

## Securing CS-MARS

---

A Security Information Management (SIM) system can contain a tremendous amount of sensitive information. This is because it receives event logs from security systems throughout a network. These logs potentially contain information that can be used to target attacks at sensitive systems. For example, intrusion detection system (IDS) logs can contain actual packets seen on the network. Some of these packets can be decoded with freely available packet analyzers to find usernames and passwords that your employees might be using to access websites, e-mail systems, and network devices.

Although security people always encourage users to select unique passwords for company networks, the reality is that many users tend to reuse passwords both for work and home activities. If an employee has decided to use his work network password as his personal web-based e-mail password, if an attacker discovers the cleartext authentication for web e-mail, he has also discovered an account on your network in which to begin nefarious activities.

As a topology-aware SIM product, the Cisco Security Monitoring, Analysis, and Response System (CS-MARS) often contains even more sensitive information. The most accurate method of maintaining the network topology awareness within MARS is by discovering each network device. This involves configuring access information for MARS to authenticate to the devices, retrieve interface information, and periodically rediscover it. From within the user interfaces, both the command-line interface (CLI) and web user interface, device authentication information is masked to prevent anyone from using the console to gain unauthorized information. However, if an attacker gains access to the base operating system, or gains physical access to the appliance, he could use that access to retrieve all information contained on the hard drives, which could include device authentication information. He can also use that access to install back doors to allow remote access at any time.

This chapter describes recommendations for securing MARS appliances, both physically and electronically. It also provides detailed insight into the TCP and User Datagram Protocol (UDP) ports that MARS requires for communication with other MARS appliances, in addition to monitored security, network, and other devices.

## Physical Security

You cannot properly address network security without also addressing physical security. This is evident with common sense and in the various regulations addressed in Chapter 2, “Regulatory Challenges in Depth.” All the network security in the world is worthless if someone with malicious intent can gain physical access to the target.

Make sure that the hosts on your security management network, and MARS specifically, reside in a protected facility. At the very least, they should be locked in a room that is inaccessible to the public and staff without a specific business need. Ideally, security management resides in a datacenter that exercises strong controls. Staff with access rights to the facility need to have a security badge and need to sign in, either on paper or electronically, before entering. In Chapter 2, the Payment Card Industry (PCI) data security standard has good recommendations that datacenters everywhere should attempt to adhere to, even if your facility is not affected by PCI requirements.

## Inherent Security of MARS Appliances

Management access to all MARS appliances is through Secure Socket Layer (SSL)–encrypted web access (HTTPS) and Secure Shell (SSH). These protocols, using TCP/443 and TCP/22, respectively, are inherently secure because they use encryption, authentication, and authorization. Unencrypted protocols that serve similar functions, such as HTTP and Telnet, are both disabled on the MARS appliance and cannot be enabled.

MARS appliances are hardened Linux servers that run a variety of services, including Oracle, Apache HTTP Server, and more. With each software update, the various services and drivers on MARS are updated with new versions or patches to mitigate against any newly discovered vulnerabilities. Additionally, unnecessary or unused services are disabled to prevent them from being potential weaknesses in the security of the appliances.

This hardening of the operating system provides a good starting level of security. However, it is not enough. You need to take into consideration the sensitivity of the information contained on the MARS appliances when considering how secure the appliances should be. You should have a well-defined written plan for preventing MARS from being used as an attack vector on your network. This includes placing the appliance in a part of your network that is protected from the rest of the network by a firewall and an IDS.

Without protecting MARS with a firewall and IDS or intrusion prevention system (IPS), a hacker can try to find vulnerabilities, either in the management protocols or in other protocols that are used to monitor security or network devices. The additional protections provided by the firewall or IDS/IPS allow you to limit the exposure to attacks while also creating an audit trail of attempted attacks.

As an example, consider SSH, the command-line method of administering MARS remotely. In the past, a number of vulnerabilities have appeared in the OpenSSH

application, which provides this service for MARS. No known vulnerabilities exist in the SSH service that MARS uses at this time. However, at some time in the future, a new vulnerability might be found. For this reason, it makes sense to restrict the capability of computers to establish an SSH connection to MARS unless they are connected to a specific network or set of networks at your location. A stateful inspection firewall is the ideal device for providing these limits. A network IDS or IPS that is regularly updated with new signatures can detect when someone is attempting to use a known vulnerability to compromise the MARS appliance.

Another example, also using SSH, involves a brute-force password attack against the MARS appliance. In this attack, an attacker repeatedly uses a dictionary of passwords, using a script, to attempt to crack the password that administers the MARS appliance. MARS is especially vulnerable to this type of attack because the administrator's username is a well-known value—pnadmin—and this is the only username that can use SSH. This example is mitigated using the same methods as the first example. First, placing MARS on a protected network, with a stateful inspection firewall separating it from the rest of your network, allows you to limit connection attempts to a limited number of devices or networks on your company's network. Additionally, a network IDS or IPS can detect multiple login attempts, whether by SSH or web-based. The detection by an IDS can notify the appropriate personnel, or an IPS can prevent further attempts.

## Security Management Network

As a best practice, you should create a network as a security management network if you don't already have one. This network should contain various servers used for administering and monitoring the security of your network. The entire network should be protected by a firewall and IDS/IPS. Access to it should be tightly restricted, and any remote access to it should be through a Virtual Private Network (VPN).

Examples of hosts that should reside on this network include the following:

- MARS global controller (GC)
- MARS local controller (LC), if practical
- MARS archive server
- Firewall management consoles, such as Cisco Security Manager or Check Point SmartCenter
- IDS/IPS/HIPS management consoles
- Any existing syslog servers
- Vulnerability scanning hosts

The systems that reside on your management networks are some of the most sensitive in your organization. They often contain the keys to the kingdom, and for this reason, the management networks are targets of attackers. After an attacker has compromised a host on

a management network, an open freeway often exists to other systems because of the trust assigned to hosts on the management network.

Don't cut corners on network hardware that you use on your security management network. Install switches that support security features. You might want to consider configuring features such as private VLANs, which provide isolation between hosts on the same network. Other switch security features, such as the capability to prevent VLAN hopping, should also be considered.

## MARS Communications Requirements

Before you can protect MARS with a firewall, you first need to understand which TCP and UDP ports MARS requires to operate properly, and which of these carry outbound or inbound traffic. Table 4-1 provides a summary of all communications when MARS and the various monitored devices are all configured with default ports. Many or all of these can be changed, and you might need to modify this table for your installation.

**Table 4-1** *MARS TCP and UDP Ports*

Port	Description	Direction
TCP/21	Used by MARS to retrieve switch and router configuration files from centralized servers. FTP uses additional TCP ports (usually TCP/20), and most firewalls allow this to occur automatically.	Outbound
TCP/22	Used for management access to MARS LCs and GCs.	Inbound
	Used by MARS to connect to devices when learning topology or investigating hosts.	Outbound
TCP/23	MARS uses Telnet as one method to connect to some network devices when learning topology or investigating hosts.	Outbound
TCP/25	Used by MARS to e-mail reports and alerts.	Outbound
UDP/53	Used by MARS to look up host name-to-IP address resolution.	Outbound
TCP/53	Used by MARS to look up host name-to-IP address resolution.	Outbound
TCP/80	Used by MARS to communicate with Cisco routers for Distributed Threat Mitigation (DTM).	Outbound
	Used by MARS to receive some events, including web logs from iPlanet and Apache web servers, as well as NetCache.	Inbound
UDP/123	Used by MARS to synchronize time with Network Time Protocol (NTP) servers.	Outbound
TCP/137	Used by MARS to pull events from Windows systems.	Outbound
UDP/161	Used for Simple Network Management Protocol (SNMP) communications from MARS to monitored devices that use SNMP as the access method.	Outbound

**Table 4-1** *MARS TCP and UDP Ports (Continued)*

Port	Description	Direction
UDP/162	Used by MARS to receive SNMP traps from monitored devices that are configured to use traps for logging.	Inbound
TCP/443	Used for management access to MARS LCs and GCs.	Inbound
	Used by MARS to pull security events from Cisco IDS 4.x and IPS 5.x sensors and Cisco IOS IPS.	Outbound
	Used by MARS GCs and LCs for communications between appliances.	Inbound and Outbound
TCP/445	Used by MARS to pull events from Windows systems.	Outbound
UDP/514	Used by MARS to receive syslog messages from monitored devices.	Inbound
UDP/2049	Used by MARS to write archive data using Network File System (NFS).	Outbound
UDP/2055	Used by MARS to receive NetFlow data from monitored devices.	Inbound
TCP/8444	Used for communications between MARS GC and LC appliances.	Inbound and Outbound
TCP/18184	Used by MARS to pull event logs from Check Point firewalls.	Outbound
TCP/18190	Used by MARS to retrieve configuration settings from Check Point firewalls.	Outbound
TCP/18210	Used by MARS to retrieve certificates from Check Point firewalls or management consoles.	Outbound
All TCP/UDP	Used for vulnerability assessment scanning by MARS if enabled.	Outbound

## Network Security Recommendations

As you can see, depending on your environment and the location of hosts, a complex set of rules can be required on your firewall. Don't let the complexity prevent you from properly configuring the firewall, however. A little work initially can mean a better, more secure monitoring solution.

The following sections discuss issues regarding firewall protection for MARS and network-based IPSs and IDSs. The suggestions given are a good place to begin, but they by no means work in every network. For example, the TCP and UDP ports described in the preceding sections are only defaults. You can configure most of these services, which are common in

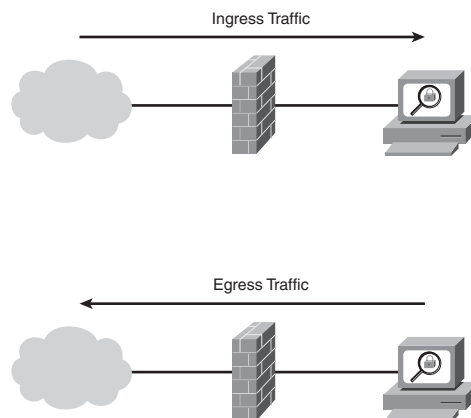
many networks, to use other ports. Check Point firewalls, for example, are commonly configured to use different ports than the defaults of TCP ports 18184, 18190, and 18210.

## Ingress Firewall Rules

To simplify the work involved, you should define some network object groups on your firewall. If you're not familiar with this term, think of object groups as variables that you can use while configuring the firewall to make life easier. Rather than referring to a large list of IP addresses or TCP/UDP ports, you can simply refer to a name instead. The following examples use an object group called CORP\_NET, which consists of all IP addresses used on your organization's network.

*Ingress traffic* refers to traffic that is inbound to a firewall (toward CS-MARS) from a less trusted network. Figure 4-1 shows both ingress traffic and *egress traffic*, or traffic that leaves CS-MARS to go toward the less trusted network.

**Figure 4-1** *Ingress and Egress Traffic*



The following ingress rules are a good starting point for most companies:

- Step 1** Permit syslog and SNMP trap traffic (UDP 162 and 514) from security operations (SecOps).
- Step 2** Permit NetFlow traffic (UDP 2049) from SecOps.
- Step 3** Permit HTTPS (TCP 443) from SecOps if a large number of people will be accessing the web console of MARS to run ad hoc reports. Otherwise, permit HTTPS to a restricted range of addresses.

- Step 4** Permit SSH (TCP 22) to a very restricted set of addresses. If the security management network has its own VPN gateway, which might be a function of the firewall, you might want to require administrators to establish a VPN connection before permitting SSH.
- Step 5** Permit HTTP (TCP 80) from any monitored web servers running iPlanet or Apache. If you're using NetCache appliances, permit HTTP from it as well.
- Step 6** If your MARS deployment consists of multiple MARS LCs that communicate to a centralized MARS GC, permit required management traffic between those systems (TCP 443 and 8444).
- Step 7** Deny all other traffic.

## Egress Firewall Rules

*Egress firewall rules* refer to filters that restrict traffic from the protected network to less trusted networks. Ideal security would restrict outbound traffic to only those ports that are necessary for proper functioning of the MARS appliance. However, in real life, this might be unmanageable. You need to determine the proper balance between security and manageability.

For example, a strict default egress policy might make sense for your company's public-facing web server. Hopefully, connectivity from the Internet to your web server (ingress rule) is permitted only on either TCP 80 or 443, depending on whether your web server uses encrypted HTTP. The egress policy should deny all traffic that originates from the web server to hosts on the Internet. In other words, someone should never be allowed to browse the Internet from your web server, to download files from the web server, or to have other communications from the web server to the Internet. By applying a proper egress rule on the firewall that denies it, an attacker is also denied that same communications path. In most instances where a web server, or any other server, is compromised by a hacker, the hacker's next steps include copying files to the web server. This is either to deface websites, install root kits, or retrieve the software needed to further hack into the network. Strict egress filters raise the difficulty level, often to a level that exceeds the capabilities of the hacker.

Depending on your environment and which MARS features you're using, strict egress filters might be unmanageable. However, you should evaluate them to see whether they are workable in your environment.

The following list of egress filters serves as a good starter set for most networks:

- Step 1** Permit traffic required for name resolution to CORP\_NET—for example, Domain Name System (DNS) and Server Message Block (SMB) for Windows hosts (TCP and UDP 53, TCP 137 and 445) to CORP\_NET.
- Step 2** Permit Network Time Protocol (NTP) to specified NTP servers, either on your network or internetwork.
- Step 3** Permit device discovery traffic on CORP\_NET for routers and switches—for example, Telnet (TCP 23), SSH (TCP 22), and SNMP (UDP 161).
- Step 4** Permit HTTPS to CORP\_NET to allow MARS to discover Cisco IDS/IPS sensors as well as to allow event retrieval from Cisco IDSs/IPSs and Cisco routers running IOS IPS, and to allow communications between MARS LCs and GCs. If possible, restrict this range to a subset of CORP\_NET.
- Step 5** Permit FTP (TCP 21) to a centralized FTP server that contains configuration files of routers and switches, if you want to take advantage of this feature.
- Step 6** Permit Simple Mail Transfer Protocol (SMTP) (TCP 25) to allow MARS to e-mail reports and alerts to your SMTP gateway.
- Step 7** Permit NFS (UDP 2049) if your MARS archive server resides on a different network (not recommended).
- Step 8** Permit TCP 8444 to allow communications between MARS LCs and GCs, if they reside in different locations.
- Step 9** Deny all other traffic.

If you want to take advantage of the MARS internal vulnerability assessment capabilities, the preceding list of rules *will not work*. Instead, use the following egress filter list:

- Step 1** Permit all TCP and UDP traffic sourced from CS-MARS or a third-party vulnerability scanner.
- Step 2** Permit NTP traffic to defined NTP servers, if they do not exist locally on SecOps.
- Step 3** Deny all other traffic.

In day-to-day use of MARS, when you choose to get more information about a specific host, the internal vulnerability assessment feature of MARS initiates a port scan of the host. You cannot accurately define an egress rule list that permits the vulnerability assessment to take place while also restricting outbound ports. If you already use a supported third-party



vulnerability assessment tool, such as QualysGuard, you do not need to use the internal tool. Otherwise, using the tool can greatly improve the accuracy of information presented to you by MARS.

## Network-Based IDS and IPS Issues

A network-based IPS offers an additional level of protection to complement that provided by a stateful inspection firewall. An IPS is closely related to an IDS. At first glance, the most obvious difference between the two is how they are deployed.


An IDS examines copies of network traffic, looking for malicious traffic patterns. It then identifies them and can sometimes be configured to take an automated response action, such as resetting TCP connections or configuring another network device to block traffic from an attacker.

### NOTE

It is important to remember that an IDS detects malicious traffic after it has already happened. Although automated response actions can take place, it is usually too late to stop the attack.

As shown in Figure 4-2, an IDS is typically deployed beside a traffic flow. It receives copies of network traffic from the network switches, hubs, taps, or routers. Because it does not sit in the flow of traffic, it does not break anything that MARS requires.

**Figure 4-2** *Intrusion Detection System*




An IDS often issues a large number of alerts based on traffic generated from MARS, especially if you're using the internal vulnerability assessment feature. You need to tune your IDS so that it does not alert on the vulnerability scans that originate from MARS. You might want to adjust the IDS tuning so that scans from MARS to your CORP\_NET are

ignored, but scans directed to the Internet trigger an alert. It is generally considered a bad practice to automatically scan hosts outside your own network; the practice might even be illegal. Make sure that MARS is not configured to scan anything that is not on your own network. Your firewall egress rules should not allow this either. However, in the case of a misconfiguration, your IDS can alert the appropriate personnel so that the configuration errors can be corrected.

An IPS sits in the path of network traffic (see Figure 4-3), usually as a transparent device (like a bridge), and watches for many of the same behaviors as an IDS. A major difference between the two, though, is the capability of the IPS to act instantly when malicious traffic is seen.

**Figure 4-3** *Intrusion Prevention System*



**NOTE**

In addition to the automated actions an IDS can take, an IPS can also prevent the malicious traffic from passing through it.

Because traffic must pass through an IPS, the IPS can prevent MARS from functioning properly if it is misconfigured. Take time to closely watch alerts generated by your IPS and tune it appropriately. Like the IDS, you should tune the IPS to allow vulnerability scanning to occur from MARS to CORP\_NET, while preventing it from scanning the Internet.

Some of the newest types of IPSs, such as the Cisco IPS, have a feature called *traffic normalization*. This feature, in particular, causes the MARS vulnerability assessment to fail. Traffic normalization enables several functions, including the following:

- Prevents illegal combinations of TCP flags from passing, or removes the illegal flags
- Prevents fragmented traffic from passing, or rebuilds it so that it is not fragmented
- Changes all packets in a traffic flow to have the same time to live (TTL)

This is just a small sampling of what a traffic normalizer does. In general, you can think of it as an engine that takes traffic that does not conform to standards, and either prevents the traffic from passing through the IPS or makes it conform to standards first.

By itself, traffic normalization breaks a large amount of attacks and reconnaissance activities. It also stops vulnerability assessment tools from being able to accurately determine information such as the operating system that a target host is running.

---

**NOTE** Cisco IPS 5.x and 6.x software, by default, does not generate alerts on most traffic normalization signatures. To properly tune the software, you need to enable alerts on that family of signatures.

---

If you're protecting your security management network with an IPS that supports traffic normalization, you need to tune it to either ignore the scans from MARS and Qualys (or other vulnerability scanners) or disable the traffic normalization capabilities.

## Summary

MARS contains sensitive information that you need to protect from malicious users. You must protect MARS with a firewall that is properly configured to allow necessary inbound and outbound traffic.

A network-based IDS or IPS can provide increased levels of security, but adds a level of complexity.