

Security Design


The Cisco Service Ready Architecture (SRA) for Schools solution is designed with security to provide a safe online environment for teaching and learning. Following the proven guidelines of the Cisco SAFE Security architecture, a series of security technologies and products are strategically deployed throughout the solution to protect minors from harmful and inappropriate content, to guarantee the confidentiality of student, staff and faculty private data, and to ensure the availability and integrity of the systems and data.

Protecting the infrastructure and keeping students and staff safe requires the implementation of security controls capable of mitigating both well-known and new forms of threats. Common threats to school environments include:

- *Service disruption*—Disruption of the administrative infrastructure and learning resources such as computer labs caused by botnets, worms, malware, adware, spyware, viruses, DoS attacks.
- *Harmful or inappropriate content*—Pornography, adult, aggressive, offensive and other type of content that could put the physical and psychological well being of minors at risk.
- *Network abuse*—Peer-to-peer file sharing and instant messaging abuse, use of non-approved applications by students, staff, and faculty.
- *Unauthorized access*—Intrusions, unauthorized users, escalation of privileges, and unauthorized access to learning and administrative resources.
- *Data loss*—Theft or leakage of student, staff and faculty private data from servers, endpoints, and while in transit, or as a result of spyware, malware, key-loggers, viruses, etc.

The solution design follows a defense-in-depth approach, whereby multiple layers of protection are built into the architecture. Different security products and technologies are combined together for enhanced security visibility and control. [Figure 1](#) illustrates the security design and its product positioning.

Figure 1 SRA Security Design



The security design 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 students and staff from harmful and inappropriate content. Enforcing E-mail and web browsing policies.
- *Network Access Security and Control*—Securing the access edges. Enforcing authentication and role-based access for students, staff and faculty residing at school sites and district office. Ensuring systems are up-to-date and in compliance with the school's network security policies.
- *Network Endpoint Protection*—Protecting students, staff and faculty from harmful and inappropriate content. Enforcing E-mail and web browsing policies.

The design guidelines and best practices for each focus area are discussed next. For more detailed information, refer to the *Cisco SAFE Reference Guide* at the following URL: http://www.cisco.com/en/US/docs/solutions/Enterprise/Security/SAFE_RG/SAFE_rg.html

Network Foundation Protection

School 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 SRA solution 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 TACACS+ or RADIUS to authenticate access.
 - Authorize actions and log all administrative access.
 - Display legal notification banners.
 - Ensure confidentiality by using secure protocols like SSH and HTTPS.
 - Enforce idle and session timeouts.
 - Disable unused access lines.
- Routing infrastructure
 - Restrict routing protocol membership by enabling 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 NTP time synchronization.
 - Collect system status and event information with SNMP, Syslog, TACACS+/RADIUS accounting.
 - Monitor CPU and memory usage on critical systems.
- 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 VLANs to reduce the size of broadcast domains.
- Protect the Spanning Tree Protocol (STP) domain with BPDU Guard, 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.
- Enable Traffic Storm Control.
- Implement Catalyst Infrastructure Security Features (CISF) including port security, Dynamic ARP Inspection, and 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 school network architecture.
 - 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 local at the District Office.
 - Secure the OOB by enforcing access controls, using dedicated management interfaces or VRFs.
 - Provide secure in-band management access for systems residing at the school sites by deploying firewalls and ACLs to enforce access controls, using Network Address Translation (NAT) to hide management addresses, and use 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.


Configurations are shown in the design chapters. For more detailed information on the NFP best practices, refer to “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 school architecture assumes the existence of a centralized Internet connection at the district office, serving students, staff and faculty residing at all school premises. Common services typically provided include E-mail for staff and faculty, Internet browsing for everyone, and a school web portal accessible over the Internet. 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 to students, staff, and faculty, and to provide public services without compromising the confidentiality, integrity, and availability of school resources and data. To that end, the Internet perimeter incorporates the following security functions:

- **Internet Border Router**—This is the Internet gateway responsible for routing traffic between the school and the Internet. The Internet border router may be administered by school 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 using the Network Foundation Protection (NFP) best practices.
- **Internet Firewall**—A Cisco Adaptive Security Appliance provides Stateful access control and deep packet inspection to protect the school resources and data from unauthorized access and disclosure. The security appliance is configured to prevent incoming access from the Internet, to protect the school web portal and other Internet public services, and to control student, staff and faculty traffic bound to the Internet. The security appliance may also implement an Advanced Inspection and Prevention Security Services Module (AIP SSM) for enhanced threat detection and mitigation. This IPS module may be configured either in inline or promiscuous mode. The security appliance may also provide secure remote access to faculty, staff, and students in the form of IPSec or SSL VPN.
- **Public Services DMZ**—The Internet school web portal, 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 school'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 (i.e., Cisco Security Agent) 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 content that may be harmful or inappropriate for minors or that are known to be the source of spyware, botnets, or other type of malware. The WSA is also responsible for the monitoring of Layer-4 traffic for rogue activity and infected systems.

The following subsections describe the design guidelines for implementing the security functions.

Note For implementation details on Remote Access VPN, IPS, CSA, and Internet border router hardening, refer to the *Cisco SAFE Reference Guide* at the following URL: http://www.cisco.com/en/US/docs/solutions/Enterprise/Security/SAFE_RG/SAFE_rg.html. Firewall and web security configurations can be found in the “School District Office Design” chapter of the *Schools SRA Design Guide*.

Internet Border Router Guidelines

The Internet border router provides connectivity to the Internet via 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, and SNMP may be used to gain visibility on traffic flows, network activity and system status. In addition, the Internet border router should be secured. This includes restricting and controlling administrative access, protecting the management and control planes, and securing the dynamic exchange of routing information.

The “School District Office Design” chapter of the *Schools SRA Design Guide* provides an example of Internet edge ACL. For more information on how to configure the Internet border router, refer to “Internet Edge” chapter of the *Cisco SAFE Reference Guide* at the following URL:

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

Internet Firewall Guidelines


The Cisco Adaptive Security Appliance (Cisco ASA) deployed at the Internet perimeter is responsible for protecting the school's 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 HTTP/HTTPS access for students, staff and faculty residing at any of the school premises.
- Allow outbound Internet SSL access for administrative updates, SensorBase, IPS signature updates, etc.
- Allow students, staff and faculty access to DMZ services such as school web portal, 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 Mail Transfer Agent, DNS to DNS server, etc.).
- Restrict connections initiated from DMZ to the only necessary protocols and sources (DNS from DNS server and mail server, SMTP from mail server, SSL for 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.

In addition, the Cisco ASA should be hardened following the NFP 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 school design, higher 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 “School District Office Design” chapter of the *Schools SRA Design Guide*.

E-mail Security Guidelines

The Cisco Ironport C Series E-mail Security Appliance (ESA) deployed at the DMZ is responsible for 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 4. 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 IPS appliances 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 school 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 inappropriate transport of sensitive information.

Figure 4 E-mail Delivery Chain





Figure 4 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.


Figure 5 Common ESA Deployments

There are multiple deployment approaches for the security appliance depending on the number of interfaces used (see [Figure 5](#)):

- *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 school architecture implements the ESA with a single interface. In addition, using a single interface leaves other data interfaces available for redundancy.

[Figure 6](#) illustrates the logical location of the ESA within the E-mail flow chain.

Figure 6 Typical Data Flow for Inbound E-mail Traffic

The following steps explain what is taking place in [Figure 6](#):

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 Internet firewall should be configured to allow communications to and from the Cisco IronPort ESA. 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

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


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 in 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 7), 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 7 Cisco IronPort ESA NIC Pairing



For more information on how to configure ESA, refer to the *Cisco SAFE Reference Guide* (http://www.cisco.com/en/US/docs/solutions/Enterprise/Security/SAFE_RG/SAFE_rg.html) and *IronPort ESA User's Guide* (<http://www.ironport.com/support/>).

Web Security Guidelines


The school architecture implements a Cisco IronPort S Series Web Security Appliance (WSA) to block access to sites with content that may be harmful or inappropriate for minors, and to protect the schools 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 school design assumes a centralized Internet connection, the WSA is implemented at the distribution layer of the district office network. This allows the inspection and enforcement of web access policies to all students, staff, and faculty residing at any of the school premises. Logically, the WSA sits in the path between web users and the Internet, as shown in Figure 8.

Figure 8 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) enable 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 school to have control over the configuration of the endpoints, which may not be always possible. For example, the school may allow students, staff and faculty to use personal laptops, smart-phones and other devices outside the school'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 school architecture 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.


Note It is recommended to enable both the Layer-4 traffic monitor and transparent proxy during the initial System Setup Wizard. Either of these services can be disabled or reconfigured after initial setup from the web interface. If you do not enable one of the features in the System Setup Wizard and then need to enable it later, you must run the System Setup Wizard again, losing all configurations added to the appliance.

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

Figure 9 WCCP Redirection

The following steps describe what takes place in [Figure 9](#):

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 block access to a specific application by preventing access to the server or block of IP addresses to which the client application must connect.

Note In the school design, the Cisco ASA is configured to allow only permitted ports (HTTP and HTTPS), so any connection attempts on other ports should be blocked by the firewall.

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 school's 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 is traffic in both directions should flow through the same stateful device. The school architecture 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>).

The following are the key guidelines when deploying the L4 Traffic Monitor:

- *Determine physical connection*—L4TM requires traffic to be directed to the WSA for monitoring. This can be done by connecting a physical network tap, configuring SPAN port mirroring on a switch, or using a hub. Network taps forward packets in hardware, while SPAN port mirroring is generally done in software. On the other hand, SPAN port mirroring can be easily reconfigured, providing further flexibility.
- *Location*—Deploy L4TM in the network where it can see as much traffic as possible before getting out to the Internet through the firewall. It is important that the L4TM be logically connected after the proxy ports and before any device that performs network address translation (NAT) on client IP addresses.
- *Action setting*—The default setting for the L4TM is monitor only. Optionally you may configure the L4TM to monitor and block suspicious traffic. TCP connections are reset with the generating of TCP resets, while UDP sessions are tear down with ICMP unreachable. The use of L4TM blocking requires that the L4TM and the Web Proxy to be placed on the same network so that all clients are accessible on routes that are configured for data traffic.

In the school architecture, L4TM is deployed by setting a SPAN session on the distribution switch to replicate all TCP and UDP traffic on the links connecting to the inside interface of the firewall. Using SPAN provides greater flexibility, and inspecting the firewall's inside links ensures traffic is monitored before NAT and before being sent out the Internet. The L4TM deployment is shown in [Figure 10](#).

Figure 10 L4TM Deployment

L4TM action is set to monitor only. Because the Internet firewall is configured to block any traffic bound to the Internet other than HTTP and HTTPS, there is no additional benefit in using L4TM blocking. If active mitigation is required, consider implementing a Cisco IPS module or appliance in in-line 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. In addition, the Cisco IPS provides multiple configurable response actions including blocking the malicious packet only, blocking the entire session, or blocking any traffic coming from the offending system.

Configuration steps and examples are included in the “School District Office Design” chapter of the *Schools SRA Design Guide*.

Network Access Security and Control

Some of the most vulnerable points of the network are the access edges where students, staff and faculty connect to the network. With the proliferation of wireless networks, increased use of laptops and smart mobile devices, the school administration cannot simply rely on physical controls hoping to prevent unauthorized systems from being plugged into the ports of the access switches. Protection should be rather embedded into the network infrastructure, leveraging the native security features available in switches and routers. 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 help reduce the potential loss of sensitive information by enabling schools 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 Schools SRA achieves access security and control by using the following technologies:

- Catalyst Integrated Security Features (CISF)
- Cisco NAC Appliance
- Cisco Identity-Based Network Networking Services (IBNS)

The 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.

In addition to using the CISF features to secure the access ports, access control solutions like IBNS or NAC appliance may be deployed to control who can access the network and ensure endpoint compliance with the schools security policies.

The Cisco IBNS solution 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 enabling enterprise policy enforcement of all users and hosts, whether managed or unmanaged. In addition to holistic access-control provided by 802.1X, IBNS also offers device-specific access-control through MAC-Authentication Bypass (MAB) and user-based access-control through Web-Auth.

The Cisco Network Admission Control (NAC) Appliance uses the network infrastructure to enforce security policy compliance on all devices seeking to access network computing resources. With NAC Appliance, network administrators can authenticate, authorize, evaluate, and remediate wired, wireless, and remote users and their machines prior to network access. It identifies whether networked devices such as laptops, IP phones, or game consoles are compliant with your network’s security policies and repairs any vulnerability before permitting access to the network. Noncompliant machines are redirected into a quarantine area, where remediation occurs at the discretion of the administrator.

The choice of which access control solution to use depends on the security goals and the direction of the network design. For networks using or moving towards 802.1x-based wired or wireless access and interested in identity-based access control, Cisco IBNS solution should be considered. For networks requiring role-based access control using posture assessments to ensure security compliance, Cisco NAC Appliance should be considered.

The Catalyst Integrated Security Features (CISF), and the Cisco Identity-Based Networking Services (IBNS) and NAC Appliance access control solutions are discussed in the “Schools Access Layer Features” chapter of the *Schools SRA Design Guide*.

Endpoint Protection

Servers, desktop computers, laptops, printers, and IP phones are examples of the diverse network endpoints commonly present in School environments. The great variety in hardware types, operating systems, and applications represents a clear challenge to security. In addition, students and staff may bring laptops and other portable devices that could also be used outside the school’s premises, potentially introducing viruses, worms, spyware, and other type of malware.

Properly securing the endpoints requires not only adoption of the appropriate technical controls but also end-user awareness. While this document focuses on the implementation of the technical controls, the school’s security strategy must include

security awareness campaigns and programs. Students and staff must be continuously educated on current threats, use best practices, and the security measures needed for keeping endpoints up-to-date with the latest updates, patches, and fixes.

Following the best practices of Cisco SAFE, the school architecture implements a range of security controls designed to protect the endpoints. These include host-based Cisco IPS, network-based intrusion prevention systems, and web and E-mail traffic security.

As Cisco host-based IPS, the architecture leverages the Cisco Security Agent (CSA). CSA takes a proactive and preventative approach, using behavior-based security to focus on preventing malicious activity on the host. Malicious activity is detected and blocked, independent of the type of malware, spyware, adware, or virus affecting the host.

Once deployed on an endpoint, whenever an application attempts an operation, the agent checks the operation against the application's security policy—making a real-time allow or deny decision on the continuation of that operation and determining whether logging the operation request is appropriate. Security policies are collections of rules that IT or security administrators assign to protect servers and desktops, either individually or organization-wide. CSA provides defense-in-depth protection against spyware and adware by combining security policies that implement distributed firewall, operating system lockdown and integrity assurance, malicious mobile code protection, and audit-event collection capabilities in default policies for servers and desktops.

CSAs are centrally managed with the CSA Management Center (CSA-MC), which in the Cisco SAFE design is placed in a secure segment in the data center. The Management Center (MC) also provides centralized reporting and global correlation.

For complete details about deploying CSA in a network, refer to the *Rapid Deployment Guide for Cisco Security Agent 6.0 for Desktops* at the following URL:

http://www.cisco.com/en/US/prod/collateral/vpndevc/ps5707/ps5057/deployment_guide_c07-501928.htm

Internet Perimeter Security

Internet Border Router

Whether the Internet border router is managed by the school or the ISP, it must be hardened following the best practices listed in the Network Foundation Protection section. 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 school is assigned the 198.133.219.0/24 address block, 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 school's 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 School'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 here provided are reserved for the exclusive use of Cisco Systems, Inc.

Internet Firewall

The mission of the Internet firewall is to protect the school'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 Schools Service Ready architecture uses a Cisco ASA appliance as illustrated in [Figure 11](#).

Figure 11 Internet Edge Firewall

- The Cisco ASA is implemented with three interface groups, each one representing a distinct security domain:
- *Inside*—The Inside is the interface connecting to the core/distribution switch that faces the interior of the network where internal users and resources reside.
- *Outside*—Interface connecting to the Internet border router. The router may be managed either by the school or a service provider.
- *Demilitarized Zone (DMZ)*—The DMZ hosts school services that are accessible over the Internet. These services may include a web portal and E-mail services.

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
- 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 the 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 OBB 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 core/distribution 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 school typically gets a limited number of public IP addresses. In addition, NAT helps shield the school'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 school 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 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 school environments, the school administration may opt to override the default rules with more stringent rules indicating exactly what ports and protocols are permitted.

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 student, staff and faculty residing at all school premises 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 students, staff, and faculty access to DMZ services such as school web portal, 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 to the only 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 school 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
```

Firewall Redundancy

The Internet perimeter of the school architecture 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 school 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.html

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 12](#) illustrates the concept of redundant interface.

Figure 12 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 12](#).

```
! 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 school architecture, 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 13](#). 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 13 Cisco ASA OSPF



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 14](#). This option does not require the configuration of the Internet border router.

Figure 14 Cisco ASA Static Route



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.

Web Security

The Schools Service Ready Architecture implements a Cisco IronPort WSA at the core/distribution layer of the district office, as illustrated in Figure 15. 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 15 WSA Deployment



Following subsections describe 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 is critical to note that some of the initial settings cannot be changed afterwards without resetting the appliance's configuration to its factory defaults. Therefore, care should be taken in choosing the right configuration options. Plan not only for the features to be implemented immediately, but also for what that might be required in the future.

The following are some guidelines when running the System Setup Wizard:

- **Deployment Options**—Step 2 of the wizard gives the user the options to enable only L4 Traffic Monitoring, enable only Secure Web Proxy, or enable both functions. Select enable both Secure Web Proxy and L4 Traffic Monitor if you plan to use both functions.
- **Proxy Mode**—If the Secure Web Proxy function has been enabled, Step 2 of the wizard requires the user to choose between Forward and Transparent mode. It should be noted that a WSA appliance initially configured in Transparent mode can

still be configured as a Forward Web Proxy, per contrary, the Transparent Web Proxy function is not available if the appliance is configured in Forward mode. Therefore, select Forward mode only if you certain that the Transparent mode will never be required.

Note The deployment and proxy mode options cannot be changed after the initial configuration without resetting the WSA appliance to its factory defaults. Plan your configuration carefully.

Interface and Network Configuration

The following need to be configured 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)

These settings are configured as part of an initial setup using the System Setup Wizard, but can be later modified by using the WSA Web-based GUI.

Configuring Network Interfaces

Independently from the model, all Cisco IronPort WSA appliances are equipped with six Ethernet interfaces as shown in Figure 16.

Figure 16 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 school architecture, interface M1 connects to the out-of-band management network. Interface M1 can optionally be used to handle data traffic in case the school does not have an out-of-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 school architecture. P1 connects to the inside subnet of the firewall.
- **L4 Traffic Monitor (L4TM)**—T1 and T2 are the L4TM interfaces. The school design uses only the T1 interfaces. The T1 interface connects to a core/distribution switch port configured as the destination of the SPAN session used to capture traffic bound to the Internet.

Figure 17 illustrates the network topology around the WSA used in the Cisco validation lab.

Figure 17 WSA Network Topology




Figure 18 illustrate 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 18 WSA Interface Configuration

Interfaces

Interfaces	Ethernet Port	IP Address	Netmask	Hostname
M1		172.26.191.105	255.255.255.0	ironport.cisco.com
P1		10.125.33.8	255.255.255.0	ironport.cisco.com

Separate Routing for Management Services: Separate routing (M1 port restricted to appliance management services only)

Appliance Management Services: HTTP on port 8080, HTTPS on port 8443, Redirect HTTP request to HTTPS

L4 Traffic Monitor Wiring: Duplex TAP: T1 (In/Out)

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 19.

Figure 19 WSA Route Configuration

Routes

Routes for Management Traffic (Interface M1: 172.26.191.105, Interface P1: 10.125.33.8)

Name	Destination Network	Gateway
Default Route	All Others	172.26.191.1

Routes for Data Traffic (Interface P1: 10.125.33.8)

Name	Destination Network	Gateway
Default Route	All Others (Including External)	10.125.33.10
Internal-10	10.0.0.0/8	10.125.33.9

Configuring DNS

The initial setup requires the configuration of a host name for the WSA appliance, and listing the DNS servers. Figure 20 shows the DNS configuration.

Figure 20 WSA DNS Configuration

DNS

DNS Server Settings

DNS Servers: Use these DNS Servers:

Priority	IP Address
0	64.102.6.247

Interface for DNS traffic: Auto

Wait Before Timing out Reverse DNS Lookups: 20 seconds

DNS Domain Search List: None

Time Settings

Time synchronization is critical for forensic analysis and troubleshooting, therefore enabling NTP is highly recommended. Figure 21 shows how the WSA is configured to synchronize its clock with an NTP server located on the OOB management network.

Figure 21 WSA NTP Configuration

Time Settings

Time Keeping Method: Using NTP Servers:

Priority	IP Address
1	172.26.129.252

Interface for NTP Server Queries: Management (172.26.191.105/24: ironport.cisco.com)

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 22 and Figure 23.

Figure 22 WSA Upgrade Settings**Upgrade Settings**

Upgrade Settings	
The following settings are used when running a System Upgrade.	
Server:	http://downloads.ironport.com/asyncoos/upgrade/ (IronPort Upgrade Server)
Interface:	Management (172.26.191.105)
HTTP Proxy Server:	http://proxy-rtp-1.cisco.com
Edit Upgrade Settings...	

Figure 23 WSA Component Updates**Component Updates**

Update Settings for Security Components	
Update Server:	https://update-manifests.ironport.com
Interface:	Management
Proxy Server:	http://proxy-rtp-1.cisco.com:80
Edit Update Settings...	

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 24 illustrates the Web Proxy settings.

Figure 24 WSA Proxy Settings**Proxy Settings**

Web Proxy Settings	
Basic Settings	
Proxy:	Enabled
HTTP Ports to Proxy:	80, 3128
Caching:	Enabled Clear Cache
IP Spoofing:	Disabled
Advanced Settings	
Reserve Timeouts:	Client Side: 300 Seconds Server Side: 300 Seconds
Persistent Timeouts:	Client Side: 300 Seconds Server Side: 300 Seconds
Simultaneous Persistent Connections:	Server Maximum Number: 2000
Headers:	X-Forwarded-For: Do Not Send VIA: Send
Edit 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 25 shows an example.

Figure 25 WSA Transparent Proxy**Edit WCCP v2 Service**

WCCP v2 Service	
Service Profile Name:	web-https-cache
Service:	<input type="radio"/> Standard service ID: 0 web-cache (destination port 80) <input checked="" type="radio"/> Dynamic service ID: 10 0-255 Port numbers: 80,443 <small>(up to 8 port numbers, separated by commas)</small> <input checked="" type="radio"/> Redirect based on destination port <input type="radio"/> Redirect based on source port (return path) <small>For IP spoofing, define two services, one based on destination port and another based on source port (return path).</small> <input checked="" type="radio"/> Load balance based on server address <input type="radio"/> Load balance based on client address <small>Applies only if more than one Web Security Appliance is in use.</small>
Router IP Addresses:	10.125.33.10 <small>Separate multiple entries with line breaks or commas.</small>
Router Security:	<input checked="" type="checkbox"/> Enable Security for Service Password: <input type="password"/> Confirm Password: <input type="password"/>
<small>Advanced: Optional settings for customizing the behavior of the WCCP v2 Router.</small>	

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 10.0.0/8 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.

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. An example is provided 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 26](#).

Figure 26 WSA HTTPS Proxy

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 27](#) 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 27 WSA Upstream Proxy Group

Edit Upstream Proxy Group

Proxy Address	Port	Reconnection Attempts	
64.102.255.40	80	2	[Add Row]
128.107.241.169	80	2	[Delete]

Load Balancing: Fewest Connections

Failure Handling: Connect directly to destination host

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 28 presents an example where all traffic is directed to the proxies in the Upstream-Lab_proxy group.

Figure 28 WSA Routing Policies

Routing Policies

Order	Members	Routing Destination	Delete
	64.102.255.40:80, 128.107.241.169:80	Upstream-Lab_proxy	[Delete]

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 school can control what Internet applications (instant messaging clients, peer-to-peer file-sharing, web browsers, Internet phone services, etc.) and URL categories students, staff and faculty 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 29 shows the global policy.

Figure 29 Global Access Policy

Access Policies

Order	Group	Applications	URL Categories	Objects	Web Reputation and Anti-Malware Filtering	Delete
	Identity: All	Allow: FTP over HTTP, HTTP	Redirect: 0 Monitor: 52 Block: 1 Allow: 0	Object Max Size: None	(enabled)	[Delete]

URL categories corresponding to content inappropriate for minors should be blocked in compliance with the school’s Internet access policies. Figure 30 provides an example on how the “Adult/Sexually Explicit” category is blocked.

Figure 30 URL Categories

Access Policies: URL Categories: Global Policy

Custom URL Category Filtering: No Custom URL Categories are defined.

Category	Monitor	Block
Adult/Sexually Explicit	[Select all]	[Select all]
Advertisements & Search	[Select all]	[Select all]

Layer-4 Traffic Monitoring (L4TM)

L4TM can be implemented in the school environment to identify rogue traffic across all network ports and detect malware attempts to bypass port 80. Additionally, L4TM is capable of identifying internal clients with malware and that attempt to phone-home across non-standard ports and protocols.

Implementing L4TM requires the following:

1. Configuring L4TM interfaces
2. Configuring WSA L4TM global settings
3. Configuring traffic monitoring

Configuring L4TM Interfaces

The wiring type depends on how traffic is directed to the WSA appliance. Network taps and SPAN can be either configured in simplex or duplex mode. If using a hub, only duplex mode can be used. The wiring type configuration is typically done during the initial setup as described earlier in this chapter. Figure 31 show the wiring options.

Figure 31 L4TM Wiring Type

L4 Traffic Monitor Wiring: Duplex TAP: T1 (In/Out)

Configuring WSA L4TM Global Settings

The ports to be monitored can be specified in the L4TM Global Settings. Options are:

- *All ports*—Monitors all 65535 TCP ports for rogue activity.
- *All ports except proxy ports*—Monitors all TCP ports except HTTP and HTTPS proxy ports.

Note The Cisco ASA in the Internet perimeter is configured to allow only permitted ports, so any connection attempts on rogue ports should be blocked by the firewall.

Figure 32 shows the options.

Figure 32 L4TM Global Settings

Edit L4 Traffic Monitor Global Settings



Configuring Traffic Monitoring

While a hub or a network tap could be used, using SPAN port mirroring provides the greatest flexibility. SPAN allows the monitoring of port traffic based on VLANs or source interfaces and it can easily be reconfigured.

The following is a configuration example of SPAN to replicate traffic:

```
! Enables port mirroring on the switch ports connecting to the firewall
inside interfaces
monitor session 10 source interface Gi4/4
monitor session 10 source interface Gi5/3
!
! Sets the interface connecting to the WSA as the destination
monitor session 10 destination interface Gi6/3
```

The SPAN configuration can be seen on the switch with the **show monitor session** command.

Activity monitored by the L4TM feature can be seen in the L4 Traffic Monitor page of the WSA web-based GUI. Figure 33 shows client activity with a website known to be the source of malware.

Figure 33 L4 Traffic Monitor

