Cisco Security Monitoring, Analysis and Response System 4.2

GENERAL

Q. What is the Cisco® Security Monitoring, Analysis and Response System?

A. The Cisco Security Monitoring, Analysis and Response System (Cisco Security MARS) is an appliance-based, all-inclusive solution that allows network and security administrators to monitor, identify, isolate, and counter security threats. High-performance, scalable threat-mitigation appliances fortify deployed network devices and security countermeasures by combining network intelligence with features such as ContextCorrelation, SureVector analysis, and AutoMitigate capability, empowering companies to readily identify, manage, and eliminate network attacks and maintain compliance.

Going beyond first- and second-generation security information management systems, Cisco Security MARS more efficiently aggregates and reduces massive amounts of network and security data from popular network devices and security countermeasures. By gaining network intelligence, it effectively identifies network and application threats through sophisticated event correlation and threat validation. Verified attacks are visualized through an intuitive, detailed topology map to augment incident identification, investigation, and workflow. Upon attack discovery, the system allows the operator to prevent, contain, or stop an attack in real time by pushing specific mitigation commands to network enforcement devices. The system supports customer-centric rule creation, threat notification, incident investigation, and a host of security posture and trend reports.

The entire solution is cost-effectively delivered in an appliance platform that affords low adoption costs and flexible use. Cisco Security MARS appliances consist of standard Intel platforms with availability features accessible through a Web-based user interface, hardened OS, embedded Oracle database, proprietary logic, and scalable architecture with different performance characteristics and price points to address a broad range of customer sizes and deployment scenarios.

Q. Is there a software alternative?

A. No. Cisco Security MARS is a hardened, purpose-built appliance that includes an Oracle database, OS, and all necessary components for a scalable, high-performance security information management and security threat management solution.

Q. What types of information does Cisco Security MARS monitor?

A. Cisco Security MARS centrally aggregates logs and events from a wide range of popular network devices (such as routers and switches), security devices and applications (such as firewalls, intrusion detection systems [IDSs], vulnerability scanners, and antivirus software), hosts (such as Windows, Solaris, and Linux syslog), server-based applications (such as databases, Web servers, and authentication servers), and network traffic (such as Cisco NetFlow).

Q. Is there an additional charge for each firewall, IDS, or other device monitored by Cisco Security MARS?

A. Cisco Security MARS does not charge for agents. Customers pay one price for an appliance. As long as the appliance is able to keep up with the performance requirement, the customer does not have to pay any additional fee. In fact, the solution can be agentless.

Q. What is the capacity of a Cisco Security MARS appliance?

A. There are six different models of the appliance:

1. Cisco Security MARS 20R can sustain 50 events per second and 1,500 NetFlow flows per second.

2. Cisco Security MARS 20 can sustain 500 events per second and 15,000 NetFlow flows per second.

3. Cisco Security MARS 50 can sustain 1000 events per second and 25,000 NetFlow flows per second.

4. Cisco Security MARS 100e can sustain 3000 events per second and 75,000 NetFlow flows per second.

5. Cisco Security MARS 100 can sustain 5000 events per second and 150,000 NetFlow flows per second.

6. Cisco Security MARS 200 can sustain 10,000 events per second and 300,000 NetFlow flows per second.

Note: The NetFlow numbers indicate flows per second, not packets per second. The numbers stated are over a sustained rate; the system can manage peaks that exceed these numbers.

Q. **What happens if the rate of events per second is higher than the rate supported?**

A. If the load of events is higher that what the system supports, a queue prioritizes the events that will be processed, and the lower-severity events are dropped while higher-severity events are still processed.

The rate of events supported officially by the system is for a sustained period of time. If the peak exceeds those numbers, the events will not be dropped. Cisco Security MARS has a strong capacity to sustain high peaks of events per second during a network attack.

Q. **Do I need extra licenses besides the appliance?**

A. All required components and licenses are included in the list price.

Q. **If I have multiple Cisco Security MARS appliances, is there a recommended deployment architecture?**

A. The Cisco Security MARS Global Controller provides a central console for multiple-box solutions. In the Global Controller architecture, the Local Controller forwards summarized data to the Global Controller where the combined "zones" (topology, incidents, etc.) are displayed. All administrative access and control can be performed from the Global Controller.

A standalone solution is for customers who have all their devices locally in one geographical location or who do not mind inundating their WAN links with Syslog, NetFlow, and Simple Network Management Protocol (SNMP) data. For geographically distributed sites and instances where customers do not want to flood their WAN links, customers should consider the distributed solution.

Q. **What is a zone?**

A. A zone is an area of a customer network related to one Local Controller. Each Local Controller represents a specific zone.

Q. **Is the appliance secured?**

A. The Cisco Security MARS appliance comes security-hardened. All unnecessary services such as Simple Mail Transfer Protocol (SMTP) have been removed.

DEVICE SUPPORT

Q. **What is the list of supported devices in Cisco Security MARS?**

A. The complete list of supported devices is available at:
http://www.cisco.com/en/US/products/ps6241/products_device_support_tables_list.html.

Q. **What is the best way to add my devices and configure the network?**

A. There are different ways to add devices to Cisco Security MARS:

- From seed file

- Manually, one by one

- Using a discovery process

For more details on how to build the seed file, how to add single devices, and how to discover devices, see "Adding Reporting and Mitigation Devices" at: http://www.cisco.com/en/US/products/ps6241/products_user_guide_chapter09186a008075038a.html.

Q. **Can Cisco Security MARS receive events from devices not officially supported?**

A. Yes. This can be done in two ways:

- Use the "custom parser" to create a specific parser to allow Cisco Security MARS to understand syslog for devices not officially supported.

- All the events that Cisco Security MARS is unable to understand are stored in the database as "unknown." You can run queries on those events using the keywords option. In this case, "unknown" events are also searched and if the string matches the keyword, these events will be reported in the result.

For information on defining a custom parser, visit:
http://www.cisco.com/en/US/products/ps6241/products_user_guide_chapter09186a008074f1e2.html.

Q. **How do I export a list of all the added devices?**

A. To get the full list of security and monitoring devices, click Admin > Security and Monitoring Devices. Go to the bottom of the page and select a range that displays all of the devices. Select all the text on the page (Ctrl+A), and copy (Ctrl+Insert) and paste (Shift+Insert) that selection to a textfile editor for manual cleanup.

Q. **How does Cisco Security MARS collect logs and events?**

A. It can collect events by either pushing or pulling, depending on the specifics of the device itself.

Most of the devices are able to send events or syslog messages, while few of them, such as Cisco Secure Access Control Server (ACS) or host servers, need to let the Cisco Security MARS obtain the log files directly from the server.

Q. **On which devices can Cisco Security MARS perform mitigation?**

A. Cisco Security MARS generates the mitigation command for Cisco, NetScreen, and CheckPoint devices. Layer 2 mitigation can be done for any device that has SNMP MIB II support. However, only Cisco devices have been tested.

For Layer 3 devices, the Cisco Security MARS recommends a command and the device that should enforce that command. The user must copy and paste the command in the command-line interface (CLI) of the recommended device.

Q. **Can specific hosts be flagged as more vulnerable than others (for example, if they do not have the latest patches installed)? If so, how is this accomplished?**

A. No. Cisco Security MARS is not a vulnerability assessment tool. However, you can use Cisco Security MARS to access the information collected by the supported vulnerability assessment tools to get information such as the OS, the service pack, patch level, and other data.

Q. **What is the preferred method of discovering devices: Telnet, Secure Shell (SSH) Protocol, or SNMP?**

A. SNMP is the preferred method for discovering many devices quickly. However, Telnet or SSH must be used with routers with Network Address Translation (NAT) or access control list (ACL) information.

Q. **How is network topology information used on Cisco Security MARS?**

A. Cisco Security MARS uses network topology in a variety of ways:

- To identify various events and flows for the same session—Such events may be generated by network devices across NAT boundaries and hence have different source or destination addresses and ports. Cisco Security MARS uses patent-pending algorithms that use the knowledge of topology and device-configuration information to correlate these diverse events into one session. This reduces the amount of event data and creates the full context of an attack.

- To reduce false positives—By identifying events for the same session and by analyzing the topological path taken by an attack from the source to the destination, Cisco Security MARS can identify whether an attack actually reached the intended destination or was dropped by an intermediate device such as a firewall or an intrusion prevention system (IPS).

- To identify the optimal mitigation point—By analyzing the path from an attacker to the intended destination, Cisco Security MARS can determine the device closest to the attacker as the optimal mitigation point. Device configuration knowledge also enables the system to generate accurate device-specific mitigation commands.

- To identify attack paths and network hotspots—Cisco Security MARS uses the knowledge of network topology to visualize the topological attack path and identify the network zones containing attackers and attacked hosts.

- For enhanced forensics—Because Cisco Security MARS knows about NAT from device configurations, it can provide greatly enhanced forensics by identifying network address translations and attacker MAC addresses. You can query the system by pre- and post-NAT addresses.

Q. **Can Cisco Security MARS work without network topology?**

A. Yes. It can be used as a basic logging appliance where it will collect and correlate events, but where the network topology connectivity, hotspot graph, and mitigation capabilities would be missing. The NATs and some sessionization can still be obtained if the firewall logging levels are set high enough.

Q. **Can I configure Cisco Security MARS to only discover an area that I want? For example, can I keep it from discovering my ISP's network?**

A. Yes. By defining valid networks, you can force Cisco Security MARS to discover only internal networks. A typical choice of valid networks for an organization includes the private addresses 10.0.0.0/8 and 192.168.0.0/16, and the public addresses assigned to that organization.

Q. **How does Cisco Security MARS keep its network topology updated?**

A. You can schedule device rediscovery on Cisco Security MARS to accomplish this. This rediscovery can be scheduled on a network basis.

Q. **How long does it take Cisco Security MARS to discover a large network?**

A. The answer depends on the size of SNMP information returned by the devices and how busy the devices are. It takes a couple of hours for an enterprise-grade, 300-node network. However, event processing can continue while the network is being discovered. Simply discover the monitored devices and activate the changes.

Q. **What are the benefits of NAT and Port Address Translation (PAT) information?**

A. Cisco Security MARS can understand NATs and PATs. When you see an incident in the system, you can ask for the address, pre- or post NAT or PAT. For example, if your ISP asks the identity of a firewall PAT IP address at a particular time, you can query by post-NAT address and you might find a DNS name or a Windows workstation name if an incident fired.

Q. **How can Cisco Security MARS perform mitigation?**

A. It performs mitigation using the following protocols:

- Telnet

- SSH

- CheckPoint Management Interface (CPMI)

- SNMPv1 (write access required)

Also, Cisco Security MARS can let you know, when an incident is fired, what command to deploy to mitigate the incident and on which device. Based on the awareness it has of the network topology, it can suggest the best place to implement a specific mitigation command.

NETFLOW AND NAT/PAT ANALYSIS

Q. **What is NetFlow?**

A. NetFlow is a technology developed by Cisco Systems® for monitoring network traffic. It is supported in many Cisco IOS® Software images. NetFlow uses a User Datagram Protocol (UDP)-based protocol to periodically report on flows seen by the router. A flow is a

Layer 7 concept consisting of a session setup, data transfer, and session teardown. For every flow, a NetFlow-enabled router records several flow parameters, including:

- Flow identifiers (source and destination addresses and ports and protocol)
- Ingress and egress interfaces
- Packets exchanged
- Bytes transferred

Periodically, a collection of flows and their associated parameters are packed in a UDP packet according to the NetFlow protocol and sent to specified collection points. Because multiple flows are packed into one UDP packet, NetFlow is a highly efficient mechanism for monitoring high volumes of flows compared to traditional mechanisms, such as syslog or SNMP.

There are also hardware-based implementations of NetFlow. The flow information contained in NetFlow packets is similar to that sent through syslog, SNMP, or Check Point Log Export API (LEA) by enterprise firewalls such as Cisco PIX® Firewall, NetScreen ScreenOS, and Check Point Firewall-1. However, NetFlow is significantly more efficient than the other protocols. To receive traffic logs, the syslog levels on the firewalls have to be set to DEBUG (quite high), which in turn causes throughput degradation on the firewalls at moderate to high loads.

NetFlow provides the following data:

- Network usage—Top users, ports, etc. via queries and reports; arranged by bytes and sessions.
- Traffic-anomaly detection—For example, when a user is sending or receiving (statistically) unusually large numbers of flows on a port, such as (hard-coded) excessive Internet Relay Chat (IRC) connections, Internet Control Message Protocol (ICMP) traffic, SMTP traffic, etc. from the same source.

Q. **From which devices should NetFlow information be collected?**

A. Ideally, NetFlow information should be collected from the distribution switches and routers. These devices, together with NetFlow from Internet-facing routers or syslog from firewalls, represent the entire network. Cisco Security MARS normalizes NetFlow and syslog so there are no issues there.

Q. **How many NetFlow records can Cisco Security MARS manage?**

A. It depends on the appliance model. Details are provided in the General section as part of the question "What is the capacity of a Cisco Security MARS appliance?"

Q. **How does analysis of Layer 2 work in Cisco Security MARS in terms of incident detection and mitigation?**

A. To identify the port to block, Cisco Security MARS performs IP to MAC address mapping and then MAC address to physical switch port mapping using the Spanning Tree Protocol for Layer 2.

HOW QUERIES, REPORTS, AND RULES WORK

Q. **What is the difference between queries, reports, and rules?**

A. Queries, reports, and rules are created in the same way. The difference between them is the following: A *query* is run in a specific moment to investigate an incident and obtain details. If after you have created a query, you want to run it often, or associate an action to it, you can save it as a report. The result of a query and of a report is exactly the same, but you can save reports to avoid having to enter the values every time you want that specific output. A report can be added to the "My Reports" page. Reports are also more flexible in their generation, format, and delivery options. They can be scheduled to run automatically at specified intervals, can be output in HTML or CSV format, and can be e-mailed to recipients upon completion of a scheduled generation.

A. A *rule* is more complex than a report. A rule analyzes the events and sessions created by Cisco Security MARS, and generates incidents based on the sequence specified. A rule cannot be deleted; therefore you should work on a rule first to tune the result you want. Only after you know that the query is exactly what you are looking for, then you should save it as a rule.

Q. **In what format can the report be exported, and how?**

A. Cisco Security MARS generates reports in CSV and HTML formats. Reports can be e-mailed or exported in either format.

Q. **What is the "My Reports" page?**

A. There is a tab called "My Reports" where you can add your own reports. This provides an easy way to view the reports you want to access more often.

Q. **Can Cisco Security MARS send an e-mail alert if a firewall rule is hit?**

A. Yes. It is possible to write a rule in Cisco Security MARS and specify the same five components as the firewall rule and an action (e-mail). The e-mail can be sent to a single e-mail address or to a group of addresses.

Q. **What type of alerts can I configure on Cisco Security MARS?**

A. Cisco Security MARS alerts administrators using the following protocols:

- E-mail
- SNMP trap
- Syslog
- Pager
- Short Message Service (SMS)
- Distributed Threat Mitigation (new in Cisco Security MARS 4.1)
- Extensible Markup Language (XML) document (new in Cisco Security MARS 4.2)

CISCO DISTRIBUTED THREAT MITIGATION WITH IPS

Q. **Where can I learn about Cisco Distributed Threat Mitigation with IPS?**

A. To learn more about Cisco Distributed Threat Mitigation with IPS and how to configure this feature in Cisco Security MARS, visit:
http://www.cisco.com/en/US/products/ps6241/products_configuration_example09186a008067a2b0.shtml.

ARCHIVING AND MAINTENANCE

Q. **Is it possible to archive data?**

A. Cisco Security MARS archives compressed data to an offline store using the network file system (NFS) protocol.

Q. **Can I reinstall while connected to the network?**

A. While using the Recovery DVD, make sure none of the interfaces are connected because this may prevent the system from setting up correctly.

Q. **How exactly does archiving work?**

A. If archiving is enabled, the events will be written twice: once to the database and once to the NFS share archive. The local database is still important because queries do not run over archived data, and can analyze only data in the database.

When you want to look at the old data to do any kind of investigation, look at the old incidents, run queries, or run reports, you must restore the old data in another appliance. The reason to use a separate box to look at old data and restore the archive is that when a restore is performed, all the data and configuration on the box is deleted.

For more information on how archiving works and its possible uses, see "Configuring and Performing Appliance Data Backups" at: http://www.cisco.com/en/US/products/ps6241/products_installation_and_configuration_guide_chapter09186a00804c4db4.html.

GLOBAL CONTROLLER AND LOCAL CONTROLLER

Q. **What is the advantage of using Global Controller and Local Controller?**

A. Users who have a large network and want a distributed architecture can implement one Global Controller and multiple Local Controllers in different locations of the network. This allows more scaling, because the load of the computation is distributed to the local units.

The Global Controller provides the summary of all Local Controller information:

- Network topologies
- Incidents
- Queries and reports result

It also provides a central point for creating rules and queries, and then applies them to multiple Local Controllers simultaneously. Changes in the Local Controller are automatically propagated to the Global Controller and vice versa. You can easily navigate to any Local Controller from the Global Controller GUI.

This deployment architecture simplifies the ability to logically divide the network based on zones (whether they are departmental or regional). This improves overall scalability and organizational workflow.

Q. **How can the Local Controller and Global Controller share device information?**

A. Devices are treated as follows:

- Device and topology information in the Local Controller is pushed to the Global Controller
- Adding devices is common on Local Controllers
- You can add devices while logged into the Global Controller, through the active zone
- Devices and topology information in the Local Controller is pushed to the Global Controller

Q. **What incidents can I see at the Global Controller?**

A. When a Global Controller and one or more Local Controllers are used, each Local Controller has all the rules needed to generate incidents (they have been defined at either the Local Controller or the Global Controller). The Local Controller processes all the events; then, only the incidents are propagated up to the Global Controller, saving significant processing on the Global Controller. This allows scaling on a more distributed level.

More specifically:

- Only the incidents fired on the Local Controller against global rules are pushed up to the Global Controller; any incidents fired against rules created on the Local Controller will not be pushed up to the Global Controller.
- Incidents on the Global Controller can be viewed based on the selected Local Controller

The Incident ID includes the zone name. If a user adds an alert to the global rule created on the Global Controller, and the rule is pushed down and fired on the Local Controller, the designated user receives the alert from the Local Controller, not the Global Controller.

Q. **What rules can I see at the Local Controller?**

A. The built-in rules shipped with the system are global and are defined both in the Local Controllers and the Global Controller.

- Rules created on the Global Controller are known as "global rules"

- Global rules are pushed down to all Local Controllers (with the exception of system rules)

- Changes made to system rules on the Global Controller are pushed down to Local Controllers

- Identical rules created on the Global Controller and Local Controller can be determined by their naming convention

- Global user-inspected rules are named as "Global Rule -*<name>*" on the Local Controller

- Global rules are automatically pushed down to a newly added Local Controller

- Global rules cannot be deleted but the state can be changed to "inactive" on both the Global Controller and the Local Controller

Q. **Can I scale running reports at the Global Controller?**

A. You can create batch queries and reports for selected zones, which are then pushed down to the Local Controllers. Global batch query and report results from each selected Local Controller are pushed up to the Global Controller. On the Global Controller, you can view reports based on:

- Aggregated *results* from all the Local Controller filters by total, peak, or recent activity

- The sum of *all* Local Controllers, list zones, or one specific zone