This chapter helps you to understand important information related to the functions and deployment of CS-MARS in your network. You learn the following:

- The Self-Defending Network and the Expanding Role of CS-MARS
- CS-MARS as an STM Solution

CHAPTER 2

Role of CS-MARS in Your Network

CS-MARS plays two major roles in your network's security solution. The first is as a critical component that enhances the self-defending network (SDN). The SDN increases the level of protection in your network by enabling additional communication between devices. The second role of CS-MARS is as a security threat mitigation (STM) system that enables substantially quicker and more accurate information for threat response. In addition to these roles, CS-MARS reporting is a valuable tool that can meet legislative and reporting requirements.

After examining the features that CS-MARS provides, you will see how the appliances can be deployed to achieve enhanced security and response.

The Self-Defending Network and the Expanding Role of CS-MARS

The self-defending network is a security deployment methodology that many of the nation's top security engineers recommend and support.

The basic concept of the self-defending network is that there are many layers of defenses to recognize and protect against malicious activity, and the devices in these layers can communicate with other layers to further enhance your network and device security. For example, if a device recognizes an attack in the core of your network, that device could notify the perimeter devices of the offending traffic, and the perimeter device could put protection in place to stop the attack from reinfecting the core or spreading to other parts of your network. As an added bonus, the core device would still have the capability to recognize the attack, so you could stop the attack at multiple layers.

In this section, in addition to getting a more detailed description of the self-defending network, you will learn how CS-MARS can help you take that defense to a new level. CS-MARS does this by correlating alerts from many different security devices manufactured by multiple security vendors and then recognizing and responding to attacks. CS-MARS identifies attacks accurately and quickly to help security responders mitigate attacks and contain malicious activity.

The following topics are discussed in this section:

- **Understanding the Self-Defending Network**—This section helps you to understand how the self-defending network expands on the defense-in-depth concept and works to mitigate attacks against networks, network devices, and PCs in your enterprise.
- **Enhancing the Self-Defending Network**—This section explains the missing links in the self-defending network. With these missing links addressed, the capability of the network to protect itself, critical assets, and critical applications; the capability of network engineers to achieve fast and accurate forensics; and the capability for fast, accurate mitigation of attacks are greatly enhanced.
- **CS-MARS: Filling the Gaps in the Self-Defending Network**—This section highlights how the CS-MARS product enhances not just the capability of the network to defend itself, but also the capability of security response teams to recognize an attack in progress and to accurately respond to that attack.

Understanding the Self-Defending Network

To understand the self-defending network, you must look at the foundation of the SDN, which is known as defense-in-depth, and explain how the self-defending mechanism helps automate and expand the security posture of your network.

Defense-in-Depth and the Self-Defending Network

Defense-in-depth is a multilayer model that defines layers of protection for your network. Each layer has network- and host-defense features and is capable of stopping a network or host attack. In addition to the layers of defense that build up the basic foundation of your security, the self-defending aspects automatically and intelligently learn about threats and communicate with other network devices to change their configurations and build a stronger network defense.

When you put these layers together, they not only have the specific functions for stopping an attack, but they also provide multiple chokepoints to contain malicious activity and keep it from spreading throughout your network.

At a high level, defense-in-depth defines four main layers of protection for your network and an abstract layer that encompasses security best practices. These layers are as follows:

- Authentication layer
- Perimeter layer
- Network intrusion-prevention layer
- Host intrusion-prevention layer
- Security best practices

We start with a discussion on these layers of protection and then discuss how the self-defending network is used to augment defense-in-depth.

Authentication Layer

The basic description of authentication is that it employs user or device credentials before allowing access to your network or to devices in your network.

Authentication is possible only if a protocol or application is designed to accept and track usernames, passwords, or certificates.

Commonly used protocols that allow authentication include the following:

- **IPSec**—For remote network and remote management access
- **SSH**—For remote management access
- **HTTPS**—For remote management access
- **HTTP**—For inbound and outbound web connections
- **SMTP, IMAP, and POP**—For e-mail access
- **FTP**—For file transfers
- **802.1x**—For device and user authentication

IPSec, HTTPS, and SSH encrypt traffic and use certificates and passwords to authenticate devices. This encryption mitigates against attackers who might be sniffing your network to glean these important bits of information.

You should always use authentication before allowing access to critical devices in your network. This includes hosts, servers, routers' switches, firewalls, load-balancers, infrastructure servers such as DNS and DHCP, and IP telephony equipment. To deploy centralized authentication on your network, you normally use a product such as the Cisco Secure Access Control Server (Cisco Secure ACS), which provides a common authentication database and works in conjunction with most devices that use the protocols listed earlier. In addition to providing authentication to local usernames and passwords, Cisco ACS can forward authentication requests to domain databases, LDAP databases, and UNIX and Microsoft password databases.

In many cases, you want to allow anyone on the Internet to access data on your web server, such as product and marketing information. Because of this, most transactions from the Internet to the web server are not authenticated, meaning that anyone can get to your web server to access public information.

Take note that although allowing anyone into a web server on your network is the correct thing for your business, it opens the door for hackers who use this access to exploit vulnerabilities in your web server, access your network, and launch other attacks or view critical information. Because of this, you must be sure to deploy the additional layers of the defense-in-depth model in your network.

In the Cisco world, network admission control (NAC) is used in conjunction with network-access devices such as these:

- Routers
- Switches
- VPN concentrators
- Wireless access point

These devices act similarly to authentication proxies and authorization proxies. They forward security posture information about your device to authentication servers to see if it has security software installed that meets your network-access security policies before giving you access to the network. Based on what is learned during the NAC process, your network devices either grant access, grant limited access, or deny access to the device.

The intelligence is also built into NAC so that if a device changes its security posture after it has authenticated, NAC will discover this and take appropriate action, such as denying the device network access.

The dynamic nature of NAC and its capability to learn and respond classify it as a self-defending technology.

Some network experts believe that NAC falls into either the defense-in-depth category of host security or authentication, as we have defined it in this book. It really has aspects of both of these layers of defense, and where it should be defined depends on an individual's viewpoint.

Perimeter-Layer Defenses

After users or devices have been authenticated (or not authenticated, as is the case for most web traffic), the next step is to determine what these users can or cannot do after they access the network. This level of access is applied at the perimeter layer.

The perimeter layer has two main functions:

- Traffic filtering
- Network perimeter-attack protection

The first function of perimeter protection, traffic filtering, enforces rules that define what traffic is allowed into the network. This traffic is defined by your security policy. Traffic filtering is deployed to ensure that outside users have access to only devices and services that you have defined.

These are common types of Internet traffic that enterprises allow in their networks:

- Clear-text web traffic (HTTP), for viewing noncritical data.
- Encrypted and authenticated web traffic (HTTPS/SSL), for secure transactions and viewing critical data.

- Domain Name Services (DNS), to handle Internet requests to translate domain names that you own to IP addresses.
- Simple Mail Transport Protocol (SMTP), for sending and receiving mail traffic to and from Internet users to your internal users.
- File Transfer Protocol (FTP), for transferring files between Internet hosts and file servers on your network. Although it is not as prevalent as it was in the early days of the Internet, FTP is still used enough to warrant mentioning.

The second function of network perimeter-attack protection is to defend against attacks on the perimeter of your network. These are examples of the most popular of those attacks:

- Denial-of-service attacks
- Unauthorized perimeter device access
- Application attacks
- Worm and virus propagation

These types of attacks are usually recognized and mitigated using various security devices, protocols, and techniques.

Denial-of-service attacks (DoS or DDoS) are generally recognized and mitigated by devices such as firewalls, routers, and specialized devices such as Cisco Guard. These various security and network devices use a combination of techniques to defend against denial-of-service attacks:

- Anomaly-based traffic recognition
- Algorithms that look at excessive flows
- Signature-based systems that recognize an excess of traffic conditions that indicate a DoS attack in progress
- Trigger points that recognize more traffic than normal passing onto the network

Denial-of-service attacks are mitigated by various device actions. Routers use rate-limiting, a technology that limits the flow of data that has been classified as malicious. Firewalls examine their state tables and tear down flows that are deemed malicious. Cisco Guard devices use a combination of rate-limiting and filtering technology to determine valid traffic and stop malicious traffic. IPS uses signatures and anomaly algorithms to determine whether an attack is underway and can report or drop the traffic.

Defending against DoS attacks is a good example of how defense-in-depth can be improved, as discussed in the section “Enhancing the Self-Defending Network.” In many cases, these devices need to be manually configured and can’t change dynamically. If you have a true self-defending network that can recognize behavior and share information to stop bad behavior, the solutions just mentioned become much more attractive and cost-effective.

Unauthorized perimeter device access is an attack that becomes possible mainly when firewalls are not correctly configured. For example, during the Slammer attack, a buffer overflow was caused by sending a crafted packet to exploit a vulnerability using TCP port 1434. Under normal circumstances, perimeter devices would block that port, but as it ended up, several thousand perimeters were opened to that port, allowing hundreds of thousands of hosts and servers to be exploited by the Slammer worm. That being said, one of the main purposes of perimeter protection is to ensure that only traffic that matches an enterprise's access security policy should be let into the enterprise network.

Note that the two types of protection provided by the perimeter layer of defense-in-depth will potentially stop thousands of attacks. But if a hacker is launching an attack that uses valid data and there is no mechanism to stop the attack, perimeter security won't stop the attack; the attack will make it to your next level of security defense.

Application attacks are attempts by hackers to use existing application protocols and pass attacks through perimeter devices. Most firewalls and perimeter security devices now have application-inspection protection engines. These devices look for violations of well-known applications such as HTTP, FTP, or SMTP. If a violation is detected, the device can stop or report the attack before it passes into the network.

It's important to recognize that the attacks we've discussed in this section are not always sourced from Internet attackers. Other sources include the following:

- Employees who have accidentally (or not accidentally) had malicious software installed on their machine
- Employees who are connected to your campus network and have picked up malicious traffic from Internet download, spyware, or e-mail
- Disgruntled employees who launch attacks from inside your network
- Misconfiguration of hosts or network devices
- Trusted networks such as VPN termination points

The following are descriptions of devices that provide protection at the perimeter:

Appliances that defend against DDoS attacks are good examples of devices that adapt to a learned threat and then change their posture to more accurately protect against that threat. For example, Cisco Guard works in conjunction with Cisco Traffic Anomaly Detectors to learn what the normal flow of traffic looks like. If that traffic varies a certain amount from the norm, Cisco Guard will analyze that traffic and then take defensive action against the threat by modifying its own tables to recognize the valid traffic and let only that traffic into your network. You can learn more about Cisco Guard from Cisco's website at http://www.cisco.com/go/guard.

Firewalls are the most common of perimeter devices. Traditionally, firewalls don't learn that a threat is active and change their configuration; they do, however, have the capability to accept commands from IPS devices called shuns. As of IPS v5.1, Cisco IPS devices also have the capability to recognize DoS and DDoS traffic and send commands to a router to rate-limit that traffic; this helps to reduce or eliminate the denial-of-service attack. IPS v5.1 also can accept signatures from the Cisco Incident Control Server to stop new Internet threat outbreaks.

Although Cisco Incident Control Server (ICS) is more of a manual process, it is another excellent example of a security service that sends learned threat information to devices to strengthen their security posture if a threat is recognized. The process for ICS starts manually. When a new threat is detected on the Internet, the ICS server is populated and sends access-control lists and signatures that are crafted to stop the new threat. The ICS server can send data to Cisco routers, firewalls, and IPS devices.

Network Intrusion Prevention

Up to this point, with the first two layers (authentication and perimeter) of defense-in-depth, you have effectively done the following:

- Granted access only to desired users
- Enforced rules specifying what traffic will traverse your network
- Provided protection against many perimeter attacks
- Verified device security posture before allowing network access
- And in many cases, changed (strengthened) your security posture based on learned threats

All this protection probably sounds good. But the problem you run into is that hackers are adjusting their exploit methods to use valid traffic and follow valid protocol standards. Because of this, you need the next layer of defense-in-depth, called network intrusion prevention.

The purpose of this layer is to look inside the traffic that you have allowed after you have applied all the previous defenses to your traffic.

This level of protection is normally achieved with one or more Cisco security devices:

- Cisco Intrusion Prevention (IPS) appliances
- Cisco Intrusion Prevention Catalyst Service Modules
- Cisco ASA appliances with Security Service Modules running IPS software
- Cisco integrated security routers

These devices recognize attacks using various different technologies:

- **Signature matching**—The traffic is matched against attack signatures. If the device finds an attack, it takes whatever protective action you have defined for that signature.

- **Anomaly detection**—Using this technology, a device establishes a baseline for normal traffic. If traffic starts to flow outside that baseline, the device takes whatever protective action you have defined for this type of alert. Anomaly protection is an effective technology to stop mutating and scanning day-zero worms.
- **Application inspection**—This is the same security feature described for the perimeter layer. IPS devices and perimeter devices are capable of stopping traffic that violates a well-known protocol. It is often desirable to do this protection in the IPS layer and free up CPU cycles on the perimeter to filter traffic and protect against high-bandwidth attacks such as DoS and DDoS.

If one of the devices mentioned earlier identifies an attack, you have the option to configure the device to either drop or drop and report the traffic. Users who elect to drop the packet must make sure that they are not dropping valid packets; therefore, the signatures shipped with the security appliance are well-known attacks and leave very little chance of valid traffic being dropped.

The main problem with signature-based network intrusion prevention is that it's only as good as the last attack. That means that signatures can stop only known attacks. Day-zero attacks, or new attacks, pose the greatest threat to network and host security. Because of this, host intrusion prevention, which protects against bad behavior prevalent in day-zero attacks, is the next layer of defense-in-depth that should be deployed.

Host Intrusion-Prevention Layer

Even though you have deployed the first three layers of defense-in-depth, the possibility exists that attack traffic that does not match these signatures or behaviors can pass through to the inside of your network. Host intrusion prevention (HIPS), the last layer of defense-in-depth, is designed to stop the remainder of attacks.

Host intrusion prevention is designed to stop the following:

- Any attack that doesn't traverse the perimeter security appliances as described in the previous layers of defense-in-depth. An example is an attack sourced by users located inside the perimeter.
- Any attack that was sourced from the outside of the security appliance but wasn't stopped by the security appliance filters or the application firewall. An example is traffic that follows network or application protocol but has an exploit built into the payload.

HIPS can stop these attacks that have bypassed signature and anomaly detection because it looks for behavior on a host and stops behavior that it recognizes as being malicious. Malicious behavior includes the following:

- Browsers acting as servers listening for incoming connections from the Internet
- Browsers trying to write data to a disk besides log, cookie, or history files

- Browsers trying to install software
- Processes trying to execute code that has been written to a system or application data stack following a buffer overflow
- Unauthorized processes or applications attempting to install software, write to the system directory, or modify the system registry

At the HIPS layer, Cisco Security Agent (CSA) is an excellent example of a self-defending technology. Not only does it protect against unknown attacks based on the behavior that it observes at the endpoint, but it also dynamically builds rules when it recognizes malicious behavior on a host and then downloads those rules to other hosts running Cisco Security Agent software. In addition, it can recognize scans against hosts and build rules to prevent those scans on other hosts.

This presence on the endpoint and the accuracy of the alerts make CSA a unique and valuable source of data to provide information to other network devices and enhance the self-defending network to a new level.

If NAC is used in conjunction with network-access devices, it can ensure that this critical layer of protection (HIPS) is activated in your network.

Security Best Practices

Even if defense-in-depth has been applied in your network as described in the previous sections of this chapter, you should follow certain network, host, and server security best practices to ensure additional protection.

Examples of those best practices include the following:

- Network device security posture hardening
- Host and server security posture hardening
- Layer 2 security posture hardening
- Management best practices (Chapter 1, “Understanding SIM and STM,” offers more information on management and response best practices.)
- Security response best practices
- Password management best practices

Security best practices are essential to having a secure network. Consider what potential damage can be done if you deploy a firewall as a perimeter device to allow only desired traffic into your network and then allow management access from the outside with a username and password of Cisco and Cisco. It would take about five minutes for someone to hack your firewall and completely compromise the perimeter of your network. Of course, this example describes an unlikely omission of procedure, but at every level of defense-in-depth, you need to consider best practices on how to deploy that layer to reduce the chance of a similar error.

For a detailed discussion on security best practices, refer to the “SAFE Enterprise Architecture” whitepaper on the Cisco website at http://www.cisco.com/go/safe.

Enhancing the Self-Defending Network

The defense-in-depth paradigm has worked quite well for several years to defend against worms and mutating viruses, in addition to attacks launched by inexperienced script kiddies.

But there are still some missing links to the self-defending network:

- **Automated log correlation**—Provides a single source for log correlation
- **Automated threat response**—Automatically learns your network topology, analyzes security alerts, and notifies security responders with up-to-date accurate threat information
- **Automated mitigation**—Automatically evaluates threats and recommends a mitigation action to your security responders that will stop or contain the attack in your network

These items are mostly in the area of automated threat response. Automated threat response is simple in concept but very complicated in delivery. The concept is simply that when an attack occurs, it needs to be automatically recognized and verified by your security devices; then a response action needs to be taken to mitigate or contain that threat.

This area of the self-defending network has been relatively ignored until recently by most security vendors and customers alike. From a customer perspective, there has always been a mentality that security equipment should be installed and then just work. Unfortunately, new threats and vulnerabilities occur on a daily basis. Because of this lightning-fast changing environment, security gear needs to be capable of responding in kind. Log correlation, threat response, threat mitigation, and threat containment need to happen automatically as much as possible. The remainder of this chapter is dedicated to those concepts and also to how CS-MARS works to help automate these tasks.

Automated Log Correlation

Customers commonly have several different vendors’ equipment in their networks to perform specific best-of-breed operations. Although this sometimes provides enhanced security, it commonly causes problems related to logging security events and alerts such as correlating logs between different systems. Most customers end up not reading or even keeping logs unless they need to because of legislative requirements.

For those customers who do check their logs, security-response engineers manually look at this data and decide whether an attack actually has occurred. If they determine that it has, they begin the manual prevention process.

Automated log integration is a function that allows devices to collect and correlate logs from almost any security device, network device, host, server, or key software system on your network. Automated log integration needs to encompass not only support for Cisco devices, but also popular third-party devices.

If a device isn't supported by a logging server, it should have the flexibility to define custom log parsing for any device that generates syslog or SNMP data.

The logging integration server should have the capability to not only collect, normalize, and correlate all logs, but also to classify and analyze each log that it receives.

This functionality is the first step in giving security responders the tools they need to increase the accuracy and speed at which they respond to threats.

Automated Threat Response

Another significant hurdle to the self-defending network is the capability to automatically recognize a threat and provide as much data as possible about that threat. An example is a system that would report to security responders that a successful attempt was made to exploit your web server at address 10.1.1.1, that the threat was sourced from address 192.169.1.10, that the source device is located in your network data center, and that the path the attack took was through the Internet firewall to the core and finally to the data center.

If a security responder had all this information and the confidence that the alerts were not false positives, it would add substantial value to the self-defending network.

Automated Mitigation

If events and logs from all systems have been correlated and analyzed and threat information is found, providing guidance or recommendations for mitigation is the next logical step.

If you know the type of attack, the source, the destination, the location of the exploited device in your network, and the path the attack traversed, what's standing in the way of either mitigating the attack or making a mitigation recommendation? As you will see in the remainder of this book, CS-MARS mitigates attacks or makes a recommendation for you. Automated mitigation is not yet achievable because many security devices still put out false-positive alerts, but CS-MARS makes a recommendation for mitigation and offers security responders a single click to deploy commands on devices that will stop offending traffic after the responder has analyzed the attack data.

CS-MARS: Filling the Gaps in the Self-Defending Network

Many network security experts feel that the self-defending network is a powerful methodology for mitigating network-based attacks. If you add to that existing SDN model log correlation, automated threat response, and automated threat mitigation, you have a stronger, more maintainable, and more robust security solution. This solution now provides not just defense in-depth and the capability to learn and respond but it also enables your security-response engineers to streamline the process of recognizing and responding to attacks. That recognition and response is exactly the purpose of the CS-MARS STM appliance.

CS-MARS Log Integration

CS-MARS is capable of collecting, correlating, analyzing, and storing data from thousands of different systems. This includes not only security systems, but also network devices, hosts, servers, and applications.

The following is a list of the hardware devices and software applications that have reporting capabilities and are supported natively by CS-MARS. Notice the supported devices aren't only security devices; the list includes operating systems, databases, web servers, web caches, antivirus servers, vulnerability scanners, authentication servers, and SNMP servers.

- Cisco routers
- Cisco switches (IOS and CATOS)
- Extreme switches
- Generic routers
- Cisco PIX
- Cisco Adaptive Security Appliance (ASA)
- Cisco Firewall Services Module (FWSM)
- Cisco IOS Firewall Feature Set
- Juniper NetScreen
- Check Point OPSEC NG/AI and Provider-1
- Nokia Firewall (running Check Point)
- Cisco VPN concentrator
- Cisco network IDS and IDSM
- Cisco intrusion-prevention system (IPS), Network IPS v5.0
- Cisco IPS ASA Security Services Module
- Cisco IOS IPS module
- McAfee Intrushield
- Juniper NetScreen IDP
- Symantec ManHunt
- ISS RealSecure

- Snort
- Enterasys Dragon
- Cisco Security Agent
- McAfee Entercept
- ISS RealSecure Host Sensor
- Symantec AntiVirus
- Cisco Incident Control System
- Network Associates VirusScan
- McAfee ePolicy Orchestrator
- eEye REM
- Qualys QualysGuard
- Foundstone Foundscan
- Windows NT, 2000, XP, 2003
- Solaris
- Red Hat Linux
- Microsoft Internet Information Server
- Sun iPlanet
- Apache
- NetApp NetCache
- Oracle
- AAA Server
- Cisco Secure Access Control Sever (ACS)
- SNMP and syslog servers
- Generic syslog server

In addition to these systems, CS-MARS has a function that enables you to write a custom parser for devices that don't appear on this list.

NOTE Custom parsing for appliance or software event data supports syslog or SNMP logs only.

Each of the event data from these different systems is sessionized with other events to formulate information about possible attacks. The sessionized data is then weighed against vulnerability-assessment information to determine whether the attack is possible or likely. The CS-MARS device then can accurately determine the probability of an attack.

After collecting the alerts and log data, CS-MARS displays it in nicely formatted graphs or reports. Strictly from a reporting standpoint, CS-MARS has taken you from legacy syslogs and IPS alerts (see Figures 2-1 and 2-2) to automated log integration and correlation (see the CS-MARS Device Summary screen shown in Figure 2-3).

Figure 2-1 *Legacy IPS Alerts*

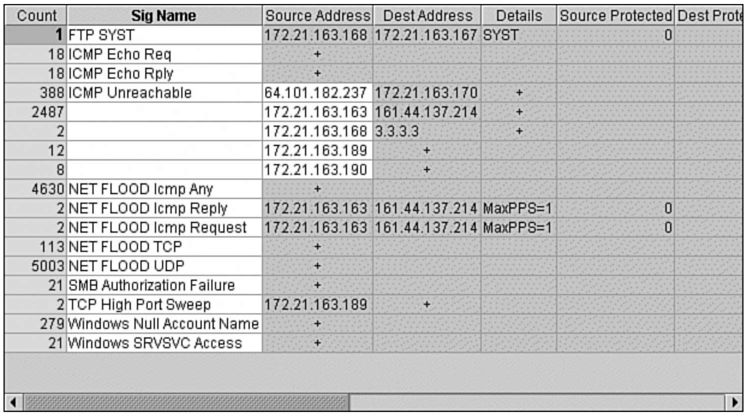

Count	Sig Name	Source Address	Dest Address	Details	Source Protected	Dest Prote
1	FTP SYST	172.21.163.168	172.21.163.167	SYST	0	
18	ICMP Echo Req	+				
18	ICMP Echo Rply	+				
388	ICMP Unreachable	64.101.182.237	172.21.163.170	+		
2487		172.21.163.163	161.44.137.214	+		
2		172.21.163.168	3.3.3.3	+		
12		172.21.163.189	+			
8		172.21.163.190	+			
4630	NET FLOOD Icmp Any	+				
2	NET FLOOD Icmp Reply	172.21.163.163	161.44.137.214	MaxPPS=1	0	
2	NET FLOOD Icmp Request	172.21.163.163	161.44.137.214	MaxPPS=1	0	
113	NET FLOOD TCP	+				
5003	NET FLOOD UDP	+				
21	SMB Authorization Failure	+				
2	TCP High Port Sweep	172.21.163.189	+			
279	Windows Null Account Name	+				
21	Windows SRVSVC Access	+				

Figure 2-2 *Legacy Syslog Alerts*

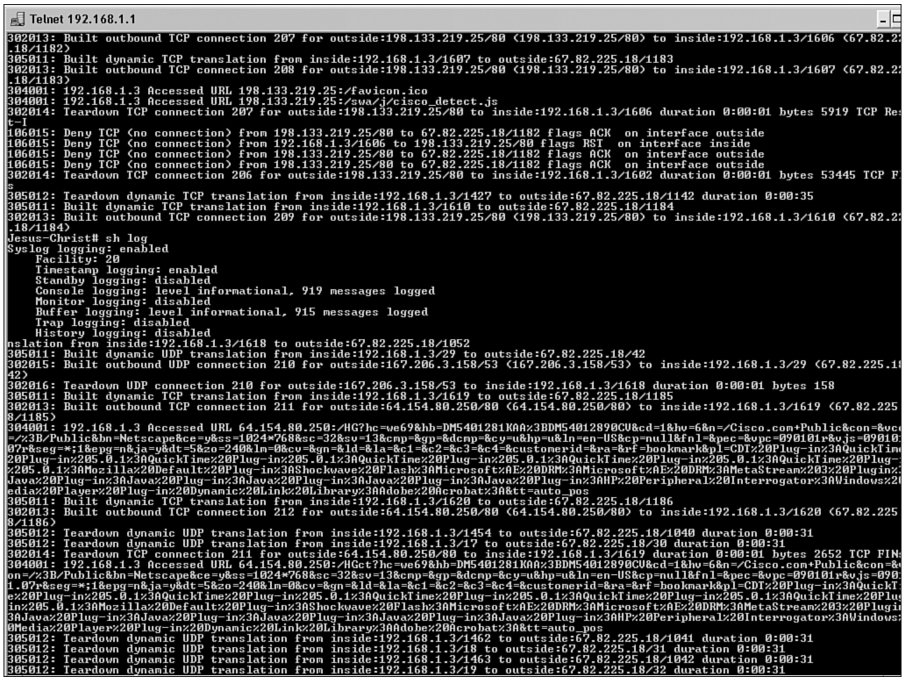

```
302013: Built outbound TCP connection 207 for outside:198.133.219.25/80 (198.133.219.25/80) to inside:192.168.1.3/1606 (67.82.2
.18/1182)
305011: Built dynamic TCP translation from inside:192.168.1.3/1607 to outside:67.82.225.18/1183
302013: Built outbound TCP connection 208 for outside:198.133.219.25/80 (198.133.219.25/80) to inside:192.168.1.3/1607 (67.82.2
.18/1183)
304001: 192.168.1.3 Accessed URL 198.133.219.25:/favicon.ico
304001: 192.168.1.3 Accessed URL 198.133.219.25:/swa/j/cisco_detect.js
302014: Teardown TCP connection 207 for outside:198.133.219.25/80 to inside:192.168.1.3/1606 duration 0:00:01 bytes 5919 TCP Re
t-I
106015: Deny TCP (no connection) from 198.133.219.25/80 to 67.82.225.18/1182 flags ACK  on interface outside
106015: Deny TCP (no connection) from 192.168.1.3/1606 to 198.133.219.25/80 flags RST  on interface inside
106015: Deny TCP (no connection) from 198.133.219.25/80 to 67.82.225.18/1182 flags ACK  on interface outside
106015: Deny TCP (no connection) from 198.133.219.25/80 to 67.82.225.18/1182 flags ACK  on interface outside
302014: Teardown TCP connection 206 for outside:198.133.219.25/80 to inside:192.168.1.3/1602 duration 0:00:01 bytes 53445 TCP F
s
305012: Teardown dynamic TCP translation from inside:192.168.1.3/1427 to outside:67.82.225.18/1142 duration 0:00:35
305011: Built dynamic TCP translation from inside:192.168.1.3/1610 to outside:67.82.225.18/1184
302013: Built outbound TCP connection 209 for outside:198.133.219.25/80 (198.133.219.25/80) to inside:192.168.1.3/1610 (67.82.2
.18/1184)
Jesus-Christ# sh log
Syslog logging: enabled
    Facility: 20
    Timestamp logging: enabled
    Standby logging: disabled
    Console logging: level informational, 919 messages logged
    Monitor logging: disabled
    Buffer logging: level informational, 915 messages logged
    Trap logging: disabled
    History logging: disabled
nslation from inside:192.168.1.3/1618 to outside:67.82.225.18/1052
305011: Built dynamic UDP translation from inside:192.168.1.3/29 to outside:67.82.225.18/42
302015: Built outbound UDP connection 210 for outside:167.206.3.158/53 (167.206.3.158/53) to inside:192.168.1.3/29 (67.82.225.1
42)
302016: Teardown UDP connection 210 for outside:167.206.3.158/53 to inside:192.168.1.3/1618 duration 0:00:01 bytes 158
305011: Built dynamic TCP translation from inside:192.168.1.3/1619 to outside:67.82.225.18/1185
302013: Built outbound TCP connection 211 for outside:64.154.80.250/80 (64.154.80.250/80) to inside:192.168.1.3/1619 (67.82.225
8/1185)
304001: 192.168.1.3 Accessed URL 64.154.80.250:/HG?hc=we69&hb=DM5401281KAA%3BDM54012890CV&cd=1&hv=6&n=/Cisco.com+Public&con=&vc
=/%3B/Public&bn=Netscape&ce=y&ss=1024*768&sc=32&sv=13&cmp=&gp=&dcmp=&cy=u&hp=u&ln=en-US&cp=null&fnl=&pec=&vpc=090101r&vjs=09010
07r&seg=*;1&epg=n&ja=y&dt=5&zo=240&lm=0&cv=&gn=&ld=&la=&c1=&c2=&c3=&c4=&customerid=&ra=&rf=bookmark&pl=CDT%20Plug-in%3AQuickTim
20Plug-in%205.0.1%3AQuickTime%20Plug-in%205.0.1%3AQuickTime%20Plug-in%205.0.1%3AQuickTime%20Plug-in%205.0.1%3AQuickTime%20Plug-
%205.0.1%3AMozilla%20Default%20Plug-in%3AShockwave%20Flash%3AMicrosoft%AE%20DRM%3AMicrosoft%AE%20DRM%3AMetaStream%203%20Plugin%
Java%20Plug-in%3AJava%20Plug-in%3AJava%20Plug-in%3AJava%20Plug-in%3AJava%20Plug-in%3AHP%20Peripheral%20Interrogator%3AWindows%2
edia%20Player%20Plug-in%20Dynamic%20Link%20Library%3AAdobe%20Acrobat%3A&tt=auto_pos
305011: Built dynamic TCP translation from inside:192.168.1.3/1620 to outside:67.82.225.18/1186
302013: Built outbound TCP connection 212 for outside:64.154.80.250/80 (64.154.80.250/80) to inside:192.168.1.3/1620 (67.82.225
8/1186)
305012: Teardown dynamic UDP translation from inside:192.168.1.3/1454 to outside:67.82.225.18/1040 duration 0:00:31
305012: Teardown dynamic UDP translation from inside:192.168.1.3/17 to outside:67.82.225.18/30 duration 0:00:31
302014: Teardown TCP connection 211 for outside:64.154.80.250/80 to inside:192.168.1.3/1619 duration 0:00:01 bytes 2652 TCP FIN
304001: 192.168.1.3 Accessed URL 64.154.80.250:/HGct?hc=we69&hb=DM5401281KAA%3BDM54012890CV&cd=1&hv=6&n=/Cisco.com+Public&con=&
on=/%3B/Public&bn=Netscape&ce=y&ss=1024*768&sc=32&sv=13&cmp=&gp=&dcmp=&cy=u&hp=u&ln=en-US&cp=null&fnl=&pec=&vpc=090101r&vjs=090
1.07r&seg=*;1&epg=n&ja=y&dt=5&zo=240&lm=0&cv=&gn=&ld=&la=&c1=&c2=&c3=&c4=&customerid=&ra=&rf=bookmark&pl=CDT%20Plug-in%3AQuickT
e%20Plug-in%205.0.1%3AQuickTime%20Plug-in%205.0.1%3AQuickTime%20Plug-in%205.0.1%3AQuickTime%20Plug-in%205.0.1%3AQuickTime%20Plu
in%205.0.1%3AMozilla%20Default%20Plug-in%3AShockwave%20Flash%3AMicrosoft%AE%20DRM%3AMicrosoft%AE%20DRM%3AMetaStream%203%20Plugi
3AJava%20Plug-in%3AJava%20Plug-in%3AJava%20Plug-in%3AJava%20Plug-in%3AJava%20Plug-in%3AHP%20Peripheral%20Interrogator%3AWindows
0Media%20Player%20Plug-in%20Dynamic%20Link%20Library%3AAdobe%20Acrobat%3A&tt=auto_pos
305012: Teardown dynamic UDP translation from inside:192.168.1.3/1462 to outside:67.82.225.18/1041 duration 0:00:31
305012: Teardown dynamic UDP translation from inside:192.168.1.3/18 to outside:67.82.225.18/31 duration 0:00:31
305012: Teardown dynamic UDP translation from inside:192.168.1.3/1463 to outside:67.82.225.18/1042 duration 0:00:31
305012: Teardown dynamic UDP translation from inside:192.168.1.3/19 to outside:67.82.225.18/32 duration 0:00:31
```

Figure 2-3 *CS-MARS Device Summary Display*

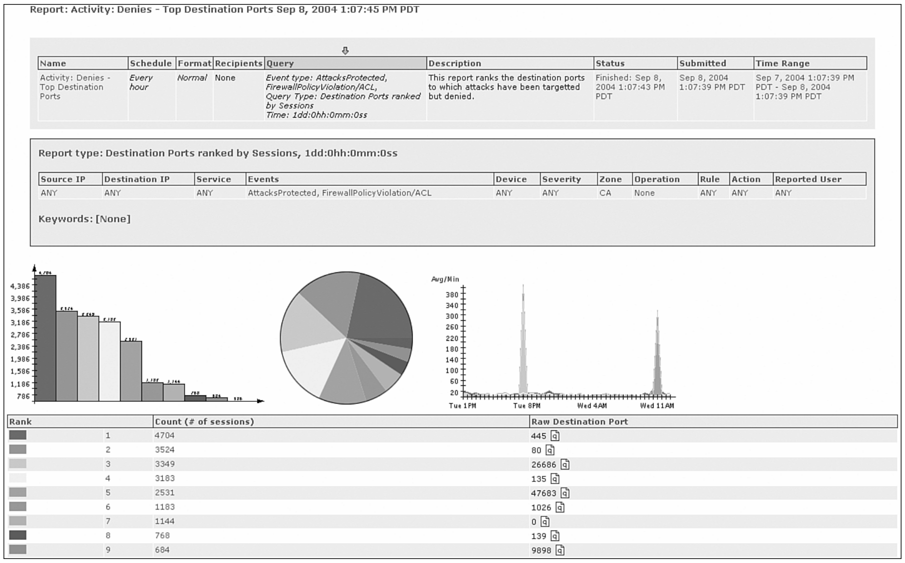

Report: Activity: Denies - Top Destination Ports Sep 8, 2004 1:07:45 PM PDT

Name	Schedule	Format	Recipients	Query	Description	Status	Submitted	Time Range
Activity: Denies - Top Destination Ports	*Every hour*	*Normal*	None	*Event type: AttacksProtected, FirewallPolicyViolation/ACL,* *Query Type: Destination Ports ranked by Sessions* *Time: 1dd:0hh:0mm:0ss*	This report ranks the destination ports to which attacks have been targetted but denied.	Finished: Sep 8, 2004 1:07:43 PM PDT	Sep 8, 2004 1:07:39 PM PDT	Sep 7, 2004 1:07:39 PM PDT - Sep 8, 2004 1:07:39 PM PDT

Report type: Destination Ports ranked by Sessions, 1dd:0hh:0mm:0ss

Source IP	Destination IP	Service	Events	Device	Severity	Zone	Operation	Rule	Action	Reported User
ANY	ANY	ANY	AttacksProtected, FirewallPolicyViolation/ACL	ANY	ANY	CA	None	ANY	ANY	ANY

Keywords: [None]

Rank		Count (# of sessions)	Raw Destination Port
■	1	4704	445
■	2	3524	80
■	3	3349	26686
□	4	3183	135
■	5	2531	47683
■	6	1183	1026
■	7	1144	0
■	8	768	139
■	9	684	9898

CS-MARS Automated Threat Response

In parallel with the data collection by the logging integration processes, CS-MARS is also querying network device routing tables, configurations, ARP tables, CAM tables, system probes, and other processes to determine the topology of your network and the location of each device.

After the log data is collected and the alert information is analyzed, it is cross-referenced with this topology information to determine the validity and calculate the attack path. CS-MARS has accurate topology and attack information, and can display this to security responders in both report and topology illustrated format.

Now, in addition to a strong security defense, you have accurate and powerful data that your responders can use for threat response.

Figure 2-4 is an example of a map generated by CS-MARS to show the topology of an attack.

Figure 2-4 *CS-MARS Attack Path Topology Map*

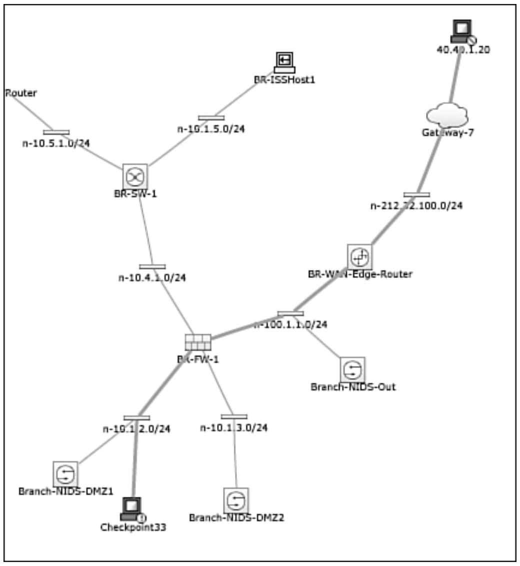

CS-MARS Automated Mitigation

CS-MARS has collected data and presented you with the information about the threat. It is now ready to take threat response a step further. Because CS-MARS has topology information and accurate threat information, it can easily determine the source of the attack and the destination of that attack. Because it has routing information, it also can determine the path that the attack took. With all this information, it's just a matter of CPU cycles to present the attack path to your security responders and to identify a network or security device that should be configured to mitigate or contain the attack. To isolate an attack and keep it from spreading, normally the network or security device closest to the source of the attack is selected for mitigation; the remainder of the components in the path of the attack are suggested as alternatives.

CS-MARS uses either SNMP, Telnet, or SSH to look at the configuration of the device it has chosen to mitigate the attack. It uses its analysis of the attack data and the device configuration to determine the commands that need to be entered to stop the attack. This mitigation suggestion shows up on the mitigation screen of the CS-MARS device. The security responder would look at the incident report, view the attack analysis, click the Mitigation icon, and select the suggested configuration changes. If the responder chose to mitigate, he or she then would click the red Push button, instructing CS-MARS to issue the suggested command to stop the attack.

Figure 2-5 is an example of a screen recommending a configuration that will mitigate an attack detected by CS-MARS.

Figure 2-5 *CS-MARS Mitigation Window*

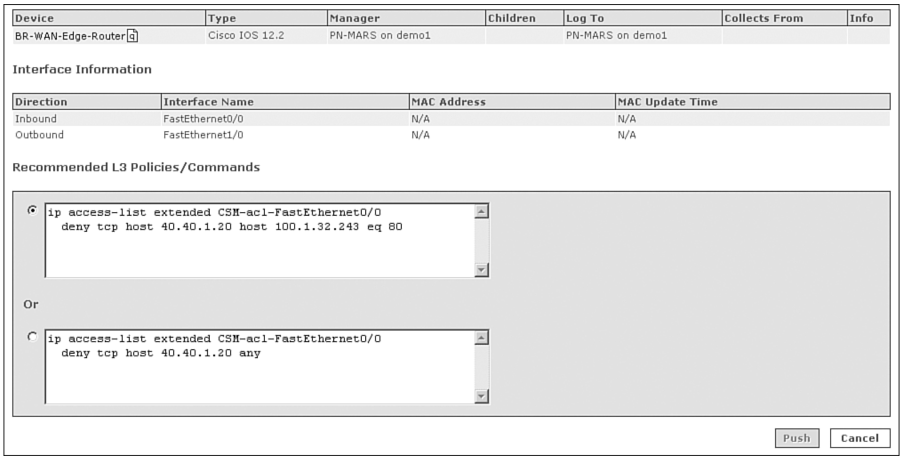

Device	Type	Manager	Children	Log To	Collects From	Info
BR-WAN-Edge-Router	Cisco IOS 12.2	PN-MARS on demo1		PN-MARS on demo1		

Interface Information

Direction	Interface Name	MAC Address	MAC Update Time
Inbound	FastEthernet0/0	N/A	N/A
Outbound	FastEthernet1/0	N/A	N/A

Recommended L3 Policies/Commands

●
```
ip access-list extended CSM-acl-FastEthernet0/0
  deny tcp host 40.40.1.20 host 100.1.32.243 eq 80
```

Or

○
```
ip access-list extended CSM-acl-FastEthernet0/0
  deny tcp host 40.40.1.20 any
```

CS-MARS as an STM Solution

This section explains the advantages CS-MARS provides beyond interaction with network and security devices. As discussed in Chapter 1, CS-MARS is a powerful solution that provides the necessary features of STM.

Reasons for an STM

In Chapter 1, you learned the benefits of moving from SIM to STM technology. Now let's look at a few of the specific advantages of an STM and take that first step to a proactive security framework. The following are some true examples of how STM gives you the advantage. The company names are fictitious, but the details of the stories are true.

Day-Zero Attacks, Viruses, and Worms

Think about Sasser, one of the more recent, notorious, fast-moving day-zero worms, before it had its name. What it did and what was affected were not apparent until the damage was already done.

You probably didn't even know it existed because there was not a known vulnerability for it to exploit. Network infrastructures were pummeled by the malicious code overnight, and 24 hours later, the worm was given a name. Now picture a system that informed you that a host on your network was behaving abnormally and opening hundreds of connections to other hosts. That system then allowed you to shut down the port that the offending host was connected to, giving you the switch name and specific interface. Because this is not normal behavior for a host on your network, you pushed the "red button" and stopped it. 24 hours later, an AV or HIPS signature was released, and news about the new worm spread. Now it's time to update those hosts with patches, and you're not scrambling around trying to do damage control.

CS-MARS stopped Sasser exactly as described. It recognized the attack using its anomaly algorithms and provided security responders with a timely and accurate mitigation suggestion, saving customers hundreds of thousands of dollars in tangible recovery costs.

Monitoring and Enforcing Security Policy

Widgets, Inc., just approved an addendum to its existing security policy and mapped out a new VLAN for its wireless access network. According to policy, users must authenticate to a Cisco VPN concentrator and tunnel into the LAN to use wireless. If a user does not have a Cisco VPN client, all it gets is a default gateway to the Internet. The policy calls for periodic audits of user access to the wireless segment. Security personnel are overwhelmed as it is with day-to-day responsibilities; however, because they have CS-MARS installed on their network, they just have reports automatically generated according to policy and sent via e-mail to the auditors.

Since the security department has configured the CS-MARS for the automated report, it has not had to spend any time on the new addendum. This has freed up the security engineering resources to work on security-related problems.

Insight, Integration, and Control of Your Network

At approximately 3:00 p.m., ABC University's help desk started to get trouble calls about failed access to servers on its network. The help desk immediately notified the network engineer on staff; she was at her desk eating lunch. When she received the call, she attempted to access the servers in question and confirmed that they could not be reached. She then attempted to access other devices from her machine and had no issues. She opened her web browser and decided to log in to her new CS-MARS appliance. She noticed a spike in port 53 activity and immediately began investigating. She discovered that her secondary DNS server on the private university network was going haywire. She made the decision to use the CS-MARS to shut it down. When the DNS server was offline, access to the other servers in the same segment was restored.

It took her 3 minutes to get it under control and 5 minutes to find out what was going on. She used CS-MARS to discover that several crafty students had hacked into the DNS server and placed a homemade Doom relay application using port 53 to circumvent the firewall policy that prevented Internet Doom contests. The students were identified and reprimanded.

The network engineer finished her lunch without getting up from her desk, and the university saved thousands of dollars in post-attack research and other potential recovery activities.

Auditing Controls

Investment firm Y is a public company and, therefore, must conform to SOX auditing. The third-party auditing firm requested access logs from its Microsoft servers, COBIT DS 9.4—Configuration Control, and Successful Object Modification logs from their Oracle database servers. The auditing company requested that the reports be prepared for review on the third day of the audit. The day the auditors arrived on the premises they were greeted, introduced to the staff, and handed the reports in a nice three-ring binder. At the end of the audit, the auditors commented on how well prepared the investment firm was and praised them on their efficiency.

It took the security department one hour to run, print, and bind the reports. This left them free to focus on security-related activities.

Monitoring Access Control

Bob is an employee of Company X and finished his project a little early on Friday. He had time on his hands before the end of the day. Bob decided to click his My Networks icon on his desktop and look around. He ran across an HR file server and clicked the icon.

Eric, the IT admin for HR, had a long day yesterday. The director for HR was having issues with getting access to the file server, and Eric just couldn't get permissions to work. To save time and embarrassment, he just set permissions to All.

Bob ran across a file marked "John D. offer letter." John D. is Bob's new teammate, and it was too irresistible not to open the file. He opened the document and, to his amazement, discovered that John earns $10,000 more per year than him and negotiated an extra week of paid vacation.

This is a good example of an intangible cost of bad security. Bob was upset and might have acted inappropriately based on what he saw, not to mention that he could have searched for other, more interesting data on the HR server. Because he was granted access to the server, there might be no audit trail, and Bob becomes a potential problem. With CS-MARS deployed appropriately, a single report showing access to the HR database would have been a great forensics tool for threat responders to use to recognize unauthorized access.

Using CS-MARS to Justify Security Investment

Eric was a one-man security shop for his company. He was again tasked with helping his ISO/IT director create the security budget for next fiscal year. Eric had his challenges last year asking for money to purchase IPS. He was given a small portion of his requested budget and told that, from a financial perspective and with the lack of security breeches on the network, they could not justify purchasing IPS. With a little investigation and help from a security analysis firm, he was advised to look at STM. Eric convinced his ISO to use the budget money to purchase CS-MARS and a server so he could use free-ware Snort on the demilitarized zone (DMZ) where clients access their extranet. With the CS-MARS in place, Eric used the existing network to send NetFlow data to the CS-MARS and correlate with the basic Snort event data from the extranet. He was able to identify numerous security threats to the network and customers' information in a matter of only one week.

Coincidently, Eric now manages three IPS sensors, two CS-MARS 100s, and a CS-MARS Global Controller. It's still the same budget year, and the company's sales media now coin a catchy security phrase identifying strong security as a reason people should do business with the company.

The STM Deployment

Chapter 1 defined the requirements for an ideal STM as a reporting and mitigation system that reduced the time and increased the accuracy of threat mitigation, threat containment, and threat reporting.

STM can be deployed many ways in your network; the choice is determined by the requirements of your organization. For example, if your organization has several remote locations and lower-bandwidth access portals, you might choose to locate a CS-MARS box in a remote location, to reduce the amount of traffic on that slow link and ensure that the network link is available for your company's revenue-based tasks.

How you deploy the CS-MARS in your network is critical to the success of achieving your organization's goals. With the CS-MARS product, there are two types of deployment scenarios: global and standalone.

To fully understand these deployment scenarios, you must first be familiar with the CS-MARS product line. At the time of this writing, Cisco offers two types of products in the CS-MARS portfolio:

- **Global Controller**—A master unit that allows for global management of one or more Local Controllers.
- **Local Controller**—A single appliance, ranging from a CS-MARS M20 to CS-MARS M200

Table 2-1 explains the Cisco offerings for the CS-MARS product family.

Table 2-1 *CS-MARS Product Portfolio*

CS-MARS	EPS	NetFlow per Sec	Storage
M20	500	15000	120 GB*
M50	1000	25000	120 GB
M100e	3000	75000	750 GB
M100	5000	150000	750 GB
M200	10000	300000	1000 GB
MARS GC	—	—	1000 GB
MARS GCm	—	—	1000 GB

*The CS-MARS 20 does not have RAID storage.

NOTE Note that there are two GC offerings. The MARS GC has an unlimited license that allows any number of CS-MARS Local Controllers to communicate with it. The MARS GCm allows only five Local Controllers, mixed between CS-MARS M20s and M50s, to communicate with it.

A global deployment simply means that one or more Local Controllers are reporting to the CS-MARS Global Controller. In this deployment, Local Controllers report summarized event and session data to the Global Controller in both text and graphical format over an HTTPS session. Additionally, all operations in the Local Controller now become globally manageable. A Global Controller does not do global correlation—that is, the data from each Local Controller is not correlated. You need a global CS-MARS deployment for several reasons:

- To conserve WAN bandwidth
- To log data security
- To facilitate distributed processing of event data
- To facilitate distributed management and reporting
- For high availability and to archive log retention

In a standalone deployment, all event-reporting devices send their respective log data to a single CS-MARS device. All capabilities discussed in this text are the function of the Local Controller, unless specifically indicated otherwise. This deployment is the most common for small to medium-size businesses. These are some reasons for deploying a single Local Controller:

- Cost
- Isolated (non-WAN) or local network with Internet or VPN
- Minimal number of reporting devices

Summary

This chapter explained defense-in-depth combined with the self-defending network.

The layers of the self-defending network are the following:

- Authentication layer
- Perimeter layer
- Network intrusion-prevention layer
- Host intrusion-prevention layer
- Security best practices

Table 2-2 outlines some the different network devices and their capability to self-defend inside your network.

Table 2-2 *Network Devices and Self-Defending Capabilities*

Cisco firewalls and ASA appliances	Accept and apply commands called shuns that stop traffic flows that Cisco IPS devices have identified. Accept and apply access control lists that Cisco Incident Control servers have generated, to block new network outbreaks such as high-priority worms and viruses. Send syslog files to CS-MARS for correlation and analysis to be used with syslogs and events from other security servers. CS-MARS uses this data to determine threat conditions and to formulate the correct response to that threat. Send SNMP data to CS-MARS to report high CPU utilization conditions, enabling CS-MARS to take defensive action to protect the CPU that might be getting attacked. Send critical data to CS-MARS to allow for network topology discovery.
IPS appliances, IPS Service Modules, ASA Security Services Modules, and integrated security routers running IPS	Send Security Device Event Exchange (SDEE) alerts to CS-MARS for correlation and analysis to be used with syslogs and events from other security servers. CS-MARS uses this data to determine threat conditions and to formulate the correct response to that threat. Recognize attacks and send shuns to firewall and Cisco IOS devices, to protect against malicious flows. Recognize attacks and send commands to rate-limit malicious traffic. Recognize attacks and drop traffic in-line to protect network assets of both hosts and network devices. Analyze destination hosts to determine the probability of an attack succeeding. Send critical data to CS-MARS to allow for network topology discovery.

Table 2-2 *Network Devices and Self-Defending Capabilities (Continued)*

Host intrusion-prevention technology (CSA)	Recognizes and stops bad behavior on a host or server. Updates itself with globally correlated data and then automatically creates and deploys resulting rules that will stop security outbreaks, network scans, and hacker reconnaissance activity. Kills applications that are behaving badly. Sends alerts to CS-MARS for correlation and analysis to be used with syslogs and events from other security servers. CS-MARS uses this data to determine threat conditions and to formulate the correct response to that threat.
Cisco Network Admission Control	Works with routers, access points, VPN concentrators, and switches to stop hosts from accessing your network if those hosts do not have the proper security posture. Takes protective action and can shut down a Layer 2 port if it's determined that a host is behaving badly. Sends alerts to CS-MARS for correlation and analysis to be used with syslogs and events from other security servers. CS-MARS uses this data to determine threat conditions and to formulate the correct response to that threat.

CS-MARS extends the self-defending network by providing a much-needed layer of automated threat identification and response.

The following features of CS-MARS were discussed:

- **Automated log integration**—Provides a single source for log aggregation
- **Automated threat response**—Automatically learns the network topology, analyzes security alerts, and provides up-to-date accurate threat information.
- **Automated mitigation**—Automatically evaluates existing threats and recommends a mitigation action to security responders that will stop or contain the threat in the network.

Now that you understand the role that CS-MARS plays in your network from a technical or engineering standpoint, you examine in the next chapter how this technology can result in cost savings.