



CHAPTER 6

Incident Investigation and Mitigation

An incident is a chain of events that are correlated by a rule to signal an attack upon your network. MARS simplifies and expedites the detection, mitigation, reporting, and analysis of the incident. The Network Summary dashboard and the Incident pages help to detect recent incidents and show the rules and the events that compose them. Mitigation refers to the ability of the MARS to isolate the attacking and compromised network devices by identifying and configuring enforcing devices that act as choke points in the network. Queries and reports reveal the scope of a problem and gather data for analysis and regulatory compliance. All this information can be captured in a case report with Case Management and escalated to the relevant personnel.

Incidents Overview

An attack can consist of a reconnaissance activity (for instance, a port scan), followed by a penetration attempt (such as, a buffer overflow), and followed by malicious activity on the target host (for example, a local privilege escalation attack or the installation of backdoors).

An incident, which is generated by a Local Controller, collects the interesting events that constitute an attack scenario and uses rules to describe them. MARS provides you with pre-defined, system rules—which you can fine tune—and gives you the ability to create your own rules.

Incidents that appear on the Global Controller are fired by global rules at the Local Controller level and are compiled at the Global Controller level. Incidents that appear on a Local Controller are fired by rules local to that Local Controller. They are used by Local Controllers for local reporting and are *not* propagated *up* to the Global Controller.

Predefined System Rules are treated as global rules. When an incident is fired by a system rule on the Local controller, it gets propagated to the Global Controller.

Incidents are sub-divided into instances to make it easier for you to investigate the attack scenario. Each instance alone is a full attack scenario.

For example, if your network is probed for a DoS attack and then attacked, a rule fires when it sees the follow up attack. The incident displays the instances of this attack.

Figure 6-1 A DoS probe followed by a DoS attack

Incident ID: 42998483

Offset	Firing Event / Session / Incident ID	Event Type	Source IP / Port
Instance 1			
3		[1906920] Net Flood TCP	+ Total: 5
Instance 2			
3	S:45754259, I:42998483, I:42998484	[1906910] Net Flood UDP	10.4.17.4
Instance 3			
1		[1905037] WWW SGI MachineInfo Info Leak	10.1.1.21
1	S:45775179, I:42998480, I:42998481, I:42998483, I:42998487, I:42998490, I:42998492, I:42998493, I:42998495	[1905110] WWW SuSE Installed Packages Info Leak	10.1.1.21

The Incidents Page

Click the **Incidents** tab to navigate to the Incidents page. The Incidents page displays recent incidents. Incidents are collections of events and sessions that meet the criteria for a rule, each having helped to cause the rule to fire. An incident’s duration only includes the events that contributed to the incident firing.

Figure 6-2 Global Controller Incidents Navigation Page

Incident ID	Event Type	Matched Rule	Action	Time	Path	Cases
1:347915126 (pluto)	Built/teardown/permitted IP connection	System Rule: Client Exploit - Sysbug Trojan		Sep 20, 2005 10:17:07 AM PDT		C:119662 (New) Security Team
1:347915135 (apollo)	Deny connection - no xlate Built/teardown/permitted IP connection PIX reserved a network state container for a host	abcRule-New2		Sep 20, 2005 10:16:42 AM PDT - Sep 20, 2005 10:17:06 AM PDT		
1:347915137 (pluto)	Deny connection - no xlate Built/teardown/permitted IP connection PIX reserved a network state container for a host	test save as rule		Sep 20, 2005 10:16:42 AM PDT - Sep 20, 2005 10:17:06 AM PDT		

Arrows 1-9 point to: 1. Incident ID, 2. Event Type, 3. Matched Rule, 4. Action, 5. Time, 6. Path, 7. Cases, 8. Path, 9. Cases.

1	Name of the Local Controller reporting the incident, also links to the Local Controller Incident page.	2	Links to the Incident Detail page of the reporting Local Controller.
3	The incident severity indication icon	4	The events that compose the Incident. Links to the Event Type Details popup window.
5	Query icon. Links to the Query page and populates the corresponding query field with the item.	6	The rule that fired to create the incident. Links to the rule page to display the details of the rule.
7	Start and end time of the incident.	8	Links to the the reporting Local Controller Incident Path and Incident Vector diagrams.
9	Links to the View Case page of the reporting Local Controller		

The Incident page's table:

- *Incident ID*

An incident's unique ID, followed by the Local Controller on which the incident occurred.

- *Severity*

Low (green), medium (yellow), and high (red) icons.

- *Event Type*

The normalized signature sent from the reporting devices.

- *Matched Rule*

The rule whose criteria were met.

- *Action*

The description of the notification taken when this rule fires (epage, email, etc.)

- *Time*

A single time or a time range (see [Time ranges for Incidents, page 6-73](#) for more information)

- *Incident Path*

The icon that takes you to the incident's path diagram on the Local Controller.

- *Incident Vector*

The icon that takes you to the source, event type, and destination diagram on the Local Controller.

Time ranges for Incidents

The time column displays both single entries for time (Sep 6, 2003 12:09:54 PM PDT), and time ranges (Sep 6, 2003 12:06:43 PM PDT - Sep 6, 2003 12:06:47 PM PDT).

A single time tells you that all of the firing events were received in the same second. The duration of the incident includes only events that have fired that incident.

Incident Details Page

Clicking the Incident ID takes you to its Incident Details page on the Local Controller. The Incident Details page is rich in information and information gathering tools. This page answers questions, such as who did it, what event types happened, when it happened, and to whom it happened.

Figure 6-3 The Incident Details Page

Rule Name: aLcTest **Status:** Active
Action: None **Time Range:** 0h:10m
Description: aLcTest

Offset	Open (Source IP	Destination IP	Service Name	Event	Device	Reported User	Keyword	Severity	Count	Close	Operation
1		ANY	ANY	ANY	ANY	cherryWall	ANY	ANY	ANY	1		

Incident ID: 200982691

Offset	Session / Incident ID	Event Type	Source IP/Port	Destination IP/Port	Protocol	Time	Reporting Device	Reported User	Path / Mitigate	False Positive
1	S:200882690, I:200982691, I:200982688, I:200982689, I:200982690	Built/teardown/permitted IP connection PIX firewall login failed, TCP access requested to the PIX firewall, TCP or UDP access permitted to the PIX firewall, SSH session disconnected for a reason	10.2.3.33 40224	10.4.5.1 22	TCP	Sep 5, 2005 11:20:05 AM PDT	cherryWall	pix		False Positive

Copyright © 2003, 2005 Cisco Systems, Inc. All rights reserved. Summary :: Incidents :: Query / Reports :: Rules :: Management :: Admin :: Help ::

On the top of this page are the tools that let you search for Incident and Session ID and view the Matched Rule.

To Search for a Session ID or Incident ID

Step 1 Enter the ID into the appropriate field.

Step 2 Click the **Show** button.

To view a partially hidden rule

Click the Show button next to the Rule Description.



Note

Incidents can only be included in a case or mitigated from the Local Controller.

Incident Details Table

When you click the Incident ID, the Incident Details table appears in a separate browser on the Local Controller. Each row of the Incident Details table represents either a session or the information common to a group of sessions. You can see all of the collapsed session information by clicking the plus signs to expand the group. You can expand or collapse all of the incident's information by clicking the **Expand All** or **Collapse All** buttons.

Figure 6-4 Expanding a Row in a Table

Offset	Session / Incident ID	Event Type	Source IP/Port	Destination IP/Port	Protocol	Time	Reporting Device	Reported User	Path / Mitigate	False Positive
1		Built/teardown/permitted IP connection	Groups: 6, Total: 12							
1		Built/teardown/permitted IP connection	Groups: 6, Total: 12							
1		Built/teardown/permitted IP connection	0.0.0.0	0	0.0.0.0	0	TCP	Sep 5, 2005 11:20:05 AM PDT	cherryWall	Total: 4
1		Built/teardown/permitted IP connection	10.2.3.42	51893	10.4.1.20	18184	TCP	Sep 5, 2005 11:20:09 AM PDT	cherryWall	Total: 2
1	S:200882705, I:200982691, I:200982688, I:200982689, I:200982690	Built/teardown/permitted IP connection	10.2.3.43	52499	10.4.1.251	443	TCP	Sep 5, 2005 11:20:05 AM PDT	cherryWall	False Positive
1		Built/teardown/permitted IP connection	10.4.1.200	1025	10.1.1.189	514	UDP	Sep 5, 2005 11:20:05 AM PDT	cherryWall	Total: 2
1	S:200882686, I:200982691, I:200982688, I:200982689, I:200982690	Built/teardown/permitted IP connection	10.4.2.11	22	10.2.3.33	40222	TCP	Sep 5, 2005 11:20:05 AM PDT	cherryWall	False Positive
1		Built/teardown/permitted IP connection	67.116.29.66	3684	Total: 2					
1	S:200882690, I:200982691, I:200982688, I:200982689, I:200982690	PIX firewall login failed, TCP access requested to the PIX firewall, TCP or UDP access permitted to the PIX firewall, SSH session disconnected for a reason	10.2.3.33	40224	10.4.5.1	22	TCP	Sep 5, 2005 11:20:05 AM PDT	cherryWall, pix	False Positive

Copyright © 2003, 2005 Cisco Systems, Inc. Summary :: Incidents :: Query / Reports :: Rules :: Management :: Admin :: Help :: Feedback

This high-density information table lets you drill deep into incidents. Click the Query icon anywhere on this page to query on a particular criteria. Click the Raw Events icon for raw events for a particular session. You can click the Tune link to tune incidents for False Positives, see The False Positive Page, page 6-78 or click the Mitigate link to mitigate an attack.

Figure 6-5 Incident Table

Incident ID: 200982691

Expand All Collapse All

Offset	Session / Incident ID	Event Type	Source IP/Port	Destination IP/Port	Protocol	Time	Reporting Device	Reported User	Path / Mitigate	False Positive
1		Built/teardown/permitted IP connection	Groups: 6, Total: 12							
1	S:200882690, I:200982691, I:200982688, I:200982689, I:200982690	PIX firewall login failed, TCP access requested to the PIX firewall, TCP or UDP access permitted to the PIX firewall, SSH session disconnected for a reason	10.2.3.33	40224	10.4.5.1	22	TCP	Sep 5, 2005 11:20:05 AM PDT	cherryWall, pix	False Positive

1	Incident ID	2	Severity icon
3	Path and Incident Vector icons. Launch popup windows to display Path and Incident Vector diagrams (L2 or L3 attack path information)	4	Offset number
5	Links to Session and Incident Detail pages of all incidents within the session	6	Links to the Event Type Details pages
7	Launches False Positive popup window	8	Link to the Device Information page
9	Query icon links to Query page	10	Click Device icon to launch popup window to display raw message information
11	Link to the Mitigation Information page	12	Link to the False Positive Tuning page

The following information describes some of the fine points of this table.

- Instances

Sometimes rows are split into instances. The *only* relationship among the different instances is that they fired the same rule in the same time frame.

- *Session/Incident ID*

This column shows the sessions that contributed to the incident, and the other incidents those sessions belong to.

- *Events column*

The Events column shows types of the firing events. Multiple firing events of the same types are shown once per session.

- *Time column*

An incident's duration only includes the events that contributed to the incident firing.

False Positive Confirmation

When investigating incidents, you will invariably come across false positive events. In some cases, firing events are classified automatically by MARS as system-confirmed false positives and unconfirmed false positives. Vulnerability scanning often identifies the false positive events, but at times you must investigate events to determine their validity.

To understand the false positive nomenclature and what tasks you are expected to perform within the user interface, we must study the possibilities among three variables surrounding possible attacks: legitimate attack, valid target, and attack detected. We examine these differences in [Table 6-1](#).

Table 6-1 Attack Type Truth Table

	Legitimate Attack	Valid Target	Attack Detected
invalid scenario	0	0	0
False Positive	0	0	1
invalid scenario	0	1	0
False Positive	0	1	1
False Negative	1	0	0
Attack/Alarm (noise)	1	0	1
True False Negative	1	1	0
Intrusion/True Alarm	1	1	1

Based on the valid cases in [Table 6-1](#), we can clearly distinguish the false positive terminology:

- A *legitimate attack* is an actual attempt by an attacker to gain access to or information about a specific host using a known exploit.
- A *valid target* is a host that is susceptible to the launched attack. A host can become an *invalid target* if it is properly patched or has some other preventative measure in place, such as a local firewall, virus scanner, or intrusion prevention software that guards against the attack.
- *Attack detected* refers to whether the monitoring device detected the attack and generated an alarm.
- A *false positive* is when the monitoring system generates an alarm for a condition that is benign. In this case, there is no legitimate attack, despite the alarm generation.

- An *unconfirmed false positive* is one where the monitoring system, based on data not available to the reporting device, has determined that an alarm is a false positive. Unconfirmed refers to the fact that the administrator must review and accept or reject the assessment of the false positive.
- A *false negative* is when the monitoring system fails to detect a legitimate attack.
- *Noise* refers to those alarms that are triggered due to attacks against invalid targets. While they can represent real attacks, the target cannot be compromised due to preventative measures. Attacks that fall within the noise category are of secondary importance in terms of investigation and mitigation.
- *Intrusion* identifies a successful attack against the host, where the host is compromised by the attacker.
- A *true false negative* identifies an intrusion that remains undetected by the monitoring system.
- A *true alarm* identifies an intrusion that is detected by the monitoring system.

When a Local Controller receives an event, it is evaluated against the conditions of the defined rules. If the event satisfies the conditions of a rule, then the incident triggers. When an event triggers an incident, we refer to that event as a *firing event*. False positive analysis is performed for such firing events to reduce the number of false alarms.

Using built-in event vulnerability data, learned topology paths, sessionized event data, ACL analysis of layer 2 and 3 reporting devices, supporting data from 3rd-party vulnerability analysis (VA) software (such as Foundstone and eEye), and information that you provide about hosts, MARS analyzes the firing events reported to it determine whether they hold up to a higher-level review.

In the case of MARS, a *system-confirmed false positive* is where, after further analysis, a firing event is determined to be invalid. Example system-confirmed false positives include:

- When an IDS device monitoring the network outside of a firewall reports an attack; however, the firewall drops that session as part of its standard access restrictions. Therefore, the attack never reaches the target.
- Cisco Security Agent detects an attack and blocks it.

An *unconfirmed false positive* is where, after further analysis, the firing event is believed to be invalid primarily due to the attack being against an invalid target. Example unconfirmed false positives include:

- A reporting device reports a valid attack against a host; however, the host is not susceptible to that attack because it targets a different operating system. You can reduce these types of false positives by employing OS fingerprinting technologies on the reporting devices.
- A reporting device reports a valid attack against a host's application; however, the host is not susceptible to that attack because it targets a different application.
- A reporting device reports a valid web attack against TCP port 80, however, dynamic probing determines that no services on the target host listen to TCP port 80.

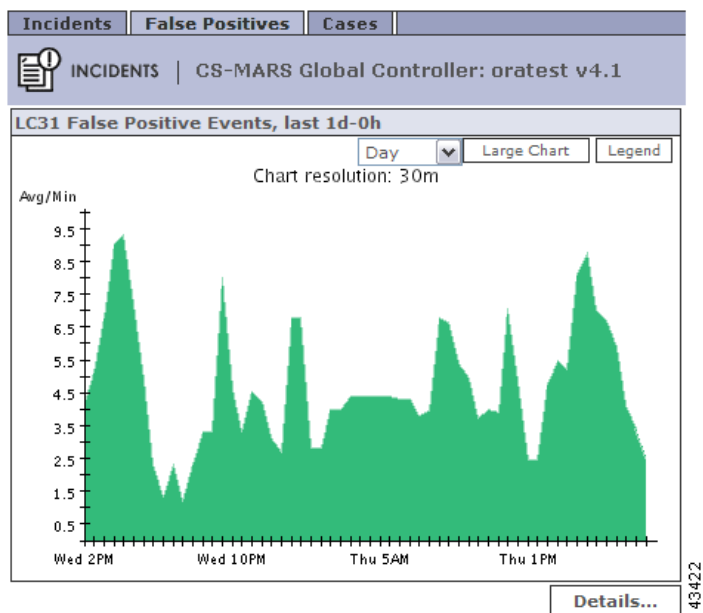
For unconfirmed false positives, you must manually investigate the alarm and specify in Global Controller whether it is an actual false positive. For actual false positives, you should define a drop rule for the event. Defining a drop rule does not mean that the event is not stored in the database, you have the option of dropping the event from incident evaluation and either shoring it in the database or not. Whether you store the event in the database or not, events matching the event type and target host can no longer act as firing events. By refining the event processing in this fashion, MARS frees up your time to focus on actual incidents by more accurately correlating events into incidents and reducing noise.

As part of your operational strategy, you should strive to refine event generation and processing to tune out the possibility for false positives. You can perform such tuning at the device level, by refining what traffic or action can generate an event, and at the Local Controller level by providing more information about your network, such as identifying the operating system of hosts attached to the network segments monitored by that Local Controller.

The False Positive Page

To navigate to the False Positives page, click **Incidents**, and click the **False Positives** sub-tab.

Figure 6-6 False Positive Graph for a Local Controller



The False Positives page is where you can see groupings of False Positives for each Local Controller or for the sum of all the Local Controller zones. You can change the graph by selecting **Hour**, **Day**, **Week**, **Month**, **Quarter**, or **Year** from the first pull-down menu, and **Sum Zones** or an individual local zone from the second menu.

If you want to see details of the false positive on the Local Controller, click the Details button.

Virtual Private Network Considerations

Currently, MARS cannot display accurate Path/Mitigation information or compute the complete route of an attack originated by a host with a source IP address on a virtual private network (VPN). MARS can identify the attacking host if the VPN IP address of the host was supplied by a Cisco 3000 Series VPN Concentrator configured as a MARS reporting device.



Note

You must be able to recognize from your knowledge of your network that the IP address of the attacking host is an IP address allocated to a VPN.

To identify a host attacking from a VPN, perform a query of “Cisco VPN User connected/disconnected” events for the Cisco VPN Concentrator device. The attacking host name or next network element is disclosed in the raw messages of the events.

