

SAFE: IDS Deployment, Tuning, and Logging in Depth

Authors

Ido Dubrawsky and Roland Saville are the primary authors of this white paper. Ido Dubrawsky is the lead architect for the reference implementation at Cisco Systems® headquarters in San Jose, California. Jason Halpern provided significant contributions to this paper. All three are network architects focused on VPN and security solutions.

Abstract

The *SAFE: A Security Blueprint for Enterprise Networks* white paper, available at the Cisco® SAFE Website (<http://www.cisco.com/go/safe>), provides high-level design guidance for Network-Based Intrusion Detection System (NIDS) sensor placement, host-based Intrusion Prevention System (IPS) implementation, and secure syslog messaging. *SAFE: IDS Deployment, Tuning, and Logging in Depth* builds upon this and other SAFE security white papers by discussing in detail IDS and logging design considerations and best practices for networks today. Information is included on network security technologies such as IPS and IDS alarm verification systems (AVSs) that have reached the marketplace since the original SAFE white papers were released.

The SAFE security white paper most appropriate to the size of the network under consideration should be read prior to reading this paper. For example, a business with a large network might look first at the SAFE enterprise white paper, which will frame the IDS and syslog conversation within the context of overall security design. SAFE represents a system-based approach to security design: it focuses on overall design goals and translates those goals into specific configurations and topologies. SAFE configuration information is based on Cisco products and those of its partners.

- This paper begins with an architecture overview and then details large, medium, and small single-site and multi-site designs now under consideration. Within each design, a network IDS, host-based IPSs, and logging are deployed across multiple modules. (The concept of modules is addressed in the SAFE security white papers.) The following topics are covered for each design (where appropriate):
 - Overall design best practices
 - Performance and scalability
 - Tuning, reporting, and management
 - Design alternatives

Appendix A is a primer on intrusion detection and Appendix B introduces syslog technology. Appendix C defines the technical terms used in this paper. Readers who are unfamiliar with security technologies are encouraged to read the appendices before the rest of the paper.



Audience

This paper is intended for security operations and network operations managers. Network managers can read the introductory sections in each area to obtain an overview of design strategies and considerations. Network engineers or designers can read the paper in its entirety for complete design information and threat analysis, including configuration scenarios for the devices involved. Because this paper covers a wide deployment range, it may be helpful to read the introductory sections first and then skip to the type of network being deployed.

Caveats

This paper assumes knowledge of the sensitivity of the data in the network under consideration so that an appropriate level of security can be provided. This paper also assumes that a pre-existing security policy is in place. Cisco does not recommend deploying any technology without an associated security policy. The greater the effort made to provide a complete and concise security policy before deploying IDS and expending effort in tuning, monitoring, and correlating the information provided will result in a larger return on investment (ROI).

Following the guidelines in this paper does not guarantee a secure network environment. It is the intention of the authors to provide information on the risks and benefits of the technologies addressed so that an informed choice can be made. Reasonable security can be achieved by establishing a good security policy, following the guidelines in this and the earlier SAFE security white papers, staying up to date on the latest developments in the hacker and security communities, and maintaining and monitoring all systems using sound practices.

This paper does not cover variations in the technologies addressed, such as Syslog over TCP, nor does it discuss whether signature-based IDS is superior to anomaly-based IDS.

Architecture Overview

Design Fundamentals

The networks discussed in this paper have common design fundamentals, including segmentation into modules, which enables network traffic to be isolated and thus simplifies the management of network security. The design objectives that are covered in this paper are listed below:

- Security implementation throughout the network
- Secure management and reporting
- Intrusion detection for critical resources, subnets, and servers
- Support for emerging network applications

The security architecture of the network must prevent attacks from affecting valuable network resources. Attacks that succeed in penetrating the first line of defense, or those that originate from within the network, must be quickly and accurately detected and contained to minimize their effect on the rest of the network.

SAFE Axioms

The following axioms apply to nearly all NIDS, host-based IPS, and logging deployments, and are thus included in the front of this paper to limit redundancy. The paper assumes conformity with the security axioms in the original SAFE white paper that is most appropriate to the size of the network under consideration.



Monitoring Throughout the Network

Because attacks originate from within the corporate network as well as from the Internet, placing security monitoring devices such as network IDS sensors, host-based IPS agents, and logging devices only at network entry points may not suffice. In addition, by placing monitoring devices throughout the network, security administrators can track the spread of attacks and can take corrective action such as modifying access control lists (ACLs) on routers and internal firewalls to prevent the attack or contagion from reaching critical resources. Suggested locations for placement follow:

- Behind firewalls that provide access to the module
- On demilitarized zone (DMZ) segments that house public servers (Web, File Transfer Protocol [FTP], Domain Name System [DNS]), data centers, or e-commerce servers
- Behind VPN concentrators, to monitor unencrypted VPN traffic
- On segments that house corporate servers or other intranet services that are sensitive according to the defined security policy
- On segments that house network and security management servers
- On the corporate intranet where critical resources are located
- At corporate extranet junction points between the campus network and branch networks as well as between the enterprise and partner networks

Switched Infrastructure

Most corporate LANs have migrated from a hub to a switch infrastructure over the past five years. Hubs broadcast all traffic to all ports, allowing network IDS sensors to automatically view and capture traffic crossing the hub. Switches only broadcast traffic to the port on which the destination MAC address(es) has been learned. Therefore, network IDS sensors will not see all traffic crossing the switch under normal circumstances. In order to overcome this limitation, the monitoring interfaces of all sensors should be connected to a switch port that has traffic from all necessary ports or all necessary VLANs mirrored to it.

The network designer should also use caution when selecting the switch to use when implementing countermeasures. The switch should support the ability for the Switched Port Analyzer (SPAN) or mirroring port to both receive and transmit traffic, because any countermeasures employed come from the network IDS sensor itself. Furthermore, because the IDS sensor spoofs the MAC address of the device that it is protecting (for example, the Web server, mail server, etc.), the switch must be configurable such that MAC address learning is disabled for the SPAN port. Otherwise, the switch will direct all legitimate traffic destined for the server to the IDS sensor until the server sends another packet and causes the MAC address to be re-learned on the correct port.

Countermeasures

Many network-based IDSs provide countermeasure capabilities. These capabilities include the ability to introduce new ACL rules into an edge router to block traffic coming in from an offending IP address as well as the ability to terminate the communication channel between the attacker and the target host through the use of TCP Reset (RST) packets. TCP resets should not be sent to known, valid devices such as network infrastructure components.

An additional network IDS countermeasure is the use of in-line communication termination. In this case, the NIDS no longer sits as a passive monitor but is instead a part of the router or firewall operating system. This in-line capability, which is quickly being included in traditional network IDS appliances, removes the need for the network



designer to be concerned with the use of SPAN or a network tap. However, the NIDS appliance can cause a network bottleneck. In addition, failure of the NIDS will leave the network in a closed state because no traffic can pass through the failed appliance.

Staffing

Threat Validation and Response (TVR) systems can dramatically reduce the time that security personnel spend investigating IDS alarms. These systems provide automated analysis and alarm escalation or elimination. As a result, security personnel can be more confident that an alarm on the IDS console is valid and requires further investigation.

Some host-based IPS implementations, such as the Cisco Security Agent, block malicious activity at the operating system level. Additionally, countermeasures on network IDS sensors provide some level of automated prevention once a signature is detected. To successfully implement network intrusion detection and syslog monitoring, a company must have the resources available to monitor and respond to incidents. Twenty-four-hour staffing requires approximately five employees (in three shifts of 8 hours each, 7 days per week, plus additional coverage for weekends, vacation, illness, etc.). These employees may not need to spend 100 percent of their time monitoring IDS alerts and syslog messages. Time spent depends on the number of network sensors and host-based IPS agents deployed, the number of alerts and syslog messages received, and the level of responsiveness desired. If the company is not prepared to expend the necessary resources, then the security monitoring function can be outsourced. Implementing only periodic monitoring of IDS and syslog reports also decreases staffing requirements. The network staff will not be in a position to respond immediately to a security incident, of course: staff will only be able to see that a security incident has occurred. The staffing decision should be based upon the value of the resources being protected, the expected loss to the company if resources are compromised, and the company's security policy.

Network IDS Tuning Methodology

Tuning sensors is critical to a successful network IDS implementation. Without tuning, IDS sensors generate alerts in response to all traffic matching an established criteria. The number of false alarms can easily overwhelm security personnel and reduce the value of the information the IDS provides. IDS sensors that are not tuned may also raise an alarm on an attack that does not impact the network resource being protected. Given enough false alarms, the security personnel will tend to ignore the IDS sensor—increasing the odds that a real attack will be successful—or will disable it, which leaves the network segment unprotected. The following are guidelines for tuning network IDS sensors.

Step 1. Identify Potential Locations for Sensors

To properly tune IDS sensors, the first step is to identify network locations where the sensors can be placed for maximum efficiency. In the public Internet segment, placing an IDS sensor in a location without traffic filtering can overwhelm the sensor's capabilities. Place IDS sensors behind a traffic filtering device such as a firewall, or in the case of a DMZ, a router with ACLs. If traffic is filtered upstream of the IDS, only traffic destined for the network resources being protected will reach the IDS, which reduces the sensor's workload.

Step 2. Apply an Initial Configuration

The objective of step 2 is to take a first pass at configuring the network IDS sensors. First, sensors are classified and grouped according to active signatures and are then configured by group with a common signature profile. The sensors in a group are managed collectively, which simplifies the management of



large groups of sensors. A decision must be made to either deploy signature profiles using the default values or to tune specific signatures. Refer to Appendix A, “Intrusion Detection Primer,” for a discussion of signature types and how they are implemented on different sensors throughout the network. A general guideline, if this is the initial deployment of network IDS, is to use the default settings at this step and then to tune them in one of the later steps. The destination for sensor alarms should also be determined during this step

Step 3. Monitor the Sensor While Tuning

The objective of steps 3 and 4 is to monitor IDS sensor alarms and tune out any alarms caused by normal background traffic rather than malicious activity. As steps 3 and 4 are executed, there should be a decrease in the number of false alarms. The monitoring period can last from several days to a week or more. If this is an initial network IDS implementation with a large number of sensors, the number of alarms may be quite large.

Step 4. Analyze Alarms, Tune Out False Positives, and Implement Signature Tuning (If Needed)

During the initial tuning period, you will need to determine the cause of every alarm in order to identify false positives. This task could be tedious, but it is necessary for your network IDS deployment to be of any use in detecting malicious activity.

How do you determine if an alarm is a false positive? A few suggestions follow:

- 1) Consult the IDS Network Security Database (NSDB) to determine whether the alarm in question is typically due to malicious activity or normal network background activity.
- 2) If the alarm appears to be the result of normal activity, determine the source of the alarm.
- 3) If the source of the alarm is a server, work with the appropriate applications and operating system support groups to determine if they think that this is normal behavior, malicious activity, or perhaps an unknown process or application running on the server. If possible, duplicate conditions in a known secure lab with a known, uncompromised server. The server may have to be built from scratch to guarantee this. It is critical that a list of all systems, their operating systems, and the applications running within the monitored environment be available in order to refine the IDS tuning as well as to determine the applicability of any patches released by vendors.
- 4) If the source is a network device (router, switch, etc.), follow step 3 procedures, but consult the appropriate network operations staff.
- 5) Deploy a threat analysis system in order to help validate IDS alarms as well as evaluate their significance, impact, and possible responses.

Once a false positive is identified, determine first if the activity that caused the alert can be modified so that an alarm is not generated. For instance, if NetBIOS is running on a server and is unnecessary on that server, disable NetBIOS before tuning the network IDS sensors to ignore alerts caused by the traffic. If the service or application generating the alarm is required, however, or cannot be disabled, an alternative is to configure the network IDS sensors to ignore the alarm. The following configurations are possible:

- Do not generate an alarm if this particular signature is seen for all sources
- Do not generate an alarm if this particular signature is seen from this particular address

It is always better to be specific when implementing security. Therefore, it is recommended that the network IDS sensor be tuned such that it does not generate an alarm for the specific signature from the specific address. However, if additional servers that also require the application or service are later



implemented on the same network segment, it may be necessary to go back and tune these sources out as well. The administrative overhead is reduced if the server IP addresses can be grouped together into a single profile.

As mentioned in step 2, IDS sensors can be grouped and collectively managed according to their alarm profiles. This tuning feature reduces the administrative impact when sensors are added to the network. Similarly, IDS signatures can be grouped into four broad categories:

- **Exploit signatures**—These signatures indicate attempts to compromise network systems through buffer overflows, Structured Query Language (SQL) injection, brute force password attacks, and other well-known exploits.
- **Connection signatures**—These signatures indicate reconnaissance activity: an attacker is enumerating systems and services on the network.
- **String-match signatures**—These signatures indicate corporate policy violations detected by sensors searching for custom text strings in the network traffic. For example, an IDS sensor could signal an alarm on any connection that transmits the phrase “Company Confidential” by e-mail or FTP.
- **Denial-of-service (DoS) signatures**—These signatures indicate attempts by attack tools such as Trinoo, tribal flood network (TFN), Stacheldrucht, and TCP SYN floods to consume bandwidth or computing resources to disrupt normal operations.

These signatures are discussed in detail in Appendix A. If the default signature tuning templates were deployed in step 2, then specific tuning should now be attempted. Implement the specific tuning in stages, rather than all at once. If the sensors are grouped into four profiles, for example, tune each profile separately. Tune each category of signature separately as well. If a group of new false positives appears after pushing a new configuration to the sensors, this tuning method helps to identify the specific tuning responsible. The potential confusion that tuning all of the sensors at once would create must be weighed against the time that it takes to tune the sensors in stages.

Step 5. Selectively Implement Response Actions

Once the false positives are tuned out and logging due to IDS tuning changes is sufficiently reduced, response actions such as TCP resets, shunning, and IP logging can be implemented. Please see the section titled “Response Actions” for usage guidelines. Again, deploy the response actions in stages, rather than all at once.

Step 6. Update Sensors with New Signatures

Automatic signature updates should be implemented for deployments with large numbers of sensors. See the subsection titled “Configuration Management” under “Scaling Network IDS” for a discussion of the benefits of implementing these updates. After a signature upgrade, repeat steps 3 through 5 to tune each new signature appropriately. The time required to monitor the deployment for false positives should be reduced to just a day or two, because the number of new signatures introduced with each signature pack is small. This is another reason to keep the signature levels on the sensors up to date: the greater the difference in signature pack levels, the larger the number of new signatures introduced and the greater the length of time required to monitor and tune out false positives.



Host-Based IPS Tuning

Host-based IPSs use system policies to differentiate between permitted actions and those that should be denied. Unlike the misuse-detection-based network IDS where each attack is represented by a signature, host-based IPSs focus on calls within the operating system to identify permissible actions. The software installs as shims in the operating system call stack that intercept and evaluate various system calls to the operating system kernel. The shims provide visibility into the process actions and represent locations in the operating system where an action can be controlled. Typically, host-based IPS software installs with various default policies that can be used to protect a range of system types from desktop hosts to application servers. Host-based IPS tuning is similar to network IDS tuning: a monitoring period ensues after IPS installation that is governed by the number of alarms seen. As tuning proceeds, there should be a decrease in the number of alarms. This period can last for several days (the NIDS initial tuning period can extend to one or more weeks). Once the final tuning is in place, the policy can be locked in and switched over to a responsive mode that will automatically react to policy violations without human confirmation or other intervention.

External Threat Verification and Response

New Threat Verification and Response (TVR) technologies reduce the workload of security personnel and increase the accuracy and reliability of the IDS. By analyzing, verifying, and then validating or cancelling IDS alarms, these technologies greatly reduce the number of false alarms and thus the need to tune IDS sensors.

The TVR system uses an external system to monitor alarms. When an alarm is generated, the TVR system identifies the target system, determines the host operating system, and initiates an in-depth check of the host. For Windows-based systems, the in-depth check can include registry key checking, file analysis, log analysis, and file capture. If the alarm is determined to be false, the TVR system downgrades or removes the alarm from the management console. However, if an attack is detected and verified, security personnel are alerted.

Scaling the Deployment

Scaling Network IDS

To scale the network IDS management infrastructure, the network designer must first understand the functions that the infrastructure provides as well as any limiting factors.

The network IDS management infrastructure provides configuration management and event monitoring and management. Event monitoring and management can be divided into 1) real-time event monitoring and management and 2) analysis based on archived information (reporting). These functions can be handled by a single server or they can be placed on separate servers to scale deployment.

Configuration Management

Configuration management includes sensors tuning as well as management of the operating system version and the signature database on the sensor. Ensuring that each sensor is active and pushing new configuration updates to these sensors consumes memory and CPU cycles, but the number of sensors configured by a single management console is



typically not the limiting factor in a network IDS deployment. The limiting factor is the number of alarms per second that the event transaction database can handle. The following are some configuration best practices that will improve IDS efficiency:

- When setting up a large deployment of sensors, automatically update signature packs rather than manually upgrading every sensor. Security operations personnel will then have more time to analyze events. When new signature packs are available, download the new signature packs to a secure server within the management network. For the SAFE large enterprise model, there exists an out-of-band (OOB) management network for sensor signature updates and retrieving sensor alarms and logs. Place the signature packs on a dedicated FTP server within the management network. If a signature update is not available, then a custom signature may be created in order to detect and mitigate a specific attack. The FTP server should be configured to allow read-only access to the files within the directory on which the signature packs are placed, and then only from the account that the sensors will use. The sensors can then be configured to automatically check the FTP server periodically, such as once a week on a certain day, to look for the new signature packs and to update themselves. A host IDS or IPS can be used to protect this server from attack by an outside party.
- Stagger the time of day when the sensors check the FTP server for new signature packs, perhaps through a pre-determined “change window.” This will prevent multiple sensors from overwhelming the FTP server by asking for the same file at the same time. The need to upgrade sensors with the latest signature packs must be balanced against the momentary downtime—and therefore the vulnerability to attack—incurred while upgrading them. Finally, the signature levels supported on the management console must remain synchronized with the signature packs on the sensors themselves.
- Group IDS sensors together under a few larger profiles. Every signature upgrade requires that all new signatures be appropriately tuned on every sensor. Tuning signatures for groups of sensors rather than for each sensor on the network significantly reduces configuration time. This administrative advantage must be balanced against the ability to finely tune sensor configuration by establishing a separate profile for each sensor.

Because different sets of sensors will see different traffic crossing the switches to which they are attached, separate signature profiles should be configured for each set of sensors. Each profile can then be tuned to generate alarms based on the traffic types seen (connection signatures), the attack signatures (exploit signatures), and the specific traffic (string signatures) that is relevant to that particular set of sensors. If more than one set of sensors will see the same traffic types, then the same signature profile may be used for both sets.

Another best practice is to place the configuration management function and the event monitoring and reporting function on separate servers for greater scalability. The event monitoring function and the reporting function can also be relegated to separate servers to achieve additional scalability.

Event Monitoring and Management

Current network IDS sensors from Cisco can generate up to approximately 25,000 alarms per minute in short bursts. Current IDS management consoles can sustain up to approximately 300 alarms per second. The number of sensors that should be forwarding alarms to a single IDS management console is a function of the aggregate number of alarms per second generated by those sensors. Experience with customer networks has shown that the number of sensors reporting to a single IDS management console should be limited to 25 or fewer. These customers use the default signature profiles, but also tune some specific signatures. The number of alarms generated by each sensor is determined by how sensitively the sensor is tuned—that is, by how carefully the IDS sensor is tuned to the particular applications on a given server and to the expected traffic patterns on the network segment being monitored. The more



sensitive the tuning, the fewer the alarms generated and the larger the number of sensors that can report to a single IDS management console. It is essential to tune out false positives in order to maximize the scalability of the network IDS deployment. Sensors that are expected to generate a large number of alarms, such as those sitting outside the corporate firewall, should log to a separate IDS management console, because the number of false alarms raised increases the noise-to-signal ratio dramatically and makes it difficult to identify otherwise valid events.

When implementing multiple IDS management consoles, implement either separate monitoring domains or a hierarchical monitoring structure.

Separate Monitoring Domains

Separate security monitoring domains can be used when separate security operations groups are responsible for different geographic areas or business units. In this implementation, there is no ability to monitor real-time activity across the entire enterprise: sensors send alarms only to the IDS management console in their geographic area or business unit. To offset lost functionality, implement a separate, centralized system for trend analysis and reporting, and have all IDS management consoles forward events to this system. This architecture is similar to a Manager of Managers (MOM) architecture.


Hierarchical Monitoring Structure

An effective hierarchical monitoring structure requires an alarm or event policy to distinguish and identify those alarms requiring a local, regional, or corporate-wide response. Local alarms indicate a small-scale, localized attack against a branch network or remote office. Regional alarms indicate an attack against several branch networks, telecommuters, or remote-office networks within a given geographic region. Regional incidents can be escalated to a more corporate-wide audience should it be determined that additional resources are necessary. Corporate-wide incidents represent a broad, enterprise-wide attack from one or more sources. In the latter situation, an enterprise security incident response team must coordinate the response. Local security personnel in regional networks may require direction to effectively coordinate resources to contain the various incidents and restore overall network integrity.

In a three-tier hierarchical monitoring structure, local IDS management consoles are implemented at each branch, regional, and corporate site. IDS sensors send alarms to their local IDS management console. Lower-priority alarms are analyzed by local security personnel. Medium-level alarms are sent up to a regional security monitoring site that provides additional expertise. High-priority alarms are sent up to an enterprise-wide security monitoring site where personnel with high-level security expertise reside. Figure 1 shows a three-tier, hierarchical IDS monitoring structure.



Figure 1
Three-Tier, Hierarchical IDS Monitoring Structure



The IDS management consoles at the local network level should forward alarms to a separate, centralized system for trend analysis and reporting.

Scaling Host-Based IPSs

Host-based intrusion prevention scales like network IDS implementations: a central management console is deployed and used to maintain a database of policies as well as system nodes, all of which have the IPS agent installed. Also like IDS implementations, host-based IPS agents on similar systems should be grouped together to streamline the process of assigning policies. Before placing hosts into groups, the security needs of those hosts should be analyzed and a security plan should be mapped out. Grouping hosts together has the added benefit of applying a consistent set of policies across multiple host systems. The following are criteria to consider when grouping hosts together:

- System functions
- Business needs
- Geographical or topological location
- Organizational value—Servers that perform mission-critical roles benefit from being grouped together even if they perform different functions: their value to the enterprise is great and they may therefore warrant special policies that are not applicable on other systems.

The same multi-tier architecture deployed for network-based IDSs can be used with host-based IPSs.

Scaling Logging

Logging provides an audit trail for network events. Network operations personnel and security personnel can identify problems and track down when or where an attack originated. For a detailed discussion of logging and the syslog facility in particular, please refer to Appendix B, “Syslog Primer.”

Scaling logging is problematic because most security systems use the syslog protocol to generate log event messages and the syslog protocol uses the connectionless User Datagram Protocol (UDP) as the message transport mechanism. The SAFE enterprise white paper and the SAFE white paper for small, midsize, and remote-user networks call for a



central log server in the management module or the management section, respectively, of the network. However, when devices in remote network modules generate log messages, these messages must traverse WAN or VPN links to the log server in the management module or section of the central network. Unnecessary bandwidth use may result, and log messages may be lost due to link outages.

To scale logging more appropriately, an additional log server should be deployed in each branch network. These branch log servers should be configured to forward log event messages to the central log server. If a WAN or VPN link is interrupted, the local log server can capture the log data for later review by network operations or security personnel. If bandwidth use across a WAN or VPN link is a problem, the local log servers can be configured not to forward log event messages. However, this will require that the monitoring and investigation of logs be performed on the local branch log servers should an event be detected that requires further analysis.

Various open source and proprietary tools can be used to convert Microsoft event logs into syslog formats as well as to transmit the information in those logs to a local or remote syslog server. One logging standard should be adhered to within a network. The network administrator can select either the Microsoft or the syslog standard.

Data Correlation

Network IDSs and host-based IPSs identify and respond to network attacks; syslog information from routers, switches, and firewalls provides a record of attack paths. Syslog can also identify targets that an attacker compromised by means not identified by the network IDS. In addition, systems that cannot support IPS software agents may be monitored through syslog messages. These messages must be sent securely to logging hosts where the information can be collated and analyzed. Logging hosts should be located within the OOB management network and should be part of the overall event analysis console. Log data, however, does not provide an efficient, real-time analysis method for ongoing network attacks.

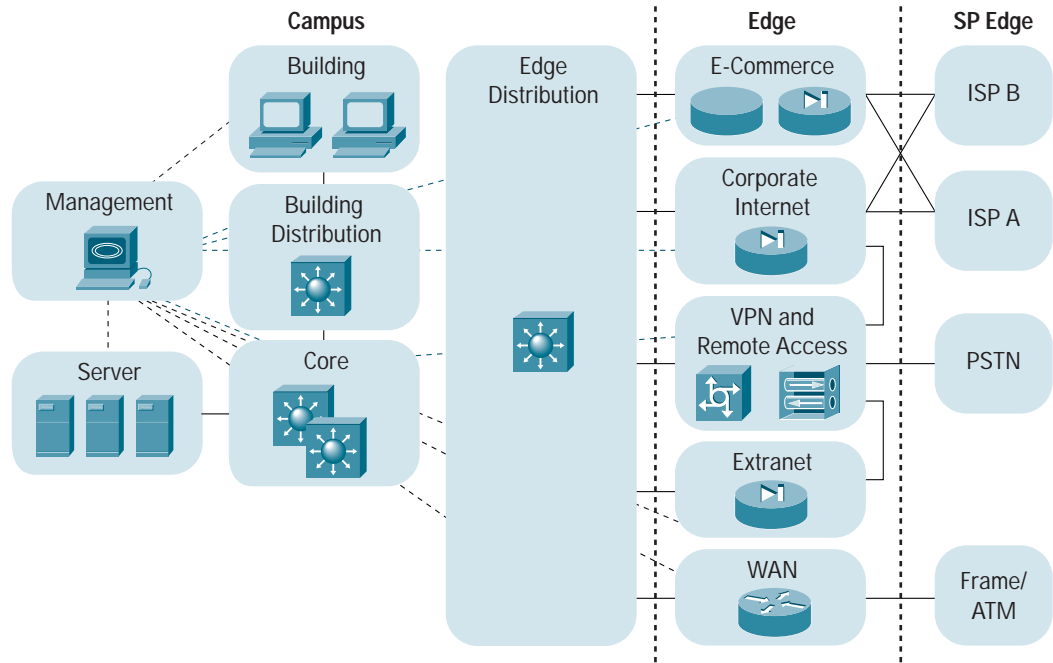
Data correlation between the network IDS and the host-based IPS provides an accurate picture of a network attack. The NIDS identifies the source, destination, and signature of the attack; the IPS validates that the attack reached the intended target and determines what the final impact was.

Large, Single-Site Design

The large, single-site design utilizes the network design from the SAFE enterprise security paper (Figure 2).



Figure 2
Enterprise Network Modules



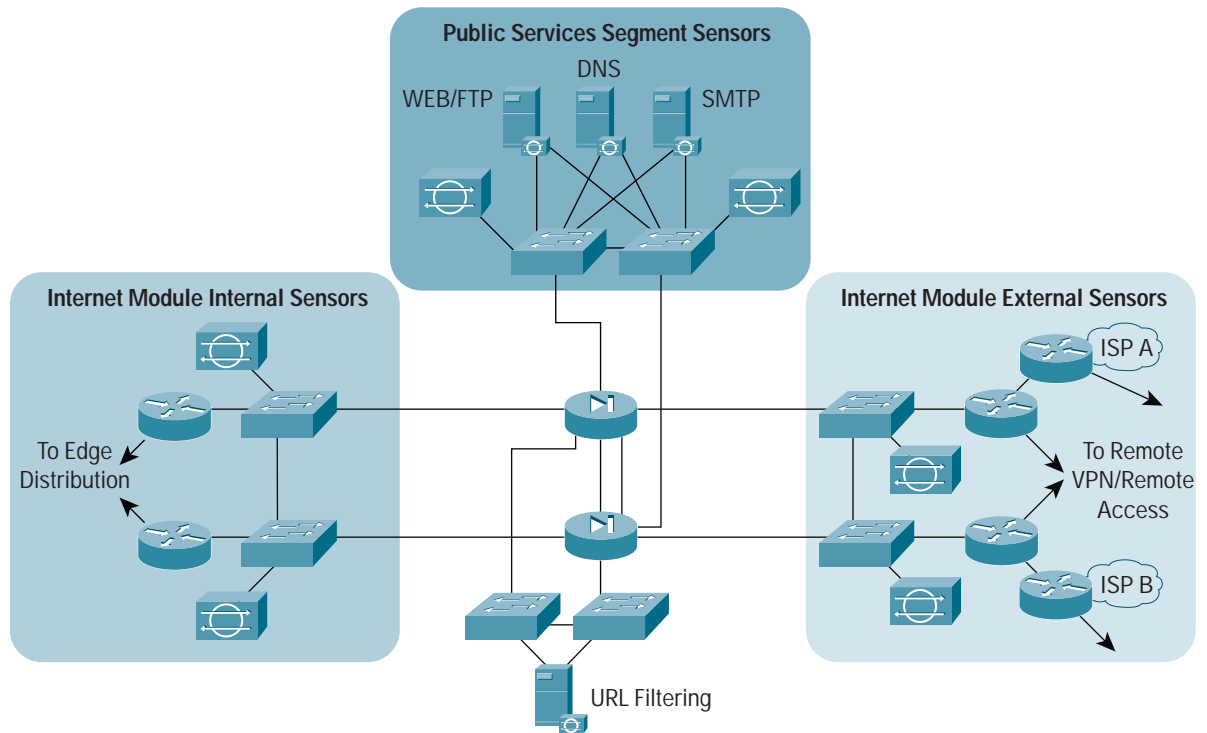
The SAFE blueprint recommends placing network IDS sensors in the Corporate Internet Module, the VPN/Remote Access Module, the Building Distribution Module, the Server Module, the Management Module, the Extranet Module, and the Data Center/E-Commerce Module. The sensors are deployed in redundant sets in each module, with the exception of the Management Module. Redundant sets help ensure that if one of the switches fails and traffic is routed to the other switch, there is a sensor attached that will still monitor network activity. The following sections detail each of the modules.

Corporate Internet Module

Figure 3 shows the recommended placement of network IDS sensors in the Corporate Internet Module. Host-based IPS agents are also deployed to protect both servers in the Public Services Segment and the URL filtering server. All network devices (routers, switches, and firewalls) are set up to log syslog messages to servers in the Management Module. The network IDS sensors are connected to a TVR server in the Management Module to reduce the number of false alarms generated by IDS sensors in this module.



Figure 3
Corporate Internet Module IDS Design



Key Intrusion Detection and Syslog Devices

- *Network IDS appliances*—Provide Layer 4–7 monitoring of key network segments, active session resetting (TCP resets), and coordination of the blocking of malicious IP addresses (shunning)
- *Host-based IPSs*—Provide monitoring and intrusion prevention for the operating system and Web server application of key servers in the module, DoS protection, and attack mitigation at the endpoint host
- *Stateful firewall*—Provides syslog messaging regarding attempted access-control violations and administrative access as well as shunning
- *Routers*—Provide flow control as well as syslog messaging regarding attempted access-control violations and administrative access
- *Layer 2 switches*—Provide syslog messaging regarding administrative access

Threats Mitigated by Intrusion Detection and Syslog

- *Application layer attacks*—These attacks are mitigated through IDS at the host and network levels.
- *Viruses and Trojan horses*—These attacks are mitigated through e-mail content filtering and the host-based IPS.
- *Password attacks*—Limited services are available to brute force attacks, and the operating system and host-based IPS can detect the threat.
- *Packet sniffers*—A switched infrastructure and the host-based IPS limit exposure.
- *Network reconnaissance*—The IDS detects recon, and protocols are filtered to limit effectiveness.



- *Port redirection*—Restrictive filtering and the host-based IPS limit attacks.
- *Worms*—These attacks are mitigated through content filtering and host-based IPS

Detailed Design

The Corporate Internet Module IDS design comprises three pairs of redundant sensors: the Intranet Module External Sensors, the Public Services Segment Sensors, and the Intranet Module Internal Sensors. The tuning, alarm logging, and attack response of each pair are detailed in the sections below.

Internet Module External Sensors

The Internet Module External Sensors sit between the routers connected to the Internet service providers (ISPs) and the external interfaces of the redundant firewall pair (Figure 3).

Objective

These sensors monitor attempts to breach the external firewall and also determine what attack types are attempted. This is the front line for monitoring the spread of an attack or contagion originating from the Internet.

The traffic and attack types that these sensors see are dependent upon whether there is protocol filtering at the routers connected to the ISPs. There is often only RFC 1918 and RFC 2827 filtering, to mitigate source address spoofing, so a wide range of protocols and attacks could be seen.

Tuning and Alarming

If protocol filtering is not performed by routers connected to the ISPs or at the enterprise edge, then tuning based on connection signatures (specific protocol types) cannot be implemented effectively. All traffic on this segment should have a registered IP address due to the filtering at the routers connected to the ISPs and the Network Address Translation (NAT) pools in the firewalls. Tune these sensors to alarm on any traffic seen with RFC 1918 addressing. Unless specific signatures matching specific attacks are turned off by default, leave the exploit signature tuning at default settings to provide a broad level of analysis outside the firewall.

Most alarms seen within network IDS implementations come from sensors outside the firewall. Most will be eliminated by the TVR system; log those that remain to a separate network IDS console to separate the important alarms that other sensors generate from the alarms that these sensors generate. This may be a good place to practice reporting by exception, in which the reporting system provides information when the number and type of attacks deviate from some statistical norm. Because these sensors are outside the Corporate Internet Module firewall, specific attack signatures seen here do not necessarily indicate that corporate resources were compromised.

Response

Generally, avoid automatic address blocking on these sensors. Although e-commerce is not a function of the Internet Module, the source address of an attacker may be spoofed. Therefore, blocking could prevent legitimate attempts to reach the servers in the Public Services Segment. TCP resets may be used; they are less disruptive than blocking, but have not been configured within the SAFE lab network reference implementation. Do not send TCP resets to known, valid devices such as network infrastructure components and extranet partners. Log IP session data only if a specific attack signature(s) is seen. Otherwise, the volume of information generated may overwhelm security operations personnel. Concentrate on logging IP packets for attacks that have breached the firewall.



Alternatives

These sensors may generate numerous alarms. If there is no time to analyze the data or no reporting system to provide detailed information, do not expend resources deploying sensors in this location.

Public Services Segment Sensors

The Public Services Segment Sensors sit on the network segment that contains the public Web, FTP, DNS, and mail servers (Figure 3).

Objective

These sensors detect both specific attacks on public service servers and the traffic types (protocols) that should not be seen on this segment.

Tuning and Alarm Logging

Tune these sensors to alarm if any traffic that should not be seen on the Public Services Segment is detected (connection signatures). Review and tune as necessary the signatures that are specific to vulnerabilities in Web, DNS, FTP, and mail servers, even though these servers may have the latest security patches and should be running host-based IPSs. In general, leave the signatures tuned at the default levels, but implement custom string match signatures for vulnerabilities such as Code Red and Nimda that are not covered in the default attack signatures. Alerting on specific connection signatures, general attack signatures, and specific string signatures provides focused segment analysis at Layers 4 through 7.

The number of attack signatures seen on this segment varies. A Code Red attack, for example, can generate a large number of alarms. If the infected servers were patched for the vulnerability targeted in this attack, the attack can be tuned to alarm at a lower priority or not at all to avoid overwhelming the alarm console and security personnel. A TVR system also reduces the workload that these types of attacks generate so that other attacks can be identified and controlled.

Attack Response

Generally, avoid automatic blocking on these sensors. Again, although e-commerce is not a function of the Internet Module, there is a possibility that the source address of an attacker has been spoofed. Automatic blocking could prevent legitimate attempts to reach the Public Services Segment. Instead, manually reconfigure the firewall or router connected to the ISP to block a particular source address. Because they are less disruptive than automatic shunning, TCP resets may be used. Resets have been configured on these sensors within the SAFE lab reference network. Please see the cautions regarding implementing TCP resets under “Response Actions” in Appendix A.

Because the attack signatures that these sensors see indicate that the firewall was penetrated, IP session data should be logged. Implement logging for a small percentage of signatures to minimize the data generated.

Alternatives

NIDS sensors should be deployed here. If they are not, visibility into some network-level attacks will be lost, and host-based IPSs should be installed on the public services servers to provide a degree of protection.



Internet Module Internal Sensors

The Internet Module Internal Sensors sit on the switches connected to the inside interface of the redundant firewall pair (Figure 3).

Objective

Sensors in this location monitor attacks that breach the external firewall. Correlate data from these sensors with that of other sensors located throughout the network to track the spread of an attack or contagion.

Tuning and Alarm Logging

The traffic that these sensors see is dependent upon the corporate security policy. For example, if the corporate security policy provides employees with open access to the Internet, then very little tuning based upon traffic type (connection signatures) can be performed on these sensors. However, if the corporate security policy is more restrictive, then tuning based upon traffic type is possible.

These sensors should see few attacks, and the attack signatures can generally be left tuned at the default levels. Attacks seen here have breached the external firewall, however, or are being generated from within the internal network, and should therefore be taken seriously. Again, during incidents such as a Code Red attack, it may be necessary to tune specific signatures at a lower level in order to keep from being overwhelmed by high-priority alarms.

Attack Response

Address blocking and TCP resets can and probably should be used here. Shunning can be used because the traffic crossing these sensors is generated or terminated by users on the internal corporate site. Thus, although the source address of the attacker may be spoofed, shunning will generally not bring down any business-critical applications.

Attacks seen by these sensors can be generated from sources internal or external to the company, so the source must be determined quickly. Block external sources at the firewall. Track down and isolate internal sources, possibly using ACLs within the distribution layer of the SAFE architecture.

Because the attack signatures that these sensors see indicate that the firewall was penetrated, IP session data should be logged. Implement logging for a small percentage of signatures to minimize the data generated.

Alternatives

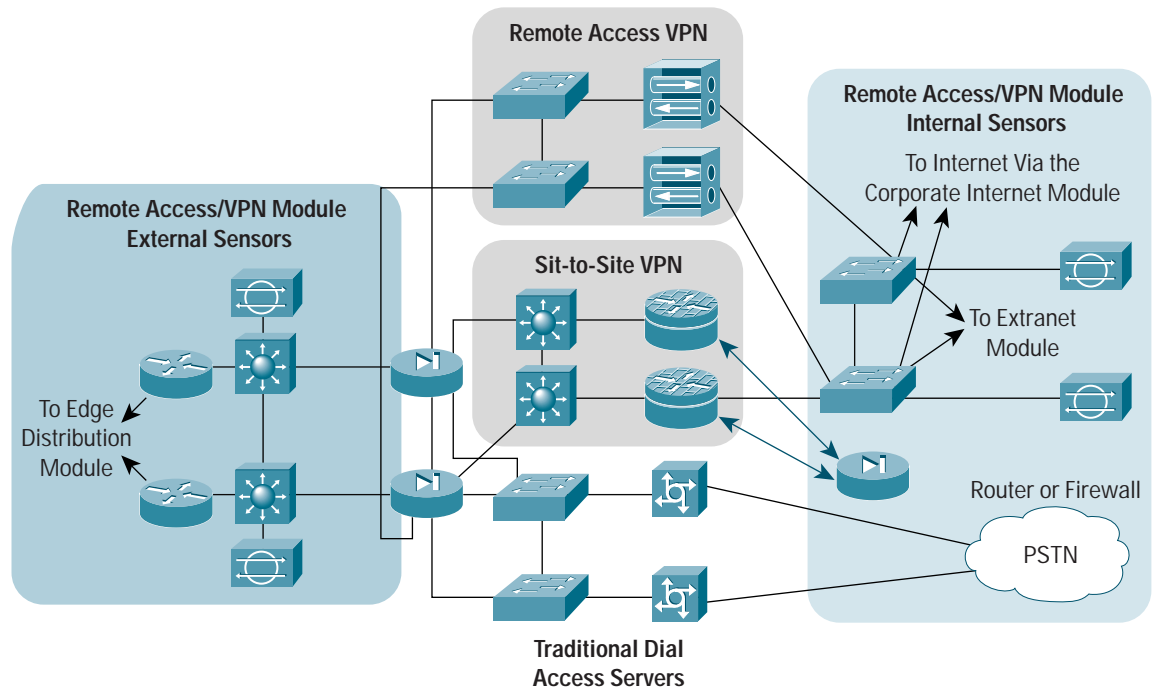
There are no alternatives to placing sensors here. If limited resources are available for implementing sensors, this is one of the locations where they should be deployed.

Remote Access/VPN Module

Figure 4 shows the recommended placement of network IDS sensors in the Remote Access/VPN Module. Because there are no servers in the module, there are no host-based IPS agents. All network devices (routers, switches, firewalls, and VPN concentrators) are set up to log syslog messages to servers located in the Management Module.



Figure 4
Remote Access/VPN Module IDS Design



Key Intrusion Detection and Syslog Devices

- *Network IDS appliances*—Provide Layer 4–7 monitoring of key network segments, active session resetting (TCP resets), and coordination of the blocking of malicious IP addresses (shunning)
- *Host-based IDSs*—Provide monitoring and intrusion prevention for the operating system and Web server application of key servers within the module
- *Stateful firewalls*—Provide syslog messaging regarding attempted access-control violations and administrative access as well as shunning
- *Routers*—Provide syslog messaging regarding attempted access-control violations and administrative access
- *Layer 2 switches*—Provide syslog messaging regarding administrative access
- *VPN concentrators*—Provide syslog messaging regarding administrative access

Threats Mitigated by Intrusion Detection and Syslog

- *Application layer attacks*—These attacks are mitigated through IDS at the network level.
- *Password attacks*—Limited services are available to brute force attacks, and the operating system and IDS can detect the threat.
- *Network reconnaissance*—The IDS detects recon, and protocols are filtered to limit effectiveness.
- *Worms*—The IDS mitigates infection of remote systems that then may transmit the worm into the enterprise network.



Detailed Design

The Remote Access/VPN Module network IDS design comprises two pairs of redundant sensors: the Remote Access/VPN Module External Sensors and the Remote Access/VPN Module Internal Sensors. The tuning, alarm logging, and attack response of each pair are detailed in the sections below.

Remote Access/VPN Module External Sensors

The Remote Access/VPN External Sensors sit between the routers connected to the ISPs and the external interfaces of the VPN concentrators of the Remote Access/VPN Module (Figure 4).

Objective

Sensors in this location help ensure that the only protocols seen are those that are specifically allowed through the router connected to the ISPs. Router filtering should limit any network reconnaissance. Internet Key Exchange (IKE [UDP 500]), Encapsulating Security Payload (ESP [IP 50]), and a user-specified UDP port or other protocol used to encapsulate IP Security (IPSec) traffic when NAT is being performed between VPN peers are the specific protocols that should be seen. Internet Control Message Protocol (ICMP) traffic that is necessary for path maximum transfer unit (PMTU) discovery to work correctly should also be permitted.

Tuning and Alarm Logging

Generally, tune these sensors (using connection signatures) to alarm if traffic other than the protocols listed above is seen. All traffic on this segment should have a registered IP address due to the filtering at the routers connected to the ISPs and the addressing of the public interfaces of the VPN concentrators. Therefore, tune these sensors to alarm on any traffic seen with RFC 1918 addressing. The sensors should raise few alarms, because all traffic other than the initial IKE traffic will be encrypted.

Attack Response

Generally, avoid shunning on these sensors. The source address of an attacker may be spoofed, and shunning could block a remote peer of a site-to-site VPN from establishing a VPN connection. If shunning is configured, do not block the addresses of known site-to-site VPN peers. Shunning could also accidentally block access for remote-access VPN users. If dynamic addressing is used, reduce the shun time to a short period. The access lists of routers connected to the ISPs may have to be reconfigured if a particular source address needs to be blocked.

Avoid TCP resets for the reasons listed above.

If attempted reconnaissance or DoS attacks are seen on VPN termination devices, implement session logging based on those attacks to determine what the attacker intends. Use this information to reconfigure router ACLs.

Alternatives

Implementing sensors in this location may not be possible if there are limited resources to monitor the consoles and review the IDS reports. Minimal resources should be required to monitor activity on these sensors, however.

Remote Access/VPN Module Internal Sensors

The Remote Access/VPN Internal Sensors sit between the inside interface of the Remote Access/VPN Module firewalls and the Edge Distribution Module routers (Figure 4).



Objective

These sensors monitor the traffic to and from remote VPN locations. A network's security perimeter should extend to all remote sites, and it does extend when the remote site is a small or medium-sized office where physical security is maintained. However, for home-office locations and particularly when the remote user has split tunneling capabilities, this perimeter cannot always be guaranteed. It is therefore wise to monitor the traffic coming in from VPN connections.

Tuning and Alarm Logging

The protocols that these sensors see are dependent upon the security policy of the corporation for both site-to-site and remote-access VPNs. Companies that allow unrestricted access to corporate resources from VPN connections will not tune sensors by traffic type (connection signatures). Companies that restrict what VPN users can access, however, will make use of connection signatures. Generally, site-to-site VPN access within corporations is unrestricted.

These sensors should raise few alarms because all traffic should come from remote corporate users. However, if the physical security of the remote site in a site-to-site or remote-access VPN connection is suspect, or if the remote site has an insecure connection to the Internet, then these sensors may see attack signatures. The severity of the alarm depends upon the signature. For example, a port sweep coming from a remote site could be a remote user testing a scanner tool and may then be treated as a low- or medium-level alarm. However, a signature match for Code Red or Nimda may be treated as a high-level alarm, because it indicates that the remote site has a contagion that could spread to the corporate headquarters.

Attack Response

Deploy shunning on these sensors because they should see only internal traffic. Even if the source address of the attack is spoofed, shunning the remote address will result only in a single remote user being unable to access the corporate headquarters for a period of time. Given that the remote user may call into the network operations center to resolve the issue, the risk of firing on false positives is well worth the return. In addition, shunning information is useful when tracking down the source of the attack. For example, coordination of the shunning logs with the authentication, authorization, and accounting (AAA) server logs may indicate that the remote user's password has been compromised. Changing the password or implementing a one-time password server will mitigate future attacks. In the case of a worm or virus, updating the signature pack or creating a custom signature will mitigate attacks. Deploying TCP resets on these sensors is also recommended. Follow the guidelines for tuning the sensors in the "SAFE Axioms" section of this paper before implementing either shunning or TCP resets. Finally, set up session data logging as needed if attack signatures are coming in from VPN connections.

Alternatives

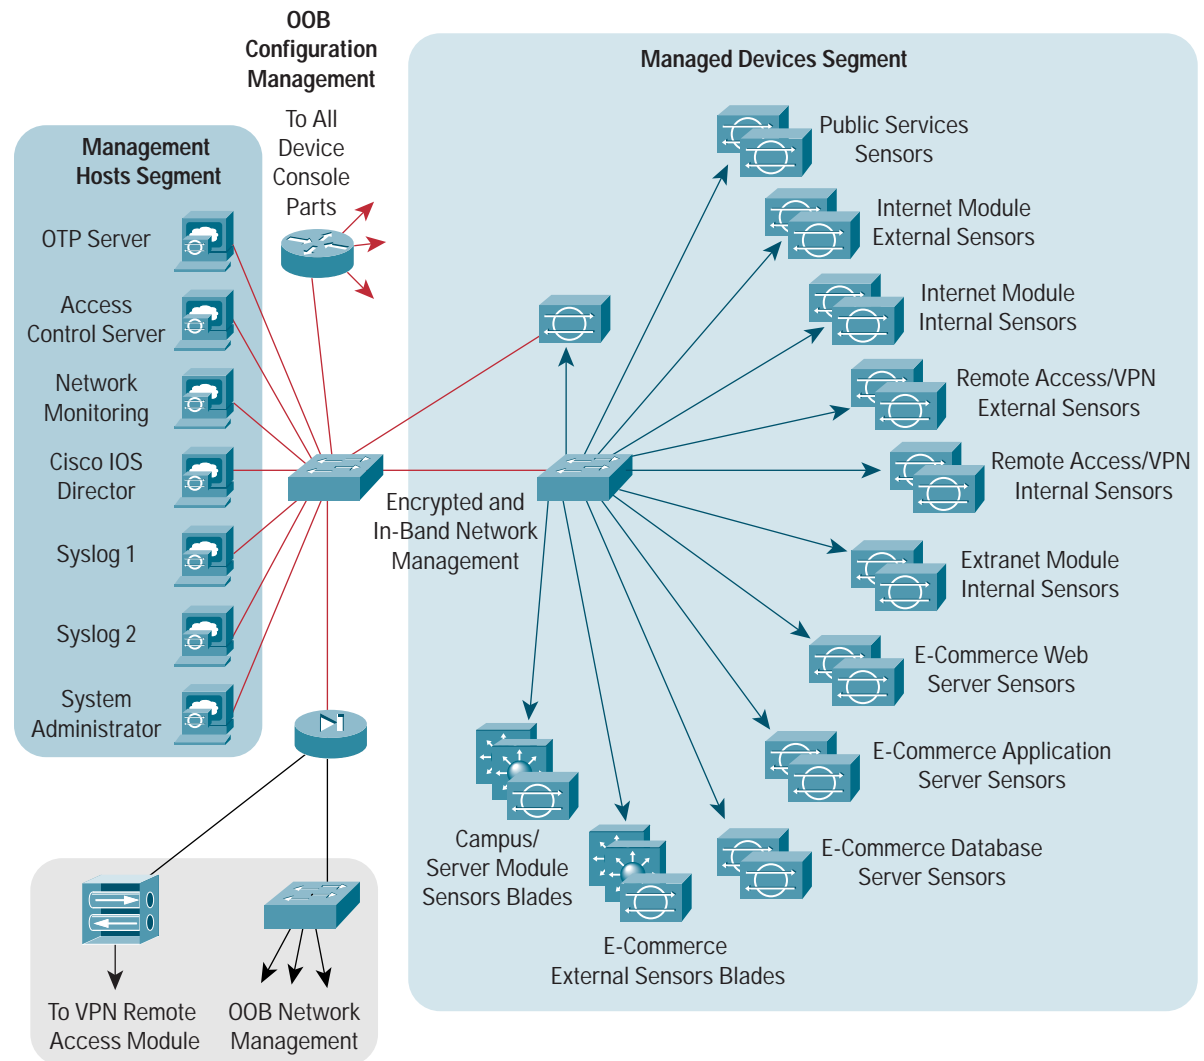
The alternative to placing sensors in this location is to implement them at all remote site-to-site VPN locations and to backhaul alarm traffic across the VPN connection. For larger, regional offices, this design is recommended; it is detailed later in this document. However, for small or home offices, this design may be costly and difficult to manage, particularly in networks of hundreds or thousands of small offices. Placing network IDS sensors just at this location, however, provides a view of attacks coming in from all of the VPN connections.



Management Module

Figure 5 shows the placement of the Management Module Sensor that is used to protect all management hosts within the Management Module. A single sensor is used because the Management Hosts Segment does not use redundant switches.

Figure 5
Management Module IDS Design



In addition to network IDS sensors, host-based IPS agents are deployed to protect servers located in the Management Hosts Segment. The TVR server is located in the Management Module and provides alarm analysis to reduce the number of false alarms as well as to escalate and validate high-priority alarms. All network devices (routers, switches, and firewalls) are set up to send syslog messages to the servers located in the Management Module, using interfaces located on the Managed Devices Segment. Figure 5 also shows how the management interfaces of all sensors deployed in the SAFE enterprise design connect back to the IDS console(s).



Key Intrusion Detection and Syslog Devices

- *Network IDS appliances*—Provide Layer 4–7 monitoring of key network segments, active session resetting (TCP resets), and coordination of the blocking of malicious IP addresses (shunning)
- *Network IDS consoles*—Provide centralized configuration management and viewing of alarms from network IDS sensors deployed across the network
- *Host-based IPS consoles*—Provide centralized configuration management and viewing of alarms from host-based IPS agents deployed on servers throughout the network
- *Threat Validation and Response Server*—Provides automated analysis of IDS alarms; decreases the number of false alarms and escalates high-priority alarms
- *Reporting servers*—Correlate events and reporting across syslog devices, network IDS sensors, and host-based IPS agents
- *Host-based IPSs*—Provide monitoring and intrusion prevention for the operating system and Web server application of key servers in the module
- *Stateful firewalls*—Provide syslog messaging regarding attempted access-control violations and administrative access as well as shunning
- *Routers*—Provide syslog messaging regarding attempted access-control violations and administrative access
- *Layer 2 switches*—Provide syslog messaging regarding administrative access

Threats Mitigated by Intrusion Detection and Syslog

- *Application layer attacks*—These attacks are mitigated through IDS at the host and network levels.
- *Viruses and Trojan horses*—These attacks are mitigated through host-based IPSs.
- *Password attacks*—Limited services are available to brute force attacks, and the operating system and IDS can detect the threat.
- *Packet sniffers*—A switched infrastructure and host-based IPS limit exposure.
- *Network reconnaissance*—The IDS detects recon, and protocols are filtered to limit effectiveness.
- *Port redirection*—Restrictive filtering and the host-based IPS limit attacks.

Detailed Design

Use a stateful firewall to limit traffic flow between the Managed Devices Segment and the Management Hosts Segment (Figure 5). In order for NIDSs, host-based IPSs, and logging to work, the firewall must allow TCP port 22 (Secure Shell [SSH] protocol), TCP port 443 (HTTPS), and UDP 514 (syslog) to pass traffic from the Managed Devices Segment to the Management Hosts Segment. Because both the Managed Devices Segment and the Management Hosts Segment are OOB segments that are isolated from regular data traffic, encryption between the consoles and the network IDS sensors, host-based IPS agents, and syslog devices is optional. If remote site management is included in the design, however, then encryption is strongly recommended. The tuning, alarm logging, and attack response of the Management Module Sensor are detailed in the section below.

Management Module Sensor

The Management Module Sensor sits with both the monitoring interface and the management interface on the Management Hosts Segment (Figure 5).



Objective

The Management Module Sensor detects all attacks on the management hosts. The following traffic is allowed on the Management Hosts Segment:

- *Trivial File Transfer Protocol (TFTP [UDP 69])*—For network device configuration files from devices on the Managed Devices Segment
- *FTP-Data (TCP 20)*—For file transfers to network devices on the Managed Devices Segment and for Internet downloads
- *FTP-Control (TCP 21)*—For file transfers to network devices on the Managed Devices Segment and for Internet downloads
- *Syslog (UDP 514)*—From network devices on the Managed Devices Segment
- *Telnet (TCP 23)*—To network devices on the Managed Devices Segment
- *SSH (TCP 22)*—To network devices on the Managed Devices Segment
- *Network Time Protocol (NTP [UDP 123])*—To synchronize the clocks of all network devices on the Managed Devices Segment
- *HTTP (TCP 80)*—To the Internet and from hosts on other segments to download the host-based IPS agent software
- *HTTPS (TCP 443)*—To network devices on the Managed Devices Segment and the Internet as well as between the host-based IPS Console and its agents
- *TACACS+ (TCP 49)*—For administrator authentication to devices on the Managed Devices Segment
- *RADIUS (UDP 1812/1813 authentication/accounting)*—For authentication of administrator remote-access VPN connections coming from the Remote Administration Segment
- *ICMP (IP Protocol 1)*—Echo request and response to reach network devices on the Managed Devices Segment and the Internet
- *DNS (UDP 53)*—For name translation services for management hosts as they access services on the Internet
- *Simple Network Management Protocol (SNMP [UDP 161])*—To query information from network devices on the Managed Devices Segment
- *SNMP-Trap (UDP 162)*—To receive trap information from network devices on the Managed Devices Segment

Tuning and Alarm Logging

Ideally, the Management Module Sensor is tuned to alarm if protocols other than those listed above are seen on the Management Hosts Segment. However, because both the number of servers and the applications deployed on those servers may change over time, implementing connection signatures on this segment can be difficult without having to constantly tune out false positives. All false positives must be filtered out, because a compromised server on the Management Hosts Segment could compromise all network devices. Because no attack signatures should be seen on this segment, investigate all signature matches as soon as possible.

Set up session logging in advance for high-priority alarms. If a specific attack signature is seen on the Management Hosts Segment, the log data may help to determine how compromised the targeted management server(s) was.

Session logging can also be set up as needed.




Alternatives

There are no alternatives to placing sensors here. If limited resources are available for implementing sensors, this is one of the locations where they should be deployed.

Extranet Module

Figure 6 shows the recommended placement of IDS sensors in the Extranet Module.

Figure 6
Extranet Module IDS Design



The Extranet Module IDS design comprises a pair of redundant sensors—the Extranet Module Sensors—that are placed on switches connecting the application servers to business partners who need network access. In addition to network IDS sensors, host-based IPS agents are deployed to protect the extranet servers. All network devices (routers, switches, and firewalls) are set up to send syslog messages to the servers located within the Management Module.

Key Intrusion Detection and Syslog Devices

- *Network IDS appliances*—Provide Layer 4–7 monitoring of key network segments, active session resetting (TCP resets), and coordination of the blocking of malicious IP addresses (shunning)
- *Host-based IPSs*—Provide monitoring and intrusion prevention for the operating system and Web server application of key servers in the module, DoS protection, and attack mitigation at the endpoint host
- *Stateful firewalls*—Provide syslog messaging regarding attempted access-control violations and administrative access as well as shunning
- *VPN concentrator routers*—Provide syslog messaging regarding attempted access-control violations and administrative access
- *Layer 2 switches*—Provide syslog messaging regarding administrative access



Threats Mitigated by Intrusion Detection and Syslog

- *Application layer attacks*—These attacks are mitigated through IDS at the host and network levels.
- *Viruses and Trojan horses*—These attacks are mitigated through host-based IPSs.
- *Password attacks*—Limited services are available to brute force attacks, and the operating system and IDS can detect the threat.
- *Packet sniffers*—A switched infrastructure and host-based IPS limit exposure.
- *Network reconnaissance*—The IDS detects recon, and protocols are filtered to limit effectiveness.
- *Port redirection*—Restrictive filtering and the host-based IPS limit attacks.

Detailed Design

The tuning, alarm logging, and attack response of the Extranet Module Sensors are detailed in the section below.

Extranet Module Sensors

The firewalls in Figure 6 serve as the VPN termination point for the site-to-site and remote-access VPN tunnels that business partners use to access application servers in the Extranet Module. The firewalls use stateful inspection to limit this access to specific protocols and addresses.

Objective

These sensors detect attacks coming from business partners' networks. Any attack signature seen indicate that a partner's network may be insecure or compromised, or that the partner's employees or contractors are attempting to exploit vulnerabilities in the application servers. Either situation is serious, because an outage caused by penetration of the Extranet Module application servers can affect the business relations of every partner with access to module resources.

Tuning and Alarm Logging

Because the firewalls are typically configured in a restrictive manner, granting access only to protocols that support the business applications, tune the Extranet Module Sensors to alarm when other traffic is seen. For example, if business partners require only a secure HTTP connection to the Extranet Module servers, then configure the firewalls to allow only TCP 443 traffic originating from the VPN tunnels onto the inside Ethernet segment. Tune the sensors to alarm on any other protocol. However, if the application servers must communicate with a back-end database server located in the Server Module, then tune the sensors not to alarm when the facilitating protocol is seen.

In addition, attack signatures specific to the types of applications running on the Extranet Module servers should be tuned appropriately if seen by the sensors. For instance, if a Windows-based Web server is running on the Extranet Module servers, then any signature that attempts to exploit vulnerabilities in IIS may be tuned to alarm at a high priority. Although the server should be both patched with the latest security fixes and running a host-based IPS, it is still wise to be alerted to exploit attempts. The new exploits that are discovered will encourage the system administrator to keep up with the latest security patches.



Attack Response

Generally, avoid automatic shunning on these sensors. With multiple business partners accessing application servers in the Extranet Module, an attack that spoofs the source address of a legitimate business partner could be disastrous if automatic shunning is deployed. For TCP-based applications, TCP resets, which are less disruptive, may be used. An alarm and possibly a TCP reset will alert the security operations group to reconfigure the firewall to prevent further attacks and to contact business partners to investigate the source of the alarm.

The secure business partner traffic on this extranet segment should raise few, if any, alarms. Set up session logging in advance for high-priority alarms. If an attack occurs, log data may help to determine how compromised servers on the segment were. Session logging can also be set up as needed.

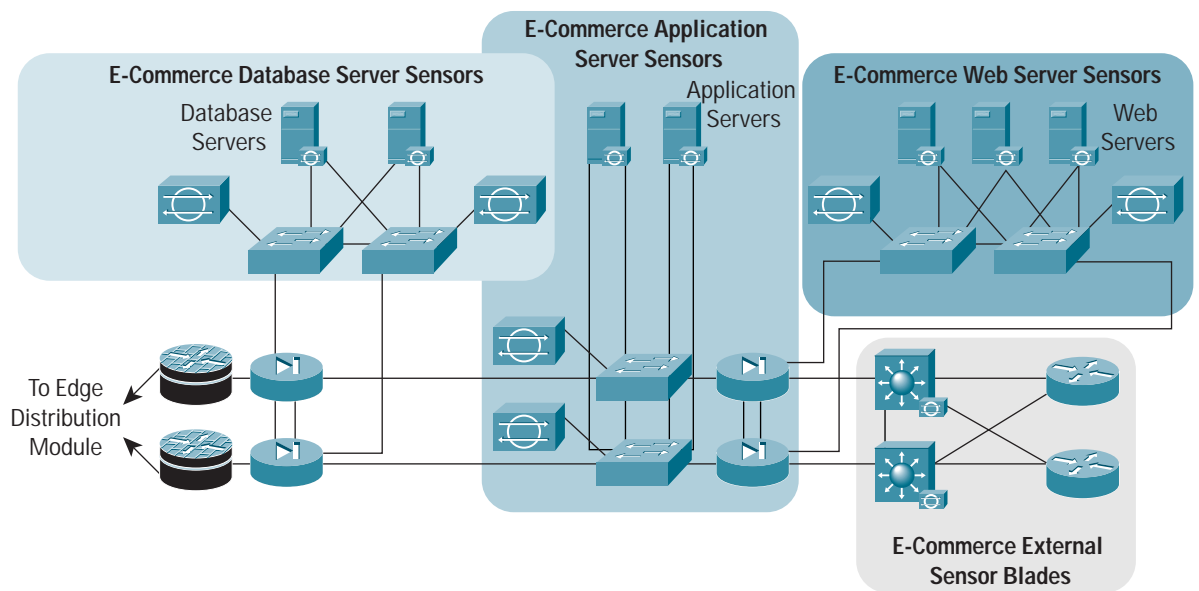
Alternatives

If network IDS sensors are not placed here, visibility into some network-level attacks will be lost. The application servers should have a host-based IPS installed to provide a degree of protection.

Data Center/E-Commerce Module

Figure 7 shows the recommended placement of IDS sensors in the E-Commerce Module.

Figure 7
E-Commerce Module IDS Design



The E-Commerce Module IDS design comprises four pairs of redundant sensors: the E-Commerce External Sensor Blades, the E-Commerce Web Server Sensors, the E-Commerce Application Server Sensors, and the E-Commerce Database Server Sensors. In addition to network IDS sensors, host-based IPS agents are deployed to protect the Web, application, and database servers. All network devices (routers, switches, and firewalls) are set up to send syslog messages to the servers located in the Management Module.



Key Intrusion Detection and Syslog Devices

- *Network IDS appliances*—Provide Layer 4–7 monitoring of key network segments, active session resetting (TCP resets), and coordination of the blocking of malicious IP addresses (shunning)
- *Layer 3 switches with IDS modules*—Provide Layer 4–7 monitoring of key network segments, shunning, and syslog messaging regarding attempted access-control violations and administrative access
- *Host-based IPSs*—Provide monitoring and intrusion prevention for the operating system and Web server application of key servers in the module, DoS protection, and attack mitigation at the endpoint host
- *Stateful firewalls*—Provide syslog messaging regarding attempted access-control violations and administrative access as well as shunning
- *Routers*—Provide syslog messaging regarding attempted access-control violations and administrative access
- *Layer 2 switches*—Provide syslog messaging regarding administrative access

Threats Mitigated By Intrusion Detection and Syslog

- *Application layer attacks*—These attacks are mitigated through IDS at the host and network levels.
- *Viruses and Trojan horses*—These attacks are mitigated through host-based IPSs. [Some “Threats” sections in this paper read “IDS” where highlighted in yellow; others read “IPS.” Is this an accurate depiction of the architecture or a discrepancy?]
- *Password attacks*—Limited services are available to brute force attacks, and the operating system and IDS can detect the threat.
- *Packet sniffers*—A switched infrastructure and host-based IPS limit exposure.
- *Network reconnaissance*—The IDS detects recon, and protocols are filtered to limit effectiveness.
- *Port redirection*—Restrictive filtering and the host-based IPS limit attacks.

Detailed Design

The tuning, alarm logging, and attack response of each pair of sensors are detailed in the sections below.

E-Commerce External Sensor Blades

The E-Commerce External Sensor Blades sit between the routers connecting to the ISP and the first set of firewalls within the E-Commerce Module (Figure 7). In a traditional, three-tier, e-commerce architecture, only Web (HHTTP and HTTPS) traffic crosses these sensors.

Objective

These sensors monitor who is attempting to breach the external firewall and determine the types of attacks. This is the front line for monitoring the spread of a contagion or attack originating from the Internet.

Tuning and Alarm Logging

If basic ACLs limiting inbound traffic to HTTP (TCP 80) and HTTPS (TCP 443) are implemented on the routers connecting to the ISP, then tune the sensors to alarm if any other protocols (connection signatures) are seen. The administrator may also allow ICMP request and response traffic to test reachability from the Internet. If access lists are not implemented on the routers connecting to the ISP, then no such tuning is recommended. Because RFC 1918 addresses should not be seen on this segment, tune the sensors to alarm if they are seen.



Because Web server access is this module's primary business requirement, the E-Commerce External Sensor Blades can be tuned to alarm at a high priority if attack signatures specific to Web server vulnerabilities are seen. Tune other signatures for the environment or leave them at default levels. The number of alarms could be high, depending on the level of filtering the routers connected to the ISP perform. Alarms can be logged to a separate console to separate the noise that the sensors generate from legitimate attacks that penetrate the firewalls.

Attack Response

Generally, avoid automatic shunning on these sensors. Because multiple customers may access the Web servers, an attack that spoofs the source address of a legitimate business customer or ISP proxy server could be disastrous if automatic shunning is deployed. TCP resets also cannot be implemented within this type of deployment. Log IP session data only if a particular attack signature is seen and forensic analysis is desired. Otherwise, IP logging could generate excessive traffic. IP data can be logged on sensors on DMZs or internal interfaces of the firewall, because these attacks have breached the firewall.

This may be a good place to practice reporting by exception, in which the reporting system provides information when the number and type of attacks deviate from some statistical norm.

Alternatives

These sensors may generate numerous alarms. If there is no time to analyze the data or no reporting system to provide detailed information, do not expend resources deploying sensors in this location.

E-Commerce Web Server Sensors

The E-Commerce Web Server Sensors sit on a DMZ segment on the first set of firewalls within the E-Commerce Module (Figure 7). Again, in a traditional, three-tier, e-commerce architecture, Web (HTTP and HTTPS) traffic and the protocols used between the Web server and the back-end application servers are all that cross these sensors. ICMP echo request and response traffic may be allowed in order to test reachability from the Internet.

Tuning and Alarm Logging

Implement access control on the firewall to limit traffic on this segment to HTTP (TCP 80) and HTTPS (TCP 443) and the protocol(s) between the Web servers and application servers. Tune the sensors to alarm if any other protocols are seen. Depending on whether NAT is implemented in the firewall, RFC 1918 addresses may or may not be used for the Web servers. If the servers are configured for non-RFC 1918 addressing, then tune them to alarm at medium priority if any RFC 1918 addresses are seen. Finally, because Web server access is this module's primary business requirement, these sensors may be tuned to alarm at a high priority if attack signatures specific to Web server vulnerabilities are seen. Tune as necessary other signatures such as those specific to application server vulnerabilities.

The number of alarms that these sensors generate could vary. During a Code Red attack, for example, the sensors could generate numerous alarms while infected servers attempt to connect to e-commerce servers. If the servers were patched, are running host-based IPSs, and are immune to such attacks, the attack can be tuned to alarm at a lower priority or not at all so that other attacks can be identified and controlled.



Attack Response

Generally, avoid automatic shunning on these sensors. Because multiple customers may access the Web servers, an attack that spoofs the source address of a legitimate business customer or ISP proxy server could be disastrous if automatic shunning is deployed. TCP resets, which are less disruptive, may be deployed here, but only after false positives are tuned out that could disrupt access to the Web servers.

Alternatives

Sensors should be placed here. If they are not, visibility into some network-level attacks will be lost; host-based IPSs should be installed on the application servers to provide a significant degree of protection.

E-Commerce Application Server Sensors

The E-Commerce Application Server Sensors sit on a segment between the first set and second set of firewalls within the E-Commerce Module (Figure 7). In a traditional, three-tier, e-commerce architecture, the only traffic crossing these sensors will be the protocols used between the Web server and the back-end application servers and the protocols used between the application servers and the database servers.

Tuning and Alarm Logging

Implement access control on both sets of firewalls to limit traffic on this segment to the protocols required between the application servers and the database servers. Tune the sensors to alarm if any other protocols are seen. Tune to the environment the signatures specific to vulnerabilities in the application server and database server. The sensors should rarely alarm, so it may be appropriate to set the alarm level to a high priority.

Attack Response

Avoid automatic shunning on these sensors. Because the addresses of the Web servers and database servers that are communicating with the application servers are known and fixed, shunning these addresses effectively cuts off access between the Web servers and the application servers or between the database servers and the application servers. TCP resets, which are less disruptive, may be deployed here. For example, if a Web server is compromised and is attempting malicious activity over its connection with a back-end application server, then resetting the particular TCP connection may protect the application server during the time that it takes the security operations personnel to identify the Web server and remove it from service.

Reporting

The sensors should see few attacks because the firewall provides access control; attacks seen have bypassed the firewall. Whether the attack affects the servers depends on what operating system upgrades and security patches were applied, what application security patches were applied, and whether host-based IPSs are running on the servers. Because few alarms are generated, investigate any alarms that are seen.

Because the attack signatures that the sensors see indicate that the firewall was penetrated, IP session data should be logged. Implement logging for a small percentage of signatures to minimize the data generated.

Alternatives

There are no alternatives to placing sensors here. If deploying sensors is not practical, then the application servers should run host-based IPSs. Visibility into some network-level attacks will be lost, but the IPSs will provide significant protection.



E-Commerce Database Server Sensors

The E-Commerce Database Server Sensors sit on a DMZ segment on the second set of firewalls within the E-Commerce Module (Figure 7). In a traditional, three-tier, e-commerce architecture, the only traffic crossing these sensors will be the protocols used between the application servers and the database servers.

Tuning and Alarm Logging

Implement access control on the firewalls to limit traffic on this segment to the protocols required between the application servers and the database servers. Tune the sensors to alarm if other protocols are seen. Tune to the environment the signatures specific to vulnerabilities in the database server. The sensors should rarely alarm, so it may be appropriate to set the alarm level to a high priority.

Attack Response

Avoid automatic shunning on these sensors. Because the addresses of the application servers that are communicating with the database servers are known and fixed, shunning these addresses cuts off access between the database servers and the application servers. TCP resets, which are less disruptive, may be deployed here with caution.

The sensors should see few attacks because both sets of firewalls provide access control; attacks seen have bypassed the firewalls. Whether the attack affects the servers depends on what operating system upgrades and security patches were applied, what application security patches were applied, and whether host-based IPSs are running on the servers. Because few alarms are generated, investigate any alarms that are seen.

Because the attack signatures that the sensors see indicate that the firewall was penetrated, IP session data should be logged. Implement logging for a small percentage of signatures to minimize the data generated.

Alternatives

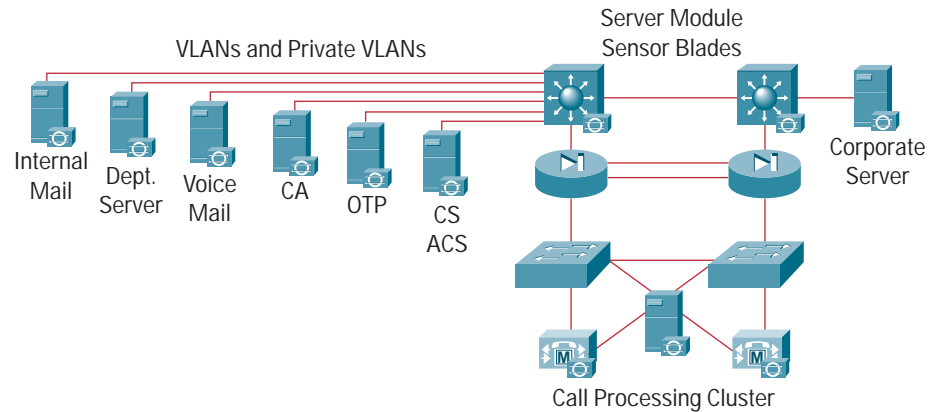
There are no alternatives to placing sensors here. If deploying sensors is not practical, then the application servers should run host-based IPSs. Visibility into some network-level attacks will be lost, but the IPSs will provide significant protection.

Server Module

Figure 8 shows the IDS sensor blades that are recommended in the Server Module. A single set of sensor blades located in the switches is recommended for all of the VLAN segments that house the corporate and department servers. Because a separate sensor is not deployed in each server VLAN, deployment costs are reduced. Host-based IPS agents are also deployed to protect the internal corporate servers. All network devices (routers, switches, and firewalls) are set up to send syslog messages to the servers located in the Management Module.



Figure 8
Server Module IDS Design



Key Intrusion Detection and Syslog Devices

- *Layer 3 switches with IDS modules*—Provide Layer 4–7 monitoring of key network segments, coordination of the blocking of malicious IP addresses (shunning), and syslog messaging regarding attempted access-control violations and administrative access
- *Host-based IPSs*—Provide monitoring and intrusion prevention for the operating system and Web server application of key servers in the module, DoS protection, and attack mitigation at the endpoint host

Threats Mitigated By Intrusion Detection and Syslog

- *Application layer attacks*—These attacks are mitigated through IDS at the host and network levels.
- *Viruses and Trojan horses*—These attacks are mitigated through host-based IPSs.
- *Password attacks*—Limited services are available to brute force attacks, and the operating system and IDS can detect the threat.
- *Packet sniffers*—A switched infrastructure and host-based IPS limit exposure.
- *Network reconnaissance*—The IDS detects recon, and protocols are filtered to limit effectiveness.
- *Port redirection*—Restrictive filtering and the host-based IPS limit attacks.

Detailed Design

The tuning, alarm logging, and attack response of the Server Module Sensor Blades are detailed in the section below.

Server Module Sensor Blades

Objective

The Server Module Sensor Blades monitor traffic across all VLAN segments, which house the corporate and department servers. Correlate data from these blades with that of other sensors located throughout the network to track the spread of an attack or contagion. Access lists on the SAFE design distribution layer can stop the spread of the attack.



Tuning and Alarm Logging

Configure the switch that houses the Server Module Sensor Blades so that only traffic from the VLANs that contain corporate and department servers is sent to the blades. Because there may be several VLANs and the server applications on the VLANs may be numerous, it may not be possible to tune the sensor blades to alarm if they see protocols other than those that are required for operation of the applications (connection signatures). This level of tuning also requires the switch to have ACLs that limit the protocols on the segments. Tuning attack signatures to alarm at specific levels may not be appropriate for the same reasons. The Cisco default signature profile is therefore recommended. The sensors should generate few alarms once the false positives caused by background protocols and services running on the servers are tuned out.

Attack Response

Implement automatic shunning for high-priority attacks, because all traffic seen here should be internal. However, choosing the VLAN interface on which to implement the shun could be difficult, depending on the source of the attack. If shunning is not appropriate, then configure TCP resets on the sensor blades. As an alternative, manually reconfigure the ACLs on the switches.

Reporting

The sensors should see few attacks because the traffic is internal. Set up IP session logging in advance based on the high-priority alarms seen here. If a specific attack signature is seen, the log data may help to determine how compromised the servers were. Session logging can also be set up as needed.

Alternatives

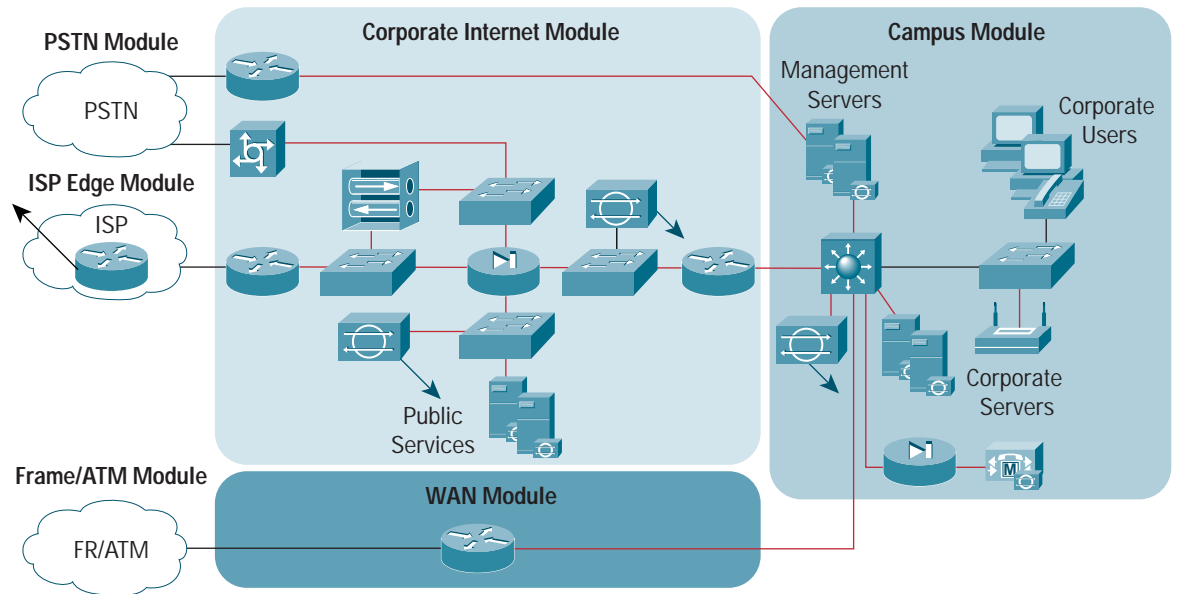
There are no alternatives to placing sensors here. If limited resources are available for implementing sensors, deploy sensors here because they can be used to determine if an attack or contagion has spread to the internal corporate servers. In addition, deploy host-based IPSs on the corporate servers.

Medium, Single-Site Design

The medium, single-site design uses the network design from the SAFE white paper for small, midsize, and remote-user networks (Figure 9).



Figure 9
SAFE Medium Network Design



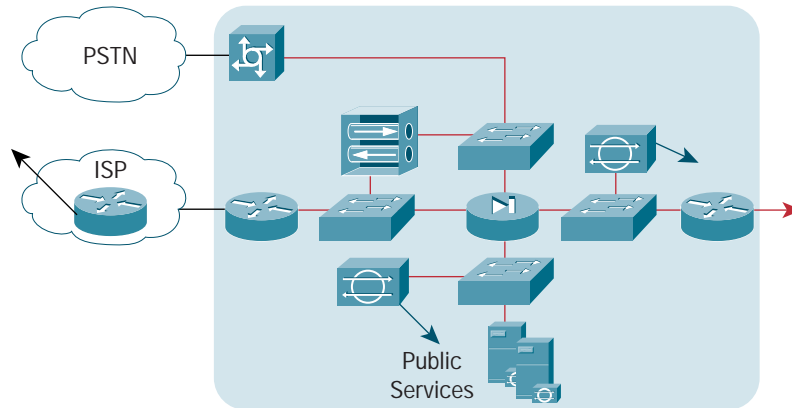
The SAFE Medium Network Design eliminates the redundancy that is present in the Enterprise Network Design: NIDS sensors are placed only at critical exposure points in the Corporate Internet Module and the Campus Module. If the switches providing connectivity between the core network and those areas fail, then the IDS devices can no longer monitor the areas. While the loss of visibility is a concern, the failure of the switches isolates the servers that the IDS devices are monitoring and results in the loss of connectivity to them. Thus, the failure of the switches represents a “fail-closed” solution. The servers are also protected by their endpoint intrusion protection software. The following sections detail each of the modules shown in Figure 9.

Corporate Internet Module

Figure 10 shows the recommended placement of network IDS sensors in the Corporate Internet Module for the medium, single-site design. Host-based IPS agents are also deployed to protect both the servers located in the Public Services Segment and the URL filtering server. All network devices (routers, switches, and firewalls) are set up to log syslog messages to servers located in the Management Module.



Figure 10
Medium, Single-Site Corporate Internet Module



Key Intrusion Detection and Syslog Devices

- *Network IDS appliances*—Provide Layer 4–7 monitoring of key network segments, active session resetting (TCP resets), and coordination of the blocking of malicious IP addresses (shunning)
- *Host-based IPSs*—Provide monitoring and intrusion prevention for the operating system and Web server application of key servers in the module, DoS protection, and attack mitigation at the endpoint host
- *Stateful firewall*—Provides syslog messaging regarding attempted access-control violations and administrative access as well as shunning
- *Routers*—Provide syslog messaging regarding attempted access-control violations and administrative access
- *Layer 2 switches*—Provide syslog messaging regarding administrative access

Threats Mitigated By Intrusion Detection and Syslog

- *Application layer attacks*—These attacks are mitigated through IDS at the host and network levels.
- *Viruses and Trojan horses*—These attacks are mitigated through e-mail content filtering and the host IDS.
- *Password attacks*—Limited services are available to brute force attacks, and the operating system and IDS can detect the threat.
- *Packet sniffers*—A switched infrastructure and host-based IPS limit exposure.
- *Network reconnaissance*—The IDS detects recon, and protocols are filtered to limit effectiveness.
- *Port redirection*—Restrictive filtering and the host-based IPS limit attacks.

Detailed Design

The Medium, Single-Site Corporate Internet Module IDS design comprises two sensors: one on the Public Services Segment and the other between the private interface of the firewall and the internal router. The primary function of the NIDS appliance on the Public Services Segment is to detect attacks on ports that the firewall is configured to permit. The primary duty of the NIDS appliance between the private firewall interface and the internal router is to provide a final layer of analysis against attacks. The tuning, alarm logging, and attack response of each sensor are detailed in the sections below.



Public Services Segment Sensor

The Public Services Segment Sensor sits on the segment that contains the public Web, FTP, DNS, and mail servers (Figure 10).

Objective

This sensor detects specific attacks on public service servers and also detects protocols that should not be seen on this segment. Configure this sensor in a restrictive manner because signatures seen here have successfully passed through the firewall. The servers in the Public Services Segment have endpoint intrusion prevention software installed. The primary function of the host-based IPS is to monitor and prevent rogue activity at the operating system level and in common server applications (HTTP, SQL, FTP, Simple Mail Transfer Protocol [SMTP], etc.).

Tuning and Alarm Logging

Tune this sensor to alarm if any traffic other than that which is required on the Public Services Segment is seen (connection signatures). Review and tune as necessary the signatures specific to vulnerabilities in Web, DNS, FTP, and mail servers, even though these servers may have the latest security patches and should be running endpoint intrusion prevention software. In general, leave the signatures tuned at the default levels, but implement custom string match signatures for vulnerabilities such as Code Red and Nimda that are not covered in the default attack signatures. Alerting on specific connection signatures, general attack signatures, and specific string signatures provides focused segment analysis at Layers 4 through 7.

The number of attack signatures seen on this segment varies. A Code Red attack, for example, can generate a large number of alarms. If the infected servers were patched for the vulnerability targeted in this attack, the attack can be tuned to alarm at a lower priority or not at all to avoid overwhelming the alarm console and security personnel, allowing both to concentrate on other, high-priority attacks.

Attack Response

Generally, avoid automatic blocking on this sensor. Although e-commerce is not a function of the Internet Module, there is a possibility that the source address of an attacker has been spoofed. Automatic blocking could prevent legitimate attempts to reach the Public Services Segment. Instead, manually reconfigure the firewall or router connected to the ISP to block a particular source address. Because they are less disruptive than automatic shunning, TCP resets may be used. Resets have been configured on this sensor within the SAFE lab reference network. Please see the cautions regarding implementing TCP resets under “Response Actions” in Appendix A.

Because the attack signatures that this sensor sees indicate that the firewall was penetrated, IP session data should be logged. Implement logging for a small percentage of signatures to minimize the data generated.

Alternatives

A sensor should be deployed here. If it is not, visibility into some network-level attacks will be lost; host-based IPSs should be installed on the public services servers to provide a significant degree of protection.

Internet Module Internal Sensor

The Internet Module Internal Sensor sits on the switch connected to the inside interface of the firewall (Figure 10).



Objective

This sensor monitors attacks that successfully breach the external firewall. Correlate data from this sensor with that of other sensors located throughout the network to track the spread of an attack or contagion. Configure this sensor in a restrictive manner because signatures seen here may be the result of a compromised server on the Public Services Segment. For example, if the SMTP server is compromised, the attacker might try to exploit the internal SMTP server over TCP port 25. This traffic is permitted between the Public Services Segment and the internal network in order to facilitate mail transfer between the two hosts.

Tuning and Alarm Logging

The traffic that this sensor sees is dependent upon the corporate security policy. For example, if the corporate security policy provides employees with open access to the Internet, then very little tuning based upon traffic type (connection signatures) can be performed on these sensors. However, if the corporate security policy is more restrictive, then tuning based upon traffic type is possible.

This sensor should see few attacks, and the attack signatures can generally be left tuned at the default levels. Attacks seen here have breached the external firewall, however, or are being generated from within the internal network, and should therefore be taken seriously. During incidents such as a Code Red attack, it may be necessary to tune specific signatures at a lower level in order to keep from being overwhelmed by high-priority alarms.

Attack Response

Address blocking and TCP resets can and probably should be used here. Shunning can be used because the traffic crossing these sensors is generated or terminated by users on the internal corporate site. Thus, although the source address of the attacker may be spoofed, shunning will generally not bring down any business-critical applications.

Attacks seen by this sensor can be generated from sources internal or external to the company, so the source must be determined quickly. Block external sources at the firewall. Track down and isolate internal sources, possibly using ACLs within the distribution layer of the SAFE architecture.

Because the attack signatures that this sensor sees indicate that the firewall was penetrated, IP session data should be logged. Implement logging for a small percentage of signatures to minimize the data generated.

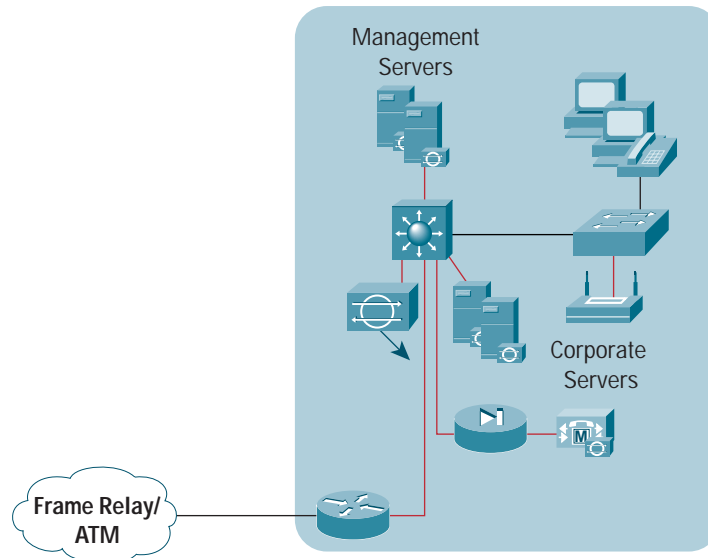
Alternatives

There are no alternatives to placing a sensor here. If limited resources are available for implementing sensors, one should be deployed here.



Campus Module

Figure 11
Medium, Single-Site Campus Module



Key Intrusion Detection and Syslog Devices

- *Network IDS appliances*—Provide Layer 4–7 monitoring of key network segments, active session resetting (TCP resets), and coordination of the blocking of malicious IP addresses (shunning)
- *Host-based IPS Console*— Provides centralized configuration management and viewing of alarms from host-based IPS agents deployed on servers throughout the network.
- *Host-based IPSs*—Provide monitoring and intrusion prevention for the operating system and Web server application of key servers in the module, DoS protection, and attack mitigation at the endpoint host
- *Stateful firewall*—Provides syslog messaging regarding attempted access-control violations and administrative access as well as shunning
- *Syslog host(s)*—Aggregates logging information for the firewall and NIDS.
- *Routers*—Provide syslog messaging regarding attempted access-control violations and administrative access
- *Layer 2 switches*—Provide syslog messaging regarding administrative access

Threats Mitigated By Intrusion Detection and Syslog

- *Application layer attacks*—These attacks are mitigated through IDS at the host and network levels.
- *Viruses and Trojan horses*—These attacks are mitigated through e-mail content filtering and the host IDS.
- *Password attacks*—Limited services are available to brute force attacks, and the operating system and IDS can detect the threat.
- *Packet sniffers*—A switched infrastructure and host-based IPS limit exposure.
- *Network reconnaissance*—The IDS detects recon, and protocols are filtered to limit effectiveness.
- *Port redirection*—Restrictive filtering and the host-based IPS limit attacks.



Detailed Design

The Medium, Single-Site Campus Module IDS design comprises a single sensor attached directly to the core switch. The tuning, alarm logging, and attack response for this sensor are detailed in the section below.

Objective

This sensor detects and analyzes attacks originating from within the campus network as well as those that get through the WAN module. For example, if a workstation is compromised because of an unknown modem connection, this sensor detects suspicious activity from that host within the campus network. Additional attack vectors could be disgruntled employees, workstations vulnerable to unauthorized employee access, and Trojan horse applications inadvertently loaded on laptops. Connect the sensor's monitoring port to a switch port configured to mirror traffic from all monitored VLANs.

All of the corporate intranet and management servers in this module have host-based IPS software loaded.

Tuning and Alarm Logging

The traffic that this sensor sees is dependent upon the corporate security policy. For example, if the corporate security policy provides employees with open access to the Internet, then very little tuning based upon traffic type (connection signatures) can be performed on these sensors. However, if the corporate security policy is more restrictive, then tuning based upon traffic type is possible.

The number of attacks seen by this sensor should be fairly minimal. Unless a specific IDS tuning is desired, a general guideline would be to leave specific attack signatures seen by this sensor at their default parameters. Attacks seen here are generated from within the internal network or originate from the corporate WAN, and should therefore be taken seriously. During incidents such as a Code Red attack, it may be necessary to tune specific signatures at a lower level in order to keep from being overwhelmed by a large number of high-priority alarms.

Attack Response

Address blocking and TCP resets can and probably should be used here. Shunning can be used because the traffic crossing these sensors is generated or terminated by users on the internal corporate site. Thus, although the source address of the attacker may be spoofed, shunning will generally not bring down any business-critical applications.

Attacks seen by this sensor can be generated from sources internal or external to the company, so the source must be determined quickly. Block sources originating from the corporate WAN at the router in the WAN module. Track down and isolate internal sources, possibly using ACLs within the distribution layer of the SAFE architecture.

Because the attack signatures that this sensor sees indicate that the firewall was penetrated, IP session data should be logged. Implement logging for a small percentage of signatures to minimize the data generated.

Alternatives

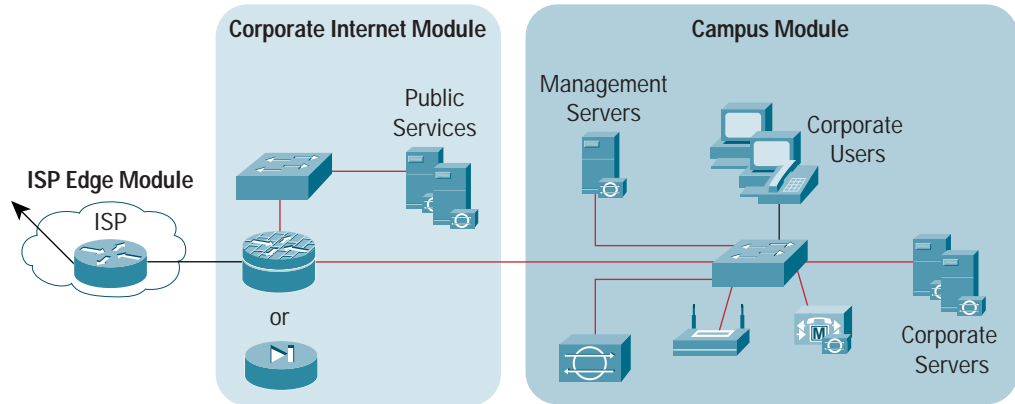
This sensor can be replaced with an integrated IDS module in the core switch. This alternative provides for higher throughput into the IDS, because it connects directly to the switch backplane rather than through a 10/100 Ethernet port. Use ACLs on the switch to control what traffic is sent to the IDS module for inspection. If limited resources are available for implementing sensors, one should be deployed here.



Small, Single-Site Design

The small, single-site uses the small network design from the SAFE white paper for small, midsize, and remote-user networks (Figure 12).

Figure 12
SAFE Small Network Design




The SAFE small network blueprint comprises two modules: the Corporate Internet Module and the Campus Module. However, this network design collapses the network security functions into a single system. The typical NIDS sensor is removed and the NIDS functions are moved to the edge router or firewall. The servers in both modules have endpoint security protection software installed. The following sections detail each of the modules.

Corporate Internet Module

Figure 13 shows the recommended placement of IDS functionality in the Corporate Internet Module for the small, single-site design. Host-based IPS agents are deployed to protect servers located in the Public Services Segment. All network devices (routers, switches, and firewalls) are set up to log syslog messages to servers located in the Campus Module where the management servers reside.

Figure 13
Small, Single-Site Corporate Internet Module





Key Intrusion Detection and Syslog Devices

- *Network IDS appliances*—Provide Layer 4–7 monitoring of key network segments, active session resetting (TCP resets), and coordination of the blocking of malicious IP addresses (shunning)
- *Host-based IPSs*—Provide monitoring and intrusion prevention for the operating system and Web server application of key servers in the module, DoS protection, and attack mitigation at the endpoint host
- *Stateful firewall*—Provides syslog messaging regarding attempted access-control violations and administrative access as well as shunning
- *Routers*—Provide syslog messaging regarding attempted access-control violations and administrative access
- *Layer 2 switches*—Provide syslog messaging regarding administrative access

The network IDS, firewall, and routing functions are provided by either the router with an enterprise operating system loaded or the firewall.

Threats Mitigated By Intrusion Detection and Syslog

- *Application layer attacks*—These attacks are mitigated through IDS at the host and network levels.
- *Viruses and Trojan horses*—These attacks are mitigated through e-mail content filtering and the host IDS.
- *Password attacks*—Limited services are available to brute force attacks, and the operating system and IDS can detect the threat.
- *Packet sniffers*—A switched infrastructure and host-based IPS limit exposure.
- *Network reconnaissance*—The IDS detects recon, and protocols are filtered to limit effectiveness.
- *Port redirection*—Restrictive filtering and the host-based IPS limit attacks.

Detailed Design

The small, single-site Corporate Internet Module IDS design comprises only the IDS functions that the edge router or the edge firewall provides (Figure 13). The NIDS process on the edge device detects attacks on ports that the device is configured to permit. The tuning, alarm logging, and attack response for this process are detailed in the section below.

Objective

The edge device detects specific attacks on public service servers (Web, FTP, DNS, etc.) and also detects protocols that should not be seen on this segment. The small-network design constraints prohibit the NIDS process on the edge device from being configured in a restrictive manner. Each of the public services servers has host-based IPS software installed, the primary function of which is to monitor and prevent rogue activity at the operating system level and in common server applications (HTTP, SQL, FTP, SMTP, etc.).

Tuning and Alarm Logging

Tuning this multi-purpose edge device is more difficult than tuning a dedicated IDS appliance. Tune the device to alarm if any traffic other than that which is required is seen (connection signatures). Review and tune as necessary the signatures that are specific to vulnerabilities in Web, DNS, FTP, and mail servers, even though these servers may have the latest security patches and should be running host-based IPSs. In general, leave the signatures tuned at the



default levels, but implement custom string match signatures for vulnerabilities such as Code Red and Nimda that are not covered in the default attack signatures. Alerting on specific connection signatures, general attack signatures, and specific string signatures provides focused segment analysis at Layers 4 through 7.

The number of attack signatures seen on this segment varies. A Code Red attack, for example, can generate a large number of alarms. If the infected servers were patched for the vulnerability targeted in this attack, the attack can be tuned to alarm at a lower priority or not at all to avoid overwhelming the alarm console and security personnel, allowing both to concentrate on other, high-priority attacks.

Attack Response

Generally, avoid automatic blocking on this edge device. Although e-commerce is not a function of the Internet Module, there is a possibility that the source address of an attacker has been spoofed. Automatic blocking could prevent legitimate attempts to reach the Public Services Segment. Instead, manually reconfigure the firewall or router connected to the ISP to block a particular source address. Because they are less disruptive than automatic shunning, TCP resets may be used. Resets have been configured on the edge device within the SAFE lab reference network. Please see the cautions regarding implementing TCP resets under “Response Actions” in Appendix A.

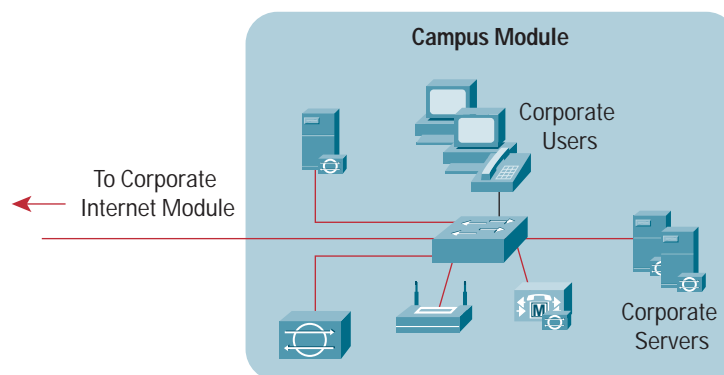
IP session data based on attack signatures seen here should be logged. Implement logging for a small percentage of signatures to minimize the data generated.

Alternatives

An edge device with IDS functionality should be deployed here. If it is not, visibility into some network-level attacks will be lost; host-based IPSs should be installed on the public services servers to provide a significant degree of protection.

Campus Module

Figure 14
Small, Single-Site Campus Module



Key Intrusion Detection and Syslog Devices

- *Host-based IPS Console*—Provides centralized configuration management and viewing of alarms from host-based IPS agents deployed on servers throughout the network
- *(Alternative) Network IDS appliance*—Provide Layer 4–7 monitoring of key network segments, active session resetting (TCP resets), and coordination of the blocking of malicious IP addresses (shunning)



- *Host-based IPSs*—Provide monitoring and intrusion prevention for the operating system and Web server application of key servers in the module, DoS protection, and attack mitigation at the endpoint host
- *Syslog host(s)*—Aggregates logging information for the firewall and NIDS.
- *Layer 2 switches*—Provide syslog messaging regarding administrative access

Threats Mitigated By Intrusion Detection and Syslog

- *Application layer attacks*—These attacks are mitigated through host-based IPS.
- *Viruses and Trojan horses*—These attacks are mitigated through e-mail content filtering and the host-based IPS.
- *Password attacks*—Limited services are available to brute force attacks, and the operating system and host-based IPS can detect the threat.
- *Packet sniffers*—A switched infrastructure and host-based IPS limit exposure.
- *Network reconnaissance*—The host-based IPS detects recon.
- *Port redirection*—Restrictive filtering and the host-based IPS limit attacks.

Detailed Design

The small, single-site Campus Module relies on deployment of the host-based IPS to provide IDS functionality in the module.

Objective

The primary function of IDS in the campus module is to protect corporate servers from attack. Any signatures seen here have passed through the IDS at the Corporate Internet Module edge device or are generated from within the network. Because of the network's size and the lower level of complexity, a dedicated NIDS appliance in this module would not benefit the overall architecture. Endpoint defense provides the greatest coverage in this module. Attack vectors include disgruntled employees, workstations vulnerable to unauthorized employee access, and Trojan horse applications inadvertently loaded on laptops. All of the corporate intranet and management servers in this module have host-based IPS software loaded.

Tuning and Alarm Logging

The host-based IPSs in this module should see few attacks. Leave specific attack signatures at their default settings.

Attack Response

Attacks seen in this module can be generated from sources internal or external to the company and represent a significant security threat. Track down and isolate internal sources quickly.

Alternatives

Adding a network IDS behind the firewall provides for Layers 4 through 7 monitoring of all traffic entering and leaving the Campus Module. TCP resets and shunning can be deployed here because signatures seen in this module have neither passed through a legitimate port in the firewall nor were generated internally. The long-term impact of deployment is minimal because no systems here are externally visible.



Large, Multi-Site Design

In a multi-site design, the enterprise management network resides at a separate location, but monitoring and management are conducted at the local, regional, and enterprise levels through secure and at times partially secure links. Management is provided through two architectures:

- A *distributed, hierarchical design*—Regional consoles communicate with local consoles in smaller networks or directly manage the IDS and syslog devices in these smaller networks. The regional consoles aggregate alarms and syslog information and summarize them for the enterprise-level management system. This design is shown in Figure 15.
- A *centralized management design*—The enterprise-level management console communicates with and manages information from every IDS and syslog device in the enterprise. This design is shown in Figure 16.

Figure 15
Distributed Management Framework

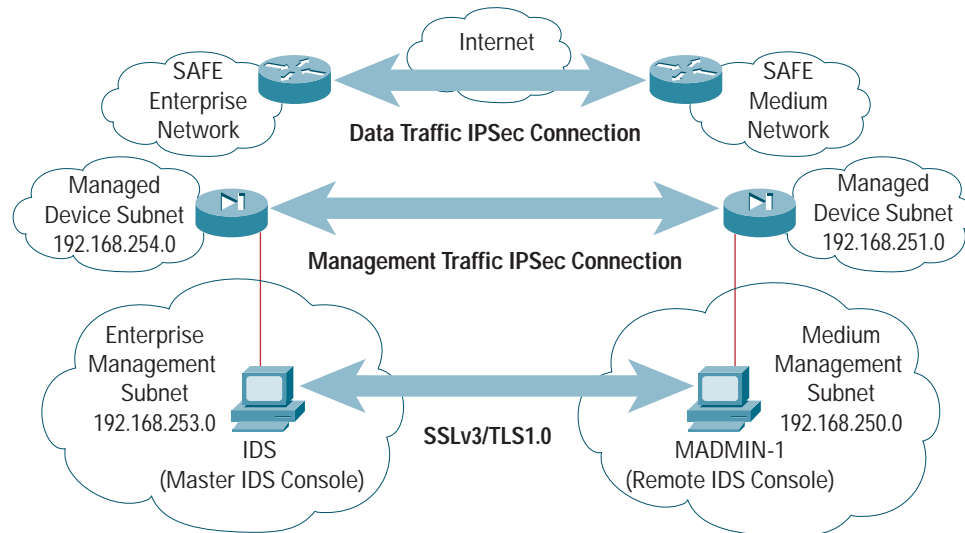



Figure 16
Centralized Management Framework





In Figure 15, a local IDS and syslog management console is present at the remote network to manage locally significant events but also to forward selected events to the enterprise-level management system. The selection of alarms to forward is based on two factors:

- Severity—The danger that an event represents to the security of the network
- Extent—The magnitude of the event and its threat to the overall security of the network

Both large, multi-site designs are discussed in detail below.

Distributed, Hierarchical Design

Detailed Design

In the distributed, hierarchical design for large, multi-site networks, communication between the central servers at headquarters and the servers at the remote and branch locations is achieved through a modified OOB method. The firewall provides an IPSec connection for secure management and monitoring of channels for the OOB network.

Modified OOB Communication

In the SAFE enterprise network, the management servers reside in their own module and communicate with devices in other modules through an OOB network. However, the management servers in the medium network all reside within the Campus Module and management is performed in band. To provide true OOB communication between the management servers in the Enterprise Management Module and the management servers in the Campus Module requires modification to the Campus Module in the medium network. The modification is shown below in Figure 17. The firewall that was added can be replaced by a router running firewall software.

In Figure 17, the firewall separates the management servers from the rest of the network and can terminate the IPSec VPN tunnel from the Enterprise Management Module. This VPN carries management traffic between the edge routers of the enterprise and medium networks. Thus, while preserving its integrity as much as possible, the OOB network present in the Enterprise Management Module is now extended to the medium network.

The additional firewall allows the following traffic to pass into the management servers:

- *TFTP (UDP 69)*—For network device configuration files from devices on the Managed Devices Segment
- *Syslog (UDP 514)*—From network devices on the Managed Devices Segment
- *NTP (UDP 123)*—To synchronize the clocks of all network devices on the Managed Devices Segment
- *HTTP (TCP 80)*—To the Internet and from hosts on other segments to download the host-based IPS agent software
- *HTTPS (TCP 443)*—To network devices on the Managed Devices Segment and the Internet as well as between the host-based IPS Console and its agents
- *TACACS+ (TCP 65)*—For administrator authentication to devices on the Managed Devices Segment
- *RADIUS (UDP 1645)*—For authentication of administrator remote-access VPN connections coming from the Remote Administration Segment
- *ICMP (IP Protocol 1)*—Echo request and response to reach network devices on the Managed Devices Segment and the Internet
- *DNS (UDP 53)*—For name translation services for management hosts as they access services on the Internet



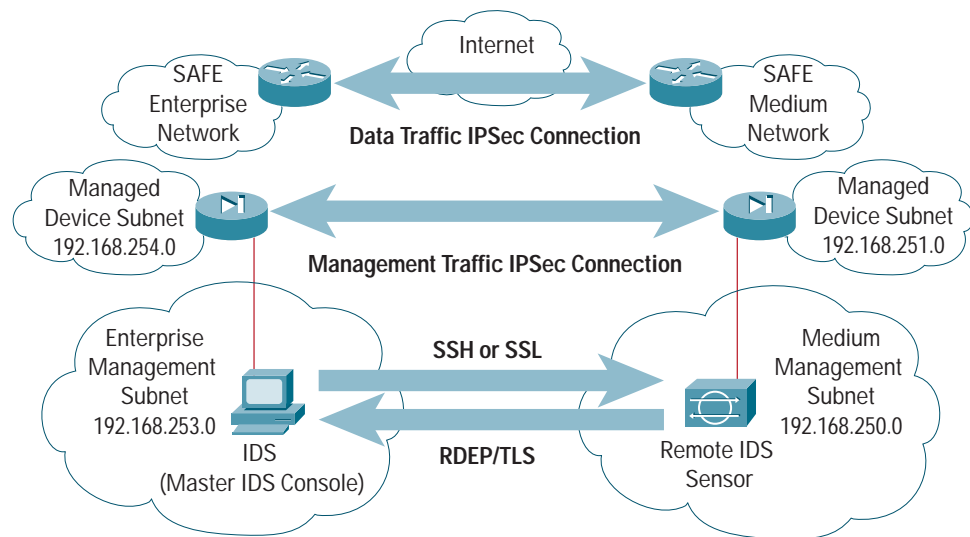
- *SNMP (UDP 161)*—To query information from network devices on the Managed Devices Segment
- *SNMP-Trap (UDP 162)*—To receive trap information from network devices on the Managed Devices Segment

Centralized Management Design

Detailed Design

In the centralized management design for large, multi-site networks, the central servers in the Enterprise Management Module configure the remote IDS sensors and receive the alarms and logs from these sensors and other network devices. There are no local IDS management servers. As with the distributed, hierarchical design, a firewall is added to the Campus Module in the centralized management design to terminate the IPsec VPN tunnel between the Enterprise Management Module and the management servers in the Campus Module of the remote or branch network. An alternative design uses in-band management over SSH or Secure Sockets Layer (SSL) between the Enterprise Management Module and the remote IDS sensors.

Figure 17
Centralized Management Framework



Medium, Multi-Site Design

The medium, multi-site design is based on the centralized management framework for the large, multi-site design. The management servers are located in the medium network's Campus Module and provide alarm capture and monitoring capabilities as well as logging for devices in the medium network. In the multi-site design, dedicated connections or VPN over the Internet connect the smaller networks to the medium, central network.



Figure 18
Centralized Management Framework in a Medium, Multi-Site Design

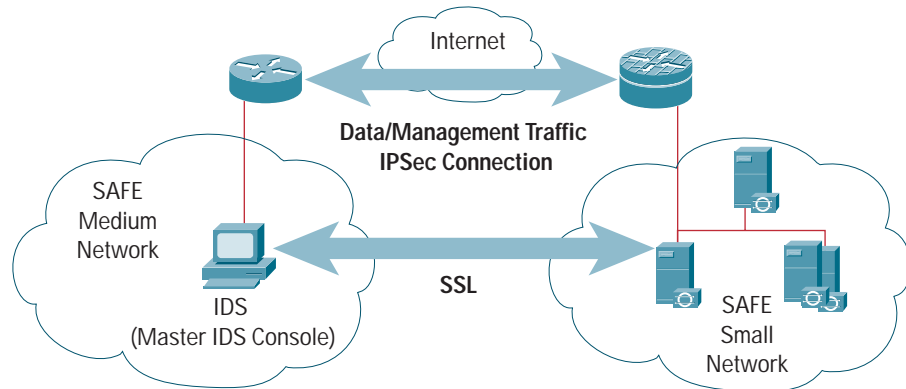
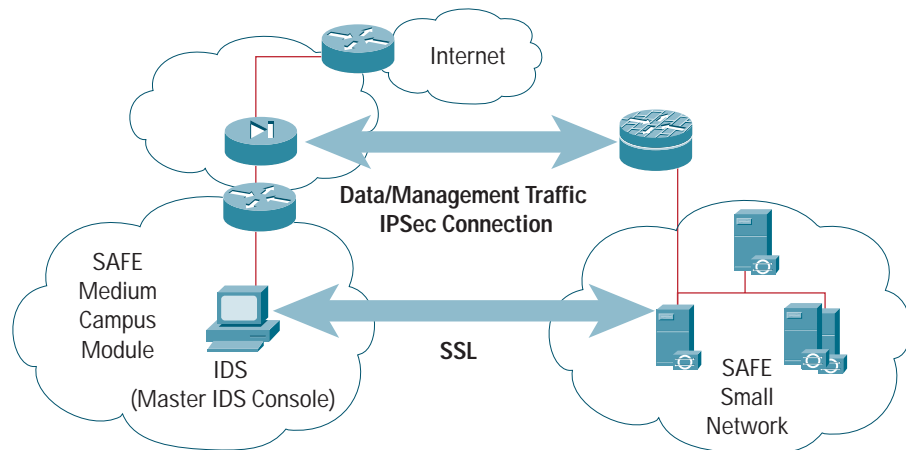


Figure 19
Alternative Centralized Management Framework in a Medium, Multi-Site Design



Detailed Design

The medium, multi-site design uses in-band communication for server and system management. VPN tunnels or dedicated paths such as leased lines enable communication between the central medium network and the small or medium branch networks. The VPN tunnels carry data and management traffic, because management and security monitoring are conducted in band. A router or a firewall provides the IPSec communications link between the central medium network and the small branch networks (Figures 18 and 19, respectively). An SSL link enables communications between the console at the central medium network and the the host IDS or IPS software on the servers in the small remote network(s). The following protocols must be able to traverse the firewall in the alternative design (Figure 19):

- *TFTP (UDP 69)*—For network device configuration files from devices on the Managed Devices Segment
- *Syslog (UDP 514)*—From network devices on the Managed Devices Segment
- *NTP (UDP 123)*—To synchronize the clocks of all network devices on the Managed Devices Segment



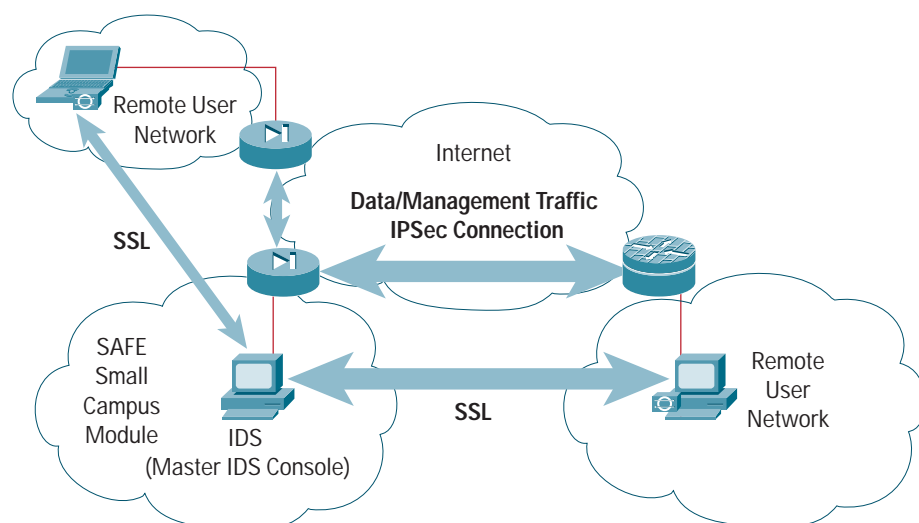
- *HTTP (TCP 80)*—To the Internet and from hosts on other segments to download the host-based IPS agent software
- *HTTPS (TCP 443)*—To network devices on the Managed Devices Segment and the Internet as well as between the host-based IPS Console and its agents
- *TACACS+ (TCP 65)*—For administrator authentication to devices on the Managed Devices Segment
- *RADIUS (UDP 1645)*—For authentication of administrator remote-access VPN connections coming from the Remote Administration Segment
- *ICMP (IP Protocol 1)*—Echo request and response to reach network devices on the Managed Devices Segment and the Internet
- *DNS (UDP 53)*—For name translation services for management hosts as they access services on the Internet
- *SNMP (UDP 161)*—To query information from network devices on the Managed Devices Segment
- *SNMP-Trap (UDP 162)*—To receive trap information from network devices on the Managed Devices Segment

For the design shown in Figure 18, where there is no firewall separating the IDS management console and any other management servers in the medium network's Campus Module, SNMP (UDP 161) and SNMP-Trap (UDP 162) should be restricted to SNMPv3 only.

Small, Multi-Site Design

The small, multi-site design is based on the centralized management framework for the large, multi-site design. The management servers are located in the small network's Campus Module and provide alarm capture and monitoring capabilities as well as logging for devices in the small network. In the multi-site design, dedicated connections or VPN over the Internet connect remote-user networks to the small, central network (Figure 20).

Figure 20
Small, Multi-Tier Design





Detailed Design

The small, multi-site design uses in-band communication for server and system management. VPN tunnels or dedicated paths such as leased lines enable communication between the central small network and the remote-user networks. The VPN tunnels carry data and management traffic, because management and security monitoring are conducted in band. This design is based on the medium design shown in Figure 18, where no firewall segments off the management server(s) from the rest of the network. The SNMP caveats discussed in the section on the medium, multi-site design apply to this small, multi-site design as well. An SSL link enables communications between the host IDS or IPS software on the servers in the remote-user network(s) and the console at the central network.

Appendix A: Intrusion Detection Primer

Intrusion Detection History

Intrusion detection is derived from mainframe system audits, which once identified potential security policy violations. When computers and networks became more sophisticated, audits no longer provided real-time protection. In the late 1970s, the U.S. Department of Defense commissioned James P. Anderson to study network security. Anderson's report noted the need for an automated audit trail that would support security goals. Anderson proposed a taxonomy differentiating internal and external risks and threats.

Between 1984 and 1986, Dorothy Denning and Peter Neumann developed one of the first IDSs, the Intrusion Detection Expert System (IDES), which used profiles to identify security policy violations. These profiles were data structures that used statistical metrics and models. The models described the behavior of system subjects such as users. The system's hybrid architecture combined an anomaly detector and an expert system. Other IDSs developed during the 1980s include the U.S. Air Force Haystack system, the Multics Intrusion Detection and Alerting System (MIDAS), the Network Audit Director and Intrusion Reporter (NADIR), and the Network System Monitor (NSM). The NSM was the first attempt to monitor network traffic and to use that traffic as the primary data source.

Until 1990, IDSs were host based and confined their monitoring to operating system audit trails and other host-oriented information sources. The first major effort to integrate host- and network-based IDSs was known as the Distributed Intrusion Detection System (DIDS). The goal of DIDS was a central security operations center that could track security violations and intrusions across networks.

Commercial IDSs were also developed in the 1980s and 1990s, but it was not until the mid- to late-1990s that this technology matured enough for large-scale deployment.

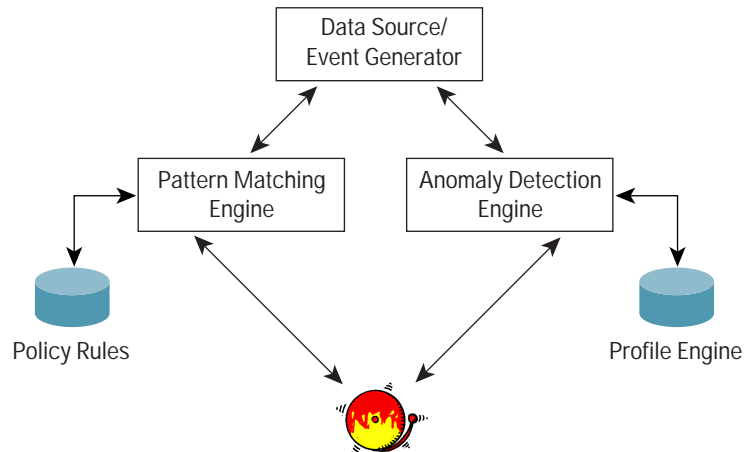
Intrusion Detection Technologies

Modern IDSs typically comprise a data source, an analysis engine, and a response component. The IDS provides a continual audit of network and system resources through either misuse detection or anomaly detection. Misuse detection relies on pattern matching to detect attacks or security violations. Anomaly detection IDSs look for deviations from normal patterns of behavior.

A generic IDS model is shown below in Figure 21. In this model, the traffic from the data source is evaluated by either a pattern matching engine or an anomaly detection engine (or in some cases a hybrid engine). A pattern matching engine compares the information from the data source against a policy rules database. This database can include specific signatures for various attacks. An anomaly detection engine compares the information against a profile engine that defines "normal" behavior.



Figure 21
Generic IDS Model



Anomaly Detection

Anomaly detection is based on learned behavior patterns. Anomaly-based IDSs monitor network activity to develop a model of normal traffic behavior and then use statistical techniques to identify pattern deviations. When traffic that is outside the norm is detected, the system raises an alarm. Anomaly-based IDSs can adapt to new, unique, or original attacks, and they are not dependent on operating system knowledge. However, this IDS does produce a high rate of false alarms. In addition, a network's traffic patterns may not be static enough for this type of IDS. The IDS then remains in a constant learning state while it tries to define the patterns well enough to identify deviations.

Misuse Detection

To define intrusion events, misuse detection IDSs use a database of known attacks and vulnerabilities as well as security policy violations. The analysis engine filters the event stream from the data source, searching for patterns that match entries in the misuse database. This system is also known as a knowledge-based or signature-based IDS and is the most common IDS deployed.

Monitoring Models

Monitoring is the action of collecting information from a data source and passing it on to an analysis engine. IDS monitoring can be network based, host based, application focused, or target focused. Each of these strategies is described in detail below.

Host-Based Monitoring

In host-based monitoring, the data source is internal to the system. This data source is often represented by system event logs such as syslog files, Windows event logs, or other audit log files (for example, Solaris BSM log files). Information in the data source is restricted to operating system events, but modern, host-based monitoring systems have introduced capabilities such as the monitoring of system calls and program execution. These features enable the monitoring system to detect attacks in real time and often to prevent an attack altogether. This IDS type is considered an IPS, because it prevents intruders from attacking the host operating system.



Network-Based Monitoring

Network-based monitoring uses as its data source the traffic passing along network links. Capturing this traffic involves placing the network interface into promiscuous mode. Normal network interface operation involves monitoring traffic across the system's network link and passing only the traffic destined for the host to the upper protocol layers. In promiscuous mode, however, the interface passes all traffic to a traffic-capturing process or program running on the system.

The network architecture affects network-based monitoring. With the wide deployment of switches throughout networks, the traffic destined for a port is reduced to only that traffic that is destined for a host connected to that port. Thus, network operations personnel employ either a port mirroring capability in the switch or use a physical network tap to capture traffic across a network. In port mirroring, the switch is configured to allow one port to see all traffic traversing another port. The mirrored port should be an uplink port where all traffic is leaving or entering the network. This could be the port to which a router or a firewall is connected.

Application-Based Monitoring

In application-based monitoring, applications running on the host are the source of information for the analysis engine. This includes application event logs and database journal logs. The application need not log to the system logging facility (for example, the Apache Web server). This type of monitoring protects the application from attack, but relies on an “after-the-fact” response.

Target-Based Monitoring

In target-based monitoring, a baseline snapshot of the host's state is created and subsequent snapshots are compared against the baseline. This type of monitoring often uses cryptographic signatures of various system files and programs and stores that database either locally on the host in a read-only fashion or on other types of media. The key feature is that an intruder cannot modify the baseline snapshot database. This type of monitoring is designed to detect changes to system-level objects.

Signatures

Three types of signatures can be implemented on network IDS sensors:

- Exploit signatures—Which match a specific known attack
- Connection signatures—Which generate an alert when a particular type of traffic (protocol) is seen
- String signatures—Which match particular string sequences seen in the data

Administrators can write custom signatures to match specific attacks that are not covered by the signature types above.

Exploit Signatures

Exploit signatures are implemented to alert security operations personnel when a specific attack is underway. There are hundreds of exploit signatures and more are added with each signature pack update; placing them in broad categories can help them to be understood and used appropriately.

Exploit signatures can be grouped by what layer of the Open Systems Interconnection (OSI) model the attack attempts to exploit. Certain exploit signatures, for example, indicate network layer—IP and ICMP—attacks. IP-based exploit signatures include fragmentation attacks (to bypass ACLs or to mount a DoS attack), attacks on the



IP-packet options field (to route packets specifically, for example, circumventing routing tables), and address-spoofing attacks (to exploit trust relationships between network hosts). ICMP-based exploits are often used for network reconnaissance or DoS attacks.

Exploit signatures can also indicate transport layer—TCP and UDP—attacks. TCP-based exploit signatures include port sweeps (for network reconnaissance) and embryonic or half-opened connections (for DoS attacks). Because UDP is stateless, UDP exploits include reconnaissance through port sweeps and denial of service through packet flooding.

Because ACLs and firewalls block many network and transport layer attacks and because these attacks often involve denial of service and reconnaissance—which are of most concern on connections facing the Internet—it is important to pay close attention to the exploit signatures that indicate these attacks when tuning the sensors outside the firewall. Be careful to set the DoS alarm threshold high enough to avoid false alarms when traffic levels fluctuate. Several days of monitoring are usually necessary to determine what constitutes normal traffic behavior.

The next group of exploit signatures identifies attacks on applications that provide networking services and on end-user applications. Applications that provide networking services include DNS, remote-procedure call (RPC), Network Information Services/Network File Sharing (NIS/NFS), FTP, TFTP, Telnet, NetBIOS, etc. The purpose of attacks on these services depends on the service itself. For example, the goal of many DNS exploits is network reconnaissance, although denial of service and escalation of privilege through the exploitation of buffer overflow vulnerabilities is also a major concern. End-user application vulnerabilities include Web server implementations, e-mail systems, and specific applications and products such as Microsoft Windows. Vulnerabilities in Web server implementations are exploited most often. The objective of these exploits is also often either escalation of privilege (to “own the system”) or denial of service through the exploitation of known buffer overflow vulnerabilities. Exploits attacking vulnerabilities in applications are contained in traffic that is allowed past firewalls and router ACLs, so the first line of defense is to patch the application with the latest security fixes. Host-based IPSs are often an effective defense against application-specific attacks and they also enforce layer security.

Be aware of which network IDS exploit signatures are relevant to the applications in use on the network and of which signatures a particular sensor should see. For example, in a DMZ segment that contains a Web server, DNS server, and FTP server, pay attention to exploit signatures relevant to DNS, Web server, and FTP traffic when tuning the sensor sitting on that segment. If there is no e-mail server sitting on that segment and no e-mail traffic crosses that segment, then time should not be spent tuning signatures on that sensor that are relevant to e-mail vulnerabilities. Tuning such exploit signatures to alarm at a low priority level will provide visibility when someone attempts to breach the network with a particular e-mail vulnerability, but will not prevent the security staff from paying attention to alerts where a real vulnerability exists.

Exploit signatures also alert when traffic that is associated with a particular attack tool is identified. Such attack tools include Trinoo, TFN, Stacheldraht, Security Administrator's Tool for Analyzing Networks (SATAN), BackOrifice, etc. Sensors should alert at a high level if traffic specific to a known attack tool is seen on the network. If security operations personnel are using scanning tools to determine vulnerabilities, however, sensors should be configured temporarily to not alarm when exploit signatures are seen originating from the IP address that is assigned to the workstation running the tools.

Exploit signatures can be categorized in many other ways. The key is always to identify and tune the signatures that are relevant to the specific applications, services, traffic, etc. on the network.



Connection Signatures

Connection signatures are used more broadly and are tuned based on the location of the network IDS sensor. Connection signatures can be tuned so that the sensor alarms if a particular traffic (protocol) type is seen. For example, a sensor that sits directly behind a firewall can be tuned to alarm if it sees NFS traffic. The security policy may state that transferring files to the Internet using NFS is not secure and is prohibited. The sensor would then alert security operations personnel of potential policy violations by employees. An alarm could also indicate that the firewall is misconfigured and is allowing NFS file transfers to the Internet.

Connection signatures can also be tuned so that the sensor alarms if traffic other than the expected is seen. For example, a sensor that sits on a DMZ segment of the firewall can be tuned to alarm if it sees anything other than Web (HTTP, HTTPS), FTP, DNS, SMTP, and IPSec encrypted VPN traffic. Traffic other than the protocols listed above could indicate one of the following:

- *A backdoor connection into the DMZ has been added*—This could be an attempt to circumvent the firewall. However, it could also indicate the misconfiguration of a network switch, an inadvertent cable connection, or a misplaced server on the segment. It is critical that network operations personnel work with security operations personnel to identify and correct such issues before an attacker exploits them.
- *The firewall is misconfigured*—This could be an attempt to circumvent firewall security or simply a configuration error or unintentional opening in the firewall. Mistakes must be identified and corrected before they are exploited.
- *Additional services have been installed or enabled on servers*—This could be the result of malicious activity or it could be that applications or services have unintentionally been left enabled on the servers. Security operations personnel must work with server administrators to identify and correct such issues. If it is not possible to disable additional, non-malicious services, then it may be necessary to tune the sensors to ignore the alarms generated by such traffic.
- *The VPN concentrator is misconfigured and is allowing unencrypted VPN traffic onto the Internet*—This could be an attempt to circumvent the security of the VPN concentrator or simply a configuration error allowing VPN peers to negotiate a tunnel with the null encryption option.

If it is possible to determine what traffic should and should not be on a particular segment, then it is possible to implement connection signatures on that segment's sensor. Because connection signatures do not represent known attack signatures, it may be advantageous to set the sensor to alarm at a level lower than it would for an attack. Security operations personnel will therefore be alerted to the issue, but will not be prevented from acting on network attacks.

String Signatures

String signatures can be deployed on sensors throughout the network to identify specific character sequences in data packets. For example, string signatures can identify the character sequence in an URL that causes a buffer overflow exploit similar to Code Red or Nimda. Security and network operations personnel can use the alerts generated to quickly identify an infected host and to isolate it through automatic blocking or by reconfiguring ACLs or firewalls. To prevent competitors from receiving confidential information over the Internet, a string signature can also be set up to alert if the words “confidential” or “proprietary” are seen by sensors that sit behind a firewall with Internet access. Obviously, if the information is encrypted, security operations personnel will not receive an alert. This type of string signature should not be deployed on all network sensors, because the corporate security policy may not restrict the flow of confidential information within the network.



Because string signatures and their deployment are customized, there are no general guidelines available for the setting of alarms.

Response Actions

Blocking

Avoid automatic blocking (shunning) on any device where the probability is high that an attacker's source address can be spoofed and where mission critical services such as e-commerce transactions can be adversely affected. Instead, implement manual blocking as needed. Blocking may also be implemented within the network to protect corporate servers that are not customer facing. IP spoofing in this environment is unlikely and any attack has a high probability of success. However, sensors should never block addresses that provide critical network resources. For example, if active routing updates provide reachability to all required devices, then sensors should be configured never to block the source IP addresses of the routers and Layer 3 switches that provide routing updates. Blocking the updates would lead to DoS attacks. Critical resources include routing sources, naming services (DNS, Windows Internet Naming Service [WINS]), etc.

On a stateful firewall, the time the address block is in place should exceed the timeout parameter for the connection. While the block is in place, there may be a legitimate connection through the firewall, but the shun blocks all traffic for that device from entering the firewall. The length of time that the shun is in place can be extended to allow the security response team to respond to an attack.

TCP Resets

TCP Reset is a blocking method whereby an IDS sensor responds to an attack by sending the source address of the attack a TCP RST packet to terminate the connection. The IDS must be able to spoof the IP address of the target host as well as determine the TCP sequence number. Be careful when implementing TCP resets: tune out false positives so that TCP resets are not sent to legitimate hosts. When upgrading the signatures in sensors, do not enable TCP resets until false positives are no longer generated by new signatures. TCP resets are appropriate on sensors where blocking is not feasible. This includes sensors monitoring traffic near publicly available servers on any Public Services Segment. Another potential deployment location is at extranet connections to partner networks.

IP Logging

Network IDS sensors can be configured to log IP packet data after an attack signature is seen. This data can help determine the extent to which the network was compromised. Configure the sensor to limit the size of log files so that performance is not degraded, and to automatically transfer the data to a dedicated server on the management network. Set up this server to write the data automatically to a non-erasable medium such as a CD to prevent the accidental or intentional erasure or corruption of data. IP logging generates a lot of information, so configure logging only if there is a plan in place for data use, and be prepared with the storage resources and manpower required to maintain the data.



Logging/Reporting

A reporting system should be flexible enough to provide the following:

- Real-time, hourly, daily, weekly, and monthly trend analysis reporting showing top attack sources, the overall number of attacks, and the specific times of day when attacks occur
- Top attack signatures seen, including the overall number of these attacks and the specific times of day when they are seen
- Administration notification in case of events

On sensors that alarm often, such as those on the outside interfaces of firewalls facing the Internet, it may be best to manage information by exception. This helps focus resources to safeguard against legitimate attacks. In such cases, the following additional information may be desired from the reporting system:

- *Deviation of a particular attack signature from the normal levels seen*—For example, if the level of Code Red attacks suddenly deviates outside a certain percentage of the normal levels, a new threat could be responsible and should be investigated.
- *New signatures not previously seen as part of the normal attacks on the site*—Again, this new signature could indicate a new threat and should be investigated.

Many existing reporting systems do not provide this kind of reporting; it may be necessary to develop custom applications.

Appendix B: Syslog Primer

Logging and Syslog

Most network devices can generate log files of operational parameters and events encountered. The content of log messages is often implemented in the operating system of the device.

All Unix and Windows operating systems support log subsystems that record messages to a local storage device. Clients (routers, switches, firewalls, etc.) that generate syslog messages are referred to as *devices*. Servers that receive syslog messages are referred to as *collectors*. Most syslog servers are Unix or Windows based.

Logging is particularly important in order to audit the application of a security policy. Security devices log by default or are configured to log when they apply a policy that could block a network connection. On most network devices that log to an external server, logging is disabled by default. On devices where log data is stored locally, logging (at some level) is enabled by default.

Products from different vendors use different log formats. While there is no industry-wide standard that defines how to log or what format log files should take, many vendors implement the Berkeley Standard Distribution (BSD) syslog. BSD syslog is a portion of the BSD Unix operating system and other Unix variants and its behavior is described in the Internet Engineering Task Force (IETF) RFC 3164. BSD syslog transports data using the TCP/IP UDP protocol on port 514. Because UDP is a connectionless protocol, syslog message delivery is a “best effort” process: packets that are lost are not retransmitted.

Other vendors have developed their own log protocols. Microsoft has implemented a proprietary event log capability in many of its Windows operating systems, for example. While these event logs capture much of the same information as syslog, the two are not compatible. Products from third-party vendors can be added to many Microsoft operating systems to convert event logs to syslog format.



Logs contain basic information (IP addresses, port numbers, ACL numbers) that can be used to analyze the actions of a device. Most log messages are difficult to understand without the log message number. Message numbers and log data are described as part of the device documentation. Each log message that is generated also has an associated timestamp. In an environment where log messages from many devices are collected at a central server, it becomes important to understand not just when messages arrived at the log server, but also when the message was generated by the device.

When collecting log messages, make sure that all devices are using the same clock. If a device is busy with high-priority tasks, it may delay transmitting a log message to the central server. If the clock used to generate the timestamp is wrong, it can be difficult to determine the order in which events on different devices occurred.

To further differentiate log messages, each message is assigned a priority level. Priority levels categorize messages based on their importance or severity, with lower priority values reflecting more important or severe log messages. The seven priority levels defined in BSD syslog are provided in Table 1 below.

Table 1 RFC-Defined Syslog Priority Messages

Severity	Numerical Code	Value
Emergency	0	System is unusable
Alert	1	Immediate action is necessary
Critical	2	Critical conditions
Error	3	Error conditions
Warning	4	Warning condition
Notice	5	Normal but significant condition
Information	6	Informational messages
Debug	7	Debug level messages

Log level setting can be configured both at the device and at the log server. Log level setting at the device defines which messages based on their level will be generated. Log level setting at the server is used to prioritize message processing. Often there are fewer messages at the lower priority levels.

Log messages are also assigned a facility identifier. Log facilities were developed in Unix implementations of syslog to differentiate log messages generated by different applications running on the same computer. Table 2 lists the various syslog log facilities.

Table 2 Syslog Log Facility Levels

Log Facility	Numerical Code	Value
KERN	0	Kernel messages
USER	1	Random user-level messages
MAIL	2	Mail system
DAEMON	3	System daemons



Table 2 Syslog Log Facility Levels (Continued)

Log Facility	Numerical Code	Value
AUTH	4	Security/Authorization messages
SYSLOG	5	Messages generated internally by syslogd
LPR	6	Line printer subsystem
NEWS	7	Net news subsystem
UUCP	8	UUCP subsystem
CRON	9	Clock daemon (i.e. cron/at) subsystem
RESERVED	10	Security/Authorization messages
RESERVED	11	FTP daemon
RESERVED	12	NTP subsystem
RESERVED	13	Log audit
RESERVED	14	Log alert
RESERVED	15	Clock daemon (i.e. cron/at) subsystem
LOCAL0	16	Local Use 0
LOCAL1	17	Local Use 1
LOCAL2	18	Local Use 2
LOCAL3	19	Local Use 3
LOCAL4	20	Local Use 4
LOCAL5	21	Local Use 5
LOCAL6	22	Local Use 6
LOCAL7	23	Local Use 7

Most logging systems by default pass traffic “in the clear”: if the data is intercepted, it is not difficult to reassemble and read. If log data is transmitted across an untrusted network, or if the data contains sensitive information, consider using a VPN.

Most organizations will implement a centralized logging system. Analyzing data from devices across the network enables the clearest picture to be drawn of a network event. Centrally locating log data also enables easier and more consistent backups, and places all log analysis tools on the server where the data is stored. One potential pitfall, however, is the lack of redundancy. Syslog enables log messages to be sent to multiple log servers, which helps ensure that messages will not be lost if one log server is down. Also, because syslog uses the connectionless UDP as a transport protocol, logging to multiple servers also helps protect against the loss of messages due to network failures.



While many devices generate log data, the value of that data across devices is not equal. It is most important to obtain log data from devices where forwarding or blocking decisions are made. Log data is important from a router or switch with an ACL that can block traffic, from a firewall that can deny connections, from a VPN concentrator or other remote-access device that can accept connections, and from an IDS that can reset a connection between hosts. Development of logging standards continues today in the IETF syslog working group.

Appendix C: Architecture Taxonomy

Anomaly detection: Detection based on whether system activity deviates outside a defined activity.

Application server: Device that provides application services directly or indirectly for enterprise end users. Services can include workflow, general office, and security applications.

False positive: Benign system activity mistakenly reported as malicious.

Firewall: Stateful packet-filtering device that maintains state tables for IP-based protocols. Traffic is allowed to cross the firewall only if it conforms to the access-control filters defined, or if it is part of an already established session in the state table.

Intrusion Prevention System (IPS): Host-based system that monitors activity on an individual host and can intervene when a given action falls into a pre-defined category of non-permissible actions. These actions can include system calls, modification of files, or login attempts.

Management server: Device that provides network management services for the operators of enterprise networks. Services can include general configuration management, monitoring of network security devices, and operation of the security functions.

Misuse detection: Detection based on whether system activity matches a pattern or signature defined as bad.

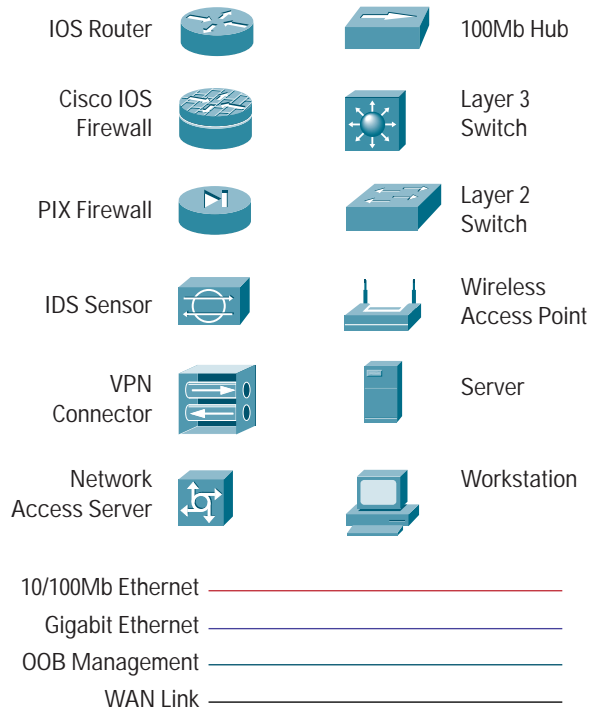
Network-Based Intrusion Detection System (NIDS): Network intrusion detection system. Typically used in a non-disruptive manner, this device captures traffic on a LAN segment and tries to match the real-time traffic against known attack signatures. Signatures range from atomic (single packet and direction) signatures to composite (multi-packet) signatures requiring state tables and Layer 7 application tracking.

Router: A wide spectrum of flexible network devices that provide many routing and security services for all performance requirements. Most devices are modular and have a range of LAN and WAN physical interfaces.

Signature: Patterns that indicate misuse of a system and define the characteristics of an attack on a system.



Diagram Legend



References

RFCs and Drafts

This section provides a list of RFCs covering various technologies and practices that are discussed in this white paper. Those who wish to understand these technologies further should refer to the relevant RFCs.

RFC 2138—Remote Authentication Dial In User Service (RADIUS)

RFC 2196—Site Security Handbook

RFC 2289—A One-Time Password System

RFC 2577—FTP Security Considerations

RFC 2827—Network Ingress Filtering: Defeating Denial of Service Attacks which employ IP Source Address Spoofing

RFC 2828—Internet Security Glossary

RFC 2350—Expectations for Computer Security Incident Response

RFC 2504—Users' Security Handbook

RFC 2903—Generic AAA Architecture

RFC 2979—Behavior of and Requirements for Internet Firewalls

RFC 3013—Recommended Internet Service Provider Security Services and Procedures

RFC 3164—The BSD Syslog Protocol

RFC 3227—Guidelines for Evidence Collection and Archiving

Miscellaneous References

Bace, Rebecca Gurley, "Intrusion Detection," MacMillan Technical Publishing, Indianapolis, 2000

Convery, Sean and Bernie Trudel, "SAFE: A Security Blueprint for Enterprise Networks,"

http://www.cisco.com/en/US/netsol/ns110/ns170/ns171/ns128/networking_solutions_white_paper09186a008009c8b6.shtml, 2000

Convery, Sean and Roland Saville, "SAFE: Extending the Security Blueprint to Small, Midsize, and Remote-User Networks," http://www.cisco.com/en/US/netsol/ns110/ns170/ns171/ns128/networking_solutions_white_paper09186a008009c8a0.shtml, 2001

Partner Product References

General Cisco AVVID (Architecture for Voice, Video and Integrated Data) Security and VPN Solution Partners Information: <http://www.cisco.com/go/securitypartners>

Acknowledgments

The authors would like to publicly thank the individuals who contributed to the SAFE architecture and the writing of this document. Certainly, the successful completion of this architecture would not have been possible without the valuable input and review feedback from all of the Cisco employees both in corporate headquarters and in the field. In addition, many individuals contributed to the lab implementation and validation of the architecture. Thank you all for your special effort.



Corporate Headquarters
Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA
www.cisco.com
Tel: 408 526-4000
800 553-NETS (6387)
Fax: 408 526-4100

European Headquarters
Cisco Systems International BV
Haarlerbergpark
Haarlerbergweg 13-19
1101 CH Amsterdam
The Netherlands
www-europe.cisco.com
Tel: 31 0 20 357 1000
Fax: 31 0 20 357 1100

Americas Headquarters
Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA
www.cisco.com
Tel: 408 526-7660
Fax: 408 527-0883

Asia Pacific Headquarters
Cisco Systems, Inc.
Capital Tower
168 Robinson Road
#22-01 to #29-01
Singapore 068912
www.cisco.com
Tel: +65 6317 7777
Fax: +65 6317 7799

Cisco Systems has more than 200 offices in the following countries and regions. Addresses, phone numbers, and fax numbers are listed on the

Cisco Web site at www.cisco.com/go/offices

Argentina • Australia • Austria • Belgium • Brazil • Bulgaria • Canada • Chile • China PRC • Colombia • Costa Rica • Croatia
Czech Republic • Denmark • Dubai, UAE • Finland • France • Germany • Greece • Hong Kong SAR • Hungary • India • Indonesia • Ireland
Israel • Italy • Japan • Korea • Luxembourg • Malaysia • Mexico • The Netherlands • New Zealand • Norway • Peru • Philippines • Poland
Portugal • Puerto Rico • Romania • Russia • Saudi Arabia • Scotland • Singapore • Slovakia • Slovenia • South Africa • Spain • Sweden
Switzerland • Taiwan • Thailand • Turkey • Ukraine • United Kingdom • United States • Venezuela • Vietnam • Zimbabwe

All contents are Copyright © 1992–2003 Cisco Systems, Inc. All rights reserved. Cisco, Cisco Systems, the Cisco Systems logo, Cisco IOS, and PIX are registered trademarks of Cisco Systems, Inc. and/or its affiliates in the U.S. and certain other countries.

All other trademarks mentioned in this document or Web site are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company.
(0304R) ETMG 203142—RD 08/03