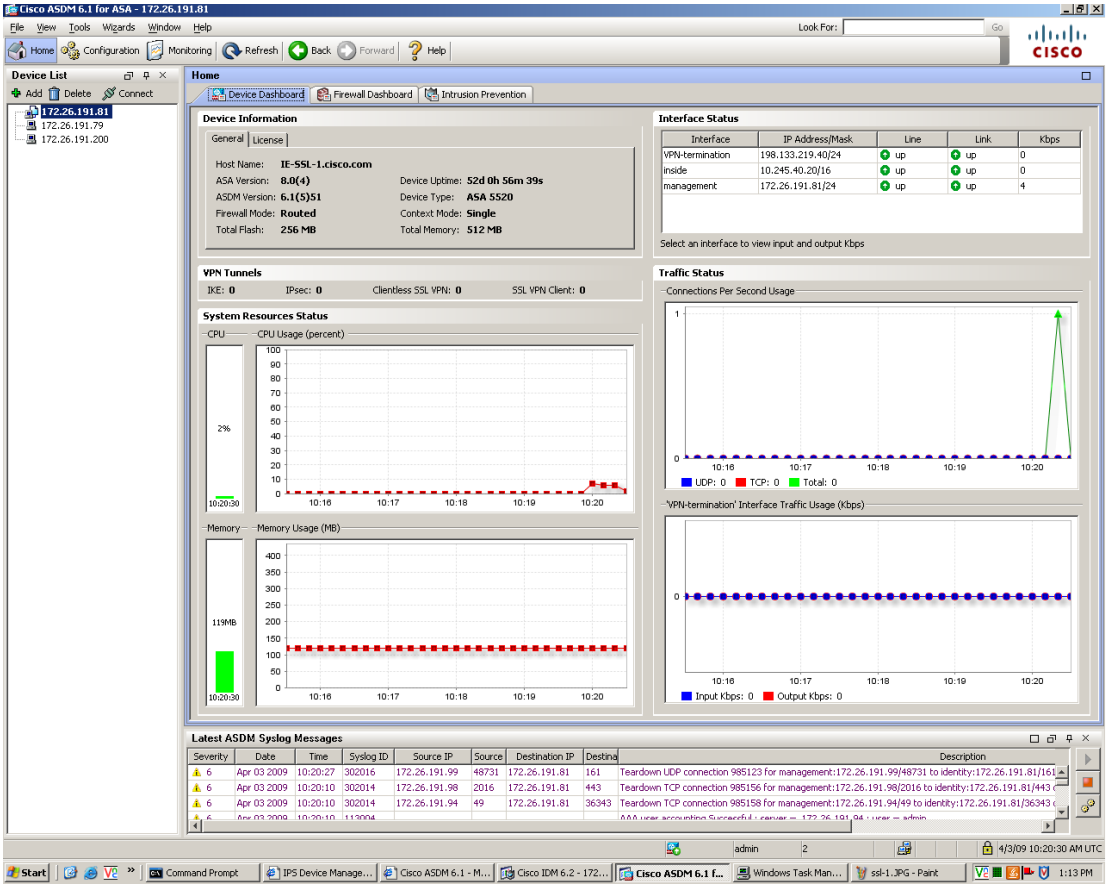

Threats Mitigated in the Internet Edge

The threats mitigated using the various platforms and features described in this chapter are summarized in Table 6-1.

Table 6-1 Internet Edge Threat Mitigation Features

	DDos/DoS/ Worms	Unauthorized Access	Spyware/ Malware/ Phishing/ Spam	Network Abuse/Intrusion	Application Layer Attack	Visibility	Control
Cisco IPS	Yes		Yes	Yes	Yes	Yes	Yes
Firewall	Yes	Yes		Yes		Yes	Yes
IronPort C-Series (ESA)			Yes			Yes	Yes
IronPort S-Series (WSA)	Yes	Yes		Yes		Yes	Yes
Cisco Application Control Engine (ACE) Web Application Firewall					Yes	Yes	Yes
Secure Routing	Yes	Yes		Yes		Yes	Yes
Secure Switching	Yes	Yes		Yes		Yes	Yes