Small Enterprise Design Profile (SEDP)—Security Design

This chapter describes how the Small Enterprise Design Profile sets the foundation for safe and secure enterprise networks by leveraging the proven design and deployment guidelines of the Cisco SAFE security architecture. The Small Enterprise Design Profile is a well-designed and validated network architecture that enables the small enterprise to deliver all of the services required for an enhanced business environment. Cisco SAFE is a security reference architecture that provides detailed design and implementation guidelines for organizations looking to build highly secure and reliable networks.

Small Enterprise Network Security Design

The Small Enterprise Design Profile was designed with built-in security to protect the infrastructure and to provide a safe environment for business. Following the proven guidelines of the Cisco SAFE security architecture, a series of network security technologies and products are strategically deployed throughout the network to protect employees and company assets, to guarantee the confidentiality of sensitive data, and to ensure the availability and integrity of systems and data.

Small enterprises are not immune to violence, theft, vandalism, and other threats. The adoption of new network collaboration and Internet-based technologies also opens the possibility for a number of cyber threats. Understanding the nature and diversity of these threats, and how they may evolve over time, is the first step towards a successful enterprise security strategy. The following are some of the common threats to enterprise environments:

- *Service disruption*—Disruption to the infrastructure, applications and other business resources caused by botnets, worms, malware, adware, spyware, viruses, denial-of-service (DoS) attacks, and Layer 2 attacks.
- *Network abuse*—Use of non-approved applications by employees; peer-to-peer file sharing and instant messaging abuse; and access to non-business related content.
- *Unauthorized access*—Intrusions, unauthorized users, escalation of privileges, and unauthorized access to restricted resources.
- *Data loss*—Theft or leakage of private and confidential data from servers, endpoints, while in transit, or as a result of spyware, malware, key-loggers, viruses, etc.
- *Identity theft and fraud*—Theft of personnel identity or fraud on servers and end users through phishing and E-mail spam.

The Small Enterprise Design Profile is based on a validated network architecture designed around both business operations and technical considerations. Recognizing the fact cost is a common limiting factor to small enterprise network designs, the architecture topologies and platforms were carefully selected to increase productivity while reducing overall costs.

The Small Enterprise Design Profile accommodates a main site and multiple remote sites of various sizes, all interconnected over a Metro Ethernet core. The architecture design is illustrated in Figure 1. At the heart of the architecture is a robust routing and switching network. Operating on top of this network are the all the services used within the enterprise. This often includes the majority of the business most critical applications such as databases, payroll, accounting and customer relationship management (CRM) applications. The core of those services are deployed and managed at the main site, allowing the enterprise to reduce the need for separate services to be operated and maintained at the various remote locations. Centralized systems and applications are served by a main site serverfarm.

Figure 1 Small Enterprise Network Design

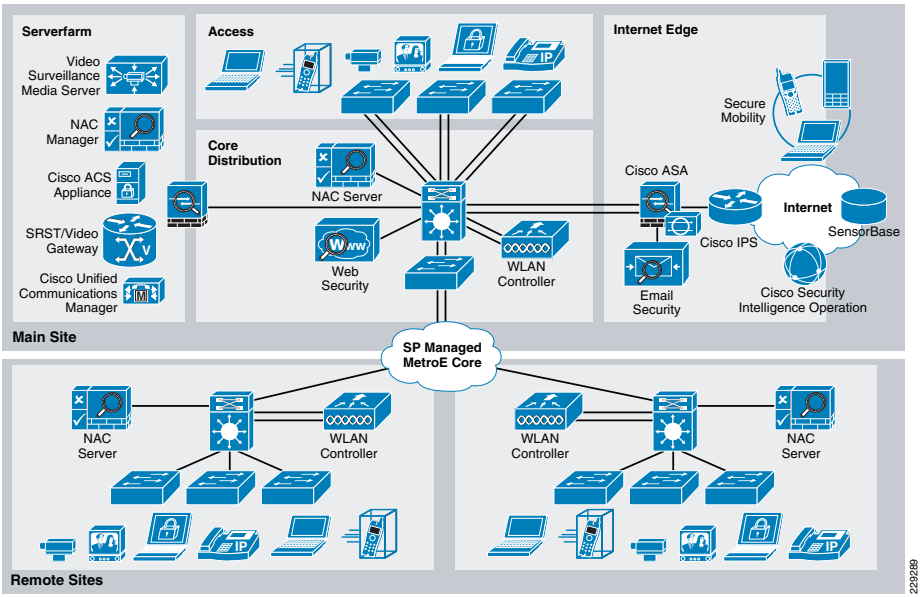

As shown in Figure 1, the small enterprise network design follows a defense-in-depth approach, where multiple layers of protection are built into the architecture. The different security tools are combined together for enhanced visibility and control.

The security design of the architecture focuses on the following key areas:

- *Network Foundation Protection (NFP)*
  - Ensuring the availability and integrity of the network infrastructure, protecting the control and management planes.
- *Internet Perimeter Protection*
  - Ensuring safe Internet connectivity, and protecting internal resources and users from malware, viruses, and other malicious software.
  - Protecting personnel from harmful and inappropriate content.
  - Enforcing E-mail and web browsing policies.
- *Serverfarm Protection*
  - Ensuring the availability and integrity of centralized applications and systems.
  - Protecting the confidentiality and privacy of sensitive data.

- *Network Access Security and Control*
  - Securing the access edges. Enforcing authentication and role-based access for users residing at the main site and remote offices.
  - Ensuring systems are up-to-date and in compliance with the enterprise network security policies.
- *Secure Mobility*
  - Providing secure, persistent connectivity to all mobile employees on laptops, smartphones and other mobile platforms. Enforcing encryption, authentication and role-based access to all mobile users.
  - Delivering consistent protection to all mobile employees from viruses, malware, botnets, and other malicious software.
  - Ensuring a persistent enforcement of enterprise network security policies to all users. Making sure systems comply with corporate policies and have up-to-date security.

The following sections discuss the key areas of the small enterprise network security design.

Network Foundation Protection

Small enterprise networks are built with routers, switches, and other network devices that keep the applications and services running. Therefore, properly securing these network devices is critical for continued operation.

The Small Enterprise Design Profile protects the network infrastructure by implementing the Cisco SAFE best practices for the following areas:

- *Infrastructure device access*
  - Restrict management device access to authorized parties and for the authorized ports and protocols.
  - Enforce Authentication, Authorization and Accounting (AAA) with Terminal Access Controller Access Control System (TACACS+) or Remote Authentication Dial-In User Service (RADIUS) to authenticate access, authorize actions and log all administrative access.
  - Display legal notification banners.
  - Ensure confidentiality by using secure protocols like Secure Shell (SSH) and HTTPS.
  - Enforce idle and session timeouts. Disable unused access lines.
- *Routing infrastructure*
  - Restrict routing protocol membership by enabling Message-Digest 5 (MD5) neighbor authentication and disabling default interface membership.
  - Enforce route filters to ensure that only legitimate networks are advertised; and networks that are not supposed to be propagated are never advertised.
  - Log status changes of neighbor sessions to identify connectivity problems and DoS attempts on routers.
- *Device resiliency and survivability*
  - Disable unnecessary services, implement control plane policing (CoPP).
  - Enable traffic storm control.
  - Implement topological, system and module redundancy for the resiliency and survivability of routers and switches and to ensure network availability.
  - Keep local device statistics.
- *Network telemetry*
  - Enable Network Time Protocol (NTP) time synchronization.
  - Collect system status and event information with Simple Network Management Protocol (SNMP), Syslog, TACACS+/RADIUS accounting.
  - Monitor CPU and memory usage on critical systems.
  - Enable NetFlow to monitor traffic patterns and flows.
- *Network policy enforcement*
  - Implement access edge filtering.
  - Enforce IP spoofing protection with access control lists (ACLs), Unicast Reverse Path Forwarding (uRPF), and IP Source Guard.
- *Switching infrastructure*
  - Implement a hierarchical design, segmenting the LAN into multiple IP subnets or virtual LANs (VLANs) to reduce the size of broadcast domains.
  - Protect the Spanning Tree Protocol (STP) domain with BPDU Guard and STP Root Guard.
  - Use per-VLAN Spanning Tree to reduce the scope of possible damage.
  - Disable VLAN dynamic trunk negotiation on user ports.
  - Disable unused ports and put them into an unused VLAN.
  - Implement Catalyst Infrastructure Security Features (CISF) including port security, dynamic ARP inspection, DHCP snooping.
  - Use a dedicated VLAN ID for all trunk ports.
  - Explicitly configure trunking on infrastructure ports.
  - Use all tagged mode for the native VLAN on trunks and drop untagged frames.
- *Network management*
  - Ensure the secure management of all devices and hosts within the enterprise network.
  - Authenticate, authorize, and keep record of all administrative access.
  - If possible, implement a separate out-of-band (OOB) management network (hardware or VLAN based) to manage systems at the main site.
  - Secure the OOB by enforcing access controls, using dedicated management interfaces or virtual routing and forwarding (VRF) tables.
  - Provide secure in-band management access for systems residing at the remote sites by deploying firewalls and ACLs to enforce access controls, using Network Address Translation (NAT) to hide management addresses, and using secure protocols like SSH and HTTPS.
  - Ensure time synchronization by using NTP.
  - Secure servers and other endpoint with endpoint protection software and operating system (OS) hardening best practices.

For more detailed information on the Network Foundation Protection best practices, refer the "Chapter 2, Network Foundation Protection" of the *Cisco SAFE Reference Guide* at the following URL:

http://www.cisco.com/en/US/docs/solutions/Enterprise/Security/SAFE_RG/chap2.html

Internet Perimeter Protection

The Small Enterprise Design Profile assumes the existence of a centralized Internet connection at the headquarters or main site serving users at all locations. Common services include E-mail and web browsing for employees, the hosting of a company's website accessible to clients and partners over the Internet, and secure remote access for mobile users and remote workers. Other services may also be provided using the same infrastructure.

The network infrastructure that provides Internet connectivity is defined as the Internet perimeter, illustrated in Figure 2.

Figure 2 Internet Perimeter

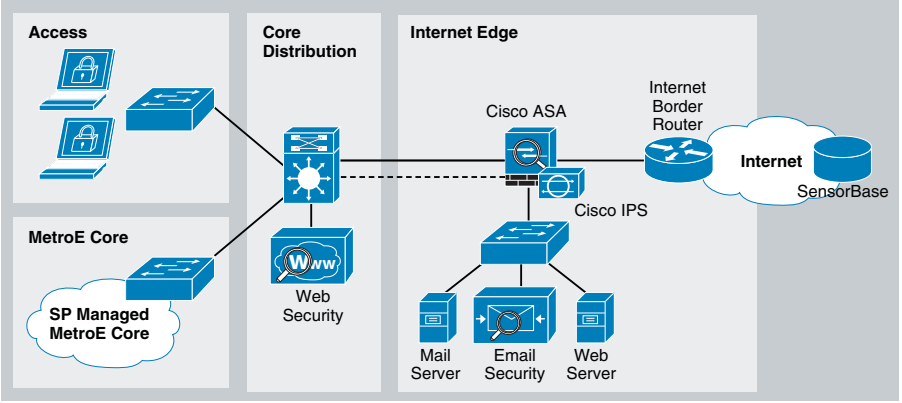

The primary functions of the Internet perimeter is to allow for safe and secure access for users residing at all locations, and to provide public services without compromising the confidentiality, integrity and availability of the company's resources and data. To that end, the Internet perimeter incorporates the following security functions:

- *Internet Border Router*—The Internet border router is the Internet gateway responsible for routing traffic between the enterprise network and the Internet. The Internet border router may be administered by the company's personnel or may be managed by the Internet service provider (ISP). The router provides the first line of protection against external threats and should be hardened following the Network Foundation Protection (NFP) best practices.
- *Internet Firewall*—A Cisco ASA provides stateful access control and deep packet inspection to protect company resources and data from unauthorized access and disclosure. The security appliance is configured to prevent incoming access from the Internet, to protect the enterprise Internet public services, and to control user traffic bound to the Internet. In addition, the Cisco ASA Botnet Traffic Filter feature can be enabled to defend the enterprise against botnet threats. Once enabled, the Botnet Traffic Filter feature monitors network traffic across all ports and protocols for rogue activity and to prevent infected internal endpoints from sending command and control traffic back to external hosts on the Internet. The security appliance may also provide secure remote access to employees with the Cisco AnyConnect Secure Mobility client.
- *Intrusion Prevention*—An Advanced Inspection and Prevention Security Service Module (AIP SSM) on the Cisco ASA or a separate IPS appliance can be implemented for enhanced threat detection and mitigation. The IPS module or appliance is responsible for identifying and blocking anomalous traffic and malicious packets recognized as well-known attacks. IPS can be deployed either in inline or promiscuous mode. The module or appliance may be configured to participate in Cisco IPS Global Correlation, allowing the IPS to gain visibility on global threats as they emerge in the Internet, and to quickly react to contain them. IPS may also be configured to help block certain Internet applications such as AOL Messenger, BitTorrent, Skype, and more.
- *Public Services DMZ*— The company's Internet website, mail server and other public-facing services may be placed on a demilitarized zone (DMZ) for security and control purposes. The DMZ acts as a middle stage between the Internet and company's private resources, preventing external users from directly accessing any internal servers and data. The Internet firewall is responsible for restricting incoming access to the public services and by limiting outbound access from DMZ resources out to the Internet. Systems residing on the DMZ are hardened with endpoint protection software and operating system (OS) hardening best practices.
- *E-mail Security*—A Cisco Ironport C Series E-mail Security Appliance (ESA) is deployed at the DMZ to inspect incoming and outgoing E-mails and eliminate threats such as E-mail spam, viruses and worms. The ESA appliance also offers E-mail encryption to ensure the confidentiality of messages, and data loss prevention (DLP) to detect the inappropriate transport of sensitive information.
- *Web Security*—A Cisco IronPort S Series Web Security Appliance (WSA) is deployed at the distribution switches to inspect HTTP and HTTPS traffic bound to the Internet. This system enforces URL filtering policies to block access to websites containing non-business related content or that are known to be the source of spyware, botnets or other type of malware. The WSA may also be configured to block certain Internet applications such as AOL Messenger, BitTorrent, Skype, etc.

Following are the design guidelines for implementing the security functions.

Internet Border Router Security Guidelines

The Internet border router provides connectivity to the Internet through one or more Internet service providers. The router act as the first line-of-defense against unauthorized access, DDoS, and other external threats. Access control lists (ACLs), uRPF, and other filtering mechanisms may be implemented for anti-spoofing and to block invalid packets. NetFlow, Syslog, SNMP may be used to gain visibility on traffic flows, network activity and system status. In addition, the Internet border router should be secured following the practices explained in the "Network Foundation Protection" section on page -2. This includes restricting and controlling administrative access, protecting the management and control planes, and securing the dynamic exchange of routing information.

The "Internet Border Router Deployment" section on page -21provides an example of Internet edge ACL. For more information on how to configure the Internet border router, refer to the "Chapter 6, Enterprise Internet Edge" of the *Cisco SAFE Reference Guide* at the following URL:

http://www.cisco.com/en/US/docs/solutions/Enterprise/Security/SAFE_RG/chap6.html

Internet Firewall Guidelines

The Cisco ASA deployed at the Internet perimeter is responsible for protecting the enterprise internal resources and data from external threats by preventing incoming access from the Internet; protecting 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 controlling user's Internet-bound traffic.

To that end, the security appliance is configured to enforce access policies, keep track of connection status, and inspect packet payloads following these guidelines:

- Deny any connection attempts originating from the Internet to internal resources and subnets.
- Allow outbound Internet access for users residing at any of the enterprise locations and for the protocols permitted by the organization's policies (i.e., HTTP and HTTPS).
- Allow outbound Internet SSL access for administrative updates, SensorBase, IPS signature updates, etc.
- Allow users access to DMZ services such as company's website, E-mail, and domain name resolution (HTTP, SMTP, POP, IMAP, and DNS).
- Restrict inbound Internet access to the DMZ for the necessary protocols and servers (HTTP to web server, SMTP to the mail transfer agent, DNS to DNS server, etc.).
- Restrict connections initiated from DMZ to the only necessary protocols and sources (DNS from DNS server, SMTP from the mail server, HTTP/SSL from Cisco IronPort ESA).
- Enable stateful inspection for the used protocols to ensure returning traffic is dynamically allowed by the firewall.
- Implement Network Address Translation (NAT) and Port Address Translation (PAT) to shield the internal address space from the Internet.

Figure 3 illustrates the protocols and ports explicitly allowed by the Cisco ASA.

Figure 3 Allowed Protocols and Ports

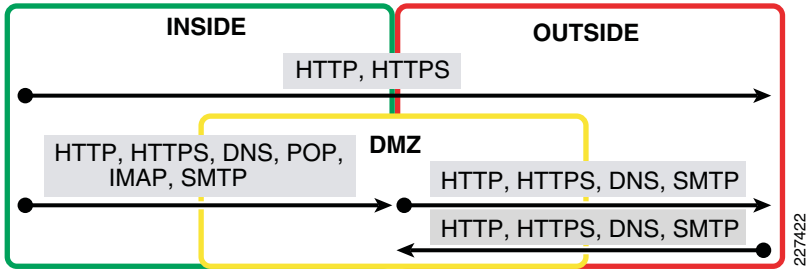

Note Figure 3 does not include any management traffic destined to the firewall. Whenever available, a dedicated management interface should be used. In case the firewall is managed in-band, identify the protocols and ports required prior to configuring the firewall ACLs.

It is also important to remember that the Cisco ASA should be hardened following the Network Foundation Protection best practices. This includes restricting and controlling administrative access, securing the dynamic exchange of routing information with MD5 authentication, and enabling firewall network telemetry with SNMP, Syslog, and NetFlow.

In the Small Enterprise Design Profile, high availability is achieved by using redundant physical interfaces. This represents the most cost-effective solution for high availability. As an alternative, a pair of firewall appliances could be deployed in stateful failover, as discussed in the "Internet Firewall Deployment" section on page -22.

Cisco ASA Botnet Traffic Filter

The Cisco ASA Botnet Traffic Filter feature can be enabled to monitor network ports for rogue activity and to prevent infected internal endpoints from sending command and control traffic back to an external host on the Internet. The Botnet Traffic Filter on the ASA provides reputation-based control for an IP address or domain name, similar to the control that Cisco IronPort SensorBase provides for E-mail and web servers.

The Cisco Botnet Traffic Filter is integrated into all Cisco ASA appliances, and inspects traffic traversing the appliance to detect rogue traffic in the network. When internal clients are infected with malware and attempt to phone home to an external host on the Internet, the Botnet Traffic Filter alerts the system administrator of this through the regular logging process and can be automatically blocked. This is an effective way to combat botnets and other malware that share the same phone-home communications pattern.

The Botnet Traffic Filter monitors all ports and performs a real-time lookup in its database of known botnet IP addresses and domain names. Based on this investigation, the Botnet Traffic Filter determines whether a connection attempt is benign and should be allowed, or is a risk and should be blocked.

The Cisco ASA Botnet Traffic Filter has three main components:

- *Dynamic and administrator blacklist data*—The Botnet Traffic Filter uses a database of malicious domain names and IP addresses that is provided by Cisco Security Intelligence Operations. This database is maintained by Cisco Security Intelligence Operations and is downloaded dynamically from an update server on the SensorBase network. Administrators can also configure their own local blacklists and whitelists.
- *Traffic classification and reporting*—Botnet Traffic Filter traffic classification is configured through the dynamic-filter command on the ASA. The dynamic filter compares the source and destination addresses of traffic against the IP addresses that have been discovered for the various lists available (dynamic black, local white, local black), and logs and reports the hits against these lists accordingly.
- *Domain Name System (DNS) snooping*—To map IP addresses to domain names that are contained in the dynamic database or local lists, the Botnet Traffic Filter uses DNS snooping in conjunction with DNS inspection. Dynamic Filter DNS snooping looks at User Datagram Protocol (UDP) DNS replies and builds a DNS reverse cache (DNSRC), which maps the IP addresses in those replies to the domain names they match. DNS snooping is configured via the Modular Policy Framework (MPF) policies

The Botnet Traffic Filter uses two databases for known addresses. Both databases can be used together, or the dynamic database can be disabled and the static database can be used alone. When using the dynamic database, the Botnet Traffic Filter receives periodic updates from the Cisco update server on the Cisco IronPort SensorBase network. This database lists thousands of known bad domain names and IP addresses.

Note The SensorBase network is an extensive network that monitors global E-mail and web traffic for anomalies, viruses, malware, and other abnormal behavior. The network is composed of the Cisco IronPort appliances, Cisco ASA, and Cisco IPS appliances and modules installed in more than 100,000 organizations worldwide, providing a large and diverse sample of Internet traffic patterns.

The Cisco ASA uses the dynamic database as follows:

1. When the domain name in a DNS reply matches a name in the dynamic database, the Botnet Traffic Filter adds the name and IP address to the DNS reverse lookup cache.
2. When the infected host starts a connection to the IP address of the malware site, the Cisco ASA sends a syslog message reporting the suspicious activity and optionally drops the traffic if the Cisco ASA is configured to do so.
3. In some cases, the IP address itself is supplied in the dynamic database, and the Botnet Traffic Filter logs or drops any traffic to that IP address without having to inspect DNS requests.

The database files are stored in running memory rather than Flash memory. The database can be deleted by disabling and purging the database through the configuration.

Note To use the database, be sure to configure a domain name server for the Cisco ASA so that it can access the URL of the update server. To use the domain names in the dynamic database, DNS packet inspection with Botnet Traffic Filter snooping needs to be enabled; the Cisco ASA looks inside the DNS packets for the domain name and associated IP address.

In addition to the dynamic database, a static database can be used by manually entering domain names or IP addresses (host or subnet) that you want to tag as bad names in a blacklist. Static blacklist entries are always designated with a Very High threat level. Domain names or IP addresses can also be entered in a whitelist,

When a domain name is added to the static database, the Cisco ASA waits one minute, and then sends a DNS request for that domain name and adds the domain name/IP address pairing to the DNS host cache. This action is a background process, and does not affect your ability to continue configuring the Cisco ASA. Cisco also recommends that DNS packet inspection be enabled with Botnet Traffic Filter snooping. When enabled, the Cisco ASA uses Botnet Traffic Filter snooping instead of the regular DNS lookup to resolve static blacklist domain names in the following cases:

- The Cisco ASA DNS server is unavailable.
- A connection is initiated during the one minute waiting period before the Cisco ASA sends the regular DNS request.

If DNS snooping is used, when an infected host sends a DNS request for a name on the static database, the Cisco ASA looks inside the DNS packets for the domain name and associated IP address and adds the name and IP address to the DNS reverse lookup cache.

If Botnet Traffic Filter snooping is not enabled, and one of the above circumstances occurs, that traffic is not monitored by the Botnet Traffic Filter.

Note It is important to realize that a comprehensive security deployment should include Cisco Intrusion Prevention Systems (IPS) with its reputation-based Global Correlation service and IPS signatures in conjunction with the security services provided by the Cisco ASA security appliance such as Botnet Traffic Filter.

For more information on the Cisco ASA Botnet Traffic Filter feature, see the following URL: http://www.cisco.com/en/US/prod/vpndevc/ps6032/ps6094/ps6120/botnet_index.html

Intrusion Prevention Guidelines

IPS is responsible for identifying and blocking anomalous traffic and packets recognized as well-known attacks. An AIP SSM module on the Cisco ASA Internet firewall or a separate IPS appliance can be implemented in the Internet perimeter for enhanced threat detection and mitigation. IPS may also be configured to help block certain Internet applications such as AOL Messenger, BitTorrent, Skype, and so on.

Integrating IPS on a Cisco ASA appliance using an AIP SSM provides a cost-effective solution for small enterprise networks. The AIP SSM is supported on Cisco ASA 5510 and higher platforms. The AIP SSM runs advanced IPS software providing proactive, full-featured intrusion prevention services to stop malicious traffic before it can affect the enterprise network.

The AIP SSM may be deployed in inline or promiscuous mode:

- *Inline mode*—The AIP SSM is placed directly in the traffic flow (see the left side of Figure 4). Traffic identified for IPS inspection cannot continue through the ASA without first passing through and being inspected by the AIP SSM. This mode is the most secure because every packet that has been identified for inspection is analyzed before being allowed through. Also, the AIP SSM can implement a blocking policy on a packet-by-packet basis. This mode, however, can affect throughput if not designed or sized appropriately.
- *Promiscuous mode*—A duplicate stream of traffic is sent to the AIP SSM. This mode is less secure, but has little impact on traffic throughput. Unlike inline mode, in promiscuous mode the AIP SSM can block traffic only by instructing the ASA to shun the traffic or by resetting a connection on the ASA. Also, while the AIP SSM is analyzing the traffic, a small amount of traffic might pass through the ASA before the AIP SSM can shun it. The right side of Figure 4 shows the AIP SSM in promiscuous mode.

Figure 4 IPS Inline and Promiscuous Modes

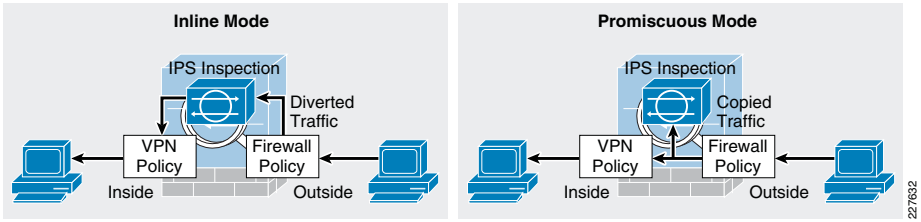

The recommended IPS deployment mode depends on the goals and policies of the enterprise. IPS inline mode is more secure because of its ability to stop malicious traffic in real-time; however, it may impact traffic throughput if not properly designed or sized. Conversely, IPS promiscuous mode has less impact on traffic throughput but is less secure because there may be a delay in reacting to the malicious traffic.

Although the AIP SSM runs as a separate application within the Cisco ASA, it is integrated into the traffic flow. The AIP SSM contains no external interfaces itself, except for the management interface on the SSM itself. When traffic is identified for IPS inspection on the ASA, traffic flows through the ASA and the AIP SSM in the following sequence:

1. Traffic enters the adaptive security appliance.
2. Firewall policies are applied.
3. Traffic is sent to the AIP SSM over the backplane.
4. The AIP SSM applies its security policy to the traffic and takes appropriate actions.
5. Valid traffic (for inline mode only) is sent back to the adaptive security appliance over the backplane; the AIP SSM might block some traffic according to its security policy and that traffic is not passed on.
6. VPN policies are applied (if configured).
7. Traffic exits the adaptive security appliance.

The AIP SSM card may be configured to fail open or close when the module becomes unavailable. When configured to fail open, the ASA allows all traffic through, uninspected, if the AIP SSM becomes unavailable. Conversely, when configured to fail close, the ASA blocks all traffic in case of an AIP SSM failure.

Cisco IPS Global Correlation

If desired, the AIP SSM module on the Cisco ASA (or IPS appliance) may participate in Cisco IPS Global Correlation for further threat visibility and control. Once enabled, the participating IPS sensor receives threat updates from the Cisco SensorBase Network at regular intervals. The Cisco SensorBase Network contains detailed information about known threats on the Internet, including serial attackers, botnet harvesters, malware outbreaks, and dark nets. The IPS uses this information to filter out the worst attackers before they have a chance to attack critical assets. It then incorporates the global threat data into its system to detect and prevent malicious activity even earlier.

IPS Global Correlation is an important improvement in the basic functions of IPS because it enables the system to understand the world in which it operates: an understanding of who the attacker is and whether the attacker has a record of bad behavior. With Global Correlation, the sensor does not have to rely on just the data in the packet or connection to make a decision about the intent of the activity and determine whether the activity is malicious. Now, the sensor can look at a ping sweep and know that the source of the ping sweep does not have a negative reputation, but later can look at another ping sweep and see that the source is a known malicious site with a history of web attacks, and the sensor can block access to and from that site. Global Correlation gives users greater confidence in the actions the sensor takes because these actions are applied to attackers that have shown a predisposition for malicious behavior.

Global Correlation provides a process through which security data is collected for IP addresses and a reputation score is developed for each IP address globally by Cisco. Cisco IPS 7.0 uses this reputation data in two ways: for its reputation filters and for Global Correlation inspection.

- Reputation filters are used to block a subset of IP networks that are owned wholly by malicious groups or were unused and have been hijacked. This first line of defense helps prevent malicious contact ranging from spam to intelligence gathering in preparation for directed attacks. Reputation filters also prevent attempts by botnets to phone home if the botnet controller machine resides in one of these networks.
- Global Correlation inspection uses reputation scores for normal IP addresses to increase the percentage of attacks that the sensor can block. First, the sensor must detect some sort of malicious activity and fire an event as a result. When an event is triggered, that event is processed to determine whether the attacker's IP address has a negative reputation and to what degree. If the event is sourced from an attacker with a negative reputation, the sensor will add risk to the event, raising its risk rating and making it more likely that the sensor will deny the event. This enables the sensor to deny packets and attackers based on the fact that the event has a negative reputation in addition to a high risk rating calculated on the sensor.

Once Global Correlation is configured, the IPS works in the following manner:

1. When a packet enters the sensor, the first check is against the preprocessor, which performs Layer-2 packet normalization and atomic signature checks. Here, the packet header is processed to help ensure that the packet is an IP packet, that the header is not incorrectly formed, and that the packet does not match any atomic signatures.
2. The packet is sent through the Cisco IPS reputation filters.
   a. Packets that match are discarded immediately, assuming that the reputation filters are enabled and not in Audit mode.
   b. Packets that do not match go to the inspection engines, starting with the Layer 3 and 4 normalization engine, then all the signature engines, and then anomaly detection.
3. If any events are triggered, alerts are sent to the Global Correlation inspection processor, where the source IP address for any alert is checked for negative reputation, and the risk rating is modified and actions are added as appropriate.
4. Any actions assigned to alerts are processed and acted upon, including event action overrides to add new actions and event action filters to remove actions.

Reputation Filters

Cisco IPS reputation filters use a list of hundreds of networks that can be safely blocked because they are not owned by any legitimate source. The reputation of the networks on this list can be considered to be –10. This list includes only IP networks consisting entirely of stolen, "zombie" address blocks and address blocks controlled entirely by malicious organizations. Individual IP addresses are never found on this list. Because there is no way that a legitimate IP address can go from a positive or neutral reputation and then, because of malicious activity, earn a place on the Cisco IPS reputation filter list, users can confidently block all activity to and from networks on this list.

The primary purpose of the IPS reputation filters is to provide protection from direct scanning, botnet harvesting, spamming, and distributed denial-of-service (DDoS) attacks originating from these malicious address blocks, and from connections being attempted back to these networks from systems already infected.

Note There is currently no capability to view the networks on this list, but the networks that are being blocked get logged by the sensor in the Statistics section of Cisco IPS Manager Express (IME).

The only user configuration required for reputation filters is enabling or disabling them and specifying whether Global Correlation is set to Audit mode (a global configuration setting for the entire sensor). In Audit mode, the sensor will report potential deny actions due to reputation filters instead of actually denying the activity.

Global Correlation Inspection

The primary activity of an IPS sensor is detection of malicious behavior. After the packet goes through the IPS reputation filter process, the signature inspection occurs. This involves inspection of packets flowing through the sensor by the various engines looking for the various types of malicious behavior. Alerts that are created are passed to the Global Correlation inspection process for reputation lookups.

When an event occurs, the Global Correlation inspection process performs a local lookup of the source (attacker) IP address of the event in its reputation database. This lookup process returns a value ranging from –.1 to –10; the more negative the value, the more negative the reputation of the source IP address. This reputation score is calculated for Cisco IPS sensors using the data in Cisco SensorBase and is sent to the sensor as a reputation update. If an IP address returns no value for reputation, than it is considered to be neutral. Cisco IPS, unlike E-mail and web security reputation applications, has no concept of positive reputation. When an event is checked for reputation, this checking occurs entirely on the sensor using data downloaded previously from Cisco SensorBase. Unlike other devices, the sensor will not send a live request for information about an IP address that it has just seen. It will look in the data that it has, and if it finds the address, it will use that data; otherwise, the sensor will assume that the address has a neutral reputation.

Global Correlation inspection has three modes of primary operation: Permissive, Standard (default), and Aggressive; you can also select Off:

- Permissive mode tells the sensor to adjust the risk rating of an event, but not to assign separate reputation only actions to the event.
- Standard mode tells the sensor to adjust the risk rating and to add a Deny Packet action due to reputation if the risk rating is greater than or equal to 86. It also adds a Deny Attacker action due to reputation if the risk rating is greater than or equal to 100.
- Aggressive mode also adjusts the risk rating due to reputation, adds a Deny Packet action due to reputation if the risk rating is greater than or equal 83, and adds a Deny Attacker action due to reputation if the risk rating is greater than or equal to 95.
- Selecting Off in the Global Correlation Inspection window prevents the sensor from using updates from Cisco SensorBase to adjust reputation.

If Global Correlation inspection is enabled and an event is generated by an attacker with a negative reputation, the risk rating for the event will be elevated by a certain amount that is determined by a statistical formula. The amount by which the risk rating is raised depends on the original risk rating of the event and the reputation of the attacker.

Network Participation and Correlation Updates

The IPS sensor pulls reputation information for addresses on the global Internet from Cisco SensorBase. When the sensor is configured initially, a DNS server needs to be configured for the sensor to use to connect to Cisco SensorBase or an HTTP or HTTPS proxy (that has DNS configured) needs to be configured. After the sensor has this information, the sensor will make an outbound connection to check for the latest updates from Cisco SensorBase. It will initiate an HTTPS request to Cisco SensorBase update servers and download a manifest that contains the latest versions of the files related to Global Correlation. The sensor will check Cisco SensorBase every 5 minutes for updates. If changes are needed, the sensor will perform a DNS lookup of the server name returned in the initial request. This lookup will return the location of the server nearest to the sensor. The sensor will then initiates an HTTP connection that will actually transfer the data. The size of a full update is about 2 MB; incremental updates average about 100 KB. If a sensor loses connection to Cisco SensorBase, Global Correlation information will begin to timeout within days, and sensor health will change accordingly.

The other component of Global Correlation is network participation. This feature sends data from events that the sensor fires back to Cisco SensorBase to adjust the reputation of IP addresses; this information is then packaged in future reputation data downloads from Cisco SensorBase. The sensor passes this information back to Cisco SensorBase according to the sensor configuration. The possible configuration options are Off, Partial, and Full.

- With the Off (default) setting, the sensor will not send back any data. The sensor will still receive reputation data, and this setting does not affect its use of that data except that the reputations of addresses attacking the network being protected will not be influenced by their generation on the sensor.
- With the Partial setting, the sensor will send back alert information. This information consists of protocol attributes such as the TCP maximum segment size and TCP options string, the signature ID and risk rating of the event, the attacker IP address and port, and Cisco IPS performance and deployment mode information.
- The Full setting adds victim IP address and port information to the information reported with the Partial setting.

Note No actual packet content information is sent to Cisco. In addition, events having RFC 1918 addresses, because they are not unique, are not considered interesting. So all events reported to Cisco SensorBase will have any such IP address information stripped from the reported data.

The mechanism used to update Cisco SensorBase with new attack information is fairly straightforward. The sensor takes event information, parses out the important pieces of data, and buffers this data for transmission back to Cisco SensorBase. It sends this data in the form of an HTTPS connection that it initiates on average every 10 minutes. The average size of an update is 2 to 4 KB, with weekly averages of about 0.5 to 1 MB. Some higher-volume sensors have average update sizes of about 50 KB, with weekly totals in the 45-MB range. Sensors with very high alert volumes can have average update sizes of about 850 KB, with weekly totals of up to 900 MB; these sensors, though, are at the extreme end of the range.

For more information on IPS Global Correlation including configuration information, see the following URL:
http://www.cisco.com/en/US/docs/security/ips/7.0/configuration/guide/cli/cli_collaboration.html.

E-mail Security Guidelines

The Small Enterprise Design Profile implements a Cisco Ironport C Series E-mail Security Appliance (ESA) at the DMZ with the purpose of inspecting E-mails and eliminating threats such as E-mail spam, viruses, and worms. The ESA can be described as a firewall and threat monitoring system 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, as illustrated in Figure 5. Upon reception, E-mails are evaluated using a reputation score mechanism based on the SensorBase network. The SensorBase network is an extensive network that monitors global E-mail and web traffic for anomalies, viruses, malware and other and abnormal behavior. The network is composed of Cisco IronPort appliances, Cisco ASA and Cisco IPS appliances and modules installed in more than 100,000 organizations worldwide, providing a large and diverse sample of Internet traffic

patterns. By leveraging the SensorBase Network, messages originating from domain names or servers known to be the source of spam or malware, and therefore with a low reputation score, are automatically dropped or quarantined by preconfigured reputation filters. Optionally, the enterprise may choose to implement some of the other functions offered by the ESA appliance, including anti-virus protection with virus outbreak filters and embedded anti-virus engines (Sophos and McAfee), encryption to ensure the confidentiality of messages, and data loss prevention (DLP) for E-mail to detect the inappropriate transport of sensitive information.

Note Alternatively, Cisco offers managed hosted and hybrid hosted E-mail security services. These services are provided through a dedicated E-mail infrastructure hosted in a network of Cisco data centers. For more information, refer to http://www.cisco.com/go/designzone.

Figure 5 E-mail Delivery Chain

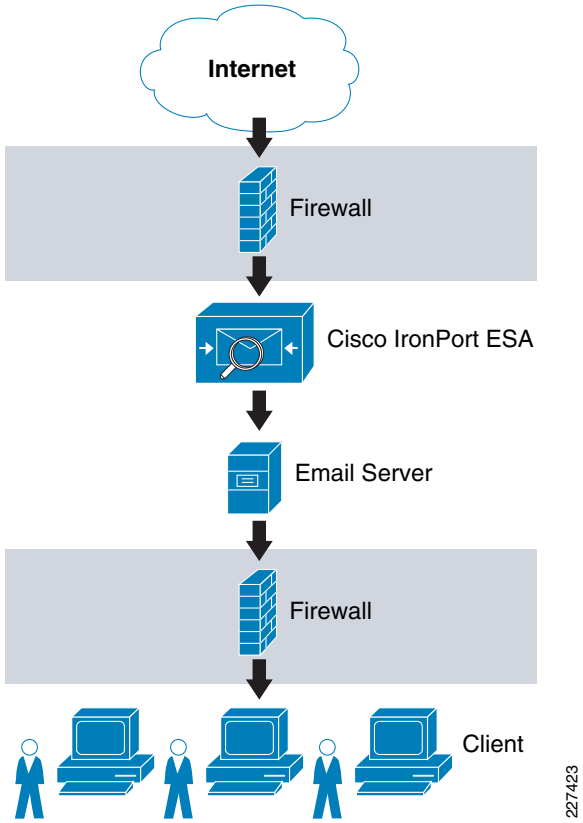

Note Figure 5 shows a logical implementation of a DMZ hosting the E-mail server and ESA appliance. This can be implemented physically by either using a single firewall or two firewalls in sandwich.

There are multiple deployment approaches for the security appliance depending on the number of interfaces used (see Figure 6):

- *Dual-armed configuration*—Two physical interfaces used to serve a public mail listener and a private mail listener, each one configured with a separate logical IP address. The public listener receives E-mail from the Internet and directs messages to the internal E-mail servers; while the private listener receives E-mail from the internal servers and directs messages to the Internet. The public listener interface may connect to the DMZ, while the private listener interface may connect to the inside of the firewall.
- *One-armed configuration*—A single ESA interface configured with a single IP address and used for both incoming and outgoing E-mail. A public mail listener is configured to receive and relay E-mail on that interface. The best practice is to connect the ESA interface to the DMZ where the E-mail server resides.

For simplicity, the Small Enterprise Design Profile implements the ESA appliance with a single interface. In addition, using a single interface leaves other data interfaces available for redundancy.

Figure 6 Common ESA Deployments

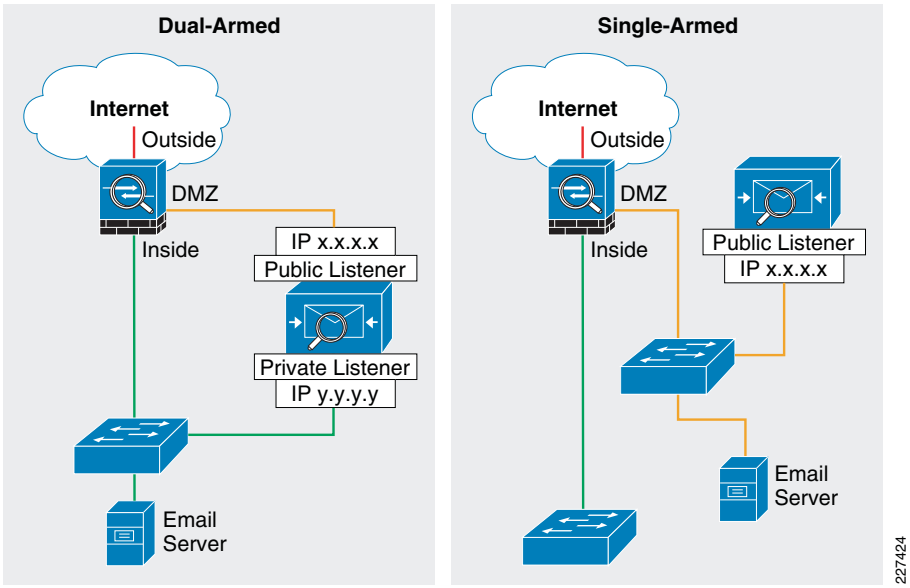

Figure 7 illustrates the logical location of the ESA within the E-mail flow chain.

Figure 7 Typical Data Flow for Inbound E-mail Traffic

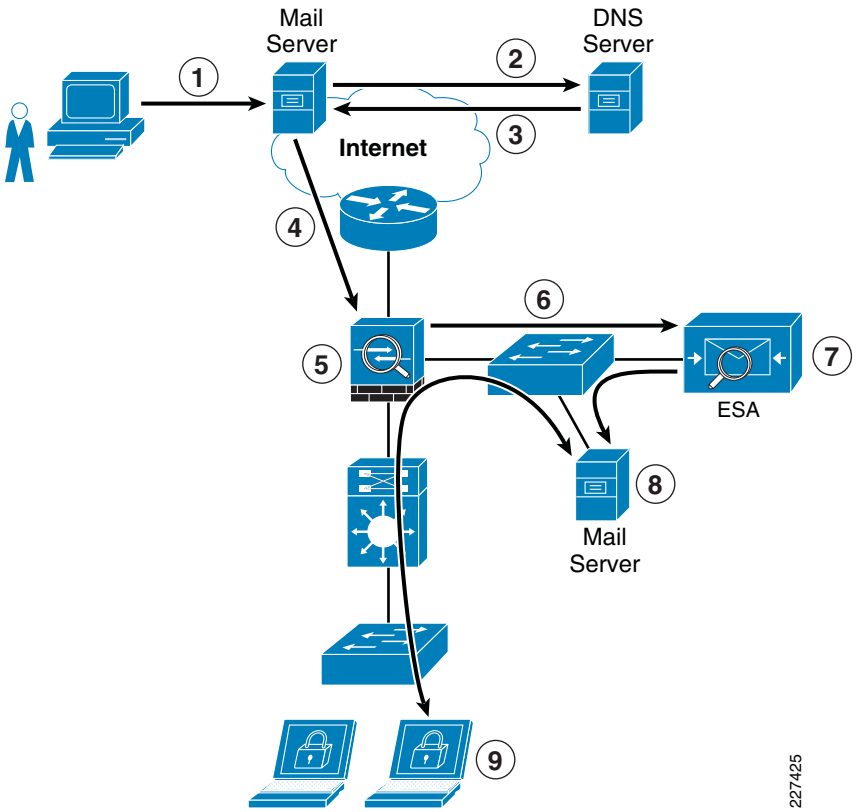

The following steps explain what is taking place in Figure 7:

1. Sender sends an E-mail to xyz@domain X.
2. What's the IP address of domain X?
3. It's a.b.c.d (public IP address of ESA).
4. E-mail server sends message to a.b.c.d using SMTP.
5. Firewall permits incoming SMTP connection to the ESA, and translates its public IP address.
6. ESA performs a DNS query on sender domain and checks the received IP address in its reputation database, and drops, quarantines E-mail based on policy.
7. ESA forwards E-mail to preconfigured inbound E-mail server.
8. E-mail server stores E-mail for retrieval by receiver.
9. Receiver retrieves E-mail from server using POP or IMAP.

The Cisco IronPort ESA appliance functions as a SMTP gateway, also known as a mail exchange (MX). The following are the key deployment guidelines:

- Ensure that the ESA appliance is both accessible via the public Internet and is the first hop in the E-mail infrastructure. If you allow another MTA to sit at your network's perimeter and handle all external connections, then the ESA appliance will not be able to determine the sender's IP address. The sender's IP address is needed to identify and distinguish senders in the Mail Flow Monitor, to query the SensorBase Reputation Service for the sender's SensorBase Reputation Service Score (SBRS), and to improve the efficacy of the anti-spam and virus outbreak filters features.
- Features like Cisco IronPort Anti-Spam, Virus Outbreak Filters, McAfee Antivirus and Sophos Anti-Virus require the ESA appliance to be registered in DNS. To that end, create an A record that maps the appliance's hostname to its public IP address, and an MX record that maps the public domain to the appliance's hostname. Specify a priority for the MX record to advertise the ESA appliance as the primary (or backup during testing) MTA for the domain. A static address translation entry needs to be defined for the ESA public IP address on the Internet firewall if NAT is configured.
- Add to the Recipient Access Table (RAT) all the local domains for which the ESA appliance will accept mail. Inbound E-mail destined to domains not listed in RAT will be rejected. External E-mail servers connect directly to the ESA appliance to transmit E-mail for the local domains, and the ESA appliance relays the mail to the appropriate groupware servers (for example, Exchange™, Groupwise™, and Domino™) via SMTP routes.
- For each private listener, configure the Host Access Table (HAT) to indicate the hosts that will be allowed to send E-mails. The ESA appliance accepts outbound E-mail based on the settings of the HAT table. Configuration includes the definition of Sender Groups associating groups or users, upon which mail policies can be applied. Policies include Mail Flow Policies and Reputation Filtering. Mail Flow Policies are a way of expressing a group of HAT parameters (access rule, followed by rate limit parameters and custom SMTP codes and responses). Reputation Filtering allows the classification of E-mail senders and to restrict E-mail access based on sender's trustworthiness as determined by the IronPort SensorBase Reputation Service.
- Define SMTP routes to direct E-mail to appropriate internal mail servers.
- If an out-of-band (OOB) management network is available, use a separate interface for administration.

A failure on the ESA appliance may cause service outage; therefore, a redundant design is recommended. There are multiple ways to implement redundancy:

- *IronPort NIC Pairing*—Redundancy at the network interface card level by teaming two of the Ethernet interfaces on the ESA appliance. If the primary interface fails, the IP addresses and MAC address are assumed by the secondary.
- *Multiple MTAs*—Consists of adding a second ESA appliance or MTA with an equal cost secondary MX record.
- *Load Balancer*—A load balancer such as Cisco ACE Application Control Engine (ACE) load-balances traffic across multiple ESA appliances.

IronPort NIC pairing is the most cost-effective solution (see Figure 8), because it does not require the implementation of multiple ESA appliances and other hardware. It does not, however, provide redundancy in case of chassis failure.

Figure 8 Cisco IronPort ESA NIC Pairing

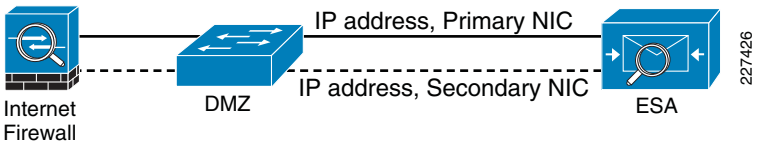

Finally, the Internet firewall should be configured to accommodate traffic to and from the Cisco IronPort ESA deployed at the DMZ. The protocols and ports to be allowed vary depending on the services configured on the appliance. For details, refer to the *Cisco IronPort User's Guide* at the following URL: http://www.ironport.com/support/

The following are some of the most common services required:

- Outbound SMTP (TCP/25) from ESA to any Internet destination
- Inbound SMTP (TCP/25) to ESA from any Internet destination
- Outbound HTTP (TCP/80) from ESA to downloads.ironport.com and updates.ironport.com
- Outbound SSL (TCP/443) from ESA to updates-static.ironport.com and phonehome.senderbase.org
- Inbound and Outbound DNS (TCP and UDP port 53)
- Inbound IMAP (TCP/143), POP (TCP/110), SMTP (TCP/25) to E-mail server from any internal client

Also remember that if the ESA is managed in-band, appropriate firewall rules need to be configured to allow traffic such as SSH, NTP, and syslog.

For more information on how to configure the Cisco Ironport ESA, see the following guides:

- *Cisco SAFE Reference Guide*—http://www.cisco.com/en/US/docs/solutions/Enterprise/Security/SAFE_RG/SAFE_rg.html
- *Cisco IronPort ESA User Guide*—http://www.ironport.com/support

Web Security Guidelines

The Small Enterprise Design Profile implements a Cisco IronPort S Series Web Security Appliance (WSA) to block access to sites with non-business related content, and to protect the enterprise from web-based malware and spyware.

Cisco IronPort WSA's protection relies in two independent services:

- *Web Proxy*—This provides URL filtering, web reputation filters, and optionally anti-malware services. The URL filtering capability defines the handling of each web transaction based on the URL category of the HTTP requests. Leveraging the SensorBase network, the web reputation filters analyze the web server behavior and characteristics to identify suspicious activity and protect against URL-based malware. The anti-malware service leverages anti-malware scanning engines such as Webroot and McAfee to monitor for malware activity.
- *Layer 4 Traffic Monitoring (L4TM)*—Service configured to monitor all Layer-4 traffic for rogue activity and to detect infected clients.

Note The SensorBase network is an extensive network that monitors global E-mail and web traffic for anomalies, viruses, malware, and other abnormal behavior. The network is composed of the Cisco IronPort appliances, Cisco ASA, and Cisco IPS appliances and modules installed in more than 100,000 organizations worldwide, providing a large and diverse sample of Internet traffic patterns.

As the small enterprise network design assumes a centralized Internet connection, the WSA is implemented at the core/distribution layer of the main site network. This allows the inspection and enforcement of web access policies to all users residing at any of the enterprise locations. Logically, the WSA sits in the path between web users and the Internet, as shown in Figure 9.

Figure 9 Cisco IronPort WSA

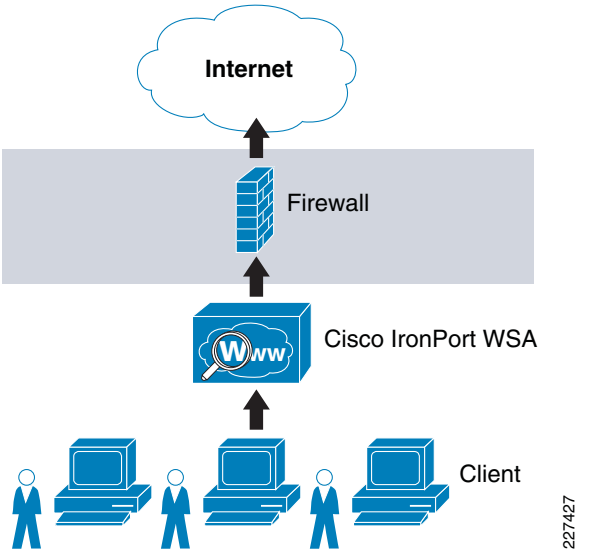

There are two deployment modes for the Web Proxy service:

- *Explicit Forward Proxy*—Client applications, such as web browsers, are aware of the Web Proxy and must be configured to point to the WSA. The web browsers can be either configured manually or by using Proxy Auto Configuration (PAC) files. The manual configuration does not allow for redundancy, while the use of PAC files allows the definition of multiple WSAs for redundancy and load balancing. If supported by the browser, the Web Proxy Autodiscovery Protocol (WPAD) can be used to automate the deployment of PAC files. WPAD allows the browser to determine the location of the PAC file using DHCP and DNS lookups.
- *Transparent Proxy*—Client applications are unaware of the Web Proxy and do not have to be configured to connect to the proxy. This mode requires the implementation of a Web Cache Communications Protocol (WCCP) enabled device or a Layer-4 load balancer in order to intercept and redirect traffic to the WSA. Both deployment options provide for redundancy and load balancing.

Explicit forward proxy mode requires the enterprise to have control over the configuration of the endpoints, which may not be always possible. For example, the enterprise may allow the use of personal laptops, smart-phones and other devices outside the company's

administration. Transparent proxy mode, on the other hand, provides a transparent integration of WSA without requiring any configuration control over the endpoints. It also eliminates the possibility of users reconfiguring their web browsers to bypass the appliance without knowledge of the administrators. For these reasons, the Small Enterprise Design Profile implements transparent proxy with WCCP. In this configuration, the Cisco ASA at the Internet perimeter is leveraged as a WCCP server while the WSA act as a WCCP Traffic Processing Entity.

The Cisco ASA uses WCCP version 2, which has a built-in failover and load balancing mechanism. Per WCCPv2 specification, multiple appliances (up to 32 entities) can be configured as part of the same service group. HTTP and HTTPS traffic is load-balanced across the active appliances based on source and destination IP addresses. The server (Cisco ASA) monitors the availability of each appliance in the group, and can identify appliance failures within 30 seconds. After failure, traffic is redirected across the remaining active appliances. In the case where no appliances are active, WCCP takes the entire service group offline and subsequent requests bypass redirection. In addition, WCCPv2 supports MD5 authentication for the communications between WCCP server and WSA appliances.

Note In the event the entire service group fails, WCCP automatically bypasses redirection, allowing users to browse the Internet without the Web controls. In case it is desired to handle a group failure by blocking all traffic, an outbound ACL may be configured on the Cisco ASA outside interface to permit HTTP/HTTPS traffic originated from the WSA appliance itself and to block any direct requests from clients. The ACL may also have to be configured to permit HTTP/HTTPS access from IPS and other systems requiring such access.

WCCPv2 supports Generic Route Encapsulation (GRE) and Layer-2-based redirection; however, the Cisco ASA only supports GRE. In addition, WCCP is supported only on the ingress of an interface. The only topology supported is one where both clients and WSA are reachable from the same interface, and where the WSA can directly communicate with the clients without going through the Cisco ASA. For these reasons, the WSA appliance is deployed at the inside segment of the Cisco ASA.

Figure 10 illustrates the how WCCP redirection works in conjunction with Cisco ASA.

Figure 10 WCCP Redirection

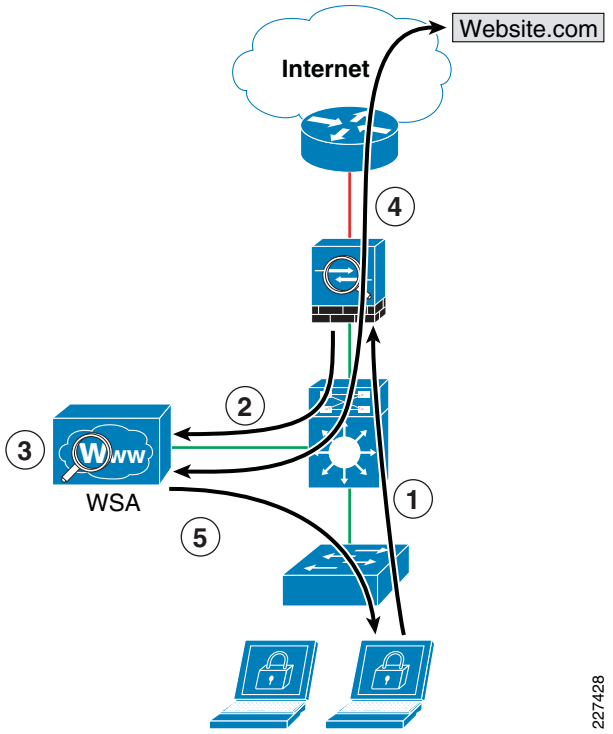

The following steps describe what takes place in Figure 10:

1. Client's browser requests connection to http://website.com.
2. Cisco ASA intercepts and redirects HTTP requests over GRE.
3. If content not present in local cache, WSA performs a DNS query on destination domain and checks the received IP address against URL and reputation rules, and allows/denies request accordantly.
4. WSA fetches content from destination web site.
5. Content is inspected and then delivered directly to the requesting client.

The WSA appliance may also be configured to control and block peer-to-peer file-sharing and Internet applications such as AOL Messenger, BitTorrent, Skype, Kazaa, etc. The way WSA handles these applications depends on the TCP port used for transport:

- *Port 80*—Applications that use HTTP tunneling on port 80 can be handled by enforcing access policies within the web proxy configuration. Application access may be restricted based on applications, URL categories, and objects. Applications are recognized and blocked based on their user agent pattern, and by the use of regular expressions. The user may also specify categories of URL to block, including the predefined *chat* and *peer-to-peer* categories. Custom URL categories may also be defined. Peer-to-peer access may also be filtered based on object and MIME Multipurpose Internet Mail Extensions (MIME) types.

- *Ports other than 80*—Applications using ports other than 80 can be handled with the L4TM feature. L4TM blocks access to a specific application by preventing access to the server or block of IP addresses to which the client application must connect.

Note The Cisco IPS appliances and modules, and the Cisco ASA (using the modular policy framework), may also be used to block peer-to-peer file sharing and Internet applications.

The following are the guidelines for implementing a Cisco IronPort WSA appliance with WCCP on a Cisco ASA:

- Deploy WSA on the inside of the firewall so that the WSA can communicate with the clients without going through the firewall.
- Implement MD5 authentication to protect the communications between the Cisco ASA and the WSA(s).
- Configure a redirect-list on the firewall to indicate what traffic needs to be redirected. Make sure the WSA is always excluded from redirection.
- Ingress ACL on the firewall takes precedence over WCCP redirection, so make sure the ingress ACL is configured to allow HTTP and HTTPS traffic from clients and the WSA itself.
- In an existing proxy environment, deploy the WSA downstream from the existing proxy servers (closer to the clients).
- Cisco ASA does not support WCCP IP source address spoofing, therefore any upstream authentication or access controls based on client IP addresses are not supported. Without IP address spoofing, requests originating from a client are sourced with the IP address of the Web Proxy, and not the one of the client.
- TCP intercept, authorization, URL filtering, inspect engines, and IPS features do not apply to redirected flows of traffic served by the WSA cache. Content requested by the WSA is still subject to all the configured features on the firewall.
- Configure WSA access policies to block access to applications (AOL Messenger, Yahoo Messenger, BitTorrent, Kazaa, etc.) and URL categories not allowed by the enterprise Internet access policies.
- If an out-of-band (OOB) management network is available, use a separate interface for administration.

Note WCCP, firewall, and other stateful features usually require traffic symmetry, whereby traffic in both directions should flow through the same stateful device. The Small Enterprise Design Profile is designed with a single Internet path ensuring traffic symmetry. Care should be taken when implementing active-active firewall pairs as they may introduce asymmetric paths.

The Layer-4 Traffic Monitor (L4TM) service is deployed independently from the Web Proxy functionality, and its mission is to monitor network traffic for rogue activity and for any attempts to bypass port 80. L4TM works by listening to all UDP and TCP traffic and by matching domain names and IP addresses against entries in its own database tables to determine whether to allow incoming and outgoing traffic. The L4TM internal database is continuously updated with matched results for IP addresses and domain names. Additionally, the database table receives periodic updates from the IronPort update server (https://update-manifests.ironport.com).

For more information on how to configure the L4TM feature on Cisco Ironport WSA, see the following guides:

- *Cisco SAFE Reference Guide*—http://www.cisco.com/en/US/docs/solutions/Enterprise/Security/SAFE_RG/SAFE_rg.html
- *Cisco IronPort WSA User Guide*—http://www.ironport.com/support

Configuration steps and examples are included in the "Web Security Deployment" section on page -31.

Serverfarm Protection

Small enterprise networks typically include a serverfarm at the main site that hosts the systems that serve business applications and store the data accessible to internal users. The infrastructure supporting it may include application servers, the storage media, routers, switches, load balancers, off-loaders, application acceleration devices and other systems. In addition, they may also host foundational services as part of the enterprise network such as identity and security services, unified communication services, mobility services, video services, partner applications, and other services.

Depending on the size of the enterprise network, the serverfarm may be constructed following different design models. Figure 11 illustrates a collapsed design, and the less common for small enterprise networks, multi-tier design. In the collapsed design all services are hosted in a shared physical server-farm, and high availability is achieved by using redundant processors and interfaces. Large enterprises may implement a more scalable multi-tier design serverfarm with chassis redundancy.

Figure 11 Serverfarm Designs

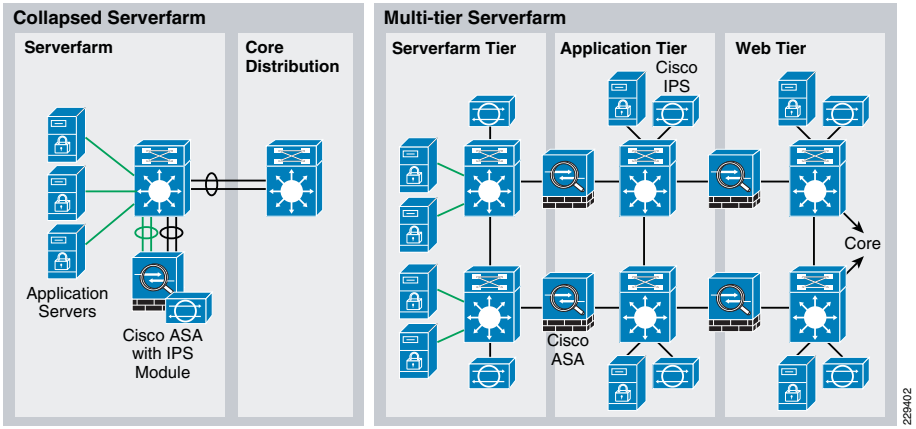

Independently from the design model adopted by the enterprise, the following are the primary security guidelines for the serverfarm design:

- *Network Foundation Protection*— All infrastructure equipment should be protected following the Network Foundation Protection best practices described earlier in this document. This includes restricting and controlling administrative access, protecting the management and control planes, and securing the switching and routing planes.
- *Firewall*—A stateful firewall may be deployed to limit access to only the necessary applications and services, and for the intended users. The firewall should be configured to control and inspect both traffic entering and leaving the server farm

segments. The firewall may also be leveraged to ensure the appropriate segregation between application layers or groups. In addition, the firewall's deep packet inspection may be used to mitigate DoS attacks and enforce protocol compliance.

- *Intrusion Prevention*— An IPS module on the Cisco ASA or a separate IPS appliance may be implemented for enhanced threat detection and mitigation. The IPS is responsible for identifying and blocking anomalous traffic and packets recognized as well-known attacks. The Cisco IPS may be configured either in inline or promiscuous mode. When deployed in inline mode, the Cisco IPS is placed in the traffic path and is capable of stopping malicious traffic before it reaches the intended target.
- *Service Isolation*— Services and applications serving different group of users or under different security requirements should be properly isolated. Isolation helps prevent data leakage and contain possible compromises from spreading across different server farm groups. Logical isolation may be achieved by separating applications and services in different VLANs and by assigning them into different firewall interfaces (physical or logical). This is illustrated in Figure 12.
- *Switch Security*—Private VLANs, port security, storm control and other switch security features may be leveraged to mitigate spoofing, man-in-the-middle, denial-of-service and other network-based attacks directed to the serverfarm applications and the switching infrastructure.
- *Endpoint Protection*— Servers residing at the different layers should be protected with host-based IPS or other endpoint security software.

Figure 12 Service Isolation

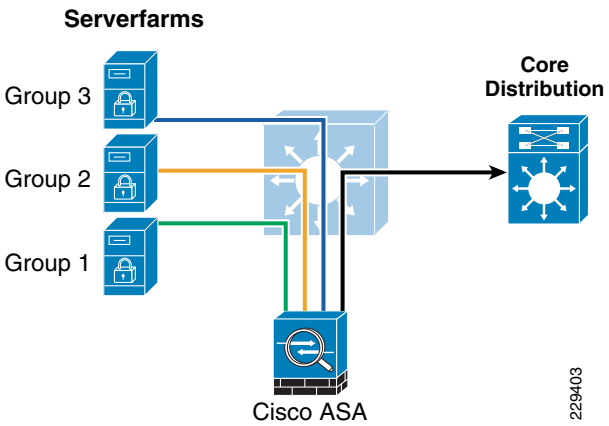

SSL termination and inspection, Web Application Firewall (WAF), Application Control Engine (ACE), and other solutions may be leveraged to complement the guidelines described above. For a more detailed discussion of serverfarm security, refer to "Chapter 4, Intranet Data Center" of the *Cisco SAFE Reference Guide* at the following URL:

http://www.cisco.com/en/US/docs/solutions/Enterprise/Security/SAFE_RG/chap4.html

Network Access Security and Control

Some of the most vulnerable points of the network are the access edges where users connect to the network. With the proliferation of wireless networks, increased use of laptops and smart mobile devices, the enterprise cannot simply rely on physical controls hoping to prevent unauthorized systems from being plugged into the ports of the access switches. Protection should rather be embedded into the network infrastructure, leveraging the native security features available in switches, routers, and WLAN systems. Furthermore, the network infrastructure should also provide dynamic identity or role-based access controls for all systems attempting to gain access.

Implementing role-based access controls for users and devices helps reduce the potential loss of sensitive information by enabling enterprises to verify a user or device identity, privilege level, and security policy compliance before granting network access. Security policy compliance could consist of requiring antivirus software, OS updates or patches. Unauthorized or noncompliant devices can be placed in a quarantine area where remediation can occur prior to gaining access to the network.

The Small Enterprise Design Profile achieves access security and control by leveraging the following technologies:

- Catalyst Integrated Security Features (CISF), wired
- Cisco Unified Wireless Network (CUWN) Integrated Security Features, wireless
- Cisco NAC Appliance, wired and wireless
- Cisco Identity-Based Network Networking Services (IBNS), wired and wireless

Catalyst Integrated Security Features

Catalyst Integrated Security Features (CISF) is a set of native security features available on Cisco Catalyst Switches and designed to protect the access infrastructure and users from spoofing, man-in-the-middle, DoS and other network-based attacks. CISF includes features such as private VLANs, port security, DHCP snooping, IP Source Guard, secure Address Resolution Protocol (ARP) detection, and Dynamic ARP Inspection (DAI). CISF features are considered to be part of a security baseline and should be deployed on all access ports.

- *Port Security*—Mitigates MAC flooding and other Layer 2 CAM overflow attacks by restricting the MAC addresses that are allowed to send traffic on a particular port. After Port Security is enabled on a port, only packets with a permitted source MAC address are allowed to pass through the port. A permitted MAC address is referred to as a secure MAC address.
- *DHCP snooping*—Inspects and filters DHCP messages on a port to ensure DHCP server messages come only from a trusted interface. Additionally, it builds and maintains a DHCP snooping binding table that contains the MAC address, IP address, lease time, binding type, VLAN number, and interface information corresponding to the local untrusted interfaces of a switch. This binding table is used by the other CISF features.
- *IP Source Guard*—Restricts IP traffic on a port based on DHCP or static IP address MAC bindings to prevent IP spoofing attacks. IP address bindings are validated using information in the DHCP Snooping binding table.
- *Dynamic ARP inspection*—Validates that the source MAC and IP address in an ARP packet received on an untrusted interface matches the source MAC and IP address registered on that interface (using the DHCP snooping binding table) to prevent ARP spoofing and MITM attacks.
- *ARP rate limiting*—Where an excessive rate of ARP request (which must be processed by network hosts CPUs), and the switch responds with access restriction if this rate is exceeded.

- *Storm Control*—Prevents broadcast and multicast storms by monitoring packets passing from an interface to the switching bus and determines whether the packet is unicast, multicast, or broadcast. The switch counts the number of packets of a specified type received within the 1-second time interval and compares the measurement with a predefined suppression-level threshold. When the suppression-level threshold is reached, the port blocks traffic until the traffic falls below the threshold level.

Cisco Unified Wireless Network (CUWN) Integrated Security Features

The Cisco Unified Wireless Network adds to the 802.11 security standards by providing additional security features. Some of these are the WLAN equivalent of CiSF features such as, Dynamic Host Configuration Protocol (DHCP) and Address Resolution Protocol (ARP) protection, peer-to-peer blocking, and access control list and firewall features. Additionally other more WLAN specific features are provided, including Enhanced WLAN security options, wireless intrusion detection system (IDS), client exclusion, rogue AP detection, management frame protection, dynamic radio frequency management, and network IDS integration.

The Cisco Unified Wireless Network solutions are discussed in the *Wireless and Network Security Integration Solution Design Guide* at the following URL:

http://www.cisco.com/en/US/solutions/ns340/ns414/ns742/ns820/landing_sec_wireless.html

Cisco NAC Appliance

The Cisco Network Admission Control (NAC) Appliance (formerly known as Cisco Clean Access) uses the network infrastructure to enforce security policy compliance on all devices seeking to access network computing resources. With Cisco NAC Appliance, network administrators can authenticate, authorize, evaluate, and remediate wired, wireless, and remote users and their machines before network access. The NAC Appliance identifies whether networked devices such as laptops, IP phones, or game consoles are compliant with your network security policies, and can repair any vulnerability before permitting access to the network.

The Cisco NAC Appliance provides the following benefits:

- Recognizes users, their devices, and their roles in the network. This first step occurs at the point of authentication, before malicious code can cause damage.
- Evaluates whether machines are compliant with security policies. Security policies can include specific anti-virus or anti-spyware software, OS updates, or patches. Cisco NAC Appliance supports policies that vary by user type, device type, or operating system.
- Enforces security policies by blocking, isolating, and repairing non-compliant machines.
- Non-compliant machines are redirected to a quarantine network, where remediation occurs at the discretion of the administrator.

The NAC solution provides the following four functions, as shown in Figure 13:

- Authenticates and authorizes
- Scans and evaluates
- Quarantines and enforces
- Updates and remediates

Figure 13 The Four Functions of the NAC Framework

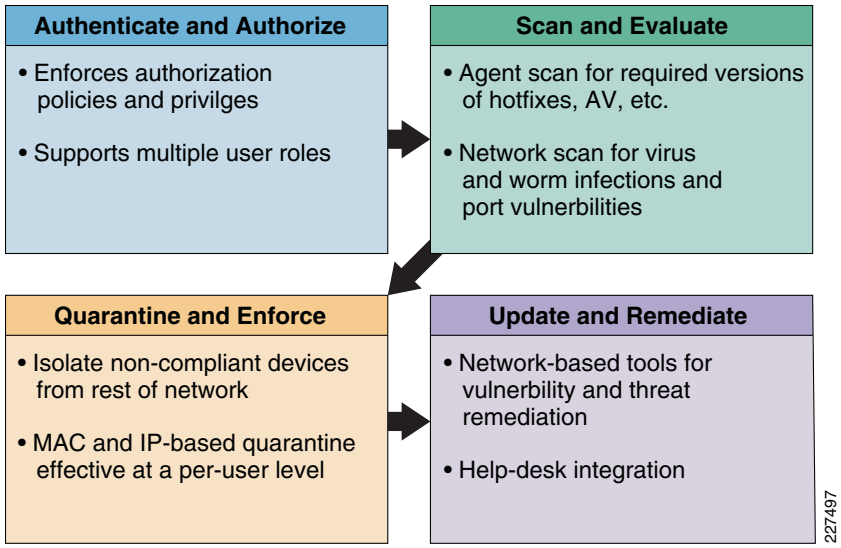

For more details of the NAC Appliance solution, see the following URL: http://www.cisco.com/go/nacappliance.

Cisco NAC Appliance is a network-centric integrated solution administered from the Cisco Clean Access Manager (NAC Manager) web console and enforced through the Clean Access Server (NAC Server) and (optionally) the Clean Access Agent or Cisco NAC Web Agent. Cisco NAC Appliance checks client systems, enforces network requirements, distributes patches and antivirus software, and quarantines vulnerable or infected clients for remediation before clients access the network. Cisco NAC Appliance consists of the components shown in Figure 14.

Figure 14 NAC Appliance Components

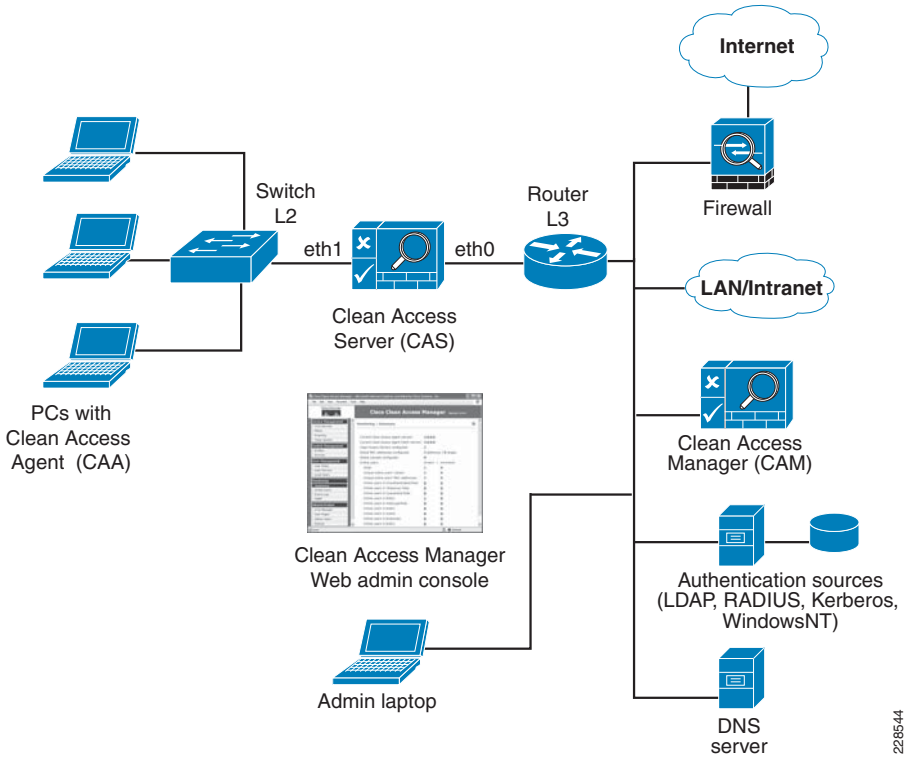

Clean Access Manager (CAM)

The Cisco CAM (also known as NAC Manager) is the administration server for Clean Access deployment. The secure web console of the Clean Access Manager is the single point of management for up to 20 Clean Access Servers in a deployment (or 40 CASs if installing a SuperCAM). For Out-of-Band (OOB) deployment, the web admin console allows you to control switches and VLAN assignment of user ports through the use of SNMP. In the Small Enterprise Design Profile, the CAM would be located at the main site.

Clean Access Server (CAS)

The Cisco CAS (a.k.a NAC Server) is the enforcement server between the untrusted (managed) network and the trusted network. The CAS enforces the policies you have defined in the CAM web admin console, including network access privileges, authentication requirements, bandwidth restrictions, and Clean Access system requirements. You can install a CAS either as a standalone appliance (like the Cisco NAC-3300 Series) or as a network module (Cisco NME-NAC-K9) in a Cisco ISR chassis and deploy it in-band (always inline with user traffic) or OOB (inline with user traffic only during authentication/posture assessment).

The CAS can also be deployed in Layer 2 mode (users are Layer-2-adjacent to the CAS) or Layer 3 mode (users are multiple Layer-3 hops away from the CAS). You can also deploy several CASs of varying size/capacity to fit the needs of varying network segments. You can install Cisco NAC-3300 Series appliances in your company headquarters core, for example to handle thousands of users and simultaneously install one or more Cisco NAC network modules in ISR platforms to accommodate smaller groups of users at a satellite office, for example.

In the Small Enterprise Design Profile, the CAS would be located at the main site and the remote locations, and it would be used to provide Layer-2 or Layer-3 OOB authentication/posture assessment.

Clean Access Agent (CAA)

CAA is optional read-only agent that resides on Windows clients. It checks applications, files, services, or registry keys to ensure that clients meets your specified network and software requirements prior to gaining access to the network.

Note There is no client firewall restriction with CAA posture assessment. The agent can check the client registry, services, and applications even if a personal firewall is installed and running.

If NAC is implemented as part of the Small Enterprise Design Profile it is recommended that the CAA be used.

Cisco NAC Web Agent

The Cisco NAC Web Agent provides temporal posture assessment for client machines. Users launch the Cisco NAC Web Agent executable, which installs the Web Agent files in a temporary directory on the client machine via ActiveX control or Java applet. When the user terminates the Web Agent session, the Web Agent logs the user off of the network and their user ID disappears from the Online Users list.

Clean Access Policy Updates

Regular updates of prepackaged policies/rules that can be used to check the up-to-date status of operating systems, antivirus (AV), antispyware (AS), and other client software. It provides built-in support for 24 AV vendors and 17 AS vendors.

NAC Appliance Modes and Positioning

NAC Appliance allows multiple deployment options and may be placed at different points in the network. The modes of operation can be generally defined as follows:

- Out-of-band (OOB) virtual gateway
- OOB IP real gateway
- In-band (IB) virtual gateway
- IB real IP gateway

Out-of-Band Modes

Out-of-Band (OOB) deployments require user traffic to traverse through the NAC Appliance only during authentication, posture assessment, and remediation. When a user is authenticated and passes all policy checks, their traffic is switched normally through the network and bypasses the appliance. See Figure 15.

Figure 15 Layer-2 OOB Topology

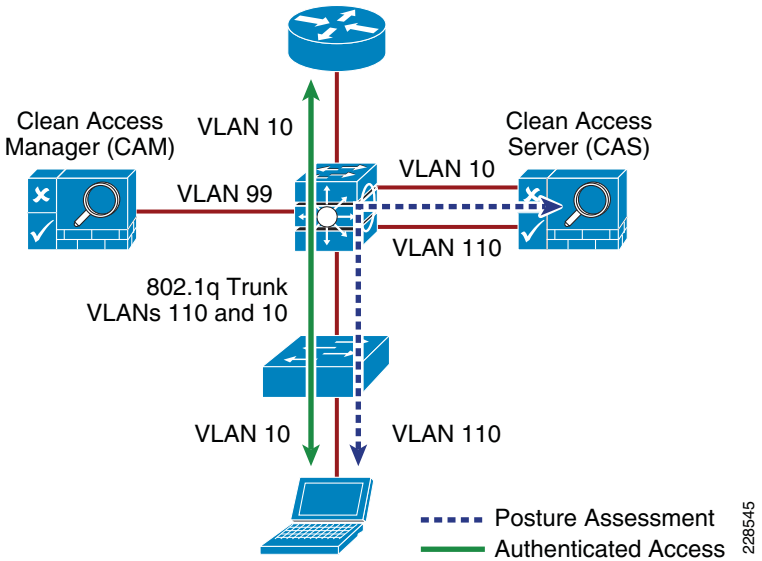

Figure 16 In-Band Virtual Gateway Topology

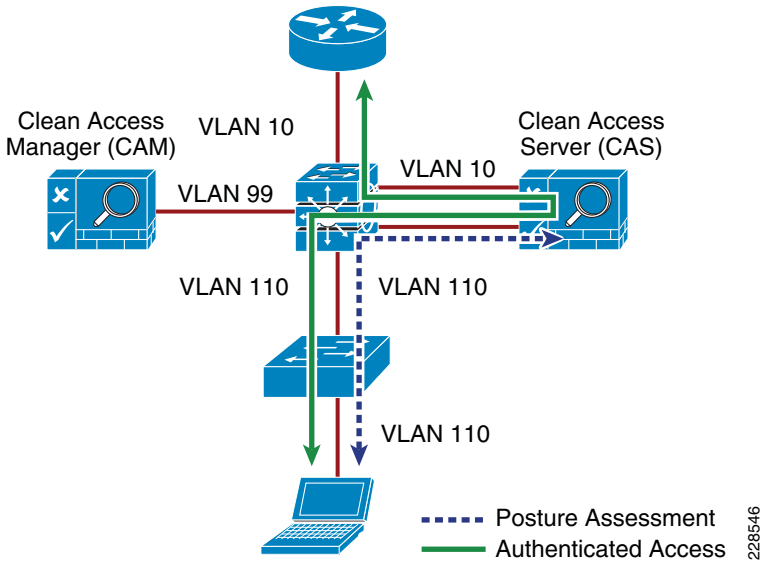

To deploy the NAC Appliance in this manner, the client device must be directly connected to the network via a Catalyst switch port. After the user is authenticated and passes posture assessment, the Clean Access Manager (CAM) instructs the switch to map the user port from an unauthenticated VLAN (which switches or routes user traffic to the NAC) to an authenticated (authorized) VLAN that offers full access privileges. For example, as shown in Figure 15, the client PC is connected through VLAN 110 to the NAC Clean Access Server for the authentication and posture assessment, and is moved to VLAN 10 once it successfully completes the authentication and authorization, scan, and evaluation phases of the NAC solution.

In-Band Modes

When the NAC Appliance is deployed in-band, all user traffic, both unauthenticated and authenticated, passes through the NAC Appliance, which may be positioned logically or physically between end users and the network(s) being protected. See Figure 16 for a logical in-band topology example and Figure 17 for a physical in-band topology example.

Figure 17 Physical In-Band Topology

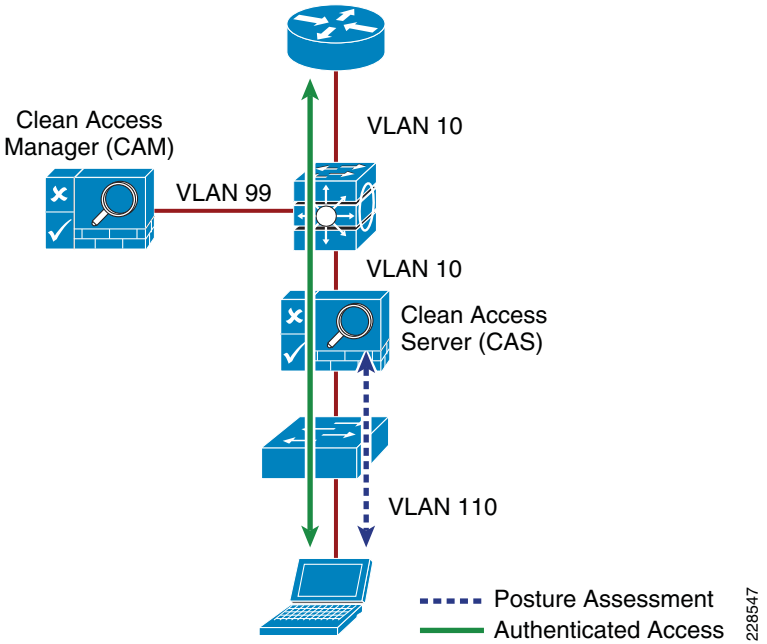

In-Band Virtual Gateway

When the NAC Appliance is configured as a virtual gateway, it acts as a bridge between end users and the default gateway (router) for the client subnet being managed. The following two bridging options are supported by the NAC Appliance:

- *Transparent*—For a given client VLAN, the NAC Appliance bridges traffic from its untrusted interface to its trusted interface. Because the appliance is aware of "upper layer protocols", by default it blocks all traffic except for Bridge Protocol Data Unit (BPDU) frames (spanning tree) and those protocols explicitly permitted in the "unauthorized" role; for example, DNS and DHCP. In other words, it permits those protocols that are necessary for a client to connect to the network, authenticate, undergo posture assessment, and remediation. This option is viable when the NAC Appliance is positioned physically in-band between end users and the upstream network(s) being protected, as shown in Figure 17.
- *VLAN mapping*—This is similar in behavior to the transparent method except that rather than bridging the same VLAN from the untrusted side to the trusted side of the appliance, two VLANs are used. For example, Client VLAN 131 is defined for the untrusted interface of the NAC Appliance. There is no routed interface or switched virtual interface (SVI) associated with VLAN 131. VLAN 31 is configured between the trusted interface of the NAC Appliance and the next-hop router interface/SVI for the client subnet. A mapping rule is made in the NAC Appliance that forwards packets arriving on VLAN 131 and forwards them out VLAN 31 by swapping VLAN tag information. The process is reversed for packets returning to the client. Note that in this mode, BPDUs are not passed from the untrusted-side VLANs to their trusted-side counterparts.

The VLAN mapping option is usually selected when the NAC Appliance is positioned logically in-band between clients and the networks being protected. This is the bridging option that should be used if the NAC Appliance is going to be deployed in the virtual gateway mode.

In-Band Real IP Gateway

When the NAC Appliance is configured as a "real" IP gateway, it behaves like a router and forwards packets between its interfaces. In this scenario, one or more client VLAN/subnets reside behind the untrusted interface. The NAC Appliance acts as a default gateway for all clients residing on those networks. Conversely, a single VLAN/subnet is defined on the trusted interface, which represents the path to the protected upstream network(s).

After successful client authentication and posture assessment, the NAC Appliance by default routes traffic from the untrusted networks to the trusted interface, where it is then forwarded based on the routing topology of the network.

The NAC Appliance is not currently able to support dynamic routing protocols. As such, static routes must be configured within the trusted side of the Layer 3 network for each client subnet terminating on or residing behind the untrusted interface. These static routes should reference, as a next hop, the IP address of the trusted interface of the NAC.

If one or more Layer-3 hops exist between the untrusted NAC interface and the end-client subnets, static routes to the client networks must be configured in the NAC Appliance. Likewise, a static default route (0/0) is required within the downstream Layer 3 network (referencing the IP address of the untrusted NAC interface) to facilitate default routing behavior from the client networks to the NAC Appliance.

Depending on the topology, multiple options exist to facilitate routing to and from the NAC Appliance, including static routes, VRF-Lite, MPLS VPN, and other segmentation techniques. It is beyond the scope of this design guide to examine all possible methods.

In-Band Versus Out-of-Band

Table 1 summarizes different characteristics of each type of deployment.

Table 1 In-Band Versus Out-of-Band Deployment Characteristics

In-Band Deployment Characteristics	Out-of-Band Deployment Characteristics
The CAS is always inline with user traffic (both before and following authentication, posture assessment and remediation). Enforcement is achieved through being inline with traffic.	The CAS is inline with user traffic only during the process of authentication, assessment and remediation. Following that, user traffic does not come to the CAS. Enforcement is achieved through the use of SNMP to control switches and VLAN assignments to ports.
The CAS can be used to securely control authenticated and unauthenticated user traffic by using traffic policies (based on port, protocol, subnet), bandwidth policies, and so on.	The CAS can control user traffic during the authentication, assessment and remediation phase, but cannot do so post-remediation since the traffic is out-of-band.
Does not provide switch port level control.	Provides port-level control by assigning ports to specific VLANs as necessary using SNMP.
In-band deployment is supported for wired and wireless clients.	OOB deployments support wired and wireless clients. Wireless OOB requires a specific network topology.[1]
Cisco NAC Appliance In-Band deployment with supported Cisco switches is compatible with 802.1x	Cisco does not recommend using 802.1x in an OOB deployment, as conflicts will likely exist between Cisco NAC Appliance OOB and 802.1x to set the VLAN on the switch interfaces/ports.

1. OOB NAC deployments for wireless require the NAC server to be deployed in Layer 2 OOB virtual gateway (bridge) mode, and the NAC server must be placed Layer 2-adjacent to the wireless LAN controller (WLC).

Out-of-Band Requirements

OOB implementation of Cisco NAC Appliance requires the switches and Wireless LAN Controllers to be supported by the Cisco NAC Appliance software. All the switches tested as part of the development of the Small Enterprise Design Profile, apart from the Cisco Catalyst 2975, are supported by the Cisco NAC OOB, and the Wireless LAN Controllers are also supported by the NAC Appliance software used in this design guide. If the Catalyst 2975 is to be used as an access switch with the Cisco NAC Appliance, the NAC solution must be an in-band solution.

Note To obtain the latest list of supported devices, check the latest version of the *Cisco NAC Appliance-Clean Access Manager Installation and Administration Guide* at the following URL: http://www.cisco.com/en/US/docs/security/nac/appliance/configuration_guide/45/cam/45cam-book.html

Out-Of-Band, Layer 2 and Layer 3

The proposed design for the Small Enterprise Design Profile is an OOB design, in order to get the highest possible performance and scalability for traffic that has passed through the authentication, posture assessment, and remediation stages of NAC. The Small Enterprise Design Profile offers two different access layer options, a Layer-2 access layer for smaller sites and a hybrid Layer-2/Layer-3 access layer for larger sites. This means that either a Layer-2 OOB solution or a Layer-3 OOB NAC solution may be deployed.

NAC Appliance Deployment in Small Enterprise Networks

Within the Small Enterprise Design Profile, a Cisco NAC Appliance is deployed at each of the site types, the headquarters or main site, and remote sites. A centralized CAM is deployed at the main site, likely within the serverfarm. A CAS is deployed at the main site and each remote site and is directly connected to the core/distribution layer at each of the locations.

The simple topology used in each site type means that a VLAN from an access layer to the untrusted interface of the NAC Appliance is always available as a standard component of the design, and untrusted traffic should never need to be tunneled to the CAS. This allows a common network configuration, to support NAC at any of the enterprise sites, regardless of whether the client devices are using a Layer-2 or Layer-3 access model. As the client can use a Layer-2 connection to the untrusted interface of the NAS in either Layer 2 or Layer 3 access mode (this requires a trunk between the Layer-3 access switch and the core/distribution. One VLAN of the trunk would carry the untrusted VLAN, and the other VLAN the IP routing for all other traffic), and the VLAN used once the client is trusted will be either be a Layer-2 access VLAN from the core/distribution switch or a Layer-3 access switch VLAN depending upon the site requirements.

This is illustrated in Figure 18 and Figure 19. In Figure 19, there is a simple Layer-2 NAC OOB connection where a client device upon initial connection to the network is given VLAN 264, which connects them directly to the untrusted interface of the NAS. VLAN 264 on the untrusted interface is mapped to VLAN 64 on the trusted interface within the NAC appliance, which allows the client to obtain an IP address that belongs on VLAN 64. Upon successful completion of the NAC authentication and validation functions, the access switch is instructed, via SNMP from the CAM, to change the client VLAN to VLAN 64. Even though the client has changed Layer-2 VLANs, its Layer-3 network connections are unchanged and the traffic from the client no longer passes through the NAC appliance.

Figure 18 Layer 2 OOB Topology

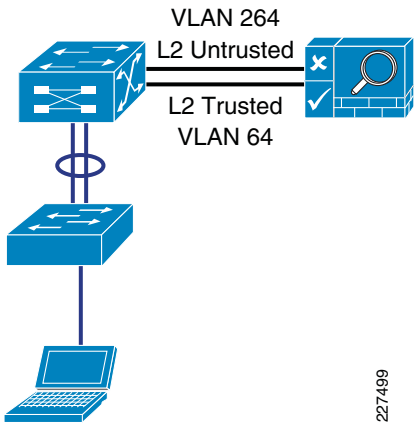

In Figure 19, the same processes are followed when the client is untrusted, but once the client has successfully completed its NAC functions the access switch is instructed via SNMP to change the client VLAN to VLAN 67—a subnet local to the access switch. As the Layer-3 information for the client has changed, the switch is also instructed to "bounce" the client switch port to initiate a new DHCP request for an IP address appropriate to VLAN 67.

Figure 19 Layer 3 OOB Topology

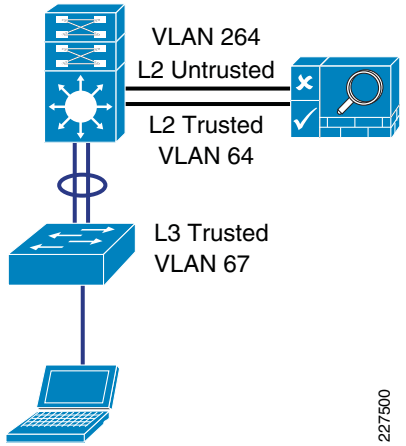

Cisco Identity-Based Network Networking Services (IBNS)

The best and most secure solution to vulnerability at the access edge is to leverage the intelligence of the network. The Cisco Identity-Based Network Networking Services solution (IBNS) is a set of Cisco IOS software services designed to enable secure user and host access to enterprise networks powered by Cisco Catalyst switches and wireless LANs. It provides standards-based network access control at the access layer by using the 802.1X protocol to secure the physical ports where end users connect. 802.1X is an IEEE standard for media-level (Layer 2) access control, offering the capability to permit or

deny network connectivity based on the identity of the end user or device. 802.1X is well-known as a way to secure wireless network access. It is equally essential in securing wired network access.

IEEE 802.1X Protocol

The IEEE 802.1X protocol allows Cisco Catalyst switches to offer network access control at the port level. Every port on the switch is individually enabled or disabled based on the identity of the user or device connecting to it. When 802.1X is first enabled on a port, the switch automatically drops all traffic received on that port. There is one exception to this rule. The only traffic a switch will accept is a request to start 802.1X authentication. Only after the 802.1X authentication has successfully completed will the switch accept any other kind of traffic on the port.

The high-level message exchange in Figure 20 illustrates how port-based access control works within an identity-based system.

Figure 20 Port-Based Access Control

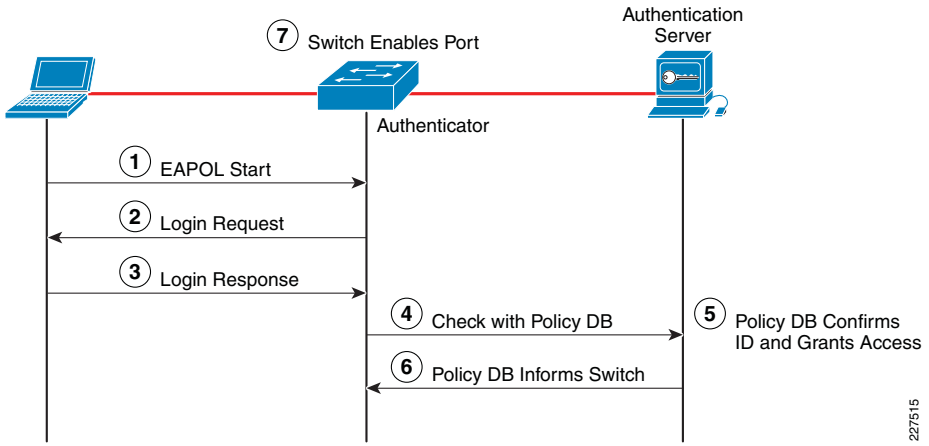

The following steps describe the port-based access control flow shown in Figure 20:

1. A client, such as a laptop with an 802.1X supplicant, connects to an IEEE 802.1X-enabled network and sends a start message to the LAN switch (the authenticator).
2. When the start message is received, the LAN switch sends a login request to the client.
3. The client replies with a login response.
4. The switch forwards the response to the policy database (authentication server).
5. The authentication server authenticates the user.
6. After the user identity is confirmed, the policy database authorizes network access for the user and informs the LAN switch.
7. The LAN switch then enables the port connected to the client.

The user and device credentials are processed by an AAA server. The AAA server is able to reference user or device policy profile information either internally, using the integrated user database, or externally, using database sources such as Microsoft Active Directory, LDAP, Novell NDS or Oracle databases. This enables the integration of the system into exiting user management structures and schemes, thereby simplifying overall deployment.

802.1X and EAP

When authenticating users for the purposes of network access control, the system must provide user and/or device identification using strong authentication technologies known to be secure and reliable. IEEE 802.1X does not by itself dictate how this is achieved. Rather, the 802.1X protocol defines an encapsulation for the transport of the Extensible Authentication Protocol (EAP) from the client to the switch. The 802.1X encapsulation is sometimes referred to as EAP over LAN (EAPoL). The switch in turn relays the EAP information to the authentication server using the RADIUS protocol (EAP over RADIUS).

EAP, which is defined by RFC 3748, is itself a framework---not a specific authentication method. EAP provides a way for the client and the authentication server to negotiate an authentication method that they both support. There are many EAP methods but the ones used more frequently for 802.1X wired authentication include EAP-TLS, EAP-PEAP, and EAP-FAST.

Impacts of 802.1X on the Network

Before enabling 802.1X in the network, it is essential to review the default security posture of a port enabled for 802.1X authentication: all traffic is dropped except 802.1X EAPoL packets. This is a fundamental change from the traditional model in which the port is enabled and all traffic is allowed from the moment that a device plugs into the port. Ports that were traditionally open will now be closed by default. This is one of the cornerstones of the strong security and network access control provided by 802.1X. However, this change in the default network access model can have a profound impact on network devices and applications. Understanding and providing for the impacts of this change will make for a smooth deployment of 802.1X network access control.

Non-802.1X-Enabled Devices

802.1X must be enabled on both the host device and on the switch to which the device connects. If a device without an 802.1X supplicant attempts to connect to a port that is enabled for 802.1X, it will be subjected to the default security policy. The default security policy says that 802.1X authentication must succeed before access to the network is granted. Therefore, by default, non-802.1X-capable devices cannot get access to an 802.1X-protected network.

Although many devices increasingly support 802.1X, there will always be devices that require network connectivity but do not and/or cannot support 802.1X. Examples of such devices include network printers, badge readers, legacy servers, and PXE boot machines. Some provision must be made for these devices.

Cisco provides two features to accommodate non-802.1X devices. These are MAC Authentication Bypass (MAB) and the Guest VLAN. These features provide fallback mechanisms when there is no 802.1X supplicant. After 802.1X times out on a port, the port can move to an open state if MAB succeeds or if the Guest VLAN is configured. Judicious application of either or both of these features will be required for a successful 802.1X deployment.

Note Network-specific testing will be required to determine the optimal values for 802.1X timers to accommodate the various non-802.1X-capable devices on your network.

802.1X in Small Enterprise Networks

As mentioned above, one of the requirements for 802.1X authentication is the requirement for a supplicant. This has typically been a challenge in enterprise environments with a wide range of devices and limited or no management of many of these devices. In many enterprises this is still the case, and this makes a company-wide 802.1X very challenging. At the same time there are pockets of an enterprise network where 802.1X may be a good choice.

For example, 802.1X protected ports may be a good choice for the network ports in the company's headquarters or main site, because these locations are more likely to have managed PCs.

Other locations in the enterprise network still need protection, but user network access may be better served by a NAC Appliance solution. For networks requiring role-based access control using posture assessments to ensure security compliance, Cisco NAC Appliance should be considered. In addition, network access ports in open areas such as lobbies and meeting rooms may use 802.1X or Cisco Clean Access NAC to protect these ports.

When considering the 802.1X deployment, the following four main 802.1X authentication options need to be considered:

- *Basic 802.1X Authentication*—An 802.1X controlled port with an 802.1X client directly connected
- *IP Phone Ports*—An IP Phone and an 802.1X controlled port with an 802.1X client connected to the phone
- *MAC Auth By-Pass*—Using the MAC address of the client to provide authentication and bypass the 802.1X authentication process. Printer and legacy device support are typical applications
- *Web Auth*—Allowing a user to authenticate by entering username and passwords in a web page. Legacy device support and guest access are typical deployment applications

For more information on the Cisco IBNS 802.1X network access solution, see the following URL: http://www.cisco.com/go/ibns.

NAC 802.1X and CISF in Combination

The three key access security features discussed above have been discussed in isolation, but can be combined. In particular, the CISF features should be considered "baseline" features that are applied on all access ports, and either NAC or 802.1X maybe overlaid on top of the CISF configuration.

The Cisco Clean Access and 802.1X configuration are also compatible (although they are not often combined in wired networks), the key consideration in combining the two is how to give the appearance of a SSO for the end user. Both 802.1X and NAC require authentication, as 802.1X authenticates the client initially, a mechanism of communicating the 802.1X authentication result to the Cisco Clean Access system is required.

If the authenticating clients join an Windows Active Directory network, the Cisco Clean Access Active Directory SSO feature allows the clients to authenticate to active directory once they have performed their 802.1X authentication. The CAM, when a client is detected, checks Active Directory to see if the client has authenticated; this allows a SSO experience for client devices that are using 802.1X and NAC.

Secure Mobility

Today workers use laptops, smartphones and other smart mobile devices to access information and applications at anytime and from anywhere there is an Internet connection. While embracing a mobile workforce clearly boosts productivity and makes the small enterprise more competitive, there are a number of challenges that arise from the use of mobile technologies. To start with, workers often use the same devices to access both business and personal information. Devices used outside the enterprise onsite controls may potentially introduce viruses, worms, spyware and other type of malware as mobile workers connect back to the corporate network. Confidential and proprietary information may also be lost or stolen while mobile users connect outside the company premises. At the other hand, the great variety in hardware types, operating systems, and applications represents a clear challenge to the enforcement of security controls and policies. In order to continue to foster innovation, enable productivity, and meet the needs of the mobile workforce, companies must adapt to the changing trends in mobility. A viable solution is one that enables access for mobile workers while ensuring that the corporate data, assets and network remain secure. At the other hand, users want the flexibility of choosing how, when, and where to access both personal and professional information to be productive without being inconvenienced by security checks.

The Small Enterprise Design Profile provides persistent and secure mobile access by implementing the Cisco AnyConnect Secure Mobility solution (see Figure 21). This solution delivers secure, persistent connectivity to all mobile employees independently from the type of mobile device used. The Cisco AnyConnect Secure Mobility solution also ensures a consistent enforcement of the network security policies to all users, no matter when, where and how they connect to the network.

The Cisco AnyConnect Secure Mobility is a collection of features across multiple Cisco products that extends control and security into borderless networks. The products that work together to provide AnyConnect Secure Mobility are as follows:

- Cisco IronPort Web Security appliance (WSA)
- Cisco ASA 5500 series adaptive security appliance (ASA)
- Cisco AnyConnect client

Cisco AnyConnect Secure Mobility addresses the challenges of a mobile workforce by offering the following features:

- *Secure, persistent connectivity*—Cisco AnyConnect (with the Cisco ASA at the headend) provides the remote access connectivity portion of AnyConnect Secure Mobility. The connection is secure because both the user and device must be authenticated and validated prior to being provided access to the network. The connection is persistent because Cisco AnyConnect is typically configured to be always-on even when roaming between networks. Although Cisco AnyConnect is always-on, it is also flexible enough to apply different policies based on location, allowing users access to the Internet in a "captive portal" situation, when users must accept terms of agreement before accessing the Internet.
- *Persistent security and policy enforcement*—The Web Security appliance applies context-aware policies, including enforcing acceptable use policies and protection from malware for all users, including mobile (remote) users. The WSA also accepts user authentication information from the AnyConnect client, providing an automatic authentication step for the user to access web content.

Figure 21 Cisco AnyConnect Secure Mobility Solution

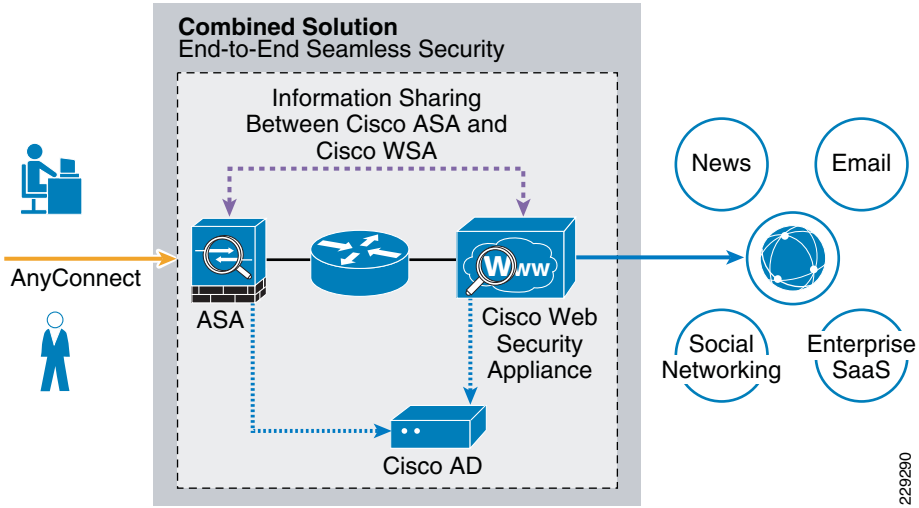

Figure 21 illustrates the relationship between the various elements of the Cisco AnyConnect Secure Mobility solution. Remote and mobile users use the Cisco AnyConnect Secure VPN client to establish VPN sessions with the Cisco ASA appliance. The Cisco ASA sends web traffic to the WSA appliance along with information identifying the user by IP address and user name. The WSA scans the traffic, enforces acceptable use policies, and protects the user from security threats. The Cisco ASA returns all traffic deemed safe and acceptable to the user.

All Internet traffic scanning is done by the WSA, not the client on the mobile device. This improves overall performance by not burdening the mobile device, some of which have limited processing power. In addition, by scanning Internet traffic on the network, the enterprise can more easily and quickly update security updates and acceptable use policies since the enterprise does not have to wait days, weeks, or even months to push the updates to the client. The WSA tracks the requests it receives and applies policies configured for remote users to traffic received from remote users.

For complete details about the Cisco AnyConnect Secure Mobility solution, refer to the documentation available at the following URL:

http://www.cisco.com/en/US/netsol/ns1049/index.html

Threats Mitigated

The success of the security tools and measures in place ultimately depends on the degree they enhance visibility and control. Simply put, security can be defined as a function of visibility and control. Without any visibility, it is difficult to enforce any control, and without any control it is hard to achieve an adequate level of security. Therefore, the security tools selected in the enterprise network design were carefully chosen not only to mitigate certain threats but also to increase the overall visibility and control.

Table 2 summarizes how the security tools and measures used in the Small Enterprise Design Profile help mitigate certain threats, and how they contribute to increasing visibility and control. Please note the table is provided for illustration purposes and it is not intended to include all possible security tools, and threats.

Table 2 Security Measures of the Small Enterprise Design Profile

	Service Disruption	Harmful Content	Network Abuse	Unauthorized Access	Data Loss	Visibility	Control
Network Foundation Protection	Yes			Yes	Yes	Yes	Yes
Stateful Firewall	Yes		Yes	Yes		Yes	Yes
IPS	Yes	Yes	Yes	Yes		Yes	Yes
Mobile Security	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Web Security		Yes	Yes	Yes	Yes	Yes	Yes
E-mail Security		Yes	Yes		Yes	Yes	Yes
Access Security and Control			Yes	Yes		Yes	Yes

Network Security Deployment

This section describes the deployment best practices for the key security platforms and features used in the Small Enterprise Design Profile. This includes deployment and setting guidelines, and configuration examples.

Internet Border Router Deployment

Whether the Internet border router is managed by the enterprise or the ISP, it must be hardened following the best practices listed in the "Network Foundation Protection" section on page -2. This includes restricting and controlling administrative access, protecting the management and control planes, and securing the dynamic exchange of routing information. In addition, the Internet border router may be leveraged as the first layer of protection against outside threats. To that end, edge ACLs, uRPF, and other filtering mechanisms may be implemented for anti-spoofing and to block invalid packets.

The following configuration snippet illustrates the structure of an edge ACL applied to the upstream interface of the Internet border router. The ACL is designed to block invalid packets and to protect the infrastructure IP addresses from the Internet. The configuration assumes the enterprise is assigned the 198.133.219.0/24 address block for its public-facing services, and that the upstream link is configured in the 64.104.10.0/24 subnet.

```
! !!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!
!--- Module 1: Anti-spoofing Denies
!--- These ACEs deny fragments, RFC 1918 space,
!--- invalid source addresses, and spoofs of
```

```
!--- internal space (space as an external source).
!
!--- Deny fragments.
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
!--- Deny special-use address sources.
!--- See RFC 3330 for additional special-use addresses.
access-list 110 deny ip host 0.0.0.0 any
access-list 110 deny ip 127.0.0.0 0.255.255.255 any
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
!--- Filter RFC 1918 space.
access-list 110 deny ip 10.0.0.0 0.255.255.255 any
access-list 110 deny ip 172.16.0.0 0.15.255.255 any
access-list 110 deny ip 192.168.0.0 0.0.255.255 any
!--- Deny packets spoofing the enterprise public addresses
access-list 110 deny ip 198.133.219.0 0.0.0.255 any
!
!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!
!--- Module 2:  Explicit Permit
!--- Permit only applications/protocols whose destination
!--- address is part of the infrastructure IP block.
!--- The source of the traffic should be known and authorized.
!
!--- Permit external BGP to peer 64.104.10.113
access-list 110 permit tcp host 64.104.10.114 host 64.104.10.113 eq bgp
access-list 110 permit tcp host 64.104.10.114 eq bgp host 64.104.10.113
!
!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!
!--- Module 3:  Explicit Deny to Protect Infrastructure
access-list 110 deny ip 64.104.10.0 0.0.0.255 any
!
!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!
!--- Module 4:  Explicit Permit for Traffic to Company's Public
!--- Subnet.
access-list 110 permit ip any 198.133.219.0 0.0.0.255
!
```

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

Internet Firewall Deployment

The mission of the Internet firewall is to protect the company's internal resources and data from external threats, secure the public services provided by the DMZ, and to control user's traffic to the Internet. The Small Enterprise Design Profile uses a Cisco ASA appliance as illustrated in Figure 22.

Figure 22 Internet Edge Firewall

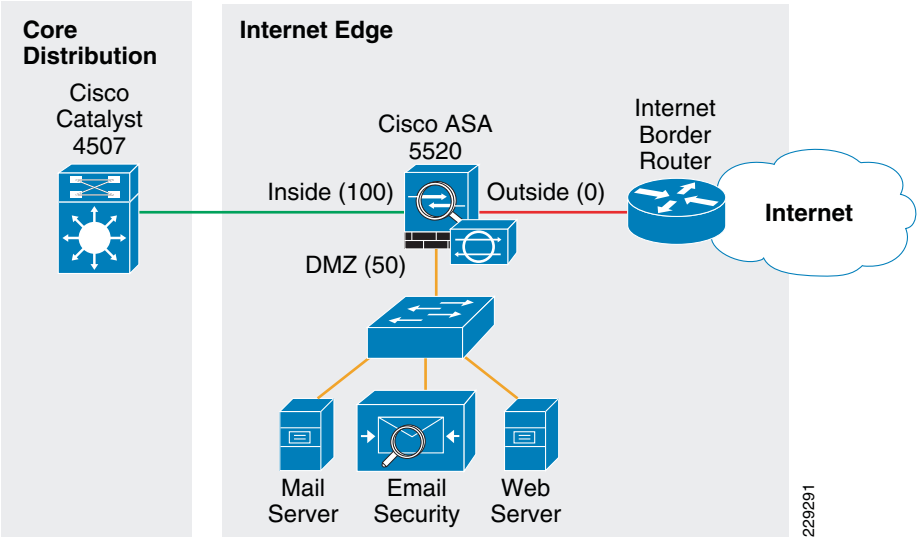

The Cisco ASA is implemented with three interface groups, each one representing a distinct security domain:

- *Inside*—The interface connecting to the core/distribution switch that faces the interior of the network where internal users and resources reside. The inside interface connects to the internal trusted networks, therefore it is given the highest security level, 100.
- *Outside*—Interface connecting to the Internet border router. The router may be managed either by the enterprise or a service provider. The outside interface connects to the Internet, hence it's given the lowest security level, 0.
- *Demilitarized Zone (DMZ)*—The DMZ hosts services that are accessible over the Internet. These services may include the company's website and E-mail services. The DMZ serves a medium level security segment, therefore should be given any security value between the ones defined for the inside and the outside interfaces, for example, 50.

The Internet firewall acts as the primary gateway to the Internet; therefore, its deployment should be carefully planned. The following are key aspects to be considered when implementing the firewall:

- Firewall hardening and monitoring
- Network Address Translation (NAT)
- Firewall access policies
- Botnet Traffic Filter

- Firewall redundancy
- Routing

Firewall Hardening and Monitoring

The Cisco ASA should be hardened in a similar fashion as the infrastructure routers and switches. According to the Cisco SAFE security best practices, the following is a summary of the measures to be taken:

- Implement dedicated management interfaces to the OOB management network.
- Present legal notification for all access attempts.
- Use HTTPS and SSH for device access. Limit access to known IP addresses used for administrative access.
- Configure AAA for role-based access control and logging. Use a local fallback account in case AAA server is unreachable.
- Use NTP to synchronize the time.
- Use syslog or SNMP to keep track of system status, traffic statistics, and device access information.
- Authenticate routing neighbors and log neighbor changes.
- Implement firewall access policies (explained in Firewall Access Policies).

The Cisco ASA 5510 and higher appliance models come with a dedicated management interface that should be used whenever possible. Using a dedicated management interface keeps the management plane of the firewall isolated from threats originating from the data plane. The management interface should connect to the OOB management network, if one is available.

The following is an example of the configuration of a dedicated management interface.

```
interface Management0/0
 nameif management
 security-level 100
 ip address 172.26.160.225 255.255.252.0
 management-only
!
```

Note Any physical interface or logical sub-interface can be configured as a management-only interface using the management-only command.

It is recommended that a legal notification banner is presented on all interactive sessions to ensure that users are notified of the security policy being enforced and to which they are subject. The notification banner should be written in consultation with your legal advisors.

The following example displays the banner after the user logs in:

```
banner motd UNAUTHORIZED ACCESS TO THIS DEVICE IS PROHIBITED.
banner motd You must have explicit, authorized permission to access or
configure this device.
banner motd Unauthorized attempts and actions to access or use this
system may result in civil and/or criminal penalties.
banner motd All activities performed on this device are logged and
monitored.
```

Management access to the firewall should be restricted to SSH and HTTPS. SSH is needed for CLI access and HTTPS is needed for the firewall GUI-based management tools such as CSM and ADSM. Additionally, this access should only be permitted for users authorized to access the firewalls for management purposes.

The following ASA configuration fragment illustrates the configuration needed to generate a 768 RSA key pair and enabling SSH and HTTPS access for devices located in the management subnet.

```
! Generate RSA key pair with a key modulus of 768 bits
crypto key generate rsa modulus 768
! Save the RSA keys to persistent flash memory
write memory
! enable HTTPS
http server enable
! restrict HTTPS access to the firewall to permitted management stations
http <CSM/ADSM-IP-address> 255.255.255.255 management
! restrict SSH access to the firewall to well-known administrative
systems
ssh <admin-host-IP-address> 255.255.255.255 management
! Configure a timeout value for SSH access to 5 minutes
ssh timeout 5
```

Administrative users accessing the firewalls for management must be authenticated, authorized, and access should be logged using AAA. The following ASA configuration fragment illustrates the AAA configurations needed to authenticate, authorize, and log user access to the firewall:

```
aaa-server tacacs-servers protocol tacacs+
 reactivation-mode timed
aaa-server tacacs-servers host <ACS-Server>
 key <secure-key>
aaa authentication ssh console tacacs-servers LOCAL
aaa authentication serial console tacacs-servers LOCAL
aaa authentication enable console tacacs-servers LOCAL
aaa authentication http console tacacs-servers LOCAL
aaa authorization command tacacs-servers LOCAL
aaa accounting ssh console tacacs-servers
aaa accounting serial console tacacs-servers
aaa accounting command tacacs-servers
aaa accounting enable console tacacs-servers
aaa authorization exec authentication-server
! define local username and password for local authentication fallback
username admin password <secure-password> encrypted privilege 15
```

As with the other infrastructure devices in the network, it is important to synchronize the time on the firewall protecting the management module using NTP.

The following configuration fragment illustrates the NTP configuration needed on an ASA to enable NTP to an NTP server located in the management network:

```
ntp authentication-key 10 md5 *
ntp authenticate
ntp trusted-key 10
ntp server <NTP-Server-address> source management
```

Syslog and SNMP can be used to keep track of system status, device access and session activity. NetFlow Security Event Logging (NSEL), now supported on all Cisco ASA models, may also be used for the monitoring and reporting of session activity. The following configuration fragment illustrates the configuration of Syslog.

```
logging trap informational
logging host management <Syslog-Server-address>
logging enable
```

The routing protocol running between the Internet firewall and the distribution/core should be secured. The following ASA configuration fragment illustrates the use of EIGRP MD5 authentication to authenticate the peering session between the inside firewall interface and the core/distribution switch:

```
interface Redundant1
 nameif inside
 security-level 100
 ip address 10.125.33.10 255.255.255.0
 authentication key eigrp 100 <removed> key-id 1
 authentication mode eigrp 100 md5
```

Network Address Translation (NAT)

NAT is required because the enterprise typically gets a limited number of public IP addresses. In addition, NAT helps shield the company's internal address space from reconnaissance and another malicious activity.

The following illustrates the NAT configuration:

```
! Static translation for servers residing at DMZ
static (dmz,outside) 198.133.219.10 10.25.34.10 netmask 255.255.255.255
static (dmz,outside) 198.133.219.11 10.25.34.11 netmask 255.255.255.255
static (dmz,outside) 198.133.219.12 10.25.34.12 netmask 255.255.255.255
static (dmz,outside) 198.133.219.13 10.25.34.13 netmask 255.255.255.255
!
! Dynamic Port Address Translation (PAT) for inside hosts going to the
Internet
global (outside) 10 interface
nat (inside) 10 10.0.0.0 255.0.0.0
!
! Static translation for inside hosts going to the DMZ and vice-versa.
The inside IP addresses are visible to the DMZ.
static (inside,dmz) 10.0.0.0 10.0.0.0 netmask 255.0.0.0
```

Firewall Access Policies

As previously explained, the Internet firewall should be configured to:

- Protect the company's internal resources and data from external threats by preventing incoming access from the Internet.
- 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.
- Control user's Internet-bound traffic.

Enforcing such policies requires the configuration of the appropriate interface security levels and the deployment of ACLs governing what traffic is allowed or prevented from transiting between interfaces.

By default, the Cisco ASA appliance allows traffic from higher to lower security level interfaces (i.e., from inside to outside). However, due to the sensitivity of enterprise environments, the security administrators are recommended to override the default rules with more stringent rules indicating exactly what ports and protocols are permitted.

In our configuration example the inside, DMZ and outside interfaces were configured with the security levels of 100, 50 and 0 respectively. With this, by default any traffic originating from the inside to the DMZ and outside, and from the DMZ to the outside interface will be allowed freely. At the same time, any traffic originating from the outside to the DMZ and inside, and from the DMZ to the inside interface will be blocked. While this may satisfy the basic access control requirements of the organization, it is always a good idea to reinforce the policies by enforcing granular ACLs.

Before designing the ACLs, it should also be noted that, as the Cisco ASA inspects traffic it is able to recognize packets belonging to already established sessions. The stateful inspection engine of the firewall dynamically allows the returning traffic. Therefore, the firewall ACLs should be constructed to match traffic in the direction in which it is being initiated. In our sample configurations ACLs are applied in the ingress direction.

The following are the guidelines and configuration examples of ACLs controlling access and traffic flows:

- Ingress Inside

  Allow Internet access to users residing at all enterprise sites for the allowed ports and protocols. This typically includes HTTP and HTTPS access.

```
access-list Outbound extended permit tcp 10.0.0.0 255.0.0.0 any eq
http
access-list Outbound extended permit tcp 10.0.0.0 255.0.0.0 any eq
https
```

Allow users access to DMZ services such as the company's website, E-mail, and domain name resolution (HTTP, HTTPS, SMTP, POP, IMAP, and DNS). Note that the previous entries in the ACL already permit HTTP and HTTPS traffic.

```
! Allow DNS queries to DNS server
access-list Outbound extended permit udp 10.0.0.0 255.0.0.0 host
10.25.34.13 eq domain
```

```
! Allow SMTP, POP3 and IMAP access to DMZ mail server
access-list Outbound extended permit tcp 10.0.0.0 255.0.0.0 host
10.25.34.12 eq smtp
access-list Outbound extended permit tcp 10.0.0.0 255.0.0.0 host
10.25.34.12 eq pop3
access-list Outbound extended permit tcp 10.0.0.0 255.0.0.0 host
10.25.34.12 eq imap4
! Apply ACL to inside interface
access-group Outbound in interface inside
```

- Ingress DMZ

  Restrict connections initiated from DMZ only to the necessary protocols and sources. This typically includes DNS queries and zone transfer from DNS server, SMTP from E-mail server, HTTP/SSL access from the Cisco IronPort ESA for updates, Sensorbase, etc.

```
! Allow DNS queries and zone transfer from DNS server
access-list DMZ extended permit udp host 10.25.34.13 any eq domain
access-list DMZ extended permit tcp host 10.25.34.13 any eq domain
!
! Allow SMTP from Cisco IronPort ESA
access-list DMZ extended permit tcp host 10.25.34.11 any eq smtp
!
! Allow update and SensorBase access to Cisco IronPort ESA
access-list DMZ extended permit tcp host 10.25.34.11 any eq http
access-list DMZ extended permit tcp host 10.25.34.11 any eq https
!
! Apply ACL to DMZ interface
access-group DMZ in interface dmz
```

- Ingress Outside

  Inbound Internet access should be restricted to the public services provided at the DMZ such as SMTP, web, and DNS. Any connection attempts to internal resources and subnets from the Internet should be blocked. ACLs should be constructed using the servers' global IP addresses.

```
! Allow DNS queries and zone transfer to DNS server
access-list Inbound extended permit udp any host 198.133.219.13 eq
domain
access-list Inbound extended permit tcp any host 198.133.219.13 eq
domain
!
! Allow SMTP to Cisco IronPort ESA
access-list Inbound extended permit tcp any host 198.133.219.11 eq
smtp
!
! Allow HTTP/HTTPS access to the company's public web portal
access-list Inbound extended permit tcp any host 198.133.219.10 eq
http
access-list Inbound extended permit tcp any host 198.133.219.10 eq
https
!
! Apply ACL to outside interface
access-group Inbound in interface outside
```

Botnet Traffic Filter

The Small Enterprise Design Profile uses the Cisco ASA Botnet Traffic Filter on the Internet Firewall to detect malware that attempts network activity such as sending private data (passwords, credit card numbers, key strokes, or other proprietary data) when the malware starts a connection to a known bad IP address. The Botnet Traffic Filter checks incoming and outgoing connections against a dynamic database of known bad domain names and IP addresses (the blacklist), and then logs or blocks any suspicious activity.

Configuring the Botnet Traffic Filter requires the following steps:

1. Configure DNS Server
2. Enable Use of the Dynamic Database
3. Enable DNS Snooping
4. Enable Traffic Classification and Actions for the Botnet Traffic Filter
5. Verify and Monitor Botnet Traffic Filter Operation

The following sections provides configuration examples for each of these steps.

Configure DNS Server

The Botnet Traffic Filter requires a DNS Server to access Cisco's dynamic database update server and to resolve entries in the static database. The following configuration illustrates this configuration.

```
! Enable DNS requests to a DNS Server out the outside interface
dns domain-lookup outside
! Specify the DNS Server Group and the DNS Servers
dns server-group DefaultDNS
 name-server 68.238.112.12
 name-server 68.238.96.12
 domain-name cisco.com
```

Enable Use of the Dynamic Database

The Botnet Traffic Filter can receive periodic updates for the dynamic database from the Cisco update server. This database lists thousands of known bad domain names and IP addresses. The following configuration enables database updates, and also enables use of the downloaded dynamic database by the adaptive security appliance.

```
! enable downloading of the dynamic database from the Cisco Update server
dynamic-filter updater-client enable
! enable use of the dynamic database
dynamic-filter use-database
```

Enable DNS Snooping

DNS Snooping enables inspection of DNS packets and enables Botnet Traffic Filter Snooping which compares the domain name with those in the dynamic or static databases, and adds the name and IP address to the DNS reverse lookup cache. This cache is then used by the Botnet Traffic Filter when connections are made to the suspicious address.

It is recommended that DNS snooping is only enabled on interfaces where external DNS requests are going. Enabling DNS snooping on all UDP DNS traffic, including that going to an internal DNS server, creates unnecessary load on the Cisco ASA. For example, if the DNS server is on the outside interface, you should enable DNS inspection with snooping for all UDP DNS traffic on the outside interface.

The following configuration example illustrates enabling DNS Snooping on the outside interface:

```
! create a class map to identify the traffic you want to inspect DNS
class-map dynamic-filter-snoop-class
 match port udp eq domain
! create a policy map to enable DNS inspection with Botnet Traffic
Filtering snooping for the class map
policy-map dynamic-filter-snoop-policy
 class dynamic-filter-snoop-class
  inspect dns preset_dns_map dynamic-filter-snoop
! activate the policy map on the outside interface
service-policy dynamic-filter-snoop-policy interface outside
```

Enable Traffic Classification and Actions for the Botnet Traffic Filter

The Botnet Traffic Filter compares the source and destination IP address in each initial connection packet to the IP addresses in the dynamic database, static database, DNS reverse lookup cache, and DNS host cache, and sends a syslog message or drops any matching traffic. When an address matches, the Cisco ASA sends a syslog message and can optionally be configured to drop the connection. You can enable Botnet Traffic filter on a subset of traffic or for all traffic by enabling an access list to classify traffic.

The following configuration example enables the Botnet Traffic Filter feature on all traffic and additionally enables dropping of connections going to IP addresses with a severity of moderate and higher.

```
! identify the traffic that you want to monitor or drop.
access-list btf-filter-acl extended permit ip any any
! enable Botnet Traffic Filter on the outside interface for traffic
classified by the btf-filter-acl access list
dynamic-filter enable interface outside classify-list btf-filter-acl
! enable automatic dropping of traffic with threat level moderate or
higher
dynamic-filter drop blacklist interface outside action-classify-list
btf-filter-acl threat-level range moderate very-high
```

Botnet Traffic Filter Verification

To monitor and verify the operation of the Botnet Traffic Filter feature, the following commands can be used:

- show dynamic-filter updater-client—Shows information about the updater server, including the server IP address, the next time the adaptive security appliance will connect with the server, and the database version last installed.

```
cr12-asa-1-ie# show dynamic-filter updater-client
Dynamic Filter updater client is enabled
Updater server URL is https://update-manifests.ironport.com
Application name: threatcast, version: 1.0
Encrypted UDI:
0bb93985f42d941e50dc8f022350d1a8a8c5097dc1d252b9cff62d26d4ec58e202883
d704fc62b85bf8629fa757fe36b
Last update attempted at 15:14:11 UTC Apr 7 2010,
  with result: Downloaded file successfully
Next update is in 00:52:14
Database file version is '1270651144' fetched at 15:14:11 UTC Apr 7
2010, size: 2097152
cr12-asa-1-ie#
```

- show dynamic-filter data—Shows information about the updater server, including the server IP address, the next time the adaptive security appliance will connect with the server, and the database version last installed.

```
cr12-asa-1-ie# show dynamic-filter data
Dynamic Filter is using downloaded database version '1270651144'
Fetched at 15:14:11 UTC Apr 7 2010, size: 2097152
Sample contents from downloaded database:
  win-antimalware2.com  firstlook.com  red-devil-sport-club.gymdb.com
mswindowsupdate.info
  zardoz.wizardz.com  exchange.bg  bisexual-photo.com  lookmbox.com
Sample meta data from downloaded database:
  threat-level: very-high,      category: Malware,
  description: "These are sources that use various exploits to
deliver adware, spyware and other malware to victim computers.  Some
of these are associated with rogue online vendors and distributors of
dialers which deceptively call premium-rate phone numbers."
  threat-level: high,   category: Bot and Threat Networks,
  description: "These are rogue systems that control infected
computers.  They are either systems hosted on threat networks or
systems that are part of the botnet itself."
  threat-level: moderate,       category: Spyware,
  description: "These are sources that distribute spyware, adware,
greyware, and other potentially unwanted advertising software.  Some
of these also run exploits to install such software."
  threat-level: low,    category: Ads,
  description: "These are advertising networks that deliver banner
ads, interstitials, rich media ads, pop-ups, and pop-unders for
websites, spyware and adware.  Some of these networks send
ad-oriented HTML emails and email verification services."
Total entries in Dynamic Filter database:
  Dynamic data: 82119 domain names , 2565 IPv4 addresses
```

```
  Local data: 0 domain names , 0 IPv4 addresses
Active rules in Dynamic Filter asp table:
  Dynamic data: 0 domain names , 2565 IPv4 addresses
  Local data: 0 domain names , 0 IPv4 addresses
cr12-asa-1-ie#
```

- show dynamic-filter statistics detail—Shows how many connections were monitored and dropped with the Botnet Traffic Filter, and how many of those connections match the whitelist, blacklist, and greylist. (The greylist includes addresses that are associated with multiple domain names, but not all of these domain names are on the blacklist.) The detail keyword shows how many packets at each threat level were classified or dropped.

```
cr12-asa-1-ie# show dynamic-filter statistics detail
Enabled on interface outside using classify-list btf-filter-acl
 Total conns classified 35, ingress 0, egress 35
 Total whitelist classified 0, ingress 0, egress 0
 Total greylist classified 16, dropped 0, ingress 0, egress 16
  Threat-level  very-high: classified        0, dropped        0,
ingress        0, egress        0
  Threat-level       high: classified        0, dropped        0,
ingress        0, egress        0
  Threat-level   moderate: classified        0, dropped        0,
ingress        0, egress        0
  Threat-level        low: classified       16, dropped        0,
ingress        0, egress       16
  Threat-level   very-low: classified        0, dropped        0,
ingress        0, egress        0
 Total blacklist classified 19, dropped 0, ingress 0, egress 19
  Threat-level  very-high: classified        9, dropped        0,
ingress        0, egress        9
  Threat-level       high: classified        0, dropped        0,
ingress        0, egress        0
  Threat-level   moderate: classified        0, dropped        0,
ingress        0, egress        0
  Threat-level        low: classified       10, dropped        0,
ingress        0, egress       10
  Threat-level   very-low: classified        0, dropped        0,
ingress        0, egress        0
cr12-asa-1-ie#
```

Note To clear the statistics, enter the clear dynamic-filter statistics [*interface name*] command.

Other commands that are useful for monitoring the Botnet Traffic filter include the following:

- show dynamic-filter reports top [malware-sites | malware-ports | infected-hosts]—Generates reports of the top 10 malware sites, ports, and infected hosts monitored. The top 10 malware-sites report includes the number of connections dropped, and the threat level and category of each site. This report is a snapshot of the data, and may not match the top 10 items since the statistics started to be collected.
- show dynamic-filter reports infected-hosts {max-connections | latest-active | highest-threat | subnet ip_address netmask | all}—Generates reports about infected hosts. These reports contain detailed history about infected hosts, showing the correlation between infected hosts, visited malware sites, and malware ports. The max-connections keyword shows the 20 infected hosts with the most number of connections. The latest-active keyword shows the 20 hosts with the most recent activity. The highest-threat keyword shows the 20 hosts that connected to the malware sites with the highest threat level. The subnet keyword shows up to 20 hosts within the specified subnet. The all keyword shows all buffered infected-hosts information. This display might include thousands of entries. You might want to use ASDM to generate a PDF file instead of using the CLI.
- show dynamic-filter dns-snoop [detail]—Shows the Botnet Traffic Filter DNS snooping summary, or with the detail keyword, the actual IP addresses and names. All inspected DNS data is included in this output, and not just matching names in the blacklist. DNS data from static entries are not included.
- show asp table dynamic-filter [hits]—Shows the Botnet Traffic Filter rules that are installed in the accelerated security path.

Firewall Redundancy

The Internet perimeter of the Small Enterprise Design Profile uses a single Cisco ASA appliance configured with redundant interfaces. The use of redundant interfaces makes the design resilient to link level failures, representing an affordable option for high availability. In cases where chassis redundancy is desirable, the enterprise may consider deploying a pair of Cisco ASA appliances configured for stateful failover. Both active/active and active/standby failover modes are supported. While stateful failover protects against chassis failures, it requires the deployment of two identical Cisco ASA appliances and the adjustment of the topologies around the firewalls, so its deployment should be carefully planned.

This guide explains the use of redundant interfaces. For information on how to configure stateful failover, refer to the *Cisco ASA 5500 Series Adaptive Security Appliances Configuration Guides* at the following URL:

http://www.cisco.com/en/US/products/ps6120/products_installation_and_configuration_guides_list.htm

A Cisco ASA redundant interface is a logical interface that pairs two physical interfaces, called active and standby interfaces. Under normal operation the active interface is the only one passing traffic. The active interface uses the IP address defined at the redundant interface, and the MAC address of the first physical interface associated with the redundant interface. When the active interface fails, the standby interface becomes active and starts passing traffic. The same IP address and MAC address are maintained so that traffic is not disrupted. Figure 23 illustrates the concept of redundant interface.

Figure 23 Cisco ASA Redundant Interface

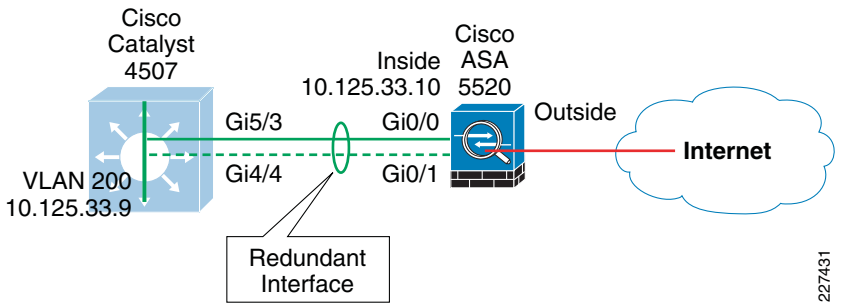

The configuration of a redundant interface consists in the configuration of the physical interface parameters and the logical redundant interface. Physical parameters such as media type, duplex, and speed are still configured within the physical interface. IP address, interface name, routing protocols, security level are configured as part of the redundant interface. The following configuration example corresponds to Figure 23.

```
! Physical interface and Ethernet parameters
interface GigabitEthernet0/0
 description Connected to cr24-4507-DO
 no nameif
 no security-level
 no ip address
!
interface GigabitEthernet0/1
 description backup to cr24-4507-DO
 no nameif
 no security-level
 no ip address
!
! Defines logical redundant interface associated with physical
! interfaces. Configures IP and logical interface parameters.
interface Redundant1
 description Connected to cr24-4507-DO
 member-interface GigabitEthernet0/0
 member-interface GigabitEthernet0/1
 nameif inside
 security-level 100
 ip address 10.125.33.10 255.255.255.0
 authentication key eigrp 100 <removed> key-id 1
 authentication mode eigrp 100 md5
!
```

Routing

An interior gateway protocol, EIGRP in our configuration examples, is used for dynamic routing. The Internet firewall may participate in routing by learning the internal routes and by injecting a default route pointing to the Internet. The default route should be removed dynamically if the Internet connection becomes unavailable.

As part of the Small Enterprise Design Profile, two different approaches were validated for the injection of the default route:

- *OSPF*—The Cisco ASA appliance learns the default route from the Internet border router using OSPF. The default route is then redistributed into EIGRP, and from there propagated into the rest of the internal network.
- *Static Route*—The Cisco ASA appliance is configured with a static default route pointing to the Internet gateway. Object tracking is configured to dynamically remove the default route when the Internet connection becomes unavailable. The default route is redistributed into EIGRP, and from there propagated into the rest of the internal network.

Injecting a default route with OSPF requires the configuration of an OSPF process between the Cisco ASA and the Internet border router, as illustrated in Figure 24. If the router is managed by the ISP, the configuration will require coordination with the service provider. This scenario also requires the default route to be propagated over OSPF. The actual default route may originate from the Internet border router itself or somewhere in the ISP network

Figure 24 OSPF Default Route Injection

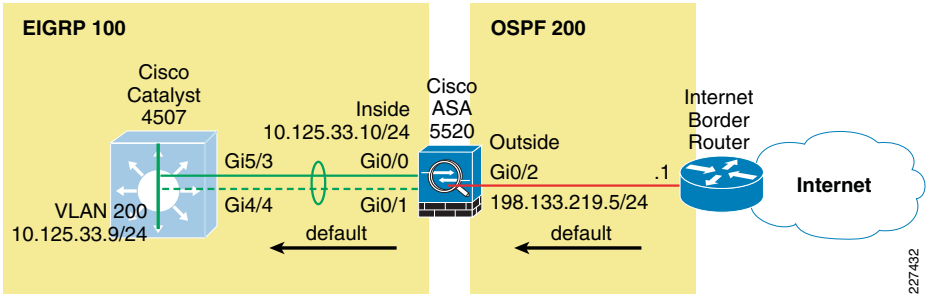

The following are the guidelines for using OSPF for the injection of a default route:

- Whenever possible, use MD5 authentication to secure the routing session between the Cisco ASA and the Internet border router.
- Since NAT is configured on the Cisco ASA and the inside address space is not visible outside the firewall, there is no need to redistribute routes from the internal EIGRP into OSPF.
- Route redistribution from OSPF into the internal EIGRP should be limited to the default route only. No other routes should be propagated into EIGRP.

The following configuration snippet illustrates the routing configuration of the Cisco ASA appliance. The configuration includes the route redistribution from OSPF into EIGRP with the enforcement of a route-map allowing only the injection of the default route. MD5 authentication is used for OSPF, and the logging of neighbor status changes is enabled.

```
! Permit default only
access-list Inbound-Routes standard permit host 0.0.0.0
```

```
!
interface GigabitEthernet0/2
 ospf message-digest-key 1 md5 <removed>
 ospf authentication message-digest
!
route-map Inbound-EIGRP permit 10
 match ip address Inbound-Routes
!
router eigrp 100
 no auto-summary
 network 10.125.33.0 255.255.255.0
 passive-interface default
 no passive-interface inside
 redistribute ospf 200 metric 1000000 2000 255 1 1500 route-map
Inbound-EIGRP
!
router ospf 200
 network 198.133.219.0 255.255.255.0 area 100
 area 100 authentication message-digest
 log-adj-changes
!
```

Note The hello-interval and dead-interval OSPF timers can be adjusted to detect topological changes faster.

The other validated alternative for the default route injection is the definition of a static default route, which then can be redistributed into the internal EIGRP process. This is shown in Figure 25. This option does not require the configuration of the Internet border router.

Figure 25 Static Default Route with Object Tracking

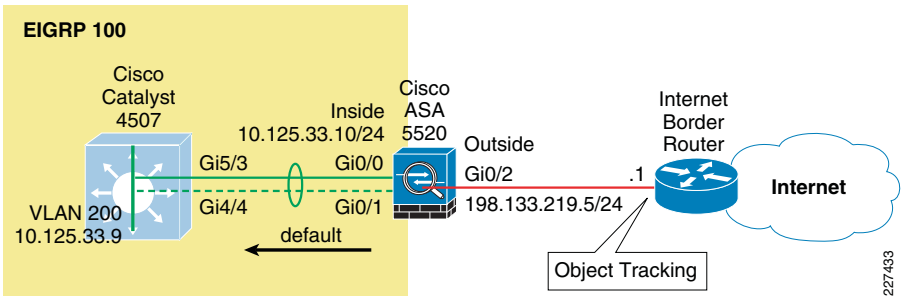

It is highly recommended to use object tracking so the default route is removed when the Internet connection becomes unavailable. Without object tracking, the default route will be removed only if the outside interface of the appliance goes down. So there is a possibility that the default route may remain in the routing table even if the Internet border router becomes unavailable. To avoid that problem, the static default route can be configured with object tracking. This consists in associating the default route with a monitoring target.

The Cisco ASA appliance monitors the target using ICMP echo requests. If an echo reply is not received within a specified time period, the object is considered down and the associated default route is removed from the routing table.

The monitoring target needs to be carefully selected. First, pick one that can receive and respond to ICMP echo requests sent by the Cisco ASA. Second, it is better to use a persistent network object. In the configuration example below the Cisco ASA monitors the IP address of the next hop gateway, which helps identifying if the Internet gateway goes down, but it will not help if the connection is lost upstream. If available, you may want to monitor a persistent network object located somewhere in the ISP network. Static route tracking can also be configured for default routes obtained through DHCP or PPPoE.

In the following configuration the IP address of the next hop gateway (198.133.219.1) is used as the monitoring target. The static default route is then redistributed into EIGRP.

```
router eigrp 100
 no auto-summary
 network 10.125.33.0 255.255.255.0
 passive-interface default
 no passive-interface inside
 redistribute static metric 1000000 2000 255 1 1500
!
route outside 0.0.0.0 0.0.0.0 198.133.219.1 1 track 10
!
sla monitor 1
 type echo protocol ipIcmpEcho 198.133.219.1 interface outside
sla monitor schedule 1 life forever start-time now
!
track 10 rtr 1 reachability
```

Note The *frequency* and *timeout* parameters of object tracking can be adjusted to detect topological changes faster.

Intrusion Prevention Deployment

The Small Enterprise Design Profile implements Intrusion Prevention using an Advanced Inspection and Prevention Security Services Module (AIP SSM) on the Cisco ASA appliance deployed at the Internet perimeter. This section describes the best practices for integrating and configuring the IPS module for maximum threat control and visibility, as well as the deployment of the IPS Global Correlation feature.

Deploying IPS with the Cisco ASA

The Advanced Inspection and Prevention Security Services Module (AIP SSM) is supported on Cisco ASA 5510 and higher platforms. The AIP SSM runs advanced IPS software that provides proactive, full-featured intrusion prevention services to stop malicious traffic, including worms and network viruses, before they can affect your network.

As described in the "Intrusion Prevention Guidelines" section on page -5, the AIP SSM may be deployed in inline or promiscuous mode. In inline mode the AIP SSM is placed directly in the traffic flow, while in promiscuous mode the Cisco ASA sends a duplicate stream of traffic to the AIP SSM.

When deploying the AIP SSM in inline mode it is particularly important to determine how traffic will be treated in case of a module failure. The AIP SSM card may be configured to fail open or close when the module becomes unavailable. When configured to fail open, the Cisco ASA appliance allows all traffic through, uninspected, if the AIP SSM becomes unavailable. Per contrary, when configured to fail close, the adaptive security appliance blocks all traffic in case of an AIP SSM failure.

The following example illustrates how a Cisco ASA can be configured to divert all IP traffic to the AIP SSM in inline mode, and to block all IP traffic if the AIP SSM card fails for any reason:

```
access-list IPS permit ip any any
class-map my-ips-class
  match access-list IPS
policy-map my-ips-policy
  class my-ips-class
    ips inline fail-close
service-policy my-ips-policy global
```

IPS Global Correlation Deployment

There are a number of aspects that need to be understood prior to the deployment of the IPS Global Correlation feature. To start with, before configuration be sure that you are using Cisco IPS Version 7.0 with the latest patch and signature updates and that Cisco IPS is configured for network connectivity in either IDS or IPS mode.

The configuration of Global Correlation can be performed using the command-line interface (CLI), Cisco IDS Device Manager (IDM), Cisco IME, or the Cisco Security Manager. The figures here provided are screenshots from Cisco IME that illustrate the basic steps in the configuration of IPS Global Correlation.

The first step in configuring the IPS sensor (module or appliance) to use Global Correlation is to add either a DNS address and/or the proxy server setup. This step, illustrated in Figure 26, enables connection to Cisco SensorBase. After you configure the DNS and proxy settings, these settings will go into effect as soon as the sensor has downloaded the latest Global Correlation updates.

Figure 26 DNS and HTTP Proxy Within the Network Setting Configuration Screen

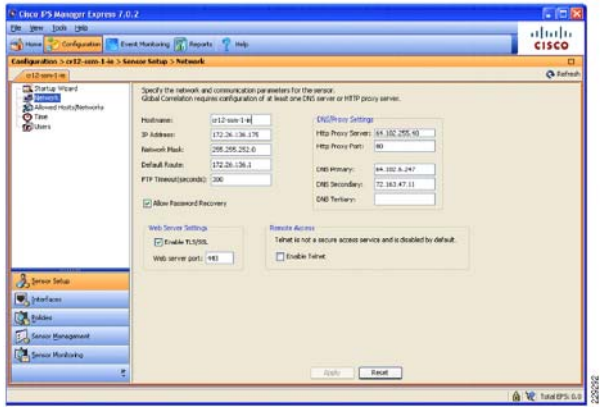

By default, a sensor runs Global Correlation Inspection in Standard mode, and enables the reputation filters, both illustrated in Figure 27. A good practice is to configure Global Correlation Inspection initially in permissive mode while monitoring the effects, and later change the configuration into Standard or Aggressive mode.

Figure 27 Global Correlation Inspection Settings

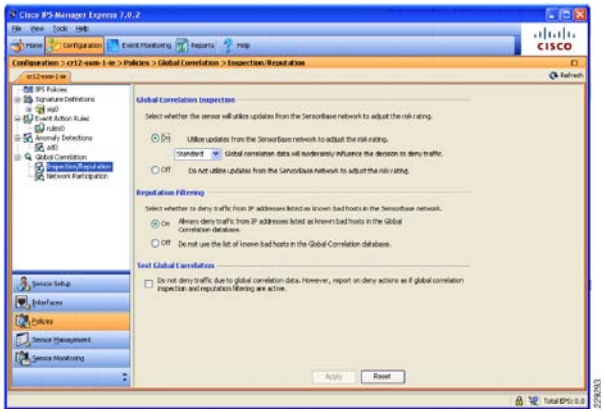

By default network participation is disabled, which means the sensor does not share any event data back to the Cisco SensorBase network. The event data provided by all devices participating in the Cisco SensorBase network is a key element that provides real-time and worldwide visibility into the threat activity, accelerating the identification and mitigation of threats propagating throughout the Internet. For this reason, Cisco recommends to configure the IPS sensors with partial or full network participation. See Figure 28.

Figure 28 Network Participation Settings (Off by Default)

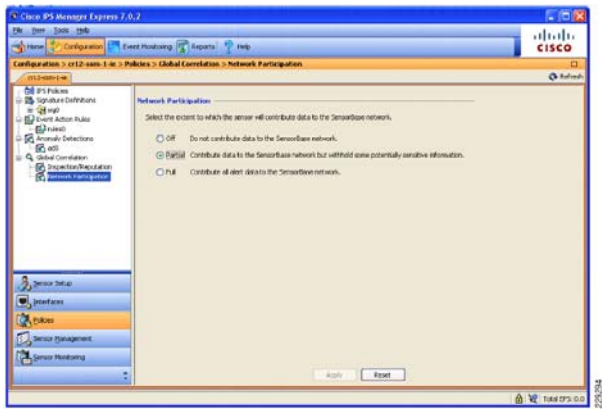

Event Monitoring with Global Correlation

Event monitoring with Global Correlation is similar to event monitoring with signature-only IPS. The primary difference is the potential addition of reputation scores representing the Global Correlation data. Figure 29 shows Cisco IPS events with reputation scores in Cisco IME.

Figure 29 Event Monitoring with Global Correlation in Cisco IME

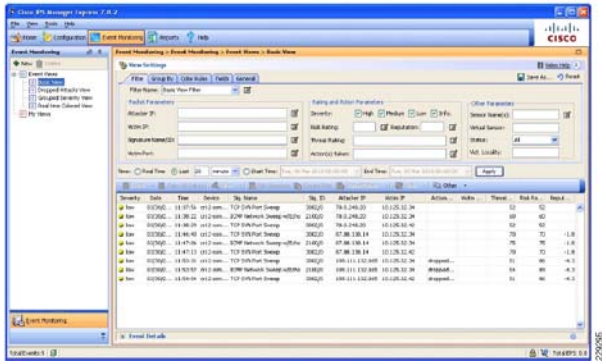

Figure 29 shows several TCP SYN Port Sweep and ICMP Network Sweep attacks that were seen by the sensor. The first three events had no reputation, and the event's risk ratings were 52 and 60 which did not meet the threshold for the packets to be dropped. The next three events were identical except that the attacker had a negative reputation of -1.8 which elevated the risk ratings to 70 and 75 which still did not meet the thresholds to be dropped in Standard Mode. The last events were also identical except this time the attacker has a negative reputation if -4.3 which elevated the risk ratings to 86 and 89. This time the risk rating was high enough for the packets to be dropped.

Figure 30 illustrates the detail view of the TCP SYN Port Sweep event coming from the attacker with a negative reputation of -4.3.

Figure 30 Detailed view of a TCP SYN Port Sweep from an Attacker with a Negative Reputation Score

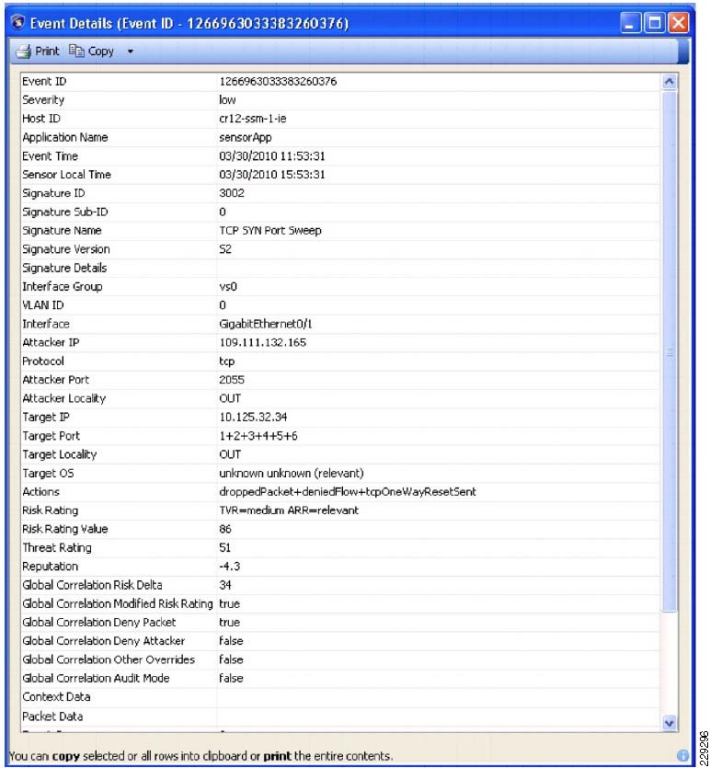

Web Security Deployment

The Small Enterprise Design Profile implements a Cisco IronPort WSA at the core/distribution layer of the main site, as illustrated in Figure 31. The WSA is located at the inside of the Cisco ASA acting as the Internet firewall. That ensures that clients and WSA are reachable over the same inside interface of the firewall, and that the WSA can communicate with them without going through the firewall. At the same time, deploying the WSA at the core/distribution layer gives complete visibility to the WSA on the traffic before getting out to the Internet through the firewall.

Figure 31 Cisco IronPort WSA deployment

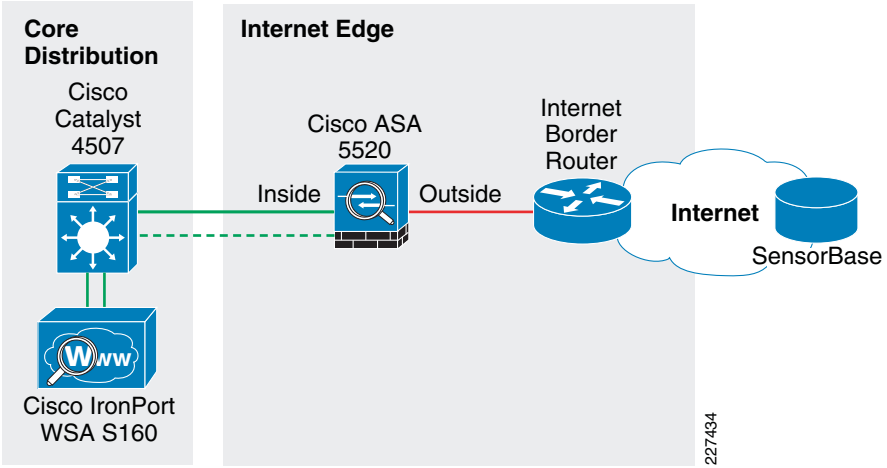

Following are the guidelines for the WSA configuration and deployment.

Initial System Setup Wizard

The WSA provides a browser-based system setup wizard that must be executed the first time the appliance is installed. The System Setup Wizard guides the user through initial system configuration such as network and security settings. It should be noted that running the initial System Setup Wizard completely reconfigures the WSA appliance and resets the administrator password. Only use the System Setup Wizard the first time you install the appliance, or if you want to completely overwrite the existing configuration.

The following are some of the default settings when running the System Setup Wizard:

- Web Proxy is deployed in transparent mode.
- The L4 Traffic Monitor is active and set to monitor traffic on all ports.

All these settings can be changed any time after the initial configuration by running the WSA web-based configuration GUI.

Interface and Network Configuration

The following need to be completed as part of the initial setup of the WSA appliance:

1. Configuring network interfaces
2. Adding routes
3. Configuring DNS
4. Setting time
5. Working with upstream proxy (if present)

Configuring Network Interfaces

Independently from the model, all Cisco IronPort WSA appliances are equipped with six Ethernet interfaces as shown in Figure 32.

Figure 32 WSA Interfaces

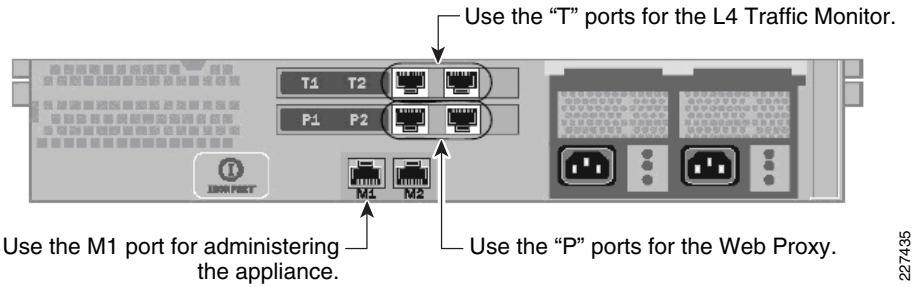

The WSA interfaces are grouped for the following functions:

- *Management*—Interfaces M1 and M2 are out-of-band (OOB) management interfaces. However, only M1 is enabled. In the Small Enterprise Design Profile, interface M1 connects to the out-of-band management network. Interface M1 can optionally be used to handle data traffic in case the enterprise does not have an out-band management network.
- *Web Proxy*—Interfaces P1 and P2 are Web Proxy interfaces used for data traffic. Only the P1 interface is used in the Enterprise Design Profile for Small Enterprise Networks. P1 connects to the inside subnet of the firewall.
- *L4 Traffic Monitor (L4TM)*—T1 and T2 are the L4TM interfaces. L4TM is not used in the Small Enterprise Design Profile, Cisco IPS Global Correlation and the Cisco ASA Botnet Filter features are used instead.

Figure 33 illustrates the network topology around the WSA used in the Cisco validation lab.

Figure 33 WSA Network Topology

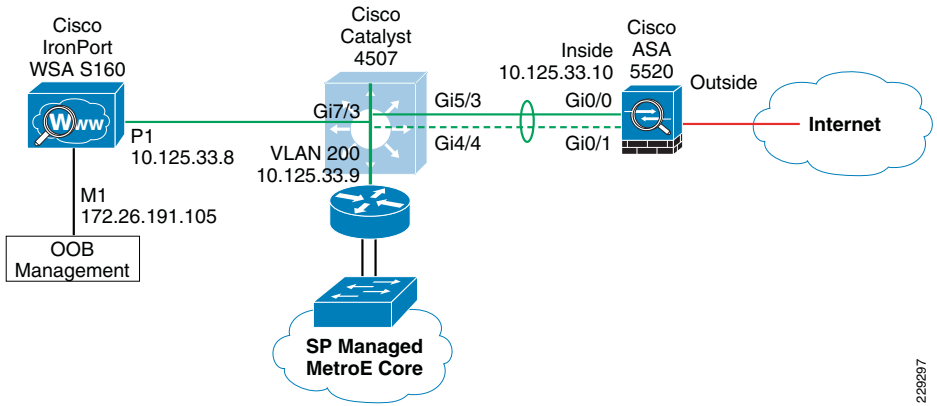

Figure 34 illustrates the IP address and hostname configurations for the interfaces used. In this case, an out-of-band management network is used; therefore the M1 port is configured with an IP address in the management subnet. In addition, the WSA is configured to maintain a separate routing instance for the M1 management interface. This allows the definition of a default route for management traffic separate from the default route used for data traffic.

Figure 34 WSA Interface Configuration

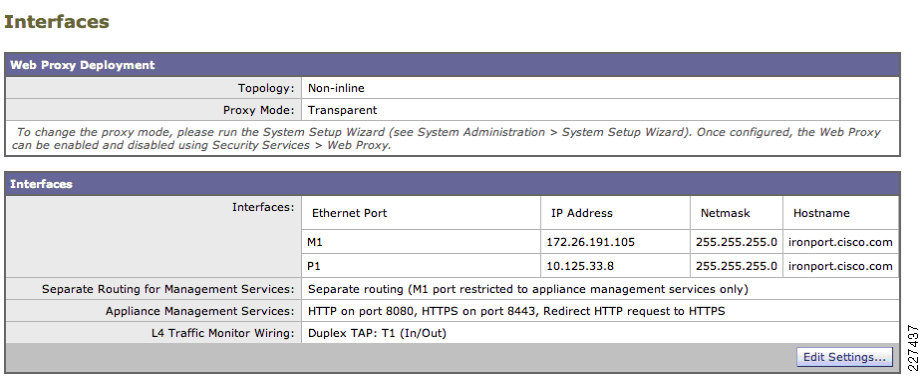

Adding Routes

A default route is defined for management traffic pointing to the OOB management default gateway (172.26.191.1). A separate default route is defined for the data traffic pointing to the inside IP address of the firewall (10.125.33.10). As all internal networks are reachable throughout the core/distribution switch, a route to 10.0.0.0/8 is defined pointing to the switch IP address (10.125.33.9) to allow the WSA to communicate with the clients directly without having to go to the firewall first. These settings are illustrated in Figure 35.

Figure 35 WSA Route Configuration

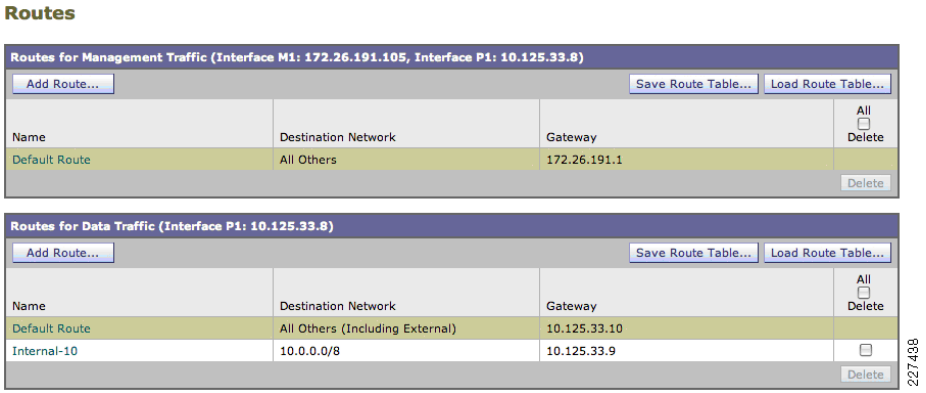

Configuring DNS

The initial setup requires the configuration of a hostname for the WSA appliance, and listing the DNS servers. Figure 36 shows the DNS configuration.

Figure 36 WSA DNS Configuration

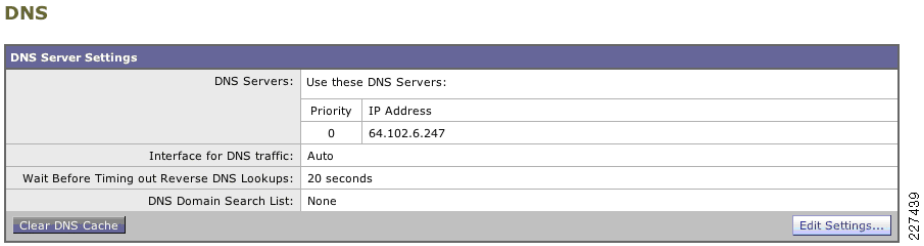

Setting Time

Time synchronization is critical for forensic analysis and troubleshooting, therefore enabling NTP is highly recommended. Figure 37 shows how the WSA is configured to synchronize its clock with an NTP server located on the OOB management network.

Figure 37 WSA NTP Configuration

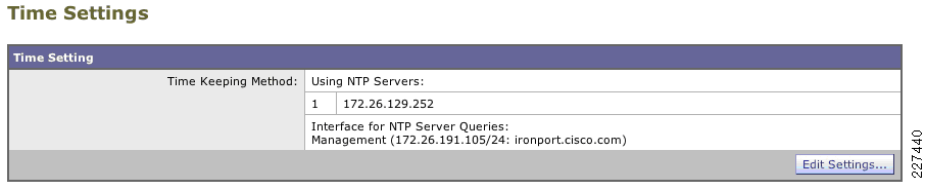

Working with Upstream Proxies

If Internet access is provided by an upstream proxy, then the WSA must be configured to use the proxy for component updates and system upgrades. This is illustrated in Figure 38 and Figure 39.

Figure 38 WSA Upgrade Settings

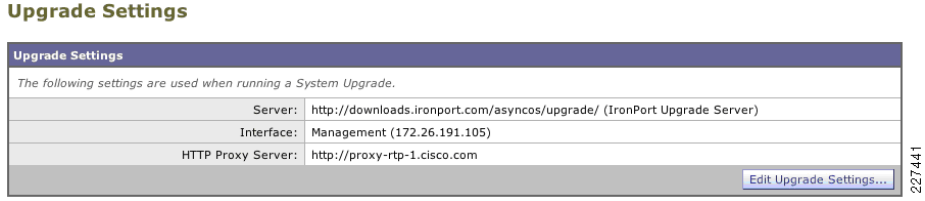

Figure 39 WSA Component Updates

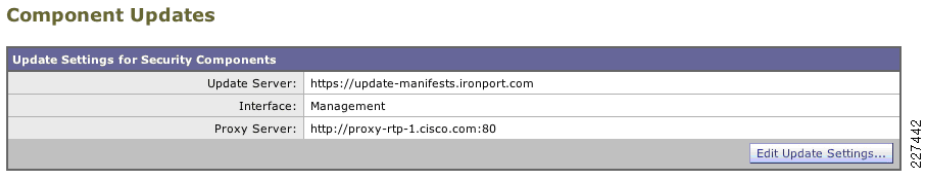

WCCP Transparent Web Proxy

The configuration of the WCCP Transparent Web Proxy includes the following:

1. Defining WSA WCCP Service Group
2. Enabling WSA Transparent Redirection
3. Enabling WCCP redirection on the Cisco ASA
4. Enabling WSA HTTPS scanning
5. Working with upstream proxy (if present)

Defining WSA WCCP Service Group

Web Proxy settings are configured as part of an initial setup using the System Setup Wizard and can be later modified with the WSA Web-based GUI. The Web Proxy setting include the following:

- *HTTP Ports to Proxy*—List the ports to be proxied. Default is 80 and 3128.
- *Caching*—Defines whether or not the WSA should cache response and requests. Caching helps reduce latency and the load on the Internet links. Default is enabled.
- *IP Spoofing* — Defines whether or not the Web Proxy should spoof IP addresses when forwarding requests to upstream proxies and servers. The Cisco ASA does not support source address spoofing.

Figure 40 illustrates the Web Proxy settings.

Figure 40 WSA Proxy Settings

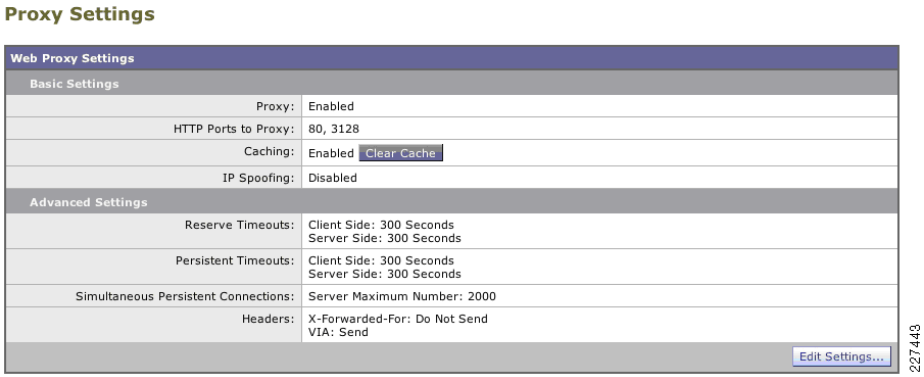

Enabling WSA Transparent Redirection

Configuring WCCP Transparent Redirection requires the definition of a WCCP service profile in the WSA. If redirecting HTTP and HTTPS, define a dynamic service ID to be used with the Cisco ASA. Use MD5 authentication to protect the WCCP communication between the WSA and Cisco ASA. Figure 41 shows an example.

Figure 41 WSA Transparent Proxy

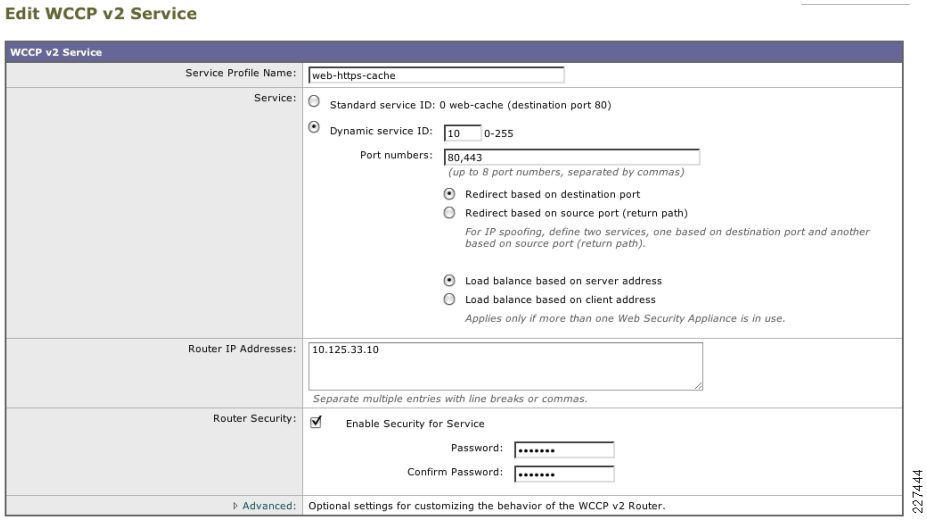

Enabling WCCP Redirection on Cisco ASA

The configuration of WCCP on the Cisco ASA appliance requires:

- A group-list indicating the IP addresses of the appliances member of the service group. In the example provided below the group-list is called wsa-farm.
- A redirect-list indicating the ports and subnets of traffic to be redirected. In the example, the ACL named proxylist is configured to redirect any HTTP and HTTPS traffic coming from the 10.0.0.0/8 subnet. It is critical to ensure traffic from the WSA(s) bypasses redirection. To that end, add an entry to the redirect-list explicitly denying traffic sourced from the WSA(s).
- WCCP service indicating the service ID. Make sure you use the same ID as defined on the WSAs. Use a password for MD5 authentication.
- Enabling WCCP redirection on an interface. Apply the WCCP service on the inside interface of the Cisco ASA.

The following is a Cisco ASA WCCP configuration example:

```
! Group-list defining the IP addresses of all WSAs
access-list wsa-farm extended permit ip host 10.125.33.8 any
!
! Redirect-list defining what ports and hosts/subnets should be
redirected
access-list proxylist extended deny ip host 10.125.33.8 any
access-list proxylist extended permit tcp 10.0.0.0 255.0.0.0 any eq www
access-list proxylist extended permit tcp 10.0.0.0 255.0.0.0 any eq https

!
! WCCP service
wccp 10 redirect-list proxylist group-list wsa-farm password cisco
```

```
!
! Applies WCCP on an interface
wccp interface inside 10 redirect in
```

The WCCP connection status and configuration can be monitored on the Cisco ASA with the show wccp command, as shown below:

```
cr26-asa5520-do# show wccp

Global WCCP information:
    Router information:
    Router Identifier:                   198.133.219.5
    Protocol Version:                    2.0

    Service Identifier: 10
    Number of Cache Engines:             1
    Number of routers:                   1
    Total Packets Redirected:            428617
    Redirect access-list:                proxylist
    Total Connections Denied Redirect:   0
    Total Packets Unassigned:            4
    Group access-list:                   wsa-farm
    Total Messages Denied to Group:      0
    Total Authentication failures:       0
    Total Bypassed Packets Received:     0
cr26-asa5520-do#
```

Enabling WSA HTTPS Scanning

To monitor and decrypt HTTPS traffic, you must enable HTTPS scanning on the WSA. The HTTPS Proxy configuration is illustrated in Figure 42.

Figure 42 WSA HTTPS Proxy

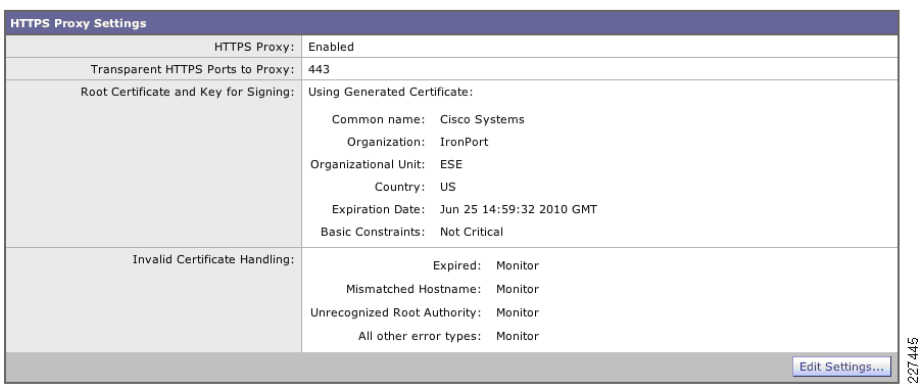

Working with Upstream Proxies

In case Internet traffic is handled by one or more upstream proxies, follow these guidelines:

- Add an Upstream Proxy Group
- Define a routing policy to direct traffic to the upstream proxies

The Upstream Proxy Group lists the IP addresses or domain names of the proxies to be used for traffic sent to the Internet. When multiple proxies are available, the WSA can be configured for failover or load balancing.

The following are the options available:

- *None (failover)*—The first proxy in the list is used. If one proxy cannot be reached, the Web Proxy attempts to connect to the next one in the list.
- *Fewest connections*—Transactions are directed to the proxy servicing the fewest number of connections.
- *Hash-based*—Requests are distributed using a hash function. The function uses the proxy ID and URL as inputs so that requests for the same URL are always directed to the same upstream proxy.
- *Least recently used*—Transactions are directed to the proxy that least recently received a transaction if all proxies are currently active.
- *Round robin*—The Web Proxy cycles transactions equally among all proxies in the group in the listed order.

Figure 43 illustrates the upstream proxy group configuration. Two upstream proxies are used, and transactions are forwarded to the proxy servicing the fewest number of connections.

Figure 43 WSA Upstream Proxy Group

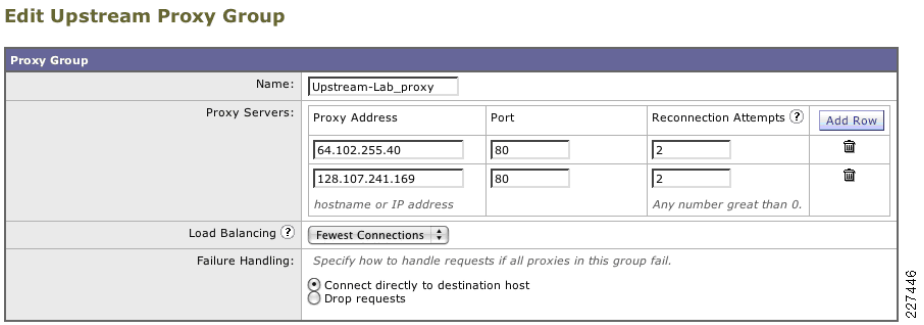

Next, a routing rule needs to be defined to indicate when and how to direct transactions to the upstream proxy group. Use the Global Routing Policy if all traffic is to be handled by the upstream proxies. If no proxies are present, then leave the routing destination of the Global Routing Policy configured as Direct Connection. Figure 44 presents an example where all traffic is directed to the proxies in the Upstream-Lab_proxy group.

Figure 44 WSA Routing Policies

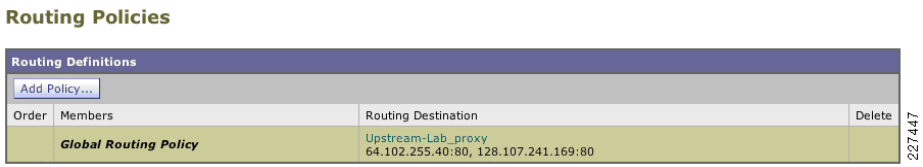

Web Access Policies

The access policies define how the Web Proxy handles HTTP requests and decrypted HTTPS connections for network users. By configuring access policies the enterprise can control what Internet applications (instant messaging clients, peer-to-peer file-sharing, web browsers, Internet phone services, etc.) and URL categories users may access. In addition, access policies can be used to block file downloads based on file characteristics, such as file size and file type.

The WSA comes with a default Global Policy that applies to all users. However, multiple policies can be defined when different policies need to be applied to different group of users. Figure 45 shows the global policy.

Figure 45 Global Access Policy

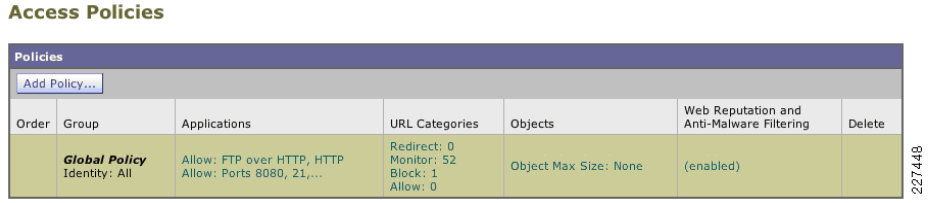

URL categories corresponding to non-business related content should be blocked in compliance with the company's Internet access policies. Figure 46 provides an example on how the "Adult/Sexually Explicit" category is blocked.

Figure 46 URL Categories

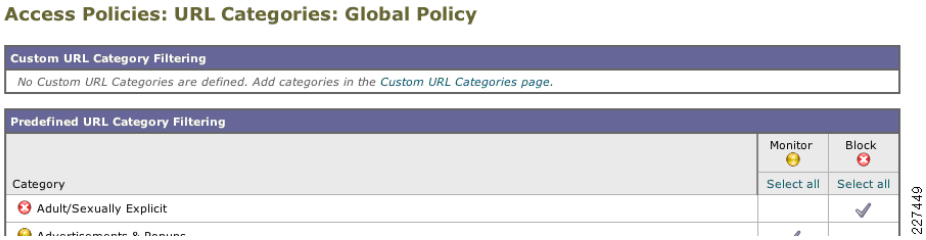

CISF Protected Ports

Catalyst Integrated Security Features (CISF) is a set of native security features available on Cisco Catalyst Switches that protect the network against attacks such as man-in-the-middle, spoofing, and infrastructure denial-of-services (DoS) attacks. CISF includes the following:

- Port Security
- DHCP snooping
- IP Source Guard
- Dynamic ARP inspection
- ARP rate limiting
- Storm Control

The following is a sample CISF port configuration:

```
! configure port security parameters
switchport port-security maximum 2
switchport port-security
switchport port-security aging time 2
switchport port-security violation restrict
switchport port-security aging type inactivity
! configure arp inspection rate limiting
ip arp inspection limit rate 100
! configure dhcp snooping rate limiting
ip dhcp snooping limit rate 100
! configure storm control parameters
storm-control broadcast level 20.00 10.00
storm-control multicast level 50.00 30.00
```

NAC Appliance Deployment

In the Small Enterprise Design Profile, a NAC Appliance solution is deployed at all locations, the main site and each of the remote offices. To that end, a centralized CAM is deployed at the main site, likely within the serverfarm. And a CAS is deployed at the main site and each remote site, and each of which directly connects to the core/distribution layer at each of the locations.

Deploying the NAC appliance solution requires the configuration of the CAM and CAS components. The initial CAM and CAS configuration is performed via console access, and this is described in the installation guide, *Cisco NAC Appliance Hardware Installation*. During the configuration stage the multiple steps must be followed in configuring the NAC Appliance with IP addresses, VLANs. passwords, etc. The installation guide contains worksheets that assist in the gathering and preparation of this information for both CAM and CAS.

After performing the initial setup through the console, the rest of the configuration of the CAM and CAS is performed using the CAM web-based GUI. The first task on the CAM before beginning configuration is the installation the licenses for the solution. A license must be installed for the CAM, and the CAS servers controlled by the CAM. The *Cisco NAC Appliance Ordering Guide*, provides information on the ordering options.

Licenses can be entered via the CAM web-based GUI, as shown in Figure 47.

Figure 47 NAC Appliance Licensing

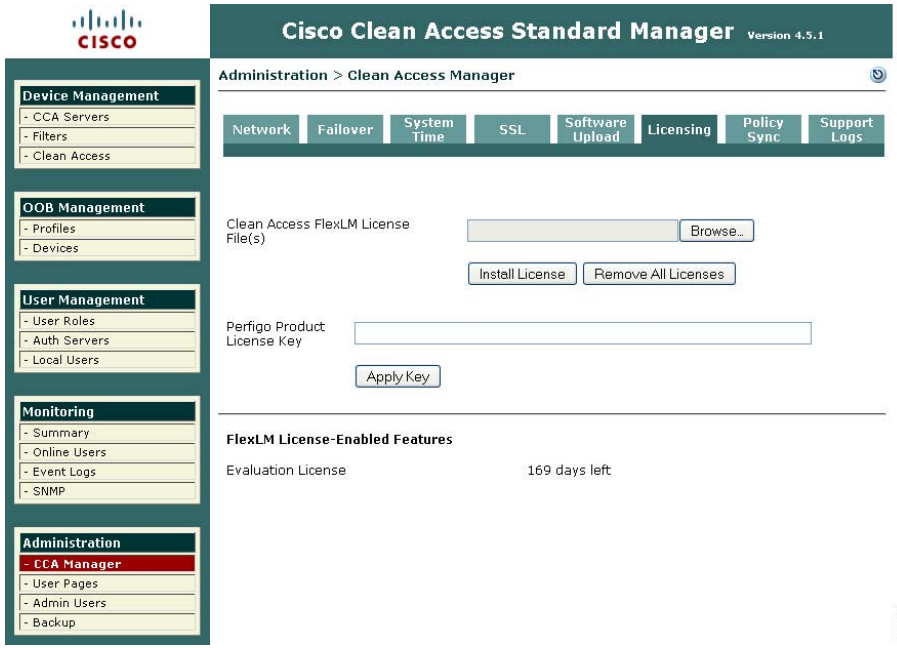

Adding a CAS to the CAM

For a CAS server to be managed by the CAM, it must be first added to the list of managed servers on the CAM. To do this the CAM needs to know the IP address of the CAS, and the Server Type (its role in the network) of the CAS. In addition to this, the CAS and the CAM must have the same shared secret. The shared secret is configured during the server installation. An example of adding a CAS to the CAM is shown in Figure 48.

Figure 48 Adding a new CAS to the CAM

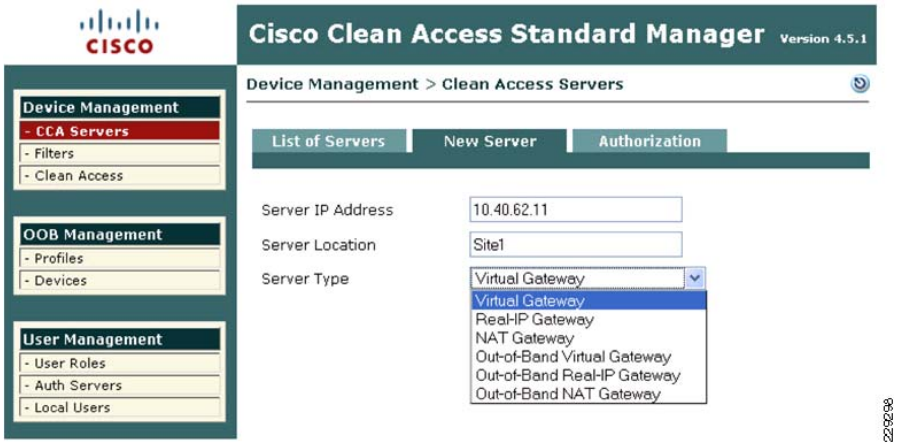

Once the CAS has been added to the CAM, it appears in the list of servers on the CAM. From this point, it can be managed directly from the CAM for almost all tasks. An example of a list of servers is shown in Figure 49.

Figure 49 List of CAS Servers

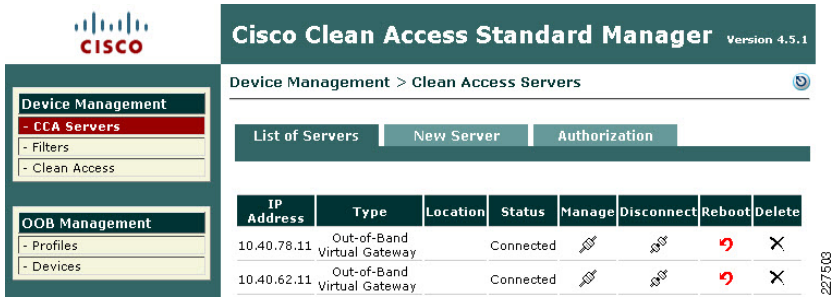

Managing the CAS

Once the CAS is in the list of servers managed by the CAM, it can be configured further for its role in the network. To manage the server click the icon under the Manage heading in the server list, this will connect you to the CAS server and present you with the summary menu shown in Figure 50.

Figure 50 AS Management Menu

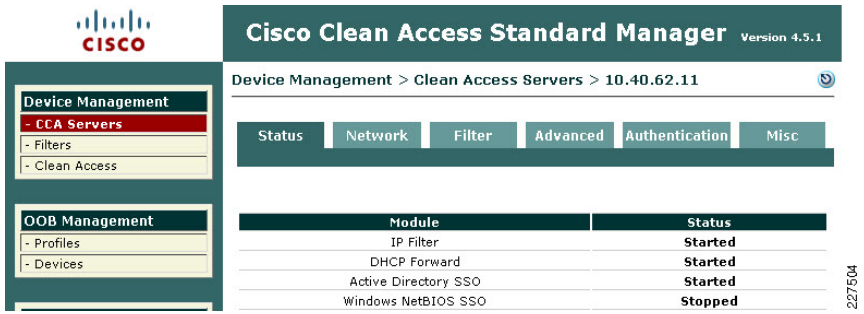

Under the CAS Network setting tab, shown in Figure 51, the basic network settings for the CAS can be seen and altered, if needed. In this example, we are keeping the network configuration performed during the server installation. The primary dialog under the Network Tab is the IP dialog, shown in Figure 51, the other dialogs allow the DHCP options to the configured—our example uses the default of DHCP passthrough— and the DNS options where host name, domain name, and DNS server information is added, as shown in Figure 52.

Figure 51 CAS Network Settings

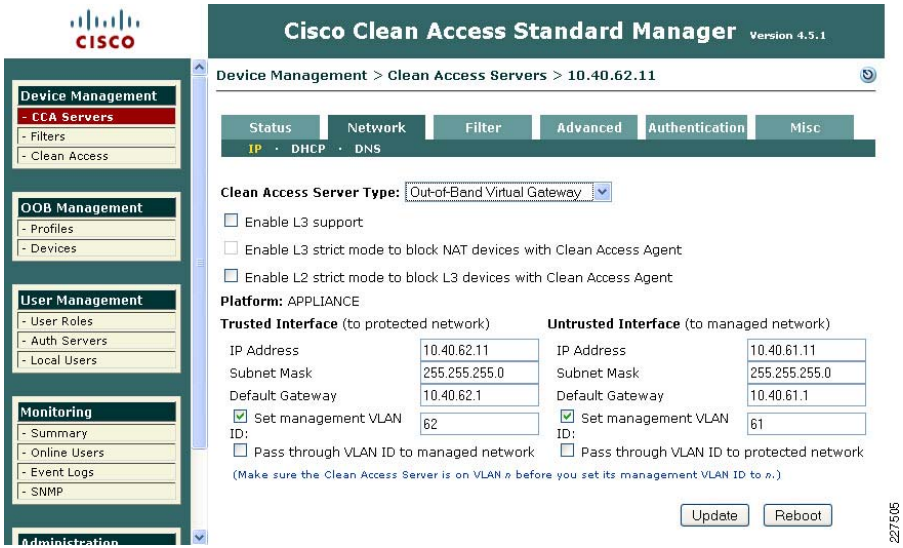

Figure 52 CAS DNS Settings

The next tab in the CAS configuration is the Filter tab (see Figure 53), for the purposes of our example the important dialog is the Roles where network traffic filters may be applied to different user Roles. The Role of interest at this moment is the default Unauthenticated Role. By default the Unauthenticated Role blocks all traffic. In this example we are allowing the Unauthenticated Role to pass Active Directory client authentication traffic to pass to the Active Directory Sever. This will allow a windows client to join the active Directory Domain, and windows users to authenticate to the domain although they have not been through the NAC process. This is often important to allow printer and drive mapping information to be sent to the winders users. As the user has already authenticated to the Active Directory Domain the user authentication information maybe learned from Active Directory, and the user does not have to reauthenticate for the NAC server.

Note The creation of Roles and their associated filters is performed in the CAM User Management -> User Roles menu.

Figure 53 CAS Filter Settings

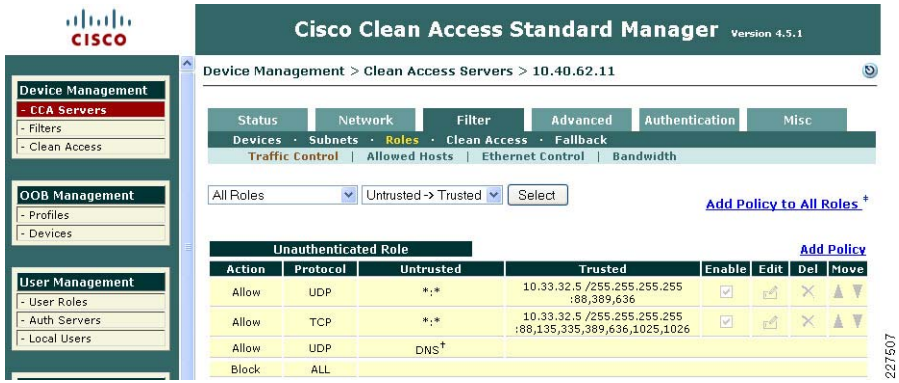

The next tab that requires configuration is the Advanced tab. This has multiple dialogs that require configuration. The first of these is the Managed Subnet dialog, where each of the trusted VLAN subnets is added to the CAS for management. An example of this shown in Figure 54.

Figure 54 CAS Managed Subnet

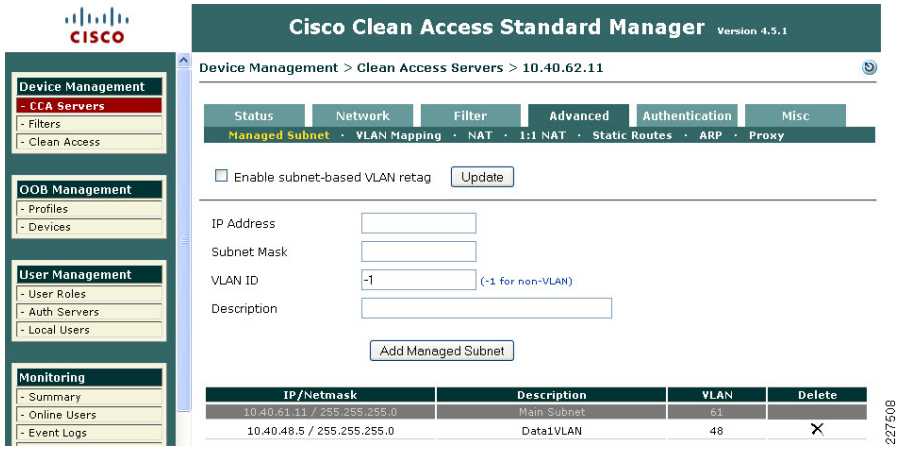

The next dialog of the Advanced Tab is the VLAN Mapping dialog, which tells the CAS which trusted VLAN to be mapped to an untrusted VLAN, an example of this is shown in Figure 55. In our example, VLAN Prunning and VLAN Mapping is also enabled.

Figure 55 VLAN Mapping

The next tab of interest is the Authentication Tab (see Figure 56), this tab has multiple dialogs for configuring different authentication options. The first dialog is the Login Page Dialog. This allows the configuration of different web login pages depending upon the untrusted subnet being used for authenticating client.

Figure 56 Authentication Login Page

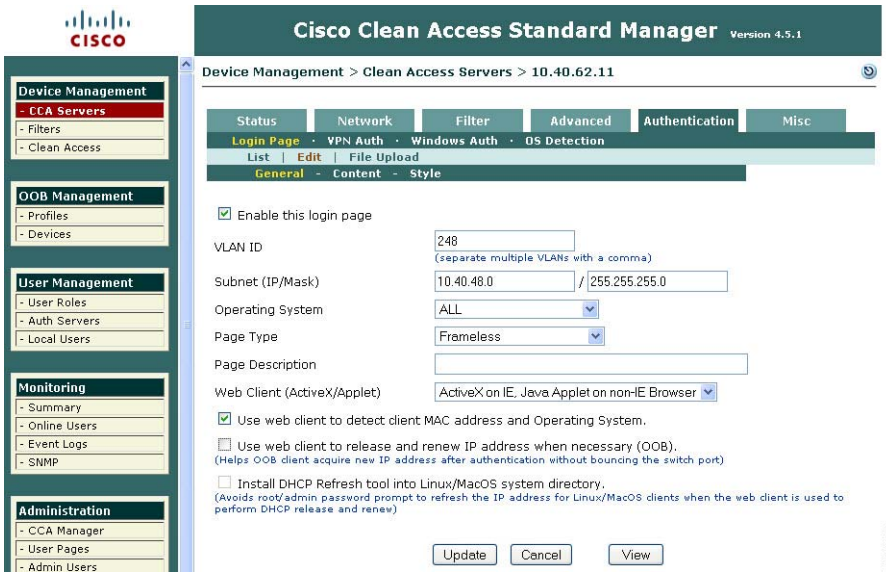

The other Authentication dialog of interest in this example is the Windows Auth dialog, as Windows Single Sign On (SSO) is used in this example. To perform Windows SSO the CAS needs to be able to communicate with Active Directory to determine the authentication state of the windows user. If Active Directory confirms that the user has authenticated to Active Directory the user doesn't need to perform additional authentication to the CAS. An example of this configuration is shown in Figure 57. There are a number of steps required to configure Active Directory SSO, as these are described in the *Cisco NAC Appliance —Clean Access Server Installation and Configuration Guide*. The key components in this configuration are:

- The creation of a Active Directory client account for the CAS
- Using the KTPass Application on Active Directory to convert the account encryption to DES encryption

Figure 57 CAS Windows Authentication

Clean Access Roles

The unauthenticated role is common to all clients, but once the client has been authenticated a different role may be applied based upon the identity of the client, different roles may be assigned for admin staff, users, etc.

User roles allow you to aggregate various policies into a user role. These policies include:

- Traffic policies
- Bandwidth policies

Note If bandwidth policies are to be enforced by the Clean Access Server it must be operating in band.

- VLAN ID retagging
- Clean Access network port scanning plugins
- Clean Access Agent/Cisco NAC Web Agent client system requirements

For example, the Admin and User roles could each have different traffic polices and VLANs, in addition the User role may enforce bandwidth policies by keeping the user traffic in band.

For more information on roles, refer to the *Cisco NAC Appliance - Clean Access Manager Installation and Configuration Guide* at the following URL:

http://www.cisco.com/en/US/docs/security/nac/appliance/configuration_guide/45/cam/45cam-book.html

Layer 2 OOB Example

Figure 58 shows an example of a Layer-2 OOB deployment. Where a wired client connected to an access switch is originally on the untrusted VLAN 264 and is switched to a trusted VLAN 64 once it has completed the NAC functions. The first NAC function is the authentication and authorization function, and this is the first design decision in implementing the NAC solution. That is, how will authentication and authorization be achieved, and what will the user experience be.

This example is focused upon the virtual gateway example, as virtual gateway provides the simplest deployment. In the virtual gateway example the original IP addressing, interfaces, and VLANs are maintained, and normal traffic flows are maintained. The only changes are the addition of the untrusted VLANs that carry client traffic during the NAC Authentication and Authorization, Scanning and evaluation, remediation, and quarantine modes.

Figure 58 Layer 2 OOB Example

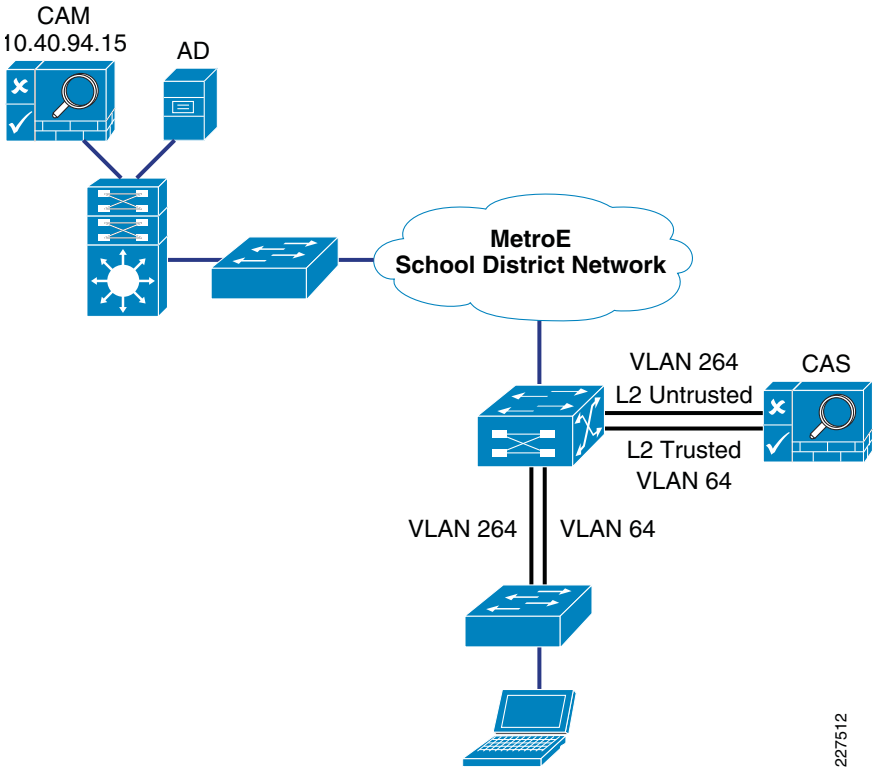

NAC Authentication Options

The authentication option in the NAC solution can be broadly categorized as NAC Authentication or NAC Single Sign On

- *NAC Authentication*—NAC authentication gives the NAC system the role of authenticating users, a user database, either local to the NAC system or a separate system such as RADIUS, or LDAP
- *NAC Single Sign On*—NAC SSO, addresses systems that already perform authentication as part of their normal operation. For example 802.1X, VPN access, or Active Directory. NAC SSO learns the authentication state of clients through RADIUS accounting, or Active Director and therefore doesn't require the user to reenter authentication.

Topology Considerations

The Layer-2 OOB solution relies upon there being a Layer-2 network connection available between the client devices and the Cisco CAS. In Figure 58 a trunk connects the access switch to the core/distribution switch. The Cisco CAS is connected to the core/distribution switch through two interfaces—trusted and untrusted. In such a simple network it is relatively easy to provide a Layer-2 connection between the client and the Cisco CAS, for larger networks Layer-3 OOB may be a better choice.

The roles of the untrusted and trusted interfaces:

- *Untrusted Interface*—The untrusted interface connects the client to the to the Cisco CAS during the NAC Authentication and Authorization, Scanning and Evaluation, Remediation, and Quarantine modes
- *Trusted Interface*—The trusted interface connects the NAC CAS to the "normal" network interface. This makes a network connection available to the CAS while it is sitting between the client and the network, thus allowing client access to services such as DHCP and DNS -and user defined services. Once a client has successfully completed its authentication and scanning phases, the CAM uses SNMP to change the client VLAN, on the access switch, from the untrusted VLAN to the trusted VLAN. Thus providing a direct connection to the network that was on the other side of the CAS (the trusted network).

Availability Considerations

Both the CAS and CAM are both highly involved in client network access, and consideration must be given to the impact on clients if either a CAS or CAM should fail or need to be taken out of service for a period of time.

The CAS is inline with client devices during the authentication, authorization, and posture assessment phases of NAC, and if "In Band NAC" is being used it may be inline at all times. A CAS outage in an OOB deployment would not impact already connected clients but would prevent network access for new clients. A CAS outage for "In Line" clients prevents access for all clients.

In situations where availability of a CAS is critical an HA CAS solution may be implemented where a pair of CAS servers are installed using a primary CAS, and a secondary in hot standby. For more information refer to the *Cisco NAC Appliance - Clean Access Server Installation and Configuration Guide*.

The CAM is also a critical part of the authentication, authorization, and posture assessment phases of NAC. Although the CAM doesn't pass client traffic, the impact of its availability needs to be considered in the network design as well. Like the CAS the CAM has a HA solution that provides for a primary server and a hot standby secondary server. In addition, each CAS may be configured with a "fallback" option (as shown in Figure 59) that defines how it will manage client traffic in a situation where the CAM is unavailable.

In both HA CAM, and HA CAS, HA licenses are use that address the HA role of the server.

The use of the HA features will be dependent upon the company unique requirements, but CAS fallback should always configured to ensure that critical network services are available in even of a network outage.

Figure 59 CAS Fallback

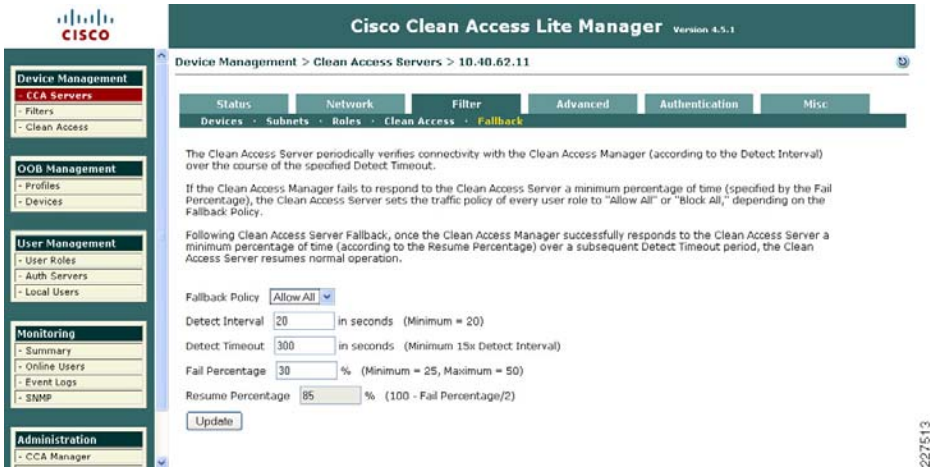

Basic Clean Access Switch Configuration

For OOB-based Clean Access some simple configuration must be performed on the switches implementing NAC This configuration is primarily to enable SNMP communication between the switches and the CAM. The table below shows a simple SNMP v1 configuration (SNMPv2c and SNMPv3 are supported).

In addition to the switch SNMP configuration, the required trusted and untrusted VLANs must exist and be operational on the switch, as illustrated in the configuration shown in Table 4. If a switch has more than one IP address the snmp-server source interface must be specified, as the CAM must be configured with the source IP address that OOB SNMP messages will originate from, alternatively all IP addresses of interfaces on the switch can be added to the CAM. If SNMP access filtering is applied on the switch (as recommended as a best practice) the CAM must be added as a trusted address.

Basic Clean Access Out-of-Band Switch Configuration

Switch Port Configuration	Global Switch Configuration

`snmp trap mac-notification change added`	`snmp-server enable traps mac-notification`
	` snmp-server enable traps snmp linkup linkdown`
	` mac-address-table aging-time 3600`
	` snmp-server host 172.16.1.61 traps version 1 cam_v1 udp-port 162 mac-notification snmp`

NAC CAS Connection

The NAC CAS connects to the core/distribution switch using two switch ports (which are not configured in EtherChannel). One is the untrusted port that serves the VLANs used by the clients prior to their authentication and authorization. The second is the trusted port that connects to the VLANs used once clients have successfully completed the NAC process. Both trusted and untrusted ports are required, even if OOB NAC is used, as the CAS requires access to the trusted VLANs during the NAC process. The following is an example NAC CAS connection configuration.:

```
interface GigabitEthernet1/0/4
description NAC Trusted Eth0
switchport trunk encapsulation dot1q
switchport trunk allowed vlan 48,57,62
switchport mode trunk
spanning-tree portfast trunk
!
interface GigabitEthernet1/0/8
description NAC Untrusted Eth1
switchport trunk encapsulation dot1q
switchport trunk allowed vlan 61,248,257
switchport mode trunk
spanning-tree portfast trunk
```

Basic 802.1X Switch Configuration

The basic 802.1X configuration controls access to an access VLAN depending upon the success or failure of the 802.1X authentication. If the 802.1X authentication is successful, there are three basic options:

- Access to the VLAN configured on the switch port
- Access to the VLAN configured on the switch port an controlled by a access list downloaded from the AAA server
- Access to a VLAN passed to the switch by the AAA server

The following is an example 802.1X configurations.

Example 3750 802.1X PC Port Configuration	Example 3750 Global Configuration
authentication port-control auto authentication periodic dot1x pae authenticator	aaa new-model aaa authentication dot1x default group radius dot1x system-auth-control ip radius source-interface Vlan300 radius-server host 10.40.62.9 auth-port 1812 acct-port 1813 key *cisco* radius-server host 10.40.94.9 auth-port 1812 acct-port 1813 key *cisco*

For more information about the 3750 802.1X configuration, refer to the following documents:

- *Catalyst 3750-E and 3560-E Switch Software Configuration Guide*, 12.2(50)SE ->Configuring IEEE 802.1x Port-Based Authentication

  http://www.cisco.com/en/US/docs/switches/lan/catalyst3750e_3560e/software/release/12.2_50_se/configuration/guide/sw8021x.html
- *Catalyst 2960 Switch Software Configuration Guide*, Rel. 12.2(50)SE Configuring IEEE 802.1x Port-Based Authentication

  http://www.cisco.com/en/US/docs/switches/lan/catalyst2960/software/release/12.2_50_se/configuration/guide/sw8021x.htm