



CHAPTER 6

Community College Security Design Considerations

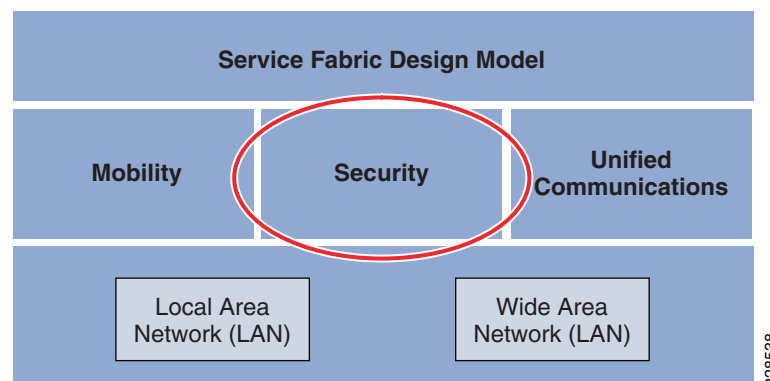
As community colleges embrace new communication and collaboration tools, transitioning from traditional classroom teaching into an Internet-based, media-rich education and learning environment, a whole new set of network security challenges arise. Community college network infrastructures must be adequately secured to protect students, staff, and faculty from harmful content, to guarantee confidentiality of private data, and to ensure the availability and integrity of the systems and data. Providing a safe and secure network environment is a top responsibility for community college administrators and community leaders.

Security Design

Within the Cisco Community College reference design, the service fabric network provides the foundation on which all solutions and services are built to solve the business challenges facing community colleges. These business challenges include building a virtual learning environment, providing secure connected classrooms, ensuring safety and security, and operational efficiencies.

The service fabric consists of four distinct components: LAN/WAN, security, mobility, and unified communications, as shown in [Figure 6-1](#).

Figure 6-1 **Service Fabric Design Model**



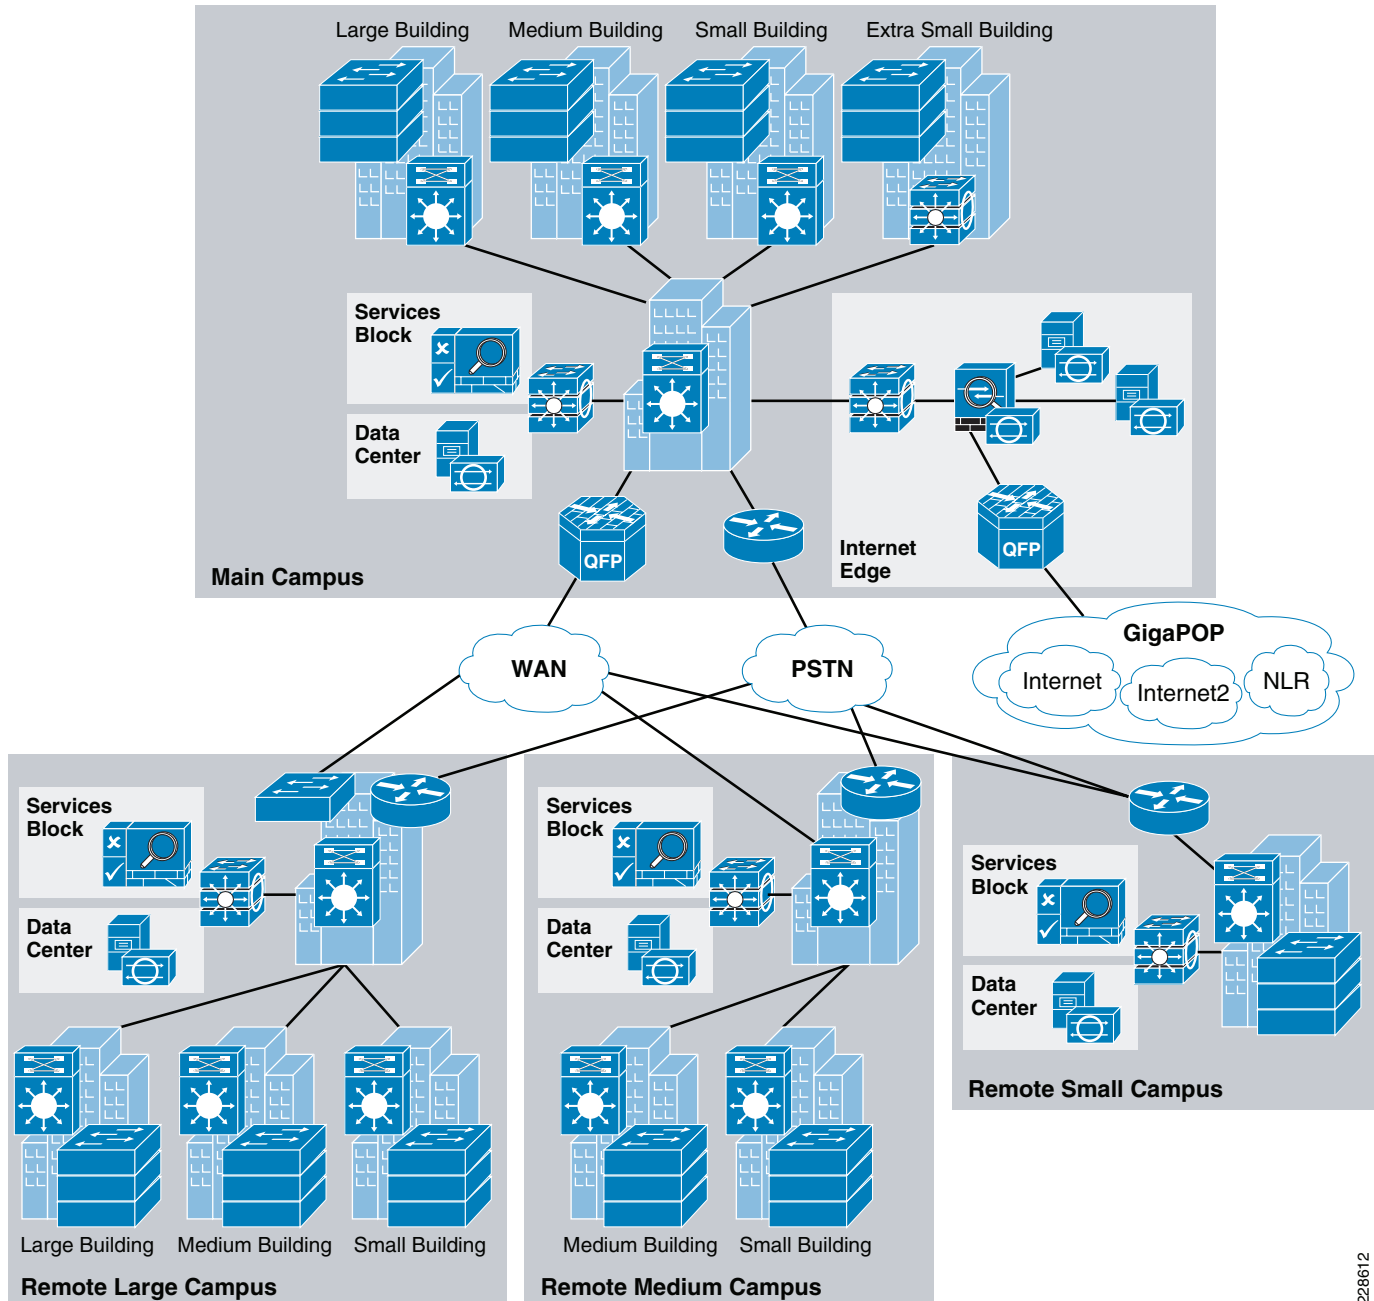
The Community College reference design includes security to protect the infrastructure and its services to provide a safe and secure online environment for teaching and learning. This design leverages the proven design and deployment guidelines of the Cisco SAFE Security Reference Architecture to secure

the service fabric by deploying security technologies throughout the entire solution to protect students and faculty from harmful 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 its services requires implementation of security controls capable of mitigating both well-known and new forms of threats. Common threats to community college environments include the following:

- *Service disruption*—Disruption to the infrastructure and learning resources such as computer labs 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 students, faculty, and staff; peer-to-peer file sharing and instant messaging abuse; and access to forbidden content
- *Unauthorized access*—Intrusions, unauthorized users, escalation of privileges, IP spoofing, and unauthorized access to restricted learning and administrative resources
- *Data loss*—Loss or leakage of student, staff, and faculty private data from servers and user endpoints
- *Identity theft and fraud*—Theft of student, staff, and faculty identity or fraud on servers and end users through phishing and E-mail spam

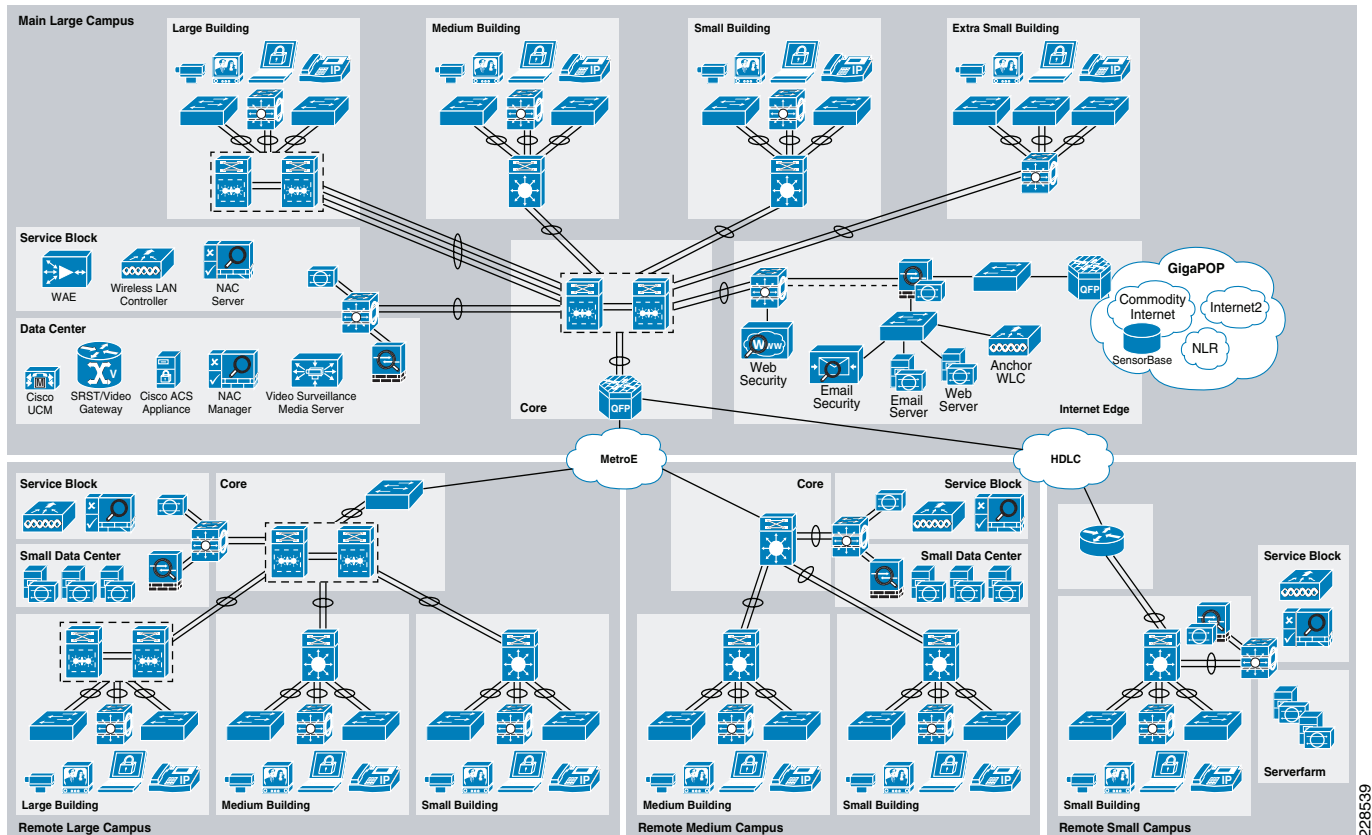
The Community College reference design accommodates a main campus and one or more remote smaller campuses interconnected over a metro Ethernet or managed WAN service. Each of these campuses may contain one or more buildings of varying sizes, as shown in [Figure 6-2](#).

Figure 6-2 Community College Reference Design Overview

228612

Operating on top of this network are all the services used within the community college environment, such as safety and security systems, voice communications, video surveillance equipment, and so on. The core of these services are deployed and managed at the main campus building, allowing each remote campus to reduce the need for separate services to be operated and maintained by community college IT personnel. These centralized systems and applications are served by a data center in the main campus.

The security design uses a defense-in-depth approach where multiple layers of security protection are integrated into the architecture. Various security products and technologies are combined to provide enhanced security visibility and control, as shown in [Figure 6-3](#).

Figure 6-3 Community College Network Security Design Overview

The following security elements should be included in the Community College Security design depicted in Figure 6-3:

- *Endpoint Security*—Desktop endpoint protection for day-zero attack protection, data loss prevention, and signature-based antivirus.
- *Network Foundation Protection*—Device hardening, control, and management plane protection throughout the entire infrastructure to maximize availability and resiliency.
- *Catalyst Integrated Security Features*—Access layer protection provided by port security, Dynamic ARP inspection, IP Source Guard, and DHCP Snooping.
- *Threat Detection and Mitigation*—Intrusion prevention and infrastructure based network telemetry to identify and mitigate threats.
- *Internet Access*—E-mail and Web Security. Stateful firewall inspection. Intrusion prevention and global correlation. Granular access control.
- *Cisco Video Surveillance*—Monitor activities throughout the school environment to prevent and deter safety incidents.
- *Enhanced Availability and Resiliency*—Hardened devices and high availability design ensure optimal service availability. System and interface-based redundancy.
- *Unified Communications*—Security and emergency services, enhanced 911 support. Conferencing and collaboration for planning and emergency response.
- *Network Access Control*—Authentication and policy enforcement via Cisco Identity-Based Networking Services (IBNS). Role-Based access control and device security compliance via Cisco Network Admission Control (NAC) Appliance.

The Community College reference design recognizes that cost and limited resources are common limiting factors. Therefore, architecture topologies and platforms are carefully selected to increase productivity while minimizing the overall cost and operational complexities. In certain cases, tradeoffs are made to simplify operations and reduce costs where needed.

The security design for the community college service fabric focuses on the following key areas.

- *Network foundation protection (NFP)*—Ensuring the availability and integrity of the network infrastructure by protecting the control and management planes to prevent service disruptions, network abuse, unauthorized access, and data loss.
- *Internet perimeter protection*
 - Ensuring safe connectivity to the Internet, Internet2, and National LambdaRail (NLR) networks
 - Protecting internal resources and users from botnets, malware, viruses, and other malicious software
 - Protecting students, staff, and faculty from harmful content
 - Enforcing E-mail and web browsing policies to prevent identity theft and fraud
 - Blocking command and control traffic from infected internal bots to external hosts
- *Data center protection*
 - Ensuring the availability and integrity of centralized applications and systems
 - Protecting the confidentiality and privacy of student, staff, and faculty records
- *Network access security and control*
 - Securing the access edges
 - Enforcing authentication and role-based access for students, staff, and faculty residing at the main and remote campuses
 - Ensuring that systems are up-to-date and in compliance with the community college's network security policies
- *Network endpoint protection*
 - Protecting servers and school-controlled systems (computer labs, school-provided laptops, and so on) from viruses, malware, botnets, and other malicious software
 - Enforcing E-mail and web browsing policies for staff and faculty

Together, these key security areas create a defense-in-depth solution for protecting community colleges from common security threats such as service disruption, network abuse, unauthorized access, data loss, and identity theft and fraud. The design guidelines and best practices for each of the above security focus areas are detailed in the following sections. For more detailed information on each of these areas, see the *Cisco SAFE Reference Guide* at the following URL: <http://www.cisco.com/go/safe>.

Network Foundation Protection

The community college network is built with routers, switches, and other infrastructure network devices that keep the applications and services running. These infrastructure devices must be properly hardened and secured to maintain continued operation and access to these services.

To ensure the availability of the community college network infrastructure, the NFP best practices should be implemented for the following areas:

- *Infrastructure device access*

- Restrict management device access to authorized parties and via only 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 such as 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, and 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 (PVST) 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 Cisco Catalyst Infrastructure Security Features (CISF) including port security, Dynamic ARP Inspection, DHCP snooping, and IP Source Guard.
- 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 community college network infrastructure.
 - Authenticate, authorize, and keep records of all administrative access.
 - If possible, implement a separate out-of-band (OOB) management network (hardware- or VLAN-based) to manage systems local to the main campus.
 - Secure the OOB management access 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 remote campus sites by deploying firewalls and ACLs to enforce access controls, using Network Address Translation (NAT) to hide management addresses, and using secure protocols such as SSH and HTTPS.
 - Ensure time synchronization by using NTP.
 - Secure management servers and endpoints with endpoint protection software and operating system (OS) hardening best practices.

For more detailed information on the NFP best practices including configuration examples, see “Chapter 2, Network Foundation Protection” in 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 Community College reference design assumes a centralized connection to the Internet, Internet2, and National LambdaRail (NLR) networks at the main campus site. This connection serves students, staff, and faculty at the main campus as well as all remote campus sites. Common services typically provided by this connection include the following:

- E-mail for staff and faculty
- Internet browsing for everyone
- Community college web portal accessible over the Internet
- Connectivity to other educational institutions over the Internet2 and NLR network
- Remote access to the community college network

The Internet2 network is a not-for-profit advanced networking consortium comprised of more than 200 U.S. universities in cooperation with 70 leading corporations, 45 government agencies, laboratories, and other institutions of higher learning as well as over 50 international partner organizations. Internet2 provides its members both leading-edge network capabilities and unique partnership opportunities that together facilitate the development, deployment, and use of revolutionary Internet technologies. The physical implementation of Internet2 network consists of an advanced IP network, virtual circuit network, and core optical network. The Internet2 network provides the necessary scalability for member institutions to efficiently provision resources to address the bandwidth-intensive requirements of their

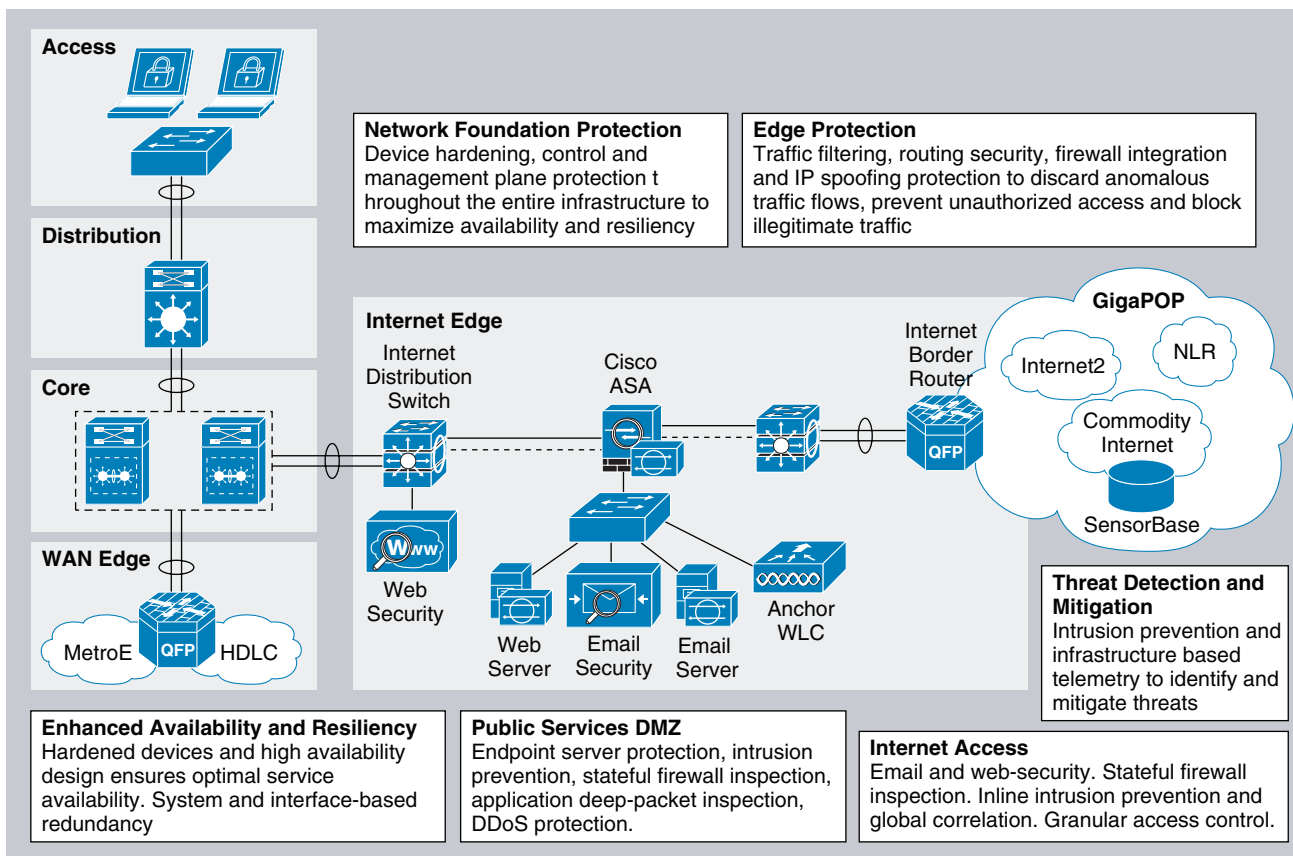
campuses, such as collaborative applications, distributed research experiments, grid-based data analysis, and social networking. For more information on the Internet2 network, see the following URL: <http://www.internet2.edu/network/>.

The National LambdaRail (NLR) network is a high-speed fiber optic network infrastructure linking over 30 cities in 21 states. It is owned by the U.S. research and education community and is dedicated to serving the needs of researchers and research groups. The NLR high-performance network backbone offers unrestricted usage and bandwidth, a choice of cutting-edge network services and applications, and customized service for individual researchers and projects. NLR services include high-capacity 10 Gigabit Ethernet LAN-PHY or OC-192 lambdas, point-to-point or multi-point Ethernet-based transport, routed IP-based services, and telepresence video conferencing services. For more information on the NLR network and its services, see the following URL: <http://www.nlr.net/>.

Community colleges typically connect to a local Gigabit point-of-presence (GigaPOP) or regional network service provider to gain access to the Internet, Internet2, and NLR networks. The same security controls are applicable regardless of whether they connect to a GigaPOP or regional network. For details on how community colleges connect to these networks, see Chapter 4, “Community College WAN Design Considerations.”

The part of the network infrastructure that provides connectivity to the Internet, Internet2, and NLR is defined as the Internet perimeter, as shown in Figure 6-4.

Figure 6-4 Internet Perimeter



The Internet perimeter provides safe and secure access to the Internet, Internet2, and NLR networks for students, staff, and faculty. It also provides access to public services such as the community college web portal without compromising the confidentiality, integrity, and availability of the resources and data of the educational institution. To provide secure access, the Internet perimeter should incorporate the following security functions:

- *Internet border router*—The Internet border router is the gateway responsible for routing traffic between the community college and the Internet, Internet2, and NLR networks. It may be administered by the community college IT staff or may be managed by the Internet, Internet2, or NLR service provider. This router provides the first line of protection against external threats and should be hardened according to the NFP best practices.
- *Internet firewall*—A Cisco Adaptive Security Appliance (ASA) provides stateful access control and deep packet inspection to protect community college resources and data from unauthorized access and disclosure. The ASA monitors network ports for rogue activity and detects and blocks traffic from infected internal endpoints, sending command and control traffic back to a host on the Internet. The ASA is configured to control or prevent incoming and outgoing access for the Internet, Internet2, and NLR networks; to protect the community college web portal and other Internet public services; and to control student, staff, and faculty traffic bound towards the Internet. The security appliance may also provide secure remote access to faculty, staff, and students in the form of a Secure Socket Layer (SSL) or IPsec virtual private network (VPN).
- *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 configured either in inline or promiscuous mode. IPS may also be configured to help block certain Internet applications such as AOL Messenger, BitTorrent, Skype, and so on.
- *Public services DMZ*—The community college external Internet web portal, mail server, and other public facing servers and services are placed on a demilitarized zone (DMZ) for security and control purposes. The DMZ acts as a middle stage between the Internet and community college 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 in the DMZ, and controls outbound access from DMZ resources to the Internet. Systems residing within the DMZ should be hardened with endpoint protection software (such as Cisco Security Agent) and OS hardening best practices.
- *E-mail security*—A Cisco IronPort C Series E-Mail Security Appliance (ESA) is deployed in 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. The WSA enforces URL filtering policies to block access to websites containing content that may be harmful for students, staff, and faculty such as sites known to be sources of spyware, adware, botnets, or other types of malware. The WSA may also be configured to block certain Internet applications such as AOL Messenger, BitTorrent, Skype, and so on, and for monitoring Layer 4 traffic for rogue activity and infected systems.
- *Guest access wireless LAN controller*—The Cisco Unified Wireless LAN Guest Access option offers a flexible, easy-to-implement method for deploying wireless guest access via Ethernet over IP (RFC3378). Ethernet over IP (EoIP) tunneling is used between two wireless LAN controller (WLC) endpoints in the centralized network design. A WLC is located in the Internet perimeter DMZ, where it is referred to as an *anchor controller*. The anchor controller is responsible for

terminating EoIP tunnels originating from centralized campus WLCs located in the services block, and interfacing the traffic from these controllers to a firewall or border router. Traffic to and from this guest access WLAN is tunneled to the DMZ transparently, with no visibility by, or interaction with, other traffic in the community college. For more information on the wireless guest access solution, see [Chapter 5, “Community College Mobility Design Considerations.”](#)

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

Internet Border Router Security

The Internet border router connects to a local GigaPOP and provides connectivity to the Internet, Internet2, and NLR networks for the community college. The router acts as the first line of defense against unauthorized access, distributed DoS (DDoS), and other external threats. ACLs, uRPF, and other filtering mechanisms should be implemented for anti-spoofing and to block invalid packets. NetFlow, Syslog, and SNMP should be used to gain visibility on traffic flows, network activity, and system status. In addition, the Internet border router should be hardened and secured following the best practices explained in [Network Foundation Protection, page 6-5](#). This includes restricting and controlling administrative access, protecting the management and control planes, and securing the dynamic exchange of routing information.

For more information on how to secure the Internet border router, see “Chapter 6, Enterprise Internet Edge” in 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

A Cisco ASA firewall should be deployed at the Internet perimeter to protect community college internal resources and data from external threats by doing the following:

- Preventing incoming access from the Internet, Internet2, and NLR networks
- Protecting public resources deployed in the DMZ by restricting incoming access to the public services and by limiting outbound access from DMZ resources out to the Internet
- Controlling user Internet-, Internet2- and NLR-bound traffic
- Monitoring network ports for rogue activity and preventing infected internal endpoints from sending command and control traffic back to a host on the Internet

The ASA should be configured to enforce access policies, keep track of connection status, and inspect packet payloads. Examples of the needed access policies include the following:

- Deny or control any connection attempts originating from the Internet, Internet2, and NLR to internal resources and subnets.
- Allow outbound Internet HTTP/HTTPS access for students, staff, and faculty residing at any of the community college campuses.
- Allow outbound SSL access to the Internet for devices requiring administrative updates such as SensorBase, IPS signature updates, and so on.
- Deny or control access between the community college internal network and the external Internet, Internet2, or NLR networks.

- Allow students, staff, and faculty access to DMZ services such as the community college web portal, E-mail, and domain name resolution (HTTP, HTTPS, Simple Mail Transfer Protocol (SMTP), point-of-presence [POP], Internet Message Access Protocol (IMAP), Domain Name Service [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 servers, and so on).
- Restrict connections initiated from the DMZ to only necessary protocols and sources (DNS from DNS server, SMTP from mail server, HTTP/SSL from Cisco IronPort ESA).
- Enable stateful inspection for the outbound protocols being used to ensure returning traffic is dynamically allowed by the firewall.
- Prevent access to the anchor WLC deployed in the DMZ for guest access except for tunneled traffic coming from the centralized campus WLCs (UDP port 16666 and IP protocol ID 97) and traffic needed to manage it (SNMP, TFTP, HTTP, HTTPS, SSH).
- Implement NAT and Port Address Translation (PAT) to shield the internal address space from the Internet.

In addition, 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 SenderBase 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.

The ASA uses this 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 ASA sends a syslog message reporting the suspicious activity and optionally drops the traffic if the 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 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 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 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 ASA. Cisco also recommends that DNS packet inspection be enabled with Botnet Traffic Filter snooping. When enabled, the ASA uses Botnet Traffic Filter snooping instead of the regular DNS lookup to resolve static blacklist domain names in the following circumstances:

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

When deploying the Internet firewall, it is important to understand the traffic and policy requirements when selecting a firewall. An appropriately sized ASA model should be chosen so that it does not become a bottleneck. The Cisco ASA should also be hardened following the NFP best practices as described in [Network Foundation Protection, page 6-5](#). This includes restricting and controlling administrative access, securing dynamic exchange of routing information with MD5 authentication, and enabling firewall network telemetry with SNMP, Syslog, and NetFlow.

Given budget and resource constraints for community colleges, high availability is achieved by using redundant physical interfaces, which provides a cost-effective solution. As an alternative, a pair of firewall appliances can be deployed in stateful failover using separate boxes at a higher cost.

Intrusion Prevention

IPS is responsible for identifying and blocking anomalous traffic and packets recognized as well-known attacks. An IPS 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 community colleges. 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 community college network.

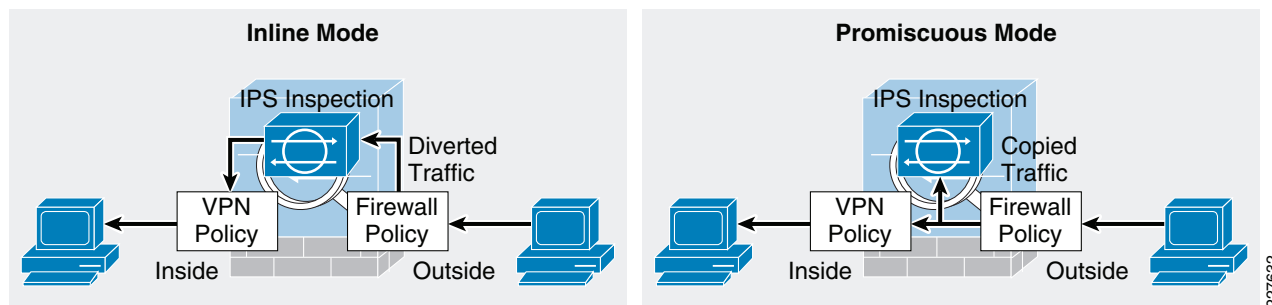
The AIP SSM module may also participate in Cisco Global Correlation for further threat visibility and control. If enabled, the participating IPS sensor receives threat updates from the Cisco SensorBase network at regular intervals. The Cisco SensorBase network contains detailed information about known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

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.

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 6-5](#)). 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 6-5](#) shows the AIP SSM in promiscuous mode.

Figure 6-5 *IPS Inline and Promiscuous Modes*

The recommended IPS deployment mode depends on the goals and policies of the community college. 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 ASA.
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. (Inline mode only) Valid traffic is sent back to the ASA over the backplane; the AIP SSM might block some traffic according to its security policy, and that traffic is not passed on.
6. Remote access VPN policies are applied (if configured).
7. Traffic exits the ASA.

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.

E-Mail Security

Cisco recommends that the Cisco IronPort C Series E-Mail Security Appliance (ESA) be deployed in the DMZ to inspect E-mails and prevent threats such as E-mail spam, viruses, and worms. The ESA acts as a firewall and threat monitoring system for SMTP traffic (TCP port 25). Logically, the ESA acts as a Mail Transfer Agent (MTA) within the E-mail delivery chain, as shown in [Figure 6-6](#).

Figure 6-6 Logical E-Mail Delivery Chain

227423

**Note**

Figure 6-6 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 a “sandwich” configuration.

When the ESA receives the E-mails, they are evaluated using a reputation score mechanism based on the SensorBase network, which is an extensive network that monitors global E-mail and web traffic for anomalies, viruses, malware, and other abnormal behavior. The SensorBase network consists of Cisco IronPort appliances, Cisco ASA, and IPS appliances installed in more than 100,000 organizations worldwide. This provides a large and diverse sample of Internet traffic patterns. By leveraging the information in 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.

In addition, a community college may optionally 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.

There are two options for deploying the ESA appliance, depending on the number of interfaces used:

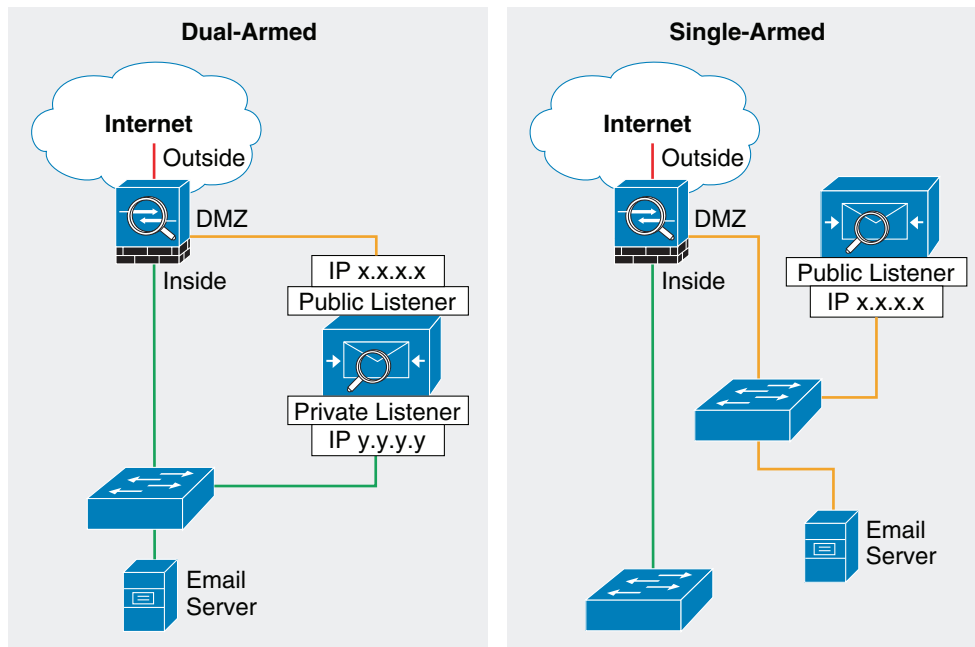
- *Dual-armed configuration*—Two physical interfaces are used to serve as a public mail listener and a private mail listener where each interface is configured with a separate logical IP address. The public listener receives E-mail from the Internet and directs messages to the internal mail servers.

The private listener receives E-mail from the internal servers and directs messages to the Internet. The public listener interface would connect to the DMZ and the private listener interface can connect to the inside of the firewall closer to the mail server.

- *One-armed configuration*—A single interface is configured on the ESA 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.

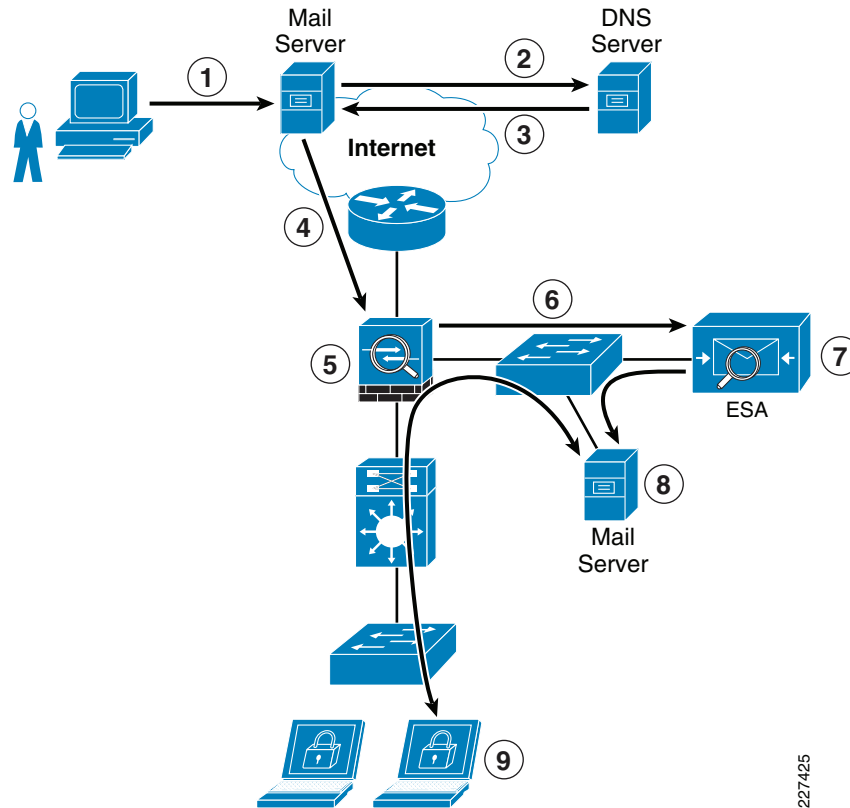
Figure 6-7 shows both configurations.

Figure 6-7 Common ESA Deployments



For simplicity, Cisco recommends that the community college network implement the ESA with a single interface in a single-armed configuration. This also leaves the other data interfaces available for redundancy.

Figure 6-8 shows the logical location of the ESA within the E-mail flow chain and the typical data flow for inbound E-mail traffic.

Figure 6-8 Typical Data Flow for Inbound E-Mail Traffic

The following steps explain what is taking place in [Figure 6-8](#):


-
- Step 1** Sender sends an E-mail to xyz@domain X.
 - Step 2** What's the IP address of domain X?
 - Step 3** It is a.b.c.d (public IP address of ESA).
 - Step 4** E-mail server sends message to a.b.c.d using SMTP.
 - Step 5** Firewall permits incoming SMTP connection to the ESA, and translates its public IP address.
 - Step 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.
 - Step 7** ESA forwards E-mail to preconfigured inbound E-mail server.
 - Step 8** E-mail server stores E-mail for retrieval by receiver.
 - Step 9** Receiver retrieves E-mail from server using POP or IMAP.
-

The ESA appliance functions as an SMTP gateway, also known as a mail exchange (MX). The following outlines some of the deployment design points for the ESA within the community college design:

- The ESA appliance needs to be accessible via the public Internet and is the first hop in the E-mail infrastructure. The IP address of the sender is needed to identify and distinguish the senders in the Mail Flow Monitor to query the SensorBase Reputation Service for the SensorBase Reputation Service Score (SBRs) of the sender. Therefore, a separate MTA should not be deployed at the network perimeter to handle the external connections.
- The ESA needs to be registered in DNS for features such as IronPort Anti-Spam, Virus Outbreak Filters, MacAfee Antivirus, and Sophos Antivirus. A DNS “A” record should be created to map the appliance hostname to its public IP address, and an MX record that maps the public domain to the appliance hostname. A priority is specified for the MX record to advertise the ESA appliance as the primary MTA for the domain.
- A static IP address translation entry on the Internet firewall should be defined to map the public IP address of the ESA to its private internal address.
- All the local domains for which the ESA appliances accept mail need to be added to the recipient access table (RAT). Inbound E-mail destined to domains not listed in the RAT is 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™, Domino™) via SMTP routes.
- For each private listener, the host access table (HAT) must be configured to indicate the hosts that are 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, and on which mail policies can be applied. Policies include the following:
 - Mail flow policies—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 restricting E-mail access based on sender trustworthiness as determined by the IronPort SensorBase Reputation Service.
- SMTP routes are defined to direct E-mail to the appropriate internal mail servers.
- If an OOB management network is available, a separate interface for administration should be used.

Because a failure on the ESA appliance may cause a service outage, a redundant design is recommended. One way to implement redundancy is to use IronPort NIC pairing, as shown in [Figure 6-9](#).

Figure 6-9 Cisco IronPort ESA NIC Pairing



IronPort NIC pairing provides redundancy at the network interface card level by teaming two of the Ethernet interfaces in the ESA appliance. If the primary interface fails, the IP address and MAC address are moved to the secondary interface. IronPort NIC pairing is the most cost-effective solution because it does not require the deployment of multiple ESA appliances and other hardware. However, it does not provide redundancy in case of chassis failure.

Alternative redundant designs include the following:

- *Multiple MTAs*—Adding a second ESA appliance or MTA and using a secondary MX record with an equal cost to load balance between the MTAs.
- *Load balancer*—Using a load balancer such as the Cisco Application Control Engine (ACE) to load balance traffic across multiple ESA appliances.

To accommodate traffic to and from the IronPort ESA provisioned in the DMZ, the Internet firewall needs to be configured to allow this communication. Protocols and ports to be allowed vary depending on the services configured on the ESA.

The following are some of the common services required to be allowed through the Internet firewall:

- 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

For more information on how to configure the 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


The Community College reference design implements a Cisco IronPort S Series Web Security Appliance (WSA) to block HTTP and HTTPS access to sites on the Internet with content that may be harmful, and to protect the community college network from web-based malware and spyware.

The following services may be enabled on the WSA:

- *Web proxy*—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)*—Monitors all Layer 4 traffic for rogue activity, and to detect infected clients.

The community college design assumes a centralized Internet connection implemented at the main campus site. The WSA should be implemented at the distribution layer in the Internet perimeter network. This allows for the inspection and enforcement of web access policies to all students, staff, and faculty located at any of the community college campuses. Logically, the WSA sits in the path between web users and the Internet, as shown in [Figure 6-10](#).

Figure 6-10 Cisco IronPort WSA



There are the following two deployment modes when enabling the Cisco IronPort WSA 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 as its proxy. The web browsers can be configured either manually or by using proxy auto configuration (PAC) files. 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 Auto-discovery 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 to intercept and redirect traffic to the WSA before going to the Internet. Both WCCP and Layer 4 load balancer options provide for redundancy and load balancing.


Explicit forward proxy mode requires administrators to have control over the configuration of the endpoints, which is often not the case in community college environments. For example, community colleges may allow students, guests, or visiting professors to use personal laptops, smart phones, or other devices outside the administration of the institution. Conversely, transparent proxy mode provides transparent integration of WSA without requiring any configuration control over the endpoints. In addition, transparent proxy also eliminates the possibility of users reconfiguring their web browsers to bypass the WSA appliance without the knowledge of the administrators. For these reasons, Cisco recommends that community colleges implement transparent proxy mode with WCCP. In the Community College reference design, the Cisco Catalyst 3750 Stackwise distribution switches deployed in the Internet perimeter can be leveraged as the WCCP server while the WSA acts as a WCCP traffic processing entity.

The Cisco Catalyst 3750 switches support WCCP version 2, which has a built-in failover and load balancing mechanism. Per the WCCPv2 specifications, 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 WSA appliances based on source and destination IP addresses. The WCCP server (Cisco Catalyst 3750 switch) monitors the availability of each appliance in the group and can identify the appliance failures in 30 seconds. After failure, the 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 communication between the WCCP server and the WSA appliances.

Figure 6-11 shows how WCCP redirection works in conjunction with the Cisco Catalyst 3750 StackWise distribution switches.

Figure 6-11 WCCP Redirection



As shown in Figure 6-11, the following steps take place:

-
- Step 1** The client browser requests a connection to `http://website.com`.
 - Step 2** The Cisco Catalyst 3750 Internet perimeter distribution switch intercepts and redirects HTTP/HTTPS requests to WSA via Layer 2 redirection.
 - Step 3** If the content is not present in the local cache, WSA performs a DNS query on the destination site and checks the received IP address against URL and reputation rules, and allows/denies the request accordingly.
 - Step 4** If allowed, WSA fetches the content from the destination website.

Step 5 The content is inspected and then delivered to the requesting client.

**Note**

In the event that the entire service group fails, WCCP automatically bypasses redirection, allowing users to browse the Internet without the web controls. If it is desired to handle a group failure by blocking all traffic, an inbound ACL may be configured on the Cisco ASA inside interface to permit only 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 direct access to the Internet without going through the WSA proxy.

WCCPv2 supports Generic Route Encapsulation (GRE) and Layer 2-based redirection. The Cisco Catalyst 6500 and 3750 switches support Layer 2-based redirection, and the redirection is supported in hardware. Therefore, the WSA must be directly connected to the switch running WCCP. In addition, WCCP is supported only on the ingress of an interface. For these reasons, WSA should connect directly to the Internet perimeter distribution switch using a VLAN that is different than the VLAN from where the client traffic is coming.

**Note**

The Cisco Catalyst 4500 does not provide the ability to create WCCP traffic redirect exception lists, which is an important component of the design. If a Cisco Catalyst 4500 is implemented as the distribution layer switch, another device, such as the Cisco ASA, should be used as the WCCP server.

The following describes some of the design considerations and caveats for implementing a Cisco IronPort WSA with WCCP on a Cisco Catalyst 3750 switch:

- The WSA must be Layer 2-adjacent to the Cisco Catalyst 3750 switch.
- The WSA and switches in the same service group must be in the same subnet directly connected to the switch that has WCCP enabled.
- Configure the switch interfaces that are facing the downstream web clients, the WSA(s), and the web servers as Layer 3 interfaces (routed ports or switch virtual interfaces [SVIs]).
- Use inbound redirection only.
- WCCP is not compatible with VRF-Lite. WCCP does not have visibility into traffic that is being used by the virtual routing tables with VRFs.
- WCCP and policy-based routing (PBR) on the same switch interface are not supported.
- WCCP GRE forwarding method for packet redirection is not supported.
- Use MD5 authentication to protect the communication between the Cisco Catalyst 3750 switches and the WSA(s).
- Use redirect-lists to specifically control what hosts/subnets should be redirected.
- Cisco Catalyst 3750 switches support switching in hardware only at Layer 2; therefore, no counters increment when the command **show ip wccp** is issued on the switch.
- In an existing proxy environment, deploy the WSA downstream from the existing proxy servers (closer to the clients).
- If an OOB management network is available, use a separate interface for WSA administration.

For more information on WCCP in relation to the Cisco Catalyst 3750 switch, see the following URL: http://www.cisco.com/en/US/docs/switches/lan/catalyst3750e_3560e/software/release/12.2_46_se/configuration/guide/swwccp.html.

**Note**

WCCP, firewall, and other stateful features typically require traffic symmetry where traffic in both directions should flow through the same stateful device. Care should be taken when implementing active-active firewall pairs because they may introduce asymmetric paths.

The WSA appliance may also be configured to control or block peer-to-peer file sharing and instant messaging applications such as AOL Messenger, BitTorrent, Skype, Kazaa, and so on. Depending on the port used for transport, the WSA handles these applications as follows:

- **Port 80**—Applications that use HTTP tunneling on port 80 can be handled by enforcing access policies within the web proxy configuration. Applications access can be controlled based on applications, URL categories, and objects. Applications are matched based on their user agent pattern, and the use of regular expressions. URLs can be blocked based on specific categories, such as predefined chat and peer-to-peer categories, or custom categories defined by the administrator. Peer-to-peer access can also be filtered based on object and Multipurpose Internet Mail Extensions (MIME) types.
- **Ports other than 80**—Applications using ports other than 80 can be handled with the L4TM feature. L4TM can block access to specific applications by preventing access to the server or IP address blocks to which the client application must connect.

In the community college design, the Internet perimeter firewall can be configured to allow only web traffic (HTTP and HTTPS) outbound to the Internet from only the WSA. This prevents users from bypassing the WSA to browse the Internet.

**Note**

Peer-to-peer file sharing and Internet instant messaging applications can also be blocked using Cisco IPS appliances and modules and the Cisco ASA firewall (using modular policy framework).


The WSA L4TM service is deployed independently from the web proxy functionality. L4TM monitors network traffic for rogue activity and for any attempts to bypass port 80. It 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 periodic updates from the Cisco IronPort update server (<https://update-manifests.ironport.com>).

The following are some of the key guidelines when deploying L4TM:

- **Physical connection**—L4TM requires a copy of the traffic to be redirected to the WSA for monitoring. This can be done by connecting a physical network tap, configuring Switch Port Analyzer (SPAN) port mirroring on a Cisco Catalyst switch, or using a hub. Network TAPs forward packets in hardware, while SPAN port mirroring is generally done in software. However, SPAN port mirroring can be easily reconfigured, providing more flexibility.
- **Location**—L4TM should be deployed in the network where it can see as much traffic as possible before going out to the Internet through the firewall. Monitoring should occur before any device that performs NAT on client IP addresses.
- **Action setting**—The default action setting for L4TM is to passively monitor only. Optionally, you can configure L4TM to monitor and actively block suspicious traffic. TCP connections are reset by the generation of TCP resets, and UDP sessions are torn down using ICMP unreachable. The use of L4TM blocking requires that the L4TM and web proxy services are placed on the same network so that all clients are accessible on routes that are configured for data traffic.

In the community college design, L4TM can be deployed by configuring a SPAN session on the Internet perimeter distribution switch to monitor all TCP and UDP traffic on the links connecting to the core distribution switches. Using SPAN provides greater flexibility. Monitoring the distribution switches link to the core switches ensures that all client traffic is inspected before NAT and before traffic is sent to the Internet. Figure 6-12 shows this L4TM deployment option.

Figure 6-12 L4TM Deployment



If the Internet perimeter firewall is configured to block all traffic bound to the Internet except HTTP and HTTPS traffic from the WSA, or the ASA Botnet Traffic Filter feature is enabled, L4TM may not provide any additional benefit. Also, if active mitigation is required, the Cisco ASA Botnet Traffic Filter Feature or a Cisco IPS appliance or module deployed in inline mode is recommended. Both the ASA Botnet Traffic Filter and inline IPS provide better mitigation by blocking traffic automatically inline, stopping malicious traffic before it reaches the intended target.

For more information on how to configure the WSA, see the following guides:

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

Data Center Protection

Community colleges typically implement a data center that hosts the systems that serve the administrative and educational applications and store the data accessible to internal users. The infrastructure supporting them may include application servers, 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 Community College reference design such as identity and security services, unified communication services, mobility services, video services, partner applications, and other services.

Depending on the need and the size of the community college, a single data center may be deployed in the main campus. Smaller data centers or server farms may also be deployed in remote campuses as required.

Securing the data center is beyond the scope of this document. For more information on the best practices for securing a data center, see “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

One of the most vulnerable points of a network is at the access edge. The access layer is where end users such as students, staff, and faculty connect to the network. In the past, network administrators have largely relied on physical security to protect this part of the network. Unauthorized users were not allowed to enter secure buildings where they could plug into the network, and students did not carry computers with them. Today, with the proliferation of wireless networks, increased use of laptops and smart mobile devices, the community college IT department cannot simply rely on physical controls to prevent these unauthorized devices from plugging into ports of the access switches. Contractors and consultants regularly have access to secure areas, and students carrying laptops are common. There is nothing preventing a contractor or student from plugging into a wall jack in a classroom, lab, or conference room to gain access to the community college network. When connected to the network, everyone has access to all resources on the network.

Protection needs to be embedded into the network infrastructure, leveraging the native security features available in switches and routers. In addition, the network infrastructure should also provide dynamic identity or role-based access controls for all devices attempting to gain access. Implementing role-based access controls for users and devices help reduce the potential loss of sensitive information by enabling administrators to verify a user or device identity, privilege level, and security policy compliance before granting access to the network. Security policy compliance can consist of requiring anti-virus software, OS updates, or patches. Unauthorized or noncompliant devices can be placed in a quarantine area where remediation can occur before network access.

The Community College reference design achieves access security and control by leveraging the following technologies:

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

Cisco Catalyst Integrated Security Features

Cisco CISF is a set of security features available on Cisco Catalyst switches designed to protect the access layer infrastructure and users from spoofing, man-in-the-middle (MITM), DoS, and other network-based attacks. CISF should be considered part of the security baseline of any network and should be deployed on all access switches and ports within the community college network architecture.

CISF includes the following features:

- *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.
- *Dynamic ARP inspection (DAI)*—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.

- *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.
- *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 Identity-Based Network Services

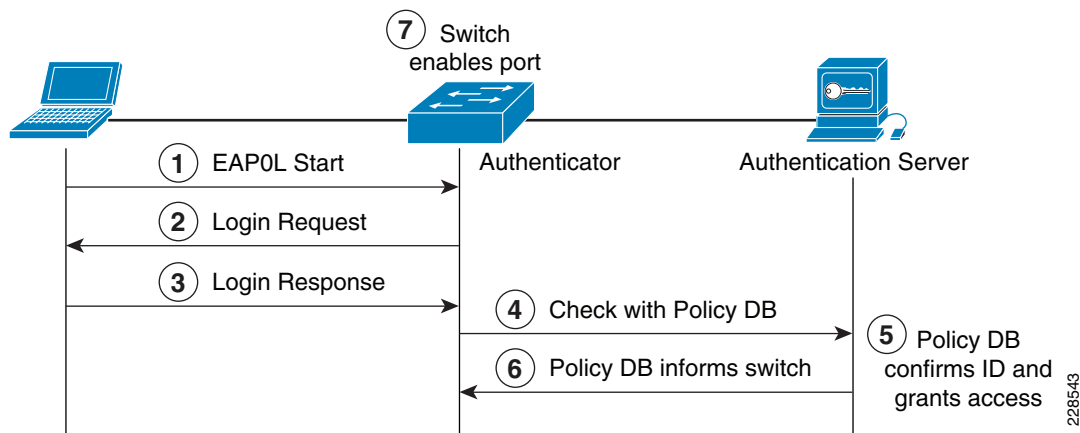
The Cisco IBNS solution is a set of Cisco IOS software services that provide secure user and host access to enterprise networks powered by Cisco Catalyst switches and wireless LANs. It provides standards-based network access control at the access layer by using the 802.1X protocol to secure the physical 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 a well-known way to secure wireless network access and is also capable of 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 the port except the request to start 802.1X authentication. After the 802.1X authentication successfully completes, the switch starts accepting other kinds of traffic on the port.

The high-level message exchange shown in [Figure 6-13](#) illustrates how port-based access control works within an identity-based system.

Figure 6-13 Port-Based Access Control



The following steps describe the port-based access control flow shown in [Figure 6-13](#):

-
- | | |
|---------------|---|
| Step 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). |
| Step 2 | When the start message is received, the LAN switch sends a login request to the client. |
| Step 3 | The client replies with a login response. |
| Step 4 | The switch forwards the response to the policy database (authentication server). |
| Step 5 | The authentication server authenticates the user. |
| Step 6 | After the user identity is confirmed, the policy database authorizes network access for the user and informs the LAN switch. |
| Step 7 | The LAN switch then enables the port connected to the client. |
-

The user or device credentials are processed by a AAA server. The AAA server is able to reference user or device profile information either internally, using the integrated user database, or externally using database sources such as Microsoft Active Directory, Lightweight Directory Access Protocol (LDAP), Novelle Directory, or Oracle databases. This enables the IBNS solution to be integrated into existing user management structures and schemes, which simplifies overall deployment.

802.1X and EAP

When authenticating users for network access, the client must provide user and/or device identification using strong authentication technologies. IEEE 802.1X does not dictate how this is achieved. Instead, 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 is defined by RFC 3748. EAP is a framework and not a specific authentication method. It 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

When 802.1X is enabled on a port, the default security posture is to drop all traffic except 802.1X EAPoL packets. This is a fundamental change from the traditional model, where traffic is allowed from the moment a port is enabled and a device is plugged into the port. Ports that were traditionally open are now closed by default. This is one of the key elements of the strong security and network access control provided by 802.1X. Understanding and accommodating for this change in access behavior facilitates 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 it connects. If a device without an 802.1X supplicant attempts to connect to a port that is enabled for 802.1X, it is subjected to the default security posture. The default security posture 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 a 802.1X-protected network.

Although an increasing number of devices support 802.1X, there are always 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 Preboot Execution Environment (PXE) boot machines. Some provisions must be made for these devices.

The Cisco IBNS solution provides two features to accommodate non 802.1X devices. These are MAC Authentication Bypass (MAB) and 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 a Guest VLAN is configured. Application of either or both of these features is required for a successful 802.1X deployment.

**Note**

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

802.1X in Community Colleges

As mentioned in the previous sections, 802.1X authentications require a supplicant on the host device. This typically poses a problem in community college environments that have a wide range of host devices and limited or no management of many of these devices. This makes a community college-wide 802.1X deployment very challenging. However, there may be pockets of a community college 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 school administration office or shared labs where PCs are managed. Other locations in the community college network still need protection, but student and faculty network access may be better served by a NAC Appliance Solution (discussed in the next section). In addition, for networks requiring role-based access control using posture assessments to ensure security compliance, Cisco NAC Appliance should be considered.

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

Cisco NAC Appliance

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.

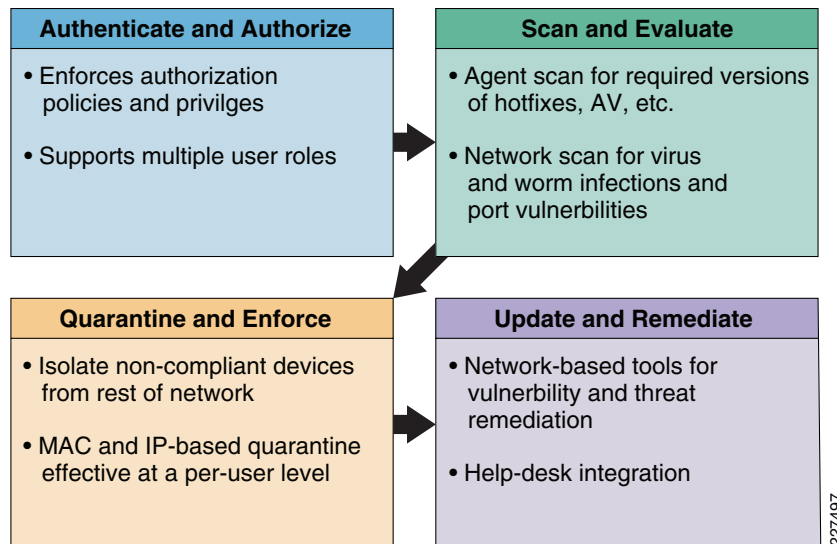
When deployed, Cisco NAC Appliance provides the following benefits:

- 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 requiring specific anti-virus or anti-spyware software, OS updates, or patches. Cisco NAC Appliance supports policies that vary by user type, device types, 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 6-14:

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

Figure 6-14 Four Functions of the NAC Solution

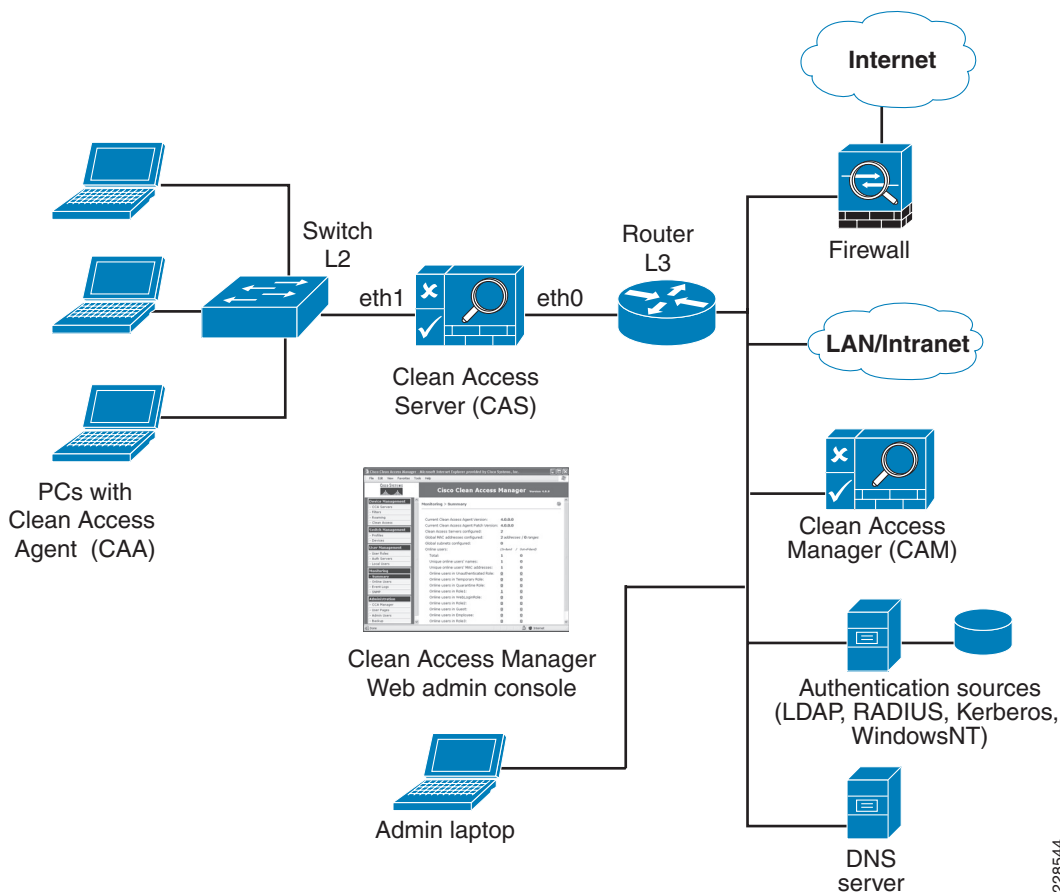


For more details of the NAC Appliance Solution, see the following URL:
<http://www.cisco.com/go/nacappliance>.

NAC Appliance Components

Cisco NAC Appliance is a network-centric, integrated solution administered from the Cisco Clean Access Manager (CAM) web console and enforced through the Cisco Clean Access Server (CAS) and (optionally) the Clean Access Agent (CAA) or 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.

Figure 6-15 shows Cisco NAC Appliance components.

Figure 6-15 NAC Appliance Components

Cisco Clean Access Manager

The Cisco CAM is the administration server for NAC Appliance deployments. The secure web console of the CAM is the single point of management for up to 20 Clean Access Servers in a deployment (or 40 CASs if using a SuperCAM). For OOB deployments, the web administration console controls the switches and VLAN assignment of user ports through the use of SNMP. In the Community College reference design, the CAM is located in the data center at the main campus site.

Cisco Clean Access Server

The Cisco CAS is the enforcement server between the untrusted network and the trusted network. The CAS enforces the policies defined by the CAM web administration console. Policies can include network access privileges, authentication requirements, bandwidth restrictions, and system requirements. The CAS can be installed as either a standalone appliance (like the Cisco NAC-3300 Series) or as a network module (Cisco NME-NAC-K9) in a Cisco ISR chassis. The CAS can be deployed in in-band (always inline with user traffic) or OOB (inline with user traffic only during authentication and posture assessment).

Additionally, the CAS can 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). Multiple CASs of varying size/capacity can be deployed to fit the needs of various network segments. For example,

Cisco NAC-3300 Series appliances can be installed in a main campus core to handle thousands of users, and one or more Cisco NAC network modules can be simultaneously installed in ISR platforms to accommodate smaller groups of users in a satellite office.

In the Community College reference design, the CAS would be located at the main campus and the remote campus sites, and deployed in Layer 2 OOB (for wireless clients) and Layer 3 OOB (for wired clients) modes for authentication and posture assessments.

Cisco Clean Access Agent

The Cisco CAA is an optional read-only agent that resides on Windows clients. It checks applications, files, services, or registry keys to ensure that clients meet the specified network and software requirements before 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.

In the community college, Cisco recommends that the CAA be used for the managed PCs, such as those for administrators and faculty.

Cisco NAC Web Agent

The Cisco NAC Web Agent provides temporal posture assessment for client machines. Using a Web browser, users launch the Cisco Web Agent executable file, 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 the network and their user ID disappears from the online users list.

In the Community College reference design, the NAC Web Agent is used for unmanaged clients such as student laptops and guest professors.

Clean Access Policy Updates

Regular updates of prepackaged policies/rules can be used to check the up-to-date status of operating systems, anti-virus (AV), anti-spyware (AS), and other client software. Built-in support is provided for 24 AV and 17 AS vendors.


NAC Appliance Modes and Positioning

The NAC Appliance can be deployed in multiple deployment options and placed at various locations in the network. The modes of operation can be generally defined as follows:

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

OOB Modes



OOB deployments require user traffic to traverse through the NAC Appliance only during authentication, posture assessment, and remediation (see [Figure 6-16](#)). When a user is authenticated and passes all policy checks, their traffic is switched normally through the network and bypasses the NAC Appliance.

Figure 6-16 Layer 2 OOB Topology

To deploy the NAC Appliance in OOB mode, the client device must be directly connected to the network via a Cisco Catalyst switch port. After the user is authenticated and passes posture assessment, the CAM instructs the switch to map the user port from an unauthenticated VLAN (which switches or routes user traffic to the CAS) to an authenticated (authorized) VLAN that offers full access privileges. For example, as shown in [Figure 6-16](#), the client PC is connected through VLAN 110 to the NAC CAS for the authentication and posture assessment, and is moved to VLAN 10 after it successfully completes the authentication/authorization and scan/evaluation phases of the NAC Appliance solution.

In-Band Modes

When the NAC Appliance is deployed in-band, all user traffic, both unauthenticated and authenticated, passes through the NAC Appliance. The CAS may be positioned logically or physically between the end users and the networks being protected. [Figure 6-17](#) shows a logical in-band topology example, and [Figure 6-18](#) shows a physical in-band topology example.

Figure 6-17 In-Band Virtual Gateway Topology**Figure 6-18 Physical In-Band Topology****In-Band Virtual Gateway**

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

- *Transparent*—For a given client VLAN, the NAC server bridges traffic from its untrusted interface to its trusted interface. The NAC server is aware of “upper layer” protocols and is able to permit those protocols that are necessary for a client to connect to the network, authenticate, and undergo

posture assessment and remediation. 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, such as DNS and DHCP. This option is viable when the NAC server is positioned physically in-band between the end users and the upstream network(s) being protected, as shown in [Figure 6-18](#).

- VLAN mapping**—This is similar in behavior to the transparent option except that rather than bridging the same VLAN from the untrusted side to the trusted side of the NAC server, two separate VLANs are used. For example, client VLAN 110 is defined for the untrusted interface of the NAC server. There is no routed interface or SVI associated with VLAN 110. VLAN 10 is configured between the trusted interface of the NAC server and the next-hop router interface (or SVI) for the client subnet. A mapping rule is made in the NAC server that forwards packets arriving on VLAN 110 and forwards them out VLAN 10 by swapping VLAN tag information. The process is reversed for packets returning to the client. Also, in this mode, BPDUs are not passed from the untrusted-side VLANs to their trusted-side counterparts.

The VLAN mapping option is typically used when the NAC server is positioned logically in-band between clients and the network(s) being protected, as shown in [Figure 6-17](#). This is the bridging option that should be used if the NAC Appliance is deployed in virtual gateway mode.

In-Band Real IP Gateway

When the NAC server is configured as a “real” IP gateway, it behaves like a router and routes packets between its interfaces. In this scenario, one or more client VLAN/subnets resides behind the untrusted interface. The NAC server 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 server 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 server is not currently able to support dynamic routing protocols. Therefore, 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 the IP address of the NAC server trusted interface as its next hop.

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

Depending on the topology, multiple options exist to facilitate routing clients to and from the NAC server, including ACLs, static routes, PBR, VRF-Lite, Multiprotocol Label Switching (MPLS) VPN, and other segmentation techniques. These options are discussed in later sections.

In-Band Versus Out-of-Band

[Table 6-1](#) summarizes various characteristics of the two deployment types.

Table 6-1 *In-Band versus Out-of-Band Characteristics*

In-Band Deployment Characteristics	Out-of-Band Deployment Characteristics
The CAS is always inline with user traffic (both before and after authentication, posture assessment, and remediation). Enforcement is achieved through being inline with traffic.	The CAS is inline with the user traffic only during the process of authentication, posture assessment, and remediation. After that, user traffic does not go to the CAS. Enforcement is achieved through the use of SNMP to control switches and VLAN assignments to end-user ports.

Table 6-1 In-Band versus Out-of-Band Characteristics (continued)

The CAS can be used to securely control authenticated and unauthenticated user traffic policies (based on port, protocol, subnet), bandwidth policies, and so on.	The CAS can control user traffic during the authentication, posture assessment, and remediation phases but cannot do so post remediation because 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 in-Band deployment with supported Cisco switches is compatible with 802.1X.	Cisco does not recommend using 802.1X in an OOB deployment, because conflicts will likely exist between Cisco NAC Appliance OOB and 802.1X in setting 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 access switches and WLCs to be supported by the NAC Appliance software for the NAC Manager to make the necessary changes to the switch ports and WLCs during the authentication, assessment, and remediation process. If access switches are to be used that are not supported, the NAC Solution must be deployed in in-band mode.

To obtain the latest list of supported devices, see 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/47/cam/47cam-book.html.

Layer 2 and Layer 3 Out-of-Band

The recommended deployment option for the Community College reference design is an OOB design. This provides the highest possible performance and scalability for traffic that has completed the authentication, posture assessment, and remediation stages of NAC. For wireless clients, a Layer 2 OOB solution should be deployed and for wired users, a Layer 2 OOB or Layer 3 OOB solution can be deployed, depending on the topology of your network.

NAC Deployment in the Community College Reference Design

Within the Community College reference design, a NAC Appliance solution is deployed at each of the site types; main campus, remote large campus, remote medium campus, and remote small campus. A centralized CAM is deployed at the main campus and is deployed within the data center at that site. A CAS is deployed at each of the campus sites (main and remote sites) and is connected within the service block connecting to the core switches at each of the sites.

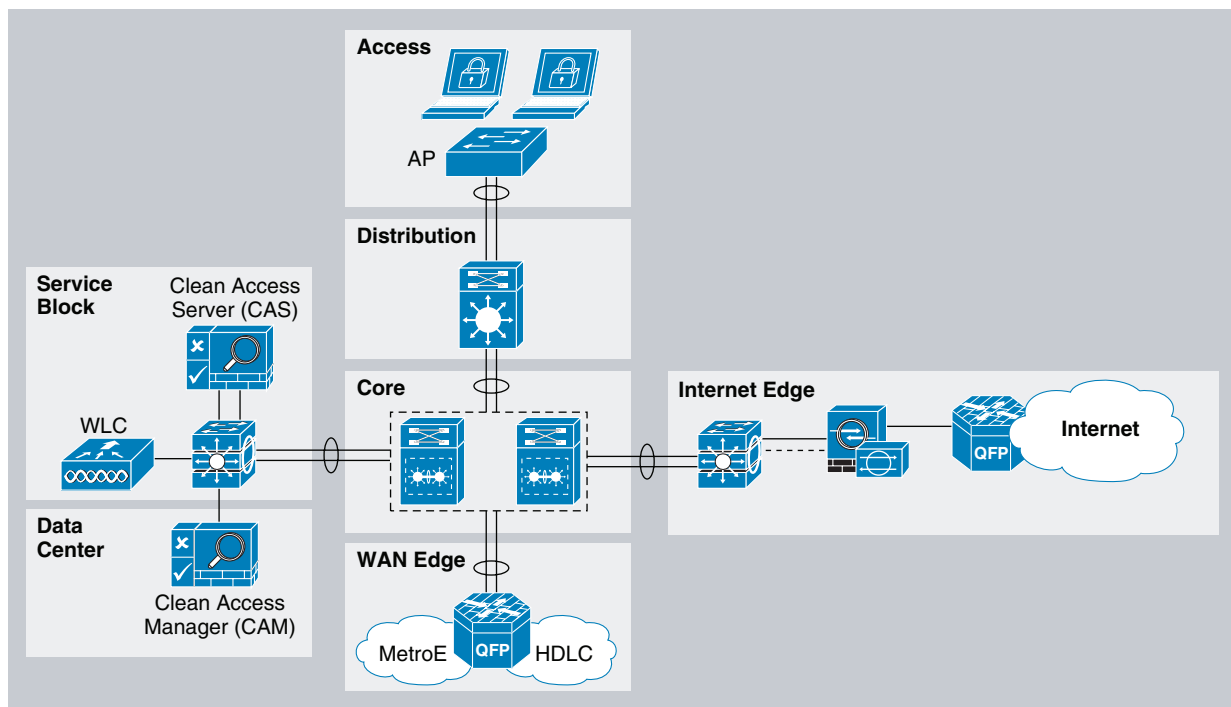
The Community College reference design provides host network connectivity using wired and wireless technologies. As such, the NAC Appliance solution must provide a solution for both connectivity options. For wireless clients, a Layer 2 OOB NAC solution is deployed, and for wired clients, a Layer 2 OOB or a Layer 3 OOB NAC solution may be deployed.

NAC Deployment for Wireless Clients

To provide network access control for wireless clients within the Community College reference design, the recommended design is the virtual gateway (bridge mode) and central deployment OOB solution. In this design, the NAC server must be placed Layer 2-adjacent to the WLC. In the Community College reference design, the WLCs are centrally deployed at each campus and are implemented in the service block off the core switches, as detailed in [Chapter 5, “Community College Mobility Design Considerations.”](#)

Therefore, the NAC server must also be implemented in the service block. The NAC Manager is implemented in the data center block, as shown in [Figure 6-19](#).


Figure 6-19 NAC OOB Deployment for Wireless Clients



228548

The WLC connects to the service block switch using a trunk port carrying the unauthenticated quarantine VLAN and authenticated access VLAN (VLAN 20 and 120). On the switch, the quarantine VLAN is trunked to the untrusted interface on the NAC server (CAS), and the access VLAN is trunked directly to the Layer 3 switch interface. Traffic that reaches the quarantine VLAN on the CAS is mapped to the access VLAN based on a static mapping configuration within the CAS.

When a wireless client associates to the WLC, it initially maps the WLAN/SSID to the quarantine VLAN interface and the client traffic flows in the quarantine VLAN (VLAN 120), which is trunked to the CAS untrusted interface. When NAC authentication, posture assessment, and remediation stages are complete and the user is certified, the NAC Manager sends an SNMP set message to the WLC that updates the VLAN ID from the quarantine VLAN to the access VLAN. After this occurs, the traffic then bypasses the NAC server and goes directly to the network. (See [Figure 6-20](#).)

Figure 6-20 Wireless NAC OOB Traffic Flow

When implementing the NAC OOB wireless solution, Cisco also recommends enabling RADIUS single sign-on (SSO), which is an option that does not require user intervention and is relatively easy to implement. This option makes use of the VPN SSO capability of the NAC solution, coupled with the Clean Access Agent software that runs on the client PC. VPN SSO uses RADIUS accounting records to notify the NAC Appliance about authenticated remote access users that connect to the network. In the same way, this feature can be used in conjunction with the WLAN controller to automatically inform the NAC server about authenticated wireless clients that connect to the network.

The most transparent method to facilitate wireless user authentication is to enable VPN SSO authentication on the NAC server and configure the WLCs to forward RADIUS accounting to the NAC server. In the event that accounting records need to be forwarded to a RADIUS server upstream in the network, the NAC server can be configured to forward the accounting packet to the RADIUS server.

**Note**

If VPN SSO authentication is enabled without the Clean Access Agent installed on the client PC, users are still automatically authenticated. However, they are not automatically connected through the NAC Appliance until their web browser is opened and a connection attempt is made. In this case, when users open their web browser, they are momentarily redirected (without a logon prompt) within the agentless phase. When the SSO process is complete, they are connected to their originally requested URL.

For more information on deploying NAC OOB for wireless environments, see the *NAC Out-Of-Band (OOB) Wireless Configuration Example* at the following URL:
http://www.cisco.com/en/US/products/ps6128/products_configuration_example09186a0080a138cc.shtml.

NAC Deployment for Wired Clients

For wired clients, the Community College reference design also uses a central OOB NAC deployment with a NAC server implemented at each of the campus sites deployed in the service block off the core switch. Depending on the type of network topology deployed, a Layer 3 OOB or Layer 2 OOB solution can be deployed. If the Layer 2 OOB solution is used, the same NAC server can be leveraged for both wired and wireless clients. However, if the Layer 3 OOB solution is deployed, separate NAC servers must be deployed for wired and wireless users.

Layer 3 Out-of-Band Deployment

Layer 3 (L3) OOB is best suited for routed access designs and has rapidly become one of the most popular deployment methodologies for NAC. By deploying NAC in an L3 OOB methodology, a single NAC Appliance can scale to accommodate more users. This deployment also allows NAC Appliances to be centrally located rather than distributed across the campus or organization. Thus, L3 OOB deployments are much more cost-effective, both from a capital and operational expense standpoint.

For the main, large, and medium remote campus locations, an L3 OOB NAC deployment is recommended, given the 3-tier hierarchical design. In the L3 OOB NAC solution, when a user connects to the access switch before being certified by the NAC server, the user is placed in the authentication VLAN (also called “dirty” VLAN). The user should not have access to any part of the network from the authentication VLAN except for the NAC server and the remediation servers in the quarantine segment. After users are certified by the NAC server, they are placed in the authenticated access VLAN, where their traffic is switched normally through the network and bypasses the NAC server.

The following are three widely deployed techniques for redirecting client traffic from the dirty VLAN to the NAC server for authentication, posture assessment, and remediation purposes:

- Access control lists—Use ACLs on the edge access switches to allow traffic from the unauthenticated VLAN only to the NAC server untrusted interface and specific infrastructure resources needed to get on the network for authentication purposes such as DHCP, DNS, and remediation servers. All other traffic from the dirty VLAN must be blocked.
- VRFs/GRE/MPLS—Use VRFs to route unauthenticated traffic to the CAS. Traffic policies configured on the NAC server (CAS) are used for enforcement on the dirty network. This approach has two sub-approaches. In the first approach, VRFs are pervasive throughout the infrastructure, in which case all Layer 3 devices participate in the tag switching. The second approach uses VRF-Lite and GRE tunnels to tunnel the VRFs through the Layer 3 devices that do not understand the tag switching. The benefit to the second approach is that minimal configuration changes are required to your core infrastructure. For more information on this approach, see the following URL: http://www.cisco.com/en/US/products/ps6128/products_configuration_example09186a0080a3a8a7.shtml.
- Policy-based routing—Use PBR to redirect all traffic in the dirty VLAN to the NAC server. PBR needs to be configured on every Layer 3 hop between the dirty VLAN and the NAC server to ensure that traffic is appropriately redirected.

The most common approach used for isolating the dirty VLAN traffic is to use ACLs. The ACLs on the Layer 3 edge access switches act as the enforcement point to ensure segregation between the “clean” and “dirty” networks. When clients first attach to the network, they are placed in a quarantine or dirty VLAN on the access switches. ACLs should be applied on the SVIs for the dirty VLAN. This ACL should block all access from the dirty VLAN going to the internal networks and allow traffic only to the untrusted interface on the NAC server, the needed remediation servers, and a few infrastructure devices needed for network access such as the DNS, DHCP, and Active Directory servers.

The clients need to communicate with the NAC server untrusted interface for the certification process. The ACLs on the access switches act as the enforcement point for path isolation for the dirty VLAN traffic. Methods for getting the dirty VLAN traffic to the untrusted interface vary, depending on whether the NAC Client Agent is used.

When the NAC agent is used, the NAC Agent communicates with the NAC server untrusted interface to initiate the login process. The NAC Agent tries to discover the NAC server based on the known discovery host value. The discovery host value in the NAC Agent points to the untrusted interface of the NAC server. In the Community College reference design, the NAC Agent can be used for managed PCs such as administrative staff and faculty.

Web login is typically required for student login sessions because student laptops are typically not managed. With the ACL isolation technique, the NAC server untrusted interface is not directly in the path of the data traffic; therefore, the user is not automatically redirected to the login page when first opening the browser. The following two options can enable the end host to get the login page:


- Option 1—Have a guest login URL known to the users (for example, *guest.nac.local*). In this case, the guest must open a browser and manually enter this URL, which redirects them to the login page.
- Option 2—Create a dummy DNS server for the unauthenticated user subnet. This dummy DNS server resolves every URL to the untrusted interface of the NAC server. When guests open a browser, regardless of which URL they are trying to reach, they are redirected to the login page. When users are then moved to the respective Role/VLAN, they get a new DNS address assignment when performing IP release/renew on a successful login.

Layer 2 Out-of-Band Deployment

For the small remote campus locations, a two-tier, collapsed core/distribution LAN design is recommended, as explained in [Chapter 3, “Community College LAN Design Considerations.”](#)

In a collapsed core/distribution design, the CAS should be deployed in the services block connected to the core/distribution switch. In this simple topology, a Layer 2 Out-of-Band (L2 OOB) NAC deployment can be used.

In the L2 OOB NAC design for the small remote campus, the unauthenticated and authenticated VLANs on the access switch (VLANs 30 and 130) are extended to the core/distribution switch using a trunk connection, as shown in [Figure 6-21](#).

Figure 6-21 Layer 2 OOB Topology

When a client device initially connects to the access switch before authentication, it is placed in the unauthenticated VLAN (VLAN 130), which connects the client directly to the untrusted interface of the CAS. The CAS maps VLAN 130 to the VLAN 30 trusted interface, allowing the client to obtain an IP address that belongs on VLAN 30. After the client is authenticated and passes the posture assessment, the access switch is instructed, via SNMP from the CAM, to change the client VLAN to the authenticated VLAN (VLAN 30), where the traffic now bypasses the CAS to access the rest of the network. Although the client has changed Layer 2 VLANs, its Layer 3 network connections are unchanged.

NAC Availability Considerations

Both the CAS and CAM are highly involved in client network access. Consideration must be given to the impact on clients if either a CAS or CAM fails or needs 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 NAC is deployed in in-band mode, it is inline even after authentication and certification. A CAS outage for inline clients prevents access for all clients. However, if NAC is deployed in OOB mode, a CAS outage does not affect already connected clients but does prevent network access for new clients.

In situations where availability of a CAS is critical, a high availability (HA) CAS solution can be implemented where a pair of CAS servers are installed using a primary CAS, and a secondary in hot standby. For more information, see the *Cisco NAC Appliance - Clean Access Server Installation and Configuration Guide* at the following URL:
http://www.cisco.com/en/US/docs/security/nac/appliance/configuration_guide/461/cas/461cas-book.html.

The CAM is also a critical part of the authentication, authorization, and posture assessment phases of NAC. Although it does not pass client traffic, the impact of its availability needs to be considered in the network design as well. Like the CAS, the CAM has an 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 that defines how it manages client traffic in a situation where the CAM is unavailable.

The use of the CAM and CAS HA features depends on the requirements of the community college. However, CAS fallback should always be configured to ensure that critical network services are available, even during a network outage.

Endpoint Protection

Servers, desktop computers, laptops, printers, and IP phones are examples of the diverse network endpoints found in community college environments. Properly securing the endpoints requires not only adoption of the appropriate technical controls but also end-user awareness. The community college security strategy must include security awareness campaigns and programs. Students, staff, and faculty must be continuously educated in current threats, Internet-use best practices, and the security measures needed for keeping endpoints up-to-date with the latest updates, patches, and fixes.

The Community College reference design implements a range of security controls designed to protect the endpoints, which include Cisco host-based IPS, network-based IPS, and web and E-mail traffic security.

For host-based IPS, the Community College reference design leverages the Cisco Security Agent on managed end-user workstations and servers. Cisco Security Agent uses behavior-based security to take a proactive approach to preventing malicious activity on the hosts. When an application attempts an operation on the host, Cisco Security Agent checks the operation against the security policy of the application, and makes a real-time decision to allow or deny the operation along with determining whether to log the operation request. Security policies are assigned by IT or security administrators individually, per department, or organization-wide.

Cisco Security Agents are centrally managed with the Cisco Security Agent Management Center, which is placed in a secure segment in the data center. Cisco Security Agent Management Center also provides centralized reporting and global correlation.

Community College Mission Relevancy

The service fabric provides the network foundation for the Community College reference design. The network service fabric is a collection of products, features, and technologies that provide a robust routing and switching foundation on which all solutions and services are built to solve the business challenges facing community colleges. These business challenges include the following:

- Virtual learning environments
- Secure connected classrooms
- Safety and security
- Operational efficiencies

The previous sections of this chapter focused on the specific design considerations for securing the community college service fabric network. This section discusses how these security design considerations relate to these business challenges.

Virtual Learning Environments

One of the key challenges that face community colleges is extending their learning environments beyond brick and mortar campuses to allow online/distance learning, professor collaboration, and anytime, anywhere access for students to obtain course materials. Maintaining a secure virtual learning environment is critical for community colleges. The community college security design helps establish the foundation for providing a secure virtual learning environment in the following areas:

- *Secure remote access*

- Allows community colleges to extend their network to anyone, anytime, anywhere by providing a secure client-based or web-based remote access solution.
- Provides granular and encrypted access to learning resources based on user or security requirements.
- *Internet perimeter*—Protects and controls access to the community college network infrastructure from remote students, faculty, staff, and guest professors by properly securing the Internet perimeter by using firewalls, intrusion prevention systems, and a DMZ to protect against unauthorized access and malicious attacks.
- *Securing Video Portal*—Secures the network and data center to prevent misdirected students or malicious intruders from hacking into restricted servers or issuing attacks on the video portal learning infrastructure.
- *Network telemetry and monitoring*—Provides visibility into attacks or malicious activity on the network by monitoring the network using NetFlow, Syslog, and SNMP.

Secure Connected Classrooms

Although providing connectivity to students while attending class is the foundation of twenty-first century learning, this poses many problems for community colleges. They must ensure that the person accessing the network should be allowed on the network, and that the computer accessing the network is free from viruses and other malware that might adversely affect the network and other users. In addition, while connectivity is provided, steps should be taken to prevent unauthorized access to restricted resources and protect against inadvertent or deliberate network attacks. The security design within the community college service fabric helps to address these challenges in the following ways:

- *Network access control*—Implementing role-based network access controls for wired, wireless, and remote users and devices to help reduce the potential loss of sensitive information by enabling administrators to verify a user or device identity, privilege level, and security policy compliance before granting access to the network.
- *Access layer security*—Enabling Cisco CISF on the access layer switches to protect the access layer infrastructure and users from spoofing, man-in-the-middle, DoS and other network-based attacks. CISF includes features such as Port Security, DHCP Snooping, Dynamic ARP Inspection, and IP Source Guard.
- *Web security*—Deploying a Cisco IronPort Web Security Appliance (WSA) to block access to sites with content that may be harmful or inappropriate, and to protect community colleges from web-based malware or spyware.
- *Network and data security*—Securing the network and data center to prevent misdirected students or malicious intruders from hacking into restricted servers or issuing attacks on the network infrastructure.

Safety and Security

Providing a safe and secure environment is a top responsibility for community college administrators and community leaders. Without adequate protection, schools may be threatened by harmful or inappropriate content that can put the well-being of the students at risk, the theft of student records and private data, the loss of school network and service availability, as well as the abuse of internal applications and network resources. A safe community college is one that successfully uses the right

tools to ensure the safety and security of students, staff, and faculty, and guarantees an immediate and effective response to security and safety incidents. The most effective strategy is one that combines physical and network controls, not in isolation but rather in collaboration and with a common purpose.

Because many of the physical security components such as video surveillance and unified communication services rely on the IP infrastructure, it is critical to ensure the availability and integrity of this infrastructure. The security design within the Community College reference design helps to ensure the availability and the integrity of the network infrastructure that the physical security components rely on by focusing 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 connectivity to the Internet, Internet2, and NLR networks and protecting internal resources and users from malware, viruses, and other malicious software.
 - Protecting students, staff, and faculty from harmful content.
 - Enforcing E-mail and web browsing policies.
- *Data center protection*—Ensuring the availability and integrity of centralized applications and systems. Protecting the confidentiality and privacy of student, staff, and faculty records.
- *Network access security and control*
 - Securing the access edges.
 - Enforcing authentication and role-based access for students, staff, and faculty residing at the main and remote campuses.
 - Ensuring systems are up-to-date and in compliance with the community colleges network security policies.
- *Network endpoint protection*
 - Protecting servers and school-controlled systems (computer labs, school-provided laptops, and so on) from viruses, malware, botnets, and other malicious software.
 - Enforcing E-mail and web browsing policies for staff and faculty.

Operational Efficiencies

Community colleges are faced with the daunting task of doing more with less, facing explosive growth as budgets and resources are reduced. The Community College reference design leverages the network as a platform to deliver expanded educational services and data center optimizations to create operational efficiencies to reduce costs and capitalize on under-used network capacity. The network as a platform goes beyond merely consolidating voice, video, and data services on a single converged network; rather it consolidates all IP-based services to use the network (wired or wireless) to extend cost reduction, improve utilization on under-used networks, and add flexibility to community colleges through business process improvements.

With these critical services relying heavily on the network infrastructure, it is imperative that the IP infrastructure remains operational at all times, and it is critical that security be implemented throughout the network infrastructure to ensure the availability and the integrity of the network.

